Tivoli® software

IBM

IBM Tivoli Storage Manager Version 5.1

Technical Guide

Server-free backup and restore for Windows 2000

Online image backup for Linux and Windows 2000

Performance enhancements and HACMP support

Barry Kadleck
Nigel Bentley
Michael Dann
Steven Pemberton
Björn Rodén

Redbooks

ibm.com/redbooks

International Technical Support Organization

Tivoli Storage Manager Version 5.1
Technical Guide

June 2002

SG24-6554-00

Take Note! Before using this information and the product it supports, be sure to read the general information in "Notices" on page xiii.

First Edition (June 2002)

This edition applies to IBM Tivoli Storage Manager Version 5 Release 1.

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

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

Contents

Figures

Tables

Notices

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

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

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

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

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

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

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:
This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AFS®
AIX®
AIX 5L™
AS/400®
DB2®
DB2 Universal Database™
DFS™
DPI®
e (logo)®
Enterprise Storage Server™
ESCON®
FICON™
FlashCopy®
HACMP/6000™
IBM®
Illustra™
Informix®
Integrated Language Environment®
iSeries™
Language Environment®
Magstar®
MVS™
OS/2®
OS/390®
OS/400®
Perform™
POWERparallel®
PowerPC®
pSeries™
Redbooks™
Redbooks(logo)™
RS/6000®
SANergy™
Scalable POWERparallel Systems®
Sequent®
SP™
Tivoli®
Tivoli Enterprise™
TotalStorage™
VM/ESA®
WebSphere®
xSeries™
z/OS™
zSeries™

The following terms are trademarks of International Business Machines Corporation and Lotus Development Corporation in the United States, other countries, or both:

Lotus®
1-2-3®
Approach®
Domino™
Passport Advantage™

The following terms are trademarks of other companies:

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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

C-bus is a trademark of Corollary, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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.

Preface

This IBM Redbook presents an overview of IBM Tivoli Storage Manager (TSM) Version 5.1, giving detailed descriptions of the changes provided in this new release.

This book is intended for customers, consultants, IBM Business Partners, IBM and Tivoli staff who are familiar with earlier releases of Tivoli Storage Manager and who want to understand what is new in Version 5.1. It should be used in conjunction with the manuals provided with the product and is not intended to replace any information contained therein.

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.

Barry Kadleck is responsible for tape and Tivoli storage projects at the International Technical Support Organization, San Jose. He has a degree in Electronic Engineering and joined the IBM UK in 1985, working with storage software and hardware for the last 10 years. Before joining the ITSO in 2001, Barry worked at the Product Introduction Consultancy in Hursley, UK as a program manager, introducing new Tivoli Storage products into Europe.

Nigel Bentley is a Senior Specialist PSR working in the UK. He has over 26 years of experience in IBM Software Support. He holds a degree in Physics and Computer Science from York University. His areas of expertise include OS/390 SSD Storage Software products (catalogs are his speciality) and over nine years working in Technical Support with the ADSM and Tivoli Storage Manager product set. He is a Tivoli Certified Tivoli Storage Manager Consultant.

Michael Dann is a Storage Consultant with the Storage Network Solutions Team at IBM Germany. He has 14 years of experience in the IT field. He holds a degree in Computer Science from the University of Applied Science at Bingen am Rhein, Germany. Before joining IBM in 2001, he worked as a Technical Manager at a Business Partner. His areas of expertise include OS/2, WINTEL, networking and storage systems and four years working with the Tivoli Storage Manager product set.

Steven Pemberton is a Consultant and Trainer with Innovative Business Knowledge (IBK) from Australia. He has seven years of experience in System Administration, and is currently on the national board of the System Administrator's Guild of Australia (SAGE-AU). His areas of expertise include Tivoli Storage Manager, AIX/SP administration, and HACMP administration. Steven is a Tivoli Certified Consultant and Instructor for Tivoli Storage Manager and an IBM Certified Specialist in AIX System Administration. He is a co-author of the redbook *Samba: Installation, Configuration and Sizing Guide*, SG24-6004.

Bjorn Roden is an AIX Expert from Sweden working as a Technical Manager and Chief Programmer for the largest IBM @server pSeries/zSeries Business Partner in Sweden (Pulsen Systems). He has 12 years of experience with AIX and has worked with ADSM/Tivoli Storage Manager since the mid-1990s. He is currently certified as an AIX Advanced Technical Expert, Mid-Range Storage Technical Specialist, SP Specialist, HACMP Specialist, ADSM/Tivoli Storage Manager Specialist, and Webserver Specialist. He is a co-author of two redbooks *RS/6000 SP System Performance Tuning Update*, SG24-5340, and *AIX 5L Performance Tools Handbook*, SG24-6039.

Figure 0-1 The team (from left to right): Steven, Nigel, Barry, Björn, and Michael

Many thanks to the following people for their contributions to this project:

International Technical Support Organization, San Jose:

Charlotte Brooks
Deirdre Hewgill
Barry Mellish
Jon Tate

IBM Advanced Technical Support, San Jose:

Randy Larson

IBM Tivoli, San Jose:

Jason Basler
Christian Bolik
Mauro Cruciani
Colin Dawson
Jim Smith
John Viksne

IBM Tivoli, Tucson:

Matt Anglin
Craig Bell
Glen Hattrick
Harry Husfelt
Naomi Iseri
Erick Kissel
Bob LaBrie
Steve Luko
Howard Martin
Dave Morton

Rosa Plaza
Harley Puckett
Andrew Raibeck
BJ Scheid

IBM, Florida:

Jason Blackburn
Forrest Wilson

IBM, Germany:

Jens Fauth
Frank Mueller

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

`ibm.com/redbooks/residencies.html`

Comments welcome

Your comments are important to us!

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

- Use the online **Contact us** review redbook form found at:

  `ibm.com/redbooks`

- Send your comments in an Internet note to:

  `redbook@us.ibm.com`

- Mail your comments to the address on page ii.

Part 1

Overview

Part 1 gives a basic introduction to the new features provided with IBM Tivoli Storage Manager V5.1. Server and client specific features are covered in more detail in later chapters.

1

IBM Tivoli Storage Manager overview

This chapter introduces an overview of the changes made in IBM Tivoli Storage Manager Version 5.1 and contains the following major areas of change:

- Overview of enhancements, additions and changes
  - Server enhancements, additions and changes
  - Client enhancements, additions and changes
  - Tivoli Data Protection Agent enhancements, additions and changes
- Packaging changes
- Terminology changes
- Pricing model changes

1.1 Overview

IBM Tivoli Storage Manager ensures availability of business applications while driving improved return on storage management investment, by providing superior data protection and resource utilization that scales with business needs. Tivoli Storage Manager is an enterprise-class recovery solution, protecting your business critical data from the laptop to the zSeries, regardless of where it resides.

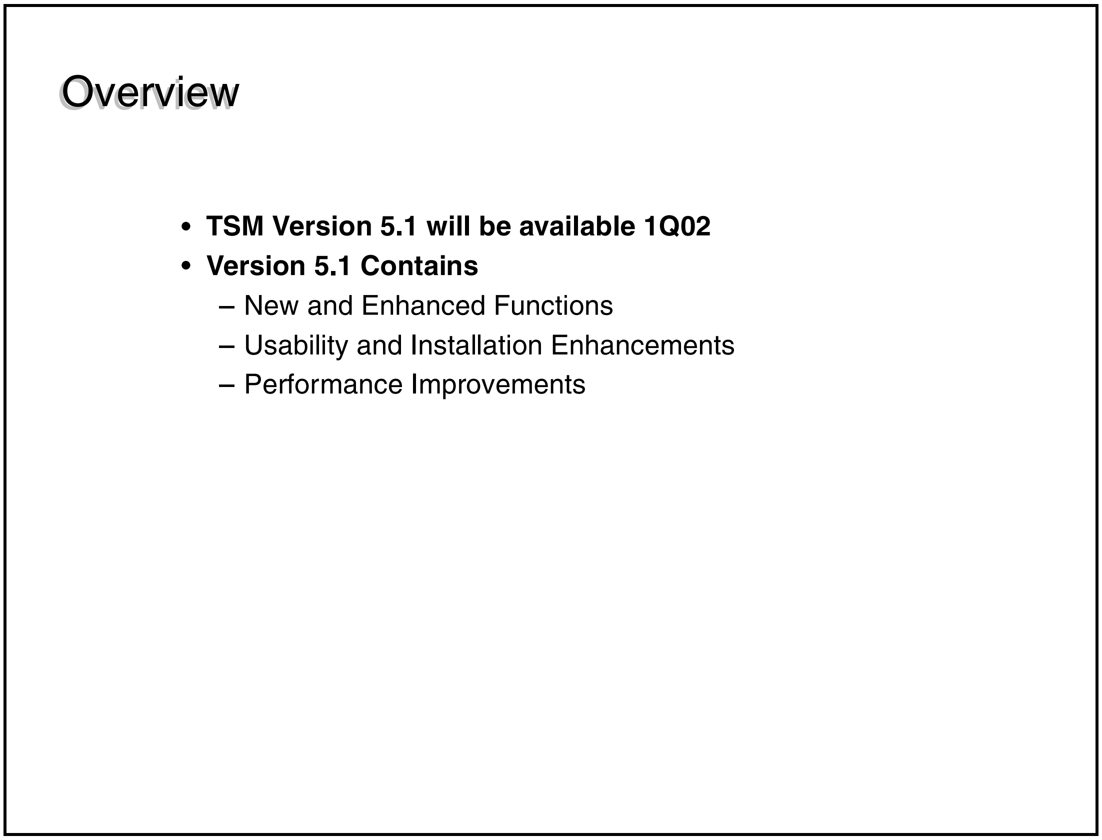

Figure 1-1 Tivoli Storage Manager V5.1 overview

Tivoli Storage Manager supports business continuance by helping to automate disaster recovery planning and recovery execution based on business priorities. Tivoli Storage Manager integrates the power of applications-aware technology in the recovery of leading database, content management, and workflow applications to ensure your entire business-process is protected.

1.2 Product positioning

IBM Tivoli Storage Manager is insurance for your data. If a disk drive fails, Tivoli Storage Manager restores your lost data. If you accidentally delete a file, Tivoli Storage Manager will bring it back. If your entire operation is hit by a disaster, Tivoli Storage Manager keeps your data safe. Tivoli Storage Manager insures your data by copying active online data to offline and off-site storage, even from 24x365 applications. If your data growth exceeds your storage capacity, Tivoli Storage Manager can effectively extend your capacity.

Tivoli Storage Manager moves inactive data from online storage to less expensive offline or nearline storage. Tivoli Storage Manager insures the data on any number of computers running any of more than a dozen different operating systems from laptops to mainframes.

1.2.1 IBM Tivoli Storage Manager capabilities

Figure 1-2 lists the capabilities of IBM Tivoli Storage Manager. This section offers a brief overview.

Figure 1-2 Capabilities

Backup-Restore

Complete data protection starts with data backups. Backups are a copy of your active online data stored on offline storage. Should an online storage device fail, a data error occur, or someone accidentally delete a file, the offline copy of that data can be copied back to online storage restored. Tivoli Storage Manager is famous for its extremely efficient backup methods. Tivoli Storage Manager uses multiple techniques to reduce data transfer sizes to the minimum possible. These techniques reduce the total time required for both data backups and more importantly, data restores.

Disaster preparation and recovery

Complete data protection also involves disaster preparation and recovery. Local copies of data will protect against discrete failures or errors in equipment, storage or people. But disasters tend to happen to entire facilities, not just a portion of the equipment inside those facilities.

Tivoli Storage Manager will prepare an additional copy of your active data for safekeeping at an off-site location to provide that extra insurance against disasters. Should a disaster strike and destroy your online storage and computers, the off-site copy of your active data can be restored to new computers to get your business up and running quickly.

Archive-Retrieve

Tivoli Storage Manager goes beyond just data backups to include data archiving. Until online storage is as inexpensive as offline storage, archiving inactive data is an effective way to reduce your online storage costs. The cost of storing a gigabyte of data on online storage is typically thirty times more expensive than the same gigabyte on offline storage. Some percentage of your data is inactive it hasn't been accessed in weeks if not months. Does that inactive data need to remain on online storage? Tivoli Storage Manager will move that inactive data to offline storage freeing up online disk space for more important active data. Should you find that the inactive data you archived must be used again, Tivoli Storage Manager will retrieve that data for you.

Hierarchical Storage Management

Tivoli Storage Manager also includes a more automated version of archive called Hierarchical Storage Management or HSM. Like archiving, HSM removes data from online storage and puts it on less expensive offline storage. But unlike archive, HSM leaves a data-stub on online storage that shows the name of the removed file. With this stub, you can easily access the offline data albeit much more slowly than if it were online. Tivoli Storage Manager's HSM capabilities are automated and will watch online data-files to see how often they are used. If not opened for an administrator specified length of time, they will be removed to offline storage leaving only the data-stub behind. For many businesses with huge amounts of data (but not all of it needed online all the time), HSM is the best way to save on storage costs.

24 x 365 application protection

Some of today’s most business-critical application programs must remain on and active 24 hours a day for 365 days a year. Fortunately, the manufacturers of some of these important programs understand the need for data protection. Some applications have a built-in capability to control external data protection applications. By directly controlling the time and method by which a data transfer to offline storage occurs, the application can continue operating with no interruption.

1.2.2 IBM Tivoli Storage Manager advantages

Figure 1-3 lists the advantages of IBM Tivoli Storage Manager. This section offers a brief overview.

IBM Tivoli Storage Manager advantages

- Enterprise management
- Management automation
- Hardware support
- Data transfer
- Unique Differentiators

Figure 1-3 Advantages

Enterprise management

Tivoli Storage Manager is a client-server application. One Tivoli Storage Manager server handles hundreds of Tivoli Storage Manager clients. Tivoli Storage Manager's administration control is extremely flexible via a hierarchy of administrators each designated with different management authority levels and specific domains defining user groups or applications.

Multiple administrators can manage Tivoli Storage Manager simultaneously with full data integrity. Multiple Tivoli Storage Manager servers can be controlled from any Tivoli Storage Manager server or via a Web interface from any computer. Data restore requests can be initiated and controlled from the Tivoli Storage Manager client, greatly reducing workload for the administrator.

The heart of the Tivoli Storage Manager server is an integrated relational database. This database catalogs each data transfer backup, archive, HSM, etc. and the complete set of data for every Tivoli Storage Manager client.

If you tell Tivoli Storage Manager to do data backups every day, the database allows the restore of data from every day in the past for as far back as you specify. Tivoli Storage Manager lets you go back in time to previous versions of your data files.You specify the frequency of backups, where the copied data is held (magnetic disk storage pool, optical disk, data tape, etc.), how long it is held there until migrating the next storage type, and how long each version is kept before it is erased and the space it used is made available for new data copies. You also specify how often data tapes should be tested for data integrity what percent of empty should a tape be before its data is consolidated to another tape (freeing complete tapes for reuse), and how quickly should data from single users be migrated from multiple tapes to a single tape to speed data restores.

Management automation

All the management capabilities of Tivoli Storage Manager can be automated. Automation control is via an extremely granular policy engine. This engine can be configured down to the individual file level to cover the who, what, where, when, and how of data transfers. Who has access to this file? What data should be transferred? Where should it be transferred to? When should Tivoli Storage Manager check to see if the file has changed and must be transferred again? How should the transfer occur? With Tivoli Storage Manager's policy-based automation engine configured and enabled, Tivoli Storage Manager really is a set-it and forget-it application. Tivoli Storage Manager reports what it has done and alerts you to anything out of the ordinary.

Hardware support

The best automation and management features for data protection mean little if they do not work with the computers you have. Tivoli Storage Manager's client software supports seventeen different operating systems. Tivoli Storage Manager's server software runs on eight operating systems. Platforms vary from laptop computers all the way up to mainframe computers.

Any server can work with any client all cross platform Solaris with Windows, AIX with Solaris, O/S390 with HP-UX any combination. Tivoli Storage Manager supports more than 400 offline storage devices including optical disk and tape.

To make sure your backups and restores can happen as efficiently as possible, Tivoli Storage Manager also supports multiple network types: dial-up, the Internet, WAN, LAN and SAN are each maximized for the data bandwidth they are capable of. Tivoli Storage Manager is also an open product. Its Application Programming Interface (API) is available to anyone through the Tivoli Ready program.

Data transfer

Even with the largest data pipes connecting servers and clients, Tivoli Storage Manager minimizes the amount of data going over those pipes. Because of Tivoli Storage Manager's integrated relational database, Tivoli Storage Manager is smart:

- Smart enough to know that if a file hasn't changed, it doesn't need to be backed up again - ever.
- Smart enough to know that if only a few bytes of a file have changed, only those bytes must be backed up again.

For those data transfers that include many bytes, Tivoli Storage Manager is optimized for disk-to-disk transfers. Instead of transferring data from the disk of a client computer, over the network, directly to a tape, Tivoli Storage Manager will instead make the first transfer to a disk storage pool. After an administrator-specified length of time (hours, days or weeks) the data on that storage pool is migrated to other less-expensive offline storage devices. This technique allows more data to be moved faster, reducing the total time of all data transfers both from and to clients. Disk-to-disk provides a data transfer technique better matched to the higher bandwidths of SANs and other Gigabit networks.

At the other of the bandwidth spectrum, if a dial-up or Internet connection is dropped, Tivoli Storage Manager is smart enough to pick up where it left off. Instead of starting a transfer all over again when the connection is restored, Tivoli Storage Manager continues from where it previously ended.

Unique differentiators

Tivoli Storage Manager has two unique features - its incremental backups that only backup the files that have changed since the last backup, and its ability to effectively utilize the disk and tapes in its storage hierarchy. The combination of these two mean less data being sent over your network, less tape drives needed to write the data and fewer tapes needed to store that data.

Tivoli Storage Manager does more than just backup your clients data - it has numerous products that integrate directly with the centralized server. These products allow you to backup your application databases, help you generate a plan for recovering from a disaster, move seldom used files off of clients local disk to near line storage, and generate in-depth reports to help you analyze and forecast your backup procedures.

Tivoli Storage Manager is great for the administrators, because there is so much flexibility in how you set up your backup policies, storage hierarchies, schedules and off-site disaster protection. In addition there is great flexibility in where you administer the product from as Tivoli Storage Manager can be administered over the Web from wherever you are.

LAN free to disk utilizing SANergy technology so as to send backups to disk via the SAN.

1.2.3 What’s new in IBM Tivoli Storage Manager Version 5.1

In IBM Tivoli Storage Manager Version 5.1 enhancements and additions have been made in both server and client functions and the number of supported platforms has increased. Major areas of change apply to disaster recovery functionality and performance, availability and throughput performance and server and client functionality, usability and platform availability of advanced functionality.

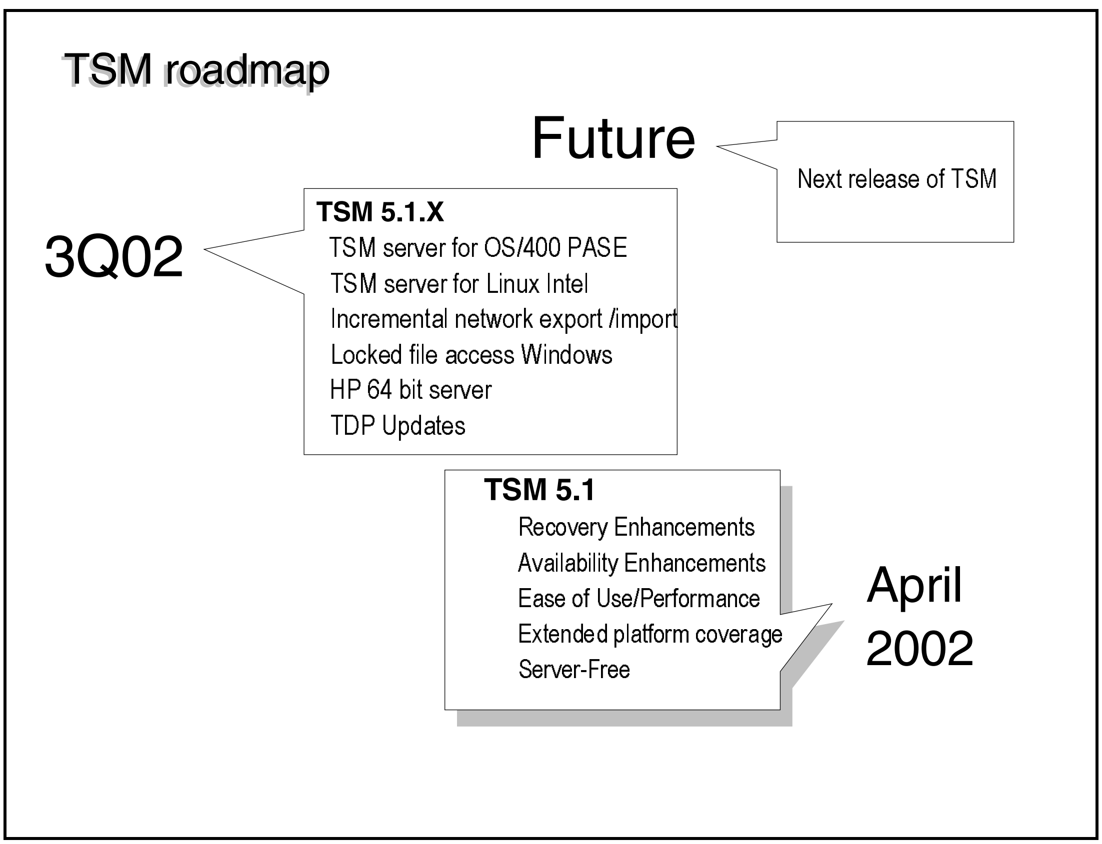

Figure 1-4 Tivoli Storage Manager roadmap

The functional changes in IBM Tivoli Storage Manager Version 5.1 can be grouped in the following areas:

- Improved disaster recovery
- Improved availability
- Other enhancements
- Improved ease of use and installation

Improved disaster recovery

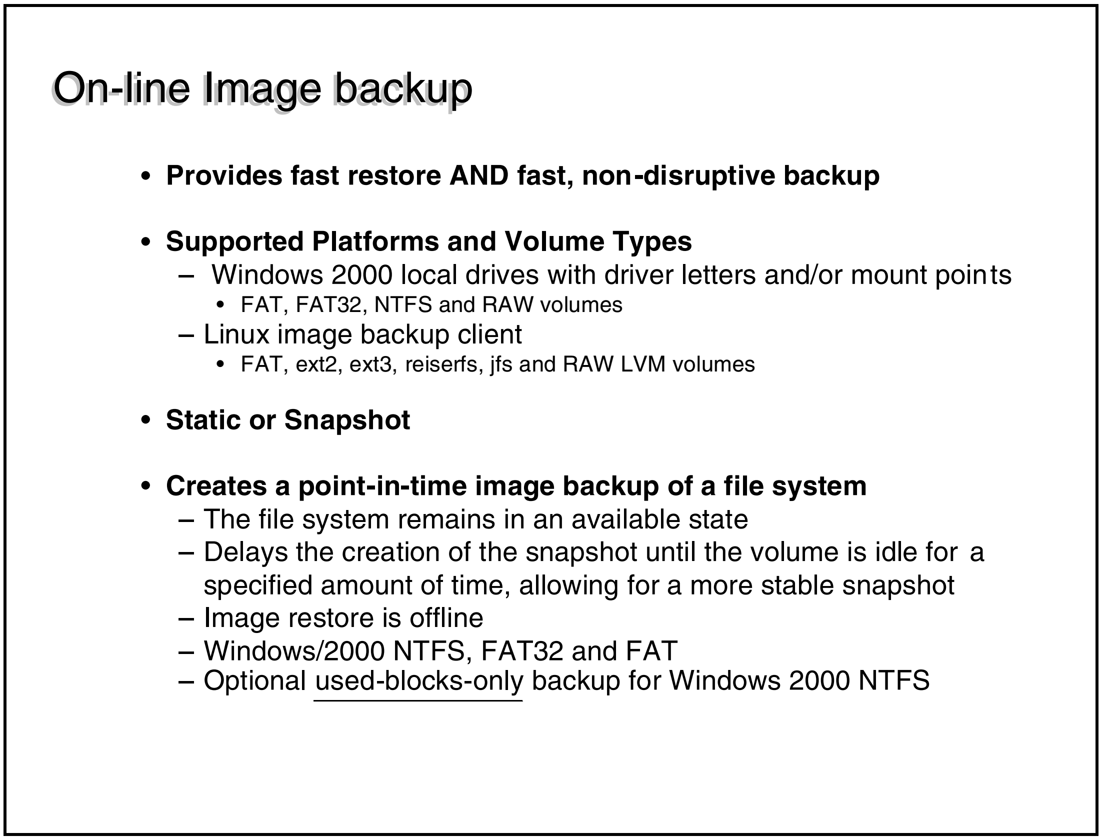

Figure 1-5 Online image backups

Online image backups

Recovery performance on Windows 2000 and Linux is enhanced by utilizing fast image restores. Image backups complement progressive incremental backups to provide full file system restores for disaster recovery or when a large percentage of a file system is to be restored.

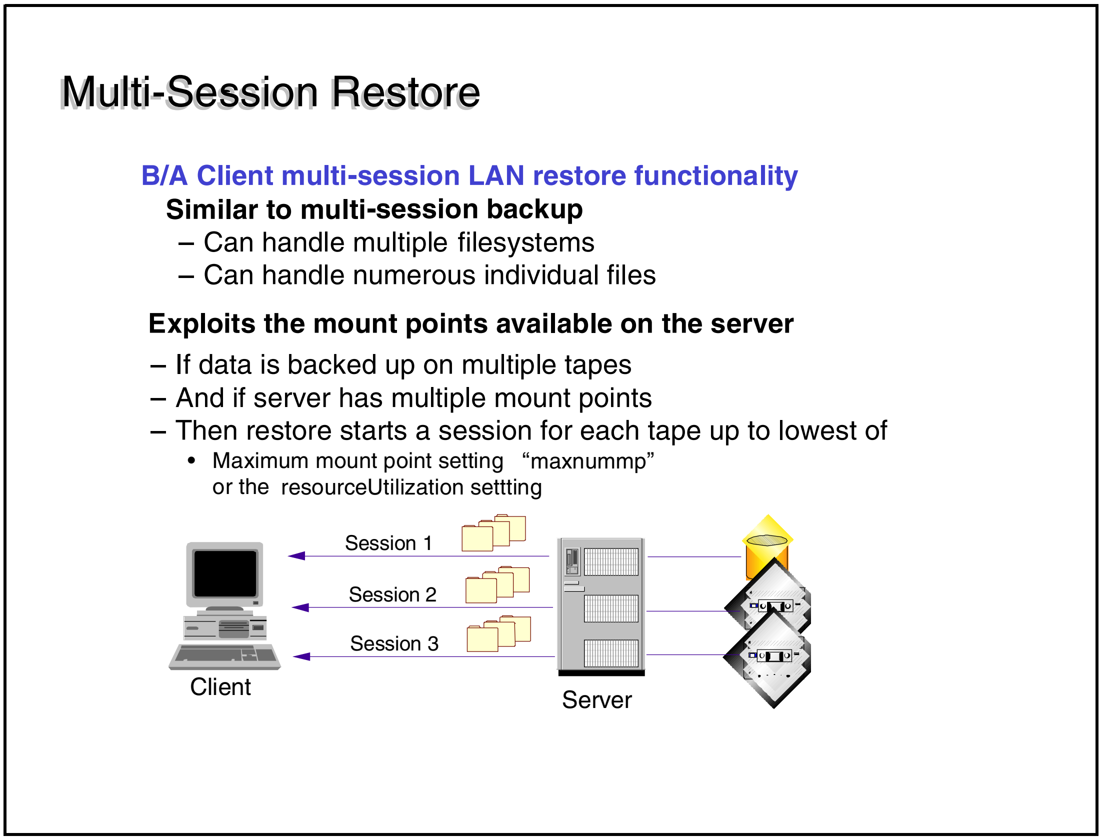

Figure 1-6 Multi-session restore

Multisession restore

Restore time is reduced when multi-session restore sessions are used with the backup-archive clients. This enable restores from tape to run in parallel when data is stored on multiple tapes and the tape drives needed are available at restore time. The number of tape drives that are available can be configured by the administrator.

More new functions

Move NODEDATA

- Moves node’s data from volumes in a sequential-access storage pool
- Allows for consolidating data for that node within the storage pool
  - **Reduces the number of volume mounts required for a restore**
- Can move data into another storage pool
  - **Move data to a disk storage pool in preparation for a client restore**

Synchronous Write to Copy Storage Pools

- Simultaneous write of data to multiple (1 to 10) copy storage pools
- These storage pools can be any currently supported devices
- Reduces the amount of time needed to create a copy storage pool

Journal Based Backup Updates

- Journal-based backup fully supports Windows 2000 cluster configurations
- Alternate journal directory location
- Dynamic addition/removal of journaled file systems
- Setup wizard changes

Figure 1-7 More new functions

Move data by node

Restore time can be improved by staging data to disk or by consolidating data on tape prior to a restore. The data movement can be done with the **`Move Nodedata`** command for data at a client node or filespace level.

Synchronous writes

The time to create recovery media for local and off-site vaulting is reduced with simultaneous writes to multiple copy storage pools. Duplicate copies of data for disaster recovery can now be created concurrently when backing up or archiving data, provided that there are sufficient mount points for each of the primary and copy storage pools. This can reduce the need for the sequential procedure of first completing the backup or archive and then creating the copies for vaulting.

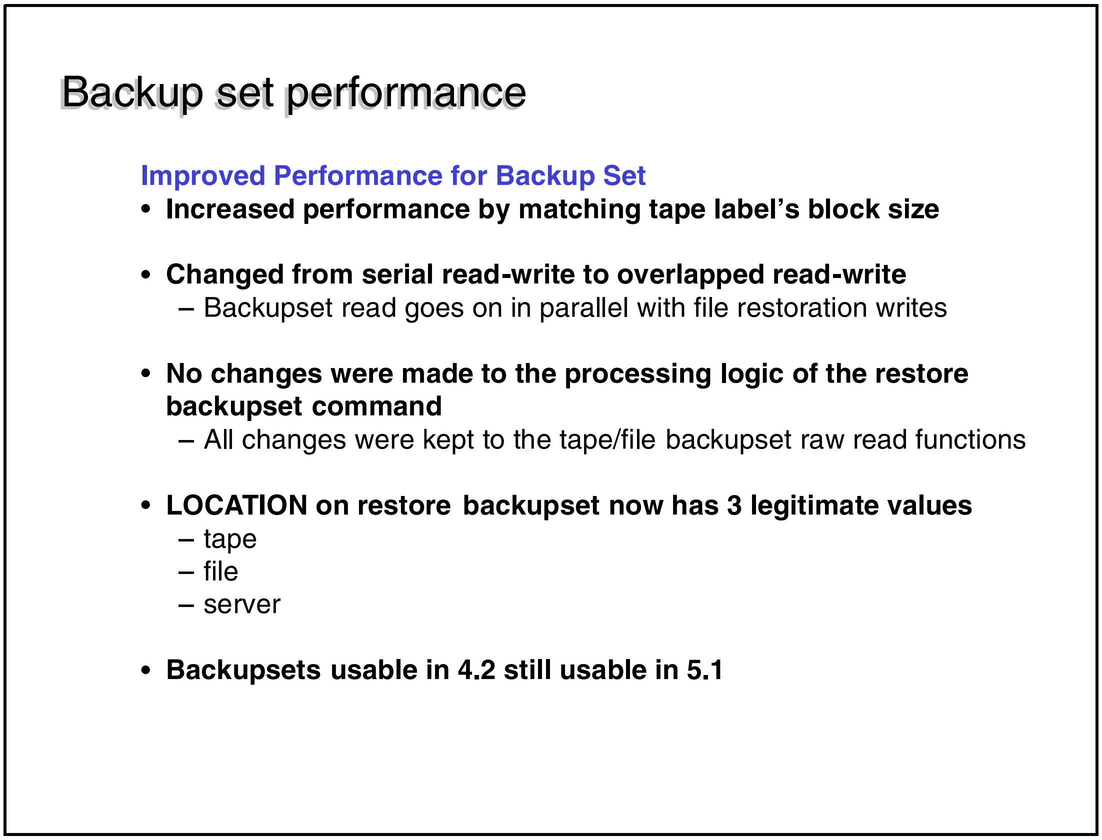

Figure 1-8 Backupset performance

Backup sets
Restore time from backup sets has been improved. Backup sets which can be used to restore data to backup-archive clients locally without a network or server, optimize the use of device block sizes to improve performance. Device support for backup sets has been expanded with the addition of LTO drives

Disaster Recovery Manager
Disaster Recovery Manager which helps you maintain business continuance is now included in enterprise edition.

Improved availability

On-line Image backup

- **Provides fast restore AND fast, non-disruptive backup**
- **Supported Platforms and Volume Types**
  - Windows 2000 local drives with driver letters and/or mount points
    - FAT, FAT32, NTFS and RAW volumes
  - Linux image backup client
    - FAT, ext2, ext3, reiserfs, jfs and RAW LVM volumes
- **Static or Snapshot**
- **Creates a point-in-time image backup of a file system**
  - The file system remains in an available state
  - Delays the creation of the snapshot until the volume is idle for a specified amount of time, allowing for a more stable snapshot
  - Image restore is offline
  - Windows/2000 NTFS, FAT32 and FAT
  - Optional used-blocks-only backup for Windows 2000 NTFS

Figure 1-9 Online image backups

Online image backups

Application availability is improved during backups by providing online, non disruptive image backups. Applications remain available while online image backups create a point-in-time image backup of the file system. Image backups complement progressive incremental backups and can be used for faster recovery of large amounts of data in file systems with large numbers of files.

Online image backups are available for Windows 2000 for FAT, FAT32, NTFS, and RAW volumes and for Linux for Ext2, ReisterFS, and RAW LVM volumes.

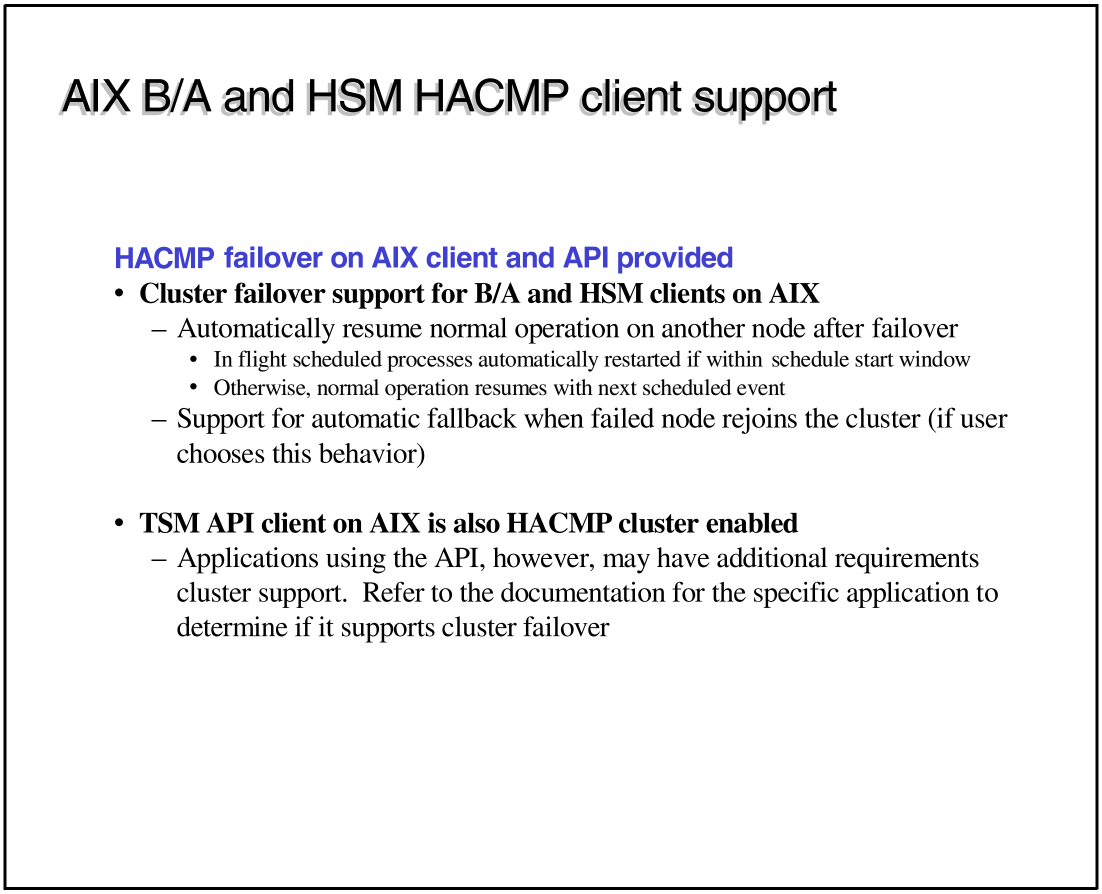

Figure 1-10 AIX backup archive client HACMP

HACMP

Support for high availability is provided with HACMP failover for the AIX backup-archive, HSM and API clients. HACMP failover automatically resumes normal operation on another node in the cluster after failover. Scheduled processes which have not completed are automatically restarted if within the schedule start window. Otherwise, normal operation resumes with the next scheduled event. Support for automatic fallback when failed node rejoins the cluster is also provided as an option.

Server-free Data Movement

- **Benefits**
  - Reduce TSM client and/or server CPU utilization
  - Eliminates data movement on the LAN
    - Implementation of the SCSI-3 extended copy command will cause data to be transferred directly between devices over the SAN or SCSI bus
  - Improve performance for data movement operations
  - Improve scalability
    - The TSM server will be able to handle more concurrent client connections and server operations because it will not be copying data
    - The TSM client processor will likewise be able to handle more application load because cycles will not be needed to read and send data buffers to the server
- **On-line server-free image backup Supports**
  - Windows NT/2000 TSM server
  - Windows 2000 B/A clients only
  - Static or Snapshot Image Backup
  - NTFS only
  - No support for Windows NT 4.0 or Windows XP b/a clients
  - IBM SAN Data Gateway or Pathlight 5000

Figure 1-11 Server free data movement

Server-free data movement

The impact on application servers during backups is minimized with server-free data movement. Server-free data movement reduces CPU utilization by removing backup and restore processing from your production machines.

Similar to LAN-free, server-free also lowers traffic on your LAN by off loading your backups and restores to the SAN. Scalability on the IBM Tivoli Storage Manager server is improved and it can handle more concurrent client connections and server operations because the server will not be copying data.

The IBM Tivoli Storage Manager client processor will likewise be able to handle more application load because cycles will not be needed to read and send data buffers to the server. Server-free will be supported initially on Windows 2000 and Windows NT IBM Tivoli Storage Manager servers and Windows 2000 backup-archive clients.

An IBM San Data Gateway will be utilized to perform the SCSI 3 copy command that facilitates the movement of the data between client disk and the servers' tape drives.

Other enhancements

Journal based backup

- The improved backup performance available with journal based backup is now supported on cluster configurations.
  - Remove query of active files on server - Thus reduce data transfer from server to client
  - Remove scan of local filesystem
  - Eliminate list processing
    - Reduce memory requirement
- Implemented on Windows NT/2000 using Win APIs
- Eliminate most queries to either the TSM server or the local file system
  - TSM monitors file changes within filesystems
  - Only new/changed files are backed up
- Fault tolerant
  - Automatically reverts to normal incremental backup if fault detected
- Dynamic addition/removal of files systems to be journaled
- Setup wizard changes

Figure 1-12 Journal based backups

Journal-based backup

Improved backup performance available with Journal-based backup is now supported on Windows cluster configurations. Journal-based progressive incremental backup has been shown to reduce backup elapsed time in many cases over standard progressive incremental backup by four times in file systems with an average daily change rate of four percent. Environments with smaller daily change rates have shown more dramatic improvement.

Actual performance results may vary significantly depending on factors such as system workload and configuration. The usability of journal-based backups is also improved by providing capabilities to dynamically add or remove file systems that will be backed up using journal based backup without recycling the journal service.

The setup wizards make theses new options easy to use. The usability of journal- based backups is also improved by providing capabilities to dynamically add or remove file systems that will be backed up using journal based backup without recycling the journal service. The setup wizards make theses new options easy to use.

AIX server direct I/O

AIX Server Performance – direct I/0

- **Helps reduce CPU utilization and improves I/O throughput**
  - Previously when using JFS file systems for DB, LOG, and STORAGE POOL volumes on AIX, the AIX Virtual Memory Manager (VMM) caches read/write requests in memory in case sequential blocks are read/written
    - In most cases, TSM server does not do sequential I/O, and if it does, it writes enough data at a time to not need the caching
  - Now TSM reduces CPU utilization by bypassing the VMM cache when writing to TSM JFS volumes
    - In AIX V4.3, a new option is available called "DIRECT I/O".
    - When a JFS volume is opened with this option, the VMM cache will be bypassed in most cases
- **AIXDIRECTIO option is added to the server options**

Figure 1-13 AIX server direct I/O

Direct I/O

Improved throughput and reduced CPU utilization on AIX servers is made possible with Direct I/O. Direct I/O enables the bypassing the AIX VMM cache when writing to JFS volumes, under 2 GB, used by IBM Tivoli Storage Manager for its database, recovery log, and storage pools.

Asynchronous I/O

Improved throughput and reduced I/O waiting on AIX servers is made possible with asynchronous I/O. asynchronous I/O enables writing to DASD by IBM Tivoli Storage Manager for its database, recovery log, and storage pools to be performed without having to wait for completion before additional writing can be initiated by the server.

Server and client Checksum

Checksum (CRC) support

- **Benefits misconfigured or malfunctioning network or storage hardware.**
- **Additional assurance to the customer that data was successfully handled.**
- **Verifies the data that is correctly transferred from the client, stored on the server, and returned to the client.**
  - New communication and SAN hardware products are more susceptible to data loss, thus the need for checksums
  - The function may also be used in problem determination to isolate possible data transfer problems
  - Optional function
  - Places checksums or CRCs in the data sent from the client and stored on the disk or tape media at the server
    - the data can later be checked during restore, retrieve, or recal operations
    - check done both at the TSM server and at the TSM client
- **Externals updated to allow for turning CRC on and off**

  REGISTER NODE, UPDATE NODE,
  DEFINE STGPOOL, UPDATE STGPOOL and AUDIT VOLUME

Figure 1-14 Server and client checksum

CRC Tivoli Storage Manager now offers optional cyclic redundancy checking capability between the IBM Tivoli Storage Manager server and client. This allows IBM Tivoli Storage Manager to identify data corruption introduced either by the network or by errors within the storage environment. The use of cyclic redundancy checking in the exchange of information and data between the client and server will allow for the detection of, and recovery from, situations that the communication protocol may not have been able to handle. This is very useful in new communication and SAN hardware products that are more susceptible to data loss.

Improved ease of use and installation

IBM Tivoli Storage Manager's usability, installation, and configuration process has been made even easier with additional wizard support and addition of consistent client return codes and Web interface updates.

Minimal IBM TSM server Backup up to local disk in minutes with a new install of the minimal configurations for IBM Tivoli Storage Manager. A minimal IBM Tivoli Storage Manager on Windows server install option aids a new customer in getting backups and restores operational in a very short time.

Wizard updates

Minimal TSM Server

- **Option to produce a minimal TSM server configuration on Windows**
  - Can complete initial backup quickly
- **The minimal configuration option involves**
  - Present Initial Configuration Dialog with choice of Default or Minimal configuration
    - **The difference between the two configurations is whether the ser ver initialization wizard is the only wizard run or whether several wizards are run in addition to the server initialization wizard**

Server Free, Journal Based backup & NAS Wizard Support added

- **Server Initialization wizard updated**
- **Node configuration wizard updated**
- **Device Configuration wizard updated**

Figure 1-15 Wizard updates

Wizard options

The configuration process is simplified for using journal-based backup, NDMP based backup for Network Attached Storage (NAS) and for server-free data movement. Wizard options have been added to make set up for these configurations easier.

Consistent client return codes

Reliable, Consistent & Documented Return Codes

- **Command line client & scheduler enhanced**
  - Facilitates automation of client operations via user-written scripts
  - Administrators can now distinguish between
    - Scheduled backups that completed successfully with no skipped files

      AND
    - Scheduled backups that completed successfully with one or more skipped files
  - If PRESCHEDULECMD command ends with non-zero return code, scheduled event will (optionally) not run
    - Ensures that scheduled events will not run if prerequisite commands do not complete successfully

Figure 1-16 Consistent client return codes

Client return codes

Automation of IBM Tivoli Storage Manager functions has been improved and made more flexible. Automation can be accomplished using the internal scheduler or an external scheduling facility. Reliable, consistent, and documented return codes have been added to the command line client and the scheduler.

This facilitates automation of client operations via user-written scripts. Administrators can now distinguish between scheduled backups that completed successfully with no skipped files and scheduled backups that completed successfully with one or more skipped files.

Also, if the `preschedulecmd` command ends with non-zero return codes, the scheduled event will not run. This ensures that scheduled events will not run if prerequisite commands do not complete successfully.

Web client updates

Usability updates to Web

- **Appearance improved by updating the images in the GUIs**
  - The HUB, splash screens, and images updated in the TSM native GUIs for the Windows and Mac platforms and in the TSM Web Client GUIs
  - The UNIX GUI updated its splash screen
  - Toolbars and tool tip help will be added to the TSM Web Client GUI
- **TSM Web Client GUI is now a Swing-based web applet**
  - Uplifted the current TSM Web Client to use Swing-based components
  - Swing components have built-in accessibility support and improved keyboard navigation
  - Utilized the native look and feel of the platform running the browser

Figure 1-17 Web client updates

Web interface Accessibility support and improved keyboard navigation have been added as part of the update to the Web interfaces for IBM Tivoli Storage Manager. The native look and feel of the platform running the browser is preserved.

1.2.4 Platforms for IBM Tivoli Storage Manager LAN-free support

Table 1-1 lists the LAN-free platform support.

Table 1-1 LAN-free support

V5.1 Servers	Backup-Archive Clients	Tivoli Data Protection Solution Clients
AIX	AIX	R3, Oracle, Domino, Informix
SOLARIS	SOLARIS	R3, Oracle, Domino, Informix
WINDOWS	WINDOWS	R3, Oracle, Domino, Exchange, SQL
HP-UX	HP-UX	R3, Oracle, Informix

Note: For more specific information on LAN-free platform support, refer to:
http://www.tivoli.com/products/index/storage_mgr/storage_mgr_concepts.html

1.2.5 Extended platform and hardware support

IBM Tivoli Storage Manager's extensive support is extended to include additional operating system, hardware, file systems, and protocols:

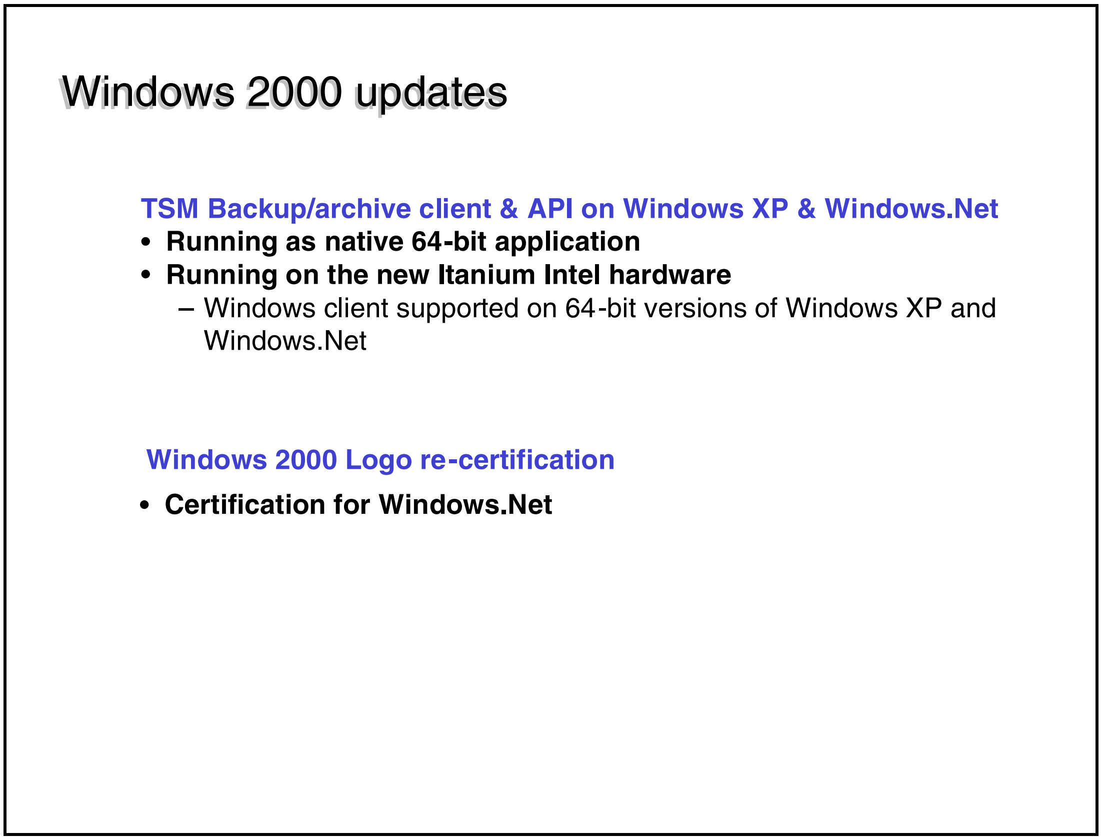

Figure 1-18 Windows 2000 client updates

WINDOWS XP and .NET The IBM Tivoli Storage Manager backup-archive client and API now support Windows XP and Windows.Net.

Macintosh client

MAC OS X support

- **As Apple puts it:**
  "Mac OS X is both a radical departure from previous Macintosh operating systems and a natural evolution from them. It carries on the Macintosh tradition of ease-of-use, but more than ever it is designed not only to be easy to use but a pleasure to use."

- **TSM Macintosh b/a client updated to run as a native Macintosh OS X client**
  – Provide backup support for both the UFS and HFS+ file systems
  – Run as a Mac OS X native executable
  – Support all file system types
  – Support new file naming conventions
  – Acquire the look and feel of a Mac OS X application
  – Utilize Apple help instead of Apple Guides

Figure 1-19 Macintosh client

MACINTOSH OS X

The IBM Tivoli Storage Manager Macintosh backup-archive client has been updated to run as a native Macintosh OS X client. It provides backup support for both the UFS and HFS+ file systems. It runs as a Mac OS X native executable and supports all file system types, new file naming conventions, and multiple users. The backup-archive client provides the look and feel of a Mac OS X application and utilizes Apple Help instead of Apple Guides.

Linux/86 support

On-line file system image backup using Linux's snapshot capability

Two new filesystems supported by Linux TSM client

- ext3
- jfs
- In addition to:
  - FAT, ext2, and reiserfs

Figure 1-20 Linux/86 support

LINUX

The IBM Tivoli Storage Manager Linux client now supports EXT3 and JFS file systems. An online file system image backup is also possible using Linux's own snap- shot capability for the EXT2, ReiserFS, and RAW LVM volumes.

NetWare client updates

NetWare 6 will be supported

Long name support for backup sets
- Backup sets can now support non-DOS names

Support for files more than 4Gb on NSS volumes
- Previously TSM only supported 32-bit architecture
  - **Limited the max file size to 4GB**
- Now support 4GB and larger files
  - **NetWare 5.x NSS file spaces**
  - **NetWare 6.x NSS file spaces**

Ability to backup/restore directory space limitations
- Novell NetWare 5 added more attributes to the files & directories
- Attributes handled through SMS will now be backed up by TSM

Figure 1-21 Netware client updates

NETWARE 6

Netware 6 is now supported by the IBM Tivoli Storage Manager Netware client. Files greater than 4 GB on NSS volumes, and the ability to backup and restore directory space limitations, are also now supported.

Sun client QFS support

B/A support for Sun QFS

- **Sun QFS**
  - High-performance file system that enables file sharing in a SAN
  - Eliminates performance bottlenecks of applications with very large file size
- **QFS Support**
  - TSM b/a client v5.1 for Solaris (2.6, 7, 8)
  - Version 3.5.0 LSCqfs it is the standalone version of QFS without using SAM
- **Limitations**
  - TSM Backup Image and Restore Image not supported
  - The actual version of QFS (3.5.0) does not support ACLs
  - QFS has two hidden system files & system directory that can't be backed up
    - A backup of these files is not needed
    - They contain internal data to manage the file system
    - This data will be automatically excluded from a backup and recreated automatically by the file system itself if a restore of files is invoked
  - The combination of QFS and SAM to archive files onto a tertiary background storage like tapes is not supported

Figure 1-22 Sun client QFS support

Sun QFS

The IBM Tivoli Storage Manager backup-archive client for Solaris (2.6, 7, 8) now supports the Sun QFS file system. Sun QFS is a high-performance file system that enables file sharing in a SAN.

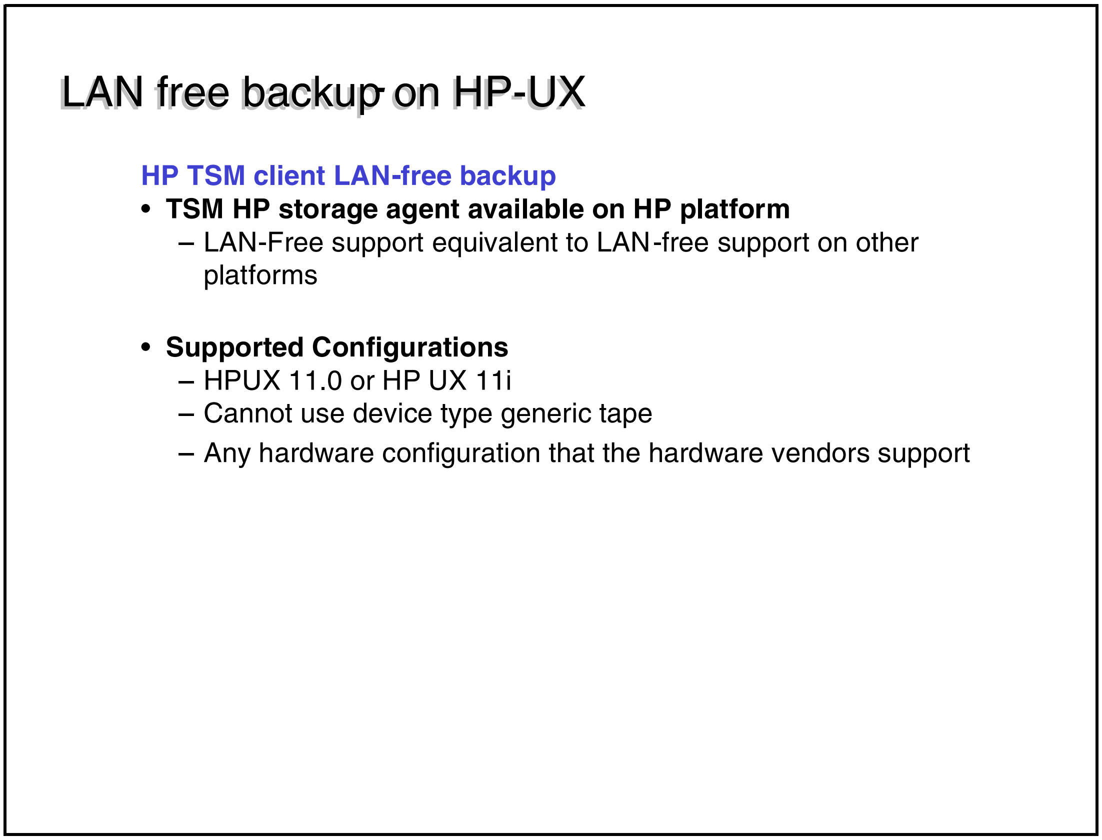

Figure 1-23 HP-UX LAN free backup

HP-UX The HP-UX clients now have LAN-free capability. This is addition to the LAN-free support available on AIX, Solaris, and Windows clients.

Tivoli Space Manager IBM Tivoli Storage Manager's HSM support which is now part of enterprise edition, supports AIX HACMP, 64-bit AIX 5.1, and 64-bit Solaris 8.

NDMP The NDMP support for NAS file servers by Tivoli Storage Manager is extended to the Tivoli Storage Managers on AIX and Solaris in addition to Windows. The NDMP support enables backup and recovery support on the Tivoli Storage Manager servers for network-attached storage (NAS) file servers from Network Appliance. The implementation of the NDMP server protocol enables the NAS file servers to be backup-ready and enables higher-performance backup to tape devices without moving the data over the LAN.

DISC IBM Tivoli Storage Manager now supports the DISC company's optical platter barcode. The IBM Tivoli Storage Manager device driver is able to get the barcode from each platter and provide the barcode to the application.

1.2.6 Server

These servers are supported in this environment.

HP-UX The HP-UX V5.1 server device support has been expanded, including, for the first time, support for optical devices.

IBM @server iSeries OS/400 PASE is an integrated runtime environment for AIX (or other UNIX-like) applications running on the iSeries system. It provides a broad subset of the Application Binary Interface (ABI) of AIX 4.3.3.

The IBM Tivoli Storage Manager 5.1 product includes a server capable of running in this environment. This server has a UNIX look and feel, but contains the necessary modifications to permit it to function in the PASE environment.

1.2.7 Tivoli Data Protection

Enhancements have been made to several of the Tivoli Data Protection components of the new IBM Tivoli Storage Manager data protection solutions.

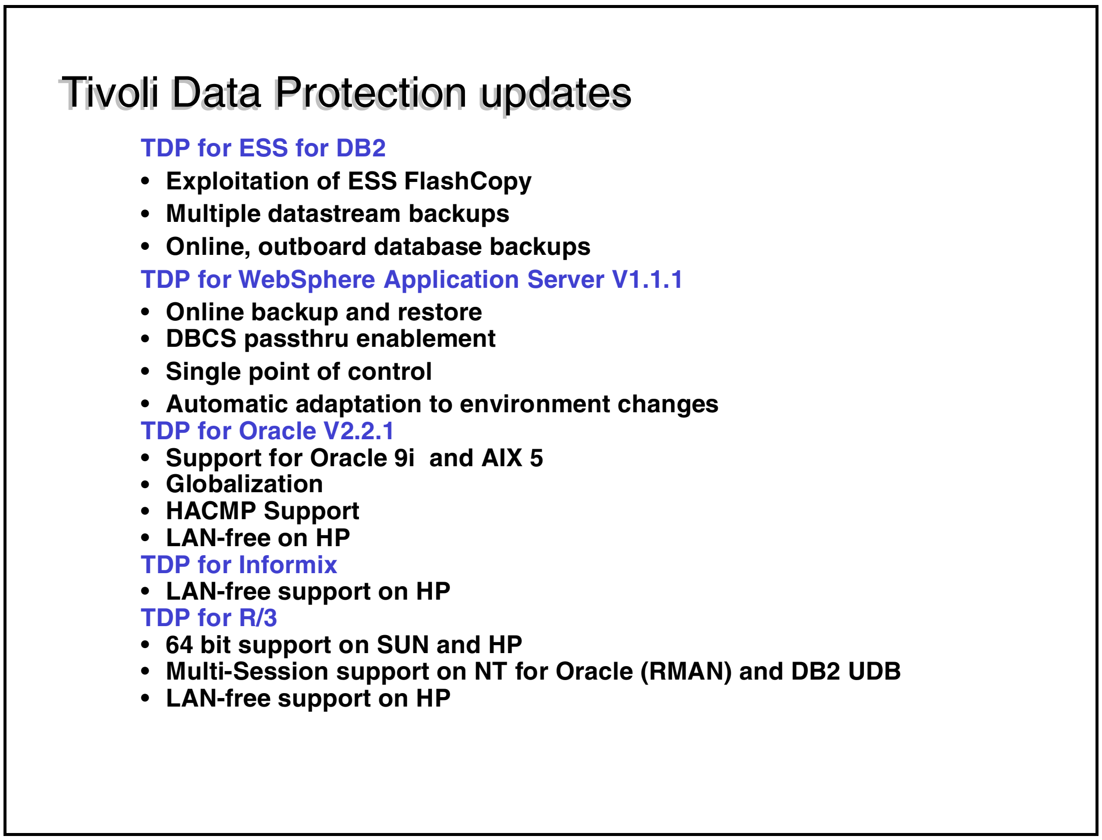

Figure 1-24 TDP updates

IBM Tivoli Storage Manager for Application Server Tivoli Data Protection for WebSphere Application Server V1.1.1 offers online backup and restore, DBCS pass through enablement, single point of control, and automatic adaptation to environment changes.

IBM Tivoli Storage Manager for Hardware Tivoli Data Protection for Enterprise Storage System for R3 now supports DB2.

IBM Tivoli Storage Manager for Mail Tivoli Data Protection for Lotus Domino is now supported on OS/400

IBM Tivoli Storage Manager for Databases Tivoli Data Protection for Oracle V2.2.1 supports Oracle 9i, AIX 5, Globalization, HACMP failover, and LAN-free on HP. Tivoli Data Protection for Informix supports LAN-free on HP.

IBM Tivoli Storage Manager for ERP Tivoli Data Protection for R/3 now supports 64-bit on Sun and HP, multi-sessions on Windows NT for Oracle (RMAN) and DB2 UDB, and LAN-free support of HP.

1.3 Packaging changes

The Tivoli portfolio of storage management products is realigned to more effectively meet your needs. Products that manage the same functional area of your environment are combined into a single product offering. Over time, functional integration, simplification of installation processes, and new enhancements will add to the power of newly combined products.

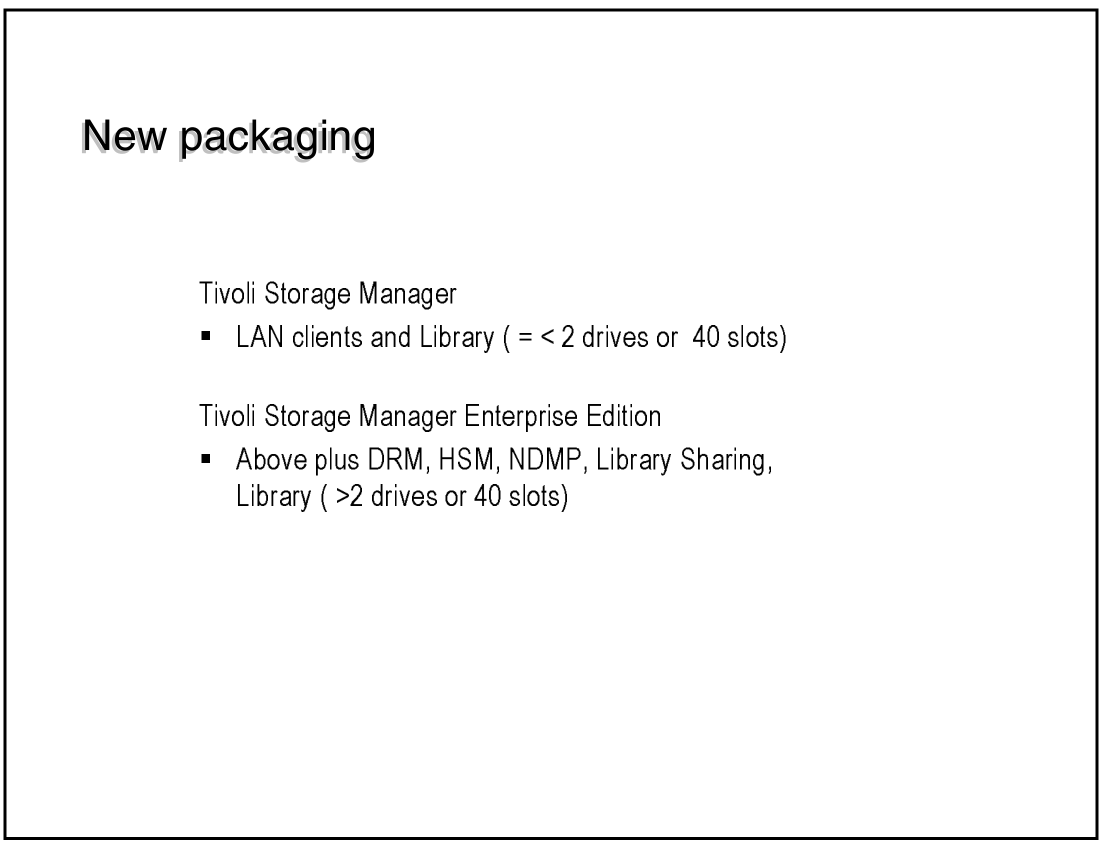

Figure 1-25 New packaging

There are two IBM Tivoli Storage Manager V5.1 offerings which can be chosen based on the business needs:

IBM Tivoli Storage Manager For basic backup-archive over a LAN.

IBM Tivoli Storage Manager Enterprise Edition Enables you to exploit advanced functions including, LAN-free backup-restore to tape and disk, library sharing, space management, disaster recovery management and NDMP support.

Note: Please read IBM Tivoli Storage Manager V5.1 announcement letter for information about products that will be discontinued based on the Tivoli end of support policy.

1.3.1 IBM Tivoli Storage Manager

IBM Tivoli Storage Manager is a multi-function storage software product that addresses the challenges of complex storage management across distributed environments. It protects and manages a broad range of data, from the laptop to the corporate server environment. Over 30 different operating-system platforms are supported. With IBM Tivoli Storage Manager you can obtain:

- Centralized backup and archive of your company's data, one of your company's most important assets
- Centralized administration for data and storage management
- Efficient management of information growth
- High-speed automated data recovery
- Compatibility with hundreds of storage devices, as well as LANs, and WANs

1.3.2 IBM Tivoli Storage Manager Enterprise Edition

IBM Tivoli Storage Manager Enterprise Edition builds upon IBM Tivoli Storage Manager's backup and data management to produce a completely integrated storage management solution that fits the needs of enterprise customers.

The Enterprise Edition expands the IBM Tivoli Storage Manager offering to include Managed Systems for SAN, which allows LAN-free backup and recovery using tape, the ability to do LAN-free backups and recovery using disk, tape library sharing, and the zSeries Tivoli Storage Manager Linux Client (SuSE OS and e390 Hardware).

If you want to use a library that has more than two drives or more than 40 tape slots, then IBM Tivoli Storage Manager Enterprise Edition is required.

The functions that were in Tivoli Disaster Recovery Manager, Tivoli Space Manager, and Tivoli Data Protection for NDMP, are included in IBM Tivoli Storage Manager Enterprise Edition.

- The Tivoli Disaster Recovery Manager function helps you maintain business continuance by:
  - Establishing and helping to automate a thorough disaster recovery plan
  - Automating vital recovery steps to bring your business back to normal
  - Managing and identifying off-site media needed for recovery
  - Tracking and reporting systems destroyed, in event of disaster
  - Performing restores in order of priority
- The Tivoli Space Manager function uses hierarchical storage management (HSM) to automatically and transparently migrate rarely accessed files to Tivoli Storage Manager storage while the files most frequently used remain in the local file systems.

By migrating rarely accessed files to the server storage, Tivoli Space Manager frees administrators and users from manual file system pruning tasks by enabling you to have sufficient free storage at your workstation or file server, deferring the need to purchase additional disk storage.

The Tivoli Space Manager function is supported on the AIX and Sun Solaris platforms.

- The Tivoli Data Protection for NDMP function provides backup and recovery support on Tivoli Storage Manager servers for network attached storage (NAS) file servers from Network Appliance.

  NAS file servers often require a unique approach to providing backup and recovery services, because these file servers typically will not run third-party software.

  Tivoli Data Protection for NDMP utilizes the Network Data Management Protocol (NDMP) to communicate with and provide backup and recovery services for NAS file servers.

  NDMP is an industry-standard protocol that allows a network storage-management application to control the backup and recovery of an NDMP-compliant file server without installing third-party software on that server.

  The implementation of the NDMP server protocol enables the NAS file servers to be backup-ready and enables higher-performance backup to tape devices without moving the data over the LAN.

1.3.3 New IBM TDP packaging

The TDP agents have been given new names in Tivoli Storage Manager V5.1.

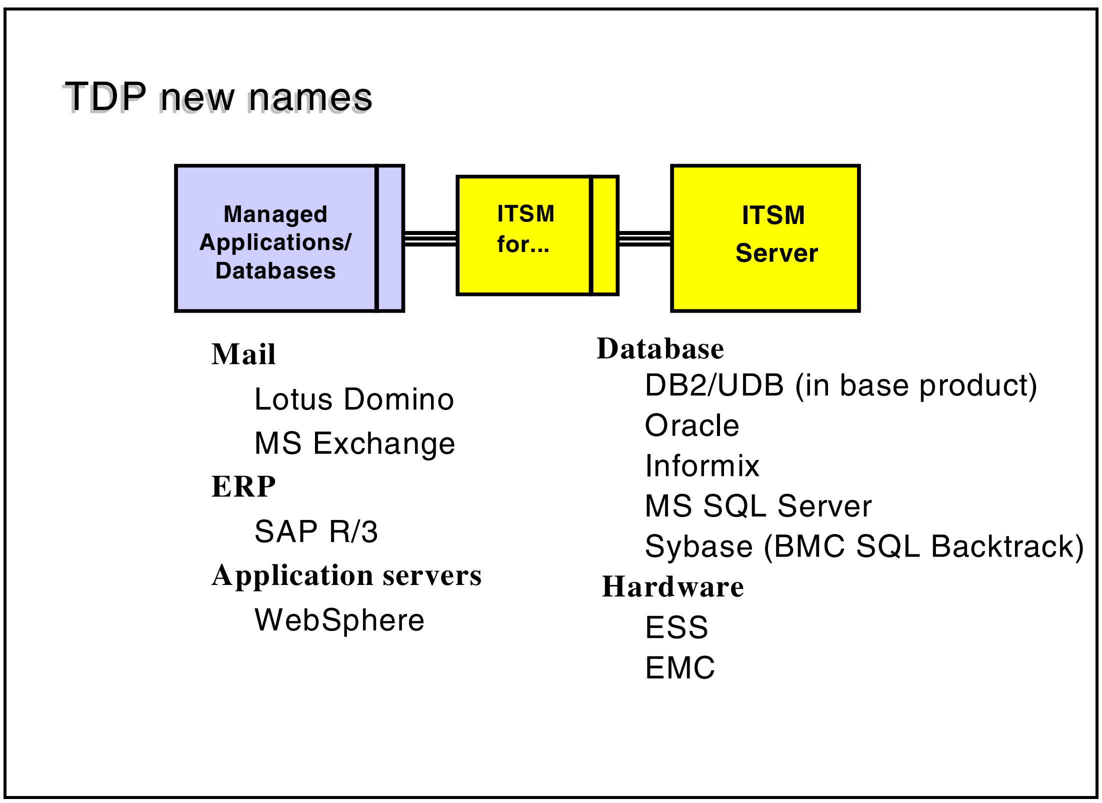

Figure 1-26 New names for TDP agents

- TDP for Domino and TDP for MSX are now IBM TSM for Mail
- TDP for Oracle Informix and MS SQL are now IBM TSM for Database

- TDP for SAP R/3 is now IBM TSM for ERP (Enterprise Resource Planning)
- TDP for WebSphere is now IBM TSM for Application servers
- TDP for ESS or EMC Symmetrix is now IBM TSM for Hardware

IBM Tivoli Storage Manager for Mail

IBM Tivoli Storage Manager for Mail complements IBM Tivoli Storage Manager to ensure that Lotus Domino and Microsoft Exchange data is safe and secure no matter where it is or how it's stored. It supports incremental and full backup as online (hot) backup and exploits the backup certified utilities and interfaces provided by Lotus Domino and Microsoft Exchange.

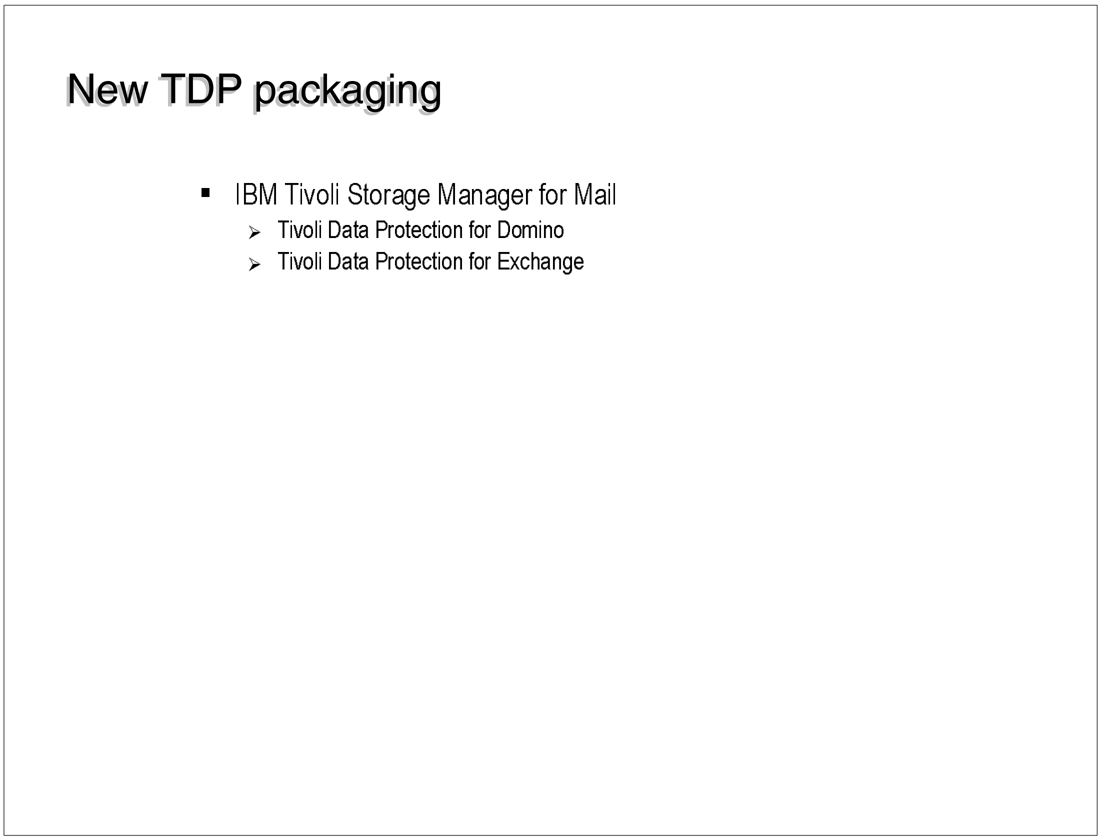

Figure 1-27 IBM Tivoli Storage Manager for Mail

This maximizes the protection of the data and the performance of the backups and restores. Archive and partial off load of mail databases and retrieve of individual mail or attachments is provided by IBM Content Manager CommonStore for Lotus Domino and IBM Content Manager CommonStore Exchange Server V7.2 which are complementary products to Tivoli Data Protection for Mail.

For more information refer to: `http://www-4.ibm.com/software/data/commonstore/`

IBM Tivoli Storage Manager for Application Servers

IBM Tivoli Storage Manager for Application Servers works in conjunction with IBM Tivoli Storage Manager server to safeguard the WebSphere Application Server environment, without disrupting the availability of the application.

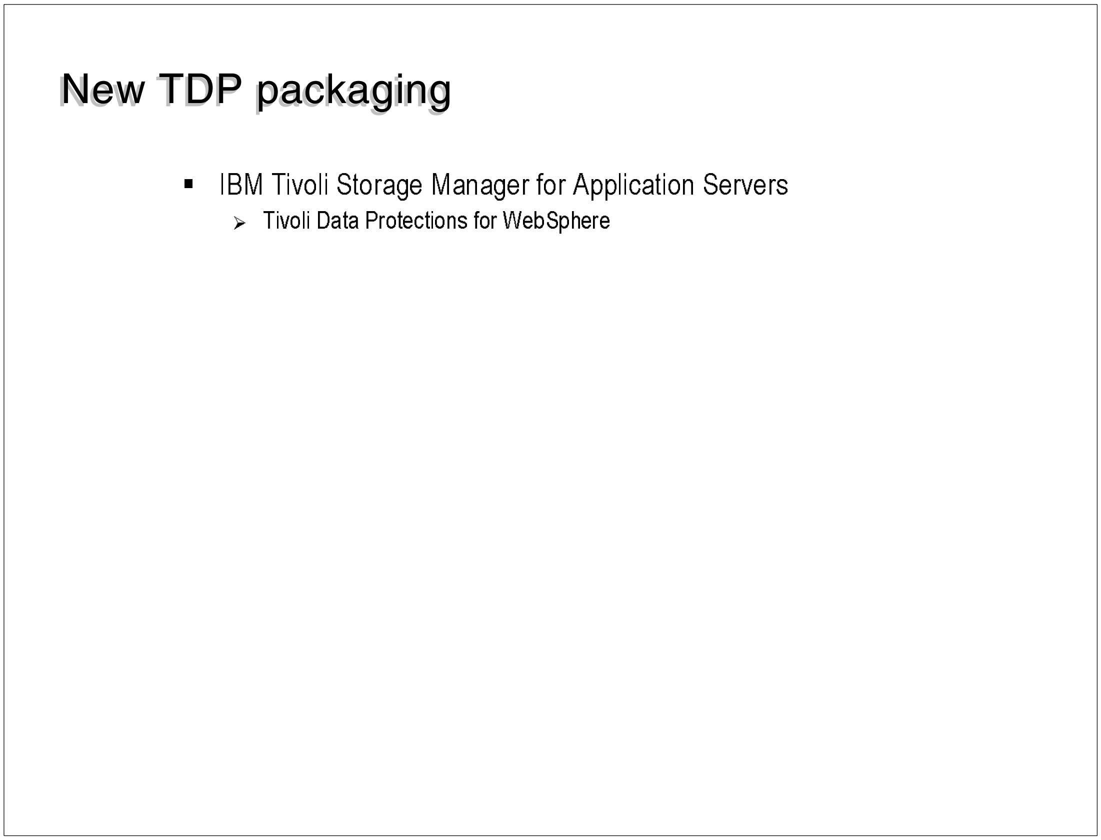

Figure 1-28 IBM Tivoli Storage Manager for application servers

It offers reproducible backup and restore of WebSphere Application Server (WAS) environments, including the WebSphere administration database, configuration data, and deployed application program files.

Changes to the WebSphere environment, such as the addition of applications are automatically detected and included in the backup to help keep backup data up-to-date.

IBM Tivoli Storage Manager for Databases

IBM Tivoli Storage Manager for Databases complements IBM Tivoli Storage Manager to ensure that Informix, Oracle, and Microsoft SQL data is safe and secure no matter where it is or how it's stored. It supports incremental and full backup as online (hot) backup and exploits the backup certified utilities and interfaces provided by Informix, Oracle, and Microsoft SQL. This maximizes the protection of the data and the performance of the backups and restores. IBM DB2 Universal Database is enabled to backup to IBM Tivoli Storage Manager and does not require IBM Tivoli Storage Manager for Databases (Figure 1-29).

New TDP packaging

- IBM Tivoli Storage Manager for Databases
  - Tivoli Data Protection for Oracle
  - Tivoli Data Protection for Informix
  - Tivoli Data Protection for MS SQL Server

Figure 1-29 IBM Tivoli Storage Manager for databases

IBM Tivoli Storage Manager for ERP

IBM Tivoli Storage Manager for ERP and its Administrative Assistant work with IBM Tivoli Storage Manager to protect your vital SAP R/3 system data. The Administration Assistant is used to assist in configuration, monitoring and administration of Tivoli Data Protection for R/3. This powerful solution enables SAP R/3 system administrators to manage the large volumes of data involved in system operations more efficiently, consistently, and reliably (Figure 1-30).

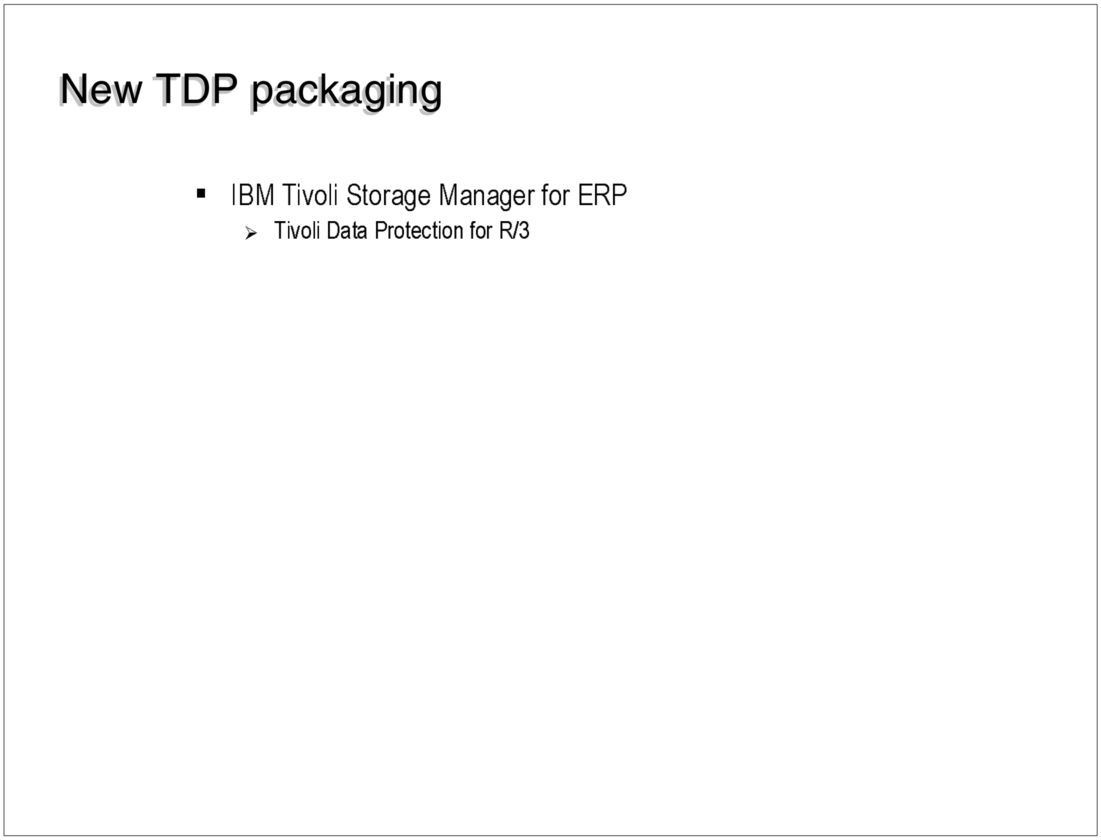

Figure 1-30 IBM Tivoli Storage Manager for ERP

IBM Tivoli Storage Manager for Hardware

IBM Tivoli Storage Manager for Hardware protects your mission-critical databases which require 24X7 availability. It offers exciting new options to implement high-efficiency backup of business-critical applications, while virtually eliminating backup-related performance impact on the production host. This is done by integrating IBM's Enterprise Storage Server (ESS) FlashCopy function with IBM Tivoli Storage Manager and its database protection capabilities for Oracle, R/3, and DB/2. Likewise, EMC's Symmetrix TimeFinder integrates with IBM Tivoli Storage Manager and its protection capabilities for Oracle and R/3 (Figure 1-31).

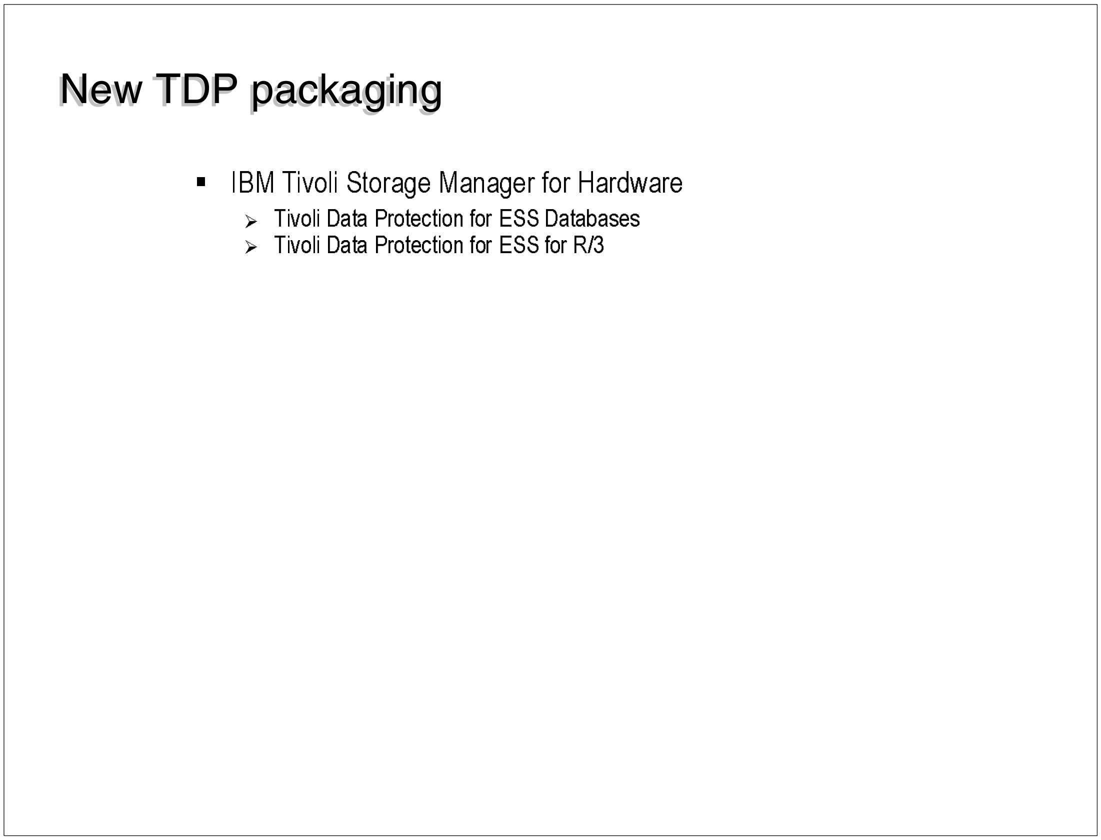

Figure 1-31 IBM Tivoli Storage Manager for hardware

1.3.4 Product migration

Figure 1-32 list the products that can be migrated.

Product migration

- Tivoli Storage Manager
- Tivoli Storage Manager Enterprise Edition
- IBM Tivoli Storage Manager for Mail
- IBM Tivoli Storage Manager for Application Servers
- IBM Tivoli Storage Manager for Databases
- IBM Tivoli Storage Manager for ERP
- IBM Tivoli Storage Manager for Hardware

Figure 1-32 Product migration

The migration table, Table 1-2, summarizes the packaging of the new Tivoli Storage Portfolio as a result of this portfolio consolidation. If you are currently licensed for the products listed in the Previous Product column in the table below, and you are current on your Tivoli Support or Passport Advantage Maintenance contracts, then you are entitled to use the all the function of the products listed in the New Product column until you migrate to the terms and conditions of the new product.

Table 1-2 Product migration table

New Product	Previous Product(s)
IBM Tivoli Storage Manager (5698-ISM)	Tivoli Storage Manager (5698-TSM) with: ► Base server - features 1685 or 1686 ► Managed System for LAN - features 1683 or 1684 ► Managed Library (no feature numbers)
IBM Tivoli Storage Manager Enterprise Edition (5698-ISE)	Tivoli Storage Manager (5698-TSM) with Base server - features 1685 or 1686 and any of: ► Managed System for SAN - features 1679 or 1680 ► Tape Library Sharing - features 1687 or 1688 Tivoli Space Manager (5698-SPM) ► Tivoli Disaster Rec Mgr (5698-DRM) ► Tivoli Data Prot for NDMP (5698-DPA)

New Product	Previous Product(s)
IBM Tivoli Storage Manager for Mail (5698-APE)	► Tiv Data Prot Lotus Domino (5698-DPD) ► Tivoli Data Prot Exchange (5698-DPX)
IBM Tivoli Storage Manager for Databases (5698-APD)	► Tivoli Data Prot Informix (5698-DPI) ► Tivoli Data Prot Oracle (5698-DPO) ► Tivoli Data Prot MS SQL (5698-DPS)
IBM Tivoli Storage Manager for Application Servers (5698-APW)	► Tiv Data Prot WebSphere App Svr (5698-DPW)
IBM Tivoli Storage Manager for Hardware (5698-APH)	► Tiv Data Prot EMC Oracle (5698-EMO) ► Tiv Data Prot EMC Symm R/3 (5698-EMR) ► Tiv Data Prot ESS Database (5698-ESO) ► Tiv Data Prot ESS for R/3 (5698-ESR)
IBM Tivoli Storage Manager for ERP (5698-APR)	► Tivoli Data Prot for R/3 (5698-DPR)

1.3.5 Database conversion

For the AIX, Sun Solaris, Windows NT, and Windows 2000 servers, the Tivoli Storage Manager database is automatically upgraded during the installation of the Tivoli Storage Manager V5.1 server. After this database upgrade, the user can no longer run a previous version of the servers with the database.

For the OS/390 and z/OS servers, the Tivoli Storage Manager database is not automatically upgraded during the installation of the Tivoli Storage Manager OS/390 and z/OS V5.1 servers. The Tivoli Storage Manager administrator must explicitly initiate the database upgrade using the `UPGRADEDB` command parameter. After this database upgrade, the user can no longer run a previous version of the server with the database.

Note: As a normal precaution, it is recommended that a full database backup be performed before installing the Version 5.1 code, as well as after completing the installation. See the README files for more detailed information.

The format of the client data header has changed in V5.1 to include new service information. Once files have been backed up by a V5.1 client to any Tivoli Storage Manager server, those files cannot be restored by any lower-level client.

As part of a migration plan from Tivoli Storage Manager V4.2 to Tivoli Storage Manager V5.1, Tivoli Storage Manager clients and servers may be upgraded at different times. To help ensure that you can continue your backup and archive activities during the migration, note the following:

- A Tivoli Storage Manager V4.2 client can perform backup, restore, archive, and retrieve functions to a Tivoli Storage Manager V5.1 server.
- A Tivoli Storage Manager V5.1 client can perform backup, restore, archive, and retrieve functions to a Tivoli Storage Manager V4.2 server.
- A Tivoli Storage Manager V5.1 client can perform V3.1 functional level backup, restore, archive and retrieve functions to a Tivoli Storage Manager Version 3.1 server on VM.

- A Tivoli Storage Manager V4.2 HSM client can perform migrate and recall functions to a Tivoli Storage Manager V5.1 server.
- A Tivoli Storage Manager V5.1 HSM client can perform migrate and recall functions to a Tivoli Storage Manager V4.2 server.
- Data that has been backed up from a Tivoli Storage Manager V5.1 client to any Tivoli Storage Manager server cannot be restored using a Tivoli Storage Manager V4.2 or lower level client.
- All command line administrative clients can administer Tivoli Storage Manager V4.2 and V5.1 servers, and the V3.1 VM servers.
- Storage agents and servers must be at the same level of code. When the server is upgraded, the storage agents which are using that particular server, must be upgraded as well.

Other migration information:

- The following server operating systems supported in V4.2 were not migrated to V5.1 and are not supported in V5.1:
  - VM/ESA V2R3
  - AS/400 native on Tivoli Storage Manager V3.1
- The following client platforms or operating systems supported in V4.2 were not migrated to V5.1 and are not supported in V5.1:
  - Apple Macintosh 8 and 9.0
  - IBM NUMA Q
  - Tru64 UNIX
  - Linux 2.2 kernel
  - NetWare 4.20
  - Windows 98
  - OS/390 Unix Systems Services V2R8
- The following client communication protocols supported on the following platform in Tivoli Storage Manager V4.2 were not migrated to V5.1 and are not supported in Tivoli Storage Manager V5.1:
  - APPC (LU6.2) with the NetWare client
  - IPX/SPX with the NetWare client

To help minimize server performance impact with use of the Unicode option on Windows NT and Windows 2000 clients operating with Version 5.1 servers, follow the recommended migration techniques found in the Administrator's Guide for the V5.1 server platform(s) being used.

> **Note:** The Tivoli Storage Manager products' performance will vary depending upon a customer's particular configuration, the client/server options, and customization.

1.4 Terminology changes

A complete list of terminology definitions are available at the following Web site: `http://www.tivoli.com/products/licensing/`. The following terminology is defined for IBM Tivoli Storage Manager V5.1.

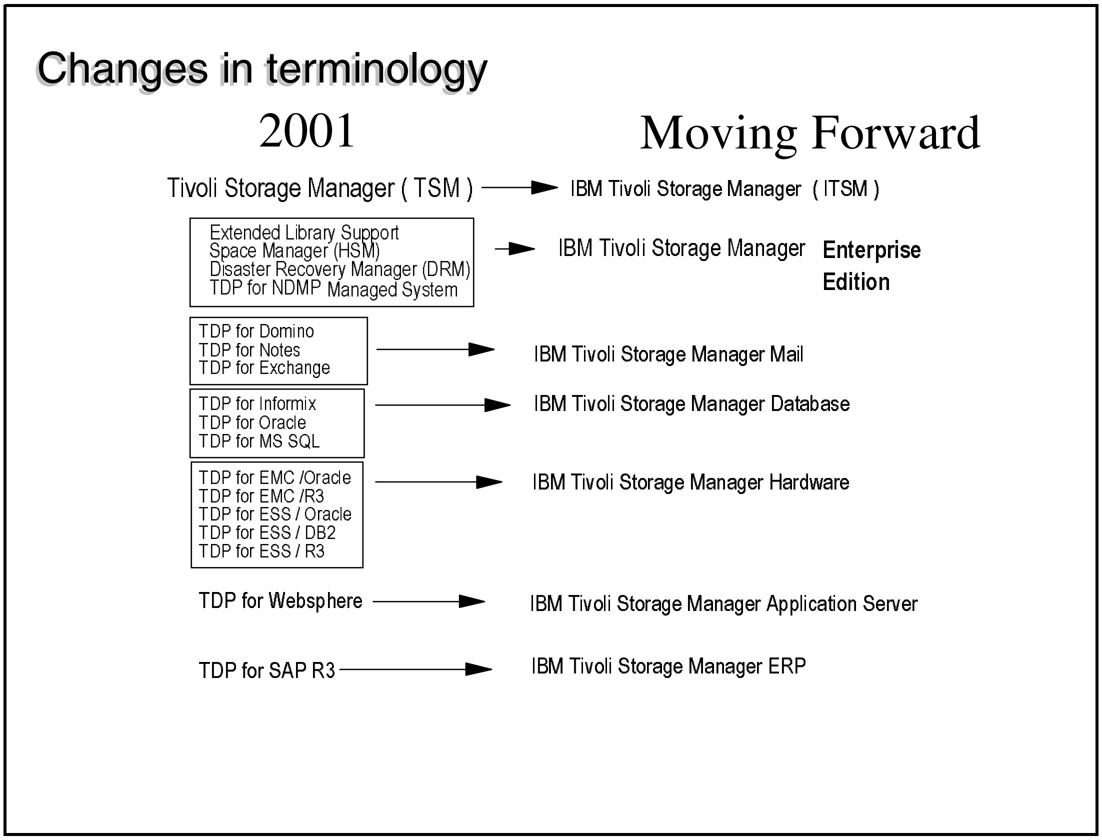

Figure 1-33 Changes in terminology

Application server

An application server is the computer system that provides services to one or more clients, servers or other devices over a network and is defined by its use within the Tivoli Application. For programs running on a distributed cluster of servers or on a S/390 sysplex, Value Units sufficient for an application server are required for each instance of the operating system rather than for each individual server.

Application instance

A single installation of an application, which may be across one or more servers.

- For SAP R/3, an instance represents an SAP R/3 database instance. Each R/3 database instance consists of processes, main memory structures (caches), and a disk-based database.
- For PeopleSoft, an instance represents a PeopleSoft database instance. Each PeopleSoft database instance consists of the database, Application Data Tables, Application Metadata, the Portal Registry, and the Batch Server.
- For Oracle applications, an instance is the Oracle Internet Application Server (as opposed to the database).
- MVS - Each instance represents one connection to the mainframe.

Client	A computer system or process that requests a service of another computer system that is typically referred to as a server. Multiple clients may share access to a common server. A client has some processing capability or is programmable to allow a user to do work. Examples include laptop computers, desktop computers, and desk side computers.
Cluster	A group of distributed servers operated as a single image, generally running in serial. Each system in a cluster is priced as a server, with the appropriate tier as determined by the Server Tier exhibit.
Device	A device is any non-client, non-server part of a network managed by a Tivoli application, including, but not limited to Palm Pilots and other pervasive devices.
Network Node	Network nodes include routers, switches, hubs, and bridges that contain a network management agent. In the Internet Protocol (IP), the managed node usually contains a Simple Network Management Protocol (SNMP) agent. A single node may contain any number of interfaces or ports.
Physical partitions	For servers capable of physical partitioning (for example, the pSeries Scalable POWERparallel Systems and p690 servers, the Sun Ultra servers, and the HP Superdome servers), the licensing is based on the number of processors in the partitioned for the application managed. Note that this only applies to fixed partitions, not variable workload. For example, the SP is a Tier 5 server if it has greater than 23 processors installed *in aggregate* in the systems. However, if a system has 2 nodes with 8 processors each (16 processors total), the system would be a Tier 4 server. As an alternate example, if the system has 3 nodes with 8 processors each (24 processors total), but the nodes are partitioned and the partition which is being managed by a Tivoli application is using 8 processors, then the partition managed is a Tier 3 server.
Registered user	A registered user is defined as anyone who is authorized to access a Tivoli program.
Server	A server is a computer system that provides services to one or more clients and/or other devices over a network. Examples include, but are not limited to, file servers, print servers, mail servers, database servers, application servers, and Web servers. A server is defined by its use in the customer's environment, not by its use within a Tivoli application.
Standby or backup systems	For programs running or resident on backup machines, IBM defines three types of situations: cold, warm, and hot. In the cold and warm situations, Value Units for the backup copy are normally not required and no additional charge applies. In a "hot" backup situation, the customer needs Value Units sufficient for that server. All programs running

in backup mode must be under the customer's control, even if running at another enterprise's location.

As a practice, the following are definitions and allowable actions concerning the copy of the program used for backup purposes:

Cold A copy of the program may be stored for backup purposes on a machine as long as the program has not been started.

Warm A copy of the program may reside for backup purposes on a machine and is started, but is "idling", and is not doing any work of any kind. Value Units are not required for this system.

Hot A copy of the program may reside for backup purposes on a machine, is started, and is doing work. However, this program must be ordered. Value Units are required for this system.

"Doing Work", includes, for example, production, development, program maintenance, and testing. It also could include other activities such as mirroring of transactions, updating of files, synchronization of programs, data or other resources (for example, active linking with another machine, program, database or other resource, and so on), or any activity or configurations that would allow an active hot switch or other synchronized switch over between programs, databases, or other resources to occur.

1.5 Pricing model

In 2000, Tivoli introduced it's industry-leading *Environment-Managed pricing model*. The unique element of this model is that pricing and licensing is based on what is managed (for example, clients, servers, and network devices), not how the customer chooses to implement the technical facets of a Tivoli solution (for example, management servers, gateways, and endpoints). With IBM Tivoli Storage Manager Version 5.1, Tivoli is not changing the fundamentals of this model, but refining some of the definitions and creating a more solution-oriented offering structure.

1.5.1 Tivoli Environment-Managed pricing model

The Environment-Managed model is used to describe the Tivoli price offering that separates licensing from architecture. In an environment managed model, the price is based on what is managed rather than the more traditional approach of pricing based on the number and type of product component. Thus, a server is defined by its use in the customer's environment, not by its use within a Tivoli application. For example, all servers (including, but not limited to, file servers, print servers, and application servers) monitored with IBM Tivoli Monitoring require *Value Units* sufficient for those servers based on the server tier definition. Some products, such as IBM Tivoli Configuration Manager, manage not only servers, but also clients and other devices.

Note: Each device or class of device requires Value Units as laid out in the Value Unit Exhibit maintained at: `http://www.tivoli.com`

For example, management capability for a certain sized environment is acquired for each Tivoli discipline required. The customer is free to implement any of the Tivoli Enterprise software components (Tivoli Management Region servers, gateways, or endpoints) in any quantity to best manage within the environment for which the discipline was acquired without cost implications. As architecture changes, the implementation can be altered as needed without necessarily affecting the licensing for the products acquired. This represents a significant change from the typical systems management licensing models that are based on purchasing the specific components and licensing them to specific systems.

As such, a server is defined by its use in the environment, not its role in a Tivoli solution. Thus, all servers in an environment are priced as servers, offering the flexibility to serve Tivoli applications from any system in the enterprise without any change in cost. Further, these servers have been stratified to better correlate cost to value while offering a simple, scalable solution. In addition to traditional clients, other devices can be managed in a server/device model, including Personal Digital Assistants (PDAs), point-of-sale terminals (tills, cash registers), Automated Teller Machines (ATMs), and appliances. In this model, only those systems or devices managed require licensing.

1.5.2 Distributed servers

Distributed server tiers are defined by the machine, *not* operating system, based on installed processors. For servers capable of physical partitioning (for example, the IBM eServer pSeries Scalable POWERparallel Systems and p690 servers, the Sun Ultra servers, and the HP Superdome servers), the licensing is based on the number of installed processors in the partition for the application managed.

Note: This only applies to physical partitions, not virtual partitions.

This is the approach for systems which have either multiple cards or multiple frames, each of which can be configured independently. For clustered environments, the licensing is determined by the number of installed processors in each server in the cluster.

The server definition has been updated since Value-Based Pricing was introduced in February 2000 to recognize the convergence of RISC/UNIX and WinTel technologies, simplify the customer's licensing requirements, and provide a smother, more scalable model.

Part 2

IBM Tivoli Storage Manager server enhancements

Part 2 describes the server-specific enhancements provided in V5.1.

2

Common server enhancements

This chapter discusses the new features and enhancements delivered in IBM Tivoli Storage Manager Version 5.1 server, which are common to all server platforms.

2.1 Changes to library and drive definition

As of Tivoli Storage Manager Version 5.1, the procedure for defining a tape library of tape drive has changed. It is now necessary to define a data path for all libraries and drives, including local libraries and drives. The path definitions are necessary for the server-free product enhancements. An example of the steps now required for library definition is shown below. Some of the options previously used with the define library/drive command have been moved to the **`path`** command, and this example shows how to define a library named 3570L and one of its drives:

```
define library 3570L libtype=scsi shared=yes
define path TSM_AIX 3570L srctype=server desttype=library device=/dev/lb0
```

2.2 Simultaneous writes to copy storage pools

IBM Tivoli Storage Manager now provides the possibility to simultaneously store a clients files into each copy storage pool specified for the primary storage pool where the clients files are written. The simultaneous write to the copy pools only takes place during backup or archive from the client. In other words, when the data enters the storage pool hierarchy. It does not take place during data migration from an HSM client nor on a LAN free backup from a Storage Agent. Up to ten copy storage pools can be specified for each primary storage pool. This is not a replacement for the `BACKUP STGPOOL` operation.The `BACKUP STGPOOL` command cannot write to multiple copy storage pools at the same time.

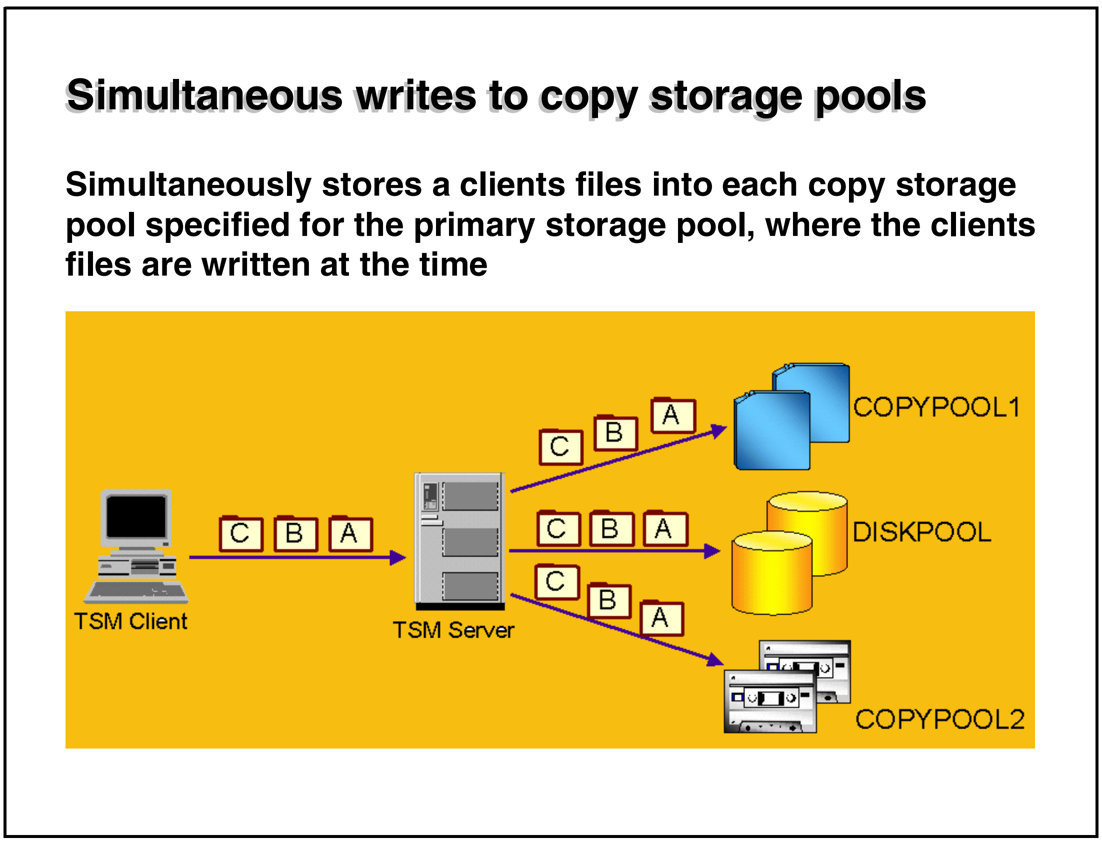

Figure 2-1 Simultaneous write

2.2.1 How to use simultaneous writes to copy storage pools

Careful consideration should be used when using simultaneous write or copy pool duplexing. The data is written to the copy storage pool and primary storage pool simultaneously. The backup performance will only be as good as the slowest device being used for any of the pools. There also has to be enough mount points available. If there are not enough mount points available, only the client, not the server, will issue the ANS1312E message:

```
'ANS1312E Server media mount not possible'
```

> **Note:** If you have multiple clients backing up to the same primary pool and the copy pools are tape based pools, each client will need multiple mount points to be able to write the data to the copy pool simultaneously. The allowed number of mountpoints for a node is restricted by the node parameter MAXNUMMP.

How to enable simultaneous write for primary storage pools

To enable simultaneous write for a primary storage pool, you must set the COPYSTGpools parameter for the storage pool (**DEFINE** or **UPDATE STGPOOL** commands). You can specify a maximum of ten (10) copy pool names separated by commas and no intervening spaces. The maximum length of each copy pool name is 30 characters. This option is only available for primary storage pools using NATIVE or NONBLOCK data format.

How to enable multiple mount points for nodes

To allow clients to utilize the simultaneous write for a primary storage pool, the client must have an appropriate MAXNUMMP setting in the node definition in the server. MAXNUMMP specifies the *maximum* number of mount points a node is allowed to use on the server and can be set to an integer from 0-999. Clients that stores its data to a storage pool that has one or more copy storage pools defined for simultaneous backup, the client may require additional mount points. As a general rule, each copy storage pool of sequential device type must be assign one mount point. If the primary storage pool is of sequential device type, then assign a mount point for the primary storage pool as well.

> **Note:** For each storage pool of sequential device type that a client will use for storing data in the Tivoli Storage Manager server, assign one mount point for each using the MAXNUMMP client node parameter.

Please refer to “MAXNUMMP” on page 173 for more examples on how to set the MAXNUMMP parameter for client nodes.

2.2.2 How to use simultaneous write for primary storage pools

In the following examples we will use one primary storage pool with DISK device class (DISKPOOL) and a copy storage pool with FILE device class (COPYPOOL). The following **dsmadmc** command create the FILE device class for the copy storage pool:

```
define devclass COPY devtype=FILE mountLimit=9 directory=/tsm/stg/c
```

The following **dsmadmc** command create the copy storage pool (the maxscratch value is the maximum value allowed but should not exceed the available space in the filesystem defined for the device class in the example above):

```
define stgpool COPYPOOL COPY pooltype=copy maxscratch=100000000
```

The following **dsmadmc** command create the primary storage pool:

```
define stgpool DISKPOOL DISK COPYSTGpools=COPYPOOL
```

The following **dsmadmc** command create the volumes used by the primary storage pool DISKPOOL:

```
define volume DISKPOOL /tsm/stg/d/disk#1.dsm formatsize=50
define volume DISKPOOL /tsm/stg/d/disk#2.dsm formatsize=50
define volume DISKPOOL /tsm/stg/d/disk#3.dsm formatsize=100
```

Example 2-1 shows how to examine the client node definitions in the IBM Tivoli Storage Manager server.

Example 2-1 Examining the updated client node definition

```
root@crete:/: dsmadmc -id=admin -password=admin query node one-on-one f=d
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSMC: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/01/02   17:52:04  Last access: 03/01/02   17:50:14

ANS8000I Server command: 'query node one-on-one f=d'

                     Node Name: ONE-ON-ONE
                      Platform: WinNT
               Client OS Level: 4.00
                Client Version: Version 5, Release 1, Level 0.0
            Policy Domain Name: STANDARD
         Last Access Date/Time: 03/01/02   17:39:44
        Days Since Last Access: <1
        Password Set Date/Time: 03/01/02   17:39:43
       Days Since Password Set: <1
         Invalid Sign-on Count: 0
                       Locked?: No
                       Contact: The Dude
                   Compression: Client
       Archive Delete Allowed?: Yes
        Backup Delete Allowed?: No
        Registration Date/Time: 03/01/02   17:39:43
     Registering Administrator: OPEN_REGISTRATION
Last Communication Method Used:
   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
                     Optionset:
                           URL:
                     Node Type: Client
    Password Expiration Period:
             Keep Mount Point?: No
  Maximum Mount Points Allowed: 1
        Auto Filespace Rename : No
             Validate Protocol:

ANS8002I Highest return code was 0.
```

In Example 2-1 for the client node ONE-ON-ONE, the policy domain is STANDARD and the MAXNUMMP value is one (1) since the primary storage pool is of DISK file type and there will only be one copy storage pool specified by the COPYSTGpools parameter, see Example 2-3. To ensure that the STANDARD policy domain will use the DISKPOOL storage pool, we change the backup copy group with the following **dsmadmc** command:

```
update copygroup standard standard standard standard destination=DISKPOOL
```

To verify we use the **query** command option with the **dsmadmc** command as shown in Example 2-2.

Example 2-2 Examining the backup copygroup

```
root@crete:/: dsmadmc -id=admin -password=admin query copygroup standard standard t=backup
f=d
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/04/02   19:47:20  Last access: 03/04/02   19:46:22

ANS8000I Server command: 'q copygroup standard standard t=backup f=d'

            Policy Domain Name: STANDARD
               Policy Set Name: STANDARD
               Mgmt Class Name: STANDARD
               Copy Group Name: STANDARD
               Copy Group Type: Backup
          Versions Data Exists: 2
         Versions Data Deleted: 1
         Retain Extra Versions: 30
           Retain Only Version: 60
                     Copy Mode: Modified
            Copy Serialization: Shared Static
                Copy Frequency: 0
              Copy Destination: DISKPOOL
Last Update by (administrator): ADMIN
         Last Update Date/Time: 03/04/02   19:07:39
              Managing profile:

ANS8002I Highest return code was 0.
```

In the output from the **dsmadmc** command above in Example 2-2, we can see that the Copy Destination is the DISKPOOL storage pool. So now let us look at the DISKPOOL storage pool information as shown in Example 2-3.

Example 2-3 Examining the DISKPOOL storage pool

```
root@crete:/: dsmadmc -id=admin -password=admin query stg diskpool f=d
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/04/02   19:32:10  Last access: 03/04/02   19:27:02

ANS8000I Server command: 'query stg diskpool f=d'
```

```
              Storage Pool Name: DISKPOOL
              Storage Pool Type: Primary
              Device Class Name: DISK
        Estimated Capacity (MB): 200.0
                       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:
           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): 0.00
Elapsed Migration Time (seconds): 0
        Reclamation in Progress?:
 Volume Being Migrated/Reclaimed:
  Last Update by (administrator): ADMIN
          Last Update Date/Time: 03/04/02   14:19:11
        Storage Pool Data Format: Native
           Copy Storage Pool(s): COPYPOOL
         Continue Copy on Error?:
                       CRC Data: No

ANS8002I Highest return code was 0.
```

The output in Example 2-3 shows that the `DISKPOOL` storage pool is a `Primary` storage pool type and that the `COPYPOOL` storage pool is its designated `Copy Storage Pool`. Let us look at the `COPYPOOL` storage pool information as shown in Example 2-4.

Example 2-4 Examining the COPYPOOL storage pool

```
root@crete:/: dsmadmc -id=admin -password=admin query stg copypool f=d
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/04/02   19:32:10  Last access: 03/04/02   19:27:02

ANS8000I Server command: 'query stg copypool f=d'

              Storage Pool Name: COPYPOOL
              Storage Pool Type: Copy
              Device Class Name: COPY
        Estimated Capacity (MB): 0.0
                       Pct Util: 0.0
                       Pct Migr:
```

```
                    Pct Logical: 100.0
                   High Mig Pct:
                    Low Mig Pct:
                Migration Delay:
             Migration Continue:
            Migration Processes:
              Next Storage Pool:
           Reclaim Storage Pool:
         Maximum Size Threshold:
                         Access: Read/Write
                    Description:
              Overflow Location:
          Cache Migrated Files?:
                     Collocate?: No
          Reclamation Threshold: 100
 Maximum Scratch Volumes Allowed: 100,000,000
   Delay Period for Volume Reuse: 0 Day(s)
         Migration in Progress?:
           Amount Migrated (MB):
Elapsed Migration Time (seconds):
       Reclamation in Progress?: No
 Volume Being Migrated/Reclaimed:
  Last Update by (administrator): ADMIN
          Last Update Date/Time: 03/04/02   14:18:43
       Storage Pool Data Format: Native
           Copy Storage Pool(s):
         Continue Copy on Error?:
                       CRC Data: No

ANS8002I Highest return code was 0.
```

The output in Example 2-4 shows that the `COPYPOOL` storage pool is a `Copy` storage pool type and that the `Device Class` is `COPY`. This option is restricted to only primary storage pools using NATIVE or NONBLOCK data format. The `COPYContinue` option specifies how the server should react to a copy storage pool write failure for any of the copy storage pools listed in the `COPYSTGpools` parameter, and can be set to `Yes` or `No`. If set to `Yes` (default) when a write failure occur, the server will exclude the failing copy storage pool from any further writes while that specific client session is active. If set to `No` then when a write failure occur, the server will fail the entire transaction including the write to the primary storage pool.

Let us look at the `COPY` device class information as shown in Example 2-5.

Example 2-5 Examining the COPY device class

```
root@crete:/tsm/pbin: dsmadmc -id=admin -password=admin query devc copy f=d
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/05/02   12:04:10  Last access: 03/05/02   11:57:43

ANS8000I Server command: 'query devc copy f=d'

             Device Class Name: COPY
        Device Access Strategy: Sequential
            Storage Pool Count: 1
                   Device Type: FILE
                        Format: DRIVE
```

```
         Est/Max Capacity (MB): 4.0
                   Mount Limit: 9
              Mount Wait (min):
         Mount Retention (min):
                  Label Prefix:
                       Library:
                     Directory: /tsm/stg/c
                   Server Name:
                  Retry Period:
                Retry Interval:
                        Shared:
Last Update by (administrator): ADMIN
         Last Update Date/Time: 03/05/02   11:47:42

ANS8002I Highest return code was 0.
```

The output in Example 2-5 shows that the `COPY` device class is of the `FILE Device Type` and the `Directory` where the files, that each will be `4.0` MB, are written is `/tsm/stg/c` and up to `9` files can be written simultaneously.

Now we can look at the progress of the backup from the client node `ONE-ON-ONE` to the Tivoli Storage Manager server. First we can examine the backup session with the **`query session`** command option to the **`dsmadmc`** command as in the following sample in Example 2-6.

Example 2-6 Examining the client session

```
root@crete:/: dsmadmc -id=admin -password=admin query session 2 f=d
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/04/02   19:14:47  Last access: 03/04/02   19:04:48

ANS8000I Server command: 'query session 2 f=d'

Sess Number: 2
             Comm. Method: Tcp/Ip
               Sess State: RecvW
                Wait Time: 0 S
               Bytes Sent: 392
              Bytes Recvd: 77.3 M
                Sess Type: Node
                 Platform: WinNT
              Client Name: ONE-ON-ONE
      Media Access Status: Current output volume(s):
                            COPYPOOL,/tsm/stg/c/00000032.BFS,(1686142 Seconds)
                User Name:
Date/Time First Data Sent: 03/04/02   19:12:10

ANS8002I Highest return code was 0.
```

In Example 2-6 we can clearly see that currently the `FILE` device class copy storage pool `COPYPOOL` is used to write the copy of the client `ONE-ON-ONE` filespace data. Since the copy storage pool `COPYPOOL` is a `FILE` device class storage pool we can examine the volumes that the Tivoli Storage Manager server has created during our clients session (the storage pool was empty when the client session started) in Example 2-7.

Example 2-7 Examining current volumes for the copy pool during backup operation

```
root@crete:/: dsmadmc -id=admin -password=admin query vol stg=copypool
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/04/02   19:17:32  Last access: 03/04/02   19:16:27

ANS8000I Server command: 'query vol stg=copypool'

Volume Name               Storage      Device      Estimated    Pct   Volume
                          Pool Name    Class Name   Capacity   Util   Status
                                                        (MB)
------------------------  -----------  ----------  ---------  -----  --------
/tsm/stg/c/00000021.BFS   COPYPOOL     COPY              0.0    0.0   Empty
...<omitted lines>...
/tsm/stg/c/00000032.BFS   COPYPOOL     COPY              0.0    0.0   Empty

ANS8002I Highest return code was 0.
```

The Example 2-7 output does not show the amount of data nor the utilization percentage, this will be updated when the operation is completed, such as after the client operation is finished. After the client session has finished the output might look as in the following sample in Example 2-8.

Example 2-8 Examining current volumes for the copy pool after backup operation

```
root@crete:/: dsmadmc -id=admin -password=admin query vol stg=copypool
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/04/02   19:56:16  Last access: 03/04/02   19:46:22

ANS8000I Server command: 'query vol stg=copypool'

Volume Name               Storage      Device      Estimated    Pct   Volume
                          Pool Name    Class Name   Capacity   Util   Status
                                                        (MB)
------------------------  -----------  ----------  ---------  -----  --------
/tsm/stg/c/00000021.BFS   COPYPOOL     COPY              4.0  100.0    Full
/tsm/stg/c/00000022.BFS   COPYPOOL     COPY              4.0  100.0    Full
...<omitted lines>...
/tsm/stg/c/00000035.BFS   COPYPOOL     COPY              4.0  100.0    Full
/tsm/stg/c/00000036.BFS   COPYPOOL     COPY              4.0  100.0    Full
/tsm/stg/c/00000037.BFS   COPYPOOL     COPY              4.0   98.6  Filling

ANS8002I Highest return code was 0.
```

All the examples in Example 2-8 shows how the simultaneous write work and how it is set up. To ensure that all data has been copied between the primary storage pool and the copy storage pool, run the `BACKUP STGPOOL` operation as shown in Example 2-9.

Example 2-9 Backing up a storage pool manually

```
root@crete:/: dsmadmc -id=admin -password=admin backup stgpool diskpool copypool wait=yes
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/05/02   13:22:38  Last access: 03/05/02   13:02:26

ANS8000I Server command: 'backup stgpool diskpool copypool wait=yes'
ANR0984I Process 2 for BACKUP STORAGE POOL started in the FOREGROUND at
13:22:38.
ANR2110I BACKUP STGPOOL started as process 2.
ANR1210I Backup of primary storage pool DISKPOOL to copy storage pool COPYPOOL
started as process 2.
ANR1212I Backup process 2 ended for storage pool DISKPOOL.
ANR0985I Process 2 for BACKUP STORAGE POOL running in the FOREGROUND completed
with completion state SUCCESS at 13:22:38.
ANR1214I Backup of primary storage pool DISKPOOL to copy storage pool COPYPOOL
has ended.  Files Backed Up: 0, Bytes Backed Up: 0, Unreadable Files: 0,
Unreadable Bytes: 0.

ANS8002I Highest return code was 0.
```

The output in Example 2-9 above shows that the `COPYPOOL` storage pool contains all files from the `DISKPOOL` storage pool.

2.2.3 How to use simultaneous write with SERVER-SERVER device classes

In the following scenario we have one Tivoli Storage Manager V5.1 server on IBM AIX and one Tivoli Storage Manager V5.1 server on Sun Solaris, the Tivoli Storage Manager name for the IBM system is `TSM_AIX` and the Tivoli Storage Manager name for the Sun system is `TSM_SUN`. The IBM system is the primary backup server, but will use the Sun system for copy pool storage. In this case we will use the simultaneous write feature of the primary storage pool for the IBM system to simultaneously write to a local copy storage pool and to a server-server storage pool, which resides on the Sun system. The IBM system is also connected to the SAN. The scenario layout is shown in Figure 2-2.

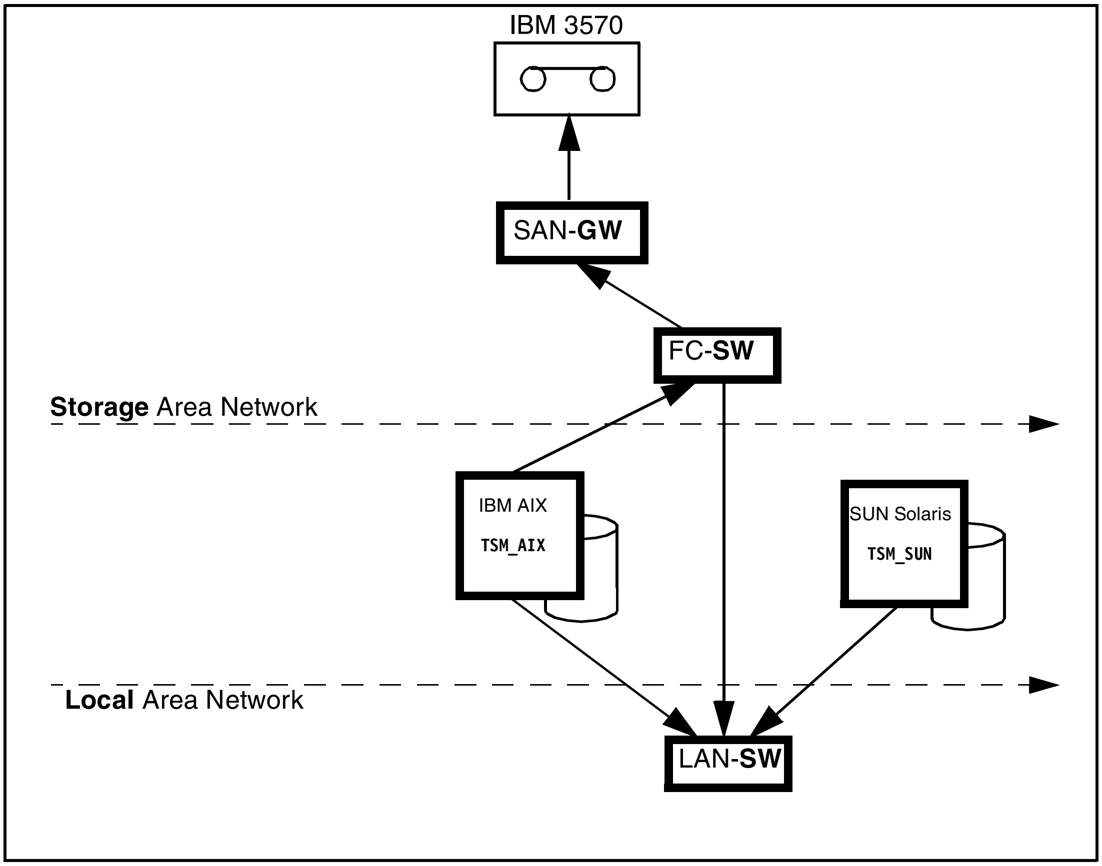

Figure 2-2 Schematic layout of the sample environment

The `TSM_AIX` server will use the following storage pools in our scenario:

- DISKPOOL
- COPYPOOL
- SUNPOOL
- TAPEPOOL

The sample in Example 2-10 shows the storage pool overview.

Example 2-10 Storage pools

Storage Pool Name	Device Class Name	Estimated Capacity (MB)	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
TAPEPOOL	3570DC	0.0	0.0	0.0	90	70	
COPYPOOL	COPY	0.0	0.0				
DISKPOOL	DISK	4,500.0	0.0	0.0	90	70	3570POOL
SUNPOOL	SUNDC	0.0	0.0				

The management class and copy groups used by the client point to the `DISKPOOL` storage pool. The `DISKPOOL` storage pool specify the `COPYSTGpools` to be `COPYPOOL` and `SUNPOOL`, and we will use `COPYContinue` set to `NO` since we want to be sure that the copy is transferred to the `SUNPOOL` properly or the backup will have to be rerun. The `Next` attribute for the `DISKPOOL` storage pool points to the tape library which is used by the `TAPEPOOL`.

Step 1: How to set up the DISKPOOL storage pool

The DISKPOOL storage pool is created with the following **dsmadmc** commands and options:

```
define stgpool DISKPOOL DISK COPYContinue=No CRCData=Yes COPYSTGpools=COPYPOOL,SUNPOOL
```

We use the `CRCData` set to `Yes` and set the `COPYContinue` to `No`. The `COPYSTGpools` parameter is set to `COPYPOOL` and `SUNPOOL` separated by a comma. Example 2-11 shows the result from the query command:

```
query stgpool DISKPOOL f=d
```

Example 2-11 query stgpool diskpool f=d

```
             Storage Pool Name: DISKPOOL
             Storage Pool Type: Primary
             Device Class Name: DISK
       Estimated Capacity (MB): 4,500.0
                      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): 0.00
Elapsed Migration Time (seconds): 0
       Reclamation in Progress?:
 Volume Being Migrated/Reclaimed:
  Last Update by (administrator): ADMIN
         Last Update Date/Time: 03/15/02   14:46:41
       Storage Pool Data Format: Native
           Copy Storage Pool(s): COPYPOOL SUNPOOL
        Continue Copy on Error?: No
                       CRC Data: Yes
```

Step 2: How to set up the TAPEPOOL storage pool

To set up the `TAPEPOOL` storage pool requires some preparation and installation of both hardware and software with the appropriate Program Temporary Fix (PTF) levels, since it is a Storage Area Network (SAN) attached 3570 tape library. In Figure 2-2 the 3570 tape library is connected to a SAN Data Gateway (`SAN-GW`). This in turn is connected to the SAN Fabric (`FC-SW`). The IBM pSeries server is connected to the SAN Fabric through an FC Adapter. The FC-SW is also connected to the Ethernet Local Area Network (LAN) for administration and monitoring purposes.

After the SAN is properly installed and connected we can start to create the prerequisite devices and definitions for the `TAPEPOOL`.To do this we first install the **Atape** software, in our case the `Atape.driver 7.0.3.0` as shown in the sample output in Example 2-12 from the `lslpp` command.

Example 2-12 lslpp -l “Atape”*

```
root@crete:/: lslpp -l "Atape*"
  Fileset                      Level  State      Description
  ----------------------------------------------------------------------------
Path: /usr/lib/objrepos
  Atape.driver               7.0.3.0  COMMITTED  IBM AIX Enhanced Tape and
                                                 Medium Changer Device Driver
```

The **Atape** drivers can be downloaded from `ftp://ftp.software.ibm.com/storage/devdrvr`.

> **Note:** Except for the Windows device drivers, there are numeric sequence numbers in each level of driver, i.e. `AIX/Atape.4.4.0.0.bin`. As newer levels of a driver are released, a higher numeric sequence number will be assigned. For Windows device drivers, one or more level of driver is included in an InstallShield image. A driver with the string "`engr`" is an engineering release driver — use engineering release drivers only as directed by IBM service.

After running **`cfgmgr`** to pick up the FC devices, we can verify this by running the AIX **`lsdev`** command with different options. The first sample in Example 2-13, shows all tape drives in the system. Note that the `mt0` is mapped to `rmt1` and `mt1` is mapped to `rmt2`.

Example 2-13 Available tape devices

```
root@crete:/: lsdev -C -s fcp -c "*tape"
rmt1 Available 3A-08-01-2,0 IBM Magstar MP Tape Subsystem (FCP)
rmt2 Available 3A-08-01-3,0 IBM Magstar MP Tape Subsystem (FCP)
mt0  Available 3A-08-01     Tape Drive
mt1  Available 3A-08-01     Tape Drive
```

Example 2-14 shows the Adstar Distributed Storage Manager (ADSM) library device.

Example 2-14 Available library devices

```
root@crete:/: lsdev -C -s fcp -c "*lib*"
lb0 Available 3A-08-01 Library/MediumChanger
```

We could have checked the entire Fibre Channel (FC) attachment with the **`lsdev`** options shown in Example 2-15, but then we also have other attached FC devices.

Example 2-15 Available fabric devices

```
root@crete:/: lsdev -C -s fcp
rmt1   Available 3A-08-01-2,0 IBM Magstar MP Tape Subsystem (FCP)
rmt2   Available 3A-08-01-3,0 IBM Magstar MP Tape Subsystem (FCP)
hdisk3 Available 3A-08-01     Other FC SCSI Disk Drive
lb0    Available 3A-08-01     Library/MediumChanger
mt0    Available 3A-08-01     Tape Drive
mt1    Available 3A-08-01     Tape Drive
```

To display more detailed information about the tape devices, you can use the **`tapeutil`** command as shown in Example 2-16, for the `rmt1` tape device.

Example 2-16 Using tapeutil

```
root@crete:/: tapeutil -f /dev/rmt1 vpd path
Getting vpd data from device...

  Manufacturer ............. IBM
  Product Identification .... 03570C12
  Revision Level ........... 5466

Querying device paths...

Primary Path Number 1
  Logical Device................. rmt1
  SCSI Parent.................... fscsi0
  Target ID...................... 2
  Logical Unit................... 0
  SCSI Bus....................... 3
  FCP SCSI ID.................... 0x31d00
  FCP Logical Unit............... 0x2000000000000
  FCP World Wide Name............ 0x2004006045160d2e
  Path Enabled................... Yes
  Path Manually Disabled......... No

Total paths configured........... 1
```

Now we can start defining our tape and library devices in Tivoli Storage Manager. First we define the tape library which we name `3570L`, with the following **`dsmadmc define library`** command, with the output shown in Example 2-17:

```
define library 3570L libtype=SCSI
```

Example 2-17 Output sample from the define library

```
ANS8000I Server command: 'define library 3570L libtype=SCSI'
ANR8400I Library 3570L defined.
```

The `3570L` library is Small Computer Interface Standard (SCSI) connected since FC emulates SCSI. Other interesting options, that we do not use now, are `LANFree` and `SHAREd`.

After defining the library we now need to use the new `PATH` command to specify that the library that we just defined, `3570L` is available from the `TSM_AIX` host, with the following **`dsmadmc define path`** command, with the output shown in Example 2-18:

```
define path TSM_AIX 3570L srctype=server desttype=library device=/dev/lb0
```

Example 2-18 Output sample from the define path

```
ANS8000I Server command: 'define path TSM_AIX 3570L srctype=server desttype=library
device=/dev/lb0'
ANR1720I A path from TSM_AIX to 3570L has been defined.
```

The path is from the host, `TSM_AIX`, to the destination which is the `3570L` library and the device we use is the ADSM library device driver `/dev/lb0`. Next we add the drivers (and this 3570 tape library has two) with the following **`dsmadmc define drive`** command; the output is shown in Example 2-19:

```
define drive 3570L 3570D1 element=16
```

Example 2-19 Output sample from the define drive

```
ANS8000I Server command: 'define drive 3570L 3570D1 element=16'
```

```
ANR8404I Drive 3570D1 defined in library 3570L.
```

This command specifies that the 3570D1 drive belongs to the 3570L library and has the library element address 16 (one drive has element address 16 and the other has element address 17 in this dual tape drive library).

After defining the first drive we now need to use the new PATH command again to specify that the drive that we just defined, 3570D1 is available from the TSM_AIX host, with the following **dsmadmc define path** command, with the output shown in Example 2-20:

```
define path TSM_AIX 3570D1 srctype=server desttype=drive library=3570L device=/dev/rmt1
```

Example 2-20 Output sample from the define path

```
ANS8000I Server command: 'define path TSM_AIX 3570D1 srctype=server desttype=drive
library=3570L device=/dev/rmt1'
ANR1720I A path from TSM_AIX to 3570L 3570D1 has been defined.
```

The path is from the host, TSM_AIX, to the destination which is the 3570D1 tape drive and the device we use is the AIX tape device driver /dev/rmt1 and not the ADSM tape device driver (mt0).

Then we do the same for the next tape driver, first we define the drive, with the output shown in Example 2-21:

```
define drive 3570L 3570D2 element=17
```

Example 2-21 Output sample from the define drive

```
ANS8000I Server command: 'define drive 3570L 3570D2 element=17'
ANR8404I Drive 3570D2 defined in library 3570L.
```

Then we define the path:

```
define path TSM_AIX 3570D2 srctype=server desttype=drive library=3570L device=/dev/rmt2
```

Now it is time to check in the volumes in the library. In this case we know that the tapes are already labeled and that there are ten tapes but that the library has room for more.

We check in the tapes with the following **dsmadmc checkin libvolume** command:

```
checkin libvolume 3570L status=scratch search=yes
```

While the Tivoli Storage Manager server now is busy checking in the tapes in the tape library, we can monitor the progress with the **dsmadmc query process** command, with the output shown in Example 2-22:

```
query process
```

Example 2-22 Output sample from the query process command

```
Process  Process Description   Status
  Number
--------  --------------------  -------------------------------------------------
      34  CHECKIN LIBVOLUME     ANR8425I Checking in volumes in search mode in
                                 library 3570L.
```

During the check in and after the check in is completed, we can query the Tivoli Storage Manager server to find how many tape volumes were found — their labels and where they are in the library with the **dsmadmc query libvol** command; the output is shown in Example 2-23.

Example 2-23 Output sample from the query libvolume

```
Library Name   Volume Name   Status       Owner        Last Use    Home Element
------------   -----------   ----------   ----------   ---------   ------------
3570L          07D9A3        Scratch                               32
3570L          080CED        Scratch                               40
3570L          081DF3        Scratch                               34
3570L          0838B5        Scratch                               38
3570L          083964        Scratch                               33
3570L          085135        Scratch                               37
3570L          085E0A        Scratch                               35
3570L          085EA7        Scratch                               39
3570L          085EA8        Scratch                               36
3570L          08659E        Scratch                               41
```

Step 3: How to set up the SUNPOOL storage pool

To set up the SUNPOOL storage pool requires, since it is a connection between two Local Area Network (LAN) attached Tivoli Storage Manager servers, access to both host systems. In the Figure 2-2 on page 57 the TSM_AIX system is connected to the Ethernet LAN switch as is the TSM_SUN system.

The basics for a server-to-server setup, in our scenario, is that one Tivoli Storage Manager server (TSM_AIX) is using the other Tivoli Storage Manager server (TSM_SUN) as an archive storage pool. However, to do so, we need to define each server at the other location with the proper passwords and addresses.

First we need to coordinate the server names for each Tivoli Storage Manager server. We make sure that the IBM Tivoli Storage Manager server in our setup is named TSM_AIX by issuing the following **dsmadmc** command:

```
set servername TSM_AIX
```

The same **dsmadmc** command is appropriate for the Sun Tivoli Storage Manager server:

```
set servername TSM_SUN
```

We also need to set the server password. In our setup we use the same password on both servers and set it on both TSM_AIX and TSM_SUN with the following **dsmadmc** command:

```
set serverpassword TSMPW
```

We then set the TSM_AIX server IP-address (a simple way to obtain the proper IP-address is to use the **ifconfig** command) and TCP port with the following **dsmadmc** commands:

```
set serverhladdress 9.1.38.191
set serverlladdress 1500
```

On the TSM_SUN server we enable other servers that define a server connection to the TSM_SUN sever to create the reverse connection specification automatically with the following **dsmadmc** command:

```
set crossdefine on
```

Now we define the TSM_SUN server on the source system TSM_AIX with the following **dsmadmc** command:

```
define server TSM_SUN serverpassword=TSMPW crossdefine=yes hladdress=9.1.38.155
lladdress=1500 comm=tcpip
```

Note that we set the password, IP-address, the server port and that we want cross definition to be performed on the TSM_SUN server. Example 2-24 shows the setting for TSM_SUN on the TSM_AIX Tivoli Storage Manager server.

Example 2-24 query server command on server TSM_AIX

```
Server   Comm.  High-level    Low-level   Days Server   Virtual  Allow
Name     Method Address       Address    Since Password Volume   Replacement
                                          Last Set      Password
                                        Access          Set
-------- ------ ------------- --------- ------ -------- -------- -----------
TSM_SUN  TCPIP  9.1.38.155    1500          <1 Yes      Yes      No
```

Since we had `crossdefine` set to `on`, the Tivoli Storage Manager server on `TSM_AIX` requested the Tivoli Storage Manager server on `TSM_SUN` to add `TSM_AIX` as a server. To verify use the **`dsmadmc query server`** command on the `TSM_SUN` system as shown in Example 2-25.

Example 2-25 query server command on server TSM_SUN

```
Server   Comm.  High-level    Low-level   Days Server   Virtual  Allow
Name     Method Address       Address    Since Password Volume   Replacement
                                          Last Set      Password
                                        Access          Set
-------- ------ ------------- --------- ------ -------- -------- -----------
TSM_AIX  TCPIP  9.1.38.191    1500          <1 Yes      Yes      No
```

Now we should be able to probe the `TSM_SUN` server from the `TSM_AIX` server with the following **`dsmadmc`** command, with the output shown in Example 2-26:

```
ping server tsm_sun
```

Example 2-26 Output sample from the ping server command

```
ANS8000I Server command: 'ping server tsm_aix'
ANR1706I Ping for server 'TSM_AIX' was able to establish a connection.
```

The following **`dsmadmc`** command will probe the `TSM_AIX` server from the `TSM_SUN` server, with the output shown in Example 2-27:

```
ping server tsm_sun
```

Example 2-27 Output sample from the ping server command

```
ANS8000I Server command: 'ping server tsm_sun'
ANR1706I Ping for server 'TSM_SUN' was able to establish a connection.
```

On the `TSM_SUN` server we need to define `TSM_AIX` as a client node of `SERVER` type with the following **`dsmadmc`** command:

```
register node TSM_AIX TSMPW domain=STANDARD type=SERVER
```

Note that we use the `STANDARD` policy domain in our sample for simplicity only. In Example 2-28, we show the definition on the `TSM_SUN` Tivoli Storage Manager server.

Example 2-28 Query node type=server on server TSM_SUN

```
ANS8000I Server command: 'q node t=s'

Node Name                 Platform Policy Domain  Days Since Days Since Locked?
                                   Name           Last Acce-   Password
                                                         ss        Set
------------------------- -------- -------------- ---------- ---------- -------
TSM_AIX                   AIX-RS/- STANDARD               <1         <1   No
```

```
                  6000
```

On the TSM_AIX server we now can create a device class, which we name SUNDC, for the server-server connection with the following **dsmamdc** command, with the output shown in Example 2-29:

```
define devclass SUNDC devtype=SERVER servername=TSM_SUN
```

Example 2-29 Query devc on server TSM_AIX

```
Device       Device       Storage    Device       Format     Est/Max     Mount
Class        Access          Pool    Type                   Capacity     Limit
Name         Strategy       Count                               (MB)
---------    ----------   -------    ---------    ------    --------    ------
SUNDC        Sequential         1    SERVER                     50.0         1
```

The next step is to create a copy storage pool on the TSM_AIX server using the SUNDC device class, with the following **dsmadmc** command:

```
define stgpool SUNPOOL SUNDC pooltype=COPY maxscratch=100000000
```

Example 2-30 shows the **dsmadmc query stgpool** command for the SUNPOOL storage pool.

Example 2-30 Query stgpool on server TSM_AIX

```
Storage      Device       Estimated    Pct    Pct  High  Low  Next Stora-
Pool Name    Class Name    Capacity   Util   Migr   Mig  Mig  ge Pool
                               (MB)                 Pct  Pct
-----------  ----------  ----------  -----  -----  ----  ---  -----------
SUNPOOL      SUNDC       38,729,238    0.0
                                 .5
```

How to monitor simultaneous write

Since we have set COPYContinue to No for the primary storage pool, DISKPOOL, then if any error occur while writing to any of the secondary copy storage pools, COPYPOOL or SUNPOOL, then the entire client session will have to be restarted when either all pools are usable or when the storage pool specifications has changed.

To verify the server-server connection between the TSM_AIX server and the TSM_SUN server, we can start a backup from a client, that uses the DISKPOOL storage pool. To find out which clients we can use the following **select** SQL query:

```
select node_name as "CLIENT NODENAME", -
bu_copygroups.destination as "STGPOOL DESTINATION", -
nodes.domain_name as "CLIENT DOMAIN", -
bu_copygroups.domain_name as "COPYGROUP DOMAIN" -
from nodes,bu_copygroups where -
nodes.domain_name = bu_copygroups.domain_name and -
bu_copygroups.destination=upper('diskpool') and -
bu_copygroups.set_name=upper('active') -
order by nodes.domain_name
```

If saved in a file with the name query1.sql, it can then be executed by the following **dsmadmc macro** command:

```
dsmadmc -id=admin -password=admin macro query1.sql
```

Note that the SQL query is a table join on the domain_name rows from each table and the destination for the copygroup has to be the DISKPOOL storage pool and only the active copygroups are interesting for our sample.

The result could be similar to the output in Example 2-31. For more examples of **select** SQL queries, please refer to “Structured Query Language” on page 300.

Example 2-31 Output example from executing dsmadmc macro

```
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM_AIX: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/15/02   19:51:35  Last access: 03/15/02   19:50:06

ANS8000I Server command: 'select node_name as "CLIENT NODENAME", bu_copygroups.destination
as "STGPOOL DESTINATION", nodes.domain_name as "CLIENT DOMAIN", bu_copygroups.domain_name a
s "COPYGROUP DOMAIN" from nodes,bu_copygroups where nodes.domain_name = bu_copygroups.domai
n_name and bu_copygroups.destination=upper('diskpool') and bu_copygroups.set_name=upper('ac
tive') order by nodes.domain_name'

CLIENT NODENAME       STGPOOL DESTINATION     CLIENT DOMAIN         COPYGROUP DOMAIN
------------------    -------------------     ------------------    ------------------
CRETE                 DISKPOOL                STANDARD              STANDARD
SOLEMIO               DISKPOOL                STANDARD              STANDARD
TSM_AIX               DISKPOOL                STANDARD              STANDARD
ONE-ON-ONE            DISKPOOL                WINDOWS               WINDOWS
...<lines omitted>
ANS8000I Server command: 'COMMIT'

ANS8002I Highest return code was 0.
```

Since the COPYPOOL storage pool is a FILE device class storage pool, it will mount local file system files. Since the SUNPOOL storage pool is a SERVER device class, it will mount remote archive files on the remote Tivoli Storage Manager server.

This means that we need to make sure that the client nodes storing data in the DISKPOOL storage pool are allowed to have two or more sequential device mountpoints, which is controlled by the MAXNUMMP client node parameter; please refer to “How to enable multiple mount points for nodes” on page 49.

The following **dsmadmc** command set the MAXNUMMP parameter for node ONE-ON-ONE to two (2):

```
update node ONE-ON-ONE maxnummp=2
```

The following examples show how to use the **dsmadmc mount** command to monitor the active volumes. The first sample, in Example 2-32, shows the current backup session from the client ONE-ON-ONE.

Example 2-32 Query session

```
ANS8000I Server command: 'q session'
Sess    Comm.    Sess       Wait     Bytes     Bytes   Sess     Platform   Client Name
Number   Method   State      Time      Sent     Recvd   Type
------   ------   ------   ------   -------   -------   -----   --------    -------------
   244   Tcp/Ip   IdleW    1.6 M     3.2 K     2.5 K   Node    WinNT       ONE-ON-ONE
   296   Tcp/Ip   RecvW      0 S       380    14.6 M   Node    WinNT       ONE-ON-ONE
```

The first **query mount** command, in Example 2-33, shows the request of one volume from device class SUNDC and one volume from device class COPY.

Example 2-33 Requested volumes

```
ANS8000I Server command: 'q mount'
ANR8376I Mount point reserved in device class SUNDC, status: RESERVED.
ANR8376I Mount point reserved in device class COPY, status: RESERVED.
```

The next **query mount** command, in Example 2-34, shows that the SERVER volume TSM_SUN.BFS.016488211 and the FILE volume /tsm/stg/c/0000014D.BFS has been mounted.

Example 2-34 Mounted volumes

```
ANS8000I Server command: 'q mount'
ANR8333I SERVER volume TSM_SUN.BFS.016488211 is mounted R/W, status: IN USE.
ANR8333I FILE volume /tsm/stg/c/0000014D.BFS is mounted R/W, status: IN USE.
```

The output, in Example 2-35, shows the sample session from the Tivoli Storage Manager client to the TSM_AIX Tivoli Storage Manager server to the primary storage pool DISKPOOL. The ANS1114I message is displayed at the client end when new volumes are requested either from the COPYPOOL or SUNPOOL storage pool during the selective backup.

Example 2-35 Selective backup with simultaneous write to local FILE and remote SERVER copy pools

```
tsm> s /usr/*
Selective Backup function invoked.

Directory-->             1,024 /usr/ [Sent]
Directory-->             1,024 /usr/4lib [Sent]
Directory-->               512 /usr/aset [Sent]
Directory-->             8,192 /usr/bin [Sent]
Directory-->               512 /usr/ccs [Sent]
Directory-->               512 /usr/demo [Sent]
...<lines omitted>...
ANS1114I Waiting for mount of offline media.
Retry # 1  Directory-->             1,024 /usr/ [Sent]
Retry # 1  Directory-->             1,024 /usr/4lib [Sent]
Retry # 1  Directory-->               512 /usr/aset [Sent]
Retry # 1  Directory-->             8,192 /usr/bin [Sent]
Retry # 1  Directory-->               512 /usr/ccs [Sent]
Retry # 1  Directory-->               512 /usr/demo [Sent]
...<lines omitted>...
ANS1114I Waiting for mount of offline media.
Retry # 1  Symbolic Link-->            15 /usr/pub [Sent]
Retry # 1  Symbolic Link-->            12 /usr/spool [Sent]
Retry # 1  Symbolic Link-->            11 /usr/src [Sent]
Retry # 1  Symbolic Link-->            10 /usr/tmp [Sent]
Retry # 1  Symbolic Link-->            16 /usr/dict [Sent]
Retry # 1  Symbolic Link-->             5 /usr/5bin [Sent]
Retry # 1  Symbolic Link-->            11 /usr/man [Sent]
Retry # 1  Symbolic Link-->             7 /usr/java [Sent]
Retry # 1  Symbolic Link-->            10 /usr/adm [Sent]
Normal File-->       31,948,800 /usr/.CPR  ** Unsuccessful **
ANS1114I Waiting for mount of offline media.
Retry # 1  Normal File-->       31,948,800 /usr/.CPR [Sent]
Selective Backup processing of '/usr/*' finished without failure.


Total number of objects inspected:      50
```

```
Total number of objects backed up:       50
Total number of objects updated:          0
Total number of objects rebound:          0
Total number of objects deleted:          0
Total number of objects expired:          0
Total number of objects failed:           0
Total number of bytes transferred:    60.94 MB
Data transfer time:                  377.01 sec
Network data transfer rate:          165.53 KB/sec
Aggregate data transfer rate:        158.66 KB/sec
Objects compressed by:                    0%
Elapsed processing time:           00:06:33
```

2.3 Cyclic Redundancy Check error detection

IBM Tivoli Storage Manager provides the option of specifying whether a Cyclic Redundancy Check (CRC) is performed during a client or storage agent session with the server or for storage pools. The server validates the data by using CRC which can help identify data corruption. The AIX **sum** command uses a similar method of performing CRC as IBM Tivoli Storage Manager does, which is based on the Berkeley Software Distribution (BSD) 4.3 **sum** command and uses a 32-bit CRC. The BSD command is based on the algorithm for calculating CRC as described in pseudo code in an article in *Communications of the ACM* by D. V. Sarwate (*Communications of the ACM, Volume 31, Issue 8, August 1988, Computation of cyclic redundancy checks via table look-up*, please refer to http://portal.acm.org).

CRC - Cyclic Redundancy Check

CRC between the TSM Server and a Client, Storage Agent or Storage pool volume

- New communication and SAN hardware products are more susceptible to data loss, thus the need for checksums
- Can be used in problem determination to isolate possible data transfer problems
- CRCs with the data backed up and restored and can be checked by TSM server and client
- Benefits misconfigured or malfunctioning network or storage hardware
- Additional assurance to the customer that data was successfully handled

Figure 2-3 CRC support

Data Validation with Cyclic Redundancy Checking

"Cyclic Redundancy Check (CRC) provides a simple yet powerful method of error detection during digital data transmission", D. V. Sarwate.

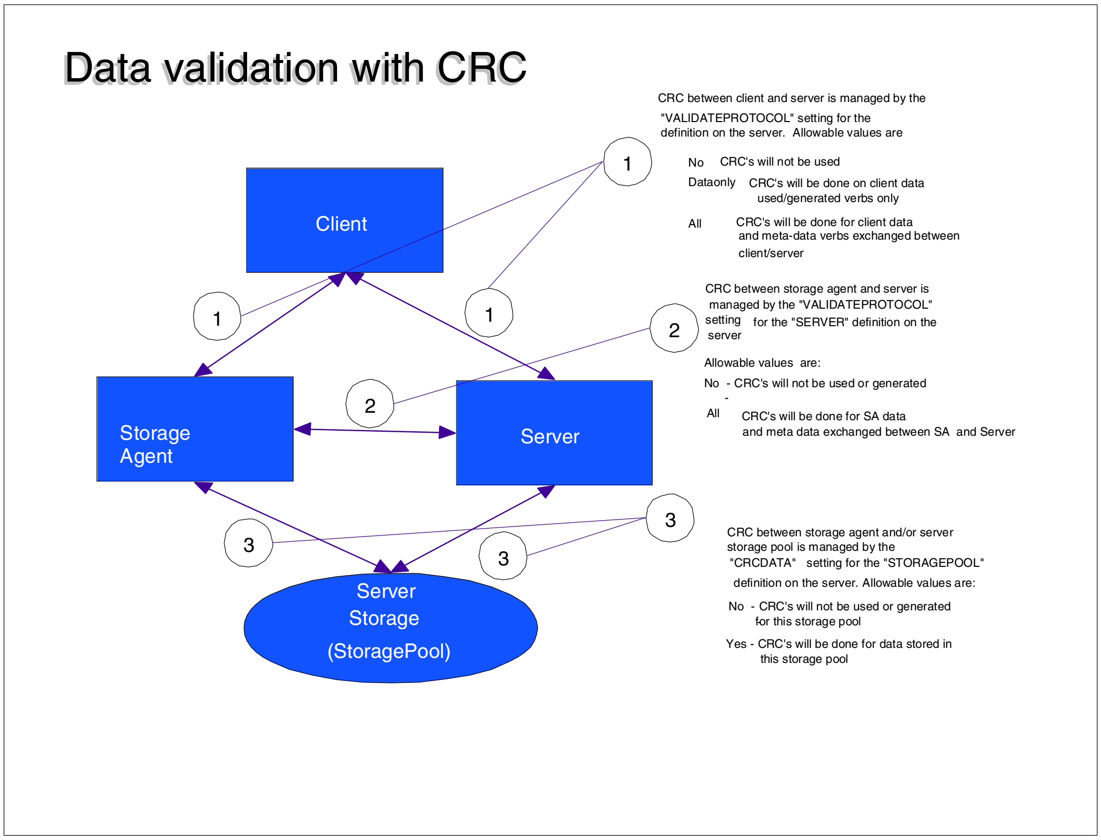

Figure 2-4 Data validation with CRC

Data validation with CRC for IBM Tivoli Storage Manager can be enabled for one or all of the following:

- Tivoli Storage Manager client nodes at Version 5.1.
- Tivoli Storage Manager storage agents at Version 5.1.
- Tivoli Storage Manager Storage pools

Tivoli Storage Manager client nodes

To enable CRC for Tivoli Storage Manager client nodes set the attribute `VALIdateprotocol` to either `All or Dataonly` (recommended). This will validate the data sent between the client and server. `Dataonly` specifies that data validation is to be performed only on file data that is sent between the client and server and does not include the file metadata. The `All` setting also includes the file metadata. However any incorrectly transmitted data between the client and the server will probably be caught by the Tivoli Storage Manager transmission protocol itself so the benefit of using the `All` setting is very small compared to the increased processing overhead it generates.

Tivoli Storage Manager storage agents

To enable CRC for Tivoli Storage Manager storage agents set the attribute `VALIdateprotocol` to `All`. This will validate all client file data, client file metadata, and Tivoli Storage Manager server metadata that is sent between the storage agent and server.

Tivoli Storage Manager Storage pools

To enable CRC for Tivoli Storage Manager storage pools set the attribute `CRCData` to `Yes`. This specifies that data is stored containing CRC information and will allow validation of storage pool data when audit volume processing occurs on the server (when running **`AUDIT VOLUME`**).

> **Note:** Data stored in storage pools is only CRC'd when it is written by a "STORE" operation. These verbs are typically client backup/archive/hsm operations. If the data is originally stored in a storage pool *without* `CRCData` set to `Yes` and then moved to a storage pool that has `CRCData` set to `Yes`, by **`MIGRATION`** or **`MOVE DATA`**, it will *not* have CRC generated.
>
> The primary storage pool for the client has to have `CRCData` set to `Yes` for the CRC operation to be performed. The CRC will then continue to exist with the clients data even it it is migrated or moved to other storage pools that does not have `CRCData` set to `Yes`.

Performance impact

The performance impact, of using the additional level of data integrity protection that CRC allows, is usually approximately five percent (5%) for servers, storage pools and clients using the `Dataonly` setting. However when more advanced protocols such as the LAN-free protocol is used, the logic gets much more complex because some of the Tivoli Storage Manager transmission protocol verbs are altered before being forwarded to the server.

Data validation impacts performance because the server requires additional CPU overhead to calculate and compare CRC values, however this is done without increasing the CPU utilization so it will usually result in the overhead causing additional elapsed time for the transmission to complete. When using CRC to validate storage pool data, there is no performance impact on the clients that own the data stored in the storage pool.

How to enable CRC for a client

In the following example we will add a client node to the IBM Tivoli Storage Manager server, then we enable CRC for the node to be used during the backup/restore transmissions.

The sample in Example 2-36 shows how to update a client node information to enable the CRC for `Dataonly`.

Example 2-36 Change the node to use CRC validation protocol for data only

```
root@crete:/: dsmadmc -id=admin -password=admin update node one-on-one VALIdate=Dataonly
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSMC: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/01/02   17:43:37  Last access: 03/01/02   17:39:48

ANS8000I Server command: 'update node one-on-one VALIdate=Dataonly'
ANR2063I Node ONE-ON-ONE updated.

ANS8002I Highest return code was 0.
```

The next sample, in Example 2-37, shows how to examine the client node definitions in the IBM Tivoli Storage Manager server.

Example 2-37 Examining the updated client node definition

```
root@crete:/: dsmadmc -id=admin -password=admin query node one-on-one f=d
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSMC: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/01/02   17:52:04  Last access: 03/01/02   17:50:14

ANS8000I Server command: 'query node one-on-one f=d'

                     Node Name: ONE-ON-ONE
                      Platform: WinNT
               Client OS Level: 4.00
                Client Version: Version 5, Release 1, Level 0.0
            Policy Domain Name: STANDARD
         Last Access Date/Time: 03/01/02   17:39:44
        Days Since Last Access: <1
        Password Set Date/Time: 03/01/02   17:39:43
       Days Since Password Set: <1
         Invalid Sign-on Count: 0
                       Locked?: No
                       Contact: The Dude
                   Compression: Client
       Archive Delete Allowed?: Yes
        Backup Delete Allowed?: No
        Registration Date/Time: 03/01/02   17:39:43
     Registering Administrator: OPEN_REGISTRATION
Last Communication Method Used:
   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
                     Optionset:
                           URL:
                     Node Type: Client
    Password Expiration Period:
             Keep Mount Point?: No
  Maximum Mount Points Allowed: 1
        Auto Filespace Rename : No
             Validate Protocol: DataOnly

ANS8002I Highest return code was 0.
```

How to enable CRC for a storage pool

In the following example we add a storage pool to the IBM Tivoli Storage Manager server with CRC enabled, then we show how to audit the volumes in the storage pool, utilizing the stored CRC information.

First we create a storage pool named DISKPOOL with CRC checking enabled.

Example 2-38 Create a storagepool with CRCData=Yes

```
root@crete:/: dsmadmc -id=admin -password=admin define stgpool DISKPOOL DISK CRCData=Yes
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
```

```
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/04/02   15:53:55  Last access: 03/04/02   15:45:39

ANS8000I Server command: 'define stgpool DISKPOOL DISK CRCData=Yes'
ANR2200I Storage pool DISKPOOL defined (device class DISK).

ANS8002I Highest return code was 0.
```

Example 2-39 shows how to examine the storagepool settings for the `DISKPOOL` storage pool, and as can be seen the `CRC Data` settings is `Yes`.

Example 2-39   Examining detailed storagepool information

```
root@crete:/: dsmadmc -id=admin -password=admin query stg diskpool f=d
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/04/02   15:46:36  Last access: 03/04/02   15:45:39

ANS8000I Server command: 'q stg diskpool f=d'

               Storage Pool Name: DISKPOOL
               Storage Pool Type: Primary
               Device Class Name: DISK
         Estimated Capacity (MB): 100.0
                        Pct Util: 92.0
                        Pct Migr: 91.9
                     Pct Logical: 100.0
                    High Mig Pct: 90
                     Low Mig Pct: 70
                 Migration Delay: 0
              Migration Continue: Yes
             Migration Processes: 1
               Next Storage Pool:
            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): 0.00
Elapsed Migration Time (seconds): 0
        Reclamation in Progress?:
 Volume Being Migrated/Reclaimed:
  Last Update by (administrator): ADMIN
Last Update Date/Time: 03/04/02   14:19:11
        Storage Pool Data Format: Native
            Copy Storage Pool(s):
         Continue Copy on Error?:
```

```
                        CRC Data: Yes

ANS8002I Highest return code was 0.
```

To perform the audit process on volumes in the storage pool we could specify the volumes we want to audit, and in this case we would need to know the volume names. This is shown in Example 2-40 on all volumes.

Example 2-40 Determining the volumes in the storage pool

```
root@crete:/: dsmadmc -id=admin -password=admin query volume stgpool=diskpool
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/04/02   16:10:48  Last access: 03/04/02   16:05:13

ANS8000I Server command: 'query volume stgpool=diskpool'

Volume Name               Storage      Device      Estimated    Pct   Volume
                          Pool Name    Class Name   Capacity   Util   Status
                                                        (MB)
------------------------  -----------  ----------  ---------  -----  --------
/tsm/stg/d/disk#1.dsm     DISKPOOL     DISK             50.0  100.0  On-Line
/tsm/stg/d/disk#2.dsm     DISKPOOL     DISK             50.0  100.0  On-Line
/tsm/stg/d/disk#3.dsm     DISKPOOL     DISK            100.0   19.8  On-Line

ANS8002I Highest return code was 0.
```

After data from clients has been stored in the storage pool, the CRC information will be used by the *audit* process when the **AUDIT VOLUME** process is initiated on the IBM Tivoli Storage Manager server as the output sample in Example 2-41 shows.

Example 2-41 Performing the audit of the volumes in the storage pool

```
root@crete:/: dsmadmc -id=admin -password=admin audit volume stgpool=diskpool fix=yes
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/04/02   16:33:38  Last access: 03/04/02   16:32:58

ANS8000I Server command: 'audit volume stgpool=diskpool fix=yes'

ANS8002I Highest return code was 0.
```

To monitor the progress of the audit process use the **query process** command options with the **dsmadmc** command as the sample output in Example 2-42 shows.

Example 2-42 Monitoring audit of the volumes in the storage pool from the process list

```
root@crete:/: dsmadmc -id=admin -password=admin query process
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.
```

```
Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/04/02   16:33:39  Last access: 03/04/02   16:32:58

ANS8000I Server command: 'query process'

 Process Process Description  Status
  Number
-------- -------------------- -------------------------------------------------
       7 Audit Volume         Volume /tsm/stg/d/disk#1.dsm (storage pool
          (Repair)             DISKPOOL), Files Processed: 0, Damaged Files
                               Deleted: 0, Partial Files Skipped: 0. Current
                               Physical File (bytes): 96,411,164

ANS8002I Highest return code was 0.
```

To examine the audit process during or after it has completed, use the **query actlog** command options with the **dsmadmc** command as the sample output in Example 2-43 shows.

Example 2-43 Monitoring audit of the volumes in the storage pool from the actlog

```
root@crete:/: dsmadmc -id=admin -password=admin query actlog s=audit
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/04/02   16:34:08  Last access: 03/04/02   16:32:58

ANS8000I Server command: 'query actlog s=audit'

Date/Time            Message
-------------------- ----------------------------------------------------------
03/04/02   16:33:38  ANR2017I Administrator ADMIN issued command: AUDIT VOLUME
                      stgpool=diskpool fix=yes quiet=no
03/04/02   16:33:38  ANR0984I Process 7 for AUDIT VOLUME (REPAIR) started in
                      the BACKGROUND at 16:33:38.
03/04/02   16:33:38  ANR2312I Audit Volume (Repair) process started for volume
                      /tsm/stg/d/disk#1.dsm (process ID 7).
03/04/02   16:33:43  ANR4132I Audit volume process ended for volume
                      /tsm/stg/d/disk#1.dsm; 1 files inspected, 0 damaged files
                      deleted, 0 damaged files marked as damaged, 0 objects
                      updated.
03/04/02   16:33:43  ANR0987I Process 7 for AUDIT VOLUME (REPAIR) running in
                      the BACKGROUND processed 1 items with a completion state
                      of SUCCESS at 16:33:43.
03/04/02   16:33:44  ANR0984I Process 8 for AUDIT VOLUME (REPAIR) started in
                      the BACKGROUND at 16:33:44.
03/04/02   16:33:44  ANR2312I Audit Volume (Repair) process started for volume
                      /tsm/stg/d/disk#2.dsm (process ID 8).
03/04/02   16:33:51  ANR4132I Audit volume process ended for volume
                      /tsm/stg/d/disk#2.dsm; 5 files inspected, 0 damaged files
                      deleted, 0 damaged files marked as damaged, 0 objects
                      updated.
03/04/02   16:33:51  ANR0987I Process 8 for AUDIT VOLUME (REPAIR) running in
                      the BACKGROUND processed 5 items with a completion state
                      of SUCCESS at 16:33:51.
03/04/02   16:33:51  ANR0984I Process 9 for AUDIT VOLUME (REPAIR) started in
```

```
                    the BACKGROUND at 16:33:51.
03/04/02   16:33:51  ANR2312I Audit Volume (Repair) process started for volume
                    /tsm/stg/d/disk#3.dsm (process ID 9).
03/04/02   16:33:53  ANR4132I Audit volume process ended for volume
                    /tsm/stg/d/disk#3.dsm; 3 files inspected, 0 damaged files
                    deleted, 0 damaged files marked as damaged, 0 objects
                    updated.
03/04/02   16:33:53  ANR0987I Process 9 for AUDIT VOLUME (REPAIR) running in
                    the BACKGROUND processed 3 items with a completion state
                    of SUCCESS at 16:33:53.
03/04/02   16:34:08  ANR2017I Administrator ADMIN issued command: QUERY ACTLOG
                    s=audit

ANS8002I Highest return code was 0.
```

2.4 Move data by node

IBM Tivoli Storage Manager 5.1 introduces the possibility to move storage pool stored data by node and node filespace in addition to moving storage pool volume data.

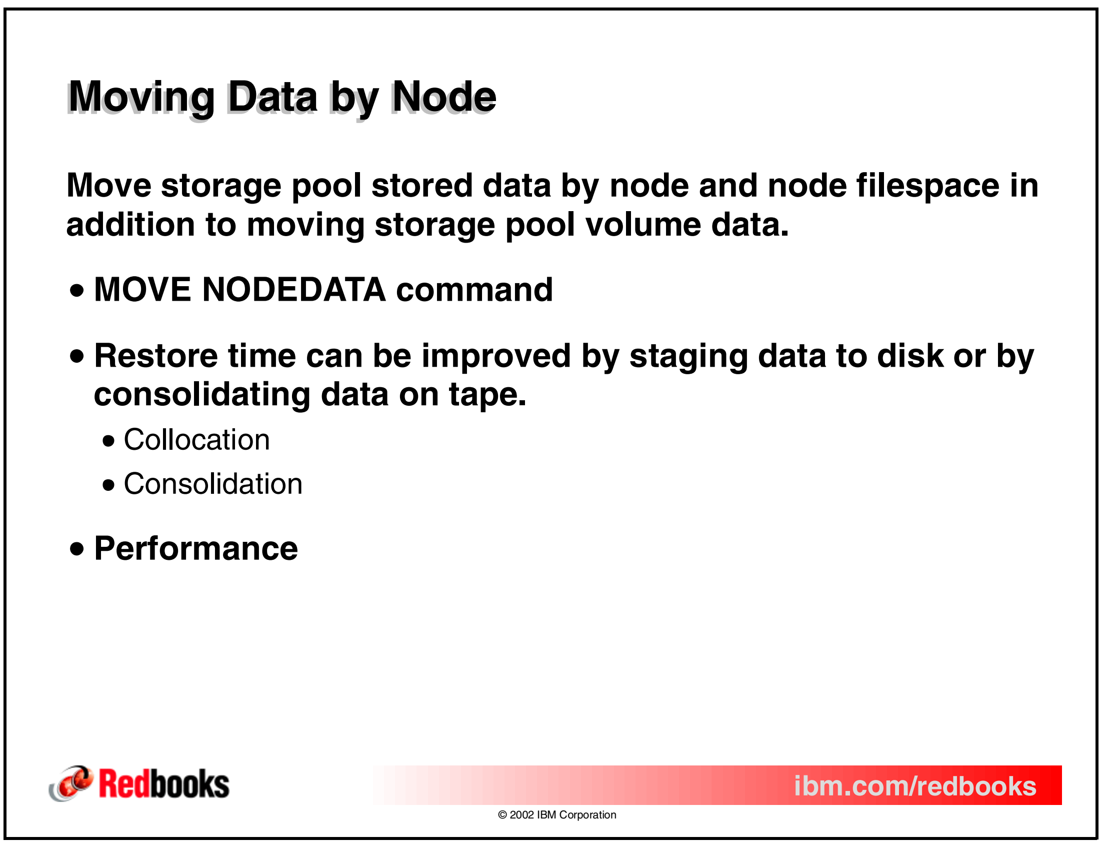

Figure 2-5 Move node data

The MOVE NODEDATA command can be used to move data from a sequential access storage pool. The destination can be either a sequential or random access storage pool. The moved storage pool data can be either all data belonging to a node or only specific filespaces belonging to that node.

Note: `MOVE NODEDATA` will move *all* data and not only the active data.

The main objectives for the `MOVE NODEDATA` command is:

Collocation If a sequential access storage pool has had `COLLOCATION` turned off during a period when the storage pool has been used to save client data, then the clients data will not be collocated on the sequential media but rather scattered on different volumes. By using `MOVE NODEDATA` this will allow manual collocation to be performed and this will reduce the number of volume mounts required during a restore operation

Consolidation In anticipation of a client restore request, `MOVE NODEDATA` can be used to move the clients data from a sequential access storage pool to a random access storage pool to eliminate the need for mountpoints during the actual restore process and also provides the fastest media to restore the client from.

How MOVE NODEDATA work

The `MOVE NODEDATA` process creates a list of nodes and filespaces to move based on the user's criteria specified in the `MOVE NODEDATA` command and starts a queue thread to determine a list of volumes to process. The volumes added to this list have to have an access of `READWRITE` or `READONLY`. The volume list created also becomes the list of volumes to exclude as output volumes. Once the queue thread is complete a process thread is started. The number of threads started is determined by the `MAXPROCESS` parameter.

Each thread started performs the following:

- Selects a volume from the volume list
- Processes each bit file in the volume to determine whether it should be moved
- Uses the move data functions to batch up the file and to move the bit file

The `MOVE NODEDATA` command uses a combination of the `FILESPACE`, `UNIFILESPACE`, and `FSID` parameters. Note that the `FSID` can be a filespace ID for a non-unicode filespace or unicode filespace.

Performance considerations for MOVE NODEDATA

How long a `MOVE NODEDATA` will take depends on a number of things such as the type of devices, the amount of data to be moved and also how many sequential access volumes, for example tapes, which are used to store the client's data.

Note: When moving data within the same storage pool note that the destination volumes cannot be the same volumes that have data to be moved. It is important to ensure that there is enough space in the other available volumes in the storage pool to contain the data that will be moved. Each volume needs a mountpoint to be accessible.

How to use MOVE NODEDATA

The following example shows how to move all data for a client named `ONE-ON-ONE` from the storage pool `SPOOLPOOL` to the random access storage pool `DISKPOOL`.

Let us first examine the occupancy for the node `ONE-ON-ONE` as shown in the sample in Example 2-44.

Example 2-44 dsmadmc query occupancy

```
root@crete:/: dsmadmc -id=admin -password=admin query occ one-on-one
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/08/02   18:39:40  Last access: 03/08/02   18:36:28

ANS8000I Server command: 'query occ one-on-one'

Node Name   Type  Filespace    FSID  Storage     Number of   Physical    Logical
                  Name               Pool Name       Files      Space      Space
                                                             Occupied   Occupied
                                                                 (MB)       (MB)
----------  ----  ----------  -----  ----------  ---------  ---------  ---------
ONE-ON-ONE  Bkup  \\a23bk41-      5  COPYPOOL            1      91.94      91.94
                   z\c$
ONE-ON-ONE  Bkup  \\a23bk41-      5  SPOOLPOOL           7     119.76     119.76
                   z\c$

ANS8002I Highest return code was 0.
```

The output in Example 2-44 shows that the node `ONE-ON-ONE` have one filespace in two storage pools. We are interested in the data in the `SPOOLPOOL` storage pool. Example 2-45 determines to which device class the storage pool belongs.

Example 2-45 dsmadmc query stg

```
root@crete:/: dsmadmc -id=admin -password=admin query stg spoolpool
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/08/02   18:42:07  Last access: 03/08/02   18:36:28

ANS8000I Server command: 'query stg spoolpool'

Storage       Device        Estimated    Pct     Pct   High   Low   Next Stora-
Pool Name     Class Name     Capacity   Util    Migr    Mig   Mig   ge Pool
                                 (MB)                   Pct   Pct
-----------   ----------   ----------  -----   -----   ----   ---   -----------
SPOOLPOOL     SPOOL        400,000,00    0.0     0.0     90    70
                                  0.0

ANS8002I Highest return code was 0.
```

The `SPOOLPOOL` storage pool belongs to the `SPOOL` device class which we examine in Example 2-46.

Example 2-46 dsmadmc query devc

```
root@crete:/: dsmadmc -id=admin -password=admin query devc spool f=d
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.
```

```
Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/08/02   18:44:29  Last access: 03/08/02   18:36:28

ANS8000I Server command: 'query devc spool f=d'

             Device Class Name: SPOOL
        Device Access Strategy: Sequential
            Storage Pool Count: 1
                   Device Type: FILE
                        Format: DRIVE
         Est/Max Capacity (MB): 4.0
                   Mount Limit: 1
              Mount Wait (min):
         Mount Retention (min):
                  Label Prefix:
                       Library:
                     Directory: /tsm/spool
                   Server Name:
                  Retry Period:
                Retry Interval:
                        Shared:
Last Update by (administrator): ADMIN
         Last Update Date/Time: 02/28/02   10:35:24

ANS8002I Highest return code was 0.
```

The output in Example 2-46 shows that the `SPOOL` device class is a `Sequential FILE` device type with the `/tsm/spool` directory as its storage area for files that each are allowed to be `4.0` MB large.

Let us now examine the `DISKPOOL` storage pool (Example 2-47) before we start performing the `MOVE NODEDATA` operation.

Example 2-47   dsmadmc query stg

```
root@crete:/: dsmadmc -id=admin -password=admin query stg diskpool
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/08/02   18:47:05  Last access: 03/08/02   18:46:55

ANS8000I Server command: 'query stg diskpool'

Storage       Device       Estimated    Pct     Pct   High   Low   Next Stora-
Pool Name     Class Name    Capacity   Util    Migr    Mig   Mig   ge Pool
                                (MB)                   Pct   Pct
-----------   ----------  ----------  -----   -----   ----   ---   -----------
DISKPOOL      DISK             200.0    0.0     0.0     90    70   SPOOLPOOL

ANS8002I Highest return code was 0.
```

The output in Example 2-47 shows that the `DISKPOOL` storage pool is empty so let us move the data for the client node `ONE-ON-ONE` from the storage pool `SPOOLPOOL` to the storage pool `DISKPOOL` with `MOVE NODEDATA` which is shown in the next sample in Example 2-48.

Example 2-48 dsmadmc move nodedata

```
root@crete:/: dsmadmc -id=admin -password=admin move nodedata one-on-one
fromstgpool=SPOOLPOOL tostgpool=DISKPOOLL filespace='*' wait=yes
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/08/02   18:30:13  Last access: 03/08/02   18:26:18

ANS8000I Server command: 'move nodedata one-on-one fromstgpool=spoolpool tostgpoo
l=diskpool wait=yes'
ANR0984I Process 22 for MOVE NODE DATA started in the FOREGROUND at 18:30:13.
ANR1284I Move node data started as process 22.
ANR1288I Move node data process 22 ended for storage pool SPOOLPOOL.
ANR0986I Process 22 for MOVE NODE DATA running in the FOREGROUND processed 7 items for a
total of 125,573,701 bytes with a completion state of SUCCESS at 18:30:41.
ANR1290I Move node data from storage pool SPOOLPOOL to storage pool DISKPOOL has ended.
Files Moved: 7, Bytes Moved: 125573701, Unreadable Files: 0, Unreadable Bytes: 0.

ANS8002I Highest return code was 0.
```

In Example 2-48, 119 MB were moved from storage pool `SPOOLPOOL` to the storage pool `DISKPOOL`. Let us now examine the storage pool `DISKPOOL` in the next sample in Example 2-49.

Example 2-49 dsmadmc query stg

```
root@crete:/: dsmadmc -id=admin -password=admin query stg diskpool
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/08/02   18:51:38  Last access: 03/08/02   18:46:55

ANS8000I Server command: 'query stg diskpool'

Storage       Device       Estimated     Pct     Pct   High   Low   Next Stora-
Pool Name     Class Name    Capacity    Util    Migr    Mig   Mig   ge Pool
                                (MB)                    Pct   Pct
-----------   ----------  ----------   -----   -----   ----   ---   -----------
DISKPOOL      DISK             200.0    59.9    59.9     90    70   SPOOLPOOL

ANS8002I Highest return code was 0.
```

In Example 2-49, the output above shows that we have approximately 60 percent utilization of 200 MB in the `DISKPOOL` storage pool which is approximately 120 MB which is what we moved from the `SPOOLPOOL` storage pool. Let us again examine the occupancy for the node `ONE-ON-ONE` as shown in the sample in Example 2-50.

Example 2-50 dsmadmc query occ

```
root@crete:/: dsmadmc -id=admin -password=admin query occ one-on-one
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.
```

```
Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/08/02   19:02:53  Last access: 03/08/02   18:57:55

ANS8000I Server command: 'query occ one-on-one'

Node Name   Type  Filespace    FSID  Storage     Number of   Physical    Logical
                  Name               Pool Name       Files      Space      Space
                                                             Occupied   Occupied
                                                                 (MB)       (MB)
----------  ----  ----------  -----  ----------  ---------  ---------  ---------
ONE-ON-ONE  Bkup  \\a23bk41-      5  COPYPOOL            1      91.94      91.94
                   z\c$
ONE-ON-ONE  Bkup  \\a23bk41-      5  DISKPOOL            7     119.77     119.77
                   z\c$

ANS8002I Highest return code was 0.
```

The output in Example 2-50 shows that the data for the `ONE-ON-ONE` node has been moved to the `DISKPOOL` storage pool.

When we want to move the data for the client `one-on-one` back to the `SPOOLPOOL` or any other pool we cannot use `MOVE NODEDATA` since the `DISKPOOL` storage pool is not of a sequential device type. We must use `MOVE DATA` or lower the migration thresholds to initiate migration. The following **dsmadmc** command will lower the high and low watermarks of utilization so that migration is initiated:

```
update stg DISKPOOL hi=0 lo=0
```

The migration will then start to move everything from the `DISKPOOL` storage pool to the next storage pool in the hierarchy which is the `SPOOLPOOL` storage pool, which Example 2-51 shows.

Example 2-51   dsmadmc query pr

```
root@crete:/: dsmadmc -id=admin -password=admin query pr
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/08/02   18:57:57  Last access: 03/08/02   18:57:55

ANS8000I Server command: 'query pr'

 Process Process Description  Status
  Number
-------- -------------------- -------------------------------------------------
      26 Migration            Disk Storage Pool DISKPOOL, Moved Files: 4, Moved
                               Bytes: 5,488,640, Unreadable Files: 0,
                               Unreadable Bytes: 0. Current Physical File
                               (bytes): 96,411,648

                               Current output volume: /tsm/spool/0000007B.BFS.


ANS8002I Highest return code was 0.
```

Let us now examine the `DISKPOOL` storage pool to verify that it is empty which is shown in Example 2-52.

Example 2-52 dsmadmc query stg

```
root@crete:/: dsmadmc -id=admin -password=admin query stg diskpool
root@crete:/tsm/pbin: d q stg diskpool
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/08/02   19:06:22  Last access: 03/08/02   18:57:55

ANS8000I Server command: 'query stg diskpool'

Storage       Device       Estimated    Pct     Pct   High   Low   Next Stora-
Pool Name     Class Name    Capacity   Util    Migr    Mig   Mig   ge Pool
                                (MB)                   Pct   Pct
-----------   ----------  ----------  -----   -----   ----   ---   -----------
DISKPOOL      DISK             200.0    0.0     0.0      0     0   SPOOLPOOL

ANS8002I Highest return code was 0.
```

The output in Example 2-52 verifies that the storagepool `DISKPOOL` is indeed empty.

Note: Reset the high and low thresholds.

Before moving any data from a storage pool defined as `DISKPOOL` and `SPOOLPOOL` above, ensure that the high and low migration values are properly set prior to initiating any `MOVE NODEDATA` operations to the storage pool defined as `DISKPOOL` above, or the Tivoli Storage Manager server will move the data and immediately migrate the data as can be seen in Example 2-53.

Example 2-53 dsmadmc query pr

```
root@crete:/: dsmadmc -id=admin -password=admin query pr
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/08/02   19:07:38  Last access: 03/08/02   18:57:55

ANS8000I Server command: 'query pr'

 Process Process Description  Status
  Number
-------- -------------------- -------------------------------------------------
      29 Move Node Data       Storage Pool SPOOLPOOL, Target Pool DISKPOOL
                               Files Moved: 4, Bytes Moved: 5,483,535,
                               Unreadable Files: 0, Unreadable Bytes: 0.
                               Current Physical File (bytes): 96,411,164
      30 Migration            Disk Storage Pool DISKPOOL, Moved Files: 0, Moved
                               Bytes: 0, Unreadable Files: 0, Unreadable Bytes:
                               0. Current Physical File (bytes): 24,576
```

```
                         Waiting for mount point in device class SPOOL
                         (16 seconds).

ANS8002I Highest return code was 0.
```

Examples

The Tivoli Storage Manager server has a storage pool `BACKPOOL2` with the volumes `VOL1`, `VOL2`, `VOL3` and `VOL5`. The data for node `NOAH` is in `VOL1`, `VOL2`, and `VOL5`.

The following **dsmadmc** command moves the data for the client `NOAH` from the `DISKPOOL` storage pool:

```
move nodedata NOAH fromstg=DISKPOOL
```

`VOL1`, `VOL2`, and `VOL5` will not be considered as candidates as destination volumes by the move nodedata process. Here are some simple examples of combinations used and the resulting data that will be moved. The filespace distribution for the node `RON` is shown in Figure 2-6.

```
TSM:NT>
q filespace ron
ANR2017I Administrator SERVER_CONSOLE issued command: QUERY FILESPACE ron

Node     Filespace  FSID Platform Filespace  Is Filespace   Capacity  Pct
Name     Name                     Type       Unicode?       (MB)      Util

-------- ---------- ---- -------- --------   ---------      -------   -----
RON      \\jee\d$   1    WinNT    NTFS       No             4,094.7   97.4
RON      \\jee\c$   2    WinNT    FAT        Yes            2,047.0   87.8
RON      \\jee\e$   3    WinNT    NTFS       Yes            6,149.9   97.1

TSM:NT>
```

Figure 2-6 Move node data filespaces

The client filespaces, shown in Figure 2-6, are stored in the `DISKPOOL` storage pool.

The following **dsmadmc** command moves all non-unicode and unicode filespaces for node `ron`:

```
MOVE NODEDATA ron FROMSTG=DISKPOOL
```

The following **dsmadmc** command will cause only non-unicode filespace `\\jee\d$` to be moved for node `ron`. No unicode filespaces will be moved:

```
MOVE NODEDATA ron  FROMSTG=DISKPOOL FILESPACE=\\jee\d$
```

The following **dsmadmc** command will cause non-unicode filespace `\\jee\d$` and unicode filespace name `\\jee\e$` to be moved:

```
MOVE NODEDATA ron FROMSTG=DISKPOOL FILESPACE=\\jee\d$ UNIFILESPACE=\\jee\e$
```

The following **dsmadmc** command will cause non-unicode filespace `\\jee\d$` and unicode filespaces `\\jee\c$` and `\\jee\e$` to be moved:

```
MOVE NODEDATA ron FROMSTG=DISKPOOL FILESPACE=\\jee\d$ FSID=2,3
```

The `MOVE NODEDATA` operation may sometimes appear to be incomplete for a number of reasons. Please refer to “Moving Data by Node” in *Tivoli Storage Manager V5.1 Administrator’s Guide* or “Move Nodedata” in *Tivoli Storage Manager V5.1 Administrator’s Reference* for further details.

2.5 Server trace enhancements

A new trace keyword has been introduced in Tivoli Storage Manager V5.1:

```
TRACE MAXSIZE nn
```

Where ‘nn’ is the number of megabytes that the trace output file can occupy. Once this limit has been reached the trace will wrap. Other changes to the Tivoli Storage Manager trace facility are:

- All server trace messages will now contain a timestamp without the need for the `SYSTIME` trace class to be enabled
- The ability to trace just a thread ID, session ID, process ID, or client node is also introduced
- The server trace now follows a fixed format

Server trace enhancements

Server Trace Enhancements

- New trace key word - "TRACE MAXSIZE nn"
  - **'nn' is the number of megabytes that the file can consume**
  - **Once this limit is reached, the trace file will "wrap"**
- All trace messages will now include time stamps
  - **Previously available only when the SYSTIME trace class was enabled**
- Traces to a file format as:
  "trace message number, timestamp, arguments for message"
  - **Messages are assigned a numeric value (trace msg id)**

Figure 2-7 Server trace enhancements

The following example is part of the result from performing a Tivoli Storage Manager trace, using either the **dsmadmc** command or the HTTP GUI command line option, during a Server-Free backup:

```
trace enable sf session tcpinfo
trace begin tsmtrace.out
```

```
trace end
```

Example 2-54 Server trace example

```
<...lines omitted...>
11:22:49.250 [37][tcpcomm.c][838]: tcpSend: 4 bytes moved into buffer
11:22:49.250 [37][tcpcomm.c][1799]: Sending 4 bytes of data to client at 9.1.38.151 on
socket 1114952, port 4063.
11:22:49.250 [37][tcpcomm.c][964]: tcpRecv: 4 bytes read
11:22:49.250 [37][tcpcomm.c][964]: tcpRecv: 8 bytes read
11:22:49.250 [37][tcpcomm.c][964]: tcpRecv: 50 bytes read
11:22:49.250 [37][smtrans.c][790]: Recv Verb: Session 52, Length=62, Code=20800,
Type=RemoteOpProgress.
11:22:49.250 [37][smtrans.c][528]: Send Verb: Session 52, Length=73, Code=20900,
Type=RemoteOpProgressResp.
11:22:49.265 [37][tcpcomm.c][838]: tcpSend: 73 bytes moved into buffer
11:22:49.265 [37][tcpcomm.c][1799]: Sending 73 bytes of data to client at 9.1.38.151 on
socket 1114952, port 4063.
11:22:51.046 [45][sfcopy.c][6694]: CCB: Calling Thread ID [45]
11:22:51.046 [45][sfcopy.c][6729]: CCB: Data from CA: Bytes Read: 0 Bytes Written:
1100218208 Blocks Read: 0 Blocks Written: 1100218368
11:22:51.046 [45][sfcopy.c][6766]: CCB: Received 0 from Copy Agent for sector 7.
11:22:51.046 [45][sfcopy.c][6792]: CCB: Sector 7 Successful.  0,671088640 bytes 0,20480
blocks written. 0,671006344 bytes 0,1310736 blocks read.
11:22:51.062 [45][sfcopy.c][6339]: Perform Copy: Attempt: 1 of 10 max attempts. Copy Agent
returned 5
11:22:51.062 [45][sfcopy.c][6341]: Perform Copy: Using DM handle:30  Source Handle:0
Target Handle:32
<...lines omitted...>
```

For additional information regarding Tivoli Storage Manager tracing, please refer to: http://www.tivoli.com/support.

3

Platform specific server enhancements

This chapter discusses the new features and enhancements provided in the IBM Tivoli Storage Manager Version 5.1 servers, which are specific to certain platforms.

3.1 AIX

The IBM Tivoli Storage Manager 5.1 for IBM AIX introduces several performance enhancements and HACMP support, specific to the IBM AIX operating system, in addition to the general Version 5.1 server enhancements.

3.1.1 Changes in IBM Tivoli Storage Manager Version 5.1 server

The following changes have been made to IBM Tivoli Storage Manager 5.1 for IBM AIX.

AIX

Changes in IBM Tivoli Storage Manger version 5.1 server

- Changes in defining drives and libraries
- Path names for external libraries
- Moving data by node
- Support for Simultaneous writes to primary and copy storage pools
- High Availability Cluster Multiprocessing
- Tivoli Data Protection for new NDMP support
- Data validation with Cyclic Redundancy Checking
- New licensing method
- Server performance enhancements

Figure 3-1 Changes to IBM Tivoli Storage Manager 5.1 for IBM AIX

- Changes in Defining Drives and Libraries

  Device names are now specified in the DEFINE PATH and UPDATE PATH commands, rather than in the DEFINE DRIVE, UPDATE DRIVE, DEFINE LIBRARY, and UPDATE LIBRARY commands

- Path names for external libraries are now specified in the DEFINE PATH command.

  See the following new commands:

  – DEFINE PATH
  – DELETE PATH
  – QUERY PATH
  – UPDATE PATH

  See the following changed commands:

  – DEFINE DRIVE
  – DEFINE LIBRARY
  – QUERY DRIVE
  – QUERY LIBRARY

- UPDATE DRIVE
- UPDATE LIBRARY

- Moving Data by Node

  You can now use the MOVE NODEDATA command to move data in a sequential-access storage pool for one or more nodes, or move data in a single node with selected file spaces. You can also use MOVE NODEDATA to consolidate data for off-site disaster recovery storage, or move data to another storage pool.

- Support for Simultaneous Writes to Primary and Copy Storage Pools

  You can specify copy storage pools in a primary storage pool definition. When a client backs up, archives, or migrates a file, the file is written to the primary storage pool and is simultaneously stored into each copy storage pool.

  See the following changed commands:

  - DEFINE STGPOOL
  - QUERY SESSION
  - QUERY STGPOOL
  - REGISTER NODE
  - UPDATE NODE
  - UPDATE STGPOOL

- High Availability Cluster Multiprocessing

  Tivoli Storage Manager can now use High Availability Cluster Multiprocessing (HACMP). HACMP provides the leading AIX-based clustering solution, which allows automatic system recovery during system failure detection. By using HACMP together with Tivoli Storage Manger, you can ensure server availability.

- Tivoli Data Protection for New Network Data Management Protocol Support

  New Network Data Management Protocol (NDMP) support now extends to the AIX (32-bit and 64-bit) Tivoli Storage Manager server platform. The new Tivoli Data Protection for NDMP product supports NDMP backup and restore for network-attached storage (NAS) file servers from IBM and Network Appliance. NDMP allows a network storage-management application to control the backup and restore of an NDMP-compliant file server without installing third-party software on that server. The NAS file server does not require installation of Tivoli Storage Manager software. The Tivoli Storage Manager server uses NDMP to connect to the NAS file server to initiate, control, and monitor a file system backup or restore operation. The NDMP support for NAS file servers enables higher performance backup to tape devices without moving the data over the LAN. TDP for NDMP is a separately priced and licensed product.

  See the following new commands:

  - BACKUP NODE
  - DEFINE DATAMOVER
  - DEFINE PATH
  - DELETE DATAMOVER
  - DELETE PATH
  - QUERY DATAMOVER
  - QUERY PATH
  - UPDATE DATAMOVER
  - UPDATE PATH
  - RESTORE NODE

- Data Validation with Cyclic Redundancy Checking

  Tivoli Storage Manager provides the option of specifying whether a cyclic redundancy check (CRC) is performed during a client session with the server, or for storage pools. The server validates the data by using a cyclic redundancy check which can help identify data corruption. Data validation can be enabled for one or all of the following:

  - Tivoli Storage Manager client nodes at Version 5.1
  - Tivoli Storage Manager storage agents at Version 5.1
  - Storage pools

  See the following changed commands:

  - DEFINE SERVER
  - DEFINE STGPOOL
  - QUERY NODE
  - QUERY SERVER
  - QUERY VOLUME
  - REGISTER NODE
  - UPDATE NODE
  - UPDATE SERVER
  - UPDATE STGPOOL

- New Licensing Method

  The new licensing method enables you to register the exact number of licenses that are required, rather than in increments of 1, 5, 10, and 50.

- Server Performance Enhancements

  There are two new Tivoli Storage Manager performance enhancements:

  - AIX Asynchronous I/O Support. This feature is available via a new option in the server options file.
  - AIX Direct I/O Support. This feature is available via a new option in the server options file.

  See the following changed commands:

  - QUERY OPTION

  See the following new server options:

  - AIXDIRECTIO
  - AIXSYNCIO

3.1.2 HACMP support for Tivoli Storage Manager server on AIX

This section describes how to install and configure a Tivoli Storage Manager server on a system in an AIX HACMP cluster so that, if the system fails, the server will be brought back up on another system in the cluster.

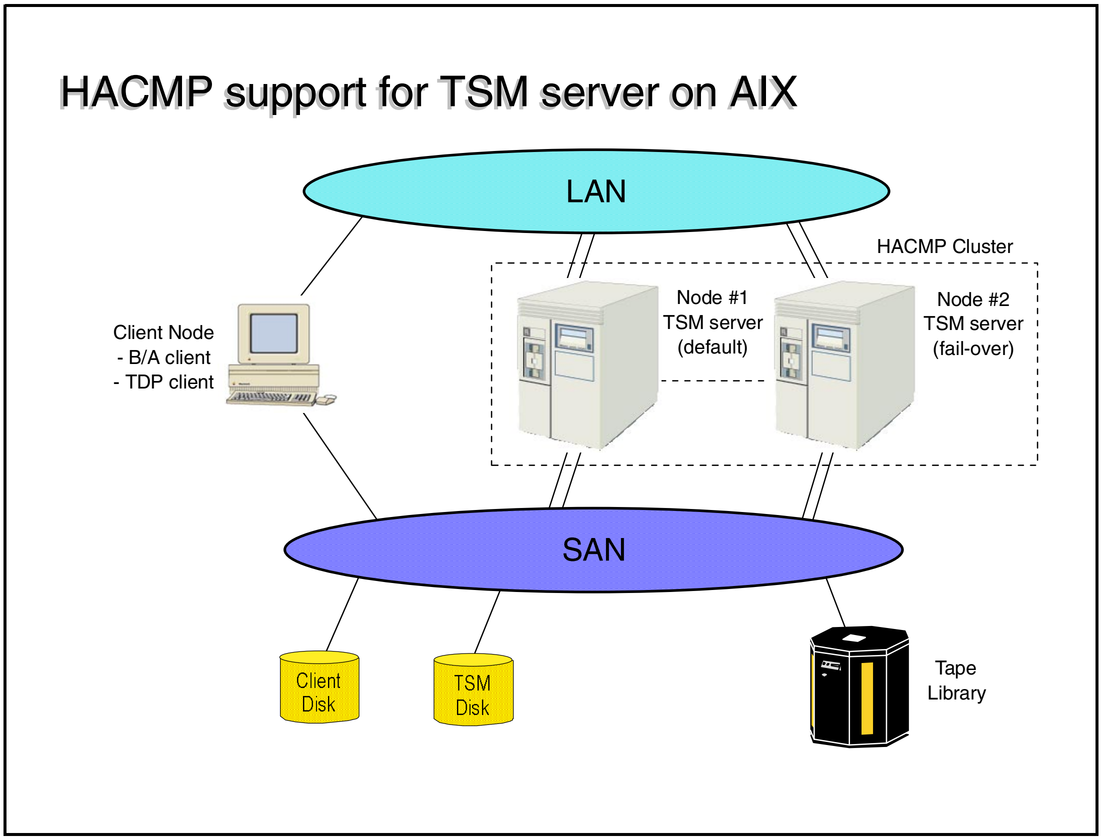

Figure 3-2 HACMP support for Tivoli Storage Manager server on AIX

The HACMP support for the Tivoli Storage Manager server on AIX has been officially supported since Tivoli Storage Manager Version 4.2, and was previously available through marketing tools. Several sample HACMP start/stop scripts are provided with the Tivoli Storage Manager server fileset. These scripts will need to be customized to suit your local environment.

If there is more than one offline storage device (tape library) for a Tivoli Storage Manager server to manage, it is possible to configure two Tivoli Storage Manager servers to run on two different systems in an HACMP cluster and to have either system run both servers if the other system fails. This can be accomplished by following the same set of directions using another file system that is accessible to both servers. It is also possible to configure the Tivoli Storage Manager Web Proxy on the second node in the cluster, configured to failover to the primary node on failure.

The installation of the Tivoli Storage Manager ADMIN client is required on both the primary and the backup cluster nodes. As of Tivoli Storage Manager Version 5.1, the B/A client itself (including the ADMIN, backup/archive, HSM and API pieces) is also supported for use in an HACMP environment. Refer to 5.1.2, “HACMP support for the B/A client on AIX” on page 181 for details on how to configure the client for use in a HACMP environment.

The support of HACMP with the Tivoli Storage Manager server has not been officially tested in either a SAN or a library sharing environment and therefore is not supported at this time in those environments. However, our environment included a SAN attached Magstar 3570 library, SAN library sharing, and a LAN-free client. Although we found this configuration to be workable, it is not a supported configuration.

Prerequisites tasks

To configure an AIX Tivoli Storage Manager server in a HACMP environment, you will need:

- Someone who has experience installing and configuring HACMP.
- Someone who has experience installing and configuring a Tivoli Storage Manager server.
- A familiarity with the following Tivoli Storage Manager documentation:
  - Tivoli Storage Manager for AIX Quick Start (GC35-0402)
  - Tivoli Storage Manager for AIX Administrator's Guide (GC35-0403)
  - Tivoli Storage Manager for AIX Administrator's Reference (GC35-0404)
  - Tivoli Storage Manager Installing the Clients (SH26-4119)
- A hardware configuration, suitable for HACMP, in which the Tivoli Storage Manager server's offline storage devices (Tape library/s) are physically connected to at least two nodes of the HACMP cluster on a shared bus.
- Sufficient shared (twin-tailed) disk space to hold the Tivoli Storage Manager database, recovery log, and disk storage pools you plan to use.

  See Chapter 18. "Managing the Database and Recovery Log" in *Tivoli Storage Manager for AIX Administrator's Guide*, GC35-0403 to:
  - Determine how much space will be required for the database and recovery log
  - Ensure the availability of the database and recovery log
  - Improve the performance of the database and recovery log
- A TCP/IP network. We will refer to the production node and the standby node. These are the two HACMP nodes on which the Tivoli Storage Manager server will run.

Step 0: Prerequisite HACMP configuration

Before configuring HACMP to support the Tivoli Storage Manager server you should have prepared a functioning HACMP environment.

The following steps assume you have already completed the following:

- Installed HACMP software on every node
- Configured HACMP Networks
- Configured HACMP Adapters

You should be able to perform failover and take-back operations before proceeding.

Step 1: Prepare the HACMP nodes

Before installing the Tivoli Storage Manager server code, you need to create the appropriate HACMP shared file systems and configure HACMP to manage the starting and stopping of the Tivoli Storage Manager server.

Define the shared filesystems

Define the shared disk space (volume group and filesystems) that will hold the Tivoli Storage Manager server configuration, database, recovery log, and storage pools you plan to use.

We chose to place the Tivoli Storage Manager server configuration, database, log, and storage pool volumes in the following shared filesystems (Example 3-1):

Example 3-1 Tivoli Storage Manager server filesystems located on shared disk

```
/tsm/server1/etc            (contains the TSM server configuration)
/tsm/server1/db             (contains the TSM server database volumes)
/tsm/server1/log            (contains the TSM server log volumes)
/tsm/server1/stgp           (contains the TSM server disk storage pool volumes)
```

You do not have to use the same names for your shared filesystems.

Note: The shared volume group/s must be set to not automatically vary on, and the shared file systems should not automatically mount at boot time. The HACMP software will manage their activation.

Remember to use the correct HACMP tools to define any shared filesystems and logical volumes. You can access the SMIT menu for this task via the command “smitty cl_lvm”.

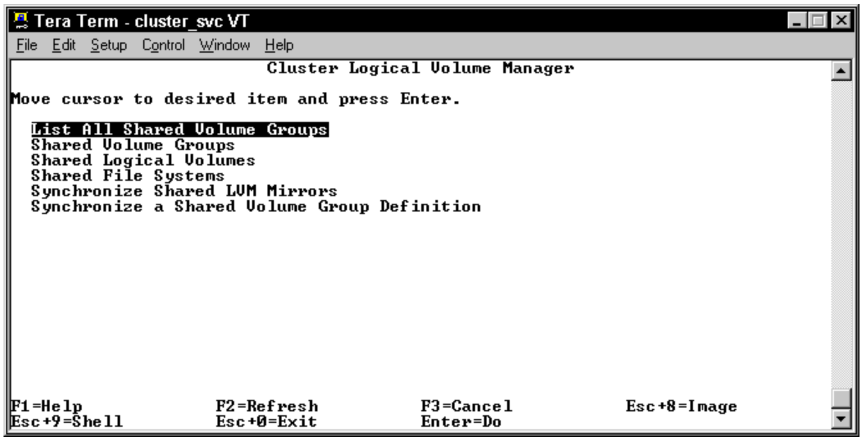

Figure 3-3 Defining shared filesystem via the smitty HACMP menu

The examples in the sample scripts are set up to use /tsm as the shared file system. If you use another name for the filesystem then you will need to modify the scripts to suit your environment.

Define the local filesystems

Define any local disk space (filesystems in the local rootvg volume group) that you plan to use with your Tivoli Storage Manager installation.

We chose to create two local filesystems on each node in the cluster. One to hold the customized HACMP start and stop scripts for the Tivoli Storage Manager server, and the other to hold any Tivoli Storage Manager database backups to a file device class (Example 3-2).

Example 3-2 Tivoli Storage Manager server filesystems located on local disk

```
/hacmp                        (contains the HACMP start/stop scripts for TSM)
/tsm/server1/dbbkup           (used to hold TSM database backup files)
```

You do not have to use that same names for your local filesystems.

Note: The local file systems should automatically mount at boot time.

Remember to create these filesystems as local filesystems on both nodes in the cluster.

Step 2: Install Tivoli Storage Manager on the production node

We recommended that you install the Tivoli Storage Manager server code on both nodes in the cluster. This ensures that both nodes have a valid LPP database, the required Tivoli Storage Manager device drivers, SMIT panels, and other files need to be installed in AIX system directories.

Install the Tivoli Storage Manager server on the production node

Install the Tivoli Storage Manager server executables as directed in *Tivoli Storage Manager for AIX Quick Start Guide*, GC35-0402. At a minimum the following filesets are required (Example 3-3):

Example 3-3 Filesets required for the Tivoli Storage Manager server

```
tivoli.tsm.devices.aix43.rte
tivoli.tsm.license.cert
tivoli.tsm.license.rte
tivoli.tsm.msg.[lang].devices
tivoli.tsm.msg.[lang].server
tivoli.tsm.msg.[lang].webhelp
tivoli.tsm.server.com
tivoli.tsm.server.rte
tivoli.tsm.server.webadmin
```

The executables should be installed on the internal disks of the production node, not on the shared Tivoli Storage Manager disk space. Tivoli Storage Manager server executables will install in the /usr/tivoli/tsm/server/bin directory by default. The required filesets will vary slightly when installing Tivoli Storage Manager server on AIX 5.1, or with non-English message support.

Copy the sample configuration files (dsmserv.opt.smp and dsmserv.dsk) to the shared filesystem configuration directory (in our example this is /tsm/server1/etc, Example 3-4). Rename the dsmserv.opt.smp file to dsmserv.opt.

Example 3-4 Move the Tivoli Storage Manager server sample configuration files to the shared disk

```
root@sicily:/tsm/server1/etc: mv /usr/tivoli/tsm/server/bin/dsmserv.opt .
root@sicily:/tsm/server1/etc: mv /usr/tivoli/tsm/server/bin/dsmserv.dsk .
```

Edit the “dsmserv.opt” file and configure Tivoli Storage Manager to use the TCP/IP communication method. The defaults in dsmserv.opt should work without modification. Use this opportunity to modify other parameters to suit your local environment.

> **Note:** TCP/IP is the only communications method the Tivoli Storage Manager server officially supports for takeover in an HACMP environment. (Even though it is not required by Tivoli Storage Manager, a line should be put in /etc/services on both hosts defining port 1500, as an indication that it cannot be used by other applications.)

Copy the sample HACMP scripts from the server directory into a directory in the local filesystem, Example 3-5:

Example 3-5 Copy the Tivoli Storage Manager server HACMP scripts to a local directory

```
cp /usr/tivoli/tsm/server/bin/startserver /hacmp/startserver
cp /usr/tivoli/tsm/server/bin/stopserver /hacmp/stopserver
```

Clean up the default server installation.

- Remove the entry from /etc/inittab that starts the Tivoli Storage Manager server. If this step is not performed then make sure this server and the server running from the /tsm/server1/etc directory are configured to use different TCP/IP ports.

  Use the following command (as the root user) to remove the /etc/inittab entry:

  ```
  rmitab "autosrvr:2:once:/usr/tivoli/tsm/server/bin/rc.adsmserv >/dev/console 2>&1
  #Start the Tivoli Storage Manager server"
  ```

- Remove the default created database, recovery log and storage pool files from the /usr/tivoli/tsm/server/bin directory. This will save a little space in the /usr directory if these files are not needed.

  Check that Tivoli Storage Manager is no longer using these files before deleting them:

  ```
  rm /usr/tivoli/tsm/server/bin/db.dsm
  rm /usr/tivoli/tsm/server/bin/log.dsm
  rm /usr/tivoli/tsm/server/bin/archive.dsm
  rm /usr/tivoli/tsm/server/bin/backup.dsm
  rm /usr/tivoli/tsm/server/bin/spcmgmt.dsm
  ```

Install the Tivoli Storage Manager admin client on the production node

Install the Tivoli Storage Manager client executables as directed in "Chapter 3. Installing UNIX Clients" of *Tivoli Storage Manager Installing the Clients,* SH26-4119.

The fileset name is tivoli.tsm.client.ba.aix43.32bit. The executables should be installed on the internal disks of the production node, not on the shared Tivoli Storage Manager disk space. Tivoli Storage Manager client executables will install in the /usr/tivoli/tsm/client/ba/bin directory by default.

> **Note:** Only the Tivoli Storage Manager ADMIN client needs to be installed. The fileset contains the backup-archive client files (command-line and GUI) and the administrative client (command-line). The command line ADMIN client is required because it is called by the stopserver script to halt the Tivoli Storage Manager server during a HACMP failover event.

Configure the dsm.sys file used by the Tivoli Storage Manager B/A client to find the Tivoli Storage Manager server we just installed. This is required for the administration client (dsmadmc) to access the Tivoli Storage Manager server. If you also require the B/A client to operate in a HACMP cluster, refer to 5.1.2, "HACMP support for the B/A client on AIX" on page 181 for details on how to configure the dsm.sys file for that environment.

Start the Tivoli Storage Manager server on the production node

We need to manually start the Tivoli Storage Manager server to create a basic configuration before continuing with the standby node installation. Issue the following commands to start the server, Example 3-6:

Example 3-6 Manually starting the Tivoli Storage Manager server on the production node

```
root@brazil:/tsm/server1/etc: cd /tsm/server1/etc
root@brazil:/tsm/server1/etc: export DSMSERV_CONFIG=/tsm/server1/etc/dsmserv.opt
root@brazil:/tsm/server1/etc: export DSMSERV_DIR=/usr/tivoli/tsm/server/bin
root@brazil:/tsm/server1/etc: /usr/tivoli/tsm/server/bin/dsmserv
...
```

Define new database and log volumes to reside on the shared filesystem. You should modify the location, size, and number of the database and log volumes to suit your local requirements (Example 3-7).

Example 3-7 Defining new database and log volumes on the shared filesystem

```
tsm: TSM> define dbvol /tsm/server1/db/db01.tsm f=1000 wait=yes
ANR0984I Process 7 for DEFINE DBVOLUME started in the FOREGROUND at 10:50:25.
ANR2240I Database volume /tsm/server1/db/db01.tsm defined.
ANR0986I Process 7 for DEFINE DBVOLUME running in the FOREGROUND processed 1
items for a total of 1,049,624,576 bytes with a completion state of SUCCESS at 10:50:35.

tsm: TSM> define logvol /tsm/server1/log/log01.tsm f=500 wait=yes
ANR0984I Process 12 for DEFINE LOGVOLUME started in the FOREGROUND at 12:38:16.
ANR2260I Recovery log volume /tsm/server1/log/log01.tsm defined.
ANR0986I Process 12 for DEFINE LOGVOLUME running in the FOREGROUND processed 1
items for a total of 525,336,576 bytes with a completion state of SUCCESS at 12:39:07.
```

After defining the new database and log volumes, delete the original volume's (located on the local filesystem) references from within the Tivoli Storage Manager server (Example 3-8). This will cause the contents of the original database and log volumes to migrate into the new volumes.

Example 3-8 Delete the original db and log references from the Tivoli Storage Manager server

```
tsm: TSM> delete dbvol /usr/tivoli/tsm/server/bin/db.dsm
ANR2243I Database volume /usr/tivoli/tsm/server/bin/db.dsm deleted.

tsm: TSM> delete logvol /usr/tivoli/tsm/server/bin/log.dsm
ANR2263I Recovery log volume /tsm/tivoli/tsm/server/log.dsm deleted.
```

This will still leave the files that contained the original volumes on the local filesystem, but they will no longer be used by Tivoli Storage Manager. You may delete the unused files if you so desire.

Expand the database and log capacity up to the required values, as shown in Example 3-9. This should probably be up to the maximum capacity of the volumes you just defined.

Example 3-9 Expanding the database and log volumes

```
tsm: TSM> q db

Available Assigned   Maximum   Maximum    Page     Total      Used   Pct  Max.
    Space Capacity Extension Reduction    Size    Usable     Pages  Util   Pct
     (MB)     (MB)      (MB)      (MB) (bytes)     Pages                  Util
--------- -------- --------- --------- ------- --------- --------- ----- -----
    1,000       16       984        12   4,096     4,096       252   6.2   6.2

tsm: TSM> extend db 984
ANR2248I Database assigned capacity has been extended.

tsm: TSM> q log

Available Assigned   Maximum   Maximum    Page     Total      Used   Pct  Max.
    Space Capacity Extension Reduction    Size    Usable     Pages  Util   Pct
     (MB)     (MB)      (MB)      (MB) (bytes)     Pages                  Util
--------- -------- --------- --------- ------- --------- --------- ----- -----
      500        8       492         4   4,096     1,536        76   4.9   5.9
```

```
tsm: TSM> extend log 492
ANR2269I Recovery log extension process initiated (process ID 6).
ANS8003I Process number 6 started.
```

Use the commands shown in Example 3-10 to verify that all of the database and log volumes are on the shared disk only.

Example 3-10 Query the database and log volumes

```
tsm: TSM> q dbvol

Volume Name       Copy    Volume Name       Copy    Volume Name       Copy
(Copy 1)          Status  (Copy 2)          Status  (Copy 3)          Status
----------------  ------  ----------------  ------  ----------------  ------
/tsm/server1/db-  Sync'd                    Undef-                    Undef-
 /db01.tsm                                   ined                      ined

tsm: TSM> q logvol

Volume Name       Copy    Volume Name       Copy    Volume Name       Copy
(Copy 1)          Status  (Copy 2)          Status  (Copy 3)          Status
----------------  ------  ----------------  ------  ----------------  ------
/tsm/server1/lo-  Sync'd                    Undef-                    Undef-
 g/log01.tsm                                 ined                      ined
```

If your shared filesystem does not reside on reliable disk volumes (RAID 0 or RAID 5) you should also create secondary copies of the database and log volumes.

Note: If you delete the original database volume, and create a new database from scratch, without first creating and migrating to a new volume, you will need to recreate the additional database structures normally created during installation.

Some symptoms of not preserving the additional database structures are the lack of a “admin” user account and a faulty Web administrative interface.

To manually add the additional database structures run the following commands:

```
cd /usr/tivoli/tsm/server/bin
export DSMSERV_CONFIG=/tsm/server1/etc/dsmserv.opt
export DSMSERV_DIR=/usr/tivoli/tsm/server/bin
./dsmserv runfile dsmserv.idl
```

Define new disk storage pool volumes on the shared filesystem (Example 3-11). To save repetition, our example only deals with one of the default storage pools. You should modify these steps to suit your local storage pool design.

Example 3-11 Define new disk storage pool volumes on the shared filesystem

```
tsm: TSM> q vol

Volume Name               Storage      Device      Estimated    Pct   Volume
                          Pool Name    Class Name   Capacity   Util   Status
                                                        (MB)
------------------------  -----------  ----------  ---------  -----  --------
/usr/tivoli/tsm/server/-  ARCHIVEPOOL  DISK              8.0    0.0   On-Line
 bin/archive.dsm
```

```
/usr/tivoli/tsm/server/-  BACKUPPOOL   DISK            8.0    0.0  On-Line
 bin/backup.dsm
/usr/tivoli/tsm/server/-  SPACEMGPOOL  DISK            8.0    0.0  On-Line
 bin/spcmgmt.dsm

tsm: TSM> define volume backuppool /tsm/server1/stgp/backup.dsm f=1000 wait=yes

ANR0984I Process 2 for DEFINE VOLUME started in the FOREGROUND at 14:57:29.
ANR2206I Volume /tsm/server1/stgp/backup.dsm defined in storage pool BACKUPPOOL
(device class DISK).
ANR0986I Process 2 for DEFINE VOLUME running in the FOREGROUND processed 1
items for a total of 1,048,576,000 bytes with a completion state of SUCCESS at 14:58:30.
```

Migrate the contents (if any) of the original disk storage pool volumes into the new volumes on the shared disk. Remove the original disk storage pool volumes when complete. Use the commands shown in Example 3-12 to migrate the data and remove the original volumes.

Example 3-12 Migrate disk storage pool contents on to the shared filesystem

```
tsm: TSM> move data /usr/tivoli/tsm/server/bin/backup.dsm wait=yes
ANR2232W This command will move all of the data stored on volume
/usr/tivoli/tsm/server/bin/backup.dsm to other volumes within the same storage
pool; the data will be inaccessible to users until the operation completes.

Do you wish to proceed? (Yes (Y)/No (N)) y
ANR2209W Volume /usr/tivoli/tsm/server/bin/backup.dsm contains no data.
ANS8001I Return code 11.

tsm: TSM> delete vol /usr/tivoli/tsm/server/bin/backup.dsm wait=yes
ANR2220W This command will delete volume /usr/tivoli/tsm/server/bin/backup.dsm
from its storage pool after verifying that the volume contains no data.

Do you wish to proceed? (Yes (Y)/No (N)) y
ANR2208I Volume /usr/tivoli/tsm/server/bin/backup.dsm deleted from storage pool
BACKUPPOOL.
```

Repeat this step for each storage pool and volume, until all your disk storage pool volumes reside on the shared filesystem.

Verify that all the disk storage pool volumes now all reside on the shared filesystem only, Example 3-13.

Example 3-13 Check the disk storage pool volume all reside on the shared filesystem

```
tsm: TSM> q vol

Volume Name               Storage      Device      Estimated    Pct   Volume
                          Pool Name    Class Name   Capacity   Util   Status
                                                        (MB)
------------------------  -----------  ----------  ---------  -----  --------
/tsm/server1/stgp/archi-  ARCHIVEPOOL  DISK          1,000.0    0.0  On-Line
 ve.dsm
/tsm/server1/stgp/backu-  BACKUPPOOL   DISK          1,000.0    0.0  On-Line
 p.dsm
/tsm/server1/stgp/spcmg-  SPACEMGPOOL  DISK          1,000.0    0.0  On-Line
 mt.dsm
```

Register an administrator that will be used by HACMP to control the Tivoli Storage Manager server and grant the administrator storage and operator privilege (Example 3-14). The HACMP stop script will use this administrator to halt the Tivoli Storage Manager server during a planned failover event.

Example 3-14 Registering an administrator for HACMP operation

```
tsm: TSM_SERVER1>register admin HACMP password
ANR2068I Administrator HACMP registered.

tsm: TSM_SERVER1>grant authority HACMP classes=operator,storage
ANR2079I Unrestricted storage privilege granted to administrator HACMP.
ANR2082I Operator privilege granted to administrator HACMP.
```

Edit the /hacmp/stopserver shell script and update the userid and password in it to the values set above. And update the servername in it to the dsm.sys stanza label set above.

Register any client licenses as required. Follow instructions in "Registering Licenses" in *Tivoli Storage Manager for AIX Quick Start*, GC35-0402.

Step 3: Install Tivoli Storage Manager on the standby node

It is recommended to install the Tivoli Storage Manager server code on both nodes in the cluster. This ensures that both nodes have a valid LPP database, the required Tivoli Storage Manager device drivers, SMIT panels, and other files need to be installed in AIX system directories.

Install the Tivoli Storage Manager server on the standby node

Make sure the Tivoli Storage Manager server is not running on the production node and bring the shared volume group and any shared Tivoli Storage Manager filesystems up on the standby node.

On the standby node install the Tivoli Storage Manager server as was performed for the production node.

Install the Tivoli Storage Manager admin client on the standby node

Install the admin client on the stand by node as described for the production node. Copy the client configuration files (dsm.ys, dsm.opt, dsm.inx) from the production node to the standby node. Remember that if the client configuration is changed on one node, it must be copied to the other node.

Start the Tivoli Storage Manager server on the standby node

Start the server on the stand by node as described in Example 3-6. Query the database, recovery log and storage pool volumes to verify they are the same as when the server was started on the production node.

Perform the clean-up of the default server installation on the standby node as described for the production node. It is very important to remove the entry from /etc/inittab that would start the Tivoli Storage Manager server at boot time.

Step 4: Define the offline storage devices (tape library)

On both the production and standby nodes, define to AIX the offline storage devices that will be used by Tivoli Storage Manager. Follow the appropriate instructions in "Using Tape and Optical Devices" in *Tivoli Storage Manager for AIX Quick Start*, GC35-0402.

Make sure the offline storage devices are configured with the same names on the standby and production nodes. You may have to define "dummy" devices on one of the nodes to accomplish this. To define a "dummy" device, issue the command 'smit devices' and go through the smit panels to define the device. Choose an unused SCSI address for the device. Instead of pressing Enter on the last panel to actually define the device, press F6 instead to obtain the command smit is about to execute. Now, exit from smit and enter the same command on the command line, adding the -d flag to the command.

If you actually attempt to define the device using smit, the attempt will, of course, fail because there is no device at the unused SCSI address you have chosen.

Step 5: Customize the HACMP start and stop scripts

You will need to customize the sample HACMP start and stop scripts to suit your local environment. In particular you will need to set the correct directory names and offline device names.

The startserver script

Edit the /hacmp/startserver script and change the VerifyDevice commands in the shell script to specify the offline storage devices that will be used by Tivoli Storage Manager. The names of these devices may be different at your site.

Change the “bindir” variable to point to your Tivoli Storage Manager server installation directory.

Example 3-15 Extracts from the HACMP start script for the Tivoli Storage Manager server

```
bindir=/usr/tivoli/tsm/server/bin
...

################################################################################
#                                                                              #
# Insert a VerifyDevice command below for every offline storage device used by #
# this TSM server and shared by more than one system in the HACMP/6000         #
# cluster.  The following examples are commented out.                          #
#                                                                              #
################################################################################
VerifyDevice lb0
VerifyDevice mt0
VerifyDevice mt1

...
################################################################################
#                                                                              #
# Update the cd command below to change to the directory that contains the     #
# dsmserv.dsk file and change the export commands to point to the dsmserv.opt  #
# file and /usr/tivoli/tsm/server/bin directory for the TSM server being       #
# started.  The export commands are currently set to the defaults.             #
#                                                                              #
################################################################################
cd /tsm/server1/etc
export DSMSERV_CONFIG=/tsm/server1/etc/dsmserv.opt
export DSMSERV_DIR=/usr/tivoli/tsm/server/bin
...
```

The startserver shell script takes care of breaking any SCSI RESERVE on the offline storage devices during a startup or failover event.

> **Note:** The VerifyDevice command will not be needed to failover an IBM 3494 Library Manager. The Library Manager is not connected via SCSI, so no special logic is required to fail it over.

Change the cd command near the end of the shell script to cd /tsm/server1/etc.

The stopserver script

Change the username and password in the sample stopserver script to be the same as for the “HACMP” administrative user we defined earlier. Use a better password than shown here.

Example 3-16 Extract from the HACMP stop script for the Tivoli Storage Manager server

```
...
echo "Stopping the TSM server now..."
/usr/tivoli/tsm/client/ba/bin/dsmadmc -servername=tsm_cluster1 -id=hacmp -password=password
-noconfirm << EOF
...
```

After customizing the scripts, test them manually before attempting to use them with HACMP.

Step 6: Configure HACMP to manage the new resources:

Configure HACMP so that the production node owns the shared volume groups and the standby node takes over the shared volume groups if the production node fails.

The following example assumes you already have a functioning HACMP environment, with correctly configured Networks, Adapters, and Nodes, etc. You should be able to successfully perform a HACMP failover and recovery before continuing.

First (if you have not already done so), define a Resource Group to manage the Tivoli Storage Manager server. You can access the SMIT menu for this task via the command “smitty cm_add_res”.

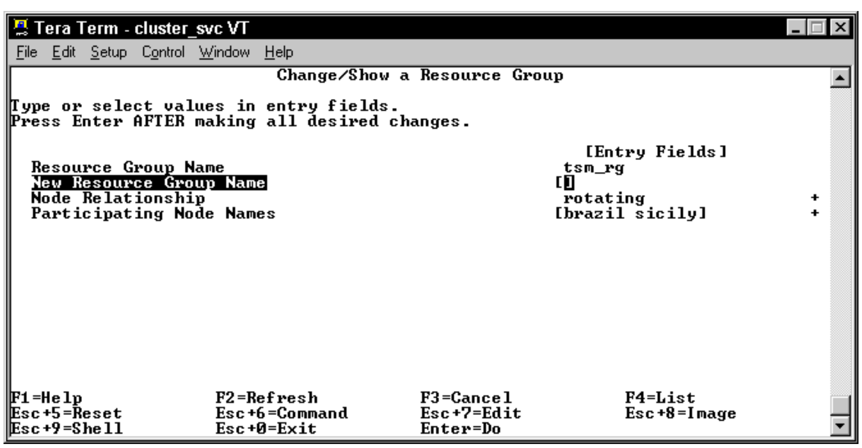

Figure 3-4 Defining a HACMP resource group for the Tivoli Storage Manager server

You will need to configure the following parameters:

Resource Group Name Enter a logical name for your resource group. Maximum length is 32 characters.

Node Relationship Toggle between Cascading/Concurrent/Rotating.

CASCADING resources are resources which may be assigned to be taken over by multiple nodes in a prioritized manner. When a node fails, the active node with the highest priority acquires the resource. When the failed node rejoins, the node with the highest priority acquires the resource.

ROTATING resources are resources which may be acquired by any node in its resource chain. When a node fails, the resource will be acquired by the highest priority standby node. When the failed node rejoins, the resource remains with its new owner.

CONCURRENT resources are not supported.

Participating Node Names
Enter the nodes in the resource chain for the resource group. These are the nodes which may acquire the resource. The order you list the nodes can determine the priority of the Resource Group allocation, depending on the Node relationship.

CASCADING RESOURCES: Priority decreases from left to right.

CONCURRENT RESOURCES: All nodes are at the same priority. Concurrent mode is not supported.

ROTATING RESOURCES: Priority decreases from left to right, but is only used to determine which standby acquires the resource in a multiple standby configuration.

Secondly, define a HACMP Application Server to manage the Tivoli Storage Manager server. This is where you define the scripts to start and stop the Tivoli Storage Manager server. You can access the SMIT menu for this task via the command “smitty cm_cfg_app”.

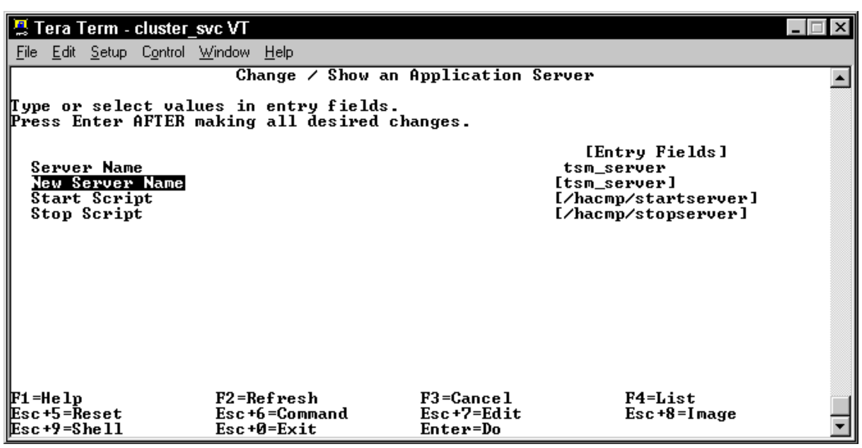

Figure 3-5 Defining a HACMP application server for the Tivoli Storage Manager server

You will need to configure the following parameters:

Server Name
Enter a logical name for your Application Server. Maximum length is 32 characters.

Start script
Script to start the application server when node joins the cluster. This script needs to be located on a local

filesystem that is available when HACMP is not yet running.

Stop script Script to stop the application server when node leaves the cluster gracefully. This script could be located on a shared filesystem, but should probably be stored on a local filesystem with the corresponding start script.

Finally, associate the application server and filesystem resources with the Resource Group you defined earlier. You can access the SMIT menu for this task via the command “smitty cm_cfg_res.select”.

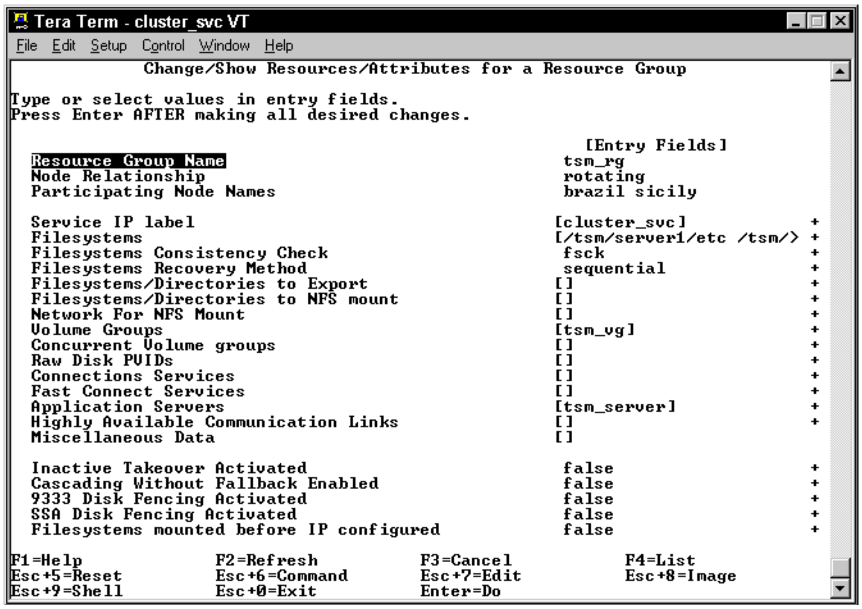

Figure 3-6 HACMP resources and application servers for Tivoli Storage Manager server

You will need to configure at least the following parameters:

Service IP label The IP label (hostname) of the adapter, associated with the numeric IP address in the /etc/hosts file.

Filesystems Enter the mount points of the filesystems which are mounted when the resource is initially acquired.

Volume Groups Enter the names of the volume groups containing raw logical volumes or raw volume groups that are varied on when the resource is initially acquired.

Application Servers Enter application servers that will be started. These are the servers defined in Figure 3-5 on page 100.

Depending on your specific HACMP configuration, there may be other parameters to configure.

Finally, synchronize your changes to the other nodes in the cluster. You can access the SMIT menu for this task via the command “smitty clsyncnode.dialog”.

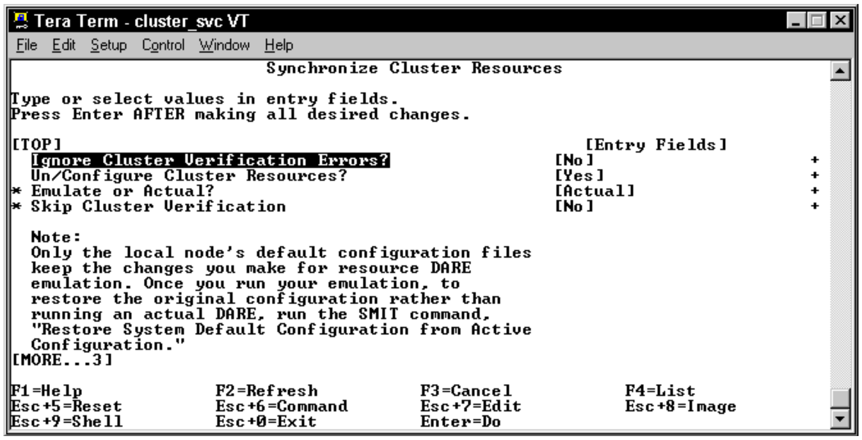

Figure 3-7 Synchronize HACMP cluster resources

Your HACMP environment should now be correctly configured and ready for testing.

Step 7: Complete the Tivoli Storage Manager server configuration

You should now be ready to start the Tivoli Storage Manager server under the control of HACMP and continue configuring the Tivoli Storage Manager server. Once the Tivoli Storage Manager server is started it can be configured exactly as if it was a stand-alone server.

Before starting the HACMP cluster for the first time, ensure you stop any services or resources normally managed by HACMP that you may have manually started during setup and installation. Repeat this check on each node in your cluster.

Start the HACMP cluster

First, start the HACMP software on all nodes in the cluster. You can access the SMIT menu for this task via the command “smitty cl_clstart.dialog

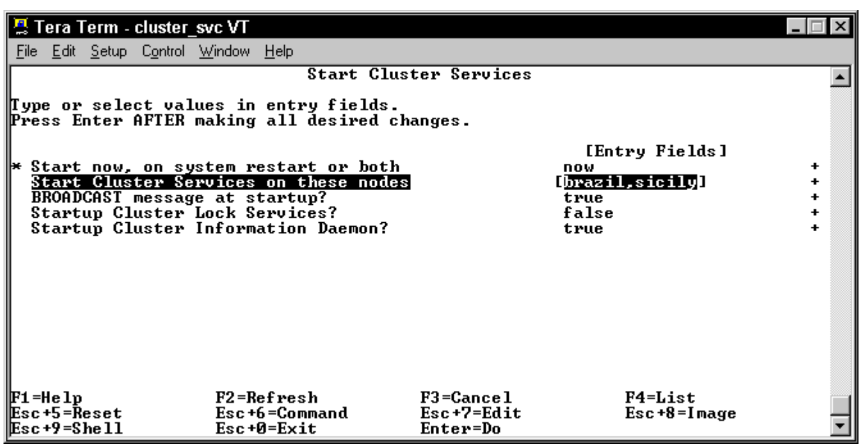

Figure 3-8 Starting the HACMP cluster on all nodes

You will need to configure the following parameters:

Start now, on system restart, or both	Choosing 'now' will start the daemons immediately. Select this option for our testing purposes. Choosing 'system restart' will start the daemons after a system reboot by adding an entry to the /etc/inittab file.

	Choosing 'both' will start the daemons immediately AND after a system reboot.
Start Cluster Services on these nodes	Identify the nodes on which to start cluster services.
Startup Cluster Information Daemon	Starts the Cluster Information daemon. This daemon needs to be started to support the B/A client.

This will cause the service IP address/s to become active, and you will lose any connections to the nodes boot IP address/s.

After the HACMP software has started, the applications and filesystems associated with your Tivoli Storage Manager Resource Group will be available on the production cluster node.

Continue with configuring the Tivoli Storage Manager server as per a normal installation. Any changes made will be saved to the Tivoli Storage Manager server database on the shared filesystem.

Step 8: Test Tivoli Storage Manager server across a failover event

The aim of configuring the Tivoli Storage Manager server to run in a HACMP cluster is to increase the availability of the Tivoli Storage Manager server. Clients should now be able to backup and restore data regardless of which cluster node is currently hosting the Tivoli Storage Manager server application.

Current client sessions should also reconnect after a failover event. The client interprets the unavailability of the Tivoli Storage Manager server during failover as a communications problem, and will retry the connection after a small delay. Hopefully this delay is long enough to allow the cluster to stabilize and the Tivoli Storage Manager server to restart on the standby node.

We now need to test the failover and client access during and after the event.

Initiate a client backup session

Register a client account on the new Tivoli Storage Manager server for a client outside of the cluster hosting the Tivoli Storage Manager server. Configure the client to connect to the new Tivoli Storage Manager server.

Initiate a client backup session. Chose a filesystem that will take some time to backup so we can see the client reconnect automatically after the cluster failover completes.

Example 3-17 Start a client backup to the HA Tivoli Storage Manager server

```
[root@tungsten /]# dsmc i -se=tsm_cluster1
Tivoli Storage Manager
Command Line Backup/Archive Client Interface - Version 5, Release 1, Level 0.0 g
3
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Node Name: TUNGSTEN_FS
Session established with server TSM_SERVER1: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 04/02/2002 14:40:29  Last access: 04/02/2002 14:21:41


Incremental backup of volume '/'
...
```

Force a failover event

Next, stop the HACMP software on the production node. This should cause the Resource Group to failover to the standby node. You can access the SMIT menu for this task via the command “smitty clstop.dialog”.

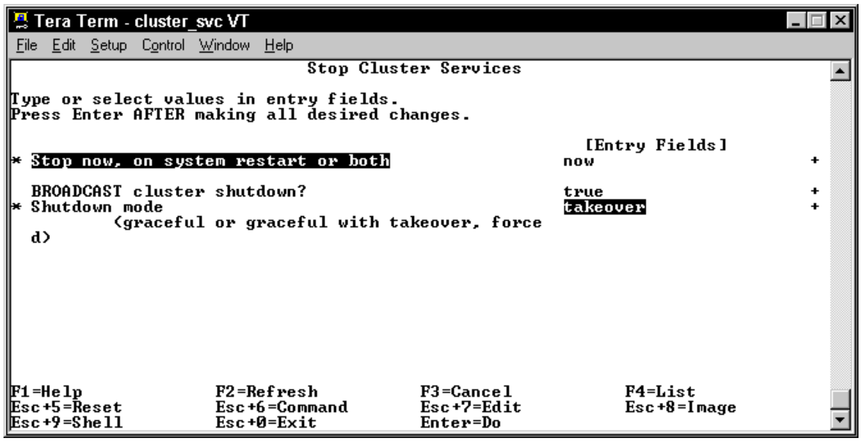

Figure 3-9 Controlling the HACMP cluster to test B/A client failover

This will cause the service IP address/s to move to another node, and you will lose any connections to the current node.

After the HACMP failover has completed, the applications and filesystems associated with your Tivoli Storage Manager Resource Group will be available on the next cluster node. Use the examples shown in Example 5-11 on page 189 to test for the correct startup of the Tivoli Storage Manager client acceptor (scheduler) on the new node.

Test server availability and continued client access

The client will lose connection to the Tivoli Storage Manager server as the failover event occurs, but will retry several times, and should reconnect automatically when the cluster stabilizes.

An ongoing client backup session will display a message similar to that shown in Example 3-18.

Example 3-18 Client backup continues after a HACMP failover

```
...
Normal File-->          39,692 /usr/bin/kspace.kss [Sent]
Normal File-->         141,352 /usr/bin/kspaceduel [Sent]
Normal File-->          33,452 /usr/bin/ksplash [Sent]

ANS1809W Session is lost; initializing session reopen procedure.
A Reconnection attempt will be made in 00:00:14

Normal File-->           3,548 /usr/bin/kspread [Sent]
Normal File-->          18,800 /usr/bin/kstart [Sent]
Normal File-->           8,180 /usr/bin/kstartperf [Sent]
...
```

Connected to TSM server on Production node

Connected to TSM server on Standby node

The client is now able to backup and restore to the Tivoli Storage Manager server regardless of which node in the cluster is currently hosting the Tivoli Storage Manager server application.

3.1.3 Direct I/O support

Normal I/O processing for Journaled File System (JFS) files goes from the application buffer to the Virtual Memory Manager (VMM) and from there to the JFS. The contents of the buffer will get cached in RAM through the VMM's use of real memory as a file buffer cache. If the file cache hit rate is high, then this type of cached I/O will improve performance of JFS I/O. However since Tivoli Storage Manager always waits until the data has passed through VMM and has been written to disk by Logical Volume Manager (LVM), the principal write behavior for the server has not changed.

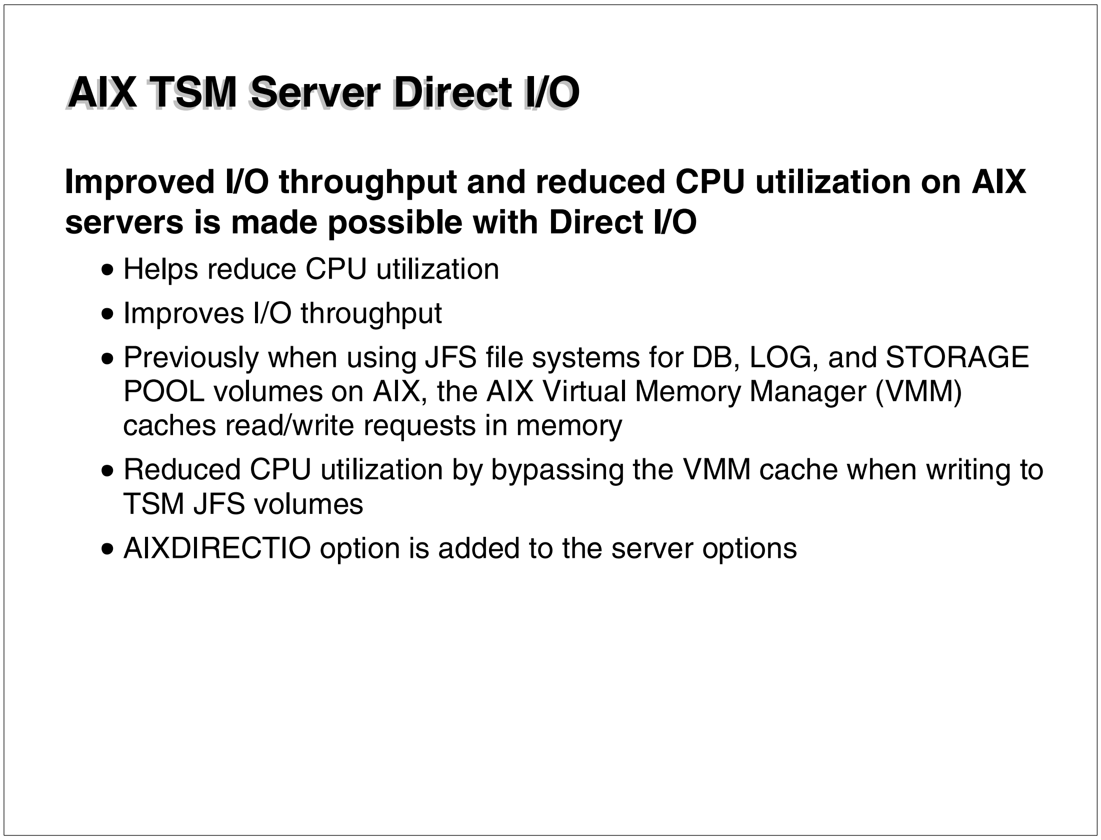

Figure 3-10 AIX direct I/O

Direct I/O (DIO) is an alternative I/O method for JFS files and is only supported for program working storage (local persistent files). The main benefit of direct I/O is to reduce CPU utilization for file reads and writes by eliminating the copy from the VMM file cache to the user buffer.

Because direct I/O writes do not get copied into memory, when a *sync* operation is performed, it will not have to flush these pages to disk, therefore reducing the amount of work the `syncd` daemon has to perform.

Direct I/O bypasses the VMM read-ahead algorithm because the I/O does not go through the VMM. The VMM read-ahead algorithm is very useful for sequential access to files because the VMM can initiate disk requests and have the pages already be resident in memory before the application has requested the pages. This also ensures that the data is physically written to disk in the event of a unexpected system halt.

For more information on asynchronous I/O please refer to the *AIX 5L Version 5.1 Performance Management Guide*, at the following URL:

http://publibn.boulder.ibm.com/doc_link/en_US/a_doc_lib/aixbman/prftungd/prftungd.htm

Applications such as IBM Tivoli Storage Manager that both do read and writes to disk storage pool volumes, the benefit of direct I/O for writing can be offset by the penalty of not having read-ahead done by VMM for read operations.

For IBM Tivoli Storage Manager to better handle multiple I/O requests while using direct I/O, you should also enable the use of asynchronous I/O as is shown below in 3.1.4, "Asynchronous I/O support" on page 106. However the benefits of direct I/O are highly dependent on the underlying file system and disk subsystem layout.

To determine if the current installation and workload benefit from direct I/O, run the IBM Tivoli Storage Manager server during a normal 24-hour workload period with direct I/O enabled and for another normal 24-hour workload period with direct I/O disabled. Monitor both AIX system and I/O usage and client throughput numbers. For more information on AIX monitoring and monitoring tools please refer to the *AIX 5L Performance Tools Handbook*, SG24-6039.

Usage of direct I/O is enable by default in IBM Tivoli Storage Manager Version 5.1 so it does *not* have to be enabled in the **dsmserv.opt** server options file. The server will display the following message when it starts if direct I/O is enabled:

```
ANR7811I Direct I/O will be used for all eligible disk files.
```

To explicitly enable direct I/O i n the dsmsrv.opt file is shown anyway in the following example.

Example 3-19 AIXDIRECTIO option in dsmserv.opt

```
AIXDIRECTIO    YES
```

Note: Please note that currently direct I/O is *only* *not* *database or log volumes*, and only on *non-compressed* and *non-large-file-enabled* file systems. It is also only supported for up to 2 GB volumes.

3.1.4 Asynchronous I/O support

Synchronous I/O occurs while the user application wait for it to be completed, this is the normal mode for Unix I/O. Applications processing cannot continue until the I/O operation is complete. Asynchronous I/O operations, on the other hand, run in the background and do not block user applications. This improves performance, because I/O operations and applications processing can run simultaneously.

Using asynchronous I/O will usually improve I/O throughput, especially when storing data in raw logical volumes (as opposed to Journaled file systems) or using Direct I/O. The actual performance, however, depends on how many server processes are running that will handle the I/O requests and the physical I/O subsystem.

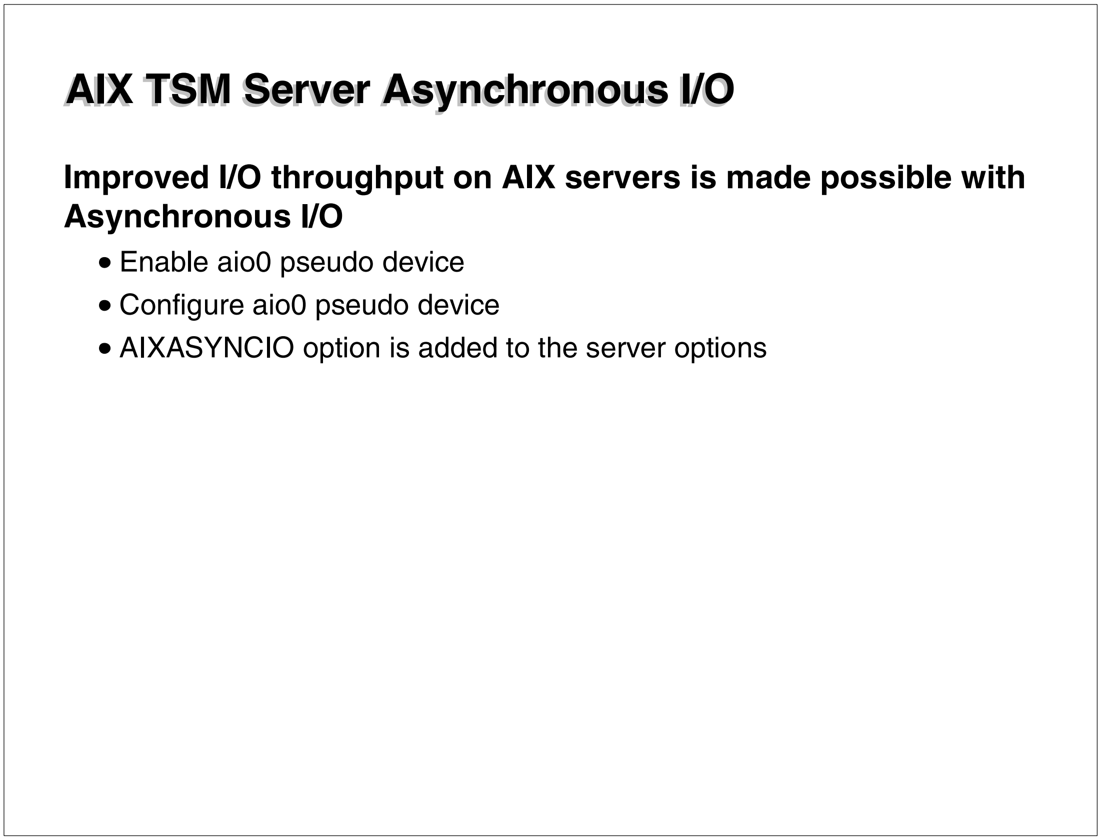

Figure 3-11 AIX asynchronous I/O

For more information on asynchronous I/O please refer to the *AIX 5L Version 5.1 Performance Management Guide*, at the following URL:

http://publibn.boulder.ibm.com/doc_link/en_US/a_doc_lib/aixbman/prftungd/prftungd.htm

Also see *AIX 5L Version 5.1 Kernel Extensions and Device Support Programming Concepts*, at the following URL:

http://publibn.boulder.ibm.com/doc_link/en_US/a_doc_lib/aixprggd/kernextc/ls_devconfig_subr.htm

Setting up Tivoli Storage Manager usage of AIXASYNCIO

The `AIXASYNCIO` server option allows the **dsmserv** process to use asynchronous I/O (AIO) support provided by AIX. Since the default is for AIO *not* to be used by **dsmserv** it must be enabled in the **dsmserv.opt** server options file, which is shown in the following example.

Example 3-20 AIXASYNCIO option in dsmserv.opt

```
AIXASYNCIO    YES
```

In addition to declaring `AIXASYNCIO` to be used for IBM Tivoli Storage Manager, AIO must also be enabled in AIX. This can be done by using SMIT or by using the **mkdev** command if the bos.rte.aio fileset has been installed. Check that the required fileset is installed with the **lslpp** command as is shown in the following example.

Example 3-21 Checking bos.rte.aio fileset on AIX 4.3.3 and AIX 5L

```
root@aussie:/: rsh aix433 "oslevel -r;lslpp -cl bos.rte.aio"
4330-09
#Fileset:Level:PTF Id:State:Type:Description
/usr/lib/objrepos:bos.rte.aio:4.3.3.75::COMMITTED:F:Asynchronous I/O Extension
/etc/objrepos:bos.rte.aio:4.3.3.75::COMMITTED:F:Asynchronous I/O Extension

roden@swede:/: oslevel -r;lslpp -cl bos.rte.aio
5100-01
#Fileset:Level:PTF Id:State:Type:Description
/usr/lib/objrepos:bos.rte.aio:5.1.0.10::COMMITTED:I:Asynchronous I/O Extension
/etc/objrepos:bos.rte.aio:5.1.0.10::COMMITTED:I:Asynchronous I/O Extension
```

Example 3-22 shows the default settings for the `aio0` pseudo device driver.

Example 3-22 Checking the settings of aio0 with the lsattr command

```
root@swede:/: lsattr -El aio0
minservers 1       MINIMUM number of servers              True
maxservers 10      MAXIMUM number of servers              True
maxreqs    4096    Maximum number of REQUESTS             True
kprocprio  39      Server PRIORITY                        True
autoconfig defined STATE to be configured at system restart True
fastpath   enable  State of fast path                     True
```

Example 3-22 shows that `autoconfig` is *defined*, this has to be changed in order for `aio0` to be enabled at IPL (Initial Program Load)[1]. To determine the current state of the `aio0` device driver use the **lsdev** command as is shown in Example 3-23.

Example 3-23 Determining the current state of the aio0 device driver

```
root@swede:/: lsdev -Cl aio0
aio0 Defined  Asynchronous I/O
```

To enable the asynchronous I/O on the running system use the **mkdev** command as shown below. However before enabling `aio0`, please change the number of the `maxservers` attribute as shown in Example 3-24.

Example 3-24 Enabling aio0 with the mkdev command

```
root@swede:/: mkdev -l aio0
aio0 Available
```

To change the `aio0` device driver to be enabled after a system restart use the **chdev** command as in Example 3-25.

Example 3-25 Changing aio0 with the chdev command

```
root@swede:/: chdev -l aio0 -a autoconfig=available
aio0 Available
```

Example 3-25 changed the `autoconfig` option from `defined` to `available`. To find out the allowable settings for `autoconfig` use the **-R** flag to the **lsattr** command as shown in Example 3-26.

[1] More commonly know as *boot* from *bootstrap*.

Example 3-26 Checking allowable settings for the autoconfig option with the lsattr command

```
root@swede:/: lsattr -Rl aio0 -a autoconfig
available
defined
```

The output above shows the only two allowable settings, `available` or `defined`. After changing the `aio0` device we can use the **`lsattr`** command again to verify that the `aio0` device will be loaded after the next system IPL which is shown in Example 3-27.

Example 3-27 Verifying settings with the lsattr command

```
root@swede:/: lsattr -El aio0
minservers 1         MINIMUM number of servers                True
maxservers 10        MAXIMUM number of servers                True
maxreqs    4096      Maximum number of REQUESTS               True
kprocprio  39        Server PRIORITY                          True
autoconfig available STATE to be configured at system restart True
fastpath   enable    State of fast path                       True
```

Now we know that asynchronous I/O will be usable after the next IPL, so now we do not have to manually enable it again.

Monitoring and tuning of AIO for IBM Tivoli Storage Manager

Asynchronous I/O is handled by a single **`kproc`**, and typically the **`kproc`** cannot process any more requests from the queue until that I/O has completed. The default minimum number of servers configured when asynchronous I/O is enabled is *1*. This is the `minservers` attribute. There is also a maximum number of asynchronous I/O servers that can get created which is controlled by the `maxservers` attribute; this has a default value of *10*. The number of servers limits the number of asynchronous I/O operations that can be in progress in the system simultaneously.

If the number of asynchronous I/O requests is high, then the recommendation is to increase `maxservers` to approximately the number of simultaneous I/Os there might be. In most cases, it is better to leave the `minservers` parameter at the default value because the AIO kernel extension will generate additional servers if needed.

By looking at the CPU utilization of the AIO servers, if the utilization is evenly divided among all of them, that means that they are all being used, in this case you may want to increasing them. To see the AIO server by name on AIX 4.3.3 run the **`pstat -a|grep aio`** command on and on AIX 5L run the **`ps -k|grep aio`** command, as shown in Example 3-28.

Example 3-28 Checking aioservers on AIX 4.3.3 and AIX 5L

```
root@aussie:/: rsh sicily "oslevel -r;pstat -a|grep -c aio"
4330-09
9

roden@swede:/: oslevel -r;ps -k|grep -c aio
5100-01
175

roden@swede:/: oslevel -r;ps laxww|sed '1p;/aio/p;/.*/d'
5100-01
     F S UID     PID    PPID   C PRI NI ADDR  SZ  RSS   WCHAN    TTY  TIME CMD
 40303 A   0  163910       1   0  39 -- 1c81  56 22092              -  0:00 aioserver
 42303 A   0  229434       1   0  39 -- 2840  48 22084              -  0:00 aioserver
 42303 A   0  417932       1   0  39 -- 1de3d  48 22080              -  0:00 aioserver
```

```
 42303 A   0  426226       1   0  39 -- 80c9  48 22084              -  0:00 aioserver
 42303 A   0  450772       1   0  39 -- 1c83e  48 22084              -  0:00 aioserver
 42303 A   0  467074       1   0  39 -- 11e71  56 22088              -  0:00 aioserver
 42303 A   0  508048       1   0  39 -- 12e32  56 22088              -  0:00 aioserver
 42303 A   0  516226       1   0  39 -- 1be3b  48 22080              -  0:00 aioserver
 42303 A   0  524422       1   0  39 -- 6e66  56 22080              -  0:00 aioserver
<...lines omitted...>
```

Since the disk drives that are being accessed asynchronously are using the Journaled File System (JFS), all I/O will be routed through the `aioserver` kprocs.

> **Note:** AIO will not relieve an overly busy disk drive. Using the **`iostat`** command with an *interval* and *count* value, you can determine if any disks are overly busy. Monitor the `%tm_act` column for each disk drive on the system. On some systems, a `%tm_act` of *35* or *higher* for one disk can cause noticeably slower performance. In this case move the data from more busy to less busy disks.

For environments in which the performance of asynchronous disk I/O is critical and the volume of requests is high, but you do not have an approximate number of simultaneous I/Os, it is recommended that `maxservers` be set to at least *10*(number of disks accessed asynchronously)*. The next example shows how to make the initial setting for a system with eight asynchronously accessed disks. However note that the change will only be applied to the Object Data Manager (ODM) database.

Example 3-29 Changing the maxservers attribute for aio0 with the chdev command

```
root@swede:/: chdev -l aio0 -a maxservers=80 -P
aio0 changed
```

In addition, you can set the maximum number of outstanding asynchronous I/O requests (`maxreqs`), and the server priority (`kprocprio`). If you have a system with a high volume of asynchronous I/O requests, it might be appropriate to increase the `maxreqs` number, but avoid lowering the `kprocprio` number since the network I/O from clients to the IBM Tivoli Storage Manger might suffer if the CPU resources are consumed by asynchronous I/O.

> **Note:** When tuning the `kprocprio` please monitor the system so that the increase in fixed priority (lower number) will not produce unwanted side effects.

To determine the maximum number of *requested* asynchronous servers set `maxservers` to 999. Monitor the number of additional servers started throughout the course of normal workload. After a 24-hour period of normal activity (or over a period of one week). To set the appropriate number of asynchronous servers, calculate the number of currently running aios adding 10 on top.

To determine the maximum number of *useful* asynchronous servers use the **`vmstat -s`** output for the metric `iodone`. Start with `maxservers` set at *10*#Disks* and then take first **`vmstat -s`** output before starting IBM Tivoli Storage Manger and then again after a full days workload (24-hours). Increase the number of `maxservers` by doubling the current value (and reboot the system) and make another reading prior to starting IBM Tivoli Storage Manager and then again after another 24-hour period. Continue this doubling/reboot/monitoring cycle until the `iodone` metric does not increase anymore. At this time check the number of maximum asynchronous servers with the **`ps -k|grep -c aio`** command. Set the `maxservers` value to the number of currently maximum asynchronous servers adding 10 on top.

3.1.5 Tivoli Space Manager (HSM) support for 64 bit AIX 5.1 on PPC

AIX 5L provides a new, scalable, 64-bit kernel that:

- Provides simplified data and I/O device sharing for multiple applications on the same system.
- Provides more scalable kernel extensions and device drivers that make full use of the kernel's system resources and capabilities.
- Allows for future hardware development that will provide even larger single image systems ideal for server consolidation or workload scalability.

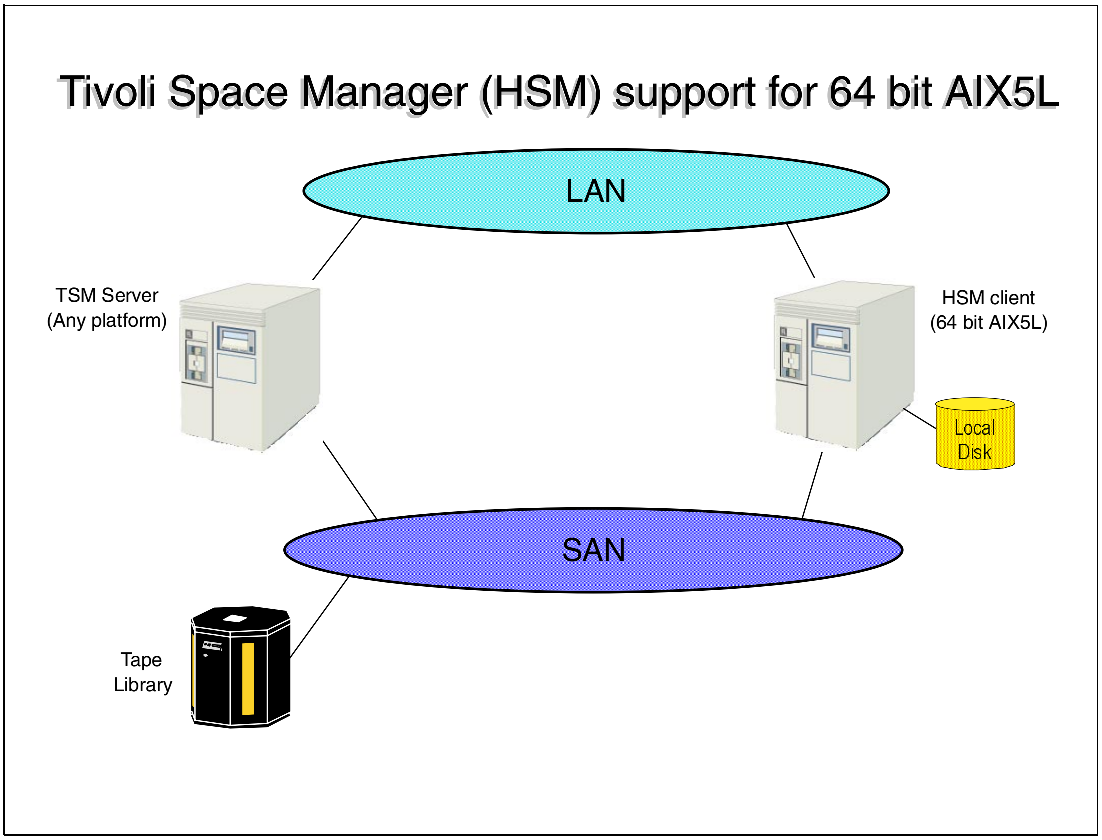

Figure 3-12 HSM support for 64bit AIX5L on PPC

Both 32-bit and 64-bit kernels are available for the POWER platform. If your system has 64-bit processors, the 64-bit kernel is automatically installed with the base operating system. However, the 64-bit kernel is only enabled if you set the "Enable 64-bit Kernel" and "JFS2" option to "yes" during the initial AIX installation.

Only 64-bit kernel extensions are supported with the 64-Bit kernel; this requires we use a 64-Bit Kernel Extension for Tivoli Space Manager (HSM). For that reason, and the dependencies between the BA and HSM clients, we now have a complete 64-Bit client to support the 64-Bit mode of AIX 5L.

Note: This does not affect the 64-bit application environment, which is supported running either the 32-bit or 64-bit kernel. The 64-bit application environment can be enabled or disabled from SMIT under **System Environments**.

Supported environment

Before attempting to configure HSM on AIX5L you should be familiar with the following Tivoli Storage Manager documentation and functions:

- Tivoli Space Manager for UNIX (Using the Hierarchical Storage Management Clients)
- The README and README.HSM of Tivoli Storage Manager Client 5.1.0.
- A working familiarity with AIX 5L
- A working familiarity with the 32-Bit HSM client

Installation of the 64-Bit HSM client it is only supported on an RS/6000 with a 64-Bit CPU, which is currently running the 64-Bit AIX 5L kernel. Refer to “Switch to 32-Bit or 64-Bit mode on AIX 5L” on page 113 to determine which kernel is currently running, and how to switch to the 64-Bit kernel.

Some limitations of the HSM support on 64-Bit AIX 5L are:

- Supports migration of JFS filesystem only (no support for Veritas VxFS)
- The HSMGUI does not support “help” in 64-Bit mode
- The “Read Without Recall” mode is not available in 64-Bit mode

There are no Tivoli Storage Manager server dependencies.

Updated command - dsmattr

The 64-Bit version of Tivoli Space Manager has nearly the same functionality as the 32-Bit version except for the lack of the “Read Without Recall” function in 64-Bit mode.

The **dsmattr** command sets or changes the recall mode for one or more migrated files. The recall mode determines how the HSM client recalls a migrated file when you access it. Because it does not support "Read Without Recall" the 64-Bit dsmattr command has one parameter less than the 32-Bit version.

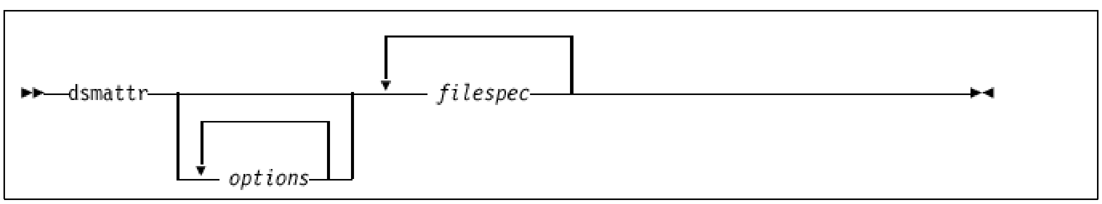

Figure 3-13 The dsmattr command syntax

The options for the dsmattr command are:

-RECAllmode or **-r**

Sets the recall mode for one or more migrated files. If you do not specify this option, the HSM client displays the current recall mode for the files that you specify.

Normal (n) Recalls the migrated file to its originating file system. This is the default. If the file is not modified, it becomes a pre-migrated file. If the file *is* modified, it becomes a resident file.

Migonclose (m) Temporarily recalls a migrated file to its originating file system. If the file is not modified, the file is replaced with a stub file on your local file system. The file returns to a migrated state when it is closed.

Readwithoutrecall (r) Reads the migrated file from storage without storing it on your local file system.

	Information is read sequentially from your file and is cached in a memory buffer on your workstation. **Note:** This parameter is only valid for the 32-Bit client.
–RECUrsive *or* **–R**	Sets or displays the recall mode for migrated files in the directory and subdirectories that you specify.

Switch to 32-Bit or 64-Bit mode on AIX 5L

Selecting the 64-bit kernel on POWER systems AIX 5L for POWER now provides a 64-bit kernel as well as the previously available 32-bit kernel.

Although the 64-bit kernel is automatically installed with the base operating system, it is only enabled if you set the “Enable 64-bit Kernel” and “JFS2” options in the “Advanced Option” screen during the installation. As shown in Figure 3-14, you need to toggle option three to enable the 64-Bit kernel.

```
                              Advanced Options

Either type 0 and press Enter to install with current settings, or type the

number of the setting you want to change and press Enter.

    1 Installation Package Set............ Default

    2 Enable Trusted Computing Base....... no

    3 Enable 64-bit Kernel and JFS2....... yes
>>> 0 Install with the settings listed above.

    88 Help ?

    99 Previous Menu

>>> Choice [0]:
```

Figure 3-14 BOS installation screen for choosing 64-bit kernel

You can also install this kernel later by installing the bos.mp64 fileset.

Only 64-bit CHRP-compliant PowerPC machines are supported for the 64-bit kernel on the POWER platform. Use the “**bootinfo -y**” command to determine the type of hardware, either 32-Bit or 64-Bit, that AIX is running on. If the command returns a 32, you cannot use the 64-bit kernel.

To manually switch from 32-Bit to 64-Bit mode, enter the following commands as “root”:

Example 3-30 Switching from 32-Bit to 64-Bit mode on AIX 5L

```
ln -sf /usr/lib/boot/unix_64 /unix
ln -sf unix_64 /usr/lib/boot/unix
bosboot -ad /dev/ipldevice
shutdown -r
```

After rebooting the machine should be running in 64-Bit mode.

To manually switch from 64-Bit to 32-Bit mode, enter the following commands as “root”:

Example 3-31 Switching from 64-Bit to 32-Bit mode on AIX 5L

```
ln -sf /usr/lib/boot/unix_mp /unix
ln -sf unix_mp /usr/lib/boot/unix
bosboot -ad /dev/ipldevice
```

```
shutdown -r
```

After rebooting the machine should be running in 64-Bit mode. These changes modify the link to the actual AIX kernel. Please be very careful doing this.

Migrating from 32-Bit to 64-Bit version of HSM

Thee following steps are necessary to migrate from 32-Bit version of HSM to 64-Bit version of HSM

1. Uninstall the 32-Bit Tivoli Storage Manager client Version 4.2 or 5.1.
2. Boot AIX 5 L in 64-Bit mode
3. Install the 64-Bit Tivoli Storage Manager client

Conversion from the 64-Bit to 32-Bit Tivoli Storage Manager client is unsupported, but should work if you simply reverse the above procedure.

Troubleshooting

When running the HSM client, the message “ANS9199S Cannot open /dev/fsm” will appear if the kernel extension (kext) is not loaded.

The kernel extension should load automatically at system start. If you installed the 64-Bit package in 64-Bit mode, but later switch the AIX 5L kernel back to 32-Bit mode, 64-Bit kernel extension will fail to load!

Check whether kext is loaded, and switch to 64-Bit mode to rectify this problem.

Example 3-32 Checking whether the HSM kernel extension is loaded

```
root@crete:/: kdb
The specified kernel file is a 64-bit kernel
Preserving 917577 bytes of symbol table
First symbol __mulh
           START              END <name>
0000000000003500 0000000001A94768 _system_configuration+000020
F00000002FF3A600 F00000002FFCF8C8 __ublock+000000
000000002FF22FF4 000000002FF22FF8 environ+000000
000000002FF22FF8 000000002FF22FFC errno+000000
F100008780000000 F100008790000000 pvproc+000000
F100008790000000 F100008794000000 pvthread+000000
F100000040000000 F100000040267488 vmmdseg+000000
F1000013B0000000 F1000083B5400000 vmmswpft+000000
F100000BB0000000 F1000013B0000000 vmmswhat+000000
F100000050000000 F100000060000000 ptaseg+000000
F100000070000000 F1000000B0000000 ameseg+000000
F100009E10000000 F100009E20000000 KERN_heap+000000
F100009C00000000 F100009C10000000 lkwseg+000000
PFT:
id....................0007
raddr.....0000000002000000 eaddr.....0000000002000000
size..............01000000 align.............01000000
valid..1 ros....0 holes..0 io.....0 seg....1 wimg...2

PVT:
id....................0008
raddr.....0000000000946000 eaddr.....0000000000000000
size..............00200000 align.............00001000
valid..1 ros....0 holes..0 io.....0 seg....1 wimg...2
```

```
(0)> lke
ADDRESS     FILE FILESIZE    FLAGS MODULE NAME

  1 0570DE00 01D3FEC0 000026B8 00080272 netintf64/usr/lib/drivers/netintf
  2 05776800 056D8000 00001A90 00180248 /unix
  3 0577F700 01D3B720 00004790 00080262 bpf64/usr/lib/drivers/bpf
  4 056DCB00 05FB6000 0007DFA8 00080272 nfs.ext64/usr/lib/drivers/nfs.ext
  5 056DC100 0604C000 00001A90 00180248 /unix
  6 056DCA00 01D3B2C0 00000438 00080272 nfs_kdes_null.ext64/usr/lib/drivers/nfs_kdes.ext
  7 056DC200 057AE000 00001A20 00180248 /unix
  8 05778F00 05E5C000 00129510 00080262 /usr/tivoli/tsm/client/hsm/bin/kext
...
```

To manually load the HSM kernel extension, enter the following command as “root”:

Example 3-33 Manually loading the HSM kernel extension

```
root@crete:/: /usr/tivoli/tsm/client/hsm/bin/installfsm -q \
> /usr/tivoli/tsm/client/hsm/bin/kext

load.path content is /usr/tivoli/tsm/client/hsm/bin/kext
Performing the sysconfig SYS_QUERYLOAD subcommand
Kernel extension /usr/tivoli/tsm/client/hsm/bin/kext is loaded at 0x05778f00
```

3.2 OS/400 PASE

The IBM eServer iSeries servers are designed and built as a total system, fully integrating all the hardware and system software components that a business demands. The iSeries architecture is a technology-neutral architecture, enabling businesses to readily exploit the latest hardware and software technologies, typically without causing disruption to existing application software.

OS/400 V5R1 provides support for Linux, Lotus Domino, Java, Microsoft Windows, Unix, and iSeries applications. High availability is combined with workload management and logical partitioning.

Portable Application Solutions Environment

OS/400

- System Licensed Internal Code
- Technology Independent Machine Interface
- Integrated File System

TSM

- Library manager and tape drive support

Figure 3-15 Portable Application Solutions Environment

3.2.1 Portable Application Solutions Environment (PASE)

The IBM eServer pSeries and iSeries share a common PowerPC chip. PASE is a fully integrated component of OS/400 that provides a subset of AIX runtime functionality. It exploits the iSeries processor's ability to switch between OS/400 and UNIX run-time modes. Applications deployed using PASE run natively on the OS/400 and take full advantage of its file systems, security and DB2 Universal Database for OS/400. Since PASE applications execute directly on the hardware in PowerPC mode, compute-intensive applications use the processor without any additional layers overhead.

Architecture

PASE is not an operating system. It uses the System Licensed Internal Code (SLIC) kernel for system services. A UNIX application is fully integrated with work management, security, backup, file systems and database.

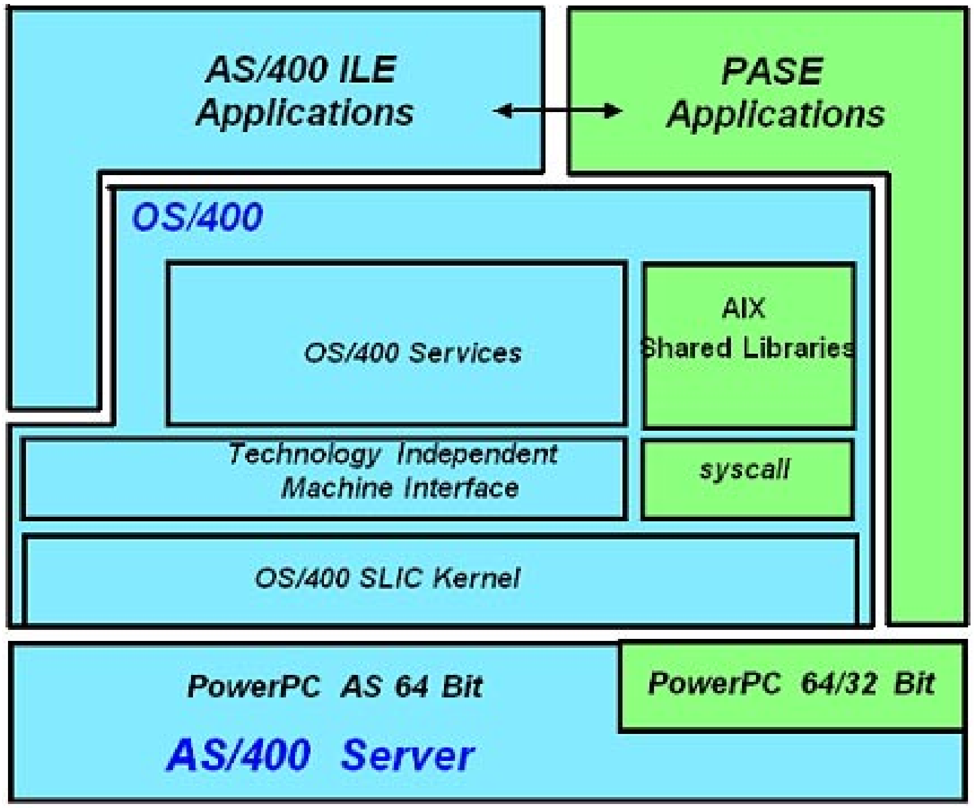

Figure 3-16 How PASE for iSeries fits with the traditional iSeries architecture

The underlying technology that makes OS/400 PASE for iSeries possible is the common Power PC processor that has been shared by pSeries and iSeries. It is a high-speed chip that provides multiple addressing modes. By using the PowerPC AS and PowerPC chip modes in a controlled, secure way, OS/400 PASE helps the iSeries harness this strength. In OS/400 V4R4 and V4R5, the PowerPC AS and 32-bit PowerPC modes are used. As of V5R1, 64-bit PowerPC mode is also used.

The SLIC kernel controls the use of the hardware and the type of address space that can be used. SLIC is a common kernel under the PASE shared libraries and OS/400. The services that SLIC provides are common to both the Integrated Language Environment (ILE) and PASE environments, creating a well-integrated system, regardless of which environment the application is using.

The environments share the same file systems, security mechanisms, threading mechanisms, and can even share memory and connect to the same sockets (with a little care to note pointer format expectations and character set being shared).

SLIC and OS/400 work together to provide an object-based OS on iSeries. PASE has not caused creation of new system objects -- it uses the ones that were already built to support standard APIs for ILE. This facilitates sharing of files and other system information between application code and also gives PASE an operational view that is just what iSeries customers expect. PASE for iSeries applications look like ILE Java or other running applications to iSeries operators.

Programs and files for new PASE applications are saved and restored in familiar ways, without the operator having to know that they are using any new system facilities.

As much of OS/400 sits architecturally on top of the Technology Independent Machine Interface (TIMI), PASE sits architecturally on top of the *syscall interface*. It looks like AIX in this respect. There is a protection layer between the kernel services for PASE and non-privileged system or user code as there is for OS/400, but the syscall is not designed to provide the same virtual machine as the TIMI. OS/400 has the capability to move applications that maintain their intermediate form, to new hardware architectures (i.e., CISC to RISC) without recompiling. This level of compatibility is not provided by the OS/400 Model or expected by Unix solution providers.

PASE applications will be recompiled when the customer moves to new hardware architecture, as one would expect with an AIX application making the same transition, since they have access directly to the hardware processor. However, within an architecture family, upgrades are often transparent to the application

PASE is called from ILE applications. This is how it is started, in a running iSeries job. Applications that use PASE are fully integrated in the customer's iSeries workload and look no different operationally. PASE for iSeries can also call ILE functions including both application code and OS/400 system services. To an iSeries programmer, PASE looks like just another program model. To a UNIX programmer, PASE provides the familiar AIX application binary interface (ABI) for their porting process.

PASE availability

PASE is available as OS/400 Option 33 in V4R4, V4R5, and V5R1. The V4R4 environment provides libraries containing more than 900 basic system APIs that are compatible with current AIX 4.2.1 releases that provide 32 bit support. V4R5 provides around 1500 system APIs plus about 120 shells and utilities that are compatible with AIX 4.3.3.

For more information about PASE please refer to the iSeries Information Center at:
`http://publib.boulder.ibm.com/pubs/html/as400`

Select **Integrated operating environment->OS/400 PASE**.

To learn more about OS/400 please refer to `http://www.as400.ibm.com`.

3.2.2 IBM Tivoli Storage Manager Version 5.1 server on PASE

The IBM Tivoli Storage Manager Version 5.1 server on PASE is installed using the standard iSeries command **RSTLICPGM** (Restore Licensed Program). This is performed like the previous OS/400 native Tivoli Storage Manager server for iSeries.

> **Note:** You must have QSECOFR authority to perform the installation.

Filesystem

The Integrated File System (IFS) incorporates all types of information stored on iSeries systems. This information includes libraries, documents, and folders (file systems that have previously existed on iSeries), new file systems similar to those found on PC and UNIX operating systems, and more.

Within the IFS, the file system object `/QOpenSys` is a case-sensitive UNIX formatted file systems, that is designed to be POSIX 1003.1 compliant and supports POSIX I/O calls. The filesystem hierarchy in `/QOpenSys` is similar to the AIX filesystem hierarchy and Tivoli Storage Manager will be installed under `/usr/tivoli/tsm`. The CD-ROM file system on OS/400 is accessed through `/QOPT`.

For more information on IFS please refer to the whitepaper at:
http://www.ibm.com/servers/eserver/iseries/whpapr/ifs.html.

When adding disk volumes to the Tivoli Storage Manager server in PASE, the volumes will be placed in the `QOpenSys` UNIX hierarchy. By default all volumes defined to the Tivoli Storage Manager server with the **`define volume`**, **`define dbvol`** or **`define logvol`** commands will be placed in the `/usr/tivoli/tsm/server/bin` directory.

Unlike AIX, there are no device drivers and the configuration is done with standard iSeries methods. Please refer to the iSeries manuals for information on how to configure and verify iSeries attached devices.

Initially, PASE did not support the use of tape devices by an application running within PASE. The new APIs for Tivoli Storage Manager tape support provided in the PASE environment is the iSeries standard tape I/O.

Library manager and tape drive support

Since the PASE interface is going through OS/400 to do all the tape functions the iSeries SLIC tape scheduler code controls which resource is used. There is currently no interface available to mount a cartridge to a specific resource.

IBM Tivoli Storage Manager 5.1 server for PASE support:

- 3590
- 3570
- LTO

> Note: Tape and library devices should be specified using the iSeries device description names.

IBM Tivoli Storage Manager 5.1 server for PASE also supports any 8 MM and QIC device that OS/400 supports, however optical devices are *not* supported.

IBM Tivoli Storage Manager 5.1 server for PASE will *not* support the following:

- Server Free
- LAN Free (Storage Agent)
- Tape Library Sharing
- Raw logical volumes
- Optical
- NAS (NDMP)
- SNMP Subagent

> Note: The PATH command is not supported and thus all library and tape definitions have to be performed, like on previous versions of Tivoli Storage Manager, with the DEFINE LIBRARY and DEFINE DRIVE command options.

Examples

The following **`dsmadmc`** command examples define a manual library tape drive on Tivoli Storage Manager in OS/400 PASE from another Tivoli Storage Manager server system. The first **`dsmadmc`** command define the manual library and names it `8MML`:

```
define library 8MML libtype=manual
```

The next step is to define the drive with the pre-5.1 syntax of the DEFINE DRIVE command option naming the drive `8MMD` in the `8MML` library:

```
define drive 8MML 8MMD device=TAP01
```

Example 3-34 shows the interaction between the command line b/a client, **dsmc**, and the Tivoli Storage Manager server running on OS/400 PASE.

Example 3-34 Performing selective backup to a OS/400 PASE Tivoli Storage Manager server

```
root@crete:/: dsmc s '*' -server=pase
Tivoli Storage Manager
Command Line Backup/Archive Client Interface - Version 5, Release 1, Level 0.0 g4
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Selective Backup function invoked.

Node Name: DUDE
Session established with server SERVER1: OS/400 PASE
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/20/02   10:21:42  Last access: 03/20/02   10:21:16

Directory-->                512 /tsm/pbin [Sent]
Normal File-->              546 /tsm/pbin/Build_from_scratch [Sent]
Normal File-->              142 /tsm/pbin/start.tsm [Sent]
Normal File-->           23,719 /tsm/pbin/crTsmServer [Sent]
<...lines omitted...>
Normal File-->              221 /tsm/pbin/more.sql [Sent]
Selective Backup processing of '/tsm/pbin/*' finished without failure.

Total number of objects inspected:       17
Total number of objects backed up:       17
Total number of objects updated:          0
Total number of objects rebound:          0
Total number of objects deleted:          0
Total number of objects expired:          0
Total number of objects failed:           0
Total number of bytes transferred:    53.89 KB
Data transfer time:                    0.00 sec
Network data transfer rate:       116,145.69 KB/sec
Aggregate data transfer rate:         13.44 KB/sec
Objects compressed by:                    0%
Elapsed processing time:           00:00:04
```

3.3 HP-UX

The IBM Tivoli Storage Manager 5.1 for HP-UX introduces LAN-Free data movement support and additional device support, in addition to the general Version 5.1 server enhancements.

3.3.1 Changes in IBM Tivoli Storage Manager Version 5.1 server

The following changes have been made to IBM Tivoli Storage Manager 5.1 for HP-UX.

HP-UX

Changes in IBM Tivoli Storage Manger version 5.1 server

- Tivoli Storage Manager device driver support
- LAN-Free Data Movement
- Changes in defining drives and libraries
- Path names for external libraries
- Moving data by node
- Support for Simultaneous writes to primary and copy storage pools
- Tivoli Data Protection for new NDMP support
- Data validation with Cyclic Redundancy Checking
- New licensing method

Figure 3-17 Changes to IBM Tivoli Storage Manager 5.1 for HP-UX

- Tivoli Storage Manager Device Driver

  The Tivoli Storage Manager device driver (tsmscsi) now supports SCSI autochangers, optical drives, and tape drives. The following device classes are now available on HP-UX: 4MM, 8MM, DLT, DTF, ECARTRIDGE, and QIC.

  See the following changed commands:

  – DEFINE DEVCLASS
  – QUERY DEVCLASS
  – UPDATE DEVCLASS

- LAN-Free Data Movement

  Tivoli Storage Manager supports LAN-free data movement in storage area network (SAN) environments. The support allows client data to move directly from the client system to a server-managed storage device on the SAN. This enhancement reduces data movement on the LAN so that more bandwidth is available to other applications.

  See the following new commands:

  – DEFINE PATH
  – QUERY PATH
  – UPDATE PATH

- Changes in Defining Drives and Libraries

  Device names are now specified in the DEFINE PATH and UPDATE PATH commands, rather than in the DEFINE DRIVE, UPDATE DRIVE, DEFINE LIBRARY, and UPDATE LIBRARY commands

- Path names for external libraries are now specified in the DEFINE PATH command.

See the following new commands:

- DEFINE PATH
- DELETE PATH
- QUERY PATH
- UPDATE PATH

See the following changed commands:

- DEFINE DRIVE
- DEFINE LIBRARY
- QUERY DRIVE
- QUERY LIBRARY
- UPDATE DRIVE
- UPDATE LIBRARY

- Moving Data by Node

  You can use the MOVE NODEDATA command to move data in a sequential-access storage pool for one or more nodes, or move data in a single node with selected file spaces. You can also use MOVE NODEDATA to consolidate data for off-site disaster recovery storage, or move data to another storage pool.

- Support for Simultaneous Writes to Primary and Copy Storage Pools

  You can specify copy storage pools in a primary storage pool definition. When a client backs up, archives, or migrates a file, the file is written to the primary storage pool and is simultaneously stored into each copy storage pool.

  See the following changed commands:

  - DEFINE STGPOOL
  - QUERY SESSION
  - QUERY STGPOOL
  - REGISTER NODE
  - UPDATE NODE
  - UPDATE STGPOOL

- Tivoli Data Protection for New Network Data Management Protocol Support

  New Network Data Management Protocol (NDMP) support now extends to the HP-UX (32-bit and 64-bit) Tivoli Storage Manager server platform. The new Tivoli Data Protection for NDMP product supports NDMP backup and restore for network-attached storage (NAS) file servers from IBM and Network Appliance. NDMP allows a network storage-management application to control the backup and restore of an NDMP-compliant file server without installing third-party software on that server. The NAS file server does not require installation of Tivoli Storage Manager software. The Tivoli Storage Manager server uses NDMP to connect to the NAS file server to initiate, control, and monitor a file system backup or restore operation. The NDMP support for NAS file servers enables higher performance backup to tape devices without moving the data over the LAN. TDP for NDMP is a separately priced and licensed product.

  See the following new commands:

  - BACKUP NODE
  - DEFINE DATAMOVER
  - DEFINE PATH
  - DELETE DATAMOVER
  - DELETE PATH
  - QUERY DATAMOVER
  - QUERY PATH
  - UPDATE DATAMOVER
  - UPDATE PATH

- RESTORE NODE
- Data Validation with Cyclic Redundancy Checking

  Tivoli Storage Manager provides the option of specifying whether a cyclic redundancy check (CRC) is performed during a client session with the server, or for storage pools. The server validates the data by using a cyclic redundancy check which can help identify data corruption. Data validation can be enabled for one or all of the following:

  - Tivoli Storage Manager client nodes at Version 5.1
  - Tivoli Storage Manager storage agents at Version 5.1
  - Storage pools

  See the following changed commands:

  - DEFINE SERVER
  - DEFINE STGPOOL
  - QUERY NODE
  - QUERY SERVER
  - QUERY VOLUME
  - REGISTER NODE
  - UPDATE NODE
  - UPDATE SERVER
  - UPDATE STGPOOL
- New Licensing Method

  The new licensing method enables you to register the exact number of licenses that are required, rather than in increments of 1, 5, 10, and 50.

3.3.2 Extended device support

With IBM Tivoli Storage Manager Version 5.1 there are additional device support for HP-UX 11 and HP-UX 11i, in both 32 and 64-bit mode.

For a complete list of all currently supported devices, please refer to `http://www.tivoli.com/support/storage_mgr/requirements.html`. For Program Temporary Fix (PTF) distributed device support, please refer to the `README.DEV` file in the server installation directory for additional information regarding supported formats for the tape drives.

For information about IBM tape and optical storage media, please refer to `http://www.storage.ibm.com/media/products.html`

HP device driver for IBM TotalStorage and Ultrium products

For IBM 3590 Magstar and IBM LTO devices, you need to use the ATDD device driver. For more detailed information regarding the current ATDD device driver support, please refer to `ftp://ftp.software.ibm.com/storage/devdrvr/HPUX/README`.

> **Note:** IBM Tivoli Storage Manager Version 5.1 will *not* support the HP Precision bus. Please note that the ATDD device driver Version 2.9 and 4.9 are not supported by Tivoli Storage Manager V5.1.

New 4mm drive support

The following additional 4mm drives are supported by IBM Tivoli Storage Manager Version 5.1 on HP-UX.

Table 3-1 4mm drives

ADIC DATa 8008	IBM 7206-110
ADIC DATa 8008i	IBM 7206-220
ADIC DATa 8008E	Seagate SP40 DDS4
Dell PowerVault 100T DDS3	Sony SDT-7000
Dell PowerVault 100T DDS4	Sony SDT-9000
Dell PowerVault 120T DDS3	Sony SDT-10000
IBM 7044	Sony SDT-11000
IBM 7206-005	

The following additional 4mm libraries are supported by IBM Tivoli Storage Manager Version 5.1 on HP-UX.

Table 3-2 4mm tape libraries

Compaq 4/16	Sony TSL-7000
Dell PowerVault120T DDS3	Sony TSL-9000
Dell PowerVault120T DDS4	Sony TSL-S9000L
Diverse Logistics Libra 8	Sony TSL-S10000L
Diverse Logistics Libra 16	Sony TSL-S11000L
IBM 7332-005 4 or 12 slot	Spectralogic 4000/20
IBM 7332-110 4 or 12 slot	Spectralogic 4000/40
IBM 7332-220 4 or 12 slot	Spectralogic 4000/60
IBM 7336-205	StorageTek 9734

New 8mm support

The following additional 8 mm drives are supported by IBM Tivoli Storage Manager Version 5.1 on HP-UX.

Table 3-3 8mm drives

Exabyte EXB-8700LT	Seagate Sidewinder 50
Exabyte EXB-8900	Sun 8505XL
Exabyte Mammoth2	Sony SDX-300
IBM @server xSeries 00N8017	Sony SDX-300C
IBM 7208-011	Sony SDX-400C
IBM 7208-341	Sony SDX-500
IBM 7208-345	Sony SDX-500C
IBM 7334-410	

The following additional 8mm libraries are supported by IBM Tivoli Storage Manager Version 5.1 on HP-UX.

Table 3-4 8mm libraries

ADIC Scalar 1000/AIT	IBM 7331-305
ADIC Scalar 220/AIT	Overland LibraryPro LP2
ADIC Scalar 480/AIT	Qualstar TLS 4212/l
Andataco Encore 10e	Qualstar TLS 4222/l
Andataco Encore 210	Qualstar TLS 412180
Andataco Encore 440	Qualstar TLS 412300
Andataco Encore 120	Qualstar TLS 412600
Andataco Encore 480	Seagate Sidewinder 200
BoxHill Cube Box	Sony Lib - 152
BoxHill Light Box	Sony TSL-A300C
BoxHill Freezer Box	Sony TSL-A400C
BoxHill Borg Box	Sony TSL-A500C
BoxHill Ice Box	Spectralogic 10000/20
EMASS AML/S	Spectralogic 10000/40
Exabyte EXB-430	Spectralogic 2000
Exabyte X-80	Spectralogic 12000
Exabyte X-200	Spectralogic 64000
IBM 7331-205	StorageTek 9711
Qualstar TLS 412360	

New Ecart (3490E) support

The following additional Ecart drives are supported by IBM Tivoli Storage Manager Version 5.1 on HP-UX.

Table 3-5 Ecart drives

Fujitsu M8100A2	StorageTek SD-3
STK T9840B	StorageTek 9940
StorageTek 9490	

The following additional Ecart libraries are supported by IBM Tivoli Storage Manager Version 5.1 on HP-UX.

Table 3-6 Ecart libraries

StorageTek 9738	StorageTek L700e
StorageTek L180	Hewlett Packard SureStore 10/180
StorageTek L700	Hewlett Packard SureStore 20/700

New DLT support

The following additional Digital Linear Tape (DLT) drives are supported by IBM Tivoli Storage Manager Version 5.1 on HP-UX.

Table 3-7 DLT tape drives

Benchmark DLT1	Dell PowerVault 110T (SDLT Drive)
DEC TZ89	Hewlett-Packard DLT1
Dell PowerVault 110T DLT 4000	Quantum SDLT
Dell PowerVault 110T DLT 7000	

The following additional Digital Linear Tape (DLT) libraries are supported by IBM Tivoli Storage Manager Version 5.1 on HP-UX.

Table 3-8 DLT tape libraries

ADIC VLS 400	DELL PowerVault 128T SDLT
ADIC FastStor 7	DELL PowerVault 130T DLT4000
ADIC FastStor 22	DELL PowerVault 130T DLT7000
ADIC FastStor 22	DELL PowerVault 136T SDLT
Compaq TL891	Exabyte 18D
Compaq TL895	Exabyte 230D
Compaq ESL 9198 DLX	Exabyte 690D
Compaq ESL 9198 SL	NEC DL0101H
Compaq ESL 9326	Overland Data LibraryXpress LXB8110
Compaq MSL5026SL	Overland Data LibraryXpress LXB8210
Compaq MSL 5026 DLX	Overland Neo LXN2000
DEC TZ875/TZ875N	Overland Neo LXN4000
DEC TZ877/TZ877N	StorageTek L40 (1 or 2 drives)
DEC TZ885/TZ885N	StorageTek L40 (3 or more drives)
DEC TZ887/TZ887N	StorageTek L80
DEC TL891	Sun StorEdge L1000
DEC TL892	Sun StorEdge L1800
DEC TL810/TL894	Sun ETL 3500
DEC TL820/TL822/TL893	Sun StorEdge L11000
DEC TL826/TL896	TL891DLX
DELL PowerVault 120T DLT	

New LTO support

The following additional Linear Tape-Open (LTO) drives are supported by IBM Tivoli Storage Manager Version 5.1 on HP-UX.

Table 3-9 LTO tape drives

Dell PowerVault 110T LTO	IBM 24P2396 Half High LTO Tape Drive
HP Ultrium 1 Tape Drive (Models 215 and 230)	Seagate LTO Tape Drive

The following additional Linear Tape-Open (LTO) libraries are supported by IBM Tivoli Storage Manager Version 5.1 on HP-UX.

Table 3-10 LTO tape libraries

ADIC Scalar 1000/LTO (with IBM LTO Drives)	Overland Neo LXN2000
ADIC Scalar 10K	Overland Neo LXN4000
ADIC FastStor 7	Plasmon FTA 21.x
ATL Products M1500	Plasmon V-100 (FTA102)
ATL Products P2000	Qualstar TLS 8111
ATL Products P3000	Qualstar TLS 8211
Exabyte 221L	Qualstar TLS 8222
HP SureStore 1/9	Qualstar TLS 8433
HP SureStore 2/20	Qualstar TLS 8466
HP SureStore 4/40	Qualstar TLS 88132
HP SureStore 6/60	Qualstar TLS 88264
HP SureStore 6/140	StorageTek L20
HP SureStore 8/80	StorageTek L40 (1 or 2 drives)
HP SureStore 10/100	StorageTek L40 (3 or more drives)
HP SureStore 10/180	StorageTek L80
HP SureStore 20/700	StorageTek L700e
NEC LL0101H	

New QIC support

The following additional Quarter Inch Cartridge (QIC) tape drives are supported by IBM Tivoli Storage Manager Version 5.1 on HP-UX.

Table 3-11 QIC tape drives

IBM 7207-012	Tandberg 1440
Tandberg SLR24	Tandberg SLR Autoloader
Tandberg SLR32	Sony DTF Tape Drives:
Tandberg SLR40	SONY GY-2120
Tandberg SLR50	SONY GY-8240
Tandberg SLR60	Sony DTF Tape Libraries:
Tandberg SLR100	Compaq SSL2000 Series
QIC Tape Libraries	GRAU IVD-340i

Tandberg 1210	Sony LIB-304
Tandberg 1420	

New optical support

The following additional optical drives are supported by IBM Tivoli Storage Manager Version 5.1 on HP-UX.

Table 3-12 Optical drives

HP Model 1300T	IBM 0632-C4B
HP Model 1716T	IBM 0632-C4D
HP Model 2600fx	IBM 7209 Model 001
HP Model 5200ex	IBM 7209 Model 002
IBM 0632-C1/AAA	IBM 7209 Model 003
IBM 0632-C2/ACA	IBM 3510 Model 001
IBM 0632-C2B/CCA	Sony SMO-F541
IBM 0632-CHA/CHB/CHC	Sony SMO-F551
IBM 0632-C4A	

The following additional optical jukeboxes are supported by IBM Tivoli Storage Manager Version 5.1 on HP-UX.

Table 3-13 Optical jukeboxes

DISC DocuStore D150U-1	HP 160ex
DISC DocuStore D245-1	HP 320ex
DISC DocuStore D255-1U	HP 400ex
DISC DocuStore D280U-1	HP 660ex
DISC DocuStore D350-1U	HP 1200ex
DISC DocuStore D510-2	IBM 3995 Model A63
DISC DocuStore D525-1	IBM 3995 Model C60
DISC DocuStore D525-1U	IBM 3995 Model 063
DISC DocuStore D1050-2	IBM 3995 Model C62
HP 20T	IBM 3995 Model 163
HP 40T	IBM 3995 Model C64
HP 40st	IBM 3995 Model C66
HP 80st	IBM 3995 Model C68
HP 100st	Plasmon M20J
HP 120T	Plasmon M32J
HP 165st	Plasmon M52J
HP 200T	Plasmon M104J

HP 300st	Plasmon M156J
HP 40fx	Plasmon M258J
HP 80fx	Plasmon M500J
HP 160fx	Sony OSL-2500
HP 200fx	Sony OSL-6000
HP 330fx	Sony OSL-10000
HP 600fx	Sony OSL-15000
HP 80ex	Sony OSL-25000
HP 125ex	

3.4 Solaris

The IBM Tivoli Storage Manager 5.1 for Sun Solaris does not introduce any Solaris specific enhancements, in addition to the general Version 5.1 server enhancements.

3.4.1 Changes in IBM Tivoli Storage Manager Version 5.1 server

The following changes have been made to IBM Tivoli Storage Manager 5.1 server for Sun Solaris.

Solaris

Changes in IBM Tivoli Storage Manger version 5.1 server

- Changes in defining drives and libraries
- Path names for external libraries
- Moving data by node
- Support for Simultaneous writes to primary and copy storage pools
- Tivoli Data Protection for new NDMP support
- Data validation with Cyclic Redundancy Checking
- New licensing method

Figure 3-18 Changes for Solaris

- Changes in Defining Drives and Libraries
  - Device names are now specified in the DEFINE PATH and UPDATE PATH commands, rather than in the DEFINE DRIVE, UPDATE DRIVE, DEFINE LIBRARY, and UPDATE LIBRARY commands
- Path names for external libraries are now specified in the DEFINE PATH command.

  See the following new commands:

  - DEFINE PATH
  - DELETE PATH
  - QUERY PATH
  - UPDATE PATH

  See the following changed commands:

  - DEFINE DRIVE
  - DEFINE LIBRARY
  - QUERY DRIVE
  - QUERY LIBRARY
  - UPDATE DRIVE
  - UPDATE LIBRARY
- Moving Data by Node

  You can use the MOVE NODEDATA command to move data in a sequential-access storage pool for one or more nodes, or move data in a single node with selected file spaces. You can also use MOVE NODEDATA to consolidate data for off-site disaster recovery storage, or move data to another storage pool.
- Support for Simultaneous Writes to Primary and Copy Storage Pools

  You can specify copy storage pools in a primary storage pool definition. When a client backs up, archives, or migrates a file, the file is written to the primary storage pool and is simultaneously stored into each copy storage pool.

  See the following changed commands:

  - DEFINE STGPOOL
  - QUERY SESSION
  - QUERY STGPOOL
  - REGISTER NODE
  - UPDATE NODE
  - UPDATE STGPOOL
- Tivoli Data Protection for New Network Data Management Protocol Support

  New Network Data Management Protocol (NDMP) support now extends to the Sun Solaris (32-bit and 64-bit) Tivoli Storage Manager server platform. The new Tivoli Data Protection for NDMP product supports NDMP backup and restore for network-attached storage (NAS) file servers from IBM and Network Appliance. NDMP allows a network storage-management application to control the backup and restore of an NDMP-compliant file server without installing third-party software on that server. The NAS file server does not require installation of Tivoli Storage Manager software. The Tivoli Storage Manager server uses NDMP to connect to the NAS file server to initiate, control, and monitor a file system backup or restore operation. The NDMP support for NAS file servers enables higher performance backup to tape devices without moving the data over the LAN. TDP for NDMP is a separately priced and licensed product.

  See the following new commands:

  - BACKUP NODE
  - DEFINE DATAMOVER
  - DEFINE PATH

– DELETE DATAMOVER
– DELETE PATH
– QUERY DATAMOVER
– QUERY PATH
– UPDATE DATAMOVER
– UPDATE PATH
– RESTORE NODE

- Data Validation with Cyclic Redundancy Checking

  Tivoli Storage Manager provides the option of specifying whether a cyclic redundancy check (CRC) is performed during a client session with the server, or for storage pools. The server validates the data by using a cyclic redundancy check which can help identify data corruption. Data validation can be enabled for one or all of the following:

  – Tivoli Storage Manager client nodes at Version 5.1
  – Tivoli Storage Manager storage agents at Version 5.1
  – Storage pools

  See the following changed commands:

  – DEFINE SERVER
  – DEFINE STGPOOL
  – QUERY NODE
  – QUERY SERVER
  – QUERY VOLUME
  – REGISTER NODE
  – UPDATE NODE
  – UPDATE SERVER
  – UPDATE STGPOOL

- New Licensing Method

  The new licensing method enables you to register the exact number of licenses that are required, rather than in increments of 1, 5, 10, and 50.

3.4.2 Installing Tivoli Storage Manager 5.1 for Sun Solaris

The IBM Tivoli Storage Manager server packages for Solaris are:

TIVsmS The base product with English language support.

TIVsmSdev The device driver package. Install this package if you are using any devices other than IBM 3490, 3570, and 3590 tape drives.

TIVsmSlic The server license module and *.lic files.

Use the `pkgadd` command to install the selected packages.

The default installation directories for IBM Tivoli Storage Manager server on Solaris are:

Tivoli Storage Manager server /opt/tivoli/tsm/server/bin

Tivoli Storage Manager Devices Drivers/opt/tivoli/tsm/server/bin

Tivoli Storage Manager server licenses /opt/tivoli/tsm/server/bin

For more detailed information about installation and configuration please refer to the *Tivoli Storage Manager for Sun Solaris Quick Start Version 5 Release 1*, GC32-0780 or the *Tivoli Storage Manager for Sun Solaris Administrator’s Guide Version 5 Release 1*, GC32-0778.

The following procedure is outlined in greater detail in the *Tivoli Storage Manager for Sun Solaris Quick Start Version 5 Release 1*, GC32-0780, but some additional screen examples are shown here.

1. Review the `README` and `README.LIC` files in the `/opt/tivoli/tsm/server directory`.
2. List the content of the installation media with the **pkginfo** command.

Example 3-35 pkginfo

```
root@solemio:/: pkginfo -d /cdrom/cdrom0
application TIVsmS         Tivoli Storage Manager/Sun Server
application TIVsmSdev      Tivoli Storage Manager/Sun Solaris Device Drivers
application TIVsmSlic      Tivoli Storage Manager/Sun Server LICENSES
```

3. Install the Tivoli Storage Manager Version 5.1 server software with the **pkgadd** command.

Example 3-36 pkgadd

```
root@solemio:/: pkgadd -d /cdrom/cdrom0

The following packages are available:
  1  TIVsmS        Tivoli Storage Manager/Sun Server
                   (sparc) Version 5 Release 1 Level 0.0
  2  TIVsmSdev     Tivoli Storage Manager/Sun Solaris Device Drivers
                   (sparc) Version 5 Release 1 Level 0.0
  3  TIVsmSlic     Tivoli Storage Manager/Sun Server LICENSES
                   (sparc) Version 5 Release 1 Level 0.0

Select package(s) you wish to process (or 'all' to process
all packages). (default: all) [?,??,q]: all

Processing package instance <TIVsmS> from </cdrom/cdrom0/TSMSRV05_01SUN.pkg>

Tivoli Storage Manager/Sun Server
(sparc) Version 5 Release 1 Level 0.0
Using </opt> as the package base directory.

Tivoli Storage Manager Server for Sun Solaris
Licensed Materials - Property of IBM
5698-ISE (C) Copyright IBM Corporation 1990, 2002. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

## Processing package information.
## Processing system information.
   5 package pathnames are already properly installed.
## Verifying package dependencies.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.

This package contains scripts which will be executed with super-user
permission during the process of installing this package.

Do you want to continue with the installation of <TIVsmS> [y,n,?] y

Installing Tivoli Storage Manager/Sun Server as <TIVsmS>

## Installing part 1 of 1.
/etc/tivready/monitorslfs/Sun_ADSM_Server.slf
/opt/tivoli/tsm/server/README
```

```
...<lines omitted>...
/opt/tivoli/tsm/server/webimages/WEBSERV2.GIF
[ verifying class <none> ]
## Executing postinstall script.

Installing 64-bit server

***************************************
UPGRADE UPGRADE UPGRADE UPGRADE UPGRADE
***************************************

Upgrading Database and Web Admin interface... Please wait.
There will be a delay after message ANR4693I, be patient.



Tivoli Storage Manager for Sun Solaris 7/8
Version 5, Release 1, Level 0.0

Licensed Materials - Property of IBM

5698-ISE (C) Copyright IBM Corporation 1990, 2002. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

ANR7800I DSMSERV generated at 11:39:12 on Feb  1 2002.
ANR7801I Subsystem process ID is 6082.
ANR9999D Error reading from standard input; console input daemon terminated.
ANR0990I Server restart-recovery in progress.
ANR0200I Recovery log assigned capacity is 8 megabytes.
ANR0201I Database assigned capacity is 12 megabytes.
ANR0306I Recovery log volume mount in progress.
ANR0353I Recovery log analysis pass in progress.
ANR0354I Recovery log redo pass in progress.
ANR0355I Recovery log undo pass in progress.
ANR0352I Transaction recovery complete.
ANR2100I Activity log process has started.
ANR4726I The NAS-NDMP support module has been loaded.
ANR1305I Disk volume /opt/tivoli/tsm/server/bin/sbkup00.dsm varied online.
ANR1305I Disk volume /opt/tivoli/tsm/server/bin/sarch00.dsm varied online.
ANR2803I License manager started.
ANR2827I Server is licensed to support Tivoli Space Manager for a quantity of 50.
ANR2860I Server is licensed to support Tivoli Disaster Recovery Manager.
...<lines omitted>...
ANR2828I Server is licensed to support Tape Library Sharing.
ANR2827I Server is licensed to support Managed System for LAN for a quantity of 50.
ANR2827I Server is licensed to support Managed System for SAN for a quantity of 50.
ANR2827I Server is licensed to support Managed Library for a quantity of 1.
ANR2828I Server is licensed to support Tivoli Data Protection for NDMP.
ANR2560I Schedule manager started.
ANR0993I Server initialization complete.
ANR0916I TIVOLI STORAGE MANAGER distributed by Tivoli is now ready for use.
ANR4693I Interface Driver information will be loaded in quiet mode: Only warning and error
messages will be displayed.
ANR4980I Auditing Interface Driver definitions.
ANR4983I Auditing Interface Driver Groups.
ANR4985I Auditing Interface Driver Group Members.
ANR4986I Auditing Interface Driver Classes.
ANR4988I Auditing Interface Driver Complex Class containers.
ANR4991I Auditing Interface Driver Tasks.
```

```
ANR4992I Auditing Interface Driver Task Members.
ANR4989I Auditing Interface Driver Operations.
ANR4990I Auditing Interface Driver Operation Parameters.
ANR4982I Interface Driver audit completed - definitions are consistent.
Installation Complete.


********************************************************************
IMPORTANT: Read the contents of file /opt/tivoli/tsm/server/README
           for extensions and corrections to printed
           product documentation.

********************************************************************

### Please wait until you get the system prompt.  Solaris is updating
    package information files....

Installation of <TIVsmS> was successful.

Processing package instance <TIVsmSdev> from </cdrom/cdrom0/TSMSRV05_01SUN.pkg>

Tivoli Storage Manager/Sun Solaris Device Drivers
(sparc) Version 5 Release 1 Level 0.0

Tivoli Storage Manager Server for Sun Solaris
Licensed Materials - Property of IBM
5698-ISE (C) Copyright IBM Corporation 1990, 2002. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

Using </opt> as the package base directory.
## Processing package information.
## Processing system information.
2 package pathnames are already properly installed.
## Verifying package dependencies.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.

This package contains scripts which will be executed with super-user
permission during the process of installing this package.

Do you want to continue with the installation of <TIVsmSdev> [y,n,?] y

Installing Tivoli Storage Manager/Sun Solaris Device Drivers as <TIVsmSdev>

## Installing part 1 of 1.
/opt/tivoli/tsm/devices/README.DEVICES
...<lines omitted>...
/opt/tivoli/tsm/devices/bin/t_parent
[ verifying class <none> ]
## Executing postinstall script.

Installing 64-bit device drivers


Installation of <TIVsmSdev> was successful.

Processing package instance <TIVsmSlic> from </cdrom/cdrom0/TSMSRV05_01SUN.pkg>
```

```
Tivoli Storage Manager/Sun Server LICENSES
(sparc) Version 5 Release 1 Level 0.0

Tivoli Storage Manager Server for Sun Solaris
Licensed Materials - Property of IBM
5698-ISE (C) Copyright IBM Corporation 1990, 2002. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

Using </opt> as the package base directory.
## Processing package information.
## Processing system information.
   4 package pathnames are already properly installed.
## Verifying package dependencies.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.

Installing Tivoli Storage Manager/Sun Server LICENSES as <TIVsmSlic>

## Installing part 1 of 1.
/opt/tivoli/tsm/server/README.LIC
/opt/tivoli/tsm/server/bin/dsmreg32.sl
/opt/tivoli/tsm/server/bin/dsmreg64.sl
/opt/tivoli/tsm/tivinv/TSMSRVSUN05-01.SIG
[ verifying class <none> ]

Installation of <TIVsmSlic> was successful.
```

4. Start the server with the **dsmserv** command.

Example 3-37 Starting the Tivoli Storage Manager server manually

```
root@solemio:/opt/tivoli/tsm/server/bin: ./dsmserv &

Tivoli Storage Manager for Sun Solaris 7/8
Version 5, Release 1, Level 0.0

Licensed Materials - Property of IBM

5698-ISE (C) Copyright IBM Corporation 1990, 2002. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

ANR7800I DSMSERV generated at 11:39:12 on Feb  1 2002.
ANR7801I Subsystem process ID is 6161.
ANR0990I Server restart-recovery in progress.
ANR0200I Recovery log assigned capacity is 8 megabytes.
ANR0201I Database assigned capacity is 12 megabytes.
ANR0306I Recovery log volume mount in progress.
ANR0353I Recovery log analysis pass in progress.
ANR0354I Recovery log redo pass in progress.
ANR0355I Recovery log undo pass in progress.
ANR0352I Transaction recovery complete.
ANR2100I Activity log process has started.
ANR4726I The NAS-NDMP support module has been loaded.
ANR1305I Disk volume /opt/tivoli/tsm/server/bin/sbkup00.dsm varied online.
ANR1305I Disk volume /opt/tivoli/tsm/server/bin/sarch00.dsm varied online.
ANR0984I Process 1 for EXPIRATION started in the BACKGROUND at 15:26:58.
ANR0811I Inventory client file expiration started as process 1.
ANR2803I License manager started.
```

```
ANR0812I Inventory file expiration process 1 completed: examined 0 objects, deleting 0
backup
objects, 0 archive objects, 0 DB backup volumes, and 0 recovery plan files. 0 errors were
encountered.
ANR8200I TCP/IP driver ready for connection with clients on port 1500.
ANR8285I Shared Memory driver ready for connection with clients on port 1510
ANR8190I HTTP driver ready for connection with clients on port 1580.
ANR2560I Schedule manager started.
ANR0984I Process 2 for AUDIT LICENSE started in the BACKGROUND at 15:26:58.
ANR2820I Automatic license audit started as process 2.
ANR0985I Process 1 for EXPIRATION running in the BACKGROUND completed with completion state
SUCCESS
at 15:26:58.
ANR0993I Server initialization complete.
ANR0916I TIVOLI STORAGE MANAGER distributed by Tivoli is now ready for use.
ANR2825I License audit process 2 completed successfully - 1 nodes audited.
ANR0987I Process 2 for AUDIT LICENSE running in the BACKGROUND processed 1 items with a
completion
state of SUCCESS at 15:26:58.
TSM:SERVER1>
```

5. Register licenses.
6. Install the API package with the **pkgadd** command.

Example 3-38 pkgadd

```
root@solemio:/: pkgadd -d /cdrom/cdrom0/TIVsmCapi.pkg TIVsmCapi

Processing package instance <TIVsmCapi> from </cdrom/cdrom0/TIVsmCapi.pkg>

Tivoli Storage Manager Solaris 2.6/7/8 API Client
(sparc) Version 5 Release 1 Level 0
Licensed Materials - Property of IBM

5655-A30
(C) Copyright International Business Machines Corp. 1990, 2001.

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

Using </opt/tivoli/tsm/client> as the package base directory.
## Processing package information.
## Processing system information.
## Verifying package dependencies.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.

The following files are being installed with setuid and/or setgid
permissions:
  /opt/tivoli/tsm/client/api/bin/dsmtca <setuid root>
  /opt/tivoli/tsm/client/api/bin64/dsmtca <setuid root>

Do you want to install these as setuid/setgid files [y,n,?,q] y

This package contains scripts which will be executed with super-user
permission during the process of installing this package.

Do you want to continue with the installation of <TIVsmCapi> [y,n,?] y
```

```
Installing Tivoli Storage Manager Solaris 2.6/7/8 API Client as <TIVsmCapi>

## Executing preinstall script.
## Installing part 1 of 1.
/opt/tivoli/tsm/client/api/README
/opt/tivoli/tsm/client/api/README.API
...<lines omitted>...
/opt/tivoli/tsm/client/api/bin64/sample <symbolic link>
[ verifying class <none> ]
## Executing postinstall script.

Completing 64 bit TSM API installation ....


WARNING:
-------
TSM API requires the following OS patche(s):
 106950-15 106327-10 106300-11

These patches are available via http://sunsolve.sun.com.
Make sure to install these patches in the given order.


Installation of <TIVsmCapi> was successful.
```

7. Install the client package with the **pkgadd** command.

Example 3-39 pkgadd

```
root@solemio:/: pkgadd -d /cdrom/cdrom0/TIVsmCba.pkg  TIVsmCba

Processing package instance <TIVsmCba> from </cdrom/cdrom0/TIVsmCba.pkg>

Tivoli Storage Manager Solaris 2.6/7/8 Backup Archive Client
(sparc) Version 5 Release 1 Level 0

This appears to be an attempt to install the same architecture and
version of a package which is already installed.  This installation
will attempt to overwrite this package.


The installation of this package was previously terminated and
installation was never successfully completed.

Do you want to continue with the installation of <TIVsmCba> [y,n,?] y
Licensed Materials - Property of IBM

5655-A30
(C) Copyright International Business Machines Corp. 1990, 2001.

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

Using </opt/tivoli/tsm/client> as the package base directory.
## Processing package information.
## Processing system information.
   7 package pathnames are already properly installed.
## Verifying package dependencies.
## Verifying disk space requirements.
```

```
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.

The following files are being installed with setuid and/or setgid
permissions:
* /opt/tivoli/tsm/client/ba/bin/dsmtca <setuid root>

* - overwriting a file which is also setuid/setgid.

Do you want to install these as setuid/setgid files [y,n,?,q] y

This package contains scripts which will be executed with super-user
permission during the process of installing this package.

Do you want to continue with the installation of <TIVsmCba> [y,n,?] y

Installing Tivoli Storage Manager Solaris 2.6/7/8 Backup Archive Client as <TIVsmCba>

## Executing preinstall script.

## Installing part 1 of 1.
/opt/tivoli/tsm/client/admin/bin/dsmadmc
/opt/tivoli/tsm/client/ba/README
/opt/tivoli/tsm/client/ba/README.NAS
/opt/tivoli/tsm/client/ba/README.WEBCLI
...<lines omitted>...
/opt/tivoli/tsm/client/ba/bin/plugins/libPiIMG.so
/opt/tivoli/tsm/client/ba/bin/plugins/libPiNAS.so
[ verifying class <none> ]
## Executing postinstall script.

Completing TSM Backup Archive Client installation .....


WARNING:
-------
TSM Backup Archive Client requires the following OS patche(s):
 106950-15 106327-10 106300-11

These patches are available via http://sunsolve.sun.com.
Make sure to install these patches in the given order.


Installation of <TIVsmCba> was successful.
```

8. Examine the installed packages with the following simple command line script.

```
pkginfo | grep TIVsm | while read dummy pkg dummy;do pkginfo -l $pkg;done
```

Example 3-40 Output example of the pkginfo command line script

```
root@solemio:/: pkginfo | grep TIVsm | while read dummy pkg dummy;do pkginfo -l $pkg;done
PKGINST:  TIVsmCapi
      NAME:  Tivoli Storage Manager Solaris 2.6/7/8 API Client
  CATEGORY:  application
      ARCH:  sparc
   VERSION:  Version 5 Release 1 Level 0
   BASEDIR:  /opt/tivoli/tsm/client
    VENDOR:  Tivoli Systems Inc.
    PSTAMP:  2002.025.11:52:15
  INSTDATE:  Mar 11 2002 15:54
```

```
  STATUS:  completely installed
   FILES:     79 installed pathnames
               6 directories
               5 executables
               2 setuid/setgid executables
           24933 blocks used (approx)

 PKGINST:  TIVsmCba
    NAME:  Tivoli Storage Manager Solaris 2.6/7/8 Backup Archive Client
CATEGORY:  application
    ARCH:  sparc
 VERSION:  Version 5 Release 1 Level 0
 BASEDIR:  /opt/tivoli/tsm/client
  VENDOR:  Tivoli Systems Inc.
  PSTAMP:  2002.025.11:51:45
INSTDATE:  Mar 11 2002 15:57
  STATUS:  completely installed
   FILES:    114 installed pathnames
               7 directories
              10 executables
               1 setuid/setgid executables
           62261 blocks used (approx)

 PKGINST:  TIVsmS
    NAME:  Tivoli Storage Manager/Sun Server
CATEGORY:  application
    ARCH:  sparc
 VERSION:  Version 5 Release 1 Level 0.0
 BASEDIR:  /opt
  VENDOR:  IBM Corporation
  PSTAMP:  2002.032.19:11:39
INSTDATE:  Mar 11 2002 15:19
  STATUS:  completely installed
   FILES:    878 installed pathnames
               5 shared pathnames
              11 directories
             827 executables
           80717 blocks used (approx)

 PKGINST:  TIVsmSdev
    NAME:  Tivoli Storage Manager/Sun Solaris Device Drivers
CATEGORY:  application
    ARCH:  sparc
 VERSION:  Version 5 Release 1 Level 0.0
 BASEDIR:  /opt
  VENDOR:  IBM Corporation
  PSTAMP:  2002.032.19:11:59
INSTDATE:  Mar 11 2002 15:24
  STATUS:  completely installed
   FILES:     25 installed pathnames
               2 shared pathnames
               4 directories
              17 executables
           12576 blocks used (approx)

 PKGINST:  TIVsmSlic
    NAME:  Tivoli Storage Manager/Sun Server LICENSES
CATEGORY:  application
    ARCH:  sparc
 VERSION:  Version 5 Release 1 Level 0.0
```

```
   BASEDIR:  /opt
    VENDOR:  IBM Corporation
    PSTAMP:  2002.032.19:11:58
  INSTDATE:  Mar 11 2002 15:24
    STATUS:  completely installed
     FILES:      9 installed pathnames
                 4 shared pathnames
                 5 directories
                 3 executables
              7106 blocks used (approx)
```

9. Create a basic `dsm.opt` and `dsm.sys` configuration files in the `/usr/bin` directory (default).

Example 3-41 dsm.opt and dsm.sys

```
root@solemio:/: cat /usr/bin/dsm.opt
SErvername        TSM_SUN

root@solemio:/: cat /usr/bin/dsm.sys
SErvername  TSM_SUN
   TCPServeraddress   localhost
```

10. Verify the connection to the local server with the **dsmadmc** command.

Example 3-42 dsmadmc

```
root@solemio:/: dsmadmc -id=admin -password=admin q admin
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM_SUN: Solaris 7/8
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/11/02   17:39:40  Last access: 03/11/02   17:39:23

ANS8000I Server command: 'q admin'

Administrator      Days Since     Days Since    Locked?     Privilege Classes
Name              Last Access   Password Set
--------------   ------------   ------------   ----------   -----------------------
ADMIN                      <1             <1       No       System
CLIENT                     <1             <1       No       Client Owner
SERVER_CONSOLE                                     No       System

ANS8002I Highest return code was 0.
```

Installation warning

If the server package (`TIVsmS`) is installed before the Sun Server LICENSES, the warning in Example 3-43 will be displayed.

Example 3-43 Installation warning message

```
ANR9613W Error loading /opt/tivoli/tsm/server/bin//./dsmreg64.sl for
  Licensing function: ld.so.1: ./dsmserv: fatal error:
  /opt/tivoli/tsm/server/bin//./dsmreg64.sl: open failure: No file nor directory exists.
```

These messages can be ignored, the order the packages are installed does not affect the success of the install.

3.4.3 Removing previous versions of Tivoli Storage Manager for Sun Solaris

If there is a version of ADstar Storage Manager (ADSM) or Tivoli Storage Manager (TSM) already installed on the Solaris system it should be removed prior to installing IBM Tivoli Storage Manager 5.1. Below is the procedure for removing Tivoli Storage Manager V4.1 and later from Solaris, for versions prior to Tivoli Storage Manager V4.1, instead of using "`TIVsm`" in the grep command, use "`IBMadsm`".

The following procedure is outlined in greater detail in the *Tivoli Storage Manager for Sun Solaris Quick Start Version 5 Release 1*, GC32-0780, but some additional screen examples are shown here.

1. Stop the server if it is running.
2. Remove any installed PTFs.
3. List the installed Tivoli Storage Manager packages with the **`pkginfo`** command.

Example 3-44 pkginfo

```
root@solemio:/: pkginfo | grep TIVsm
application TIVsmS          Tivoli Storage Manager/Sun Server (sparc) Version 4 Release 2
Level 0.0
application TIVsmCapi      Tivoli Storage Manager Solaris 2.6/7/8 API Client
application TIVsmCba       Tivoli Storage Manager Solaris 2.6/7/8 Backup Archive Client
application TIVsmSdev      Tivoli Storage Manager/Sun Solaris Device Drivers
application TIVsmSlic      Tivoli Storage Manager/Sun Server LICENSES
application TIVsmSsta      Tivoli Storage Manager/Sun Storage Agent
```

4. Remove the previous version of the server software with the **`pkgrm`** command.

Example 3-45 Removing software with the pkgrm command

```
root@solemio:/: pkgrm TIVsmS TIVsmCba TIVsmCapi TIVsmSdev TIVsmSlic TIVsmSsta

The following package is currently installed:
   TIVsmS          Tivoli Storage Manager/Sun Server
                   (sparc) Version 4 Release 2 Level 0.0

Do you want to remove this package? y

## Removing installed package instance <TIVsmS>
## Verifying package dependencies.
## Processing package information.
## Removing pathnames in class <none>
/opt/tivoli/tsm/server/webimages/WEBSERV2.GIF
/opt/tivoli/tsm/server/webimages/WEBSERV.GIF
...<lines omitted>...
## Updating system information.

Removal of <TIVsmSsta> was successful.
```

3.5 Windows

The IBM Tivoli Storage Manager 5.1 for Windows introduces Server-Free Data Movement, specific to the Windows operating system, in addition to the general Version 5.1 server enhancements.

3.5.1 Changes in IBM Tivoli Storage Manager Version 5.1 server

The following changes have been made to IBM Tivoli Storage Manager 5.1 for Windows.

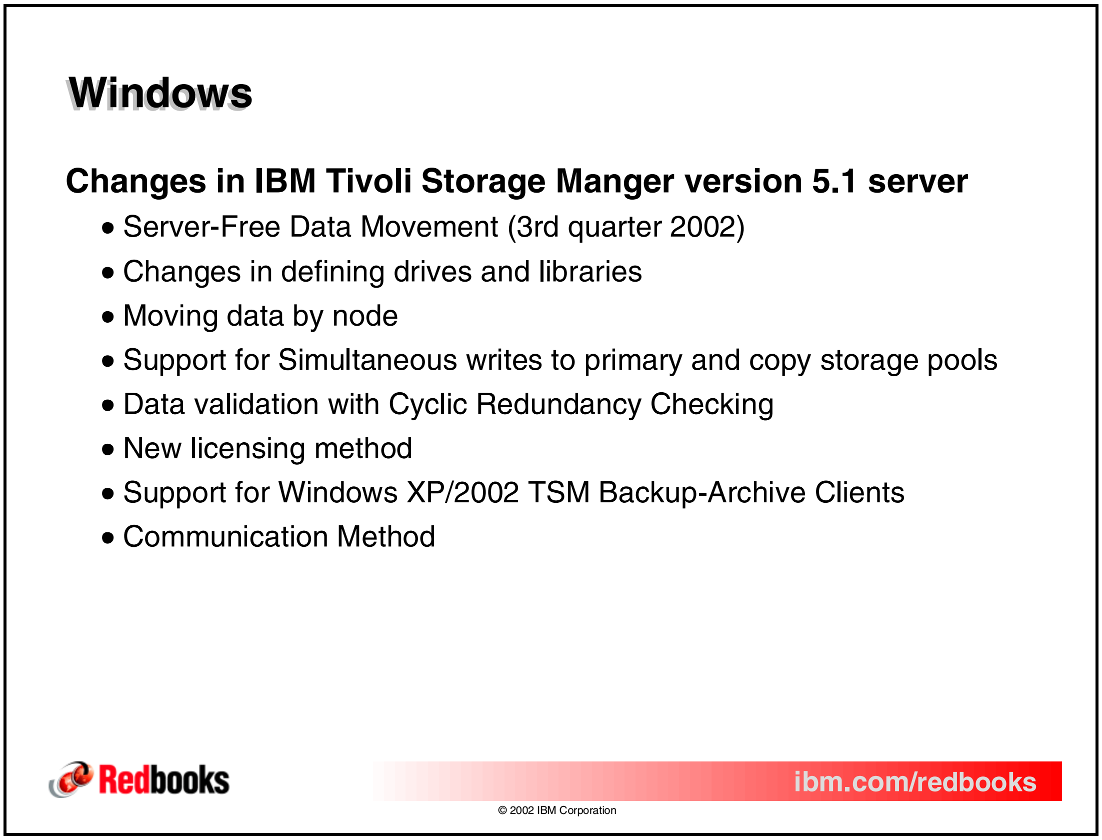

Figure 3-19 Changes for Microsoft Windows

- Server-Free Data Movement

  Server-free data movement uses the SCSI-3 extended copy command to do full-volume backup and restore. It provides the following advantages:

  – Reduces Tivoli Storage Manager client and server CPU utilization.
  – Eliminates almost all data movement on the LAN.
  – Improves performance for data movement operations.
  – Allows the server to handle more concurrent client connections and server operations.

- Changes in Defining Drives and Libraries

  Device special file names and external library managers are now specified in the DEFINE PATH and UPDATE PATH commands, rather than in the DEFINE DRIVE, UPDATE DRIVE, DEFINE LIBRARY, and UPDATE LIBRARY commands.

- Moving Data by Node

  You can use the MOVE NODEDATA command to move data in a sequential-access storage pool for one or more nodes, or move data in a single node with selected file spaces. You can also use MOVE NODEDATA to consolidate data for off-site disaster recovery storage, or move data to another storage pool.

- Support for Simultaneous Writes to Primary and Copy Storage Pools

  You can specify copy storage pools in a primary storage pool definition. When a client backs up, archives, or migrates a file, the file is written to the primary storage pool and is simultaneously stored into each copy storage pool.

- Data Validation with Cyclic Redundancy Checking

  Tivoli Storage Manager provides the option of specifying whether a cyclic redundancy check (CRC) is performed during a client session with the server, or for storage pools. The server validates the data by using a cyclic redundancy check which can help identify data corruption. Data validation can be enabled for one or all of the following:

  – Tivoli Storage Manager client nodes at Version 5.1.
  – Tivoli Storage Manager storage agents at Version 5.1.
  – Storage pools.

- New Licensing Method

  The new licensing method enables you to register the exact number of licenses that are required, rather than in increments of 1, 5, 10, and 50.

- Support for Windows XP/2002 Tivoli Storage Manager Backup-Archive Clients

  Support is now added for Windows XP and Windows 2002 Tivoli Storage Manager backup-archive clients that run in 64-bit mode. These two Windows platforms are a Unicode-only version of the client.

- Communication Method

  Tivoli Storage Manager for Windows no longer supports the Internetwork Packet Exchange (IPX)/SPX and NETBIOS communication methods. The COMMETHOD server option is updated to reflect this change.

3.5.2 Installation/migration

Before installing/migrating your Tivoli Storage Manager server, ensure that the new server provides support for your current storage devices. Refer to the Device Support section of the Tivoli Storage Manager Web site at:
`http://www.tivoli.com/support/storage_mgr/tivolimain.html`.

The following should show you, how to migrate a Tivoli Storage Manager Server to Version 5.1.

- TSM 4.2.x.x —> TSM 5.1

  – Install the new version and accept the default path, which will point to the existing server location. The setup program will automatically upgrade your Tivoli Storage Manager database and administrative Web interface.

  – Install the Tivoli Storage Manager device driver and the Tivoli Storage Manager client.

- TSM 4.1.x.x —> TSM 5.1

  – Write down the directory path of your current Tivoli Storage Manager server.

  – Use the Add/Remove Programs dialog in the Windows control panel to uninstall the current Tivoli Storage Manager server, Tivoli Storage Manager device driver and Tivoli Storage Manager client (if installed).

  – Install the new Tivoli Storage Manager server to the same location as the original. The setup program will automatically upgrade your Tivoli Storage Manager database and administrative Web interface.

  – Install the Tivoli Storage Manager device driver and the Tivoli Storage Manager client.

- TSM 3.7.x.x or ADSM 3.1.x.x —> TSM 5.1

  – Write down the directory path of your current Tivoli Storage Manager server.

  – Use the Add/Remove Programs dialog in the Windows control panel to uninstall the current Tivoli Storage Manager server and Tivoli Storage Manager client.

– Install the new Tivoli Storage Manager server to the same location as the original. After you complete the installation, the Initial Configuration Task List dialog will appear. Close this dialog.
– Update and run the script tsmfixup.cmd, located in the console directory. (Refer to the script header for update instructions.) The script will automatically upgrade your Tivoli Storage Manager database and administrative Web interface.
– Install the Tivoli Storage Manager device driver and the Tivoli Storage Manager client.

Note: Uninstalling the older version of the Tivoli Storage Manager server with the Add/Remove Programs dialog does not remove the servers Database, Recovery Log, or any Storage Pools. Thus all of your data already backed up or archived is still kept and any definitions you have made remain intact. Only the old Tivoli Storage Manager code is removed.

If you are migrating from ADSM and have created a disaster recovery plan file using Disaster Recovery Manager (DRM), be aware that Tivoli Storage Manager does not use the same default installation directories as ADSM. Disaster recovery installation path references may no longer be valid.

After you have migrated to Tivoli Storage Manager, you should back up your storage pools and database and create a new disaster recovery plan file. For the sequence and details of the procedure, refer to the Administrator's Guide.

The look and feel of the Tivoli Storage Manager main installation screen has changed, as shown in Figure 3-20.

Figure 3-20 Main installation window of Tivoli Storage Manager

At the main installation window choose **Install Products** to go to the Install Products screen of Tivoli Storage Manager Server.

Figure 3-21 Install products window of Tivoli Storage Manager

- **Tivoli Storage Manager Server**

  Server code includes the Tivoli Storage Manager database and tools to help you configure and manage Tivoli Storage Manager.

- **Tivoli Storage Manager Server Licenses**

  The license package includes support for all Tivoli Storage Manager licensed features. After you install the license package, you must apply the licenses you have purchased.

- **Tivoli Storage Manager Backup-archive Client**

  Backup-archive client code transfers data to, and receives data from, the Tivoli Storage Manager server component. It is recommended that you install the Tivoli Storage Manager Backup-archive client on the Tivoli Storage Manager server machine to help you validate your server configuration.

- **Tivoli Storage Manager Device Driver**

  The Tivoli Storage Manager device driver extends Tivoli Storage Manager media management capability.

- **Tivoli Storage Manager ODBC Client**

  The ODBC client can be used by applications that want to query the Tivoli Storage Manager server.

> **Note:** After each step a dialog prompting you to restart your computer. So if you plan to install any other Tivoli Storage Manager components at the same time, return to the Install Products screen to complete those installations before restarting your computer.

For further informations about the installation of the Tivoli Storage Manager Server, Tivoli Storage Manager Licenses, Tivoli Storage Manager Backup-archive Client and the Tivoli Storage Manager Device Driver, refer to *Tivoli Storage Manager for Windows Quick Start,* GC32-0784.

3.5.3 Configuration

The Tivoli Storage Manager for Windows server provides a graphical interface called the Tivoli Storage Manager Management Console, which includes a set of wizards that help you configure and manage your Tivoli Storage Manager system. One or more of these wizards is presented each time you add a new Tivoli Storage Manager server instance. You can choose from two wizard-based configuration paths.

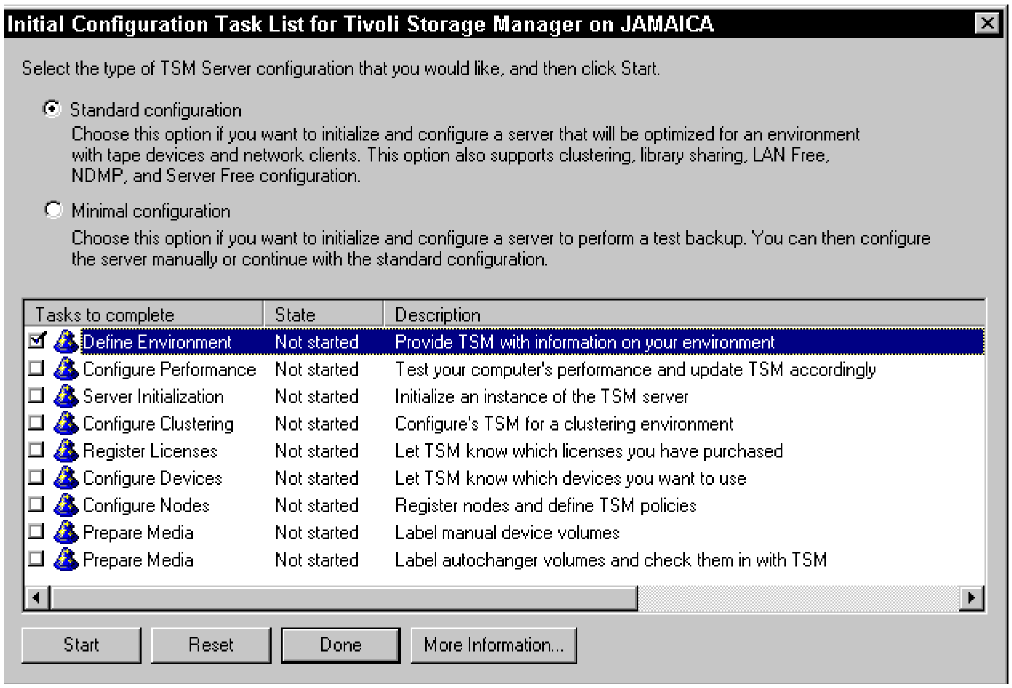

Figure 3-22 Initial configuration window for standard configuration

Choose this option to initialize and configure a server that can be set up to use storage devices and network clients. A series of wizards is presented in sequence to guide you through the initial configuration process. This is the recommended configuration path for setting up a functional production environment.

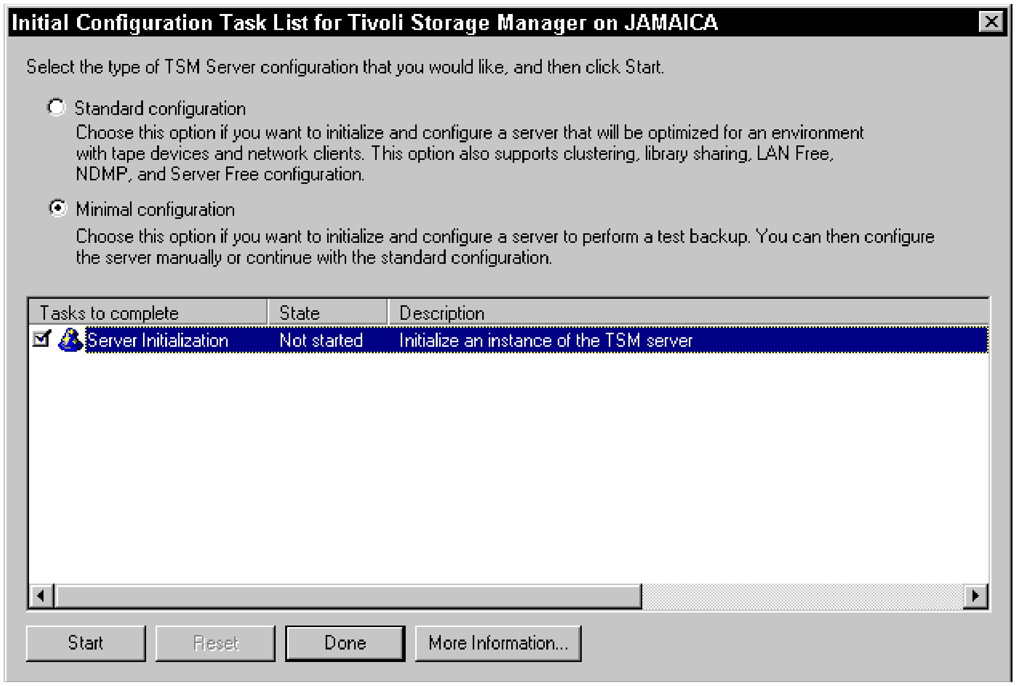

Figure 3-23 Initial configuration for minimal configuration

The Minimal configuration option is new with Tivoli Storage Manager V5.1. You can choose this option to quickly initialize a Tivoli Storage Manager server instance and perform a test backup of data located on the Tivoli Storage Manager server machine. This configuration allows you to quickly evaluate basic function, but does not support attached storage devices or network clients.

Note: While all Tivoli Storage Manager configuration and management tasks can also be performed using the command-line interface, the wizards are the preferred method for initial configuration. You can return to individual wizards after the initial configuration to update settings and perform management tasks.

For further informations about the use of the wizards, refer to *Tivoli Storage Manager for Windows Quick Start,* GC32-0784.

3.5.4 Server-Free Data Movement feature

Please note that Server-Free Data Movement will not be available in the first release of Tivoli Storage Manager V5.1 which has a GA of April 2002. We expect it to be delivered via a PTF to be made available by the 3rd quarter of 2002.

Overview

IBM Tivoli Storage Manager Version 5 Release 1 introduces Server-Free Data Movement. This is one of the major enhancements provided in Tivoli Storage Manager V5.1. The feature is currently only available with a Windows NT or 2000 server and a Windows 2000 client. We expect it to become available on more platforms later. Windows XP is likely to be the next platform.

Before we discuss Server-Free Data Movement it is appropriate to remind ourselves how we have evolved from traditional LAN based data movement through LAN-free to server-free.

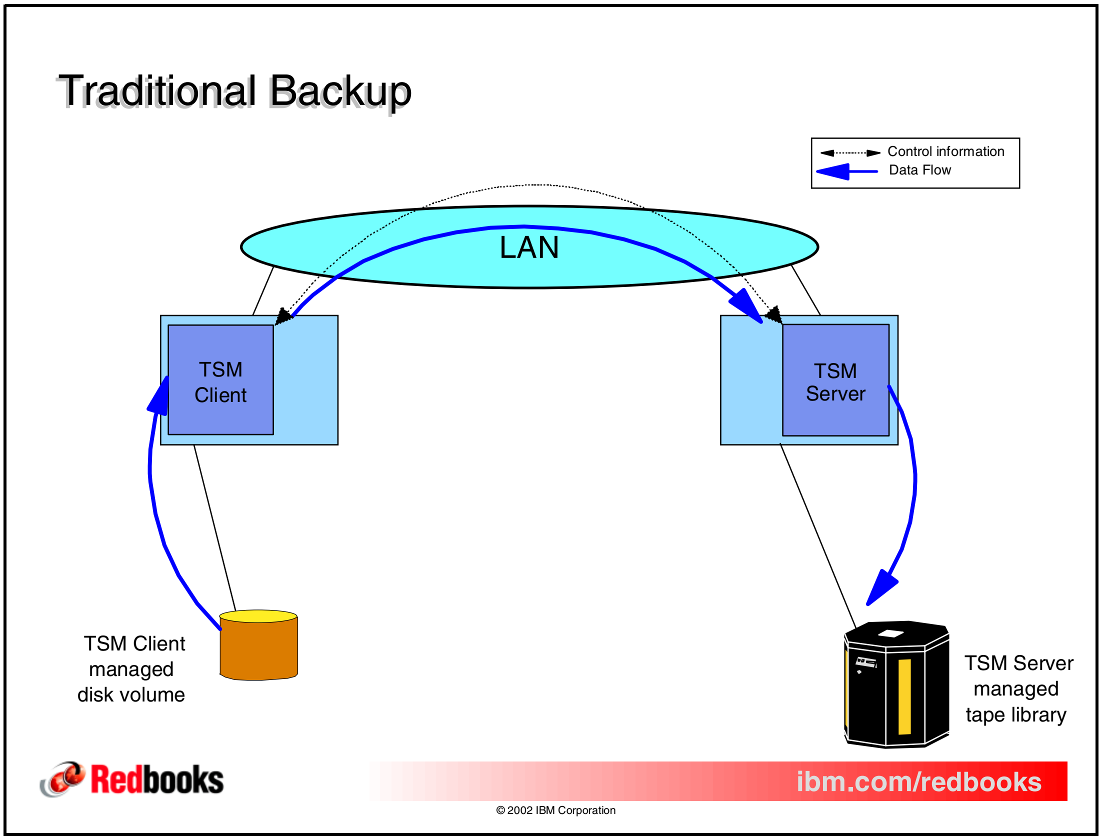

Figure 3-24 Traditional backup

Originally the Tivoli Storage Manager client would read data from locally attached disks and send it over the LAN to the Tivoli Storage Manager server. The server would receive the data then write it to the storage pool hierarchy. Notice that the data is read and written by both the Tivoli Storage Manager client and Tivoli Storage Manager server machines. In addition the control information (the metadata) is also sent over the LAN to the Tivoli Storage Manager server.

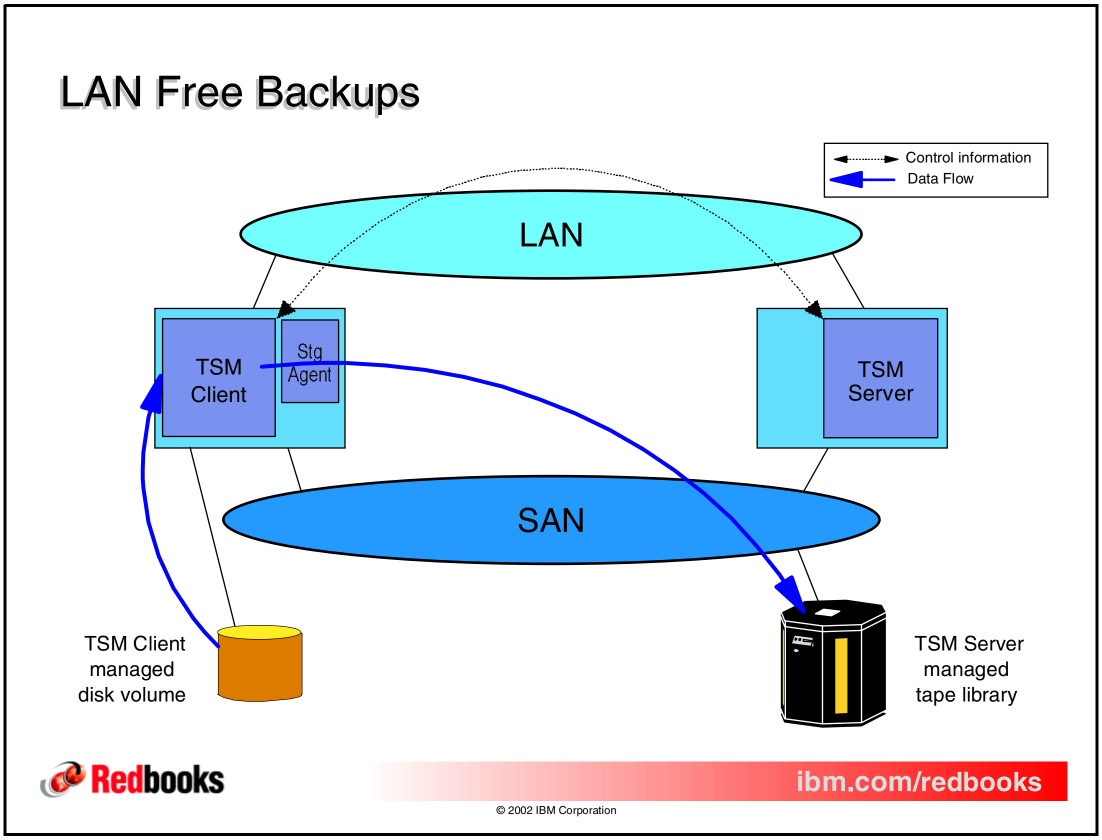

Figure 3-25 LAN-free backups

LAN-free backups introduced in Tivoli Storage Manager V3.7 relieve the load on the LAN by introducing the Storage Agent. This is a small Tivoli Storage Manager server (without a Database or Recovery Log) which is installed and run on the Tivoli Storage Manager client machine. It handles the communication with the Tivoli Storage Manager server over the LAN but sent the data directly to SAN attached tape devices, relieving the Tivoli Storage Manager server from the actual I/O transfer.

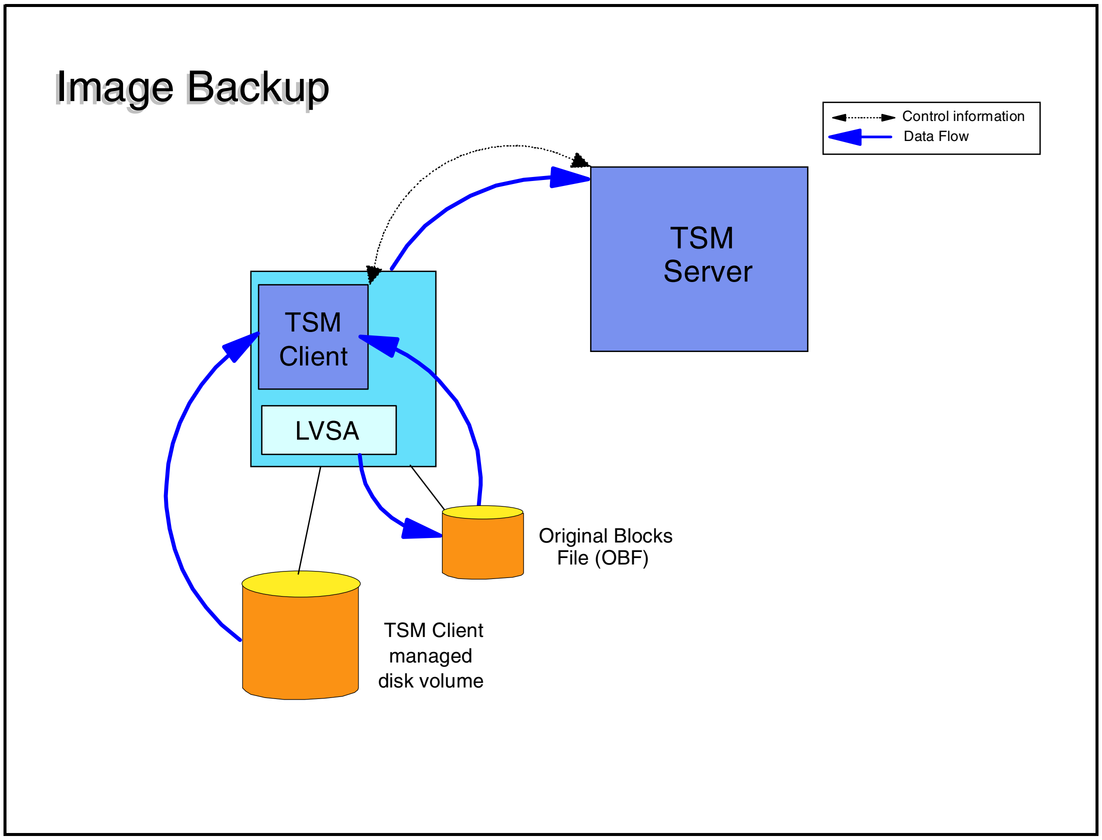

Figure 3-26 Image backup

Image Backup, which was introduced with Tivoli Storage Manager V3.7 for AIX, HP and Sun client platforms, is now also available on Windows 2000. On the Windows 2000 client platform the Logical Volume Storage Agent (LVSA) has been introduced which is capable of taking a snapshot of the volume while it is online. See Chapter 5, “Platform specific client enhancements” on page 179 for more details. This image backup is a block by block copy of the data. Optionally only occupied blocks can be copied. If the snapshot option is used (rather than static) then any blocks which change during the backup process are first kept unaltered in an Original Block File. In this way the client is able to send a consistent image of the volume as it was at the start of the snapshot process to the Tivoli Storage Manager server.

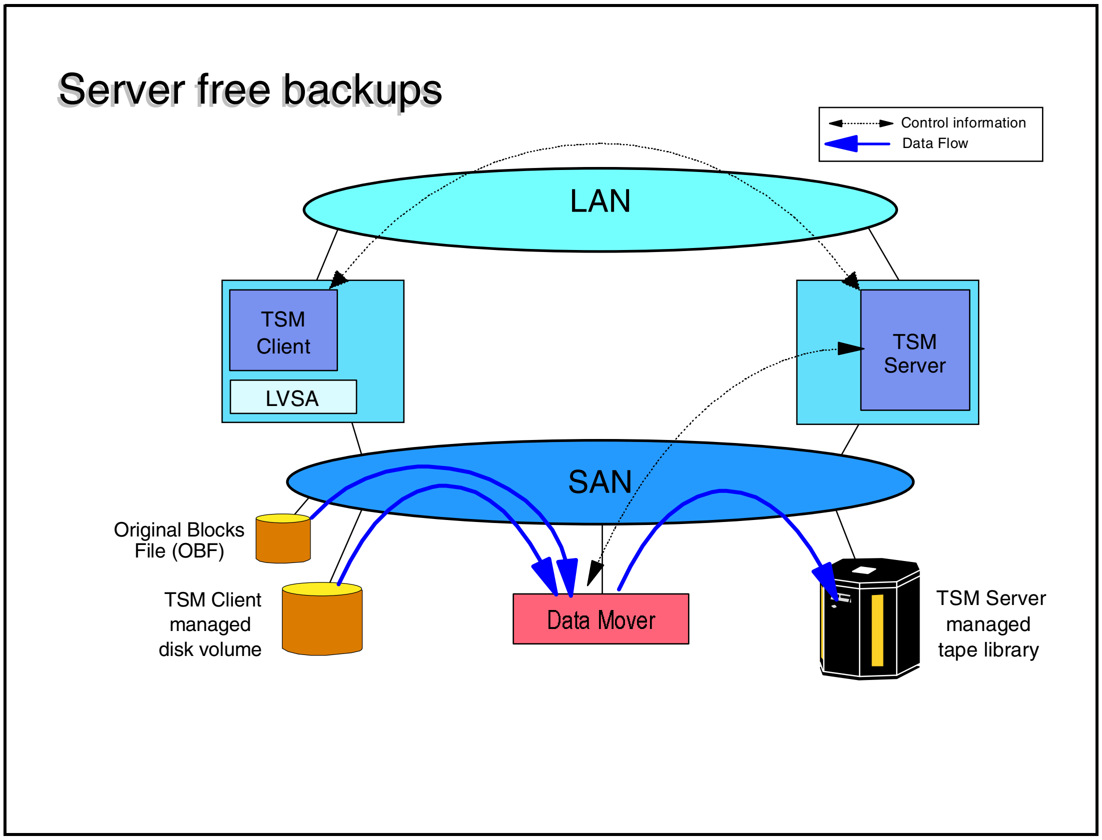

Figure 3-27 Server-free backups

Server-Free Data Movement is just a special form of standard Windows 2000 image backup. The Storage Agent used in LAN-free backups is no longer utilized. The data movement is actually done by a SAN Data Gateway or similar device on the SAN. If snapshot is being used then an OBF file is created to hold the original blocks if any are changed during the backup. This OBF is sent to the server with the image which contains any changes made during the backup. On restore it is the client's job to merge the restored image with the blocks in the OBF to produce the volume as it was at the start of the snapshot backup. Thus both Tivoli Storage Manager client and Tivoli Storage Manager server machines do not have to read and write the data at all. The Tivoli Storage Manager server sends commands to the SDG device to tell it which blocks to move from which SAN attached disk to which SAN attached tape device.

Important: It is worth pointing out that although many references are made to a Data Mover or datamover device in server-free backups that actually no data movement takes place. The data is actually copied rather then moved from one location to another. Once a backup has completed using Server-Free Data Movement for example, the original data is still there on the client disk!

This provides a way to back up and restore large volumes of data between client-owned disks and storage devices using a method that considerably reduces overhead on the Tivoli Storage Manager server and the Tivoli Storage Manager client. Remember that only volume images and not individual files can be moved by Server-Free Data Movement. It minimizes data transfer on the LAN. The data is transferred block by block rather than that doing file I/O. Both Raw and NTFS volumes can be backed up using Server-Free Data Movement, but not FAT volumes.

Data that has been backed up using Server-Free Data Movement can be restored over a server-free path, or over a LAN-free path or indeed over the LAN itself.

All data movement is out board of the client machine. In this way data is copied directly from the SAN attached Tivoli Storage Manager client disk to the SAN attached tape drive via the SAN Data Gateway data mover. No Storage Agent is involved in the data movement as in LAN-free backups. It's important to realize that it is not Tivoli Storage Manager server-free. The Tivoli Storage Manager server discovers devices in the SAN, ensures all the required paths exist, transforms extents to xcopy commands, mounts and positions the tape, records the information about the backup/restore in its database, etc. The metadata is written via the traditional LAN path.

Notice the client data is written directly from SAN attached disk to SAN attached tape via the SDG data mover. Metadata still goes to the Tivoli Storage Manager server via the LAN path. However, the client machine (application server) no longer has to read the data from disk and write it to the SAN attached tape as done previously on LAN-free backups. In this way the data movement is **application** server-free.

The impact on application servers is now minimized with Server-Free Data Movement. It reduces both Tivoli Storage Manager client and server CPU utilization. The use of a SCSI-3 extended copy command causes data to be transferred directly between devices over the SAN or SCSI bus.

The SCSI-3 command is issued and initiated by the Tivoli Storage Manager server and carried out by the data mover device that exists on the SAN. The Tivoli Storage Manager client builds the extent lists and sends them to the Tivoli Storage Manager server which controls the conversation with the data mover. The data mover device is responsible for moving the data either from a SAN attached client-owned disk to a SAN attached tape drive (backup), or from a SAN attached tape drive to a SAN attached client-owned disk (restore).The data mover device must support the SCSI-3 EXTENDED COPY command, which conforms to the ANSI T10 SPC-2 standard. For a current list of supported data mover devices, visit:

http://www.tivoli.com/support/storage_mgr/tivolimain.html

The data mover device can be anywhere in the SAN, but it has to be able to address the LUNs for both the disk and tape devices it is moving data between. It has to be able to see the LUNs in the same way that the Tivoli Storage Manager server and client see them. The SAN attached disk must be able to respond to a query for its serial number or WWN so that the Tivoli Storage Manager server can verify it is sending data to the correct device. Both SCSI and fibre connectivity work.

The Tivoli Storage Manager server will be able to handle more concurrent client connections because it will not be copying data. Similarly the Tivoli Storage Manager client will be able to handle more application load because cycles will not be needed to send any data buffers to the server.

Online server-free image backup is currently supported with a Windows NT/2000 server and Windows 2000 clients. By default it uses Tivoli Storage Manager's Logical Volume Storage Agent, which is installed via a Tivoli Storage Manager GUI wizard, and a SAN Data Gateway for the data movement.

Related Options

You enable the server to perform server-free data movement operations by using the SET SERVERFREE command:

```
set serverfree status=on
```

You can set the amount of data to be moved during an instance of an extended copy command by specifying the BATCHSIZE parameter when you issue the SET SERVERFREE command. The BATCHSIZE parameter specifies the maximum amount of data, in megabytes, that will be copied in a single instance of an extended copy command. The minimum value is 10 megabytes, the maximum value is 1000 megabytes, and the default value is 80 megabytes. For example, if you wanted to set a batchsize of 320 megabytes, you would issue the following:

```
set serverfree batchsize=320
```

You can specify the number of concurrent copy operations that the data mover can support by including the COPYTHREADS parameter when you issue the DEFINE DATAMOVER command. The COPYTHREADS parameter is used to govern how many copy commands can run concurrently using this same data mover. For example, if you wanted to be able to perform 4 concurrent server-free data movement operations, you would issue the following when defining a data mover:

```
define datamover scsi1 type=scsi wwn=2002006045160d2a serial=21081300957
dataformat=nonblock copythreads=4
```

The default value is 1. Please refer to the documentation for your device to obtain the maximum value. The number must be greater than 0. If other programs or servers will be using the extended copy capabilities of the data mover device, you may want to specify fewer than the maximum number of concurrent copy commands supported by the device for this parameter.

Data Format for Server-Free Data Movement

Data that is transferred using server-free data movement must be stored in a different format than the format that has traditionally been used for Tivoli Storage Manager. The new data format is called NONBLOCK and must be specified when you define a data mover for server-free data movement, and when you define a storage pool that is used for server-free data movement.

Attention: The terms NONBLOCK (or NONblock) and ‘Native Without’ mean the same thing. If you look at a Storage Pool which has been defined to hold server-free backups with a Web Administrator GUI you will see NONblock in the Storage Pool Data Format field. However, if you display a SCSI Data Mover device (such as the SAN Data Gateway) then the Storage Pool Data Format there will show ‘Native without’.

Data that has been backed up or archived using server-free data movement, and stored in a storage pool with a dataformat of NONBLOCK, can be moved to other storage pools. Likewise, you can also move an uncompressed image object backed up or archived by a client from a traditional storage pool to a NONBLOCK storage pool and restore it using Server-Free.

Server operations such as migration, move data, and storage pool backup can cause this data to be stored in storage pools that do not have a data format of NONBLOCK. If this occurs, it is no longer possible to restore the data using Server-Free. However, the data can still be restored over the LAN.

Requirements

You must meet the following requirements when setting up Server-Free Data Movement operations:

- Operating System for the Tivoli Storage Manager Server and Client

Windows NT 4.0 with Service Pack 6 or later, or Windows 2000 with Service Pack 2 or later are required for the Tivoli Storage Manager server.

Windows 2000 with Service Pack 2 or later is required for the Tivoli Storage Manager client.

- Data Mover Device

  The data mover device must support the version of the SCSI-3 EXTENDED COPY command, which conforms to the ANSI T10 SPC-2 standard.

- Fibre Channel Host Bus Adapter (FC HBA)

  You will need to install a Host Bus Adapter DLL that supports the Storage Network Industry Association's (SNIA) common API for purposes of device discovery. Consult the HBA documentation, or the HBA's manufacturer for the location of the required DLL (for QLogic HBAs refer to `http://www.qlogic.com/support/home_support.asp`).

> **Attention:** Device addressing in a SAN can be tricky. Especially in a Windows environment where addition or removal of devices without disruption can cause a change to the SCSI address or device name (//./Tape0) of existing devices. If possible use persistent naming (or binding) with a static device naming convention. For more details, see the Redpaper *Managing the device addressing of SAN attached tape for use with Tivoli Storage Manager*, REDP0150.

- Tape Libraries

  SCSI-attached or Fibre Channel-attached tape libraries that are supported by the Tivoli Storage Manager device driver.

- Tape Drives

  If a tape library contains one or more drives, the data mover must have addressability to all drives that will be used for server-free data movement.

- Disks

  The client-owned disks must be accessible from the data mover and the Tivoli Storage Manager server machine's HBA.

Verify the compatibility of specific combinations of data mover devices, tape devices, and SAN devices with the manufacturers of the hardware. Not all data movers support all drives. Before attempting to set up Server-Free Date Movement operations, first verify that Tivoli Storage Manager supports the devices you wish to use.

For more information about supported Data Mover Devices, FC HBAs, Tape Libraries, Tape Drives and Disks, visit `http://www.tivoli.com/support/storage_mgr/tivolimain.html`.

Preparation for Server-Free Data Movement

Careful preparation is everything, we found. Make sure you give yourself plenty of time to setup Server-Free Data Movement.

In our lab setup we used a SAN Data Gateway 2108-G07 as data mover. The Microcode of the SAN Data Gateway has to be at least 3.43.13.

First we needed to install current microcode level on the 2108-G07 SDG itself. After that we enabled the SDG as a data mover. A license key is required to enable data movement. We have used the StorWatch SAN Data Gateway Specialist to license the SAN Data Gateway as a data mover.

IBM SAN Data Gateway

The 2108-G07 SDG needs to be at current firmware levels in order to be able to perform Server-Free Data Movement as directed by the Tivoli Storage Manager server.

The following procedure describes how to update the firmware of the SAN Data Gateway Model 2108-G07, enable the Data Mover and setup the Gateway for Server-Free Data Movement. For downloading the latest version of the firmware use `www.storage.ibm.com/hardsoft/products/sangateway/support/cdr/Firmware/sdgfw.htm.`

1. If you have not already done so, install the IBM StorWatch SAN Data Gateway Specialist. For a description how to install and how to use the IBM StorWatch SAN Data Gateway Specialist please refer to the *IBM Storage Area Network Data Gateway Installation and User's Guide,* SC26-7304. Both StorWatch server and client code need to be installed on your machine. The client automatically connects to the server and then the server connects to the SAN Data Gateway of your choice (by IP address).
2. Start the IBM SAN Data Gateway Specialist and connect to the SAN Data Gateway, you want to update, by using the IP-address of the Gateway. The following window appears:

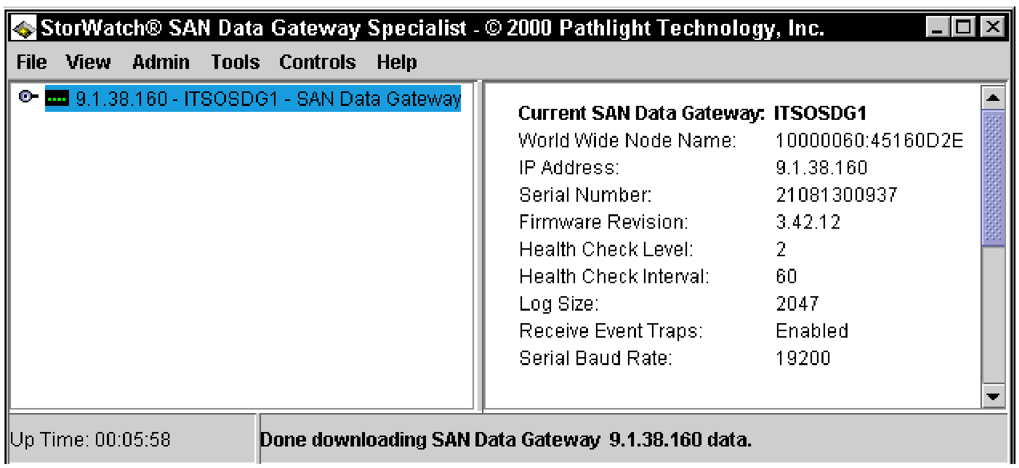

Figure 3-28 Specialist main window

3. In the right part of the window you see the current firmware revision (e.g. 3.42.12). Click **Controls —> Update Firmware** to get the file selection window. Choose the new firmware file and press **Open**.

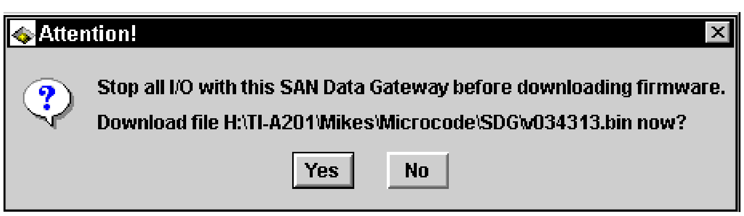

Figure 3-29 Attention! message

4. Downloading the firmware into the SAN Data Gateway is a disruptive process, so the Specialist displays a warning message. Click **Yes** for starting download.

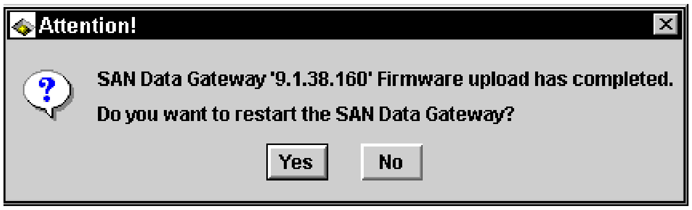

Figure 3-30 Upload completed message

5. Once all I/O activity is stopped, you can click **Yes** to restart the SAN Data Gateway.

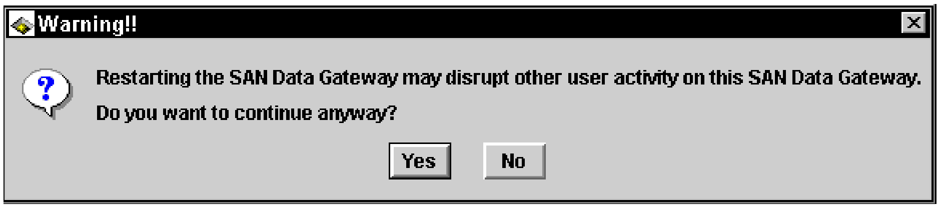

Figure 3-31 Warning! message

6. Click **Yes** to continue anyway. After reconnecting to the SAN Data Gateway you can enable the Data Mover. Therefore click **Controls —> Feature Enable —> Data Mover**. The following window appears:

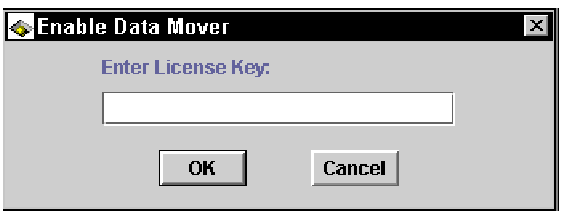

Figure 3-32 Enable Data Mover window

7. Enter your license key to activate the Data Mover.
8. After updating the firmware and activating the Data Mover you should check if the setup of the SAN Data Gateway for Server-Free Data Movement is correct. Therefore select the Fibre Channel port which is connected to the SAN and click **Controls —> Fibre Channel** as shown in Figure 3-33.

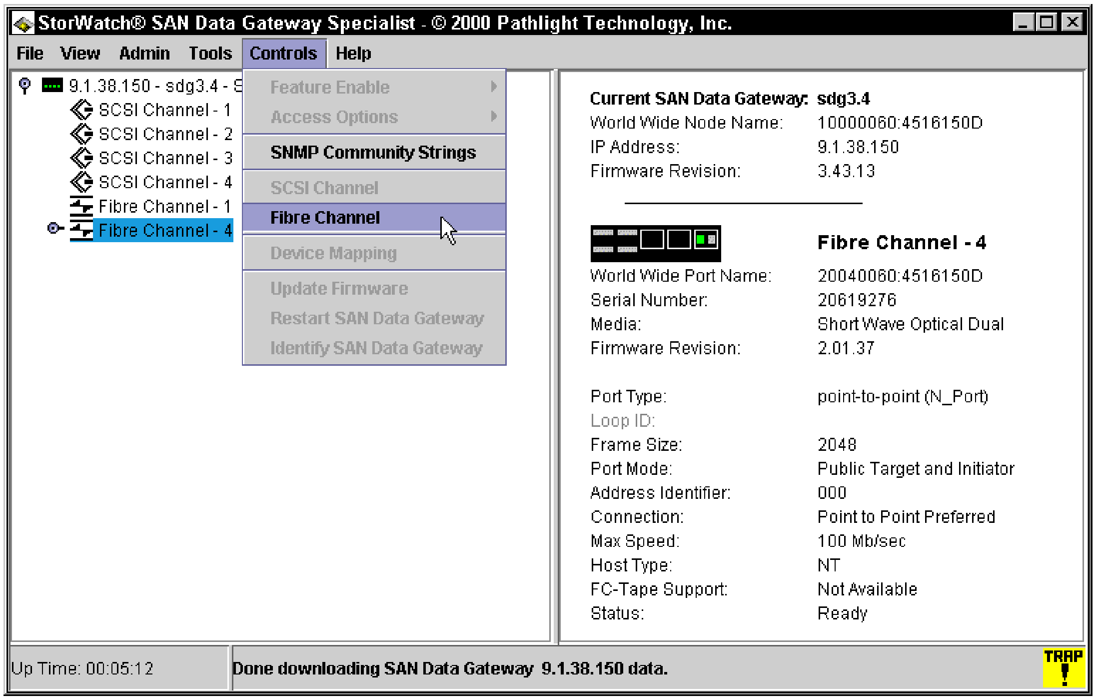

Figure 3-33 Fibre Channel selection

9. The following window appears.

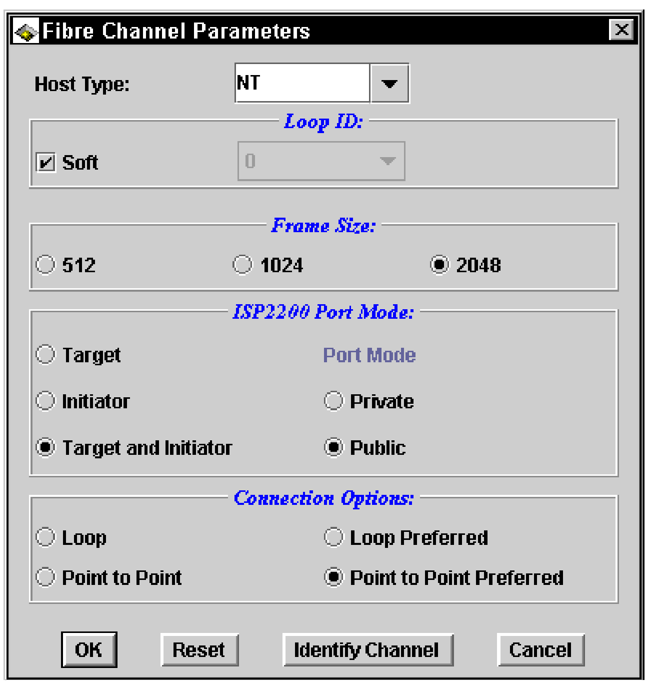

Figure 3-34 Fibre Channel parameters window

10. Select **NT** as Host Type to be sure the Tivoli Storage Manager server can see the Gateway via the SAN. Select **Target and Initiator** and **Public** Port Mode to be sure the Gateway can see the disk subsystem as well as the tape drives. For Connection Options use **Loop** or **Point to Point** dependent on if you use a FC Hub or a FC Switch. Click **OK** to apply all changes.

Note: The port settings **Public** and **Target and Initiator** can only be selected if the FC ports of the SDG uses 2200-type HBAs. With 2100-type HBAs it will not work.

11. To check if the Gateway is correctly connected to the SAN and sees all necessary devices, expand the Fibre Channel section of the Channel you use for your SAN connection as shown in Figure 3-35.

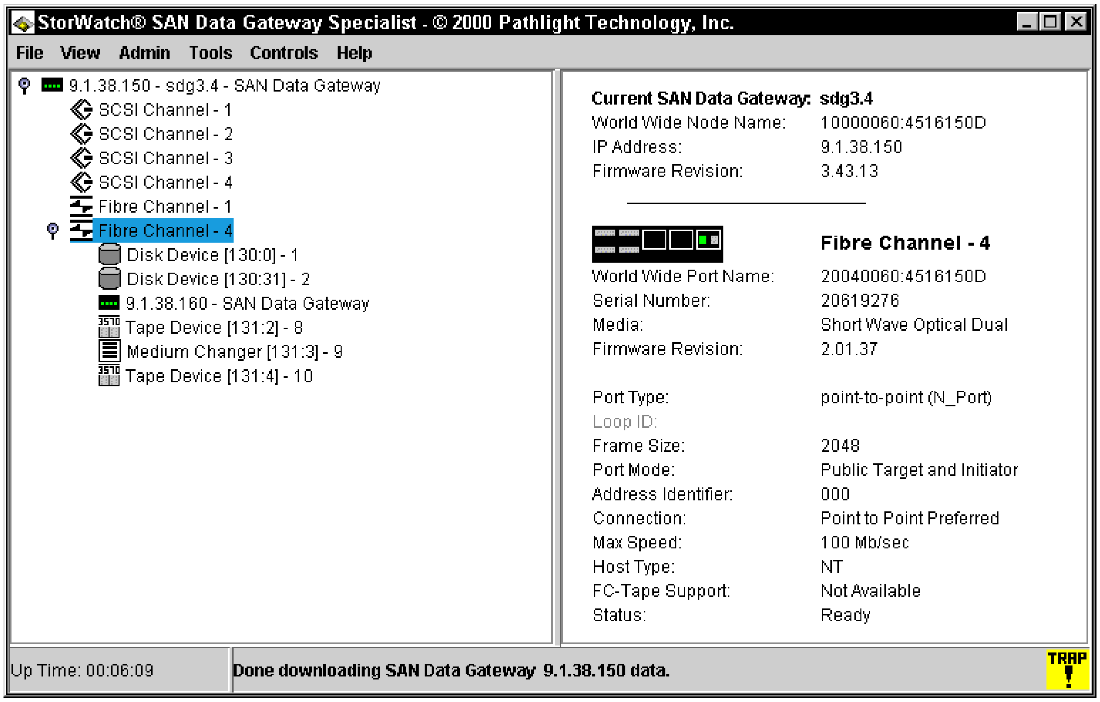

Figure 3-35 Visible devices for Fibre Channel - 4

12. If your setup is correct you will see the disk devices, the SAN Data Gateway of the tape library as well as the medium changer and his tape devices.

Environment

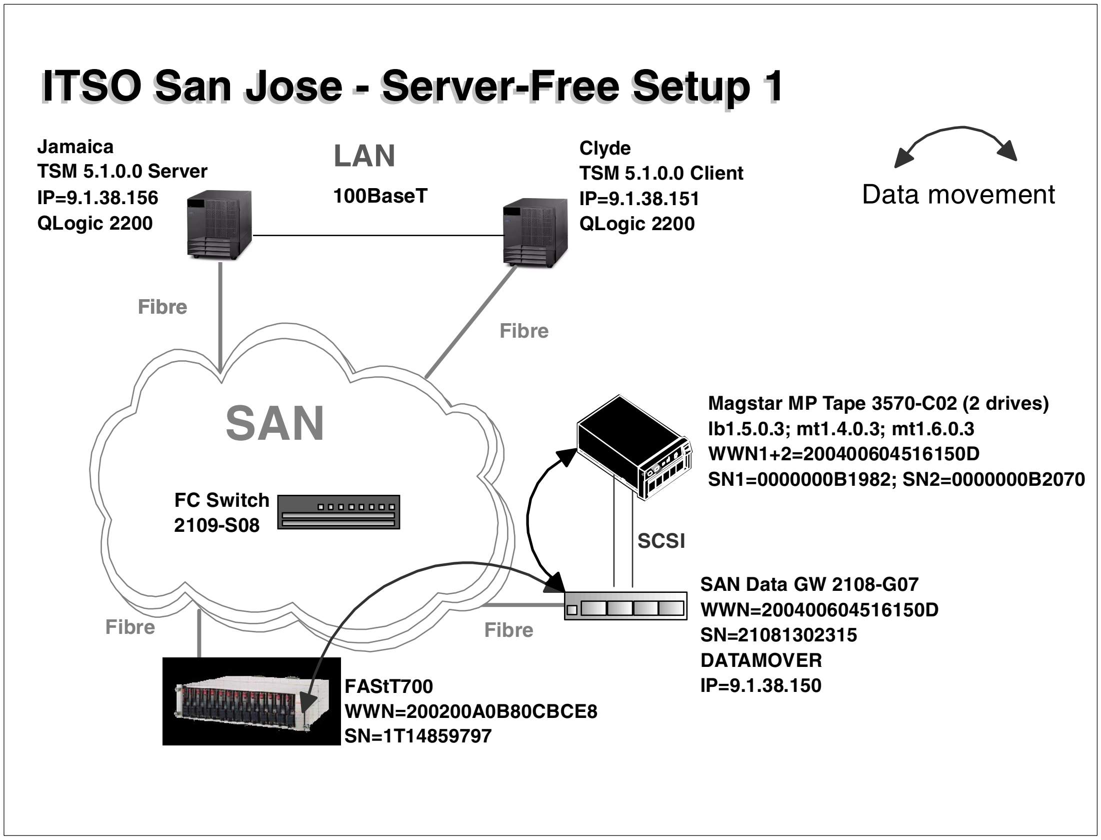

Figure 3-36 ITSO San Jose - Server-Free setup 1

Here's how our setup in the lab looked for our original test of Server-Free Data Movement. We have indicated the World Wide Names and serial numbers of the devices which were needed during the Path and Device definitions we used. Notice than in this test the tape library was directly attached to the SCSI port on the SAN Data Gateway data mover device itself.

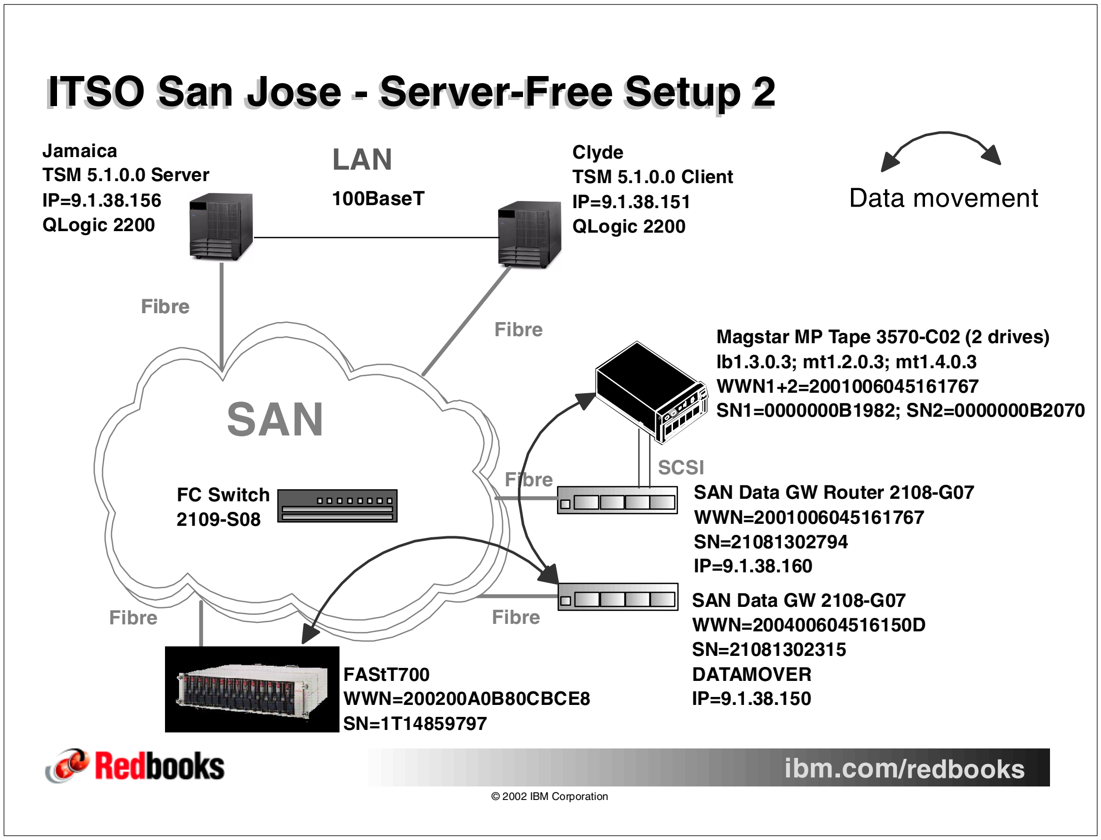

Figure 3-37 ITSO San Jose - Server-Free setup 2

We also tested Server-Free Data Movement in our lab with the tape library attached to a second SAN Data Gateway device which was **not** enabled as a data mover. The diagram above shows how we had a Tivoli Storage Manager server (Jamaica) and Tivoli Storage Manager client (Clyde) both at Tivoli Storage Manager V5.1 and both running on Windows 2000 systems. They were connected to a SAN as shown with a FAStT700 disk subsystem and a 2108-G07 SAN Data Gateway which had been enabled as our data mover. The second 2108-G07 SAN Data Gateway was required only because our Magstar library had only got SCSI connectivity. We have shown the World Wide Names and serial numbers of all the devices we needed during the preparations for Server-Free Data Movement.

Tivoli Storage Manager Preparation

For a better understanding of the following steps it will help to know some of the definitions we used for our Server-Free setup:

Server name = JAMAICA
Node name = CLYDE
Library name = MAGLIB
Tape drive 1 = MAGDRV1
Tape drive 2 = MAGDRV2
Device class = MAGCLASS
Storage pool = SRVFREEPOOL
Management class = SRVFREEMC
Data Mover = SRVFREEDM

Configuring Tivoli Storage Manager for Server-Free Data Movement

Here is a step-by-step description of our Server-Free Setup 2:

1. Connect Devices to the SAN

   Before you begin configuring Tivoli Storage Manager for server-free data movement, you must connect all devices to the SAN that are needed for this function, including:

   – Data mover device
   – Tape libraries
   – Tape drives
   – Disk devices

   The devices must exist on the SAN and must be visible for the Tivoli Storage Manager Server prior to defining them to Tivoli Storage Manager. If you are using zoning, be sure the following devices can see each other:

   – Tivoli Storage Manager Server - Data mover device - Tape libraries - Tape drives - Disk devices

     And,

   – Tivoli Storage Manager Client - Disk devices

2. Determining World Wide Name, Serial Number and LUN

   Before you can configure Tivoli Storage Manager for Server-Free Data Movement, you must first identify and record the World Wide Name, serial number, and LUN of each disk device and tape drive used, and the World Wide Name and serial number of the data mover. This information will be used when defining devices and device paths. To do this, issue the QUERY SAN command. The QUERY SAN command will display the devices that can be detected on the SAN, as well as the information for those devices. To get an overview of all available devices on the SAN type **query san** at your Web admin command line and you will get output similar to the following:

Example 3-46 Output of 'query san' command

```
Operation Results

Device      Vendor      Product         Serial          Device
Type                                    Number
--------    --------    ------------    ------------    --------------
DISK        IBM         1742            1T14859797
DISK        IBM         1742            1T14859797
GATEWAY     PATHLGHT    SAN Gateway     21081302794
DRIVE       IBM         03570C12        0000000B1982    mt1.2.0.3
LIBRARY     IBM         03570C12        0000000B1982    lb1.3.0.3
DRIVE       IBM         03570C12        0000000B2070    mt1.4.0.3
GATEWAY     PATHLGHT    SAN Gateway     21081302315
```

   To get more information including the WWN type **query san format=detailed** at your Web admin command line and you will get the following kind of output:

Example 3-47 Output of 'query san format=detailed' command

```
Operation Results

  Device Type: DISK
       Vendor: IBM
      Product: 1742
Serial Number: 1T14859797
       Device:
   Data Mover: No
     Node WWN: 200200A0B80CBCE7
     Port WWN: 200200A0B80CBCE8
          LUN: 0
```

```
      SCSI Port: 3
       SCSI Bus: 0
    SCSI Target: 0

    Device Type: DISK
         Vendor: IBM
        Product: 1742
  Serial Number: 1T14859797
         Device:
     Data Mover: No
       Node WWN: 200200A0B80CBCE7
       Port WWN: 200200A0B80CBCE8
            LUN: 1
      SCSI Port: 3
       SCSI Bus: 0
    SCSI Target: 0

    Device Type: GATEWAY
         Vendor: PATHLGHT
        Product: SAN Gateway
  Serial Number: 21081302794
         Device:
     Data Mover: No
       Node WWN: 1000006045161767
       Port WWN: 2004006045161767
            LUN: 0
      SCSI Port: 3
       SCSI Bus: 0
    SCSI Target: 1

    Device Type: DRIVE
         Vendor: IBM
        Product: 03570C12
  Serial Number: 0000000B1982
         Device: mt1.2.0.3
     Data Mover: No
       Node WWN: 1000006045161767
       Port WWN: 2004006045161767
            LUN: 2
      SCSI Port: 3
       SCSI Bus: 0
    SCSI Target: 1

    Device Type: LIBRARY
         Vendor: IBM
        Product: 03570C12
  Serial Number: 0000000B1982
         Device: lb1.3.0.3
     Data Mover: No
       Node WWN: 1000006045161767
       Port WWN: 2004006045161767
            LUN: 3
      SCSI Port: 3
       SCSI Bus: 0
    SCSI Target: 1

    Device Type: DRIVE
         Vendor: IBM
        Product: 03570C12
  Serial Number: 0000000B2070
```

```
         Device: mt1.4.0.3
     Data Mover: No
       Node WWN: 1000006045161767
       Port WWN: 2004006045161767
            LUN: 4
      SCSI Port: 3
       SCSI Bus: 0
    SCSI Target: 1

    Device Type: GATEWAY
         Vendor: PATHLGHT
        Product: SAN Gateway
  Serial Number: 21081302315
         Device:
     Data Mover: Yes
       Node WWN: 100000604516150D
       Port WWN: 200400604516150D
            LUN: 0
      SCSI Port: 3
       SCSI Bus: 0
    SCSI Target: 2
```

When you record the information, be sure to use Port WWN as the World Wide Name to be used in the path definitions.

For example, the information we recorded for our first tape drive was:

– World Wide Name: 2004006045161767
– Serial number: 0000000B1982
– LUN: 2

You can also get all these information out of the Tivoli Storage Manager Management Console, see Figure 3-38 and Figure 3-39.

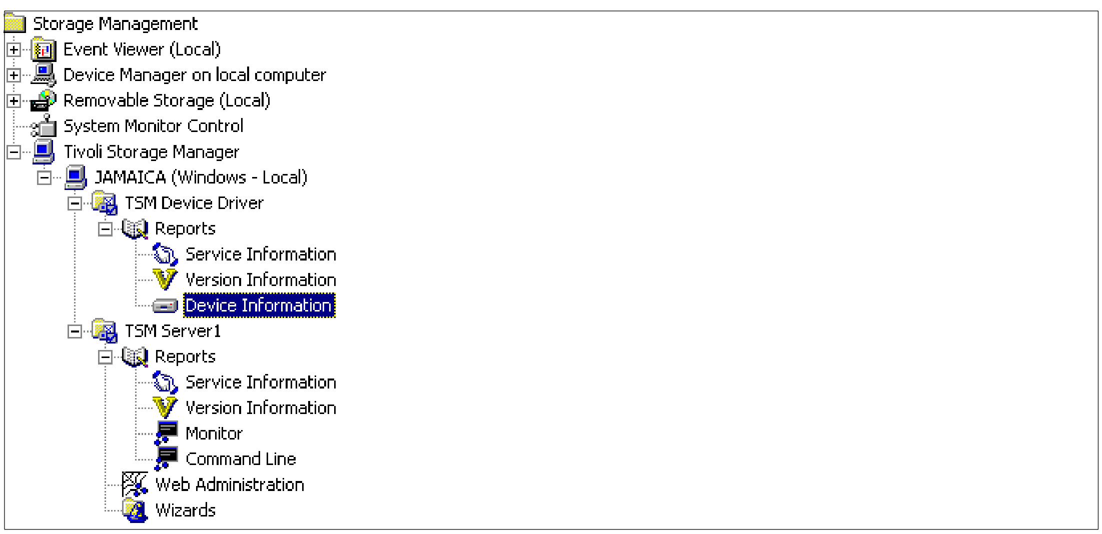

Figure 3-38 Device information

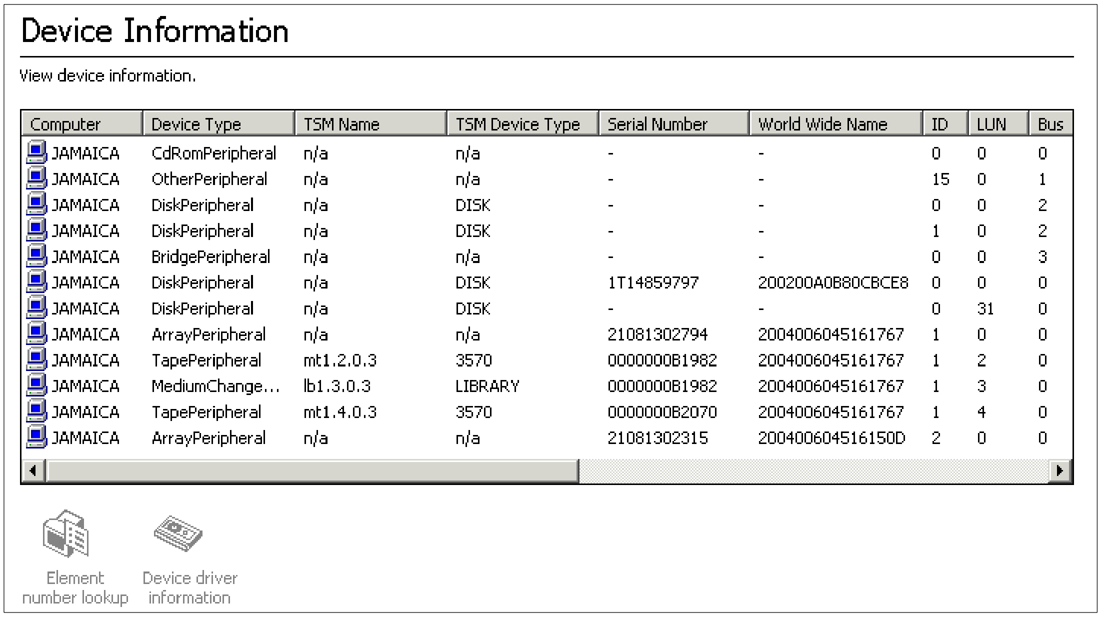

Computer	Device Type	TSM Name	TSM Device Type	Serial Number	World Wide Name	ID	LUN	Bus
JAMAICA	CdRomPeripheral	n/a	n/a	-	-	0	0	0
JAMAICA	OtherPeripheral	n/a	n/a	-	-	15	0	1
JAMAICA	DiskPeripheral	n/a	DISK	-	-	0	0	2
JAMAICA	DiskPeripheral	n/a	DISK	-	-	1	0	2
JAMAICA	BridgePeripheral	n/a	n/a	-	-	0	0	3
JAMAICA	DiskPeripheral	n/a	DISK	1T14859797	200200A0B80CBCE8	0	0	0
JAMAICA	DiskPeripheral	n/a	DISK	-	-	0	31	0
JAMAICA	ArrayPeripheral	n/a	n/a	21081302794	2004006045161767	1	0	0
JAMAICA	TapePeripheral	mt1.2.0.3	3570	0000000B1982	2004006045161767	1	2	0
JAMAICA	MediumChange...	lb1.3.0.3	LIBRARY	0000000B1982	2004006045161767	1	3	0
JAMAICA	TapePeripheral	mt1.4.0.3	3570	0000000B2070	2004006045161767	1	4	0
JAMAICA	ArrayPeripheral	n/a	n/a	21081302315	200400604516150D	2	0	0

Figure 3-39 Device information window

3. Obtaining and Correlating the Disk Information

   Once you have detected and recorded the information for all of the devices detectable on the SAN, you must obtain and correlate the information for the disks that you will be backing up.

   Log into the client machine that owns the disk and start the Tivoli Storage Manager client command line interface. Type **query diskinfo** to get the information for disks associated with the client as shown in Example 3-48.

Example 3-48 Disk information of client Clyde

```
Disk name    LUN    Serial number
----------------------------------
Harddisk0     0      TEF76061
Harddisk1     0      1T14859797
```

Note: The output of the QUERY DISKINFO command will display the information of every disk defined on the client, SAN attached as well as direct attached.

   Compare the information from the QUERY DISKINFO command output to the output of the QUERY SAN command. In our example, the serial number and LUN from the disk, Harddisk1, match the information for the disk displayed in the output from the QUERY SAN command, indicating that this is the disk to be backed up.

   The information for the disk that will be backed up was:

   – Node name: Clyde
   – Device name: Harddisk1
   – World wide name: 200200A0B80CBCE8
   – Serial number: 1T14859797
   – LUN: 0

4. Enable the server to perform Server-Free Data Movement operations using the SET SERVERFREE command:

```
set serverfree status=on
```

To check if the server supports Server-Free Data Movement and the Server-Free status is switched on type **query status** at your server command line. As you can see in the last two rows of the output in Figure 3-40, Server-Free is switched on and the Server-Free Batch Size is 80 MB.

```
                 Last License Audit: 03/18/2002 16:41:08
          Server License Compliance: Valid
                  Central Scheduler: Active
                   Maximum Sessions: 25
         Maximum Scheduled Sessions: 12
      Event Record Retention Period: 10 Day(s)
             Client Action Duration: 5 Day(s)
  Schedule Randomization Percentage: 50
              Query Schedule Period: 6 Hour(s)
            Maximum Command Retries: 10
                       Retry Period: 5 Minute(s)
                   Scheduling Modes: Any
                           Log Mode: Normal
            Database Backup Trigger: Not Defined
                        BufPoolSize: 131,072 K
                   Active Receivers: CONSOLE ACTLOG NTEVENTLOG
             Configuration manager?: Off
                   Refresh interval: 60
             Last refresh date/time:
                  Context Messaging: On
                 Server-free Status: On
             Server-free Batch Size: 80
```

Figure 3-40 Query status output

5. Define the library to the server by using the DEFINE LIBRARY command:

   `define library maglib libtype=scsi shared=no`

6. Define the library path from the Tivoli Storage Manager server to the library by using the DEFINE PATH command:

   `define path jamaica maglib srctype=server desttype=library device=lb1.3.0.3`

7. Define the tape drive(s) to the Tivoli Storage Manager server by using the DEFINE DRIVE command.

   `define drive maglib magdrv1 element=16 wwn=2004006045161767 serial=0000000B1982`

   `define drive maglib magdrv2 element=17 wwn=2004006045161767 serial=0000000B2070`

8. Define the drive paths from the Tivoli Storage Manager server to the tape drive by using the DEFINE PATH command:

   `define path jamaica magdrv1 srct=server desttype=drive library=maglib device=mt1.2.0.3`

   `define path jamaica magdrv2 srct=server desttype=drive library=maglib device=mt1.4.0.3`

9. Define the device class to the Tivoli Storage Manager server by using the DEFINE DEVCLASS command:

   `define devclass magclass devtype=3570 library=maglib`

10. Define the storage pool(s) to the Tivoli Storage Manager server by using the DEFINE STGPOOL command:

    `define stgpool srvfreepool magclass maxscratch=5 dataformat=nonblock`

Note: The dataformat must be defined as NONBLOCK for server-free.

11. Define the management class by using the DEFINE MGMTCLASS command:

```
define mgmtclass standard standard srvfreemc
```

12. Define the backup copygroup with the server-free storage pool, SRVFREEPOOL, by issuing the DEFINE COPYGROUP command:

```
define copygroup standard standard srvfreemc type=backup
destination=srvfreepool
```

13. Activate a STANDARD policy set in the STANDARD policy domain by issuing the ACTIVATE POLICYSET command:

```
activate policyset standard standard
```

14. Register the client node CLYDE, and the password of the node, as a Tivoli Storage Manager node by issuing the REGISTER NODE command:

```
register node clyde clydepwd
```

15. Define each disk using the DEFINE DISK command:

```
define disk clyde Harddisk1 wwn=200200a0b80cbce8 serial=1T14859797
```

Attention: The disk device name is case sensitive.

16. Define the data mover by using the DEFINE DATAMOVER command:

```
define datamover srvfreedm type=scsi wwn=200400604516150d serial=21081302315
dataformat=nonblock
```

17. Define the paths from the data mover to the tape drive using the DEFINE PATH command:

```
define path srvfreedm magdrv1 srctype=datamover desttype=drive library=maglib
lun=2
```

```
define path srvfreedm magdrv2 srctype=datamover desttype=drive library=maglib
lun=4
```

Note: You may wonder why you have to use a LUN number in the 'Define Path' command for the data mover when you were able to use the device name (mt1.2.0.3) when doing the 'Define Path' from the Tivoli Storage Manager server to the same device earlier in step 8. The LUN is required because the SAN Data Gateway data mover uses that to directly address the device on the SAN. It does not need to know the device name. The WWN and LUN combination uniquely identify the device. You can find out the LUN for the device from the Device Information window shown in Figure 3-39 on page 164.

18. Define the paths from the data mover to each disk using the DEFINE PATH command:

```
define path srvfreedm Harddisk1 srctype=datamover desttype=disk node=clyde
lun=0
```

Attention: The disk device name is case sensitive.

Part 3

IBM Tivoli Storage Manager client enhancements

Part 3 describes enhancements provided in the V5.1 Tivoli Storage Manager clients.

4

Common client enhancements

This chapter discusses the new features and enhancements delivered in IBM Tivoli Storage Manager Version 5.1, which are common to all client platforms.

4.1 Enhanced Web client interface

The Web client interface is enhanced to support a JRE 1.3.1 Swing-enabled browser. This type of browser facilitates the use of assisting devices for users with disabilities. The native look and feel of the platform running the browser is preserved. When you first access a Tivoli Storage Manager V5.1 server via your Web browser you will get a prompt to download and install a Java (TM) 2 runtime environment plug-in like this.

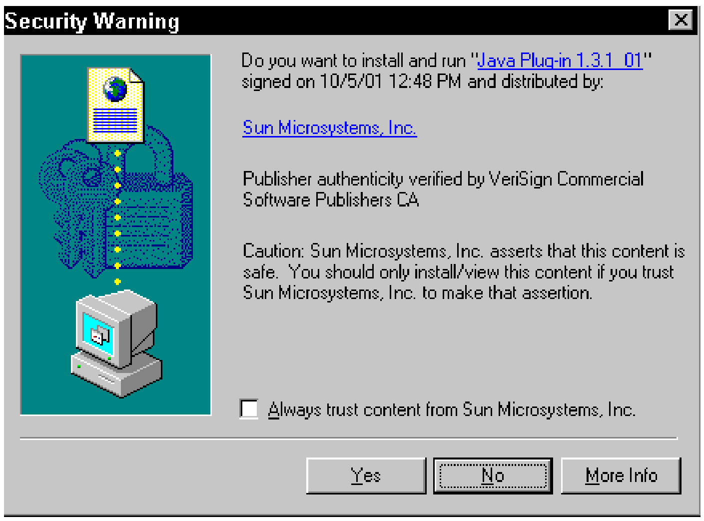

Figure 4-1 Security warning for Java plug-in

The Tivoli Storage Manager V5.1 Web Client requires Java Swing support at JRE (Java Runtime Environment) 1.3.1. Netscape 4.7 and Internet Explorer browsers come with Java support at JRE 1.1.x, which is a lower level than Web Client needs. Browsers must use the Java plug-in at Version 1.3.1. The plug-in's purpose is load JRE 1.3.1 for the browser.

- Browser requirements:
  - Netscape 6.x with the Swing plug-in installed
  - Netscape 4.7 with the JRE and plug-in at Version 1.3.1
  - Microsoft Internet Explorer 5.0, 6.0 with the JRE and plug-in at Version 1.3.1
  - JavaScript (which is the default) must be enabled in the browser
  - Netscape: **Preferences -> Advanced -> Enable JavaScript**
  - Internet Explorer: **Tools -> Internet Options -> Security -> Choose zone the Web Client runs in -> Custom Level -> Enable Scripting of Java applets**

Tivoli Storage Manager Web Client automates getting the plug-in and JRE as much as possible.

- Plug-in and JRE Installation:

  The first time the user invokes Tivoli Storage Manager, if the required JRE support is not present, the Tivoli Storage Manager Web Client will prompt the user based on which platform the browser is running. It only needs to be done the first time the user invokes the Tivoli Storage Manager Web Client.

  - Windows (32-bit), Solaris, Linux

- IE browsers: Tivoli Storage Manager Web Client automatically downloads & installs plug-in and JRE from http://java.sun.com/j2se/1.3/jre
- Netscape browsers: Tivoli Storage Manager gives above link for manual download & installation

– AIX (32-bit), HP, SGI IRIX

- User must register, download, and install the JRE manually
- Tivoli Storage Manager Web Client gives a link to the appropriate Web site for the registration:

  AIX: http://www-106.ibm.com/developerworks/java/jdk/aix/index.html

  HP: http://www.hp.com/products1/unix/java/java2/sdkrte1_3/downloads/index.html

  IRIX: http://www.sgi.com/developers/devtools/languages/javaplugin131.html

► JRE Detection and Toleration:

Sun recommends only installing one plug-in per machine.

To support this, Tivoli Storage Manager Web Client detects previously installed plug-in versions to avoid installing the plug-in if a compatible one already exists.

If a browser machine contains the following version of the JRE:

– 1.1.2.x : Tivoli Storage Manager will prompt user to install JRE 1.3.1

– 1.3.0: Tivoli Storage Manager will use the 1.3.0 JRE on Windows, Solaris, Linux, HP platforms.

– 1.3.1: Tivoli Storage Manager will not prompt to reinstall JRE for any version of JRE 1.3.1_0x.

– 1.4: Tivoli Storage Manager is not supported at JRE 1.4. Tivoli Storage Manager will prompt to install JRE at Version 1.3.1.

Note: AIX platforms require JRE 1.3.1.

► Proxy Settings:

In JRE 1.2 and above, Sun changed the Java security model. If the browser is configured to use a proxy server, Tivoli Storage Manager may return an error such as “ANS2603S Browser trying to establish connection to client; received exception: java.security.AccessControlException: access denied (java.net.SocketPermission socks2.server.ibm.com resolve)” This error occurs due to going through the proxy server instead of directly to the machine, which served the Tivoli Storage Manager Web Client applet.

To fix: Disable the proxy server for the Web Client.

Control Panel -> Java Plugin -> Proxies -> Disable “Use browser settings”

Or,

IE browsers:

Tools -> Internet Options -> Connections -> LAN Settings -> Disable “Use a proxy server”

Netscape browsers:

Edit -> Preferences -> Advanced -> Proxies -> Enable “Direct connection to the Internet”

In V3 and V4 of Tivoli Storage Manager (which used JRE 1.1) if the proxy was not disabled the applet would not load and you would get no error message. Now, you get the error message.

4.2 Multi-session restore

IBM Tivoli Storage Manager 5.1 introduces Multi-Session restore which allows the backup-archive clients to perform multiple restore sessions for *no query* restore operations, increasing the speed of restores. This is similar to the multiple backup session support.

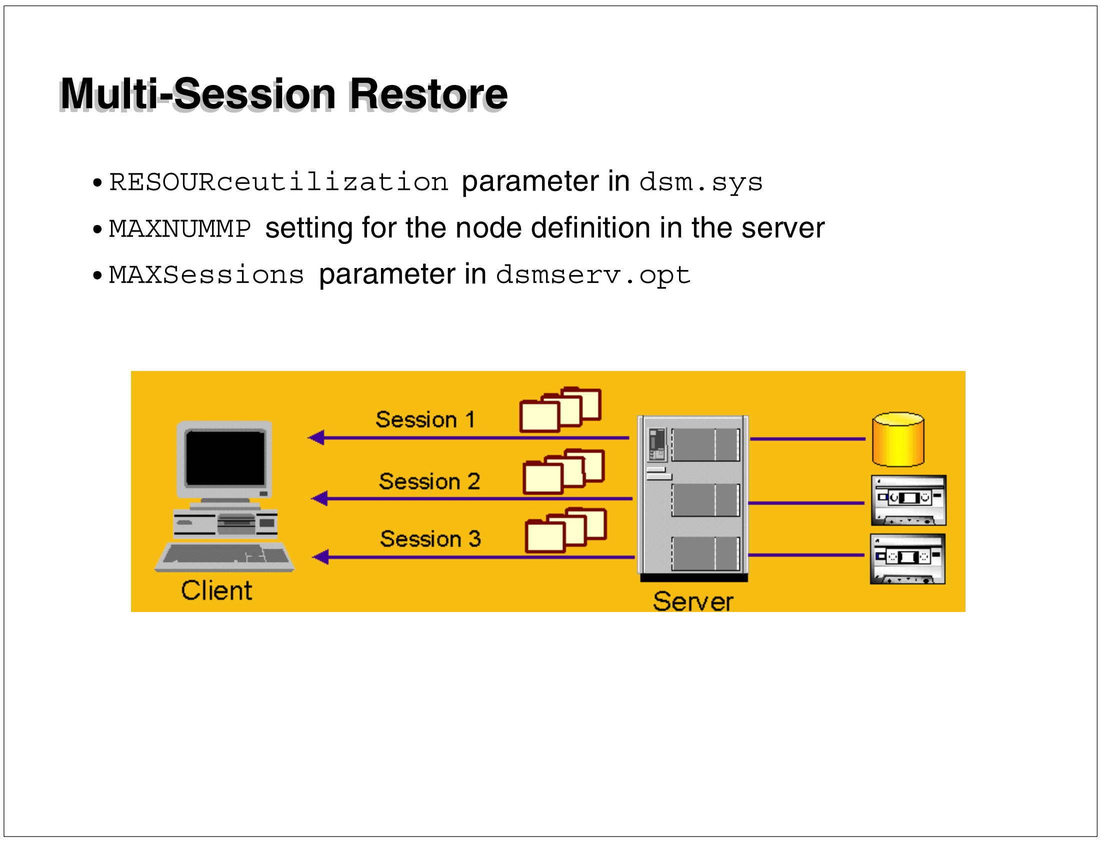

Figure 4-2 Multi-session restore

Please refer to the *Tivoli Storage Manager Backup-Archive Installation and User's Guide Version 5 Release 1*, GC32-0789, for each supported client platform, for more detailed information about the no query concept.

The Multi-Session restore exploits the mount points available on the server. If the data to be restored resides on several tapes, and there are sufficient mount points available, and the restore is done using the no query restore protocol, then multiple sessions can be used to restore the data. If there are not enough mount points available, only the client, not the server, will issue the `ANS1312E` message:

```
'ANS1312E Server media mount not possible'
```

Note: Multiple restore sessions is only allowed for *no query* restore operations.

There are potentially undesirable aspects of running multiple sessions. The client could produce multiple accounting records or the server may not start enough concurrent sessions. To avoid this, the servermaxsessions parameter must be reviewed and possibly changed.

The Multi-Session restore is a follow-on work to the LAN-Free restore work completed in Tivoli Storage Manager Version 4.2. With Lan-Free, we had multiple restore sessions (one for the storage agent and one for the server). Knowledge of Lan-Free data movement is not essential for understanding Multi-Session restore, but it is helpful.

How to use multi-session restore

The number of sessions used during the restore is dependant on the number of mount points available, the `MAXNUMMP` setting for the client node in the server, and the `RESOURceutilization` parameter setting in the `dsm.sys` file for the client. Since the number of sessions increase when performing the Multi-Session restore, please use a proper setting for the `dsmserv.opt` file parameter `MAXSessions`.

RESOURceutilization

When a restore is requested, the default is to use a maximum of two sessions, based on how many tapes the requested data is stored on, how many tape drives are available, and the maximum number of mount points allowed for the node.

The default value for the `RESOURceutilization` parameter is one (1) and the maximum value is ten (10).

For example, if the data to be restored are on five different tape volumes, and the maximum number of mount points is five for the node requesting the restore, and `RESOURceutilization` is set to three, then three sessions will be used for the restore. If the `RESOURceutilization` setting is increased to five, then five sessions will be used for the restore. There is a one-to-one (1:1) relationship to the number of restore sessions allowed and the `RESOURceutilization` setting.

The following server stanza in Example 4-1 for the `dsm.sys` client file, set the `RESOURceutilization` parameter to the maximum value of ten.

Example 4-1 Sample dsm.sys

```
servername  TSM_AIX
        tcpserveraddress    localhost
        passwordaccess      generate
        nodename            DUDE
        inclexcl            /tsm/etc/dude-tsmaix.ix
        resourceutilization 10
```

MAXNUMMP

The `MAXNUMMP` setting for the node, in the server, specifies the maximum number of mount points a node is allowed to use on the server and can be set to an integer from 0-999. Zero specifies that the node cannot acquire any mount point for a backup or archive operation. However, the server will still allow the node a mount point for a restore or retrieve operation. If the client stores its data to a storage pool that has copy storage pools defined for simultaneous backup, the client may require additional mount points. As a general rule, you must assign one mount point for each copy storage pool of sequential device type. If the primary storage pool is of sequential device type, then assign a mount point for the primary storage pool as well.

The following **dsmadmc** command set the `MAXNUMMP` parameter for node `DUDE` to the maximum value of 999:

```
update node DUDE maxnummp=999
```

Restore session example

After setting the `RESOURceutilization` parameter in the `dsm.sys` file for the client and the `MAXNMMP` setting for the node in the server, backup and restore operations attempt to use as many mountpoints that are allowed and available. The following **dsmc** command restores previously backed up files to the `/home` directory (note the trailing slash):

```
dsmc rest /home/
```

By using the following **dsmadmc** command, the usage of mount points can be monitored in the server during a clients multi-session restore operation:

```
query mount
```

The following output sample in Example 4-2 illustrates the mounting of files in a FILE device class (`SPOOL`) where the client `DUDE` has stored files that are requested for restored.

Example 4-2 Mounting files

```
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM_AIX: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/26/02   17:40:38  Last access: 03/26/02   17:38:25

ANS8000I Server command: 'query mount'
ANR8376I Mount point reserved in device class SPOOL, status: RESERVED.
ANR8376I Mount point reserved in device class SPOOL, status: RESERVED.
ANR8333I FILE volume /tsm/spool/00001AE0.BFS is mounted R/O, status: IN USE.
ANR8377I FILE volume /tsm/spool/00001AF5.BFS is mounted R/O, status: DISMOUNTING.
ANR8376I Mount point reserved in device class SPOOL, status: RESERVED.
ANR8376I Mount point reserved in device class SPOOL, status: RESERVED.
ANR8376I Mount point reserved in device class SPOOL, status: RESERVED.
ANR8333I FILE volume /tsm/spool/00001AF2.BFS is mounted R/O, status: IN USE.
ANR8379I Mount point in device class SPOOL is waiting for the volume mount to complete,
status: WAITING FOR VOLUME.
ANR8333I FILE volume /tsm/spool/00001AF4.BFS is mounted R/O, status: IN USE.
ANR8334I         10 matches found.

ANS8002I Highest return code was 0.
```

The Example 4-2 output sample shows that three files are mounted and `IN USE`, one file is `DISMOUNTING`, five mount point are `RESERVED` and one mount point is `WAITING FOR VOLUME`. The output in Example 4-3 shows ten sessions being established by the `DUDE` node.

Example 4-3 Establishing sessions

```
ANS8000I Server command: 'query session'

  Sess Comm.  Sess     Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State    Time    Sent   Recvd Type
------ ------ ------ ------ ------- ------- ----- -------- --------------------
   431 Tcp/Ip SendW    0 S  504.6 K     396 Node  AIX      DUDE
   432 Tcp/Ip SendW    0 S  504.6 K     397 Node  AIX      DUDE
   436 Tcp/Ip MediaW   0 S  365.9 K     314 Node  AIX      DUDE
   437 Tcp/Ip SendW    0 S  302.1 K     314 Node  AIX      DUDE
   438 Tcp/Ip Run      0 S  101.1 K     314 Node  AIX      DUDE
   439 Tcp/Ip Run      0 S  301.9 K     314 Node  AIX      DUDE
```

```
   440 Tcp/Ip Run      0 S  101.1 K     314 Node  AIX      DUDE
   441 Tcp/Ip SendW    0 S  201.8 K     314 Node  AIX      DUDE
   442 Tcp/Ip Run      0 S  201.4 K     314 Node  AIX      DUDE
   446 Tcp/Ip SendW    0 S  302.1 K     315 Node  AIX      DUDE
   455 Tcp/Ip Run      0 S      124     158 Admin AIX      ADMIN
```

Sample configuration setup

The next sample in Example 4-4 shows how a include-exclude file can be used to point a directory to a specific management class, in this case a management class named MSESS.

Example 4-4 Sample inclexcl file

```
include /home/* MSESS
exclude /unix/
exclude.dir /unix/
exclude /.../core
```

The MSESS management class pointed to by the include-exclude file in the previous Example 4-4, will store all backup files in the SPOOLPOOL storage pool.

Example 4-5 Sample backup copygroup for the MSESS management class

```
ANS8000I Server command: 'query copygroup * active msess f=d'

            Policy Domain Name: STANDARD
               Policy Set Name: ACTIVE
               Mgmt Class Name: MSESS
               Copy Group Name: STANDARD
               Copy Group Type: Backup
          Versions Data Exists: 2
         Versions Data Deleted: 1
         Retain Extra Versions: 30
           Retain Only Version: 60
                     Copy Mode: Modified
            Copy Serialization: Shared Static
                Copy Frequency: 0
              Copy Destination: SPOOLPOOL
Last Update by (administrator): ADMIN
         Last Update Date/Time: 03/21/02   14:48:29
              Managing profile:
```

The next sample in Example 4-6 shows the SPOOLPOOL storage pool and that it belongs to the SPOOL device class.

Example 4-6 SPOOLPOOL storage pool

```
ANS8000I Server command: 'query stgpool spoolpool'

Storage      Device       Estimated    Pct    Pct  High  Low  Next Stora-
Pool Name    Class Name    Capacity   Util   Migr   Mig  Mig  ge Pool
                               (MB)                 Pct  Pct
-----------  ----------  ----------  -----  -----  ----  ---  -----------
SPOOLPOOL    SPOOL       9,721,908.    0.0    0.0    90   70
                                  6
```

The SPOOL device class, shown in the Example 4-7, allow up to 99 mount points, or in this case files, to be mounted simultaneously.

Example 4-7 SPOOL device class

```
ANS8000I Server command: 'query devclass spool'

Device       Device       Storage    Device      Format     Est/Max    Mount
Class        Access          Pool    Type                   Capacity   Limit
Name         Strategy       Count                             (MB)
---------    ----------   -------    ---------   ------     --------   ------
SPOOL        Sequential         1    FILE        DRIVE           0.1       99
```

4.3 Consistent return codes

The Command Line Interface (CLI) for IBM Tivoli Storage Manager Version 5.1 exits with a valid, documented return code. This allows customer-written scripts to invoke the CLI and determine the success or failure of the Tivoli Storage Manager operation. Administrators can now distinguish between scheduled backups that completed successfully with no skipped files and scheduled backups that completed successfully with one or more skipped files.

The return codes are listed in Table 4-1.

Table 4-1 Return codes

Return code	Meaning
0	Operation completed successfully
4	Operation completed successfully, except for one or more skipped files
8	Operation completed with one or more warning messages issued
12	Operation failed with one or more error messages (except for skipped files) issued

Please note that:

- When running client macros, a message is issued indicating the return code for each command within the macro. When the macro is finished, a message indicating the over-all return code for the macro is issued.
- Return codes for scheduled operations are propagated to the Tivoli Storage Manager server event record for the scheduled operation.
- Scheduled operations will not run if the `PRESCHEDULECMD` command ends with a non-zero return code. This ensures that scheduled events will not run if prerequisite commands do not complete successfully.

4.4 Improved backupset restore speed

Performance for using backup sets has been improved by optimization of device block sizes.

Backup sets on portable media were introduced in client Version 3.7. The Tivoli Storage Manager V5.1 clients improve the performance of restore from local tape backup sets by doing separate read and write threads to overlap input and output and by obtaining the read block size from tape label. This has resulted in an approximate doubling of the throughput.

The number of possible values for the location option of `QUERY BACKUPSET` and `RESTORE BACKUPSET` have been *reduced* to simplify configuration and usage of backupsets.

Prior to Tivoli Storage Manager V5.1 you could specify the following device classes:

- SERVER
- FILE
- TAPE
- 8 MM
- 4 MM
- 3570
- 3590
- DLT

Now you only need to specify the following device classes:

- SERVER
- FILE
- TAPE

In addition to the previous specific tape device support, LTO support has been added.

4.5 New TDP functions

Here are the Tivoli Data Protection updates.

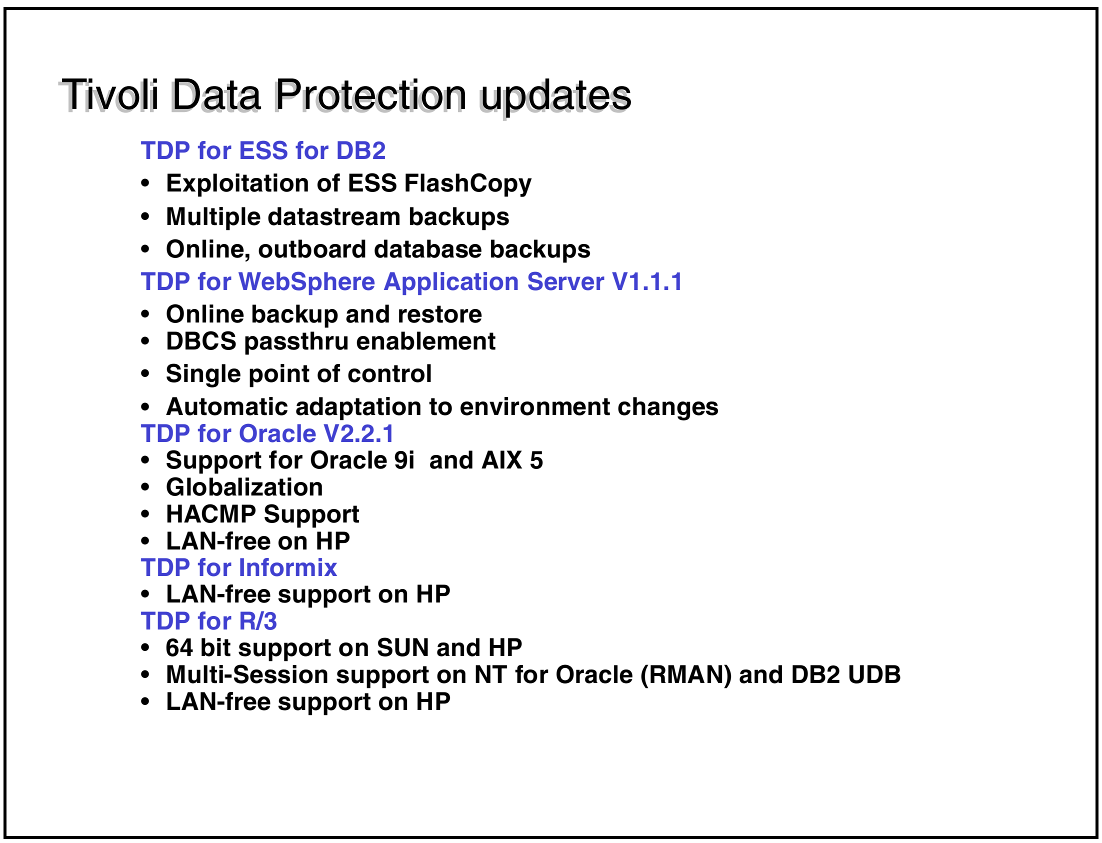

Figure 4-3 New functions in the TDP agents

TDP for ESS for DB2 now exploits ESS FlashCopy and multiple datastream backups. This allows online outboard backups.

TDP for WebSphere Application Server V1.1.1 now allows online backup and restore as well as automated adaptation to environment changes.

TDP for Oracle V2.2.1 now supports Oracle 9i and AIX 5. HACMP support is also provided as well as LAN-free backups on HP-UX platform.

TDP for Informix now does LAN-free backups on the HP-UX platform.

TDP for SAP R/3 now has 64 bit support on Sun and HP-UX. Multi-session support is now provided on NT platform for Oracle (RMAN) and DB2 UDB. LAN-free support is provided on the HP-UX platform.

5

Platform specific client enhancements

This chapter discusses the new features and enhancements delivered in IBM Tivoli Storage Manager Version 5.1, which are specific to only certain client platforms.

5.1 AIX

The IBM Tivoli Storage Manager 5.1 client for AIX introduces some new features and additional support.

5.1.1 Changes in IBM Tivoli Storage Manger Version 5.1 client

Following is a summary of changes since Tivoli Storage Manager Version 5.1:

AIX

Changes in IBM Tivoli Storage Manger version 5.1 client

- Enhanced Web Client Interface
- Support for High Availability Cluster Multi Processing (HACMP) on AIX Client
- Multiple session No Query Restore
- Consistent Client Return Codes

Figure 5-1 Changes in IBM Tivoli Storage Manger Version 5.1 client

- Enhanced Web Client Interface

  The Web client interface is enhanced to support a JRE 1.3.1 Swing-enabled browser. This type of browser facilitates the use of assistive devices for users with disabilities. The native look and feel of the platform running the browser is preserved.

- Support for High Availability Cluster Multi Processing (HACMP) on AIX Client

  Tivoli Storage Manager supports HACMP failover on AIX. This allows the client to continue operating in the event of an HACMP node failover and fallback.

- Multiple session No Query Restore

  The backup-archive clients can now perform multiple restore sessions for no query restore operations, increasing the speed of restores. This is similar to the multiple backup session support. It exploits the mount point available on the server. If data is backed up on multiple tapes, and if the server has multiple mount points available, then the restore starts a session for each tape, up to the number your administrator configures.

- Consistent Client Return Codes

  Reliable, consistent, and documented return codes have been added to the command line client and the scheduler. This facilitates automation of client operations via user-written scripts. Administrators can now distinguish between scheduled backups that completed successfully with no skipped files and scheduled backups that completed successfully with one or more skipped files. Also if the preschedulecmd command ends with non-zero return codes, the scheduled event will not run. This ensures that scheduled events will not run if prerequisite commands do not complete successfully.

5.1.2 HACMP support for the B/A client on AIX

The objective of B/A client support for HACMP on AIX is to provide support for scheduled Tivoli Storage Manager client operations to continue processing during an AIX HACMP node failover and fallback. Both local and shared filesystems will be backed up to the appropriate Tivoli Storage Manager client account, regardless of which cluster node currently has the filesystems mounted.

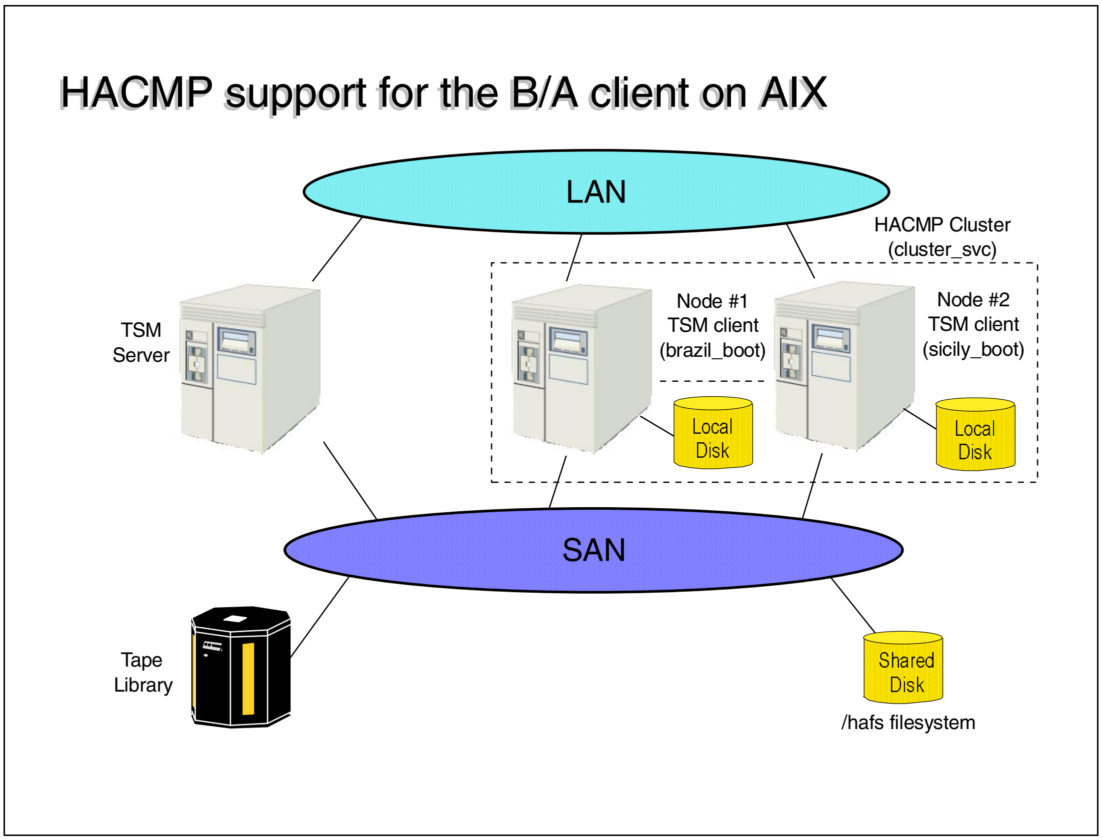

Figure 5-2 HACMP support for the B/A client on AIX

Tivoli Storage Manager supports a wide variety of configurations as well as the following takeover relationships in an HACMP environment with only one node controlling the shared resources at a time.

- Cascading (with or without fallback)
- Rotating

There is no support for concurrently mounted volumes or resource groups.

Prerequisite tasks and supported environment

Configuring the Tivoli Storage Manager B/A client to operate in a HACMP cluster requires understanding of both Tivoli Storage Manager and HACMP configuration.

The user should be familiar with the following:

- Tivoli Storage Manager documentation:

  “Appendix B. Configuring the Backup-Archive Client in a HACMP Takeover Environment” in *IBM Tivoli Storage Manager for UNIX Backup-Archive Client Installation and User’s Guide Version 5.1*, GC32-0789

- Tivoli Storage Manager options file and options
- AIX file systems
- The user should already have a functioning HACMP cluster, able to Start / Stop / Takeover / Take back an HACMP environment

The following hardware/software configurations are required for HACMP B/A support:

- AIX 4.3.3 or later, must be identical on all nodes
- HACMP 4.4 or later (or) HACMP/ES 4.4 or later
- No Tivoli Storage Manager server dependencies
- HACMP Cluster Info Daemon must be running
- An external hard disk or an nfs mounted hard disk on both systems for shared access
- Two network adapters per cluster node
- Two identical RS6000 systems

Note: In our environment we used two different model RS6000 systems with a single network adapter in each node. This is not the ideal configuration and should not be considered for a production environment. It was however suitable to demonstrate the new function.

New client option - clusternode

The clusternode option specifies whether the node is to participate in a HACMP environment.

The effect of enabling this option is to cause the Tivoli Storage Manager client node name to default to the name of the HACMP cluster, not to the hostname, as is usually the case.

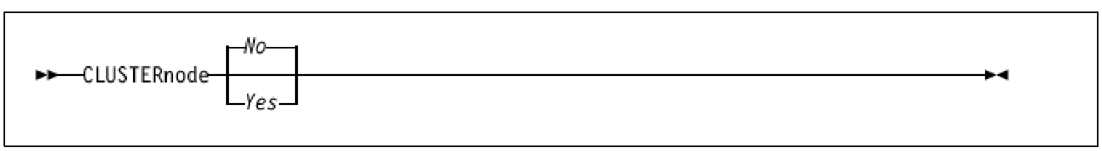

Figure 5-3 CLUSTERnode command syntax

Yes Specifies that you want to backup cluster resources and participate in cluster failover for high availability

No Specifies that you not want the B/A client to participate in cluster failover.

This option should be set in the dsm.sys file for each node participating in the cluster.

Step 1: Install the B/A client code

Locate the latest version of the B/A client code and make the media available to AIX. This may involve mounting the original CD-ROM or downloading the latest maintenance level code from the Tivoli Web site. The most up to date version is always available at the following URL:

http://www.tivoli.com/tsm

If you are installing the B/A client code for the first time, switch to the root user, and use either the “installp” or “smitty install” commands to install the required filesets. The default installation directory for the B/A client is /usr/tivoli/tsm/client/ba/bin/. Several sample configuration files can be found in this directory.

If you are upgrading over an existing installation, switch to the root user, and use either “installp” or “smitty update_all” to update the already installed filesets, Example 5-1. Your existing configuration files will be retained.

Example 5-1 Minimum filesets required to install the B/A client on AIX

```
tivoli.tsm.client.api.aix43.32bit
tivoli.tsm.client.ba.aix43.32bit.base
tivoli.tsm.client.ba.aix43.32bit.common
tivoli.tsm.client.web.aix43.32bit
```

If you are upgrading over an existing installation, ensure you stop any currently running Tivoli Storage Manager services before performing the upgrade. For example, kill any running client scheduler or client acceptor tasks, and remove any Tivoli Storage Manager client related entries from the /etc/inittab file.

Repeat this set for each node in the cluster.

Step 2: Create client node accounts

Before using the new B/A clients we need to define several client node accounts on the Tivoli Storage Manager server.

If you plan to backup both local (per node) filesystems and shared (HACMP) filesystems, you will require:

- A client account related to each node in the cluster (in our example: brazil_fs, siciliy_fs)
- A client node related to each resource group in the cluster (in our example: cluster1)

If you only plan to backup the shared (HACMP) filesystems, you will only require:

- A client node related to each resource group in the cluster (in our example: cluster1)

Start an administrative session and define the new client node accounts. You could use a command similar to that shown in Example 5-2.

Example 5-2 Registering a client node account

```
register node newnode newpassword userid=none domain=mydomain url=http://newnode:1581
```

If you plan to use the Web client interface, you should set the URL parameter as follows:

- On the node accounts, set the URL to the nodes reliable host address (if any), and set the TCP port to match what you later set in the client dsm.sys file (the default is 1581)
- On the resource group account/s, set the URL to the cluster service address, and set the TCP port to match what you set in the shared dsm.sys file (good values might be 1582+)

Example 5-3 Output from “select nodename, url from nodes” command

```
NODE_NAME          URL
------------------ ------------------
```

```
BRAZIL_FS         HTTP://BRAZIL_RELIABLE:1581
CLUSTER1          HTTP://CLUSTER_SERVICE:1582
SICILY_FS         HTTP://SICILY_RELIABLE:1581
```

There are no other special considerations regarding defining client node accounts for use in a HAMCP environment.

Step 3: Create the HACMP shared filesystems

Before configuring the B/A client we need to define some local and shared filesystems required to hold our configuration files.

Remember to use the correct HACMP tools to define any shared filesystems and logical volumes. You can access the SMIT menu for this task via the command “smitty cl_lvm”.

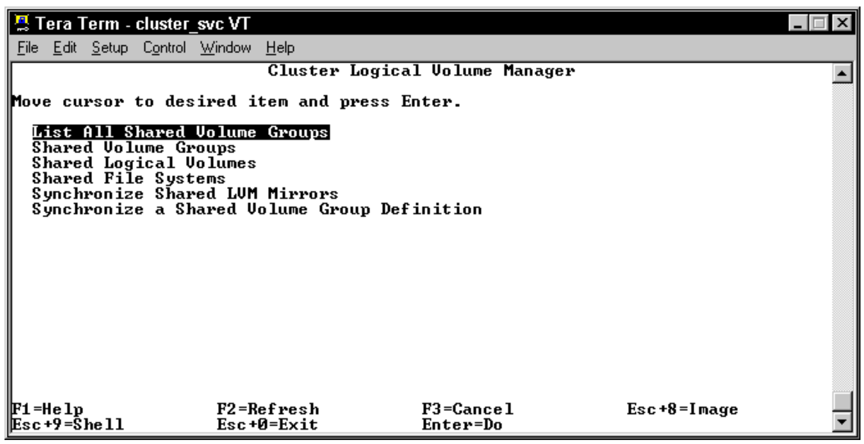

Figure 5-4 Defining shared filesystem via the smitty HACMP menu

In our environment we created a local filesystem (/hacmp) in the rootvg volume group to hold our HACMP startup scripts. We also created a shared filesystem (/hafs) in the tsm_vg volume group to hold some of our B/A client configuration files. We also created a sub-directory (/hafs/baclient) to isolate the B/A client shared configuration files.

Example 5-4 Filesystems defined for the B/A client in HACMP

```
root@brazil:/hacmp: lsvg -o | lsvg -i -l

tsm_vg:
LV NAME             TYPE      LPs   PPs   PVs  LV STATE      MOUNT POINT
loglv00             jfslog    1     1     1    open/syncd    N/A
lv01                jfs       32    32    1    open/syncd    /hafs

rootvg:
LV NAME             TYPE      LPs   PPs   PVs  LV STATE      MOUNT POINT
hd5                 boot      1     1     1    closed/syncd  N/A
hd6                 paging    34    34    1    open/syncd    N/A
hd8                 jfslog    1     1     1    open/syncd    N/A
hd4                 jfs       1     1     1    open/syncd    /
hd2                 jfs       40    40    1    open/syncd    /usr
hd9var              jfs       1     1     1    open/syncd    /var
hd3                 jfs       8     8     1    open/syncd    /tmp
hd1                 jfs       1     1     1    open/syncd    /home
paging00            paging    30    30    1    open/syncd    N/A
lv00                jfs       1     1     1    open/syncd    /hacmp
```

Manually mount the shared filesystems to allow us to configure the B/A client before starting the HACMP services. See Example 5-5.

Example 5-5 mount example

```
root@brazil:/hacmp: mount
  node       mounted        mounted over    vfs       date        options
-------- ---------------  ---------------  ------ ------------ ---------------
         /dev/hd4         /                jfs    Mar 19 16:05 rw,log=/dev/hd8
         /dev/hd2         /usr             jfs    Mar 19 16:05 rw,log=/dev/hd8
         /dev/hd9var      /var             jfs    Mar 19 16:05 rw,log=/dev/hd8
         /dev/hd3         /tmp             jfs    Mar 19 16:05 rw,log=/dev/hd8
         /dev/hd1         /home            jfs    Mar 19 16:06 rw,log=/dev/hd8
         /dev/lv00        /hacmp           jfs    Mar 19 16:06 rw,log=/dev/hd8
         /dev/lv01        /hafs            jfs    Mar 19 17:09 rw,log=/dev/loglv00
```

Remember to unmount any manually mounted shared filesystems, and vary off the tsm_vg volume group, before attempting to start the HACMP services.

Step 4: Configure the B/A client code

The HACMP support for the B/A client allows for the scheduled backup of both local (generally rootvg only) and shared (HACMP managed non-rootvg) filesystems. This is achieved by configuring multiple server stanzas in the dsm.sys file, and then running multiple client acceptor daemons with differing dsm.opt client options files.

It is necessary to configure the B/A client dsm.sys file on each node in the cluster. The dsm.sys file should contain at least two server stanzas. The different stanzas are required for:

- One server stanza to backup local filesystems
- One server stanza for each cluster resource group to be backed up

It is also necessary to configure one version of the B/A client dsm.opt file on each node and each resource group in the cluster. The different dsm.opt files will each control:

- One dsm.opt file on each node to backup local (non-clustered) filesystems
- One dsm.opt file on each resource group's shared filesystem its own backup

Configure dsm.sys for local and shared filesystem backup

The dsm.sys file used to control local and shared filesystem backup should exist on every node in the cluster. Each node will have its own instance of this configuration file.

Copy the sample /usr/tivoli/tsm/client/ba/bin/dsm.sys.smp file to dsm.sys, and edit it to resemble Example 5-6. You will need to modify some parameters to suit your local environment.

Example 5-6 Example dsm.sys file for the B/A client in a HACMP environment

```
************************************************************************
* Local filesystem backup

SErvername              tsm_local
COMMmethod              TCPip
TCPPort                 1500
TCPServeraddress        tsmserver

nodename                sicily_fs
*TCPClientaddress       sicily_reliable
HTTPport                1581
```

```
domain                  / /usr /var /home /hacmp
inclexcl                /usr/tivoli/tsm/client/ba/bin/dsm.inx

passwordaccess          generate

managedservices         schedule webclient
schedmode               polling

errorlogname            /var/dsmerror.log
schedlogname            /var/dsmsched.log

************************************************************************
* Shared filesystem backup

SErvername              tsm_shared
COMMmethod              TCPip
TCPPort                 1500
TCPServeraddress        tsmserver

clusternode             yes
*TCPClientaddress       cluster_svc
HTTPport                1582

domain                  /hafs
inclexcl                /hafs/baclient/dsm.inx

passwordaccess          generate
passworddir             /hafs/baclient

managedservices         schedule webclient
schedmode               polling

errorlogname            /hafs/baclient/dsmerror.log
schedlogname            /hafs/baclient/dsmsched.log
```

You should set the schedlogname and errorlogname parameter to save the logging information to consistent files. In the case of the clusternode's stanza this should be to somewhere on the shared filesystem.

Note: By setting the DOMAIN parameter in each stanza we can ensure that only the appropriate filespaces are processed for each client node account on this node in the cluster. Only filespaces named in the DOMAIN parameter will be considered for backup.

If you later add a new filesystem it will be excluded from the backup unless you update the DOMAIN parameter to include the new filespace name.

Repeat this step for every node in your cluster. The server stanza used for shared filespace backup should be identical between all the nodes.

Configure dsm.opt for local filesystem backup

The dsm.opt file used to control local filesystem backup should be placed on the local filesystem on each node in the cluster.

Copy the sample /usr/tivoli/tsm/client/ba/bin/dsm.opt.smp file to dsm.opt (in the same directory), and edit it to resemble Example 5-7. You may add other parameters to suit your local environment.

Example 5-7 Example dsm.opt file for local filesystem backup in a HACMP environment

```
SErvername              tsm_local
```

Replace the servername parameter with name of the stanza used in your dsm.sys file for local filespace backup.

Configure dsm.opt for shared filesystem backup

The dsm.opt file used to control shared filesystem backup should be placed on the shared filesystem itself. This allows us to maintain only a single instance of this configuration file.

Copy the sample /usr/tivoli/tsm/client/ba/bin/dsm.opt.smp file to a directory in the shared filesystem (in our example /hafs/baclient/dsm.opt), and edit it to resemble Example 5-8. You may add other parameters to suit your local environment.

Example 5-8 Example dsm.opt file for shared filesystem backup in a HACMP environment

```
SErvername              tsm_shared
```

Replace the servername parameter with name of the stanza used in your dsm.sys file for shared filespace backup.

Cache the password for the B/A client

Before continuing to start any Tivoli Storage Manager services or configure HACMP we need to manually run the dsmc client. This will allow us to cache the Tivoli Storage Manager client node password for both the local and shared client node accounts. These passwords will then be used automatically when connecting to the Tivoli Storage Manager server.

To cache the password for the local client node account, enter the command shown in Example 5-9.

Example 5-9 Caching the local client password

```
root@sicily:/hafs/baclient: dsmc q sess -se=tsm_local
Tivoli Storage Manager
Command Line Backup/Archive Client Interface - Version 5, Release 1, Level 0.0 g
4
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Node Name: SICILY_FS
Please enter your user id <SICILY_FS>:

Please enter password for user id "SICILY_FS":

Session established with server TSM_SERVER1: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/19/02   11:14:05  Last access: 03/19/02   11:14:05

TSM Server Connection Information
```

```
Server Name.............: TSM_SERVER1
Server Type.............: AIX-RS/6000
Server Version..........: Ver. 5, Rel. 1, Lev. 0.0
Last Access Date........: 03/19/02   11:14:05
Delete Backup Files.....: "No"
Delete Archive Files....: "Yes"

Node Name...............: SICILY_FS
User Name...............: root
```

Run the `htis` command a second time to ensure the password was successfully cached. On the second attempt you should not be prompted to enter the password.

Repeat this step on each node in your cluster. This is necessary because the cached password is saved on the local filesystem.

To cache the password for the shared client node account, enter the command shown in Example 5-10.

Example 5-10 Caching the shared client password

```
root@sicily:/hafs/baclient: dsmc q sess -se=tsm_shared
Tivoli Storage Manager
Command Line Backup/Archive Client Interface - Version 5, Release 1, Level 0.0 g
4
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Node Name: CLUSTER1
Please enter your user id <CLUSTER1>:

Please enter password for user id "CLUSTER1":

Session established with server TSM_SERVER1: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/19/02   11:14:05  Last access: 03/19/02   11:14:05

TSM Server Connection Information

Server Name.............: TSM_SERVER1
Server Type.............: AIX-RS/6000
Server Version..........: Ver. 5, Rel. 1, Lev. 0.0
Last Access Date........: 03/19/02   11:14:05
Delete Backup Files.....: "No"
Delete Archive Files....: "Yes"

Node Name...............: CLUSTER1
User Name...............: root
```

Run this command a second time to ensure the password was successfully cached. On the second attempt you should not be prompted to enter the password.

You only need to perform this step on one node in your cluster. This is because the cached password is saved to the shared filesystem, and the shared filesystem will be available to any node that assumes control of the resource group.

Start the Tivoli Storage Manager client acceptor

The Web client for the local filesystems needs to be started at system boot time. Use the command shown in Example 5-11 to add the Tivoli Storage Manager client acceptor to the AIX /etc/inittab file. This will cause the client scheduler to start automatically at boot time. If you want it to run immediately you will need to manually start the client acceptor daemon. Starting the daemon manually is also shown in Example 5-11.

Example 5-11 Setting the Tivoli Storage Manager client acceptor to start at system boot time

```
root@sicily:/hacmp: mkitab "dsmcad:2:once:/usr/tivoli/tsm/client/ba/bin/dsmcad
-optfile=/usr/tivoli/tsm/client/ba/bin/dsm.opt >/dev/null 2>&1"

root@sicily:/hacmp: (dsmcad:2:once:/usr/tivoli/tsm/client/ba/bin/dsmcad
-optfile=/usr/tivoli/tsm/client/ba/bin/dsm.opt >/dev/null 2>&1 &)

root@sicily:/hacmp: ps -ef | grep dsmcad | grep -v grep
    root 270074      1   0 14:29:04      -  0:00 /usr/tivoli/tsm/client/ba/bin/dsmcad
-optfile=/usr/tivoli/tsm/client/ba/bin/dsm.opt
```

The Web client for the shared filesystems will stop and start with the appropriate HACMP resource group.

The screens shown in Figure 5-5 describe the Web client restore selection box for both local and shared filesystems. The Web client for shared filesystems will not be available until the HACMP services are started.

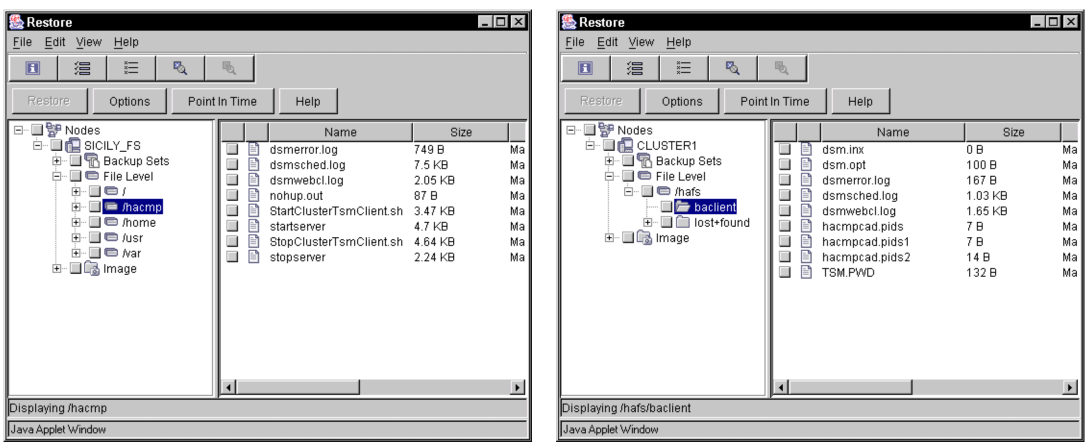

Figure 5-5 Two Web client sessions, for local and shared file systems

After completing the configuration of the B/A client and HACMP services, you should be able to use the Web client interface to backup and restore both local and shared filesystems.

Remember to restart the Tivoli Storage Manager client acceptor after any changes to the B/A client configuration files.

Step 5: Define a client schedule

In order to automate the backup of the local and shared filesystems we need to create a Tivoli Storage Manager client schedule. You may use a single client schedule to backup all the local nodes in your cluster and also the cluster shared resource groups.

Define the client schedule

Start an administrative session and define a client schedule of the "incremental" type.

Example 5-12 Defining a client schedule from the administrative command line

```
define schedule mydomain filesystem_incr action=incr starttime=18:00
```

Alternatively, you could use the Web server administrative interface to define the client schedule, as shown in Figure 5-6.

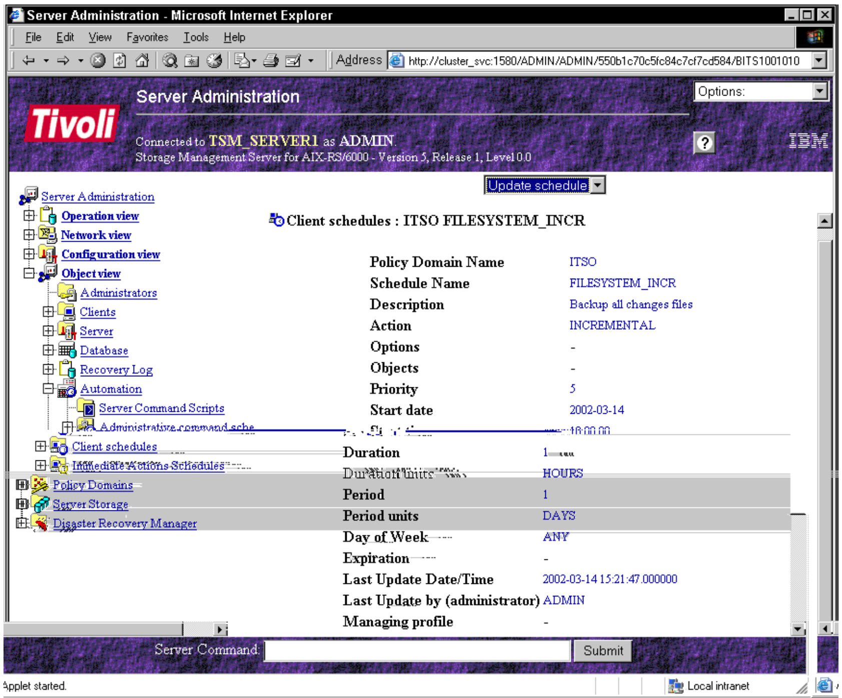

Figure 5-6 Web admin client - define schedule

Pay special consideration to the "duration" parameter when defining the client schedule. This parameter controls the window past the start time within which the clients can successfully start their scheduled backup sessions.

You should probably set the duration parameter to be equal to the estimated time of a typical incremental backup. On failover, this should allow the takeover node to restart the scheduled backup if the typical backup would have still been running at that time.

Define the schedule associations

Next, we need to associate the client schedule with all the client node accounts in the cluster, and also with the resource group account in the cluster.

Start an administrative session and associate the client accounts with the new client schedule.

Example 5-13 Associating nodes to a client schedule from the administrative command line

```
define association mydomain filesystem_incr brazil_fs sicili_fs cluster1
```

Alternatively, you could use the Web server administrative interface to associate the clients to the new schedule.

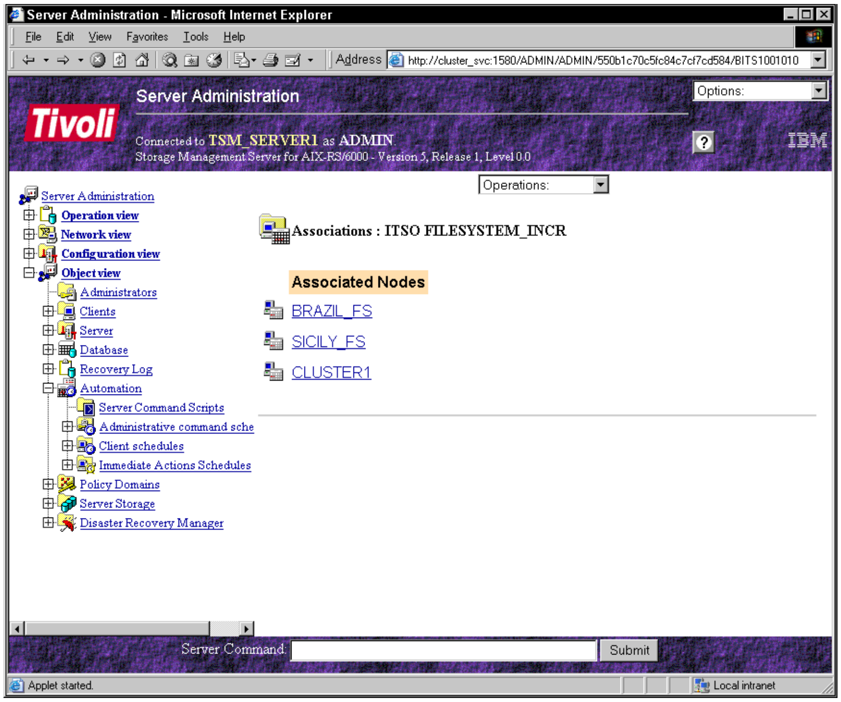

Figure 5-7 Web admin client - associate client nodes

Set the randomize parameter

The only supported client scheduling method in a HACMP environment is “client polling”. This method uses the “randomize” server parameter to spread the scheduled client backup sessions in some percentage of the “duration” start window. If we want to accurately set the “duration” start window to improve the chances of a take-over node restarting an interrupted backup, we should consider setting the randomize parameter to zero.

Start an administrative session and set the randomize parameter.

Example 5-14 Setting randomize to zero

```
tsm: TSM_SERVER1>set randomize 0
ANR2522I Randomization set to 0 percent.
```

You should also consider what effect changing this parameter will have on your other scheduled backups and on the load on your Tivoli Storage Manager server.

Step 6: Customize the sample HACMP scripts

Two sample HACMP scripts are provided with the B/A client code. These scripts can be used to control the starting and stopping of the Tivoli Storage Manager client acceptor (dsmcad) process. The dsmcad process is will be configured to initiate scheduled backups for the shared filesystem. You will need to slightly modify the sample scripts to suit your environment.

> **Note:** Because these scripts are used to start/stop the HACMP application servers, they must be available on local, non-HACMP managed, filesystems on each node in the cluster.
>
> You will need to update the scripts on each node in the cluster, and remember to propagate any future changes to all nodes.

The sample startup script

The startup script is used to start the dsmcad process as the resource group comes online. An identical copy of this script needs to exist on each node in the cluster.

Copy the sample /usr/tivoli/tsm/client/ba/bin/StartClusterTsmClient.sh.smp file to an area on the local filesystem (in our example /hacmp/StartClusterTsmClient.sh), and edit it to resemble Example 5-15. Modify the following section of the startup script to refer to the shared area for HACMP B/A client configuration files.

Example 5-15 Extract from the HACMP startup script for the dsmcad schedule client

```
# Set default HACMP DIRECTORY if environment variable not present
if [[ $HADIR = "" ]]
then
     HADIR=/hafs/baclient
fi
```

When executed the startup script will start the dsmcad process, referring to the dsm.opt configuration file located in the $HADIR directory. It will also save the process ids (PID) of the dsmcad process into files in the $HADIR directory.

Duplicate this file to every node in your cluster.

The sample stop script

The startup script is used to stop the dsmcad process when the resource group halts gracefully. An identical copy of this script needs to exist on each node in the cluster.

Copy the sample /usr/tivoli/tsm/client/ba/bin/StopClusterTsmClient.sh.smp file to an area on the local filesystem (in our example /hacmp/StopClusterTsmClient.sh) and edit it to resemble Example 5-16. Modify the following section of the startup script to refer to the shared area for HACMP B/A client configuration files.

Example 5-16 Extract from the HACMP stop script for the dsmcad schedule client

```
# Set default HACMP DIRECTORY if environment variable not present
if [[ $HADIR = "" ]]
then
     HADIR=/hafs/baclient
fi
```

When executed the stop script will stop the dsmcad process. It does this by retrieving the process ids (PID) of the dsmcad process from the files in the $HADIR directory and killing the returned process ID.

Duplicate this file to every node in your cluster.

> **Note:** Before testing the script under control of HACMP, remember to run the dsmc command manually using the appropriate dsm.opt configuration file.
>
> This will allow you to enter the correct Tivoli Storage Manager client node name and password. The password will then be cached on the shared filesystem and be available to any node after a failover.

After executing the startup script the following files should exist in the shared filesystem.

Example 5-17 Typical contents of the shared area for HACMP B/A client configuration

```
root@sicily:/hacmp: ls -l /hafs/baclient/
```

```
total 32
-rw-------   1 root     sys          132 Mar 14 18:07 TSM.PWD
-rw-r--r--   1 root     sys            0 Mar 14 18:07 dsm.inx
-rw-r--r--   1 root     sys          100 Mar 14 15:43 dsm.opt
-rw-r--r--   1 root     sys            0 Mar 14 18:10 dsmerror.log
-rw-r--r--   1 root     sys            0 Mar 14 18:10 dsmsched.log
-rw-r--r--   1 root     sys            7 Mar 14 18:10 hacmpcad.pids
-rw-r--r--   1 root     sys            0 Mar 14 18:10 hacmpcad.pids1
-rw-r--r--   1 root     sys            7 Mar 14 18:10 hacmpcad.pids2
```

Remember, changes made to the start/stop scripts will need to be duplicated on each node in the cluster, but changes made to the HACMP B/A client configuration in the shared area are persistent across failover events.

Step 7: Configure HACMP

After installing the B/A client and customizing the HACMP start/stop scripts, you next need to configure HACMP.

The following example assumes you already have a functioning HACMP environment, with correctly configured Networks, Adapters, and Nodes, etc. You should be able to successfully perform a HACMP failover and recovery before continuing. The following examples build on the HACMP environment we configured in section 3.1.2, “HACMP support for Tivoli Storage Manager server on AIX” on page 88 for the Tivoli Storage Manager server.

First (if you have not already done so), define a Resource Group to manage the Tivoli Storage Manager B/A client. You can access the SMIT menu for this task via the command “smitty cm_add_res”.

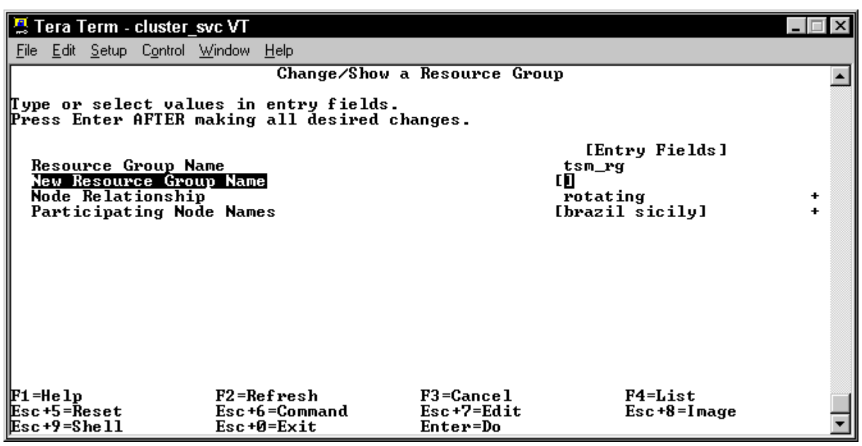

Figure 5-8 Defining a HACMP resource group for the B/A client

You will need to configure the following parameters:

Resource Group Name Enter a logical name for your resource group. Maximum length is 32 characters.

Node Relationship Toggle between Cascading/Concurrent/Rotating.

CASCADING resources are resources which may be assigned to be taken over by multiple nodes in a prioritized manner. When a node fails, the active node with the highest priority acquires the resource. When the failed node rejoins, the node with the highest priority acquires the resource.

ROTATING resources are resources which may be acquired by any node in its resource chain. When a node fails, the resource will be acquired by the highest priority standby node. When the failed node rejoins, the resource remains with its new owner.

CONCURRENT resources are not supported by the B/A client.

Participating Node Names
Enter the nodes in the resource chain for the resource group. These are the nodes which may acquire the resource. The order you list the nodes can determine the priority of the Resource Group allocation, depending on the Node relationship.

CASCADING RESOURCES: Priority decreases from left to right.

CONCURRENT RESOURCES: All nodes are at the same priority. The B/A client does not support Concurrent resources.

ROTATING RESOURCES: Priority decreases from left to right, but is only used to determine which standby acquires the resource in a multiple standby configuration.

Secondly, define a HACMP Application Server to manage the Tivoli Storage Manager client scheduler. this is where you define the scripts to start and stop the client scheduler. You can access the SMIT menu for this task via the command “smitty cm_cfg_app”.

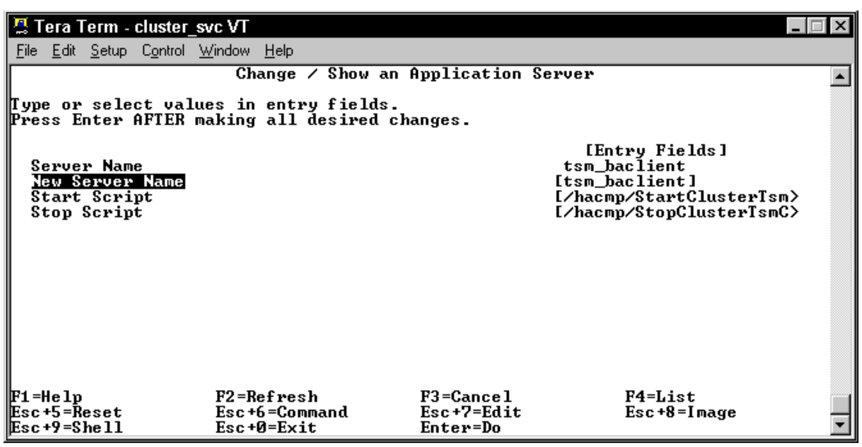

Figure 5-9 Defining a HACMP application server for the B/A client

You will need to configure the following parameters:

Server Name
Enter a logical name for your Application Server. Maximum length is 32 characters.

Start script
Script to start the application server when node joins the cluster. This script needs to be located on a local filesystem that is available when HACMP is not yet running.

Stop script Script to stop the application server when node leaves the cluster gracefully. This script could be located on a shared filesystem, but should probably be stored on a local filesystem with the corresponding start script.

Finally, associate the application server and filesystem resources with the Resource Group you defined earlier. You can access the SMIT menu for this task via the command “smitty cm_cfg_res.select”.

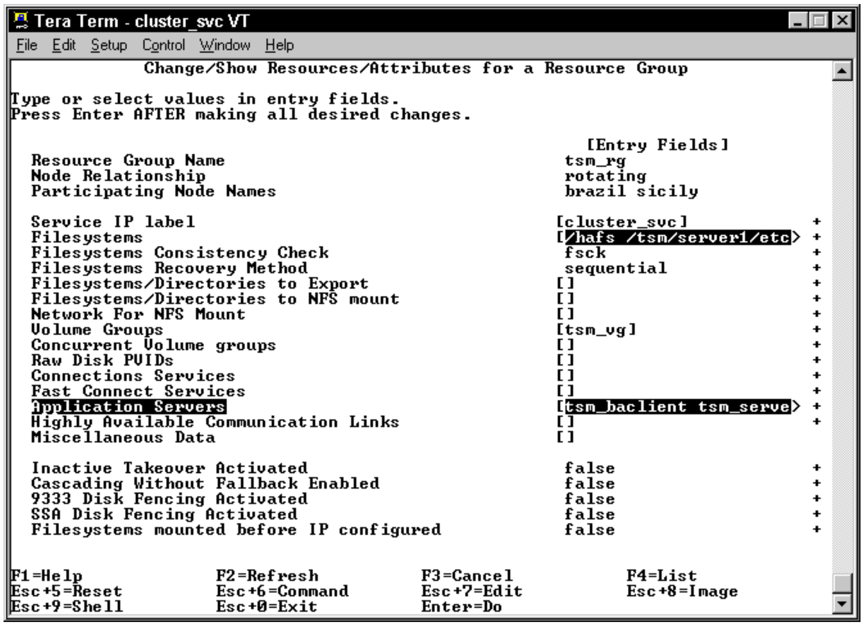

Figure 5-10 Associating HACMP resources and application servers for the B/A client

You will need to configure at least the following parameters:

Service IP label The IP label (hostname) of the adapter, associated with the numeric IP address in the /etc/hosts file.

Filesystems Enter the mount points of the filesystems which are mounted when the resource is initially acquired.

Volume Groups Enter the names of the volume groups containing raw logical volumes or raw volume groups that are varied on when the resource is initially acquired.

Application Servers Enter application servers that will be started. These are the servers defined in the "Define Application Servers" section.

The example shown in Figure 5-10 also includes the resources and application server for the Tivoli Storage Manager server, as defined in 3.1.2, “HACMP support for Tivoli Storage Manager server on AIX” on page 88. In this configuration our Tivoli Storage Manager server and B/A clients are running on the same HACMP cluster. This will probably not be the case in your environment.

Depending on your specific HACMP configuration, there may be other parameters to configure.

Your HACMP environment should now be correctly configured and ready for testing.

Step 8: Testing the failover event

After installing and customizing the B/A client and configuring HACMP, you can test the HACMP failover and interaction with the Tivoli Storage Manager client scheduler.

Before starting the HACMP cluster for the first time, ensure you stop any services or resources normally managed by HACMP that you may have manually started during setup and installation.

Example 5-18 Stopping manually started HACMP resources

```
root@sicily:/hacmp: ps -ef | grep dsm | grep -v grep
    root 269402      1   0   Mar 18      -  0:01 /usr/tivoli/tsm/client/ba/bin/dsmcad
    root 270074      1   0   Mar 18      -  0:01 /usr/tivoli/tsm/client/ba/bin/dsmcad
-optfile=/usr/tivoli/tsm/client/ba/bin/dsm.opt

root@sicily:/hacmp: kill 269402

root@sicily:/hacmp: mount
  node       mounted        mounted over    vfs       date        options
-------- ---------------  ---------------  ------ ------------ ---------------
         /dev/hd4         /                jfs    Mar 12 13:24 rw,log=/dev/hd8
         /dev/hd2         /usr             jfs    Mar 12 13:24 rw,log=/dev/hd8
         /dev/hd9var      /var             jfs    Mar 12 13:24 rw,log=/dev/hd8
         /dev/hd3         /tmp             jfs    Mar 12 13:24 rw,log=/dev/hd8
         /dev/hd1         /home            jfs    Mar 12 13:25 rw,log=/dev/hd8
         /dev/lv00        /hacmp           jfs    Mar 12 13:25 rw,log=/dev/hd8
         /dev/lv01        /hafs            jfs    Mar 14 11:58 rw,log=/dev/loglv00

root@sicily:/hacmp: umount /hafs
```

Notice that we leave the local client scheduler running as it is not under the control of HACMP.

Repeat this check on each node in your cluster.

Start the HACMP cluster

First, start the HACMP software on all nodes in the cluster. You can access the SMIT menu for this task via the command “smitty cl_clstart.dialog

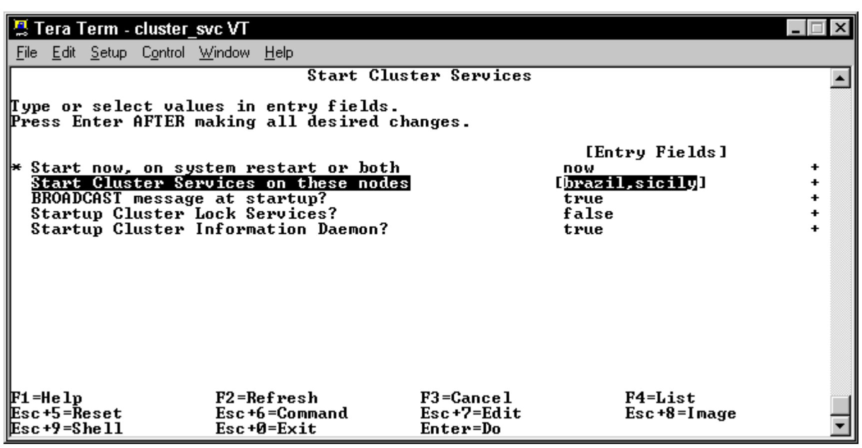

Figure 5-11 Starting the HACMP cluster on all nodes

You will need to configure the following parameters:

Start now, on system restart, or both Choosing 'now' will start the daemons immediately. Select this option for our testing purposes.

Choosing 'system restart' will start the daemons after a system reboot by adding an entry to the /etc/inittab file.

Choosing 'both' will start the daemons immediately AND after a system reboot.

Start Cluster Services on these nodes Identify the nodes on which to start cluster services.

Startup Cluster Information Daemon Starts the Cluster Information daemon. This daemon needs to be started to support the B/A client.

This will cause the service IP address/s to become active, and you will lose any connections to the nodes boot IP address/s.

After the HACMP software has started, the applications and filesystems associated with your Tivoli Storage Manager Resource Group will be available on one of your cluster nodes. Use the following examples to test for the correct startup of the Tivoli Storage Manager client acceptor daemon (which is responsible for the client scheduler and Web client interface).

Example 5-19 Testing for the startup of the client acceptor after HACMP startup

```
root@sicily:/hacmp: lssrc -g cluster
Subsystem         Group            PID     Status
 clstrmgr         cluster          5918    active
 clsmuxpd         cluster          16784   active
 clinfo           cluster          20124   active

root@sicily:/hacmp: ps -ef | grep dsmcad | grep -v grep
    root 270074      1   0   Mar 18      -  0:01 /usr/tivoli/tsm/client/ba/bin/dsmcad
-optfile=/usr/tivoli/tsm/client/ba/bin/dsm.opt
    root 270390      1   0 18:10:00      -  0:00 /usr/tivoli/tsm/client/ba/bin/dsmcad

root@sicily:/hacmp: ls -l /hafs/baclient/
total 32
-rw-------   1 root     sys          132 Mar 14 18:07 TSM.PWD
-rw-r--r--   1 root     sys            0 Mar 14 18:07 dsm.inx
-rw-r--r--   1 root     sys          100 Mar 14 15:43 dsm.opt
-rw-r--r--   1 root     sys            7 Mar 14 18:10 hacmpcad.pids
-rw-r--r--   1 root     sys            0 Mar 14 18:10 hacmpcad.pids1
-rw-r--r--   1 root     sys            7 Mar 14 18:10 hacmpcad.pids2

root@sicily:/hacmp: cat /hafs/baclient/hacmpcad.pids
270030

root@sicily:/hacmp: df -k
Filesystem    1024-blocks      Free %Used    Iused %Iused Mounted on
/dev/hd4            16384      8060   51%     1415    18% /
/dev/hd2           655360    107772   84%    18386    12% /usr
/dev/hd9var         16384     12732   23%      158     4% /var
/dev/hd3           131072     77448   41%       62     1% /tmp
/dev/hd1           229376     86884   63%       37     1% /home
/dev/lv00           16384     15776    4%       27     1% /hacmp
/dev/lv01         2097152   2031256    4%       24     1% /hafs
```

Stop the HACMP cluster

Next, stop the HACMP software on whichever node has the Tivoli Storage Manager Resource Group. This should cause the Resource Group to failover to another node. You can access the SMIT menu for this task via the command “smitty clstop.dialog”.

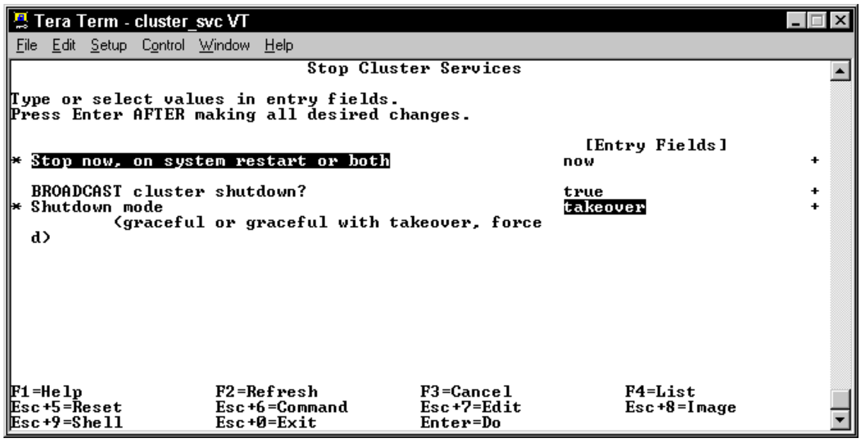

Figure 5-12 Controlling the HACMP cluster to test B/A client failover

This will cause the service IP address/s to move to another node, and you will lose any connections to the current node.

After the HACMP failover has completed, the applications and filesystems associated with your Tivoli Storage Manager Resource Group will be available on the next cluster node. Use the examples shown in Example 5-19 to test for the correct startup of the Tivoli Storage Manager client acceptor (scheduler) on the new node.

Test filespace ownership

After running a scheduled backup you should query the Tivoli Storage Manager server to check that the filespaces representing the local and shared filesystems are owned by the correct client node accounts. Use the command shown in Example 5-20 to check for filesystem ownership:

Example 5-20 Checking for filespace ownership in a HACMP cluster

```
root@brazil:/hacmp: dsmadmc -id=admin -pa=admin "q files"
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM_SERVER1: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/16/02   15:40:47  Last access: 03/16/02   15:40:24

ANS8000I Server command: 'q files'

Node Name        Filespace   FSID Platform Filespace Is Files- Capacity   Pct
                 Name                      Type        pace        (MB)  Util
                                                     Unicode?
---------------  ----------- ---- -------- --------- --------- -------- -----
BRAZIL_FS        /              1 AIX      JFS          No         16.0  49.6
BRAZIL_FS        /usr           2 AIX      JFS          No        640.0  83.6
BRAZIL_FS        /var           3 AIX      JFS          No         16.0  13.8
BRAZIL_FS        /home          4 AIX      JFS          No         16.0   3.4
BRAZIL_FS        /hacmp         6 AIX      JFS          No         16.0   3.5
CLUSTER1         /hafs          2 AIX      JFS          No      2,048.0   3.1
```

```
SICILY_FS       /               1 AIX       JFS           No          16.0  52.5
SICILY_FS       /usr            2 AIX       JFS           No         640.0  83.6
SICILY_FS       /var            3 AIX       JFS           No          16.0  23.2
SICILY_FS       /home           4 AIX       JFS           No         224.0  62.1
SICILY_FS       /hacmp          5 AIX       JFS           No          16.0   3.7

ANS8002I Highest return code was 0.
```

After a failover event the scheduled backup should be run again to check that the shared filesystem is still sending data to the cluster node account.

An alternative configuration for the B/A client in a HACMP cluster

In some situations it may be preferable to exclude the B/A client scheduler from the control of HACMP. In this configuration the scheduler is started on each node at boot time, and continues to run regardless of the status of the HACMP services or whichever node currently owns the Tivoli Storage Manager Resource Group.

The advantages of running the B/A client scheduler in this mode are:

- There is no need to use an explicit DOMAIN setting in the dsm.sys file (allowing new filesystems to be included in the backup by default)
- You can use either client polling or server prompted scheduling modes
- There is no need to stop/start the B/A client scheduler during HACMP failover events
- Much simplified configuration

The limitations of running the B/A client scheduler in this mode are:

- Requires a “reliable” network interface, not under the control of HACMP
- All backup tasks need to be performed via a customized backup script to ensure the correct client account destination on the Tivoli Storage Manager server
- Every node in the cluster requires the same Tivoli Storage Manager client account password

You should be familiar with the normal HACMP configuration for the B/A client before considering using the alternate configuration.

Step 1: Install the B/A client

Install the B/A client on each node in the cluster, exactly as you would for a stand-alone node. Refer to Section , “Step 1: Install the B/A client code” on page 182 for further details.

Step 2: Configure the B/A client

Create a client system options file and specify two stanzas to identify which Tivoli Storage Manager server to connect to. Both stanzas should connect to the same Tivoli Storage Manager server, but will identify themselves with different nodenames. One stanza will be used for local (rootvg) filesystems, and the other stanza will be used for HACMP shared (non-rootvg) filesystems.

The system options file should be located on the local filesystem. The default location is /usr/tivoli/tsm/client/ba/bin/dsm.sys

Example 5-21 dsm.sys file for alternate HACMP filesystem backup method

```
**************************************************************************
SErvername           tsm_local
COMMmethod           TCPip
```

```
TCPPort                 1500
TCPServeraddress        tsmserver

nodename                sicily_fs

inclexcl                /usr/tivoli/tsm/client/ba/bin/dsm.inx

passwordaccess          generate
schedmode               prompted

************************************************************************
SErvername              tsm_shared
COMMmethod              TCPip
TCPPort                 1500
TCPServeraddress        tsmserver

nodename                cluster1

inclexcl                /usr/tivoli/tsm/client/ba/bin/dsm.inx

passwordaccess          generate
schedmode               prompted
```

Create a client options file to specify the Tivoli Storage Manager servername to use by default. You may configure other parameters if you desire.

The client options file should be located on the local filesystem. The default location is /usr/tivoli/tsm/client/ba/bin/dsm.opt

Example 5-22 dsm.opt file for alternate HACMP filesystem backup method

```
SErvername              tsm_local
```

After updating the B/A client configuration files, manually run the "dsmc" command twice on each node. The first time identifying yourself as the local node, and the second time identifying yourself as the cluster node. This will cause the B/A client to cache the client password in a file on the local filesystem.

Note: With this alternate configuration, every node in the cluster requires the same Tivoli Storage Manager client account password.

In the alternate configuration we want to start the B/A client scheduler on each node at boot time, and leave it running regardless of the state of the HACMP services. Use the following command to add the B/A client scheduler to the AIX /etc/inittab file. This will cause the client scheduler to run immediately, start at boot time, and restart automatically if it fails or is killed.

Example 5-23 Adding the B/A client scheduler to /etc/inittab

```
root@sicily:/hacmp: mkitab "dsmcad:2:respawn:/usr/tivoli/tsm/client/ba/bin/dsmc sched
-quiet >/dev/null 2>&1"

root@sicily:/hacmp: ps -ef | grep dsmc | grep -v grep
root 268564 1 0 13:05:47 -  0:00 /usr/tivoli/tsm/client/ba/bin/dsmc sched -quiet
```

Remember to restart (in this case, kill and automatically restart) the dsmc scheduler after any changes to the B/A client configuration files.

Step 3: Create a HACMP-aware backup script

We need to create a script to control which parameters to use with the local and shared filesystems during a scheduled backup.

The backup script should be located on the local filesystem. In our example we saved the script to /usr/tivoli/tsm/scripts/backup_hacmp.sh.

Example 5-24 Script for alternate HACMP filesystem backup method

```
#!/bin/ksh

case $1 in
        local)  VG_LIST="rootvg"
                SERVER="tsm_local"
                ;;
        shared) VG_LIST=`lsvg -o | grep -v rootvg`
                SERVER="tsm_shared"
                ;;
        both)   backup_hacmp.sh local
                backup_hacmp.sh shared
                ;;
        *)      echo "Usage: backup_hacmp.sh < local | shared | both >"
                ;;
esac

for VG in $VG_LIST
do
        LV_LIST=`echo $VG | lsvg -i -l | grep "jfs[2 ] " | grep open | grep -v "/tmp" |
 awk '{print $7}'`

        for FS in $LV_LIST
        do
                dsmc i $FS -se=$SERVER
        done
done
```

It may be necessary to modify the script to suit your local cluster configuration. For example, if your cluster nodes contain local, non-rootvg filesystems, you will need to include them in the local VG_LIST variable.

> **Note:** Running a manual “dsmc” backup or restore command will only connect to the default servername (as specified in the dsm.opt file) and will be default incorrectly include any currently mounted shared filesystems.
>
> You will need to specifically identify the cluster stanza (with -se=tsm_cluster) when manually performing backup or restore of shared filesystems.

Step 4: Test the B/A client backup script

When executed the “backup_hacmp.sh” script will perform one “dsmc” command for each mounted JFS filesystem. By setting which Tivoli Storage Manager servername to use, the script can control which stanza is used from the B/A client configuration files, and indirectly, which nodename to identify ourselves as to the Tivoli Storage Manager server. This will cause local (rootvg) filesystems to be sent to the local node account, and shared (non-rootvg) filesystems to be sent to the cluster node account.

Example 5-25 Sample activity for the alternate HACMP backup script

```
root@sicily:/usr/tivoli/tsm/scripts: ./backup_hacmp.sh rootvg
dsmc i / -se=sicily_fs
dsmc i /usr -se=sicily_fs
dsmc i /var -se=sicily_fs
dsmc i /home -se=sicily_fs
dsmc i /hacmp -se=sicily_fs

root@sicily:/usr/tivoli/tsm/scripts: ./backup_hacmp.sh shared
dsmc i /hafs -se=cluster1_fs
```

The backup_hacmp.sh script can be safely run regardless of whether the HACMP services are running. If a shared (non-rootvg) filesystem is mounted, it will be backed up to the cluster node account regardless of which node currently has the filesystem mounted. If no shared filesystems are mounted, then the script will exit without performing any action.

Step 5: Create a client backup schedule

To automate the backup of the local and shared filesystems we need to create a Tivoli Storage Manager client schedule. In the alternative method we wish to perform a “command” schedule, not an “incremental” schedule. Use the “Objects” field to identify the “backup_hacmp.sh” script to execute. Remember to make the backup script executable.

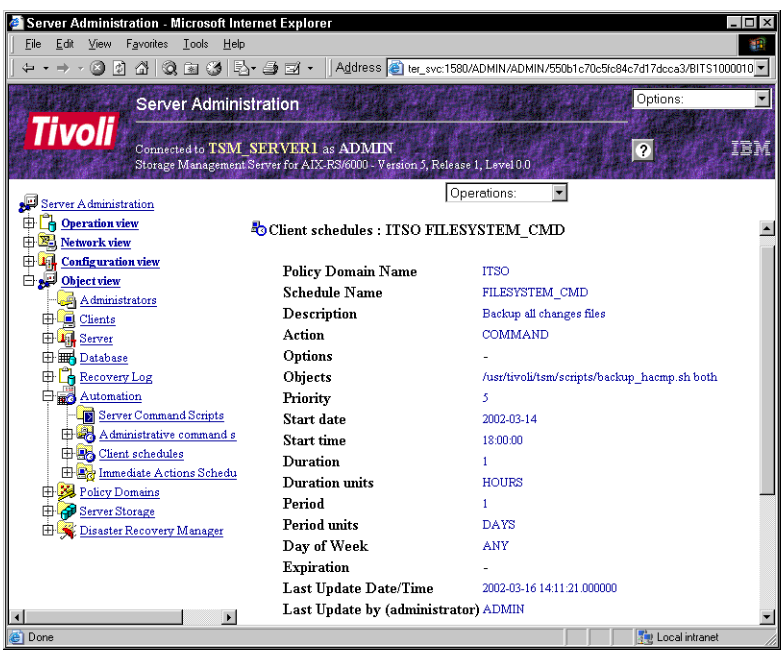

Figure 5-13 Web admin client - command schedule

In the alternate method we only need to associate the command schedule with the “real” node names. We do not need to associate the schedule to the cluster node name. Backup data will still be sent to the cluster node name, but only as a result of the script selecting which client option server stanza to use.

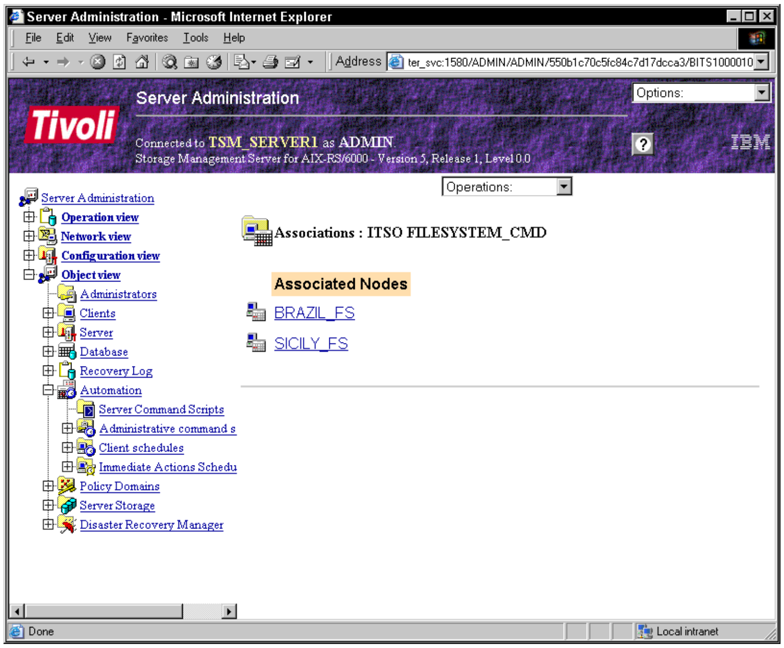

Figure 5-14 Web admin client - cluster node

After running the backup_hacmp.sh script on both nodes you should query the Tivoli Storage Manager server to check that the cluster node filesystems are owned by the correct client accounts. Use the command shown in Example 5-26 to check for filesystem ownership.

Example 5-26 Checking for filespace ownership in a HACMP cluster

```
root@brazil:/hacmp: dsmadmc -id=admin -pa=admin "q files"
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM_SERVER1: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/16/02   15:40:47  Last access: 03/16/02   15:40:24

ANS8000I Server command: 'q files'

Node Name        Filespace   FSID Platform Filespace Is Files- Capacity   Pct
                 Name                      Type         pace       (MB)  Util
                                                     Unicode?
---------------- ----------- ---- -------- --------- --------- -------- -----
BRAZIL_FS        /              1 AIX      JFS          No         16.0  49.6
BRAZIL_FS        /usr           2 AIX      JFS          No        640.0  83.6
BRAZIL_FS        /var           3 AIX      JFS          No         16.0  13.8
BRAZIL_FS        /home          4 AIX      JFS          No         16.0   3.4
BRAZIL_FS        /hacmp         6 AIX      JFS          No         16.0   3.5
CLUSTER1         /hafs          2 AIX      JFS          No      2,048.0   3.1
SICILY_FS        /              1 AIX      JFS          No         16.0  52.5
SICILY_FS        /usr           2 AIX      JFS          No        640.0  83.6
SICILY_FS        /var           3 AIX      JFS          No         16.0  23.2
```

```
SICILY_FS       /home          4 AIX      JFS          No        224.0   62.1
SICILY_FS       /hacmp         5 AIX      JFS          No         16.0    3.7

ANS8002I Highest return code was 0.
```

After a failover event the script should be run again to check that the shared filesystem is still sending data to the cluster node account.

5.1.3 HACMP support for the HSM client on AIX

The objective of HACMP support for the HSM client on AIX is to provide support for HACMP failover on AIX so that HSM managed filesystems can continue to operate in the case of an HACMP node failover and fallback. Migration and Recollection of both local and shared filesystems will be processed correctly regardless of which cluster node currently has the filesystems mounted.

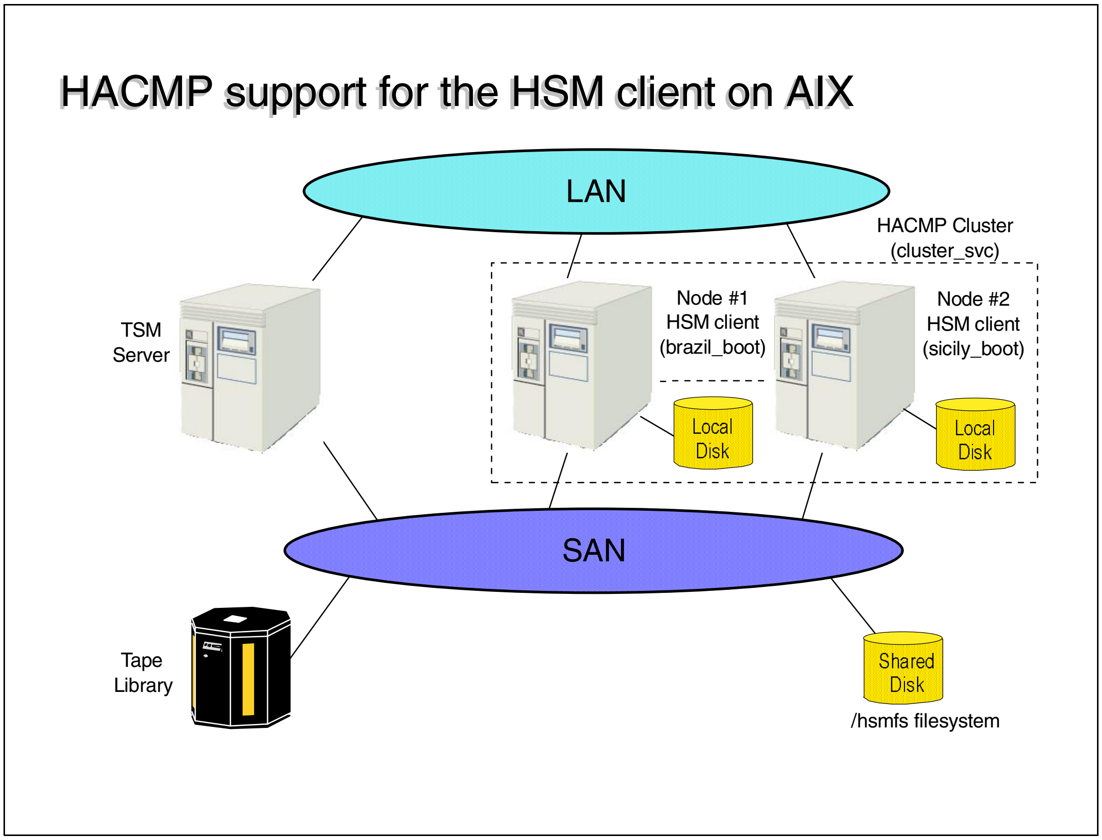

Figure 5-15 HACMP support

Tivoli Storage Manager supports a wide variety of configurations as well as the following takeover relationships in an HACMP environment with only one node controlling the shared resources at a time.

- Cascading (with or without fallback)
- Rotating

There is no support for concurrently mounted volumes or resource groups.

Prerequisite tasks and supported environment

Configuring the Tivoli Storage Manager HSM client to operate in a HACMP cluster requires understanding of both Tivoli Storage Manager and HACMP configuration.

The user should be familiar with the following:

- Tivoli Storage Manager documentation:

  "Configuring the HSM Client in a HACMP Takeover Environment" in *Tivoli Space Manager for UNIX: Using the Hierarchial Storage Management Clients*, GC32-0794
- Tivoli Storage Manager options file and options
- AIX file systems
- The user should already have a functioning HACMP cluster, able to Start / Stop / Takeover / Take back an HACMP environment
- The user should already have a functioning Tivoli Storage Manager HSM configuration.

The following hardware/software configurations are required for HACMP HSM support:

- AIX 4.3.3 or later, must be identical on all nodes
- HACMP 4.4 or later (or) HACMP/ES 4.4 or later
- HACMP Cluster Info Daemon must be running
- An external hard disk on both systems for shared access
- Two network adapters per cluster node
- Two identical RS6000 systems
- There are no Tivoli Storage Manager server dependencies

In our environment we used two different RS6000 systems with a single network adapter in each node. This is not the ideal configuration and should not be considered for a production environment. It was however suitable to demonstrate the new function.

Limitations

The following limitations refer specifically to HACMP HSM support:

- The HSM support for the HACMP environment was only tested on a two node cluster.
- The same HSM nodename has to be used for both systems.
- The management class information is only shown for the default migration server.
- The server options information is only shown for the default migration server.
- The HSM support for the HACMP Environment is not seamlessly integrated with the Backup/Archive client support for the HACMP Environment.
- The HSM GUI does not support the HACMP commands.
- Every HACMP-HSM machine has to run the same HSM version!

As is normal with the Tivoli Storage Manager HSM client, you cannot add space management to your root (/), /tmp, /usr, or /var file systems.

Updated HSM command - dsmmigfs

The "dsmmigfs" command has gained two new parameters to assist with moving a HSM managed JFS filesystem between nodes in a HACMP cluster.

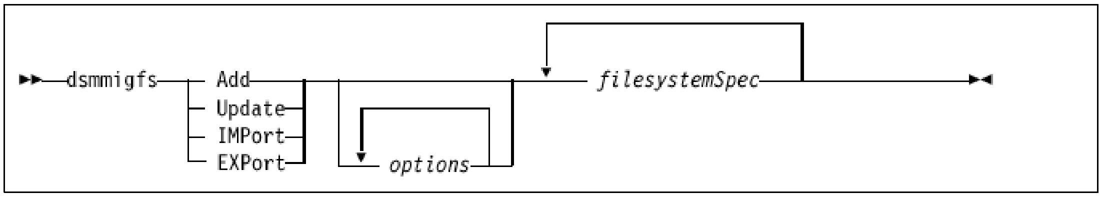

Figure 5-16 The dsmmigfs command syntax

The parameters relevant to our use in a HACMP environment are as follows:

add [filesystem] -t__ -l__
Adds space management to your file systems. This parameter is not valid in an AIX HACMP environment. It should only be used once to initialize the .SpaceMan directory.

-t {value} sets the high migration threshold
-l {value} sets the low migration threshold

update [filesystem] -t__ -l__
Updates space management parameters for your filesystems.

-t {value} sets the high migration threshold
-l {value} sets the low migration threshold

import [filesystem]
Import imports a filesystem to HSM. The import will over-mount the "fsm" filesystem and add an entry to /etc/tsm/SpaceMan/config/dsmmigfstab. A link is also created for the status file in /etc/tsm/SpaceMan/status.

After an import, a reconcile has to be run, so as to update the status files.

export [filesystem]
Exports the filesystem from HSM. The export unmounts the “fsm” filesystem and removes the entry from /etc/tsm/SpaceMan/config/dsmmigfstab. The link is removed from /etc/tsm/SpaceMan/status.

- This does not remove the .SpaceMan directory.
- A reconcile is not performed on the exported filesystem.

Step 0: Before we begin

Since the B/A client is generally configured alongside the HSM client, and having a backup of the data is generally required prior to space migration, much of the configuration of the HSM client is derived from configuring the B/A client in HACMP.

To prevent duplicating much of the configuration detail here, you should refer to 5.1.2, “HACMP support for the B/A client on AIX” on page 181. The following instruction only identify steps unique to using the HSM client in HACMP.

Step 1: Create the client node accounts

Skip to Step 2 if you have already created the client accounts as for the B/A client in HACMP.

To use the HSM client on a HACMP shared filesystem we need to define several client node accounts on the Tivoli Storage Manager server.

If you plan to backup local (per node) filesystems, and backup and migrate shared (HACMP) filesystems, you will require:

- A client account related to each node in the cluster (in our example: brazil_fs, siciliy_fs)
- A client node related to each resource group in the cluster (in our example: cluster1)

Start an administrative session and define the new client node accounts. You could use a command similar to that shown in Example 5-2 on page 183.

Example 5-27 Registering a client node account

```
register node newnode newpassword userid=none domain=mydomain url=http://newnode:1581
```

There are no other special considerations regarding defining client node accounts for use in a HAMCP environment.

Step 2: Tivoli Storage Manager server policy configuration

Skip to Step 3 if you have already configured your server policy for Space Management

Before you can migrate HSM data to the Tivoli Storage Manager server you must configure the Space Management parameters in whichever Management Class you intend to use.

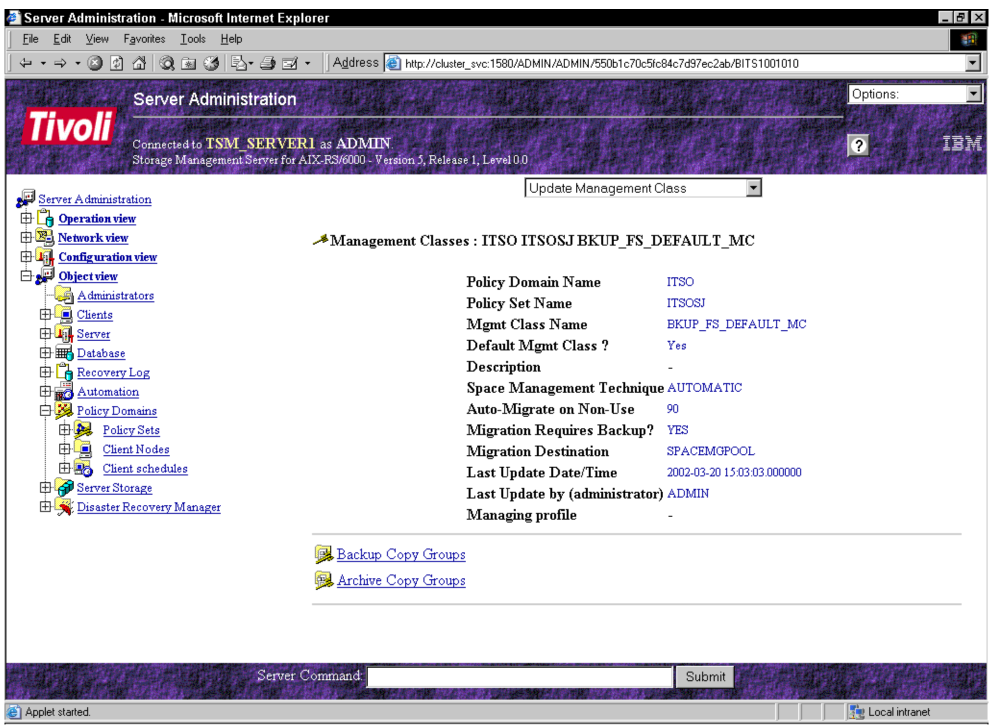

Figure 5-17 Setting the management class parameters for space management

Remember to activate the Policy Set after making any changes.

Step 3: Create a shared filesystem for HSM

Skip to step 4 if you are adding HSM support to an existing filesystem.

Before using the HSM client we need to define a shared filesystems to use for our demonstration.

Remember to use the correct HACMP tools to define any shared filesystems and logical volumes. You can access the SMIT menu for this task via the command “smitty cl_lvm”.

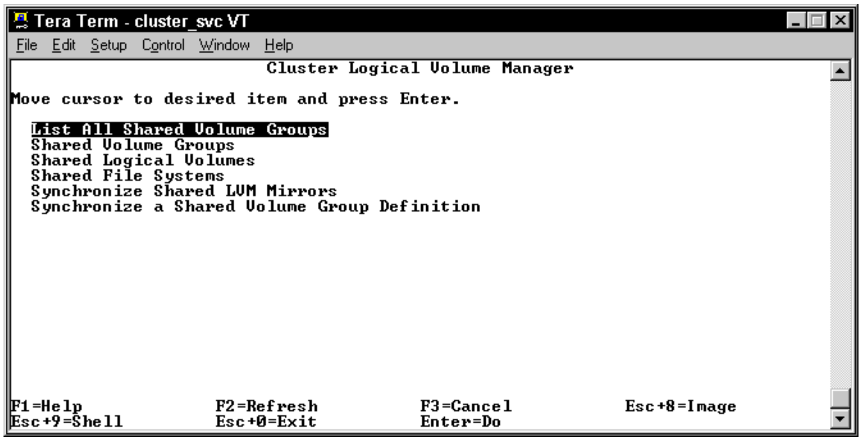

Figure 5-18 Defining shared filesystem via the smitty HACMP menu

Earlier, when we configured the B/A client in HACMP (refer to 5.1.2, “HACMP support for the B/A client on AIX” on page 181), we created a local filesystem (/hacmp) in the rootvg volume group to hold our HACMP startup scripts, and a shared filesystem (/hafs) in the tsm_vg volume group to demonstrate the B/A client.

Now we need to create a new shared filesystem to use in our HSM demonstration. In our example we created a shared filesystem (/hsmfs), also in the tsm_vg volume group.

Example 5-28 Filesystems defined for the HSM client in HACMP

```
root@brazil:/hacmp: lsvg -o | lsvg -i -l

tsm_vg:
LV NAME             TYPE       LPs   PPs   PVs  LV STATE      MOUNT POINT
loglv00             jfslog     1     1     1    open/syncd    N/A
lv01                jfs        32    32    1    open/syncd    /hafs
tsm_hsm             jfs        8     8     1    open/syncd    /hsmfs

rootvg:
LV NAME             TYPE       LPs   PPs   PVs  LV STATE      MOUNT POINT
hd5                 boot       1     1     1    closed/syncd  N/A
hd6                 paging     34    34    1    open/syncd    N/A
hd8                 jfslog     1     1     1    open/syncd    N/A
hd4                 jfs        1     1     1    open/syncd    /
hd2                 jfs        40    40    1    open/syncd    /usr
hd9var              jfs        1     1     1    open/syncd    /var
hd3                 jfs        8     8     1    open/syncd    /tmp
hd1                 jfs        1     1     1    open/syncd    /home
paging00            paging     30    30    1    open/syncd    N/A
lv00                jfs        1     1     1    open/syncd    /hacmp
```

Manually mount the shared filesystems to allow us to configure the HSM client before starting the HACMP services.

Example 5-29 Filesystems mounted for HSM configuration

```
root@brazil:/hacmp: mount
node       mounted        mounted over    vfs       date         options
```

```
-------- ---------------  ---------------  ------ ------------ ---------------
         /dev/hd4         /                jfs    Mar 19 16:05 rw,log=/dev/hd8
         /dev/hd2         /usr             jfs    Mar 19 16:05 rw,log=/dev/hd8
         /dev/hd9var      /var             jfs    Mar 19 16:05 rw,log=/dev/hd8
         /dev/hd3         /tmp             jfs    Mar 19 16:05 rw,log=/dev/hd8
         /dev/hd1         /home            jfs    Mar 19 16:06 rw,log=/dev/hd8
         /dev/lv00        /hacmp           jfs    Mar 19 16:06 rw,log=/dev/hd8
         /dev/lv01        /hafs            jfs    Mar 19 17:09 rw,log=/dev/loglv00
         /dev/tsm_hsm     /hsmfs           jfs    Mar 20 14:53 rw,log=/dev/loglv00
```

Remember to unmount any manually mounted shared filesystems, and vary off the tsm_vg volume group, before attempting to start the HACMP services.

Step 4: Install the HSM client code

Skip to Step 5 if you already have the HSM client code installed and configured.

In addition to the B/A client filesets installed in 5.1.2, “HACMP support for the B/A client on AIX” on page 181, you must install the following:

Example 5-30 Filesets required to install the HSM client on AIX

```
tivoli.tsm.client.hsm.jfs.aix43.32bit
```

If you are upgrading over an existing installation, ensure you stop any currently running Tivoli Storage Manager services before performing the upgrade. For example, kill any running HSM recall or monitor processes, client scheduler or client acceptor tasks.

Note: The installation of the HSM fileset will add an entry to the nodes /etc/inittab and immediately start the HSM daemons. You should leave this entry in the /etc/inittab file so that it can load the HSM kernel extension at system boot time.

Later, we shall configure HACMP to stop and restart the HSM daemons using the correct options files.

Repeat this set for each node in the cluster. The same level of HSM code must be running on all the cluster nodes.

Step 5: Configure the HSM client

The HACMP support for the B/A and HSM client allows for the backup both local and shared filesystems, and for the migration of shared filesystems only. This is achieved by configuring multiple server stanzas in the dsm.sys file, and stating the various daemons with the appropriate dsm.opt configuration file.

It is necessary to configure the B/A client dsm.sys file on each node in the cluster. The dsm.sys file should contain at least two server stanzas. The different stanzas are required for:

- One server stanza for backup of local filesystems
- One server stanza for backup and migration of each cluster resource group/filesystem

It is also necessary to configure one version of the B/A client dsm.opt file on each node and each resource group in the cluster. The different dsm.opt files will each control:

- One dsm.opt file on each node to backup local (non-clustered) filesystems
- One dsm.opt file on each resource group’s shared filesystem its own backup

Configure dsm.sys for shared filesystem migration

The dsm.sys file used for local and shared filesystem backup needs to be modified slightly to work with the HSM client daemons. These changes must be replicated on every node in the cluster.

Modify the dsm.sys file we created for the B/A client, adding another server stanza and some HSM specific entries at the top. The extra server stanza is required because:

- The B/A client and HSM client encrypt the client password in different ways, and therefore cannot share the same TSM.PWD password cache file.
- The HSM client does not honour the “clusternode” parameter, and therefore requires the “nodename” parameter to be set.

Refer to Example 5-6 on page 185 for a typical dsm.sys configuration supporting HSM in a HACMP environment. You need to modify some parameters to suit you local environment.

Example 5-31 Example dsm.sys file for the HSM client in a HACMP environment

```
DEFAULTServer         tsm_hsm
MIGRATEServer         tsm_hsm
CHECKThresholds       2
CANDIDATESInterval    12
RECOncileinterval     12
MAXRECOncileproc      5
MAXThresholdproc      5
MINMIGFILESize        8192
MIGFILEEXPiration     10
MINRECAlldaemons      5
MAXRecalldaemons      15
CHECKFororphans       no
MAXMIGRators          1
KERNelmessages        no

************************************************************************
* Local filesystem backup

SErvername            tsm_local
COMMmethod            TCPip
TCPPort               1500
TCPServeraddress      tsmserver

nodename              brazil_fs
*TCPClientaddress     brazil_boot
HTTPport              1581

domain                / /usr /var /home /hacmp
inclexcl              /usr/tivoli/tsm/client/ba/bin/dsm.inx

passwordaccess        generate

managedservices       schedule webclient
schedmode             polling

errorlogname          /var/dsmerror.log
schedlogname          /var/dsmsched.log

************************************************************************
* Shared filesystem backup

SErvername            tsm_shared
```

```
COMMmethod            TCPip
TCPPort               1500
TCPServeraddress      tsmserver

clusternode           yes
*TCPClientaddress     cluster_svc
HTTPport              1582

domain                /hafs /hsmfs
inclexcl              /hafs/baclient/dsm.inx

passwordaccess        generate
passworddir           /hafs/baclient

managedservices       schedule webclient
schedmode             polling

errorlogname          /hafs/baclient/dsmerror.log
schedlogname          /hafs/baclient/dsmsched.log

************************************************************************
* Shared filesystem migration

SErvername            tsm_hsm
COMMmethod            TCPip
TCPPort               1500
TCPServeraddress      tsmserver

nodename              cluster1

domain                /hafs /hsmfs
inclexcl              /hafs/hsmclient/dsm.inx

passwordaccess        generate
passworddir           /hafs/hsmclient

errorlogname          /hafs/baclient/dsmerror.log
schedlogname          /hafs/baclient/dsmsched.log
```

You should set the “passworddir” for the shared migration stanza to save the HSM client password to somewhere on the shared filesystem. You will need to create this directory manually.

Repeat this step for every node in your cluster. The server stanzas used for shared filespace backup and migration should be identical between all the nodes.

Configure dsm.opt for local filesystem backup

The dsm.opt file used to control local filesystem backup should remain unchanged from the configuration used when setting up the B/A client in a HACMP environment.

Refer to “Configure dsm.opt for local filesystem backup” on page 211 for details on how to configure this file.

Configure dsm.opt for shared filesystem migration

The dsm.opt file used to control shared filesystem migration should be a separate file from that used for shared filesystem backup. It should be placed on the shared filesystem itself. This allows us to maintain only a single instance of this configuration file.

Copy the sample /usr/tivoli/tsm/client/ba/bin/dsm.opt.smp file to a directory in the shared filesystem (in our example /hafs/hsmclient/dsm.opt) and edit it to resemble Example 5-8 on page 187. You may add other parameters to suit you local environment.

Example 5-32 Example dsm.opt file for the shared filesystem backup in a HACMP environment

```
SErvername            tsm_hsm
```

Replace the servername parameter with name of the stanza used in your dsm.sys file for shared filespace migration.

Manually add the filesystem to HSM

Before attempting to start configure HACMP we need to manually identify the filesystem to HSM. This will create the ".SpaceMan" directory in the filesystem root, add and entry to the "dsmmigfstab" files, and cache the HSM client password for the shared filesystem migration.

To identify a new filesystem to HSM, enter the following command:

Example 5-33 Adding a filesystem to HSM

```
root@sicily:/hacmp: cat /etc/adsm/SpaceMan/config/dsmmigfstab
# Filesystem  High(%)  Low(%)   Premig(%)  Age    Size   Quota  stubsize Server
# Name        Thrshld  Thrshld  Percent    Factor Factor                 Name
# ---------------------------------------------------------------------------------

root@sicily:/hacmp: ls -l /hsmfs
total 16
drwxrwx---   2 root     system       512 Mar 20 10:18 lost+found
drwxr-sr-x   2 root     sys          512 Mar 26 12:00 test

root@sicily:/hacmp: ls -l /hafs/hsmclient/
total 8
-rw-r--r--   1 root     sys            0 Mar 26 12:58 dsm.inx
-rw-r--r--   1 root     sys           96 Mar 26 12:59 dsm.opt

root@sicily:/hacmp: dsmmigfs add /hsmfs
Tivoli Storage Manager
Space Management Interface - Version 5, Release 1, Level 0.0 g4
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Please enter password for node "CLUSTER1":

ANS9309I Mount FSM: TSM space management mounted on /hsmfs
ANS9087I Space management is successfully added to file system /hsmfs.

root@sicily:/hacmp: cat /etc/adsm/SpaceMan/config/dsmmigfstab
# Filesystem  High(%)  Low(%)   Premig(%)  Age    Size   Quota  stubsize Server
# Name        Thrshld  Thrshld  Percent    Factor Factor                 Name
# ---------------------------------------------------------------------------------
/hsmfs  90      80      -       -       -       -       -       -       -

root@sicily:/hacmp: ls -l /hsmfs
total 24
drwxrwsr-x   3 bin      bin          512 Mar 26 15:07 .SpaceMan
drwxrwx---   2 root     system       512 Mar 20 10:18 lost+found
drwxr-sr-x   2 root     sys          512 Mar 26 12:00 test

root@sicily:/hacmp: ls -l /hafs/hsmclient/
total 16
```

```
-rw-------   1 root     sys           127 Mar 26 15:07 TSM.PWD
-rw-r--r--   1 root     sys             0 Mar 26 12:58 dsm.inx
-rw-r--r--   1 root     sys            96 Mar 26 12:59 dsm.opt
```

To update the migration parameters at any time, enter the following command:

Example 5-34 Updating the HSM parameters

```
root@sicily:/hacmp: dsmmigfs update -ht=80 -lt=60 /hsmfs
Tivoli Storage Manager
Space Management Interface - Version 5, Release 1, Level 0.0 g4
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

ANS9068I dsmmigfs: dsmmigfstab file updated for file system /hsmfs.

root@sicily:/hacmp: cat /etc/adsm/SpaceMan/config/dsmmigfstab
# Filesystem  High(%)  Low(%)   Premig(%)  Age    Size   Quota  stubsize Server
# Name        Thrshld  Thrshld  Percent    Factor Factor                 Name
# ------------------------------------------------------------------------------
/hsmfs  80      60      -       -       -       -       -       -       -
```

When a “dsmmigfs import” command is run the filesystem is configured with the default parameters. If you vary the per filesystem migration parameters you will need to modify the “start_HSM.sh” script to reproduce the changes after each failover and filesystem import.

Step 6: Customize the HACMP scripts

We need to create some scripts to start and stop the HSM services under the control of HACMP. The scripts provided here are only samples and may need to be slightly customized to suit your environment.

The *Tivoli Space Manager for UNIX 5.1* manual provides some sample HACMP scripts. The scripts shown in the next few pages of this Redbook are derived from those scripts.

The sample start script

The startup script is used to restart the various HSM daemons and import the HSM filesystems as the resource group comes online. An identical copy of this script needs to exist on each node in the cluster.

Copy the sample shown in Example 5-35 to an area on the local filesystem (in our example /hacmp/start_HSM.sh) and if necessary modify it to suit your local environment.

Example 5-35 Sample script to start HSM under the control of HACMP

```
#!/usr/bin/sh
#-------------------------------------------------------------------
# call with
#    start_HSM [filesystems]
#-------------------------------------------------------------------

LOG="/tmp/HSM.log"

DSM_DIR="/usr/tivoli/tsm/client/hsm/bin"
DSM_CONFIG="/hafs/baclient/dsm.opt"

HSM_DAEMONS="dsmrecalld dsmmonitord dsmscoutd"
```

```
#-----------------------------------------------------------------
# Redirect all output to the log file

exec 2>&1 >> $LOG

echo "****************************************************"
echo "* Starting HSM                                     *"
echo "****************************************************"

echo `date`
echo $DSM_DIR $DSM_CONFIG

#-----------------------------------------------------------------
# kill all running HSM demons

echo "Stopping the HSM processes ..."

for d in $HSM_DAEMONS
do
        pid=`ps -aef|grep $d |grep -v grep |awk '{print \$2}'`

        if [[ $pid != "" ]]; then
                echo "- killing $d ..."
                kill -15 $pid
        fi
done

#-----------------------------------------------------------------
# starting the daemons with the right dsm.sys and dsm.opt

echo "Starting the HSM processes ..."

for d in $HSM_DAEMONS
do
        echo "- starting $d ..."
        $DSM_DIR/$d
done

#-----------------------------------------------------------------
# import all filesystems

echo "Starting to import filesystems ..."

for fs in $*
do
        echo "- importing $fs"

        dsmmigfs import $fs
        dsmreconcile $fs
done

#-----------------------------------------------------------------
# done.

echo "HSM is Started"
```

Duplicate this file to every node in your cluster.

The sample stop script

The stop script is used to kill the various HSM daemons and any user level processes still accessing the HSM filesystem, and export the HSM filesystems when the resource group halts gracefully. An identical copy of this script needs to exist on each node in the cluster.

Copy the sample shown in Example 5-36 to an area on the local filesystem (in our example /hacmp/stop_HSM.sh) and if necessary modify it to suit your local environment.

Example 5-36 Sample script to stop HSM under the control of HACMP

```
#!/usr/bin/sh
#-------------------------------------------------------------------
# call with
#    stop_HSM [filesystems]
#-------------------------------------------------------------------

LOG="/tmp/HSM.log"

DSM_DIR="/usr/tivoli/tsm/client/hsm/bin"
DSM_CONFIG="/hafs/baclient/dsm.opt"

HSM_DAEMONS="dsmrecalld dsmmonitord dsmscoutd"

#-------------------------------------------------------------------
# Redirect all output to the log file

exec 2>&1 >> $LOG

echo "**************************************************"
echo "* Stopping HSM                                   *"
echo "**************************************************"
echo `date`

#-------------------------------------------------------------------
# before you can export the filesystems you have to stop all daemons

echo "Stopping HSM processes ..."

for d in $HSM_DAEMONS
do
        pid=`ps -aef|grep $d |grep -v grep |awk '{print \$2}'`

        if [[ $pid != "" ]]; then
                echo "- killing $d ..."
                kill -15 $pid
        fi
done

#-------------------------------------------------------------------
# to export the filesystems you have to stop all processes using them

for fs in $*
do
        echo "Looking for user processes running on $fs ..."

        for pid in `fuser -c $fs 2>/dev/null`
        do
                echo "- killing `ps -p $pid | grep $pid | awk '{print $4}'` ..."
                kill -9 $pid
        done
```

```
done

#-------------------------------------------------------------------
# now you can export the filesystems

echo "Starting to export filesystems ..."

for fs in $*
do
        echo "- exporting $fs ..."

        dsmmigfs export $fs
done

#-------------------------------------------------------------------
# Done.

echo "HSM is Stopped"
```

Duplicate this file to every node in your cluster.

Note: Do not kill the HSM daemons with the -9 option.

Test the start and stop scripts

It is important to test the function of the start and stop scripts before configuring HACMP to use them during a failover event.

Check that the filesystem to be managed by HSM is mounted, then manually run the start script as shown in Example 5-37. The example also shows the checking of the HSM configuration file for entries related to the managed filesystem.

Example 5-37 Manually running the HSM start script

```
root@brazil:/hacmp: mount | grep hsmfs
         /dev/tsm_hsm     /hsmfs           jfs    Mar 20 14:53 rw,log=/dev/loglv00

root@brazil:/hacmp: cat /etc/adsm/SpaceMan/config/dsmmigfstab
# Filesystem  High(%)  Low(%)   Premig(%)  Age    Size   Quota  stubsize Server
# Name        Thrshld  Thrshld  Percent    Factor Factor                 Name
# -----------------------------------------------------------------------------

root@brazil:/hacmp: ./start_HSM.sh /hsmfs

root@brazil:/hacmp: mount | grep hsmfs
         /dev/tsm_hsm     /hsmfs           jfs    Mar 20 14:53 rw,log=/dev/loglv00
         /hsmfs           /hsmfs           fsm    Mar 25 13:41

root@brazil:/hacmp: cat /etc/adsm/SpaceMan/config/dsmmigfstab
# Filesystem  High(%)  Low(%)   Premig(%)  Age    Size   Quota  stubsize Server
# Name        Thrshld  Thrshld  Percent    Factor Factor                 Name
# -----------------------------------------------------------------------------
/hsmfs  90      80      -       -       -       -       -       -       -
```

When you run the start script, messages similar to those shown in Example 5-38 should appear in the log file (in our example the log file is /tmp/HSM.log).

Example 5-38 Log file entries from starting the HSM services

```
root@brazil:/tmp: tail -f HSM.log
****************************************************
* Starting HSM                                     *
****************************************************
Mon Mar 25 13:41:20 PST 2002
/usr/tivoli/tsm/client/hsm/bin /hsmfs/hsmclient/dsm.opt
Stopping the HSM processes ...
Starting the HSM processes ...
- starting dsmrecalld ...
- starting dsmmonitord ...
- starting dsmscoutd ...
Starting to import filesystems ...
- importing /hsmfs
ANS9309I Mount FSM: TSM space management mounted on /hsmfs
Tivoli Storage Manager
Space Management Interface - Version 5, Release 1, Level 0.0 g4
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

ANS9087I Space management is successfully added to file system /hsmfs.

Reconciling '/hsmfs' file system:
  Reading the current premigration database...
         Read 2 entries
```

Example 5-40 Log file entries from stopping the HSM services

```
**************************************************
* Stopping HSM                                   *
**************************************************
Mon Mar 25 13:44:17 PST 2002
Stopping HSM processes ...
- killing dsmrecalld ...
- killing dsmmonitord ...
- killing dsmscoutd ...
Looking for user processes running on /hsmfs ...
- killing sh ...
Starting to export filesystems ...
- exporting /hsmfs ...
Tivoli Storage Manager
Space Management Interface - Version 5, Release 1, Level 0.0 g4
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

  Unmounting FSM from file system /hsmfs.
ANS9128I dsmmigfs: exported space management file system /hsmfs.
HSM is Stopped
```

The HSM start and stop scripts will need to exist on every node in the HACMP cluster.

Step 7: Configure HACMP

After installing the B/A and HSM client code and customizing the HSM start/stop scripts, you next need to configure HACMP.

The following example assumes you already have a functioning HACMP environment, with correctly configured Networks, Adapters, and Nodes, etc. You should be able to successfully perform a HACMP failover and recovery before continuing. The following examples build on the HACMP environment we configured in 3.1.2, “HACMP support for Tivoli Storage Manager server on AIX” on page 88 for the Tivoli Storage Manager server.

First (if you have not already done so), define a Resource Group to manage the Tivoli Storage Manager client. You can access the SMIT menu for this task via the command “smitty cm_add_res”.

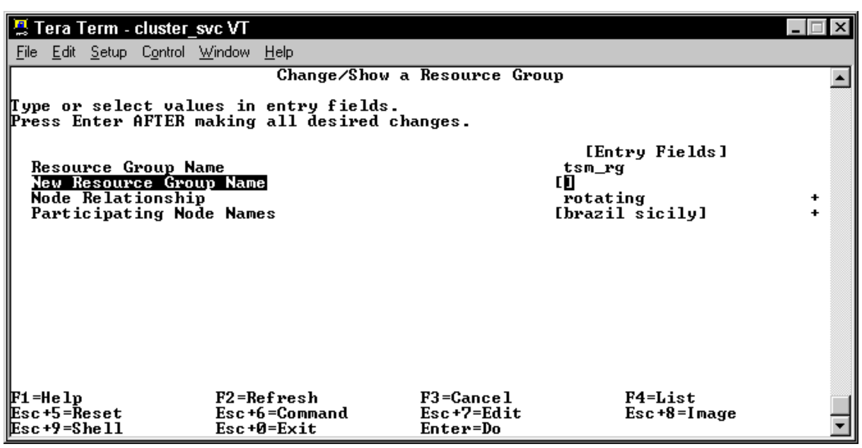

Figure 5-19 Defining a HACMP resource group for the HSM client

You will need to configure the following parameters:

Resource Group Name Enter a logical name for your resource group. Maximum length is 32 characters.

Node Relationship Toggle between Cascading/Concurrent/Rotating.

CASCADING resources are resources which may be assigned to be taken over by multiple nodes in a prioritized manner. When a node fails, the active node with the highest priority acquires the resource. When the failed node rejoins, the node with the highest priority acquires the resource.

ROTATING resources are resources which may be acquired by any node in its resource chain. When a node fails, the resource will be acquired by the highest priority standby node. When the failed node rejoins, the resource remains with its new owner.

CONCURRENT resources are not supported by the B/A client.

Participating Node Names Enter the nodes in the resource chain for the resource group. These are the nodes which may acquire the resource. The order you list the nodes can determine the priority of the Resource Group allocation, depending on the Node relationship.

CASCADING RESOURCES: Priority decreases from left to right.

CONCURRENT RESOURCES: All nodes are at the same priority. The B/A client does not support Concurrent resources.

ROTATING RESOURCES: Priority decreases from left to right, but is only used to determine which standby acquires the resource in a multiple standby configuration.

Secondly, define a HACMP Application Server to manage the Tivoli Storage Manager client scheduler. this is where you define the scripts to start and stop the client scheduler. You can access the SMIT menu for this task via the command “smitty cm_cfg_app”.

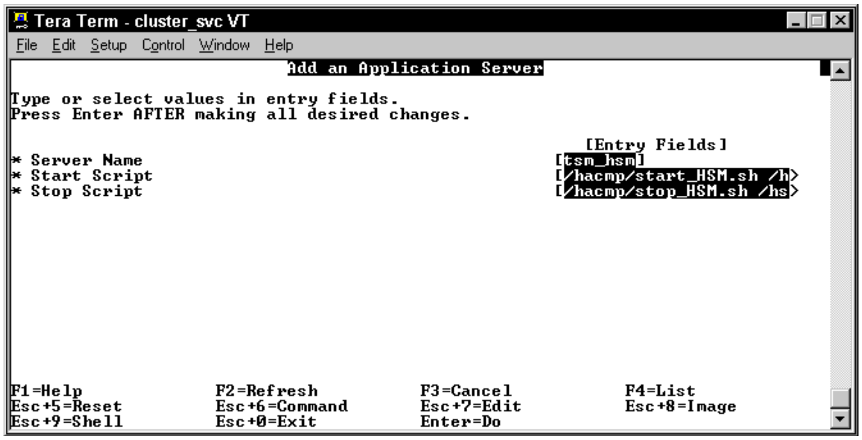

Figure 5-20 Defining a HACMP application server for the HSM client

You will need to configure the following parameters:

Server Name
Enter a logical name for your Application Server. Maximum length is 32 characters.

Start script
Script to start the application server when node joins the cluster. This script needs to be located on a local filesystem that is available when HACMP is not yet running.

Note: Remember to specify the HSM managed filesystems as parameters to the script.

Stop script
Script to stop the application server when node leaves the cluster gracefully. This script could be located on a shared filesystem, but should probably be stored on a local filesystem with the corresponding start script.

Note: Remember to specify the HSM managed filesystems as parameters to the script.

Finally, associate the application server and filesystem resources with the Resource Group you defined earlier. You can access the SMIT menu for this task via the command “smitty cm_cfg_res.select”.

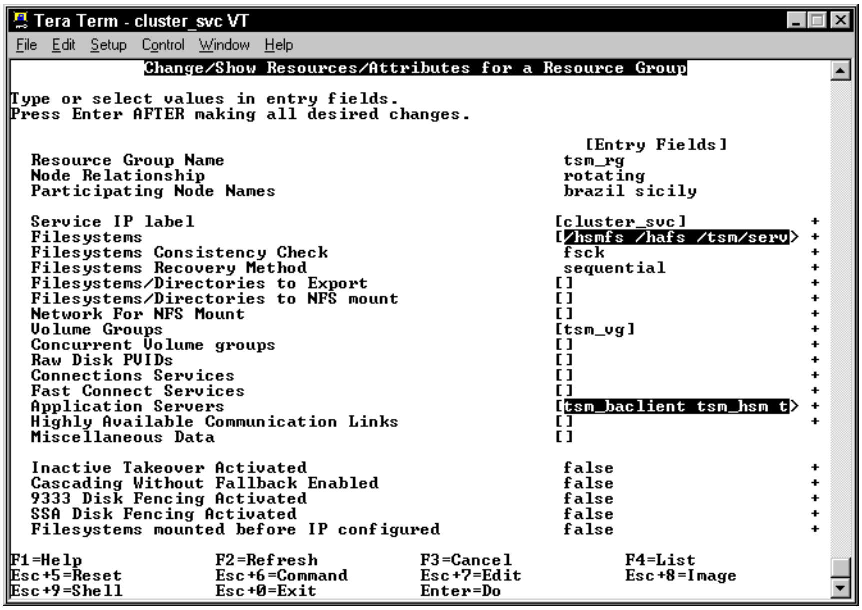

```
Tera Term - cluster_svc VT
File  Edit  Setup  Control  Window  Help
                Change/Show Resources/Attributes for a Resource Group

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

                                                        [Entry Fields]
  Resource Group Name                                  tsm_rg
  Node Relationship                                    rotating
  Participating Node Names                             brazil sicily

  Service IP label                                    [cluster_svc]           +
  Filesystems                                         [/hsmfs /hafs /tsm/serv> +
  Filesystems Consistency Check                        fsck                   +
  Filesystems Recovery Method                          sequential             +
  Filesystems/Directories to Export                   []                      +
  Filesystems/Directories to NFS mount                []                      +
  Network For NFS Mount                               []                      +
  Volume Groups                                       [tsm_vg]                +
  Concurrent Volume groups                            []                      +
  Raw Disk PVIDs                                      []                      +
  Connections Services                                []                      +
  Fast Connect Services                               []                      +
  Application Servers                                 [tsm_baclient tsm_hsm t> +
  Highly Available Communication Links                []                      +
  Miscellaneous Data                                  []

  Inactive Takeover Activated                          false                  +
  Cascading Without Fallback Enabled                   false                  +
  9333 Disk Fencing Activated                          false                  +
  SSA Disk Fencing Activated                           false                  +
  Filesystems mounted before IP configured             false                  +

F1=Help             F2=Refresh          F3=Cancel           F4=List
Esc+5=Reset         Esc+6=Command       Esc+7=Edit          Esc+8=Image
Esc+9=Shell         Esc+0=Exit          Enter=Do
```

Figure 5-21 Associating HACMP resources and application servers for the B/A client

You will need to configure at least the following parameters:

Service IP label — The IP label (hostname) of the adapter, associated with the numeric IP address in the /etc/hosts file.

Filesystems — Enter the mount points of the filesystems which are mounted when the resource is initially acquired.

Volume Groups — Enter the names of the volume groups containing raw logical volumes or raw volume groups that are varied on when the resource is initially acquired.

Application Servers — Enter application servers that will be started. These are the servers defined in Figure 5-9 on page 194.

> **Note:** The example shown in Figure 5-10 on page 195 also includes the resources and application server for the B/A client and Tivoli Storage Manager server, as defined in 3.1.2, “HACMP support for Tivoli Storage Manager server on AIX” on page 88. In this configuration our Tivoli Storage Manager server and B/A clients are running on the same HACMP cluster. This will probably not be the case in your environment.

Depending on your specific HACMP configuration, there may be other parameters to configure.

Finally, synchronize your changes to the other nodes in the cluster. You can access the SMIT menu for this task via the command “smitty clsyncnode.dialog”.

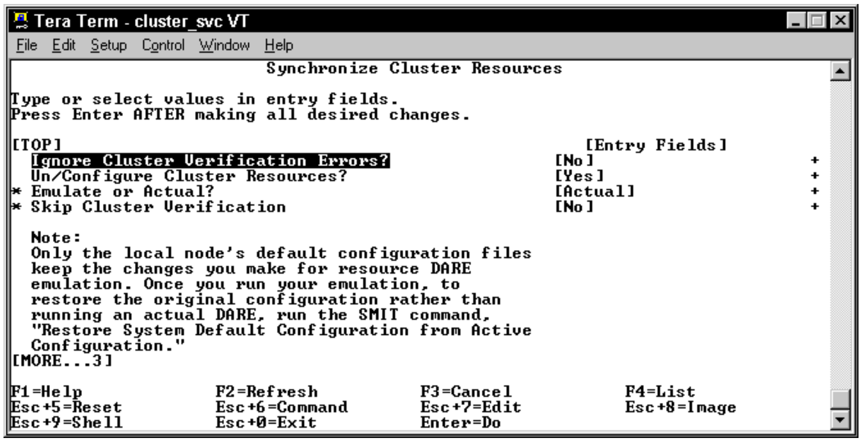

Figure 5-22 Synchronize HACMP cluster resources

Your HACMP environment should now be correctly configured and ready for testing.

Step 8: Test migration, failover, and recollection

After installing and customizing the B/A and HSM clients and configuring HACMP, you need to test migration and recalling data across a failover event.

Before starting the HACMP cluster for the first time, ensure you stop any services or resources normally managed by HACMP that you may have manually started during setup and installation. Repeat this check on each node in your cluster.

Start the HACMP cluster

First, start the HACMP software on all nodes in the cluster. You can access the SMIT menu for this task via the command “smitty cl_clstart.dialog

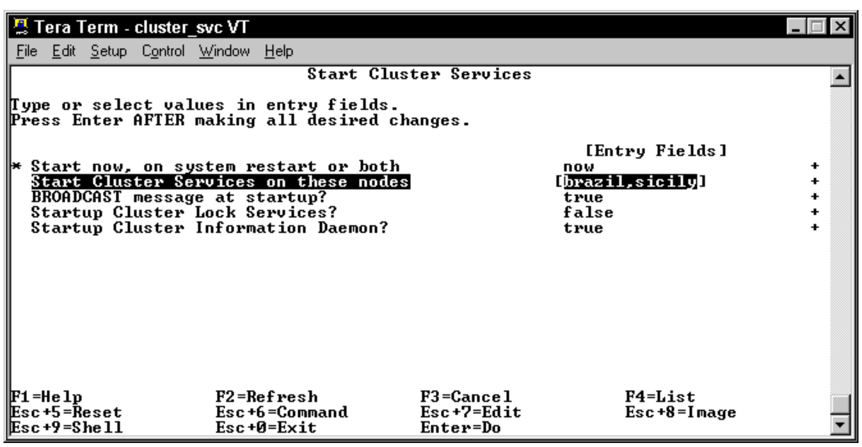

Figure 5-23 Starting the HACMP cluster on all nodes

You will need to configure the following parameters:

Start now, on system restart, or both Choosing 'now' will start the daemons immediately. Select this option for our testing purposes.

Choosing 'system restart' will start the daemons after a system reboot by adding an entry to the /etc/inittab file.

	Choosing 'both' will start the daemons immediately AND after a system reboot.
Start Cluster Services on these nodes	Identify the nodes on which to start cluster services.
Startup Cluster Information Daemon	Starts the Cluster Information daemon. This daemon needs to be started to support the B/A client.

This will cause the service IP address/s to become active, and you will lose any connections to the nodes boot IP address/s.

After the HACMP software has started, the applications and filesystems associated with your Tivoli Storage Manager Resource Group will be available on one of your cluster nodes. Use the following examples to test for the correct startup of the HSM daemons and correct mounting of the “fsm” filesystem.

Example 5-41 Testing for the startup of the HSM services after HACMP startup

```
root@brazil:/hacmp: ps -ef | grep hsm | grep -v grep
    root  5780     1   0 15:07:25      -  0:00 /usr/tivoli/tsm/client/hsm/bin/dsmrecalld
    root 11292     1   0 15:07:26      -  0:00 /usr/tivoli/tsm/client/hsm/bin/dsmmonitord
    root 21488     1   0 15:07:26      -  0:05 /usr/tivoli/tsm/client/hsm/bin/dsmscoutd

root@brazil:/hacmp: mount | grep hsmfs
         /dev/tsm_hsm     /hsmfs           jfs    Mar 20 14:53 rw,log=/dev/loglv00
         /hsmfs           /hsmfs           fsm    Mar 25 15:07

root@brazil:/hacmp: cat /etc/adsm/SpaceMan/config/dsmmigfstab
# Filesystem  High(%)  Low(%)   Premig(%)  Age    Size   Quota  stubsize Server
# Name        Thrshld  Thrshld  Percent    Factor Factor                 Name
# -------------------------------------------------------------------------------
/hsmfs  90      80      -       -       -       -       -       -       -
```

Migrate some data to the Tivoli Storage Manager server

Copy some non-critical data to the /hsmfs filesystem to test the migration function. You will need to perform a backup of this data before testing the migration. See Example 5-42.

Example 5-42 Testing the migration function

```
root@brazil:/hacmp: mkdir /hsmfs/test
root@brazil:/hacmp: cp /smit* /hsmfs/test

root@brazil:/hacmp: dsmmigrate /hsmfs/test/*
Tivoli Storage Manager
Space Management Interface - Version 5, Release 1, Level 0.0 g4
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

root@brazil:/hacmp: dsmls /hsmfs/test/smit*
Tivoli Storage Manager
Space Management Interface - Version 5, Release 1, Level 0.0 g4
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

      Actual     Resident     Resident  File   File
        Size         Size     Blk (KB)  State  Name
      209828         4095            4   m      smit.log
       14264         4095            4   m      smit.scrip
```

As the files migrate you should see entries in the Tivoli Storage Manager server activity log.

Force a failover to another node

Next, stop the HACMP software on whichever node has the Tivoli Storage Manager Resource Group. This should cause the Resource Group to failover to another node. You can access the SMIT menu for this task via the command “smitty clstop.dialog”.

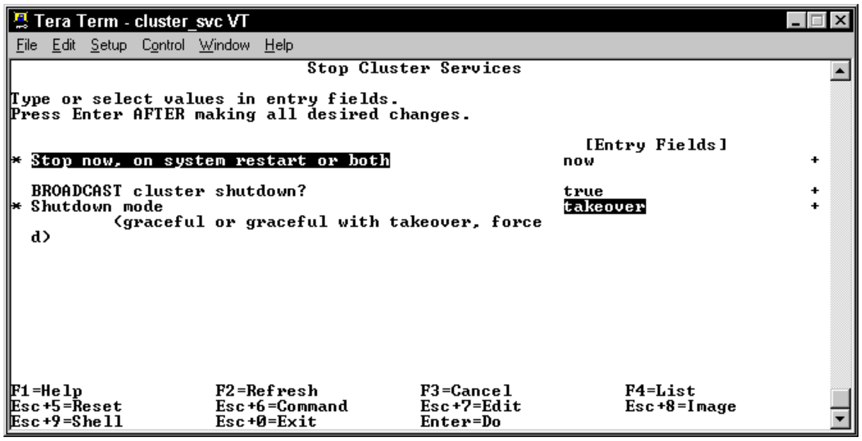

Figure 5-24 Controlling the HACMP cluster to test B/A client failover

This will cause the service IP address/s to move to another node, and you will lose any connections to the current node.

After the HACMP failover has completed, the applications and filesystems associated with your Tivoli Storage Manager Resource Group will be available on the next cluster node. Use the examples shown in Example 5-41 to test for the correct startup of the HSM services on the new node.

Recall the migrated data

Finally, with the filesystem mounted on the alternate node, we need to test both manual and automatic recollection of the migrated data. Use the following commands to test recollection of the migrated data:

Example 5-43 Checking for filespace ownership in a HACMP cluster

```
root@sicily:/hacmp: vi /hsmfs/test/smit.log
"/hsmfs/test/smit.log"ANS9283K Tivoli Space Manager is recalling a migrated file.
  7948 lines, 209828 characters
...

root@sicily:/hacmp: dsmrecall /hsmfs/test/smit.script
Tivoli Storage Manager
Space Management Interface - Version 5, Release 1, Level 0.0 g4
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

root@sicily:/hacmp: dsmls /hsmfs/test/smit*
Tivoli Storage Manager
Space Management Interface - Version 5, Release 1, Level 0.0 g4
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

      Actual     Resident     Resident  File   File
        Size         Size     Blk (KB)  State  Name
      209828       209828          208   p      smit.log
       14264        14264           16   p      smit.script
```

As the files recall you should see entries in the Tivoli Storage Manager server activity log.

5.2 HP-UX

The IBM Tivoli Storage Manager 5.1 client for HP-UX introduces some new features and additional support.

5.2.1 Changes in IBM Tivoli Storage Manger Version 5.1 client

Following is a summary of changes since Tivoli Storage Manager Version 5.1:

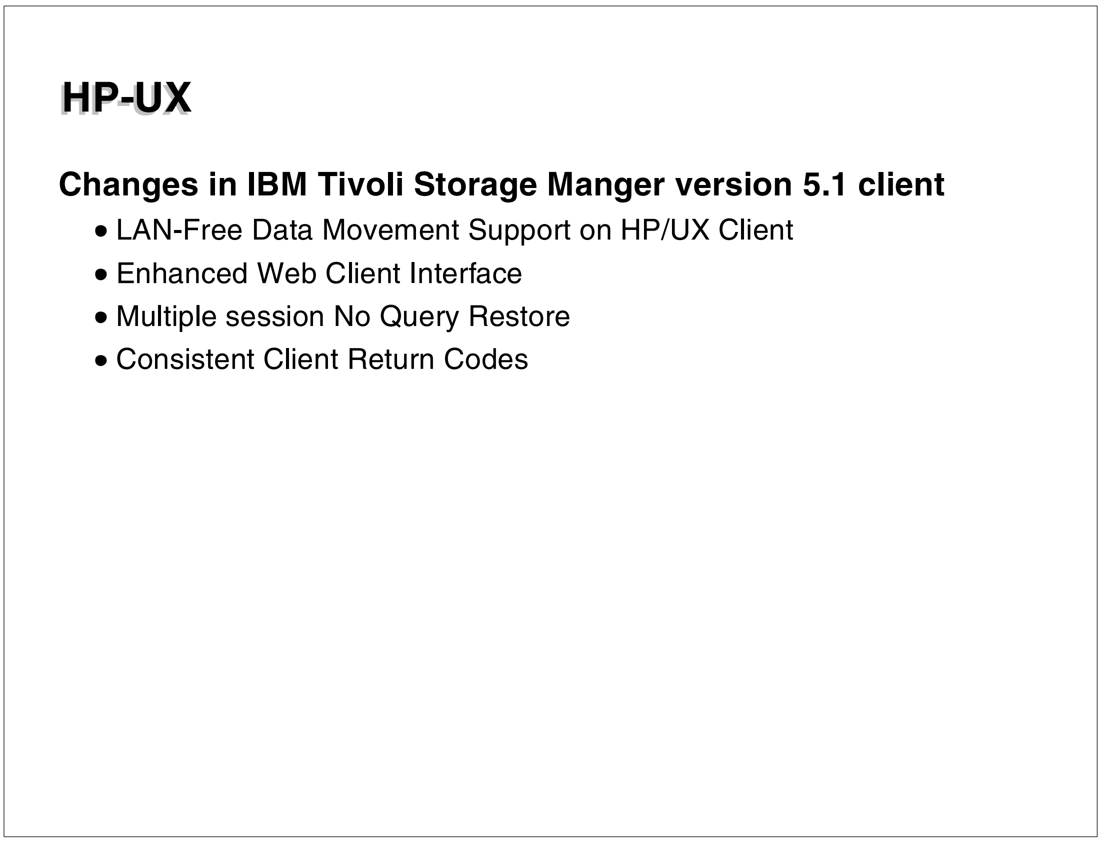

Figure 5-25 Changes in IBM Tivoli Storage Manger Version 5.1 client

- LAN-Free Data Movement Support on HP-UX Client

  Tivoli Storage Manager supports LAN-Free data movement in a Storage Area Network (SAN) environment for the HP-UX client. LAN-Free data movement allows client data to move directly from the client to a SAN-attached storage device. Shifting the client data movement from the communications network to a SAN decreases the load on the server. This allows the server to support a greater number of simultaneous client connections.

- Enhanced Web Client Interface

  The Web client interface is enhanced to support a JRE 1.3.1 Swing-enabled browser. This type of browser facilitates the use of assistive devices for users with disabilities. The native look and feel of the platform running the browser is preserved.

- Multiple session No Query Restore

  The backup-archive clients can now perform multiple restore sessions for no query restore operations, increasing the speed of restores. This is similar to the multiple backup session support. It exploits the mount point available on the server. If data is backed up on multiple tapes, and if the server has multiple mount points available, then the restore starts a session for each tape, up to the number your administrator configures.

- Consistent Client Return Codes

  Reliable, consistent, and documented return codes have been added to the command line client and the scheduler. This facilitates automation of client operations via user-written scripts. Administrators can now distinguish between scheduled backups that completed successfully with no skipped files and scheduled backups that completed successfully with one or more skipped files. Also if the preschedulecmd command ends with non-zero return codes, the scheduled event will not run. This ensures that scheduled events will not run if prerequisite commands do not complete successfully.

5.2.2 LAN-free support for B/A client on HP-UX

As of Version 5.1, the Tivoli Storage Manager storage agent is now available on the HP-UX platform. The Storage Agent introduces LAN-Free support on HP-UX. This support is equivalent to LAN-Free support on other platforms.

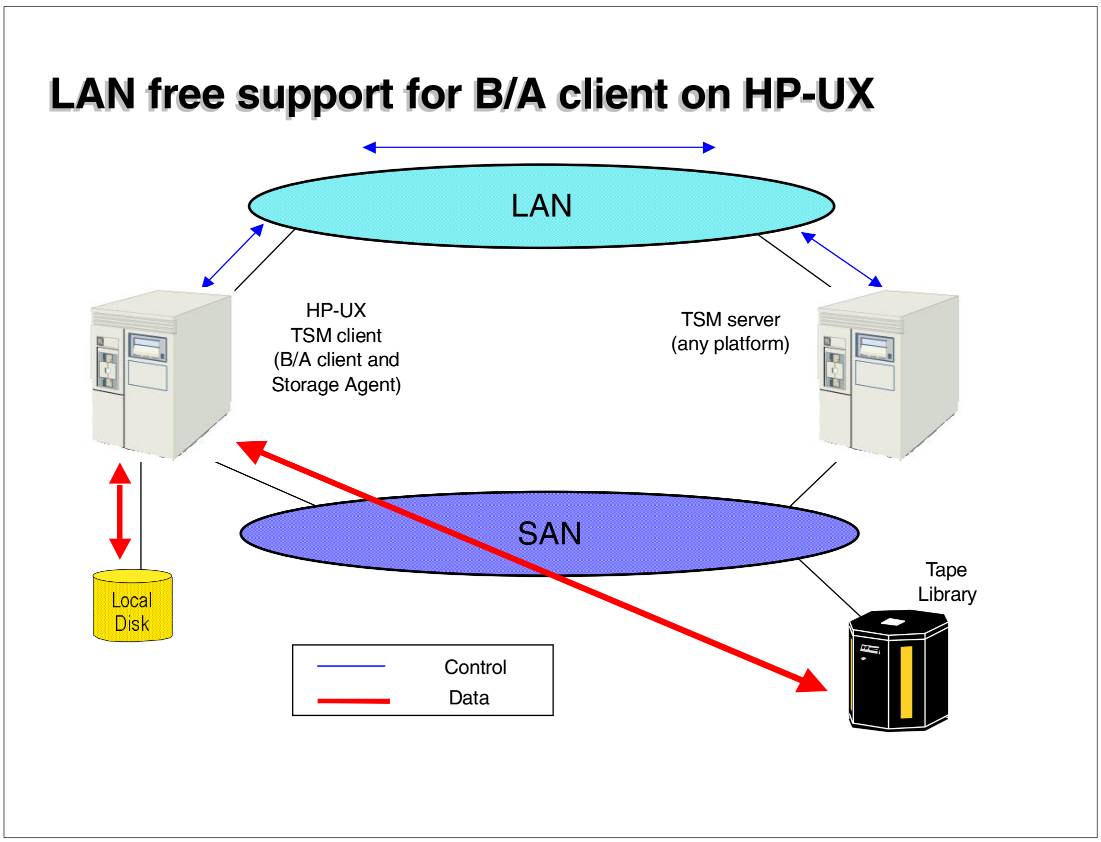

Figure 5-26 LAN free support

The supported configurations on HP-UX are:

- HPUX 11.0 or HP UX 11i
- Cannot use device type generic tape
- Any hardware configuration that the hardware vendors support

For further details regarding LAN-Free support, refer to the *Managed System for SAN Storage Agent User's Guide*, GC32-0727.

Upgrading from an earlier version

If a previous version of ADSM or Tivoli Storage Manager has been installed, you should remove it before installing the new packages.

- For older ADSM versions execute the following command as the root user:

  ```
  /usr/sbin/swremove -x mount_all_filesystems=false -v IBMADSM
  ```

- For older Tivoli Storage Manager versions use the following procedure:
  - To remove the BA client, execute the following as the root user:

    ```
    /usr/sbin/swremove -x mount_all_filesystems=false -v TIVsm.CLIENT
    ```

  - If you also installed the CLIENT_API and CLIENT_DOC filesets, execute the following command to remove them:

    ```
    /usr/sbin/swremove -x mount_all_filesystems=false -v TIVsm.CLIENT_API
    TIVsm.CLIENT_DOC
    ```

  - If you have installed additional languages, execute the following command to remove them (where xx_XX may be replaced with language you want to remove.):

    ```
    /usr/sbin/swremove -x mount_all_filesystems=false -v TIVsm.CLIENT_msg_xx_XX \
    TIVsm.CLIENT_API_xx_XX
    ```

  Note you can only uninstall the Web Client together with the regular client.

Install the device driver

The Tivoli Storage Manager device driver, called tsmscsi, and newly supported device classes are now available to replace the SCSI pass-through device driver, the HP-UX generic device drivers, and the GENERICTAPE device class.

Attention: For libraries that can use tsmscsi, you must migrate to the tsmscsi device driver. The SCSI pass-through device driver is no longer supported.

For tape drives, tsmscsi is optional but recommended. However, tracing and error reporting is better with tsmscsi than with the standard HP-UX device driver, and tsmscsi supports additional device classes. These device classes let you specify format types for tapes, allow greater server control of SCSI devices, and enable new server functions.

The latest version of the IBM device driver is available online at the following URL:

`ftp://ftp.software.ibm.com/storage/devdrvr`

Versions of this driver are identified by the following levels:

ATDD 1.x.x.x HP PCI Bus - HP-UX 11.00 (64 bit). This driver supports IBM TotalStorage and Ultrium devices.

ATDD 2.9.x.x HP Precision Bus - HP-UX 10.20 (32 bit). This driver supports IBM TotalStorage devices only.

ATDD 3.x.x.x HP PCI Bus - HP-UX 11i (64 bit). This driver supports IBM TotalStorage and Ultrium devices.

ATDD 4.9.x.x HP Precision Bus - HP-UX 11.00 (32 bit). This driver supports IBM TotalStorage devices only.

Note: After installing the IBM device driver, the HP-UX client will need to reboot.

Install and configure the Tivoli Storage Manager B/A client

Perform the following tasks to install the B/A client. Examples of command syntax, options, and typical responses are shown in Example 5-44 through Example 5-59 on page 232.

As a root user, do one of the following:

- If you are installing from CD-ROM, first mount the CD-ROM and change directory to:

  `/cdrom/tsmcli/hp11`

- If you have downloaded from ftp, go to the directory where the installable image is located.

To install the BA client, execute the following as the root user:

```
/usr/sbin/swinstall -x mount_all_filesystems=false -v -s `pwd`/TIVsmC TIVsm
```

Example 5-44 Typical dsm.sys file configured for LAN-free backup

```
SErvername              tsm_server1
COMMmethod              TCPip
TCPPort                 1500
TCPServeraddress        tsmserver

nodename                easter_fs
inclexcl                /opt/tivoli/tsm/client/ba/bin/dsm.inx
passwordaccess          generate

schedmode               prompted
schedlogname            /var/log/dsmsched.log
schedlogretention       7,t

errorlogname            /var/log/dsmerror.log
errorlogretention       7,t

enablelanfree yes
```

Example 5-45 Typical dsm.opt configuration file

```
SErvername              tsm_server1
```

Install and configure the Storage Agent client

Example 5-46 shows the command to install the Storage Agent. The screen is shown in Figure 5-27.

Example 5-46 The swinstall command syntax

```
# swinstall -s /tmp/tsm510stagent_beta2_refreshed.hp.img
...
```

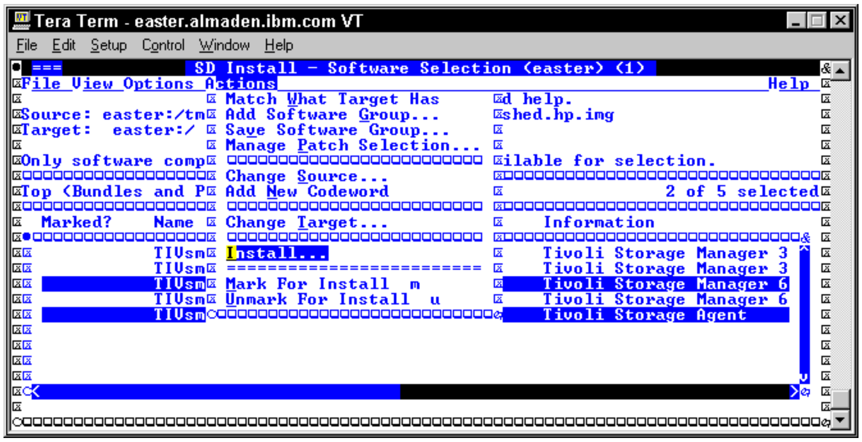

Figure 5-27 The swinstall example screen

Note: If you install the Storage Agent device driver, the HP-UX client will need to reboot.

Example 5-47 Starting the Storage Agent

```
# ./dsmsta setstorageserver myname=easter_sa mypassword=password myhladdress=easter
servername=tsm_server1 serverpassword=password hladdress=tsmserver lladdress=1500

Tivoli Storage Manager for HP-UX
Version 5, Release 1, Level 0.0

Licensed Materials - Property of IBM

5698-ISE (C) Copyright IBM Corporation 1999,2002. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

ANR7800I DSMSERV generated at 14:26:03 on Feb  1 2002.
ANR7801I Subsystem (master) process ID is 4636.
ANR0900I Processing options file dsmsta.opt.
ANR1432I Updating device configuration information to defined files.
ANR1433I Device configuration information successfully written to devconfig.
ANR2119I The SERVERNAME option has been changed in the options file.
ANR0467I The SETSTORAGESERVER command completed successfully.
```

Example 5-48 Typical Storage Agent devconfig file

```
SET STANAME EASTER_SA
SET STAPASSWORD 1804bdc565234d
SET STAHLADDRESS EASTER
DEFINE SERVER TSM_SERVER1 HLADDRESS=9.1.38.163 LLADDRESS=1500 SERVERPA=18830fd367cfa8
```

Example 5-49 Minimal dsmsta.opt file (stripped of comments)

```
COMMmethod TCPIP
COMMmethod SHAREDMEM

devconfig devconfig
SERVERNAME TSM_SERVER1
```

Configure the Tivoli Storage Manager server

Example 5-50 Register client

```
REGISTER NODE easter_fs ?***? DOMAIN=ITSO AUTOFSRENAME=NO ARCHDELETE=YES BACKDELETE=NO
FORCEPWRESET=NO TYPE=CLIENT KEEPMP=NO MAXNUMMP=2 URL=http://easter:1581 USERID=none
```

Example 5-51 Define server

```
DEFINE SERVER easter_sa COMMMETHOD=TCPIP SERVERPASSWORD=?***? HLADDRESS=easter
LLADDRESS=1500 VALIDATEPROTOCOL=NO NODENAME=easter_fs PASSWORD=?***?
```

> **Note:** If you intend is to send client data via the Storage Agent directly to a tape storage pool, you need to modify the copygroup parameter "destination" in the appropriate management class.
>
> If you intend to send client data via the Storage Agent to a shared file storage pool, you will need to configure Tivoli SANergy.

Example 5-52 Set the Copygroup destination to a tape device

```
UPDATE COPYGROUP ITSO ITSOSJ BKUP_FS_DEFAULT_MC STANDARD DESTINATION=PRI_FS_TAPE
FREQUENCY=0 VEREXISTS=2 VERDELETED=1 RETEXTRA=30 RETONLY=60 MODE=MODIFIED
SERIALIZATION=SHRSTATIC
```

Example 5-53 Update library to "shared=yes"

```
UPDATE LIBRARY LIBRARY0 SHARED=YES
```

To see the special file names, enter the command: **ioscan -kfn.** This produces output as shown in Example 5-54.

Example 5-54 Output from "ioscan -kfn" command

```
Class       I  H/W Path                Driver    S/W State   H/W Type     Description
========================================================================================
...
ctl        58  0/4/0/0.3.27.0.0.0.0     sctl      CLAIMED     DEVICE       PATHLGHTSAN
                                       /dev/rscsi/c17t0d0
tape      108  0/4/0/0.3.27.0.0.0.6     atdd      CLAIMED     DEVICE       IBM     03570C12
                                       /dev/rmt/108m           /dev/rmt/c17t0d6BEST
                                       /dev/rmt/108mb          /dev/rmt/c17t0d6BESTb
                                       /dev/rmt/108mn          /dev/rmt/c17t0d6BESTn
                                       /dev/rmt/108mnb         /dev/rmt/c17t0d6BESTnb
...
tape      141  0/4/0/0.3.27.0.0.1.0     atdd      CLAIMED     DEVICE       IBM     03570C12
                                       /dev/rmt/141m           /dev/rmt/c17t1d0BEST
                                       /dev/rmt/141mb          /dev/rmt/c17t1d0BESTb
                                       /dev/rmt/141mn          /dev/rmt/c17t1d0BESTn
                                       /dev/rmt/141mnb         /dev/rmt/c17t1d0BESTnb
...
```

Example 5-55 Define drive path for the HP-UX storage agent

```
DEFINE PATH easter_sa DRIVE0 SRCTYPE=SERVER DESTTYPE=DRIVE LIBRARY=LIBRARY0
DEVICE=/dev/rmt/??? ONLINE=YES

DEFINE PATH easter_sa DRIVE1 SRCTYPE=SERVER DESTTYPE=DRIVE LIBRARY=LIBRARY0
DEVICE=/dev/rmt/??? ONLINE=YES
```

Operation

Example 5-56 Starting the Storage agent in the foreground

```
# ./dsmsta

Tivoli Storage Manager for HP-UX
Version 5, Release 1, Level 0.0

Licensed Materials - Property of IBM

5698-ISE (C) Copyright IBM Corporation 1999,2002. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corporation.

ANR7800I DSMSERV generated at 14:26:03 on Feb  1 2002.
ANR7801I Subsystem (master) process ID is 4693.
ANR0900I Processing options file dsmsta.opt.
ANR8200I TCP/IP driver ready for connection with clients on port 1500.
ANR8285I Shared Memory driver ready for connection with clients on port 1510
ANR0993I Server initialization complete.
ANR0916I TIVOLI STORAGE MANAGER distributed by Tivoli is now ready for use.
TSM:EASTER_SA>
```

Example 5-57 shows the output from a typical LAN-free backup.

Example 5-57 Output from a typical LAN-Free B/A client backup

```
# dsmc s '/opt/tivoli/*' -subdir=yes
Tivoli Storage Manager
Command Line Backup/Archive Client Interface - Version 5, Release 1, Level 0.0 g3
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Selective Backup function invoked.

Node Name: EASTER_FS
Session established with server TSM_SERVER1: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 04/05/02   15:57:26  Last access: 04/05/02   15:54:44

Directory-->                  96 /opt/tivoli [Sent]
Directory-->                  96 /opt/tivoli/tsm [Sent]
Directory-->                  96 /opt/tivoli/tsm/StorageAgent [Sent]
Directory-->                  96 /opt/tivoli/tsm/client [Sent]
Directory-->                  96 /opt/tivoli/tsm/devices [Sent]
Directory-->                  96 /opt/tivoli/tsm/tivinv [Sent]
Directory-->               1,024 /opt/tivoli/tsm/StorageAgent/bin [Sent]
Normal File-->            32,137 /opt/tivoli/tsm/StorageAgent/bin/README.STA [Sent]
ANS1114I Waiting for mount of offline media.
...
Normal File-->                 7 /opt/tivoli/tsm/tivinv/TSMBACHPX05_01.SIG [Sent]
```

```
Normal File-->                 7 /opt/tivoli/tsm/tivinv/TSMSACHPX05_01.SIG [Sent]
Selective Backup processing of '/opt/tivoli/*' finished without failure.

Total number of objects inspected:      367
Total number of objects backed up:      367
Total number of objects updated:          0
Total number of objects rebound:          0
Total number of objects deleted:          0
Total number of objects expired:          0
Total number of objects failed:           0
Total number of bytes transferred:    93.41 MB
  LanFree data bytes:                 91.76 MB
Data transfer time:                   37.66 sec
Network data transfer rate:        2,539.50 KB/sec
Aggregate data transfer rate:        915.12 KB/sec
Objects compressed by:                    0%
Elapsed processing time:           00:01:44
```

When a LanFree operation has completed the command line will report the 'LanFree Data Bytes:' as one of the final statistics. The GUI will update the 'LanFree Data Bytes:' field at the end of every transaction.

ANR8337I 3570 volume 07D9A3 mounted in drive DRIVE0 (/dev/rmt0).

Example 5-58 Output from dsmsta console

```
ANR0400I Session 11 started for node EASTER_FS (HPUX) (ShMem).
ANR8337I 3570 volume 0838B5 mounted in drive DRIVE1 (/dev/rmt/c17t1d0BEST).
ANR8336I Verifying label of 3570 volume 0838B5 in drive DRIVE1 (/dev/rmt/c17t1d0BEST).
ANR8468I 3570 volume 0838B5 dismounted from drive DRIVE1 (/dev/rmt/c17t1d0BEST) in library
LIBRARY0.
ANR0403I Session 11 ended for node EASTER_FS (HPUX).
```

Example 5-59 Relevant messages in TSM server activity log

```
ANR0406I Session 110 started for node EASTER_FS (HPUX) (Tcp/Ip 9.1.38.152(52001)).
ANR0406I Session 111 started for node EASTER_FS (HPUX) (Tcp/Ip 9.1.38.152(52003)).
ANR8337I 3570 volume 07D9A3 mounted in drive DRIVE0 (/dev/rmt0).
ANR0408I Session 112 started for server EASTER_SA (HP-UX) (Tcp/Ip) for library sharing.
ANR8337I 3570 volume 0838B5 mounted in drive DRIVE1 (/dev/rmt2).
ANR0409I Session 112 ended for server EASTER_SA (HP-UX).
ANR0403I Session 111 ended for node EASTER_FS (HPUX).
ANR0403I Session 110 ended for node EASTER_FS (HPUX).
ANR0408I Session 113 started for server EASTER_SA (HP-UX) (Tcp/Ip) for library sharing.
ANR8336I Verifying label of 3570 volume 0838B5 in drive DRIVE1 (/dev/rmt2).
ANR0409I Session 113 ended for server EASTER_SA (HP-UX).
ANR8468I 3570 volume 07D9A3 dismounted from drive DRIVE0 (/dev/rmt0) in library LIBRARY0.
ANR8468I 3570 volume 0838B5 dismounted from drive DRIVE1 (/dev/rmt2) in library LIBRARY0
```

5.3 Linux

The IBM Tivoli Storage Manager 5.1 client for Linux introduces some new features and additional support.

5.3.1 Changes in IBM Tivoli Storage Manger Version 5.1 client

Following is a summary of changes since Tivoli Storage Manager Version 5.1:

Linux x86

Changes in IBM Tivoli Storage Manger version 5.1 client

- Support for a Logical Volume Backup as a Single Object (Image Backup)
- Support for Online Image Backup of File Systems and Raw Logical Volumes
- Enhanced Web Client Interface
- Multiple session No Query Restore
- Consistent Client Return Codes

Figure 5-28 Changes in IBM Tivoli Storage Manger Version 5.1 client

- Support for a Logical Volume Backup as a Single Object (Image Backup)

  The Linux86 client is enhanced to support a logical volume image backup of file systems and raw volumes. The Tivoli Storage Manager server does not track individual files in the file system image. File system images are tracked as individual objects and management class policy will be applied to the file system image as a whole.

- Support for Online Image Backup of File Systems and Raw Logical Volumes

  The traditional offline image backup prevents access to the volume by other system applications during the operation. For Linux86 only: Tivoli Storage Manager performs an online image backup of file systems residing on a logical volume created by the Linux Logical Volume Manager, during which the volume is available to other system applications.

- Enhanced Web Client Interface

  The Web client interface is enhanced to support a JRE 1.3.1 Swing-enabled browser. This type of browser facilitates the use of assistive devices for users with disabilities. The native look and feel of the platform running the browser is preserved.

- Multiple session No Query Restore

  The backup-archive clients can now perform multiple restore sessions for no query restore operations, increasing the speed of restores. This is similar to the multiple backup session support. It exploits the mount point available on the server. If data is backed up on multiple tapes, and if the server has multiple mount points available, then the restore starts a session for each tape, up to the number your administrator configures.

- Consistent Client Return Codes

  Reliable, consistent, and documented return codes have been added to the command line client and the scheduler. This facilitates automation of client operations via user-written scripts. Administrators can now distinguish between scheduled backups that completed successfully with no skipped files and scheduled backups that completed successfully with one or more skipped files. Also if the preschedulecmd command ends with non-zero return codes, the scheduled event will not run. This ensures that scheduled events will not run if prerequisite commands do not complete successfully.

5.3.2 Logical volume image backup for Linux

Since Version 4.2 the Tivoli Storage Manager B/A client for Linux has been able to perform offline image backups of filesystem block devices. The advantage of an image backup is quicker backup and restore than a file-by-file backup, as there is no overhead involved in creating individual files. This also conserves resources on the Tivoli Storage Manager server as only one object (the image file) needs to be tracked in the server database.

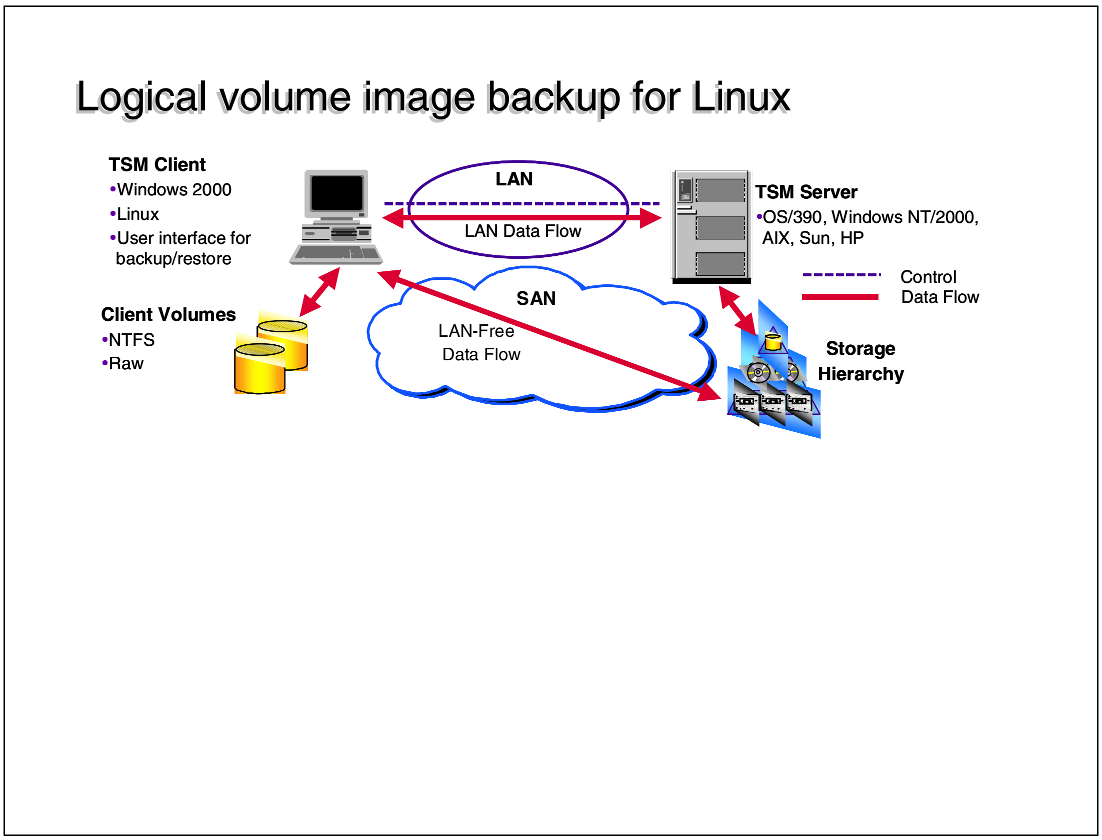

Figure 5-29 Logical volume image backup for Linux

On offline image backup has the unfortunate restriction of having to quiesce and unmount a filesystem before performing the image backup. The filesystem would remain unavailable until the image backup was completed and it was again mounted in read-write mode.

In Version 5.1 a new online “snapshot” image backup mode is available for the Linux B/A client. This allows the filesystem to remain online and active during the snapshot image backup. This function leverages the “snapshot” mode available in the Linux logical volume manager.

There are some restrictions regarding when the online “snapshot” mode can be used:

- The online “snapshot” image backup is only supported in Linux (and Windows)
- The online “snapshot” image backup is only supported on Linux LVM volumes
- The Tivoli Storage Manager API code must be installed to use the “backup image” command
- The volume group containing the logical volume being backed up must contain sufficient free space (logical extents) to hold the snapshot cache data.

Other features of the Linux B/A client remain unchanged from the previous version.

Linux logical volume manager terms

Most Linux distributions now come with the Logical Volume Manager function enabled. The LVM allows for a much more flexible allocation of disk resources in comparison to traditional PC disk partitioning. Those familiar with the LVM in commercial UNIX’s will find the Linux LVM comfortingly familiar. Below are a few terms you will need to understand before working with the Linux LVM.

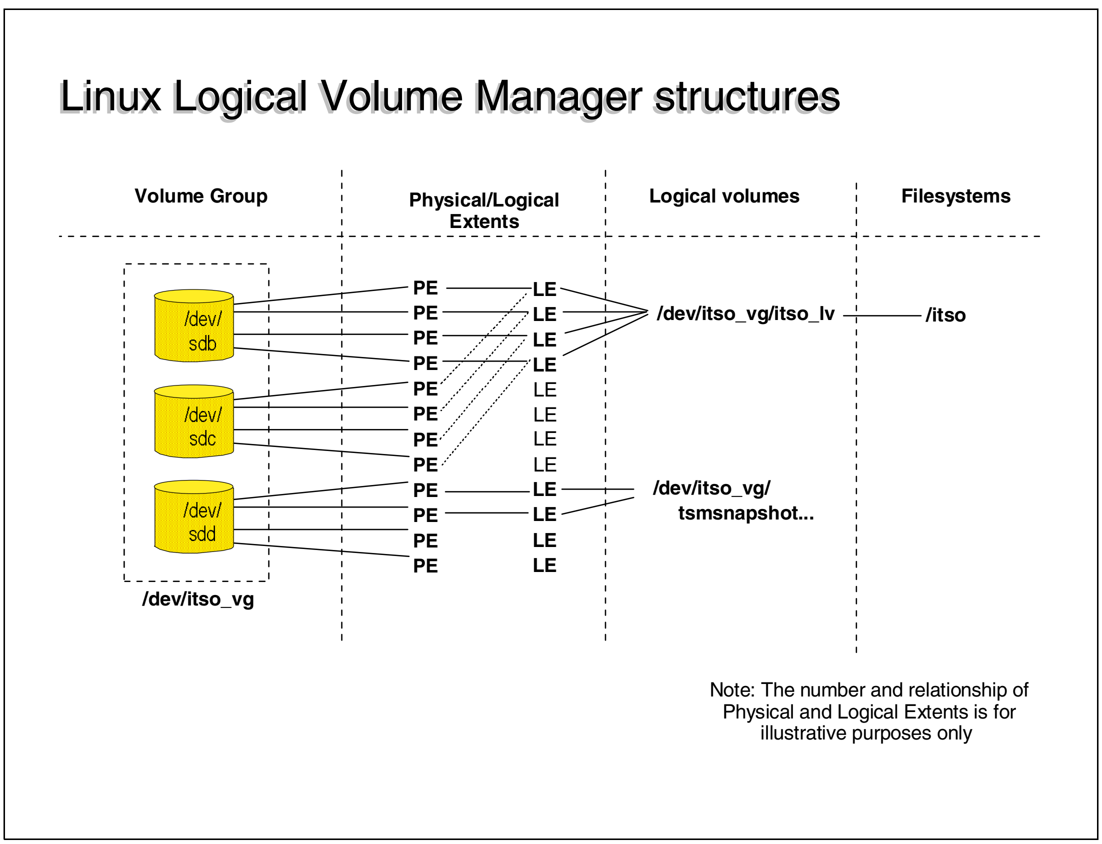

Figure 5-30 Linux Logical Volume Manager structures

Physical Volume (PV) A physical volume is typically a hard disk, though it may be any block device that 'looks' like a hard disk. This is equivalent to a Physical Volume (PV) in the AIX logical volume manager.

Physical Extents (PE) Each physical volume is divided chunks of data, known as physical extents. This is equivalent to a Physical Partition (PP) in the AIX logical volume manager.

Logical Extents (LE) Each physical extent is mapped to a logical view, known as logical extents. This is equivalent to a Logical Volume (LP) in the AIX logical volume manager.

Volume Group (VG) The Volume Group bundles one or more Physical Volumes into one administrative unit. This is equivalent to a Volume Group (VG) in the AIX logical volume manager.

Logical Volume (LV) The equivalent of a disk partition in a non-LVM system (/dev/itso_vg/itso_lv). A logical volume is created within a Volume Group. This is equivalent to a Logical Volume (LV) in the AIX logical volume manager.

Filesystem (FS) A perfectly normal filesystem (ext3 for example). A filesystem is created on a Logical Volume. You may create any supported filesystem on an LVM logical volume.

Updated commands

Two commands have been updated to support the new snapshot image backup functionality. The "backup image" command has gained an "imagetype" parameter and there are some new parameters for the "include.image" option in the system include/exclude file.

Backup image command

An image backup can be performed from either the Tivoli Storage Manager B/A command line client "dsmc" or via the GUI client interface "dsm". Example 5-17 illustrates the syntax of the new client command.

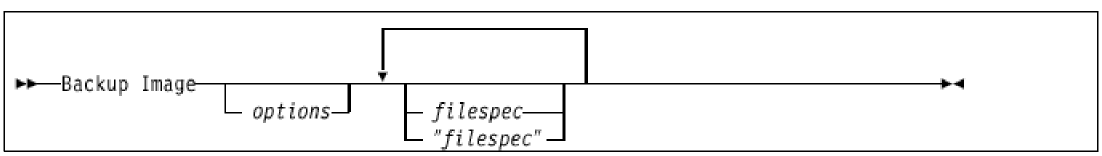

Figure 5-31 Syntax of the "dsmc backup image" command

The various options are as follows:

-mode= Determines whether to backup the entire logical volume image, or only those files which have changed.

selective: (Default) Creates and sends a full backup copy of the logical volume to the Tivoli Storage Manager server. The entire logical volume is backed up, not just the used blocks.

incremental: Performs and incremental backup of only those files which have changed since the data of the last selective image backup. This is similar in effect and limitations to a filesystem "incremental by date" backup.

-imagetype=	Determines whether to perform an online or offline image backup. **static:** Performs an offline image backup during which the volume is unmounted and remounted read-only. The filesystem will be available in read-only mode during the course of the backup. When the backup completes the volume will be remounted in re-write mode. **dynamic:** Performs an online image backup during which the volume remains mounted and writable. Use this option only if the volume cannot be unmounted and remounted read-only. Corruption of the backup may occur if applications write to the volume while the backup is in progress. In this case, run fsck after a restore. **snapshot:** Performs an online image backup during which the volume remains mounted and writable. Changes to the filesystem are cached in the snapshot cache and synchronized after the backup has completed. This is only valid for file systems residing on a logical volume created by the Linux Logical Volume Manager.
-snapshotcachesize=	Specifies the size of the snapshot area as a percentage of the logical volume you wish to backup. The snapshot area will be created as a new logical volume in the same volume group as the logical volume to be backed up. The snapshot area should be large enough to store any changes to the filesystem during the backup process. If the snapshot area is too small the image backup will fail.
filespec	Specify the mount point of the filesystem you wish to include in the image backup.

There are some restrictions regarding when you can perform an incremental image backup. You cannot perform an incremental image backup if you already have a complete incremental filesystem backup of the same filesystem.

Include.image option

The default behavior of the backup image command can be modified by editing the system include/exclude file. You should add an entry similar to Example 5-60 for each filesystem on which you plan to perform an image backup.

Example 5-60 New include.image parameters in system include/exclude file

```
include.image /itso mgmtclass imagetype=snapshot snapshotcachesize=50
```

The parameters for the include.image option correspond to those set on the command line "dsmc backup image" command.

Tivoli Storage Manager server policy considerations

In previous versions of the Linux B/A client the copy group parameter “Copy Serialization” was used to determine whether to perform a static or dynamic image backup. This parameter is ignored in the new version, being replaced by the “-imagetype=” command line parameter.

It is recommended to create a specific management class for image backups. This allows the Tivoli Storage Manager administrator to retain a reduced number of versions of the large image backup data. Table 5-1 shows an example of the settings.

Table 5-1 Typical management class settings

Backup copy group parameters	Typical file management class (default_mc)	Image backup management class (image_mc)
Versions exist	10	4
Versions deleted	3	4
Retain extra	28	28
Retain only	365	28

Consideration should be given to the storage pool destination of the image backup data. The “maxsize” parameter on a storage pool may force the image data to skip to the next storage pool in the storage hierarchy.

Step 1: Install and configure the Tivoli Storage Manager client

Locate the latest version of the B/A client code and make the media available to Linux. This may involve mounting the original CD-ROM or downloading the latest maintenance level code from the Tivoli Web site. The most up to date version is always available at the following URL:

http://www.tivoli.com/tsm

If you are installing the B/A client code for the first time, switch to the root user, and use “rpm -i filename” command to install the required filesets. The default installation directory for the B/A client is /opt/tivoli/tsm/client/ba/bin/. Several sample configuration files can be found in this directory.

If you are upgrading over an existing installation, switch to the root user, and use “rpm -U filename” to update the already installed filesets. Your existing configuration files will be retained.

Example 5-61 Installing the Tivoli Storage Manager B/A client on Linux

```
[root@azure tsm]# tar xvf ../TIVsm-BETA-Linux86.tar
README
README.API
TIVsm-API.i386.rpm
TIVsm-API.msg.de_DE.i386.rpm
TIVsm-API.msg.es_ES.i386.rpm
TIVsm-API.msg.fr_FR.i386.rpm
TIVsm-API.msg.it_IT.i386.rpm
TIVsm-API.msg.ja_JP.i386.rpm
TIVsm-API.msg.ko_KR.i386.rpm
TIVsm-API.msg.pt_BR.i386.rpm
TIVsm-API.msg.zh_CN.i386.rpm
TIVsm-API.msg.zh_TW.i386.rpm
TIVsm-BA.GPFS.i386.rpm
TIVsm-BA.i386.rpm
TIVsm-BA.msg.de_DE.i386.rpm
```

```
TIVsm-BA.msg.es_ES.i386.rpm
TIVsm-BA.msg.fr_FR.i386.rpm
TIVsm-BA.msg.it_IT.i386.rpm
TIVsm-BA.msg.ja_JP.i386.rpm
TIVsm-BA.msg.ko_KR.i386.rpm
TIVsm-BA.msg.pt_BR.i386.rpm
TIVsm-BA.msg.zh_CN.i386.rpm
TIVsm-BA.msg.zh_TW.i386.rpm

[root@azure tsm]# rpm -i TIVsm-API.i386.rpm
Postinstall of the API

TSM Linux API installation complete.

Be sure to set up the configuration files!

[root@azure tsm]# rpm -i TIVsm-BA.i386.rpm
Postinstall of the Backup Archive client

TSM Linux client installation complete.

Be sure to set up the system configuration file
before starting the client!
```

If you are upgrading over an existing installation, ensure you stop any currently running Tivoli Storage Manager services before performing the upgrade. For example, kill any running client scheduler or client acceptor tasks.

Copy the sample /opt/tivoli/tsm/client/ba/bin/dsm.sys.smp file to dsm.sys and edit it to resemble Example 5-62. You will need to modify some parameters to suit your local environment.

Example 5-62 Typical dsm.sys configuration file for the Linux B/A client

```
************************************************************************
SErvername              tsm_server1
COMMmethod              TCPip
TCPPort                 1500
TCPServeraddress        tsmserver

nodename                azure_fs

domain                  /itso
inclexcl                /opt/tivoli/tsm/client/ba/bin/dsm.inx

*compression            yes

passwordaccess          generate

managedservices         schedule webclient
schedmode               prompted

errorlogname            /var/log/dsmerror.log
schedlogname            /var/log/dsmsched.log
```

Copy the sample /opt/tivoli/tsm/client/ba/bin/dsm.opt.smp file to dsm.opt (in the same directory) and edit it to resemble Example 5-63. You may add other parameters to suit your local environment.

Example 5-63 Typical dsm.opt configuration file for the Linux B/A client

```
SErvername              tsm_server1
```

Create a new file to hold your system's include/exclude configuration. The location of this file is specified in the dsm.sys file. Example 5-64 shows a simple include/exclude file configured for the /itso filesystem.

Example 5-64 Typical dsm.inx configuration file for the Linux B/A client

```
include.image /itso image_mc imagetype=snapshot snapshotcachesize=50
```

If you wish to start the Tivoli Storage Manager client acceptor daemon (dsmad) automatically at system boot time, copy the script shown in Example 5-65 to /etc/rc.d/init.d/dsmcad and make it executable by the root user.

Example 5-65 Script to start Linux Tivoli Storage Manager client acceptor daemon at system boot time

```
#!/bin/bash
#
# chkconfig: 2345 55 45
# description: TSM client acceptor start/stop
#

#--------------------------------------------------
# Functions and variables

. /etc/rc.d/init.d/functions

DSM_DIR=/opt/tivoli/tsm/client/ba/bin
DSM_CONFIG=$DSM_DIR/dsm.opt
RETVAL=0

export DSM_DIR DSM_CONFIG

#--------------------------------------------------
# What to do?

case "$1" in
        start)
                gprintf "Starting TSM client acceptor: "
                daemon $DSM_DIR/dsmcad
                RETVAL=$?
                echo
                [ $RETVAL -eq 0 ] && touch /var/lock/subsys/dsmcad
                ;;
        stop)
                gprintf "Stopping TSM client acceptor: "
                killproc dsmcad
                RETVAL=$?
                echo
                [ $RETVAL -eq 0 ] && rm -f /var/lock/subsys/dsmcad
                ;;
        restart)
```

```
                $0 stop
                $0 start
                ;;
        status)
                status dsmcad
                ;;
        *)
                gprintf "Usage: %s {start|stop|restart|status}"
                exit 1
esac

exit $?
```

Use the commands shown in Example 5-66 to configure Linux to start the Tivoli Storage Manager client acceptor at system boot time.

Example 5-66 Configuring Linux to start the Tivoli Storage Manager client acceptor at system boot time

```
[root@azure init.d]# chmod 744 /etc/rc.d/init.d/dsmcad

[root@azure init.d]# chkconfig --add dsmcad

[root@azure init.d]# chkconfig --list | grep dsmcad
dsmcad          0:off   1:off   2:on    3:on    4:on    5:on    6:off

[root@azure init.d]# /etc/rc.d/init.d/dsmcad start
Starting TSM client acceptor:                                  [  OK  ]

[root@azure init.d]# /etc/rc.d/init.d/dsmcad status
dsmcad (pid 3548) is running...
3544 (pid 3543) is running...
3542 (pid 3538) is running...
```

The Tivoli Storage Manager client acceptor will now start and stop automatically when your system boots and moves between runlevels.

Step 2: Create a filesystem on an LVM volume

Before we can demonstrate the new online “snapshot” image backup mode we need to create a LVM volume to contain the test file system. The “static” and “dynamic” backup modes will function on a traditional Linux disk partition, but the “snapshot” image backup mode is only possible on an LVM volume.

Install the Logical Volume Manager tools

Before you can create an LVM volume you need to ensure the version of Linux you are using supports LVM and that the LVM tools are installed. Most Linux distributions now come with LVM support built-in, though they may not install the user level tools by default.

Our examples were performed on Mandrake 8.2, which has support for LVM built-in. We manually loaded the LVM tools from the distribution CD-ROM with the series of commands shown in Example 5-67.

Example 5-67 Loading LVM tools

```
[root@azure tmp]# mount /mnt/cdrom
[root@azure tmp]# cd /mnt/cdrom/Mandrake/RPMS
[root@azure RPMS]# rpm -i lvm-1.0.1-1mdk.i586.rpm
```

This step will vary slightly between Linux versions and distributions.

Create the physical volume

You need to identify at least one hard disk or partition to define as a LVM physical volume. In our example we used entire unpartitioned disks. The commands shown in Example 5-68 will create the LVM structures on the identified hard disks.

Example 5-68 Creating LVM Physical Volumes

```
[root@azure root]# pvcreate /dev/sdb
pvcreate -- physical volume "/dev/sdb" successfully created

[root@azure root]# pvcreate /dev/sdc
pvcreate -- physical volume "/dev/sdc" successfully created

[root@azure root]# pvcreate /dev/sdd
pvcreate -- physical volume "/dev/sdd" successfully created
```

It is very important to document your disk volume usage when using LVM physical volumes on entire unpartitioned disks. An novice administrator may incorrectly see the unpartitioned disk as unassigned space and attempt to partition the “unassigned” disk. This will destroy your LVM structure... For this reason it may be preferable to create the LVM physical volumes inside a previously defined entire disk partition. Partition type “0x8e” is reserved for disk partitions containing Linux LVM structures.

Create the volume group

Next we need to group the LVM physical volumes into Volume Groups. Later, this will allow us to span logical volumes across any number of physical volumes in the volume group. The commands shown in Example 5-69 will define a new volume group.

Example 5-69 Creating LVM volume groups

```
[root@azure root]# pvscan
pvscan -- reading all physical volumes (this may take a while...)
pvscan -- ACTIVE   PV "/dev/scsi/host0/bus0/target1/lun0/disc" is in no VG  [2.00 GB]
pvscan -- ACTIVE   PV "/dev/scsi/host0/bus0/target2/lun0/disc" is in no VG  [2.00 GB]
pvscan -- inactive PV "/dev/scsi/host0/bus0/target3/lun0/disc" is in no VG  [2.00 GB]
pvscan -- total: 3 [6.00 GB] / in use: 2 [4.00 GB] / in no VG: 1 [2.00 GB]

[root@azure root]# vgcreate itso_vg /dev/scsi/host0/bus0/target1/lun0/disc
vgcreate -- INFO: using default physical extent size 4.00 MB
vgcreate -- INFO: maximum logical volume size is 255.99 Gigabyte
vgcreate -- doing automatic backup of volume group "itso_vg"
vgcreate -- volume group "itso_vg" successfully created and activated
```

Create a logical volume

A logical volume is the LVM equivalent of a traditional Linux disk partition. Using commands similar to those shown in Example 5-70, create a new logical volume in the new volume group.

Example 5-70 Creating LVM logical volumes

```
[root@azure root]# lvcreate -L 100M -n itso_lv itso_vg
lvcreate -- doing automatic backup of "itso_vg"
lvcreate -- logical volume "/dev/itso_vg/itso_lv" successfully created
```

Note: Be sure to leave enough free space in the volume group to allow for the TSM snapshot cache data. This should be some percentage of the size of the logical volume you created. The default snapshot size is 40% of the logical volume size.

Use the command shown in Example 5-71 to display the volume group structure and list of logical volumes it contains.

Example 5-71 Display LVM volume group structure

```
[root@azure /]# vgdisplay -v
--- Volume group ---
VG Name               itso_vg
VG Access             read/write
VG Status             available/resizable
VG #                  0
MAX LV                255
Cur LV                1
Open LV               0
MAX LV Size           255.99 GB
Max PV                255
Cur PV                2
Act PV                2
VG Size               3.98 GB
PE Size               4.00 MB
Total PE              1020
Alloc PE / Size       25 / 100.00 MB
Free  PE / Size       995 / 3.89 GB
VG UUID               nEkHgU-oOHG-bHrc-RSx8-Y4qd-QeOr-Olif44

--- Logical volume ---
LV Name                /dev/itso_vg/itso_lv
VG Name                itso_vg
LV Write Access        read/write
LV Status              available
LV #                   1
# open                 0
LV Size                100.00 MB
Current LE             25
Allocated LE           25
Allocation             next free
Read ahead sectors     120
Block device           58:0
```

Later, we will see the snapshot image backup will create a new logical volume in the volume group to hold the snapshot cache data.

Create the file system

We now have a LVM logical volume on which we can format a new filesystem. You may use any Linux supported filesystem you desire. The command shown in Example 5-72 details creating an “ext3” filesystem on the new logical volume.

Example 5-72 Creating a filesystem on a LVM logical volume

```
[root@azure root]# mke2fs /dev/itso_vg/itso_lv
mke2fs 1.26 (3-Feb-2002)
Filesystem label=
OS type: Linux
```

```
Block size=1024 (log=0)
Fragment size=1024 (log=0)
25688 inodes, 102400 blocks
5120 blocks (5.00%) reserved for the super user
First data block=1
13 block groups
8192 blocks per group, 8192 fragments per group
1976 inodes per group
Superblock backups stored on blocks:
        8193, 24577, 40961, 57345, 73729

Writing inode tables: done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 35 mounts or
180 days, whichever comes first.  Use tune2fs -c or -i to override.

[root@azure /]# tune2fs -j /dev/itso_vg/itso_lv
tune2fs 1.26 (3-Feb-2002)
Creating journal inode: done
This filesystem will be automatically checked every 21 mounts or
180 days, whichever comes first.  Use tune2fs -c or -i to override.

[root@azure root]# mkdir /itso
[root@azure root]# mount -t ext3 /dev/itso_vg/itso_lv /itso
```

You should probably add this filesystem to your /etc/fstab file to make it mount automatically at system boot time.

Step 3: Perform an Image backup

Finally, after mounting the filesystem on the new LVM logical volume, we can perform an online “snapshot” image backup.

Refer to “Backup Image” in *IBM Tivoli Storage Manager for UNIX Backup-Archive Clients Installation and User’s Guide V5.1*, GC32-0789, for more information regarding the various methods of performing an image backup.

Perform a snapshot image backup via the command line

The purpose of an image backup is to improve backup and recovery time by avoiding the need to create individual files during the backup and recovery process. The online “snapshot” backup has the further advantage of leaving the filesystem available for normal use during the image backup process.

Enter commands similar to those in Example 5-73 to perform an online snapshot image backup.

Example 5-73 Performing a snapshot image backup via the command line

```
[root@azure bin]# dsmc backup image /itso -imagetype=snapshot
Tivoli Storage Manager
Command Line Backup/Archive Client Interface - Version 5, Release 1, Level 0.0 g
3
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Node Name: AZURE_FS
Session established with server TSM_SERVER1: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
```

```
  Server date/time: 03/21/2002 15:31:18  Last access: 03/21/2002 15:27:55


Backup Image Function Invoked.

ANS1303E (RC3)    Client ended transaction
ANS1424W Retrying failed image operation for volume /itso.
Volume -->    101,529,600 /itso [Sent]
Selective Backup processing of '/itso' finished without failure.


Total number of objects inspected:        1
Total number of objects backed up:        1
Total number of objects updated:          0
Total number of objects rebound:          0
Total number of objects deleted:          0
Total number of objects expired:          0
Total number of objects failed:           0
Total number of bytes transferred:   160.00 MB
Data transfer time:                  143.00 sec
Network data transfer rate:        1,145.75 KB/sec
Aggregate data transfer rate:      1,116.63 KB/sec
Objects compressed by:                    0%
Elapsed processing time:           00:02:26
Aborted (core dumped)
```

The snapshot cache area will cache the original disk blocks as any changes occur to the filesystem during the backup process. This cached data will be synchronized after the backup completes.

Perform a snapshot image backup via the GUI interface

The Tivoli Storage Manager B/A client GUI interface contains several new options to support snapshot image backup and recovery.

Enter the **dsm** command to start the Tivoli Storage Manager B/A client GUI. You will need to have an X-Window server to view the display.

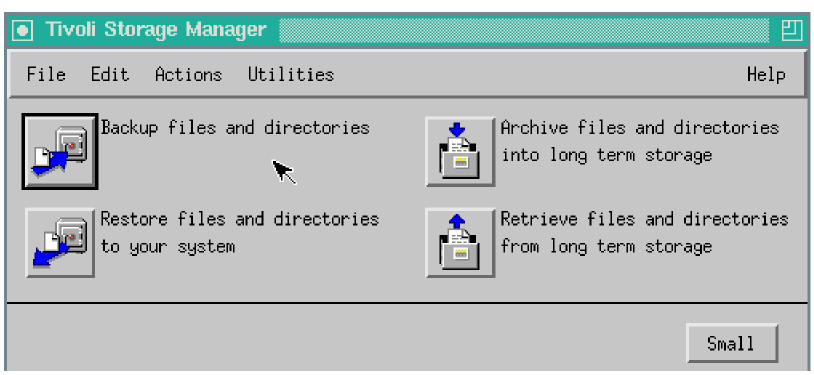

Figure 5-32 The Tivoli Storage Manager B/A client GUI interface on Linux

Select **Backup files and directories** from the initial menu.

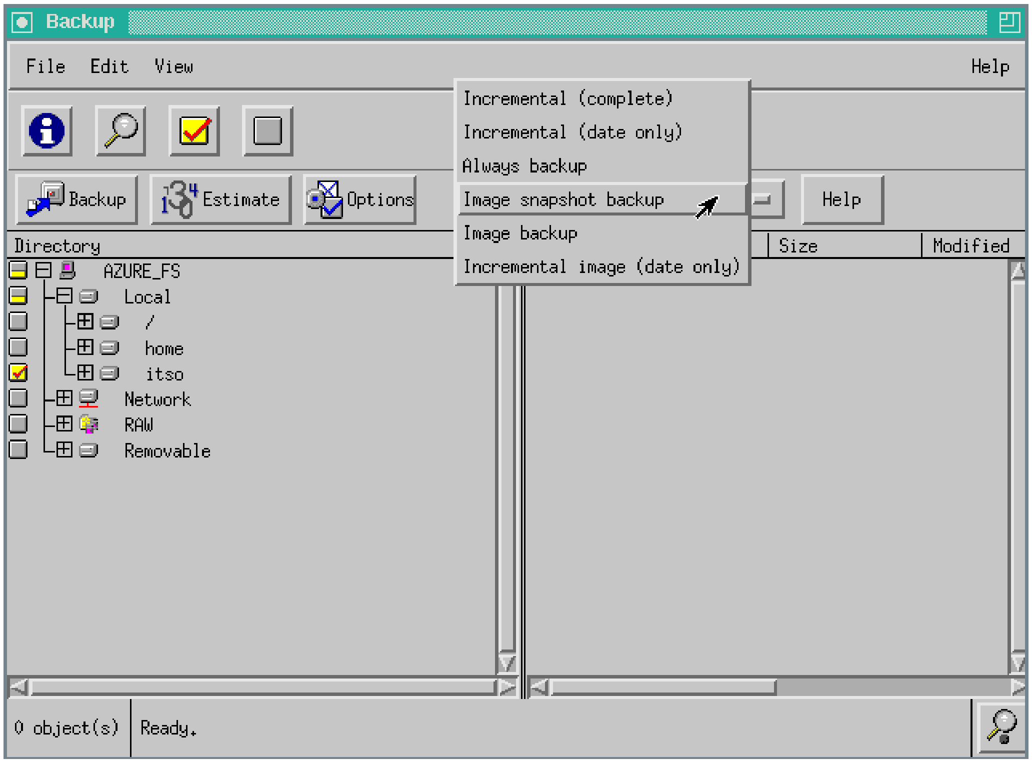

Figure 5-33 Selecting a filesystem to include in the snapshot image backup

Change the backup type drop down menu to display "Image snapshot backup", then select the filesystem/s you wish to include in the backup.

A progress indication screen will appear. When the backup completes the backup status screen will appear, shown in Figure 5-34.

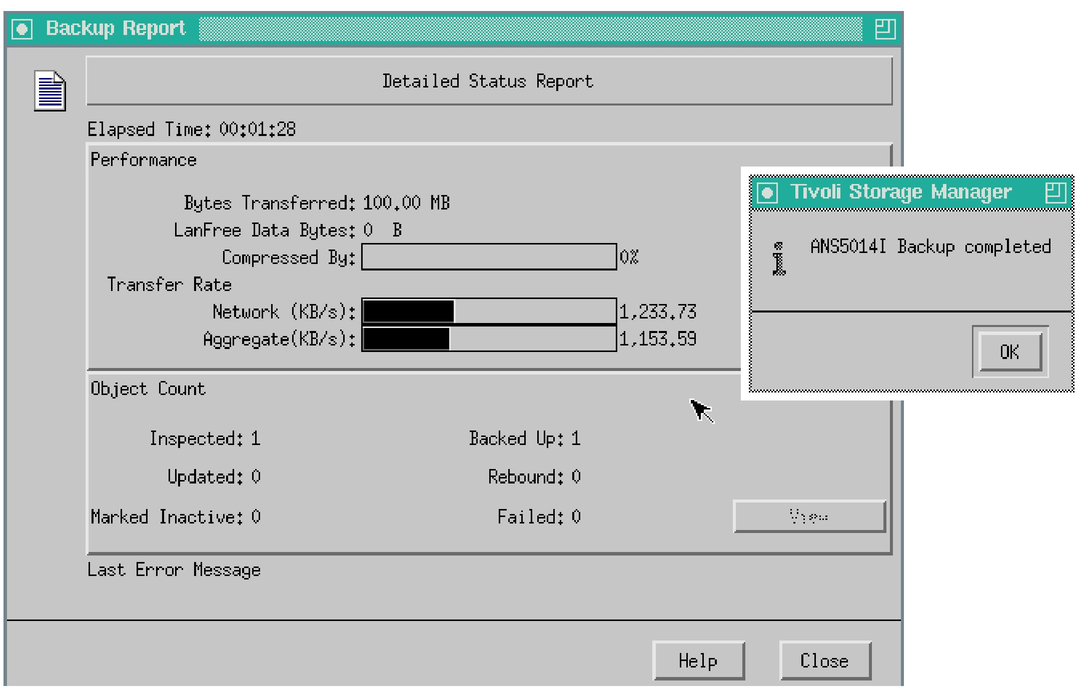

Figure 5-34 The backup status screen

Congratulations, you just performed an online “snapshot” image backup!

Checking the image backup status

You can use the B/A client to query the Tivoli Storage Manager server for a list of available backup images.

Example 5-74 Listing image backups on the Tivoli Storage Manager server

```
[root@azure tmp]# dsmc q image
Tivoli Storage Manager
Command Line Backup/Archive Client Interface - Version 5, Release 1, Level 0.0 g
3
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Node Name: AZURE_FS
Session established with server TSM_SERVER1: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/22/2002 11:32:11  Last access: 03/21/2002 17:18:09

   Image Size FSType     Backup Date     Mgmt Class A/I Image Name
   ---------- ------ ------------------- ---------- --- -------
 1  100.00 MB  EXT3  03/21/2002 16:05:15 IMAGE_MC    A  /itso
Aborted (core dumped)
```

You can also view the LVM volume group structure, which should list the new snapshot cache logical volume.

Example 5-75 Display the snapshot logical volume

```
[root@azure itso_vg]# vgdisplay -v
--- Volume group ---
VG Name               itso_vg
VG Access             read/write
VG Status             available/resizable
VG #                  0
MAX LV                255
Cur LV                2
Open LV               0
MAX LV Size           255.99 GB
Max PV                255
Cur PV                2
Act PV                2
VG Size               3.98 GB
PE Size               4.00 MB
Total PE              1020
Alloc PE / Size       38 / 152.00 MB
Free  PE / Size       982 / 3.84 GB
VG UUID               nEkHgU-oOHG-bHrc-RSx8-Y4qd-QeOr-Olif44

--- Logical volume ---
LV Name                /dev/itso_vg/itso_lv
VG Name                itso_vg
LV Write Access        read/write
LV snapshot status     source of
                       /dev/itso_vg/tsmsnapshot_date_3_21_2002_time_16.0.18 [active]
LV Status              available
LV #                   1
# open                 0
LV Size                100.00 MB
```

```
Current LE            25
Allocated LE          25
Allocation            next free
Read ahead sectors    120
Block device          58:0

--- Logical volume ---
LV Name               /dev/itso_vg/tsmsnapshot_date_3_21_2002_time_16.0.18
VG Name               itso_vg
LV Write Access       read only
LV snapshot status    active destination for /dev/itso_vg/itso_lv
LV Status             available
LV #                  2
# open                0
LV Size               100.00 MB
Current LE            25
Allocated LE          25
snapshot chunk size   4.00 KB
Allocated to snapshot  0.19% [100.00 KB/51.59 MB]
Allocated to COW-table 416.00 KB
Allocation            next free
Read ahead sectors    120
Block device          58:1


--- Physical volumes ---
PV Name (#)          /dev/scsi/host0/bus0/target1/lun0/disc (1)
PV Status            available / allocatable
Total PE / Free PE   510 / 472

PV Name (#)          /dev/scsi/host0/bus0/target2/lun0/disc (2)
PV Status            available / allocatable
Total PE / Free PE   510 / 510
```

Step 4: Perform an Image restore

The purpose of an image backup is to improve backup and recovery time by avoiding the need to create individual files during the backup and recovery process.

Note: Restoring an image backup will replace the **entire** volume contents. This is not an incremental process and will overwrite any data remaining on the target volume.

Several other restrictions apply when restoring an image backup:

- The filesystem to restore must already exist and be mounted before starting the restore process.
- The filesystem to be restored will be automatically unmounted during the image restore process. It will be remounted then the restore is complete.
- Ensure the filesystem or volume to which you are restoring the image is at least the same size as when the backup image was created.
- If you created an image of the root (‘/’) or usr (/usr) filesystem you cannot restore it to the same location.
- You cannot restore an image to the filesystem where the Tivoli Storage Manager B/A client is installed.

Restoring an image backup should only be performed in case of major data loss. It's main purpose is to accelerate recovery in a disaster recovery scenario.

Perform an image restore via the command line

Enter commands similar to those shown in Example 5-76 to restore an image backup.

Example 5-76 Restoring an image backup

```
[root@azure /]# dsmc restore image /itso
Tivoli Storage Manager
Command Line Backup/Archive Client Interface - Version 5, Release 1, Level 0.0 g
3
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Restore Image Function Invoked.

Node Name: AZURE_FS
Session established with server JAMAICA: Windows
  Server Version 5, Release 1, Level 1.0
  Server date/time: 03/24/2002 13:53:22  Last access: 03/24/2002 13:52:55

***************************** WARNING ********************************
Restoring a file system or raw logical volume will replace any data that
currently resides there. Are you sure you want to replace
    File System/Volume: '/itso'?  (Yes (Y)/No (N)) y
Restoring     106,896,384 /itso [Done]

Restore processing finished.

Total number of objects restored:         1
Total number of objects failed:           0
Total number of bytes transferred:   101.94 MB
Data transfer time:                   16.00 sec
Network data transfer rate:        6,524.43 KB/sec
Aggregate data transfer rate:      3,234.24 KB/sec
Elapsed processing time:           00:00:32
```

It is also possible to update the restored filesystem with files modified or deleted since the backup image was created. Enter commands similar to those shown in Example 5-77 to restore an image backup and update the restored filesystem to a point-in-time.

Example 5-77 Restoring an image backup and updating the filesystem automatically

```
[root@azure /]# dsmc restore image /itso -incre -del
...
```

The filesystem will be remounted upon completion of the restore process and should now be ready for use.

Perform an image restore via the GUI interface

The Tivoli Storage Manager B/A client GUI interface contains several new options to support snapshot image backup and recovery.

Enter the **dsm** command to start the Tivoli Storage Manager B/A client GUI. You will need to have an X-Window server to view the display.

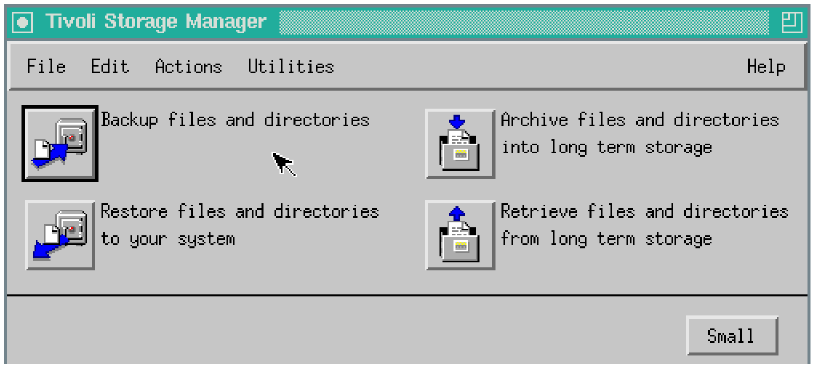

Figure 5-35 The Tivoli Storage Manager B/A client GUI interface on Linux

Select “Restore files and directories” from the initial menu.

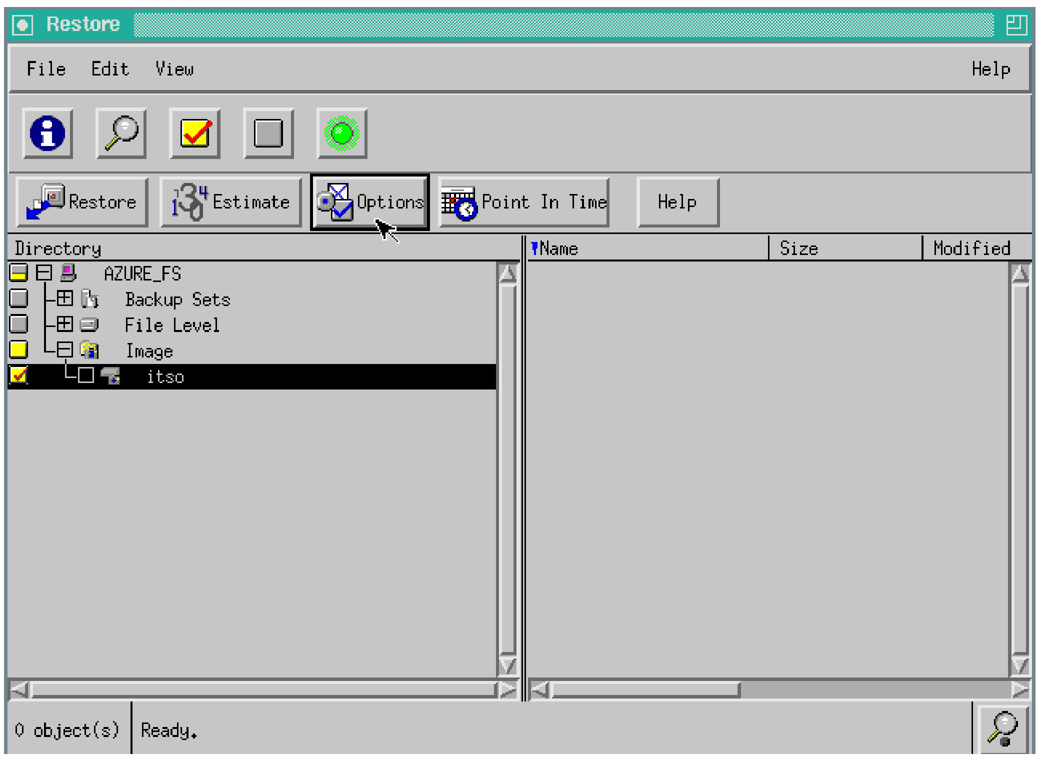

Figure 5-36 Selecting a backup image to restore

Under the Images section, select the backup image you wish to restore. This will overwrite any data on the target volume.

Click the **Options** button to select whether to update the restored filesystem with files modified or deleted since the backup image was created.

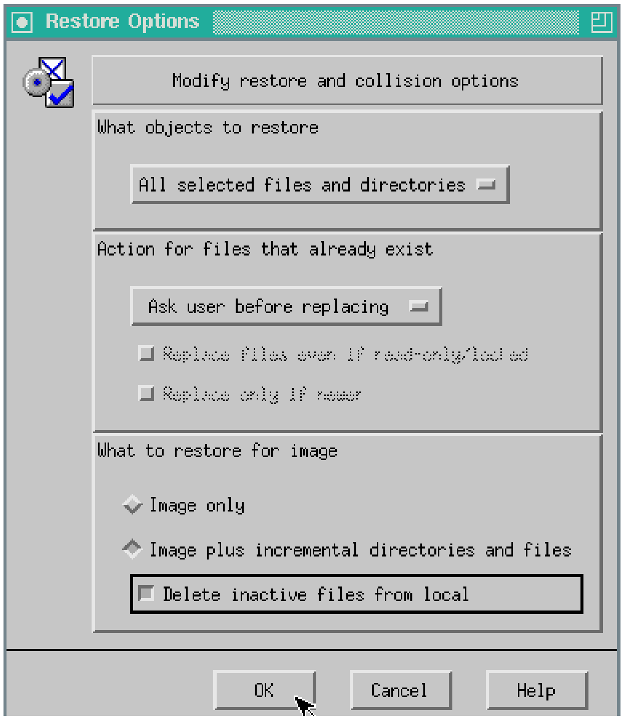

Figure 5-37 Chose whether to update the restored filesystem to a point in time

From the previous screen, click **Restore** to start the restore of the backup image. You will be prompted with the dialog shown in Figure 5-38. The options on this dialog are ignored for the purposes of an image restore. Select **Original location** and click **Restore** to continue.

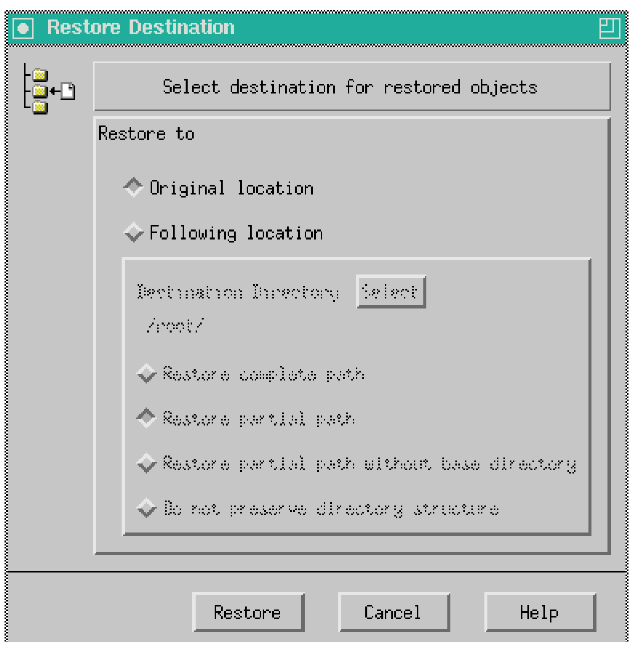

Figure 5-38 Target location dialog for the image restore process

A warning dialog will appear, providing your last change to avoid overwriting the target filesystem with the contents of the backup image.

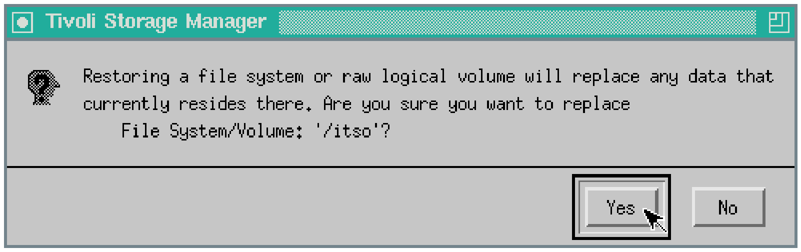

Figure 5-39 Warning dialog for the image restore process

The progress indication screen will appear, shown in Figure 5-40.

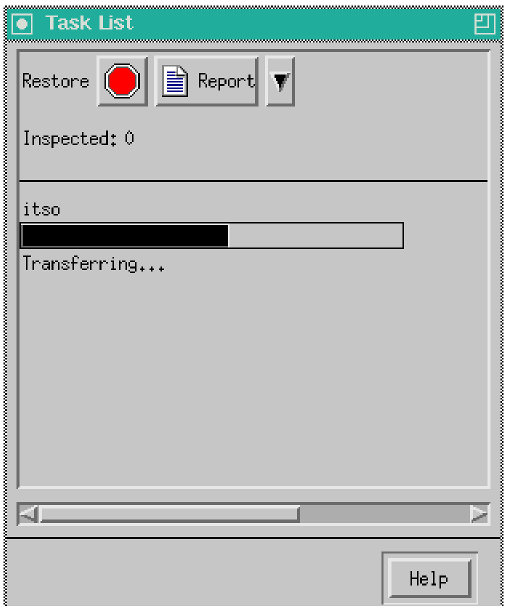

Figure 5-40 Progress indicator for the image restore process

When the restore completes, the status screen will appear, shown in Figure 5-41.

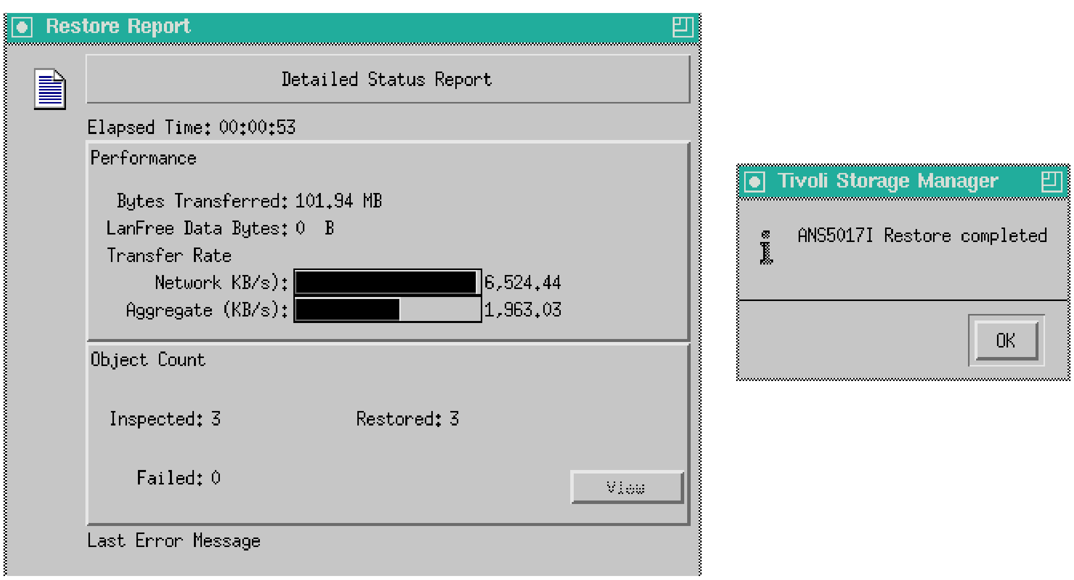

Figure 5-41 The restore status window

Congratulations, you just performed an online "snapshot" image restore!

Planning for system recovery

When using the backup image command, it is important to be aware of some special considerations if you want to do either of the following:

- Perform point-in-time restores in the future.
- Perform an image backup using mode=selective or mode=incremental.

Restoring to a point in time

To ensure that you can perform point-in-time restores of your file systems, including deleting original files which no longer exist on the logical volume, use a combination of full image backups and the incremental command as described in these steps:

1. Perform a full incremental backup of the logical volume, for example:

   ```
   dsmc i /itso
   ```

2. Perform an image backup of the same logical volume, for example:

   ```
   dsmc backup image imagetype=snapshot
   ```

3. Periodically, perform incremental backups, for example:

   ```
   dsmc i /itso
   ```

You must follow these steps in the order shown to ensure that the server records additions and deletions accurately. The following command restores the file system to its exact state as of the last incremental backup:

Example 5-78 Performing a point in time image restore

```
[root@azure root]# dsmc restore image /itso -incre -del
Tivoli Storage Manager
Command Line Backup/Archive Client Interface - Version 5, Release 1, Level 0.0 g3
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Restore Image Function Invoked.

Node Name: TUNGSTEN_FS
Session established with server TSM_SERVER1: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/24/2002 16:06:00  Last access: 03/24/2002 15:56:50

***************************** WARNING ********************************
Restoring a file system or raw logical volume will replace any data that
currently resides there. Are you sure you want to replace
    File System/Volume: '/itso'?  (Yes (Y)/No (N)) y
Restoring     106,896,384 /itso [Done]

Restore processing finished.
Restoring           1,024 /itso/ --> /itso/ [Done]
Restoring             208 /itso/hosts --> /itso/hosts [Done]

Restore processing finished.

Total number of objects restored:         3
Total number of objects failed:           0
Total number of bytes transferred:   101.94 MB
Data transfer time:                   16.00 sec
Network data transfer rate:        6,524.44 KB/sec
Aggregate data transfer rate:      3,142.03 KB/sec
Elapsed processing time:           00:00:33
```

This will allow you to later restore the filesystem image and have the Tivoli Storage Manager client automatically add changed files and remove deleted files.

Using incremental image backups

By using a combination of “backup image” commands, first with “-mode=selective” and later with “-mode=incremental” it is possible to perform a type of “incremental by date” image backup. This will later allow us to restore the original image, plus any files included in the incremental image backups. As with a normal “incremental by date” backup, no information is saved regarding when files are deleted, so any deleted files will also be restored.

The following restrictions apply:

- If logical volumes are running at or near capacity, an out-of-space condition could result during the restore.
- The file system can have no previous full incremental backups produced by the incremental command.
- You can perform the backup only on volumes with mounted file systems; not on raw logical volumes.
- Incremental-by-date image backup does not inactivate files on the server; therefore, when the image is restored, deleted files are restored too.
- If this is the first image backup for the file system, a full image backup is performed.
- Using “-mode=incremental” backs up only files with a changed date, not files with changed permissions.

When restoring from an incremental image backup, the “-deletefiles” option is allowed on the restore image command, but the server ignores it because the server is not aware of any deleted files.

5.4 Solaris

The IBM Tivoli Storage Manager 5.1 client for Solaris introduces some new features and additional support.

5.4.1 Changes in IBM Tivoli Storage Manger Version 5.1 client

Following is a summary of changes since Tivoli Storage Manager Version 5.1:

Solaris

Changes in IBM Tivoli Storage Manger version 5.1 client

- Enhanced Web Client Interface
- Support for the Sun Solaris Quick File System (QFS) file system
- Support for 64-bit Solaris 8 by IBM Tivoli Storage Manager HSM
- Multiple session No Query Restore
- Consistent Client Return Codes

Figure 5-42 Changes in IBM Tivoli Storage Manger Version 5.1 client

- **Enhanced Web Client Interface**

  The Web client interface is enhanced to support a JRE 1.3.1 Swing-enabled browser. This type of browser facilitates the use of assistive devices for users with disabilities. The native look and feel of the platform running the browser is preserved.

- **Support for the Sun Solaris Quick File System (QFS) file system**

  Client Tivoli Storage Manager supports backup, restore, archive and retrieve of the OFS file system on the Solaris client. QFS is a high-performance file system that enables file sharing in a SAN. It eliminates performance bottlenecks resulting from applications using very large file sizes.

- **Support for 64-bit Solaris 8 by IBM Tivoli Storage Manager HSM**

  IBM Tivoli Storage Manager's HSM support for Sun Solaris, which is now part of enterprise edition, now also supports 64-bit Solaris 8. The Tivoli Space Manager function uses hierarchical storage management (HSM) to automatically and transparently migrate rarely accessed files to Tivoli Storage Manager storage while the files most frequently used remain in the local file systems. By migrating rarely accessed files to the server storage, Tivoli Space Manager frees administrators and users from manual file system pruning tasks by enabling you to have sufficient free storage at your workstation or file server, deferring the need to purchase additional disk storage.

- **Multiple session No Query Restore**

  The backup-archive clients can now perform multiple restore sessions for no query restore operations, increasing the speed of restores. This is similar to the multiple backup session support. It exploits the mount point available on the server. If data is backed up on multiple tapes, and if the server has multiple mount points available, then the restore starts a session for each tape, up to the number your administrator configures.

- **Consistent Client Return Codes**

  Reliable, consistent, and documented return codes have been added to the command line client and the scheduler. This facilitates automation of client operations via user-written scripts. Administrators can now distinguish between scheduled backups that completed successfully with no skipped files and scheduled backups that completed successfully with one or more skipped files. Also if the preschedulecmd command ends with non-zero return codes, the scheduled event will not run. This ensures that scheduled events will not run if prerequisite commands do not complete successfully.

5.4.2 Sun Solaris QFS support

Sun Quick File System (QFS) is a high-performance file system that enables file sharing in a storage area network (SAN). QFS was initially developed by LSC, Inc., a company acquired by Sun. Tivoli Storage Manager Backup/Archive 5.1 client is now able to recognize, to backup and to restore a QFS file system.

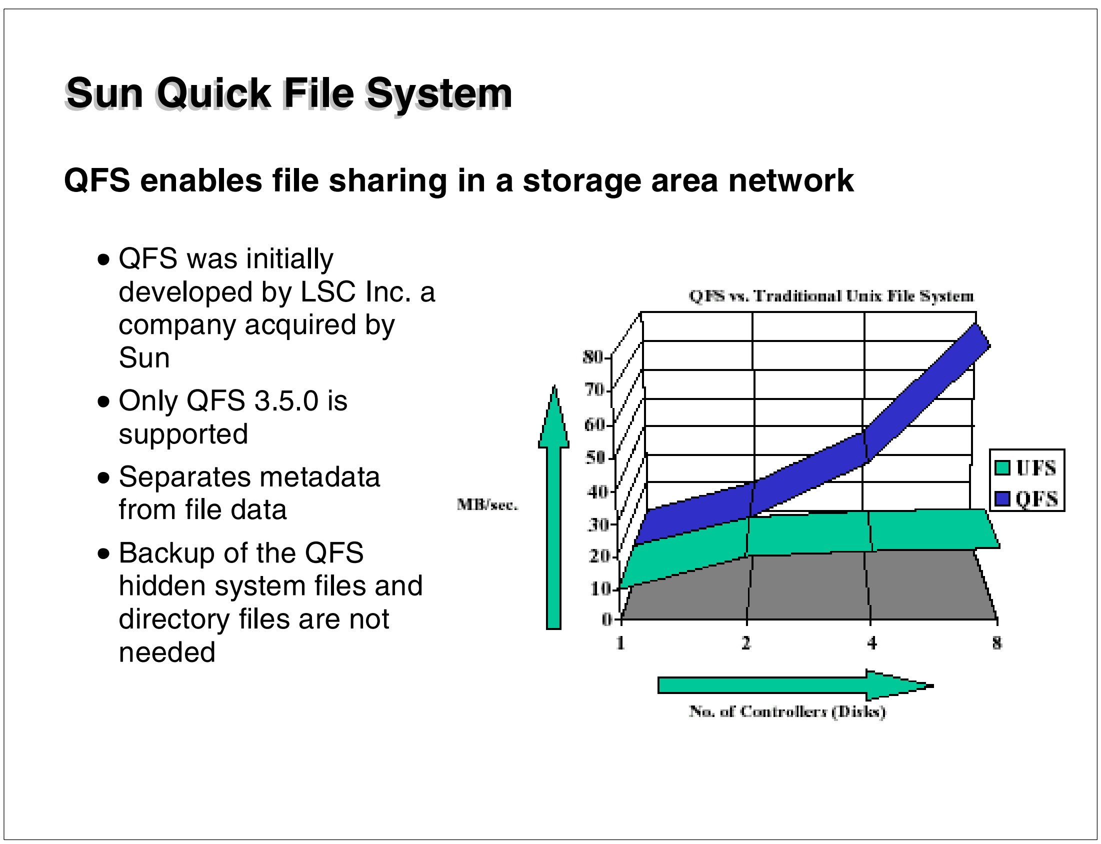

Figure 5-43 Sun QFS

Notes:

- Image Backup of a QFS file system is currently **not** supported.
- Tivoli Storage Manager V5.1 **only** provides support for QFS V3.5.0

Some of the key features of the QFS filesystem is that it separates metadata from file data, allowing drives to complete read/write requests without repositioning; metadata can be accessed and updated simultaneously. It also allow the usage of variable block sizes from 16 KB to 32 MB.

Storage and Archive Manager (SAM-FS)

SAM-FS is a software package provided by Sun (LSC Inc.) that provides storage management, archival and retrieval services. SAM-FS automatically and transparently copies files from expensive online disk to less expensive automated storage media, and restores the files back online as needed. It is essentially a counterpart of the HSM component of Tivoli Storage Manager.

This is very important to consider, because SAM-FS has the same file system type as QFS (both are samfs). Therefore a Tivoli Storage Manager backup of this file system is not recommended because files would have to be recalled from tape media to disk in order to perform a Tivoli Storage Manager backup. This could be cause performance problems in customer environments. For this reason, a warning is issued at Tivoli Storage Manager installation time if SAM-FS is detected.

Important: Tivoli Storage Manager does not support the combination of QFS and SAM-FS.

Setting up QFS

Here are the steps to set up QFS.

1. Verify if the QFS package is installed, using the **pkginfo** command, which is shown in Example 5-79.

Example 5-79 Verify if the QFS package is installed

```
root@solemio:/: pkginfo LSCqfs
system      LSCqfs         QFS Solaris 2.8
```

2. Determine which version of QFS is installed, using the **pkgparam** command, which is shown in Example 5-80.

Example 5-80 Verify which QFS version is installed

```
root@solemio:/: pkgparam -v LSCqfs | grep VERSION
VERSION='3.5.0-31'
```

Note: Only QFS 3.5.0 is supported.

3. Configure a QFS file system, in the following examples we will use the name `qfs1`. At least two disk device or partitions are required for each QFS file system, one for metadata, and one or more for file data. In Example 5-81 we examine the QFS configuration file `/etc/opt/LSCsamfs/mcf`.

Example 5-81 Examining the QFS configuration

```
root@solemio:/: cat /etc/opt/LSCsamfs/mcf
```

```
# QFS disk cache configuration
#
# Equipment     Eq      Eq      Fam.    Dev.    Additional
# Identifier    Ord     Type    Set     State   Parameters
# ----------    ---     ----    ----    -----   ----------
#
        qfs1        10      ma      qfs1
/dev/dsk/c0t1d0s6       11      mm      qfs1        on      /dev/rdsk/c0t1d0s6
/dev/dsk/c0t1d0s7       12      mr      qfs1        on      /dev/rdsk/c0t1d0s7
```

The `Eq Type` in the Example 5-81 can be either `ma`, `mm` or `mr`:

ma A QFS file System
mm A metadata device
mr A round robin or striped data device

4. Add the filesystem. The next sample in Example 5-82 shows the result after it has been added to the `vfstab`.

Example 5-82 Examining the filesystem mount configuration

```
root@solemio:/: cat /etc/vfstab
#device         device          mount           FS      fsck    mount   mount
#to mount       to fsck         point           type    pass    at boot options
#
#/dev/dsk/c1d0s2 /dev/rdsk/c1d0s2 /usr          ufs     1       yes     -
fd      -       /dev/fd fd      -       no      -
/proc   -       /proc   proc    -       no      -
/dev/dsk/c0t1d0s1       -       -       swap    -       no      -
/dev/dsk/c0t1d0s0       /dev/rdsk/c0t1d0s0      /       ufs     1       no      -
/dev/dsk/c0t1d0s4       /dev/rdsk/c0t1d0s4      /usr    ufs     1       no      -
/dev/dsk/c0t1d0s5       /dev/rdsk/c0t1d0s5      /home   ufs     2       yes     -
/dev/dsk/c0t1d0s3       /dev/rdsk/c0t1d0s3      /opt    ufs     2       yes     -
#------------------------------------------------------
# Warning: The following devices are handle by QFS.
#/dev/dsk/c0t1d0s6 - reserved for metadata of qfs1
#/dev/dsk/c0t1d0s7 - reserved for data     of qfs1
#------------------------------------------------------
qfs1        -       /qfs1       samfs   -       yes     stripe=0
```

5. Mount the filesystem with the **`mount`** command:

```
mount qfs1
```

Backing up QFS data

To test filesystem backup, use the **`dsmc`** command with the **`select`** command option as in the following sample in Example 5-83.

Example 5-83 Backup using the dsmc command

```
root@solemio:/: dsmc sel "/qfs1/*" -subdir=yes
Tivoli Storage Manager
Command Line Backup/Archive Client Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Selective Backup function invoked.

Node Name: MOONSHINE
Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/05/02   14:29:26  Last access: 03/05/02   14:29:04
```

```
Directory-->              4,096 /qfs1/ [Sent]
Directory-->              4,096 /qfs1/DirAndFiles [Sent]
Normal File-->          624,447 /qfs1/DirAndFiles.tar.Z [Sent]
Directory-->              4,096 /qfs1/DirAndFiles/Dir1 [Sent]
Directory-->              4,096 /qfs1/DirAndFiles/Dir2 [Sent]
Directory-->              4,096 /qfs1/DirAndFiles/Dir3 [Sent]
Normal File-->        1,048,576 /qfs1/DirAndFiles/file1 [Sent]
...<lines omitted>...
Selective Backup processing of '/qfs1/*' finished without failure.
Total number of objects inspected:      15
Total number of objects backed up:      15
Total number of objects updated:         0
Total number of objects rebound:         0
Total number of objects deleted:         0
Total number of objects expired:         0
Total number of objects failed:          0
Total number of bytes transferred:   26.59 MB
Data transfer time:                   4.30 sec
Network data transfer rate:       6,325.21 KB/sec
Aggregate data transfer rate:     4,482.80 KB/sec
Objects compressed by:                   0%
Elapsed processing time:          00:00:06
```

A QFS file system contains two hidden system files and a system directory that cannot be backed up. A backup of these files is not needed. They contain internal data to manage the file system. This data will be automatically excluded from a backup and recreated automatically by the file system itself if a restore of files is invoked. The directory listing in Example 5-84 shows some of the QFS files and directories. The hidden files are prefixed with a dot (.).

Example 5-84 Examining the content of the sample QFS filesystem

```
root@solemio:/: ls -al /qfs1
total 2212
drwxr-xr-x   4 root     root        4096 Aug 14 20:40 .
drwxr-xr-x  38 root     other       1536 Aug 14 20:41 ..
drwx------   2 root     root        4096 Aug  7 16:34 .archive
-r--------   1 root     root      442368 Aug  8 12:28 .inodes
-r--r--r--   1 root     root           0 Aug  7 16:34 .ioctl
drwxr-x---   5 root     other       4096 Jan 30  2002 DirAndFiles
-rw-r-----   1 root     other     624447 Aug 14 20:39 DirAndFiles.tar.Z
```

The Sun node (`MOONSHINE`) filespace on the Tivoli Storage Manager server is marked as QFS system type and can be viewed on the server with the **dsmadmc** command as the next sample in Example 5-85 shows.

Example 5-85 Examining the filespace of the sample QFS filesystem client node

```
root@crete:/: dsmadmc -id=admin -password=admin query fi moonshine f=d
Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/05/02   14:39:57  Last access: 03/05/02   14:39:53
```

```
ANS8000I Server command: 'query fi moonshine f=d'

                     Node Name: MOONSHINE
                Filespace Name: /qfs1
      Hexadecimal Filespace Name:
                          FSID: 4
                      Platform: SUN SOLARIS
                Filespace Type: QFS
           Is Filespace Unicode?: No
                 Capacity (MB): 501.6
                      Pct Util: 7.2
     Last Backup Start Date/Time:
  Days Since Last Backup Started:
Last Backup Completion Date/Time:
Days Since Last Backup Completed:
Last Full NAS Image Backup Completion Date/Time:
Days Since Last Full NAS Image Backup Completed:

ANS8002I Highest return code was 0.
```

Restoring QFS data

Restoring the data can be performed with the **dsmc** command and the **restore** command option as shown in Example 5-86.

Example 5-86 Restore QFS data with the dsmc command

```
root@solemio:/: dsmc res "/qfs1/*" -su=yes
Tivoli Storage Manager
Command Line Backup/Archive Client Interface - Version 5, Release 1, Level 0.0
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Restore function invoked.

Node Name: MOONSHINE
Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 1, Level 0.0
  Server date/time: 03/05/02   14:29:26  Last access: 03/05/02   14:29:04

ANS1247I Waiting for files from the server...
Restoring          4,096 /qfs1/ [Done]
Restoring          4,096 /qfs1/DirAndFiles [Done]
Restoring          4,096 /qfs1/DirAndFiles/Dir1 [Done]
Restoring          4,096 /qfs1/DirAndFiles/Dir2 [Done]
Restoring          4,096 /qfs1/DirAndFiles/Dir3 [Done]
Restoring        624,447 /qfs1/DirAndFiles.tar.Z [Done]
Restoring      1,048,576 /qfs1/DirAndFiles/file1 [Done]
Restoring      1,048,576 /qfs1/DirAndFiles/file2 [Done]
Restoring      1,048,576 /qfs1/DirAndFiles/file3 [Done]
Restoring      1,048,576 /qfs1/DirAndFiles/Dir1/file1 [Done]
Restoring      1,048,576 /qfs1/DirAndFiles/Dir2/file1 [Done]
Restoring      5,242,880 /qfs1/DirAndFiles/Dir2/file2 [Done]
Restoring      1,048,576 /qfs1/DirAndFiles/Dir3/file1 [Done]
Restoring      5,242,880 /qfs1/DirAndFiles/Dir3/file2 [Done]
Restoring     10,485,760 /qfs1/DirAndFiles/Dir3/file3 [Done]

Restore processing finished.

Total number of objects restored:       15
Total number of objects failed:          0
```

```
Total number of bytes transferred:    26.59 MB
Data transfer time:                    2.38 sec
Network data transfer rate:        11,432.91 KB/sec
Aggregate data transfer rate:      2,998.55 KB/sec
Elapsed processing time:           00:00:11
```

The sample output in Example 5-86 shows that all data is correctly restored. The hidden files are recreated automatically by the file system itself.

This section has described an overview of QFS. For further detailed information on QFS, see:

http://www.sun.com/storage/software/storage_mgmt/qfs/index.html

5.5 Windows

The IBM Tivoli Storage Manager 5.1 client for Windows introduces some new features and additional support.

5.5.1 Changes in IBM Tivoli Storage Manger Version 5.1 client

The following is a summary of changes introduced by Tivoli Storage Manager Version 5.1 Windows client:

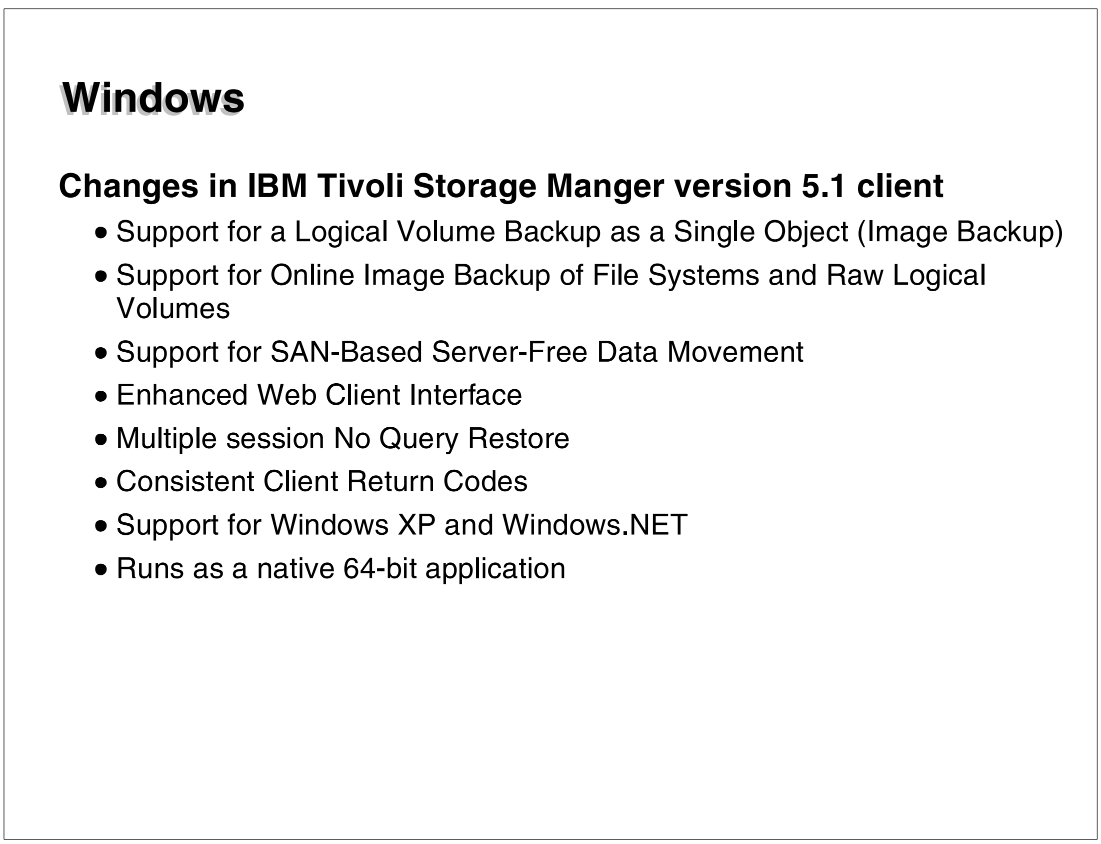

Figure 5-44 Changes in IBM Tivoli Storage Manger Version 5.1 client

- Support for a Logical Volume Backup as a Single Object (Image Backup)

  The Windows 2000 client is enhanced to support a logical volume image backup of file systems and raw volumes. You can use the native GUI, the command line interface, and the Web client GUI to create an image of file systems and raw logical volumes on the Tivoli Storage Manager server. NTFS, FAT, FAT32 file systems and raw logical volumes are supported. The Tivoli Storage Manager server does not track individual files in the file system image. File system images are tracked as individual objects and management class policy will be applied to the file system image as a whole.

- Support for Online Image Backup of File Systems and Raw Logical Volumes

  The traditional offline image backup prevents access to the volume by other system applications during the operation. For Windows 2000 only: If the Logical Volume Snapshot Agent (LVSA) is installed and configured, Tivoli Storage Manager by default performs an online image backup, during which the volume is available to other system applications.

- Support for SAN-Based Server-Free Data Movement

  The Windows 2000 client supports SAN-based Server-Free Data Movement that off-loads data movement processing from the client and server processor and from the LAN during image backup and restore operations. Data is transferred directly between client disks and SAN-attached storage devices by a third-party copy function initiated by the Tivoli Storage Manager server. This functionality is expected to be made available by 3rd quarter 2002.

- Enhanced Web Client Interface

  The Web client interface is enhanced to support a JRE 1.3.1 Swing-enabled browser. This type of browser facilitates the use of assistive devices for users with disabilities. The native look and feel of the platform running the browser is preserved.

- Multiple session No Query Restore

  The backup-archive clients can now perform multiple restore sessions for no query restore operations, increasing the speed of restores. This is similar to the multiple backup session support. It exploits the mount point available on the server. If data is backed up on multiple tapes, and if the server has multiple mount points available, then the restore starts a session for each tape, up to the number your administrator configures.

- Consistent Client Return Codes

  Reliable, consistent, and documented return codes have been added to the command line client and the scheduler. This facilitates automation of client operations via user-written scripts. Administrators can now distinguish between scheduled backups that completed successfully with no skipped files and scheduled backups that completed successfully with one or more skipped files. Also if the preschedulecmd command ends with non-zero return codes, the scheduled event will not run. This ensures that scheduled events will not run if prerequisite commands do not complete successfully.

- Support for Processing Network Attached Storage (NAS) File System Images

  Tivoli Storage Manager supports backup and restore of network attached storage (NAS) file system images to tape drives or libraries that are locally attached to NAS file servers. Tivoli Data Protection for NDMP enables backup and restore support on the Tivoli Storage Manager Windows NT, 2000, XP, and Windows.Net servers for NAS file servers from Network Appliance. Tivoli Data Protection for NDMP is a separately priced and licensed product.

5.5.2 New look and feel

The Windows client GUI and Web GUI have a new look as shown in Figure 5-45.

Figure 5-45 Tivoli Storage Manager 5.1 GUI and Web GUI

5.5.3 Logical volume backup (image backup)

The Tivoli Storage Manager Windows 2000 client is enhanced to support a logical volume image backup of file systems and raw volumes. The image backup and restore functionality is new to Tivoli Storage Manager on Windows 2000, but has been part of some of the Tivoli Storage Manager UNIX clients since Version 3.7 (AIX, Solaris, HP).

Overview

The traditional *offline* image backup prevents access to the volume by other system applications during the operation. If the Logical Volume Snapshot Agent (LVSA) is installed and configured, the Tivoli Storage Manager client performs an *online* image backup, during which the volume is available to other system applications.

An image backup provides the following benefits:

- Improves the speed with which Tivoli Storage Manager backs up and restores file systems containing many small files.
- Conserves resources on the server during backups since only one entry is required for the image.
- Provides a point-in-time picture of your file system, which may be useful if your enterprise needs to recall that information.
- Restores a corrupt file system or raw logical volume. Data is restored to the same state it was when the last logical volume backup was performed.

Supported volumes are:

- Formatted (NTFS, FAT32 and FAT)
- Unformatted (RAW)

The Tivoli Storage Manager server does not track individual files in the file system image. File system images are tracked as individual objects and management class policy will be applied to the file system image as a whole.

Note: On Windows 2000 NTFS formatted volumes, Tivoli Storage Manager will optionally transfer just the used blocks of the volume.

There are three different types of Image Backup available:

- **Image Backup (offline)**

  The volume being backed up is made unavailable to all other applications such as Windows Explorer, CMD shells, Database Software, etc. It is not even available for reading operations on the volume like in its UNIX offline (static) counterpart. The Backup-Archive client is the only one with access to the volume. Once the image backup is complete the Backup-Archive client makes the volume available to system and other applications again.

- **Image Snapshot Backup (online)**

  The volume being backed up remains available to all applications during backup. The backup consists of a point in time snapshot of the entire volume. Online is the default used by the Backup-Archive client unless otherwise specified.

- **Incremental Image (date-only**)

  Backs up only those files whose modification dates and times are newer than the date and time of the last image backup. You can perform incremental image (date-only) of last image backup regardless of whether the full image was backed up offline or online.

You can use the native GUI, the command line interface, and the Web client GUI to create an image of file systems and raw logical volumes on the Tivoli Storage Manager server.

Related Options

Here are some related options and values.

- **DOMAIN.IMAGE**

  This is not a new option, but it now applies to Windows 2000 Clients. This Option allows users to enter a set of volumes which are backed up as part of the client machine's image domain. I used when the user enters the command without any volume specified (dsmc backup image) and from GUIs via the Backup Domain Image action menu.

  ```
  DOMAIN.IMAGE d: e:
  ```

- **INCLUDE.IMAGE**

  This is also not a new option, but it now applies to Windows 2000 Clients. It is enhanced to allow users to set different image options per volume. This option will accept the following values:

  – **IMAGETYPE**

    • SNAPSHOT: Specifies that you want to perform an online image backup during which the volume is available to other system applications. This is the default if the Logical Volume Snapshot Agent (LVSA) is installed and configured.

    • STATIC: Specifies that you want to perform an offline image backup during which the volume is locked and unavailable to other system applications.

    ```
    IMAGETYPE=static
    ```

  – **SNAPSHOTCACHESIZE**

Specifies the maximum size of the OBF file (Old Blocks File), in which old data blocks are stored during an online image backup. The value is a percent of the total size of the volume being backed up. The range of values is one through 100 percent; the default is 100 percent.

```
SNAPSHOTCACHESIZE 40
```

– **SNAPSHOTCACHELOCATION**

Specifies the valid path to the location where the LVSA will place OBF file (Old Blocks File) during an online image backup. The default location is C:\TSMLVSA.

```
SNAPSHOTCACHELOCATION c\temp
```

– **SNAPSHOTFSIDLEWAIT**

Use this option with the backup image command or the include.image option when performing an online image backup to specify the amount of time that must pass in which there is no write activity on a volume before a snapshot is taken. The range of values is zero through 999; the default is 1. You may specify ms (milliseconds) or s (seconds) qualifiers with the value.

```
SNAPSHOTFSIDLEWAIT=500ms
```

– **SNAPSHOTFSIDLERETRIES**

Use this option with the backup image command or the include.image option when performing an online image backup to specify the number of additional times the LVSA should try to achieve the Snapshot FS Idle Wait time before the online image backup operation fails. The range of values is zero through 99; the default is 10.

```
SNAPSHOTFSIDLERETRIES=20
```

– **PRESNAPSHOTCMD**

Allows you to quiesce an application before the LVSA starts the snapshot during an online image backup.

```
PRESNAPSHOTCMD “<command to set application to quiesce state>”
```

– **POSTSNAPSHOTCMD**

Allows you to bring up an application after the LVSA starts a snapshot during an online image backup. See Postsnapshotcmd for more information.

```
POSTSNAPSHOTCMD “<command to restore application to normal state>”
```

– **IMAGEGAPSIZE**

Specifies the minimum size of empty regions in a formatted logical volume that should be skipped during an image backup. This option is only valid only for NTFS formatted volumes. You may specify k (kilobytes) m (megabytes) or g (gigabytes) qualifiers with the value. Without a qualifier, the value is interpreted in kilobytes. Valid values are zero through 4294967295. If you specify a value of 0, all blocks, including unused blocks at the end of the volume, will be backed up. For LAN-based and LAN-free image backup the value is 32k. For server-free image backup the default value is 5m.

```
IMAGEGAPSIZE 1m
```

All options of the INCLUDE.IMAGE statement can be used:

- In the dsm.opt file as a global option itself (except the IMAGETYPE option!), then that value overrides the default value.
- In the INCLUDE.IMAGE statement itself, then that value overrides the global options of the dsm.opt and the default value for that volume only.

- In the DSMC BACKUP IMAGE command, then that value overrides the INCLUDE.IMAGE, dsm.opt and default values for the volumes being backed up in that command (Example 5-88).

Example 5-87 Setup dsm.opt file for Logical Volume Backup

```
...
IMAGETYPE static
SNAPSHOTCACHESIZE 60
INCLUDE.IMAGE g: lvbmc IMAGETYPE=snapshot SNAPSHOTCACHESIZE=40
...
```

Example 5-88 CLI command to perform a logical image backup

```
dsmc backup image g: -imagetype=snapshot -snapshotcachesize=40
```

Policy for Logical Volume Backups

Consider defining a management class specifically for logical volume backups. To enable clients to restore a logical volume and then reconcile the results of any file backup operations since the logical volume backup was made, you must set up management classes with the backup copy group set up differently from the STANDARD. The Versions Data Exists, Versions Data Deleted, and Retain Extra Versions parameters work together to determine over what time period a client can restore a logical volume image and reconcile later file backups. Also, you may have server storage constraints that require you to control the number of backup versions allowed for logical volumes.

Backups of logical volumes are intended to help speed the restoration of a machine. One way to use the capability is to have users periodically (for example, once a month) perform a logical volume backup, and schedule daily full incremental backups. If a user restores a logical volume, the program first restores the logical volume backup and then any files that were changed since the backup (incremental or other file backup processes). The user can also specify that the restore process reconcile any discrepancies that can result when files are deleted.

For example, a user backs up a logical volume, and the following week deletes one or more files from the volume. At the next incremental backup, the server records in its database that the files were deleted from the client. When the user restores the logical volume, the program can recognize that files have been deleted since the backup was created. The program can delete the files as part of the restore process. To ensure that users can use the capability to reconcile later incremental backups with a restored logical volume, you need to ensure that you coordinate policy for incremental backups with policy for backups for logical volumes.

For example, you decide to ensure that clients can choose to restore files and logical volumes from any time in the previous 30 days. You can create two management classes, one for files and one for logical volumes. Table 5-2 shows the relevant parameters. In the backup copy group of both management classes, set the Retain Extra Versions parameter to 30 days.

In the management class for files, set the parameters so that the server keeps versions based on age rather than how many versions exist. More than one backup version of a file may be stored per day if clients perform selective backups or if clients perform incremental backups more than once a day. The Versions Data Exists parameter and the Versions Data Deleted parameter control how many of these versions are kept by the server. To ensure that any number of backup versions are kept for the required 60 days, set both the Versions Data Exists parameter and the Versions Data Deleted parameter to NOLIMIT for the management class for files. This means that the server retains backup versions based on how old the versions are, instead of how many backup versions of the same file exist.

Note: For logical volume backups, the server ignores the frequency attribute in the backup copy group.

Table 5-2 Example of a backup policy for files and logical volumes

Parameter (backup copy group in the management class)	Management Class for Files	Management Class for Logical Volumes
Versions Data Exists	NOLIMIT	3 versions
Versions Data Deleted	NOLIMIT	1 version
Retain Extra Versions	30 days	30 days
Retain Only Version	60 days	60 days

Installation of the Logical Volume Snapshot Agent (LVSA)

To install the Logical Volume Snapshot Agent on your Windows 2000 client system start the Tivoli Storage Manager client GUI and click **Utilities —> Setup Wizard**, as shown in Figure 5-46.

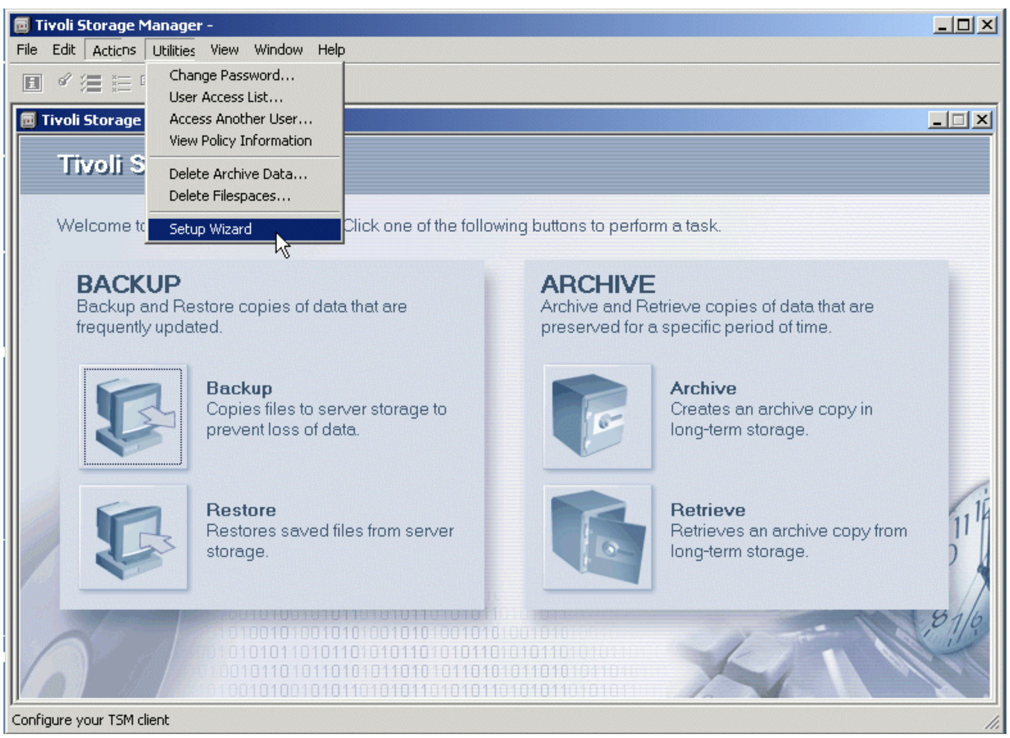

Figure 5-46 Tivoli Storage Manager GUI

Select the checkbox at the bottom and click **Next** (Figure 5-47).

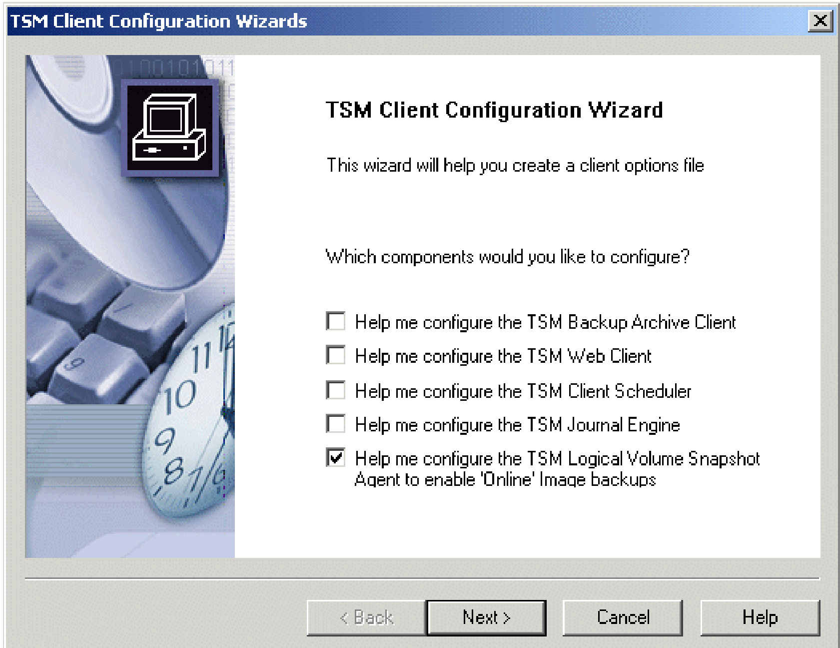

Figure 5-47 Tivoli Storage Manager client configuration wizards window

We want to perform an initial installation of the LVSA, therefore click **Next**. See Figure 5-48.

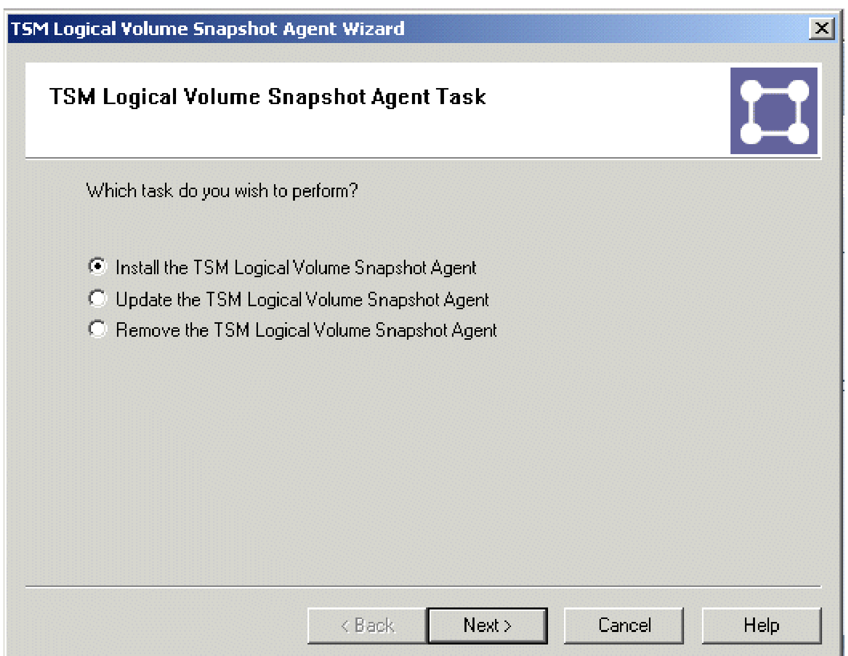

Figure 5-48 Tivoli Storage Manager logical volume snapshot agent wizard window

To finish the installation click **Finish**.

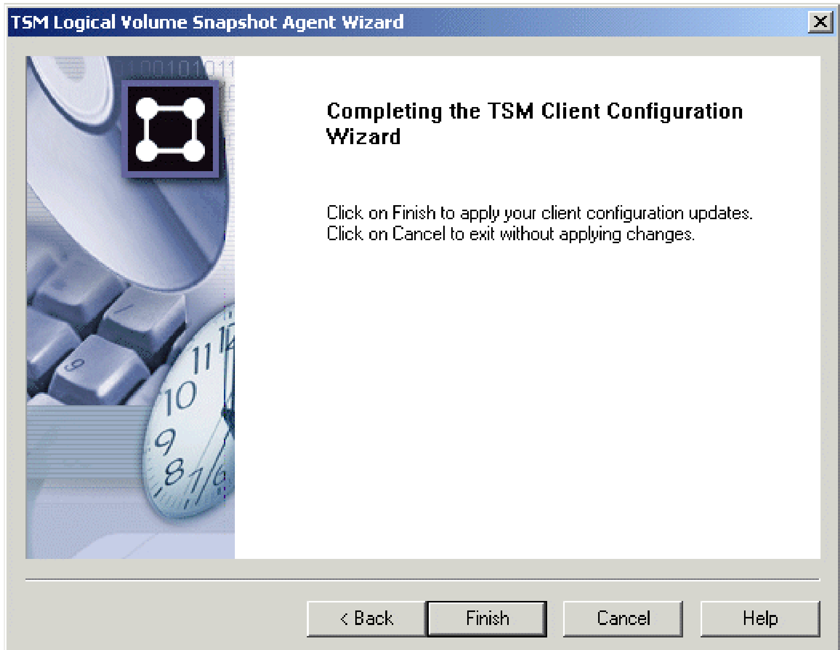

Figure 5-49 Tivoli Storage Manager logical volume snapshot agent wizard window

After the successful installation of the LVSA the following message window appears (Figure 5-50).

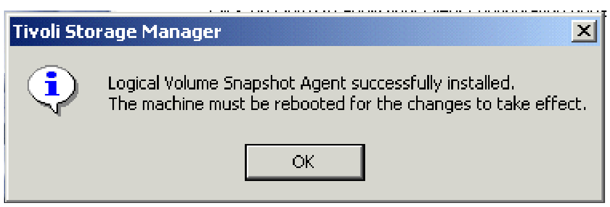

Figure 5-50 Message window

The machine has to be rebooted in order to load the LVSA kernel filter driver 'tsmlvsa.sys'.

Logical Volume Backup via GUI

Before performing a backup via GUI please check if all the backup preferences, especially for the Logical Volume Backup, are correct. Therefore open the preferences window in the GUI as shown in Figure 5-51.

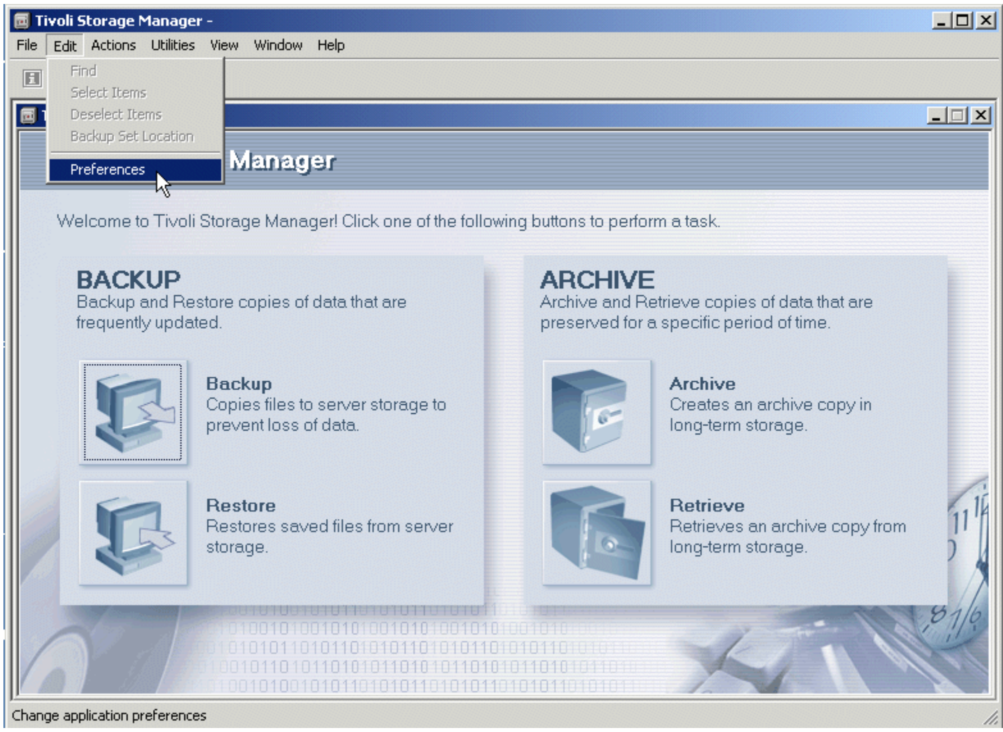

Figure 5-51 Setup backup preferences

For Logical Volume **Backup** and the **Image-Snapshot** tabs are of interest. At the **Backup** tab you have to select the Domain for the backup, as shown in Figure 5-52. This setting is necessary, otherwise you are not able to change the settings at the **Image-Snapshot** tab.

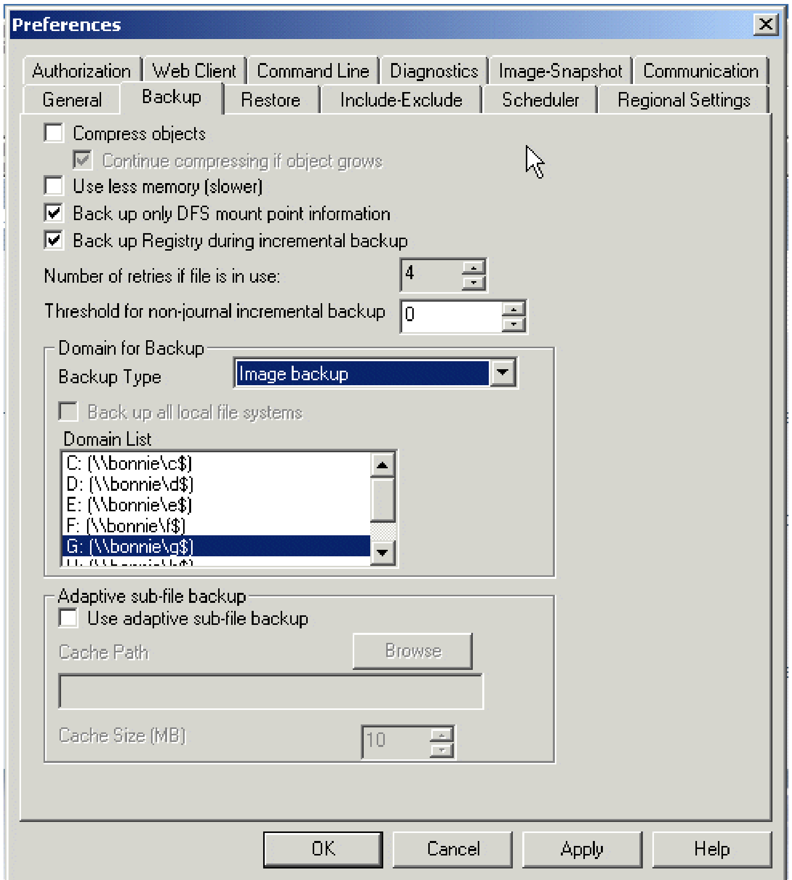

Figure 5-52 Domain setting for logical image backup

In the **Image-Snapshot** tab you can setup the options which are related to the Logical Volume Backup functions image backup (offline) and image snapshot backup (online). See Figure 5-53.

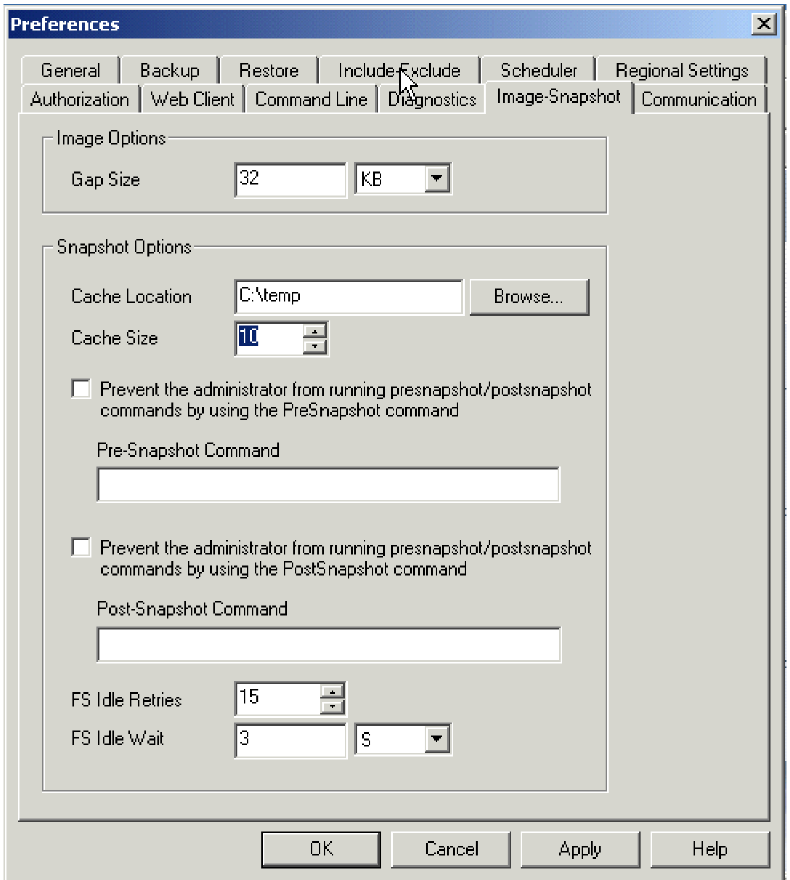

Figure 5-53 Logical volume backup settings

After you have changed the settings you have to restart the GUI to take effect of your new preferences.

In the backup window you can choose:

- Image snapshot backup (online)
- Image backup (offline)
- Incremental image (date-only)

To perform an image backup (offline) please select **Image backup**, as shown in Figure 5-54.

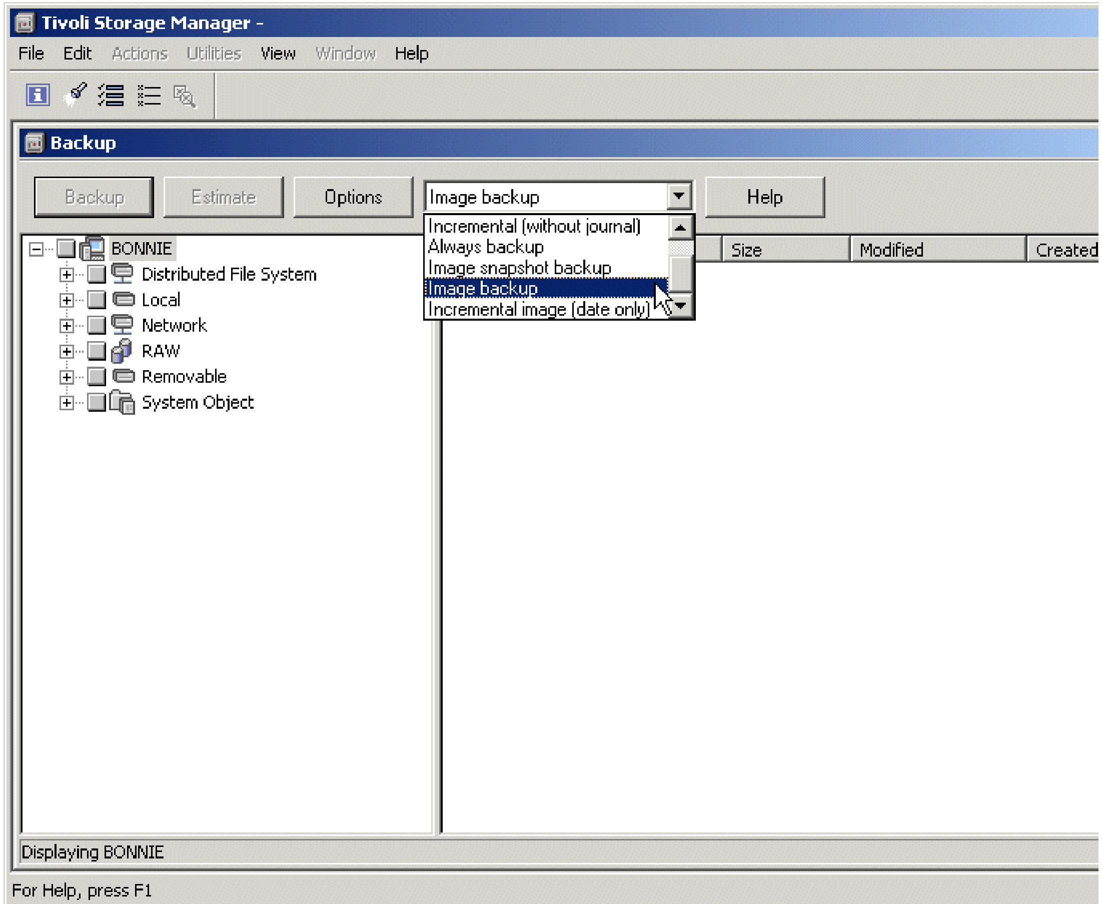

Figure 5-54 Selection of backup type

Now it is time to decide which of the volumes you want to backup. In this example we want to perform an image backup of volume g:. It is a NTFS formatted volume with a size of 500 MB. You also have the choice to select a RAW device, as shown in Figure 5-55.

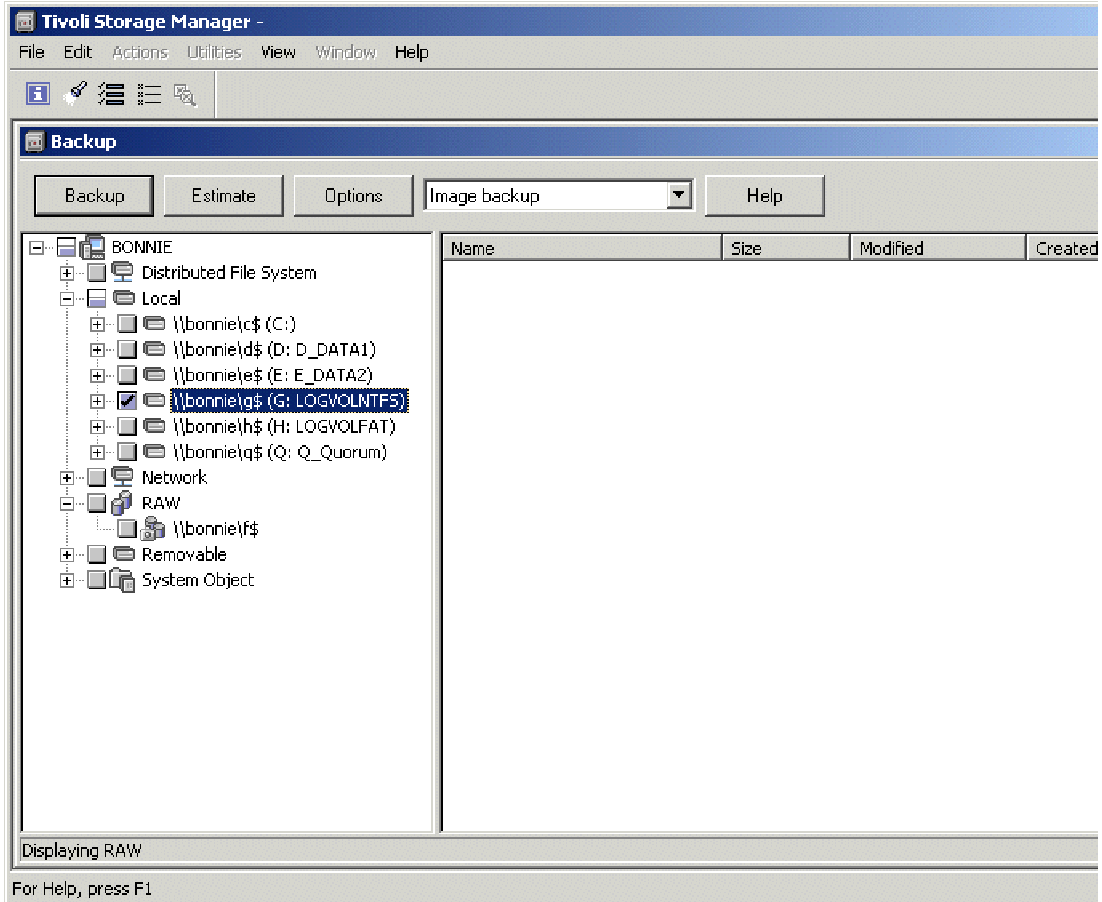

Figure 5-55 Volume selection for Logical Volume Backup

The backup Task List window shows the state of the backup of volume g: (Figure 5-56).

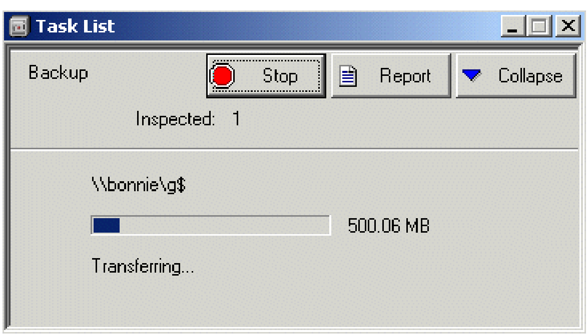

Figure 5-56 Backup execution window

As mentioned before, on Windows 2000 NTFS formatted volumes, Tivoli Storage Manager transfers just the used blocks of the volume (by default). In this example just 82 MB data were transferred to the Tivoli Storage Manager server (Figure 5-57).

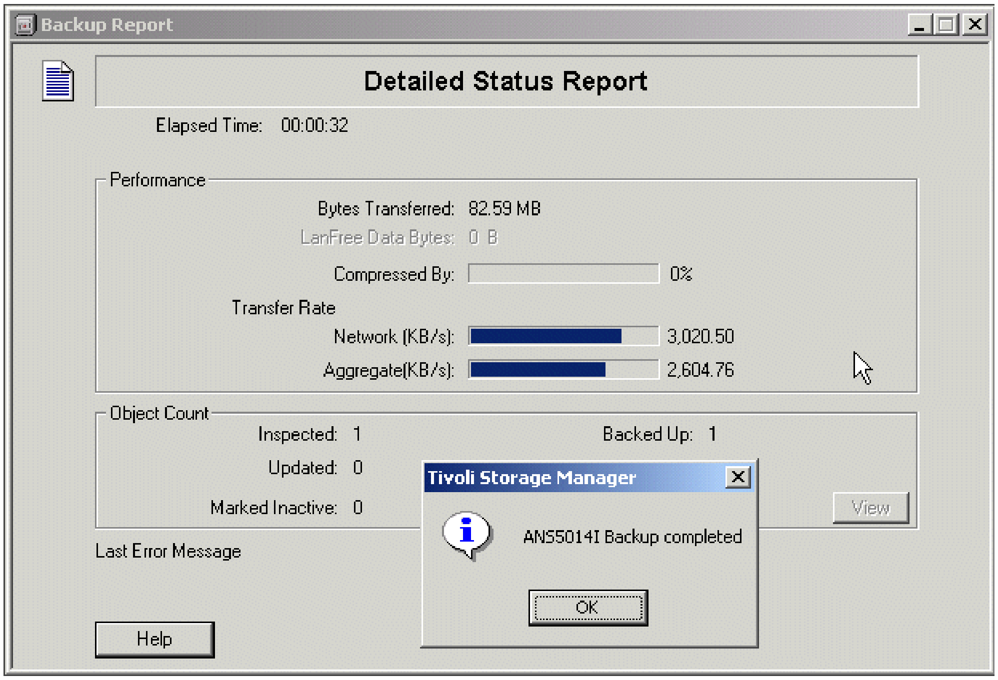

Figure 5-57 Backup completed window

Logical Volume Backup via Command Line Interface (CLI)

To perform an image snapshot backup, with a snapshotcachesize of 15%, from the command line, type:

Example 5-89 The dsmc backup image command

```
dsmc backup image g: -imagetype=snapshot -snapshotcachesize=15
```

After Tivoli Storage Manager has finished the data transfer you will get the following output:

Example 5-90 Output of the dsmc backup image command

```
Tivoli Storage Manager
Command Line Backup/Archive Client Interface - Version 5, Release 1, Level 0.0 g4
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

Node Name: BONNIE
Session established with server JAMAICA: Windows
  Server Version 5, Release 1, Level 2.0
  Server date/time: 03/12/2002 10:12:47  Last access: 03/12/2002 10:08:44

Backup Image Function Invoked.

Volume -->     524,353,024 \\bonnie\g$ [Sent]
Selective Backup processing of '\\bonnie\g$' finished without failure.

Total number of objects inspected:        1
Total number of objects backed up:        1
Total number of objects updated:          0
Total number of objects rebound:          0
Total number of objects deleted:          0
Total number of objects expired:          0
Total number of objects failed:           0
Total number of bytes transferred:    82.59 MB
```

```
Data transfer time:                 69.00 sec
Network data transfer rate:      1,225.71 KB/sec
Aggregate data transfer rate:      574.17 KB/sec
Objects compressed by:                  0%
Elapsed processing time:         00:02:27
```

Logical Volume Restore

When performing an image restore, the following things are to be considered:

- You can restore NTFS to FAT32; FAT to NTFS, RAW to NTFS, NTFS to RAW and so on, but it is recommended that the destination volume be RAW (not formatted) or formatted with the same type of the filesystem as the original one.
- Restoring the image of a volume will restore the volume to the same state that it was in when you performed your last image backup. Be absolutely sure that you need to restore an image, for it will replace your entire current file system or raw volume with the image on the server.
- Ensure that the file system or volume to which you are restoring the image is at least the same size as the image that is being restored.
- When restoring an image to a file system, ensure that the file system exists on your machine before you begin the restore. Otherwise, the restore operation will fail.
- Image restores are always offline. Ensure that the file system is not in use. The client will lock the file system before starting the restore. If the file system is in use when the client attempts to lock the file system, the restore will fail. The client will unlock the file system after the restore completes.
- You cannot restore an image to where the Tivoli Storage Manager client program is installed.
- If you created an image of the system drive, you cannot restore the image to the same location because the client cannot have an exclusive lock of the system drive. Also, because of different system component configurations, the system image may not be consistent across components (such as Active Directory). Some of these components can be configured to use different volumes where parts are installed on the system drive and others to non-system volumes.
- If you have run progressive incremental backups and image backups on your file system, you can perform an incremental restore of the file system. This process updates the original image with individual files backed up after the last image backup. Optionally, if files were deleted after the original backup, the incremental restore can delete those files from the base image. Incremental backups and restores can be performed only on mounted file systems, not on raw logical volumes.

To ensure that you can perform point-in-time restores of your file systems, including deleting original files which no longer exist on the logical volume, use a combination of full image backups and the incremental command as described in these steps:

1. Perform a full incremental backup of the logical volume, for example:

   ```
   dsmc incremental g:
   ```

2. Perform an image backup of the same logical volume, for example:

   ```
   dsmc backup image g:
   ```

3. Periodically, perform incremental backups, for example:

   ```
   dsmc incremental g:
   ```

You must follow these steps in the order shown to ensure that the server records additions and deletions accurately. The following command restores the file system to its exact state as of the last incremental backup:

```
dsmc restore image h: -incr -del
```

If you do not follow the steps exactly, two things can occur:

- After the original image is restored, all files backed up with the incremental command are restored individually.
- If you perform a backup image before performing an incremental, files deleted from the original image are not deleted from the final restored file system.

Following is an example of how an image backup and restore might work.

Table 5-3 Backup sequence

	Client	Server
Saturday	Files 1, 2 and 3 reside here	A full incremental backup is performed.
Sunday	Files 1, 2 and 3 reside here	An image backup is performed. A list of files in the image is created on the server.
Monday	File 4 is created	Daily incremental is performed - including incremental for file 4.
Tuesday	File 2 is deleted; File 3 is changed	Daily incremental is performed - including incremental for file 3; list of files in image backup shows file 2 was deleted.
Wednesday	File 1 is deleted	Daily incremental is performed; list of files in image backup shows file 1 was deleted.

Suppose on Thursday morning you have a crash of your disk subsystem and you have to restore it as soon as possible. Therefore you want to use the image created on Sunday to restore your file system as it appeared just after Wednesday's incremental backup. Tivoli Storage Manager would use the following process:

- The server sends the image backup to the client. This image replaces the current content of the target file space. The file system now contains files 1, 2 and 3.
- Based on the information from the server about subsequent incremental backups, the client removes files 1 and 2.
- The client restores file 4.
- The client restores the latest version of file 3.

The Tivoli Storage Manager client restores the image from Sunday and just the incremental which were made between Sunday and Wednesday. The client also takes care of the deleted files between Sunday and Wednesday.

Logical Volume Restore via GUI

You can use the following procedure to restore an image of your file system or raw logical volume:

Click **Restore** from the main window of the Tivoli Storage Manager Client GUI or Web GUI. The Restore window appears, as shown in Figure 5-58.

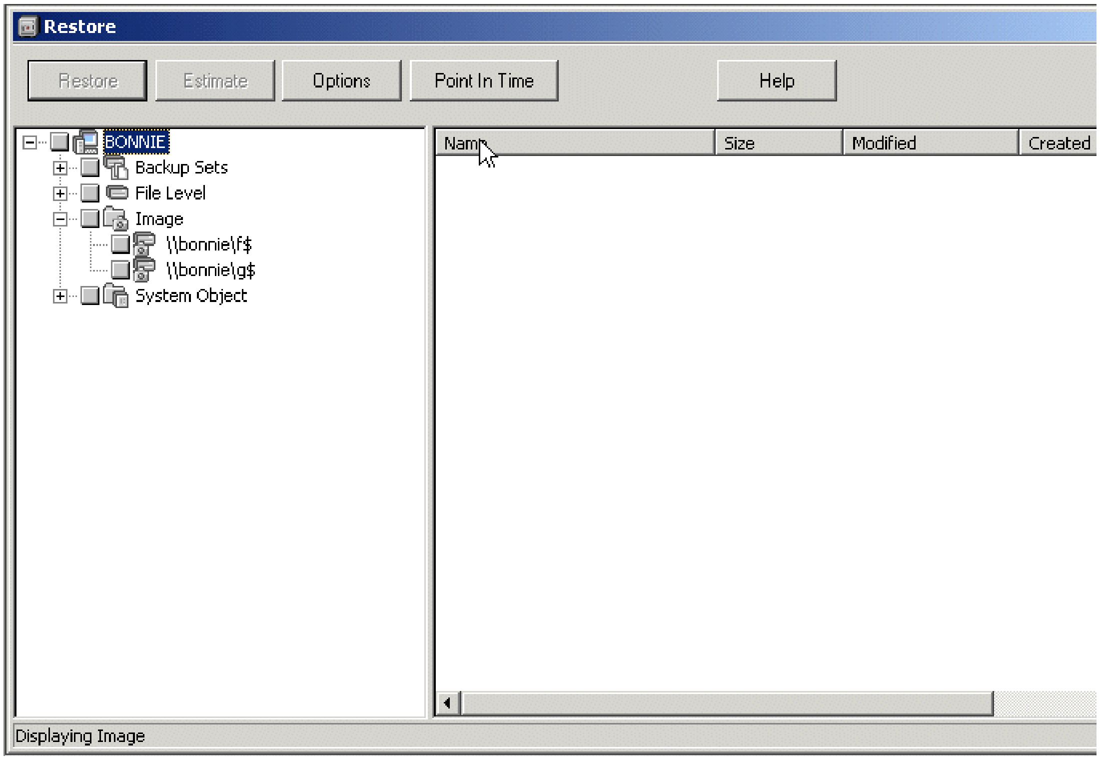

Figure 5-58 Restore window

Locate the object in the tree named **Image** and expand it. Click the selection box next to the image you want to restore.

To modify specific restore options, click the **Options** button. See Figure 5-59.

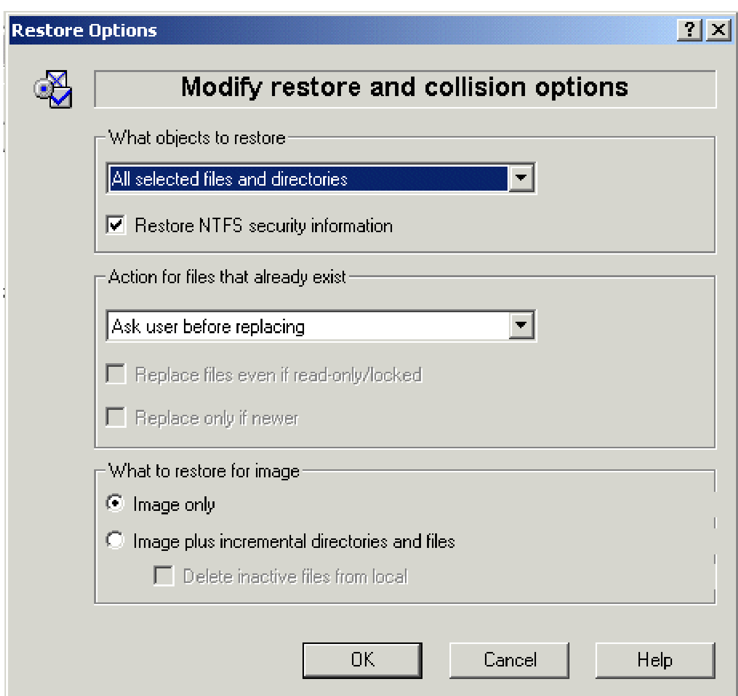

Figure 5-59 Restore options window

The following settings can be made in the Restore Options window. Remember though that most of these do not apply to an Image restore:

- What objects to restore:
  - **All selected files and directories**: Recovers both files and directories.
  - **Files only**: Recovers files only
  - **Directories only**: Recovers Directories only
  - **Restore NTFS security information**: Select the Restore NTFS security information check box if you want to restore NTFS security information like file permissions and compute the NT security CRC for files. This is applicable only to Windows NT files on NTFS drives. Note: Choosing this option might slow down your restore because Tivoli Storage Manager has to retrieve all the NTFS security information.
- Action for files that already exist:
  - **Ask user before replacing**: Prompts you to choose whether to overwrite an existing file on your workstation when it restores or retrieves a file from the server.
  - **Do not Replace**: Does not overwrite existing files on your workstation.
  - **Replace**: Overwrites existing files on your workstation. If you select this option, the Replace Files Even if Read-Only/Locked and Replace Only if Newer check boxes become available.

    If you select the Replace item in the drop-down list, the following check boxes become available:

    - **Replace Files Even if Read-Only/Locked**: Overwrites read-only or locked files on your workstation with the files from the server.
    - **Replace Only if Newer**: Replaces existing files on your workstation with the files from the server only if the files on the server are newer.
- What to restore for image:
  - **Image Only**: Select this option to restore the image from your last image backup only.
  - **Image and Incremental Directories and Files**: Select this option to restore the image from your last image backup and the backup versions from your incremental backups. Use this option if you ran progressive incremental and image backups on a file system, or if you ran incremental by date image backups and full image backups.
  - **Delete Inactive Files from Local**: Select this check box to delete the inactive files that are restored to your local file system. This check box is available only if you selected Image and Incremental Directories and Files. If incremental by date image backup was the only form of incremental backup performed on the file system, deletion of files will not occur.

After setting up the restore options and selecting the image you want to restore, the following window appears (Figure 5-60):

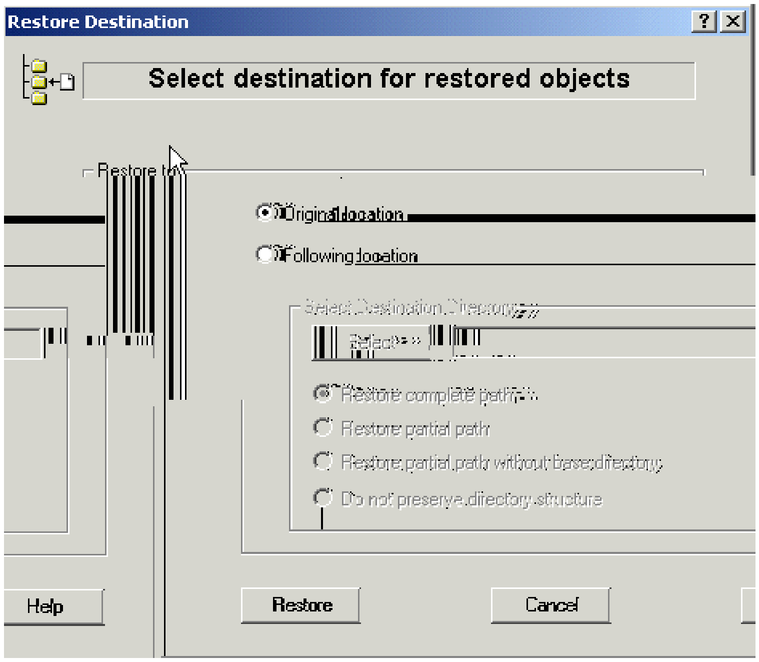

Figure 5-60 Restore destination window

In this window you can select if the image destination should be the original or an other destination. The other options do not have effects on the image restore.

Click the **Restore** button to begin the restore. A Tivoli Storage Manager message box appears, as shown in Figure 5-61.

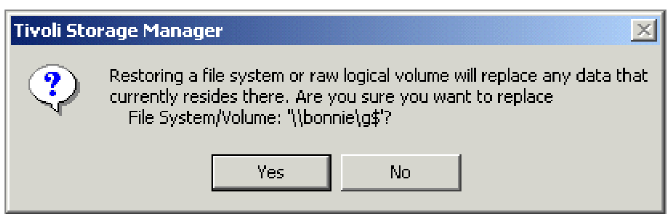

Figure 5-61 Tivoli Storage Manager message box

If you are sure to replace the File System/Volume press **Yes**. Now Tivoli Storage Manager will transfer all the image data to the destination File System/Volume. After transfer, the Restore Report window appears, see Figure 5-62.

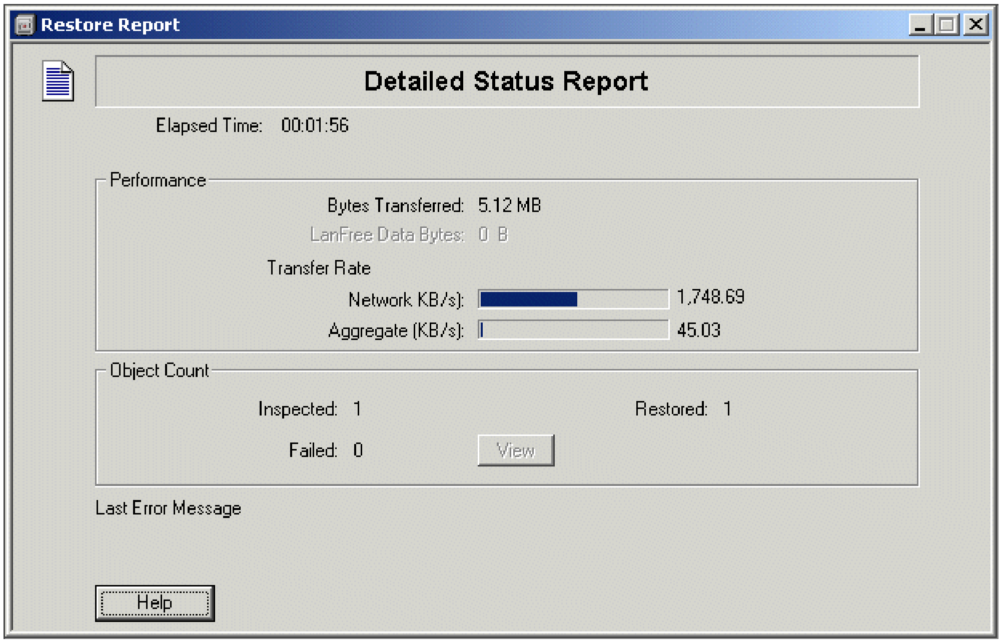

Figure 5-62 Restore report window

5.5.4 SAN-based server-free data movement

The Windows 2000 client supports SAN-based Server-Free Data Movement that offloads data movement processing from the client and server processors and from the LAN during image backup and restore operations.

Data is transferred directly between client disks and SAN-attached storage devices by a third-party copy function initiated by the Tivoli Storage Manager server.

Server-free backup is only available for NTFS-formatted volumes and raw volumes; FAT and FAT32 file systems are not supported. The following volume types are supported:

- Simple volumes
- Spanned volumes
- Mirrored volumes

Enabling the client for Server-Free backups

You enable the client for server free backups by checking the box in the preferences editor under the General tab as shown in Figure 5-63.

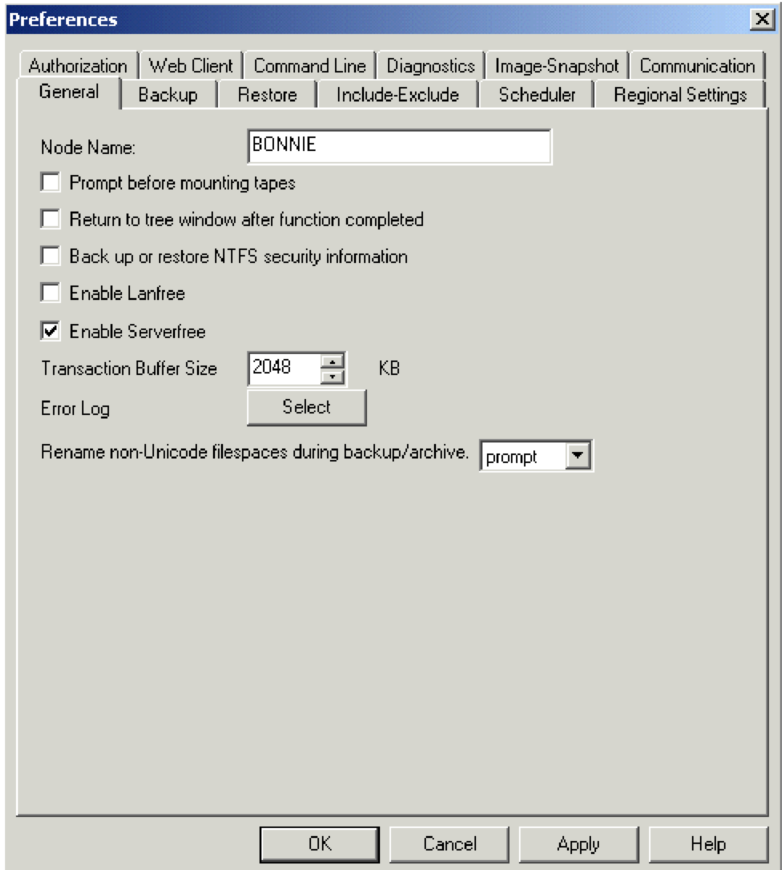

Figure 5-63 Enabling server free backups on the client

You can manually edit the Tivoli Storage Manager client options file 'dsm.opt' if you prefer. Here is an example of what our client 'dsm.opt' file looked like for Server-Free Data Movement. You need a line to enable Server-Free Data Movement and an 'INCLUDE.IMAGE' statement to identify the target Management Class (in this case SERVERFREEMAGMC) and hence the target Storage Pool for any drive you want to backup using Server-Free Data Movement.

Example 5-91 Tivoli Storage Manager client options file for Server-Free Data Movement

```
NODENAME CLYDE
PASSWORDACCESS GENERATE
TCPSERVERADDRESS 9.1.36.156
ENABLESERVERFREE YES
INCLUDE.IMAGE D: SERVERFREEMAGMC
```

Remember, you have to have a SAN Data Mover device enabled and defined to the Tivoli Storage Manager server for Server-Free Data Movement to take place. If you attempt a server free image backup from the client without a Data Mover configured you will see an error like this:

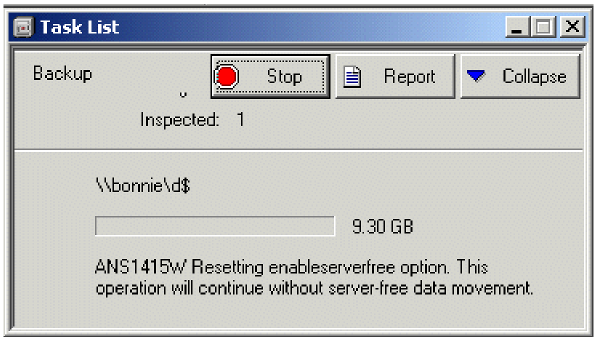

Figure 5-64 Warning if Data Mover not available on server free backup

Server-Free backup and restore

Once both the client and server have been configured and enabled for server-free image backups you can perform server free backups and restores from the Tivoli Storage Manager client GUI or command line interface. To perform a server free backup from the GUI you simply select the SAN attached drive you want to backup and then select 'Image Snapshot Backup' or 'Image Backup' from the drop down and then click on the backup button.A panel like this appears to confirm that the image backup is taking place Server-Free.

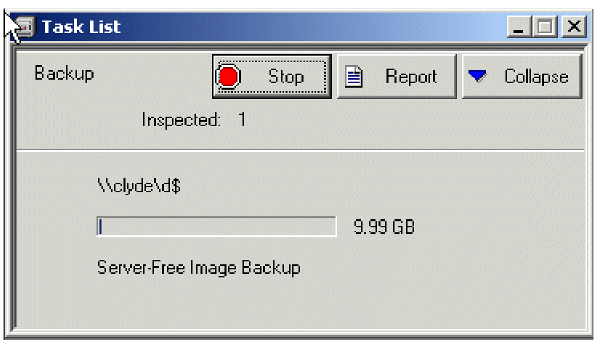

Figure 5-65 Server free image backup in progress

Note that the full drive capacity is displayed but you can choose to backup only the occupied blocks and skip the empty ones by specifying a non zero value for the Image Gap Size option if you are backing up from an NTFS volume. This is the default behavior. The default for the Image Gap Size option is 32 MB. Note also that in this case the blue progress bar which indicates how much data has been transferred will only report on the actual data transferred so will not completely fill the space when the server free backup has completed. For example, in the case here a 9.99GB drive was being backed up of which there was only 971 MB occupied. Therefore the bar only went approximately 10% of the way across. Once the backup has completed, a status report shows how much data was actually transferred and the transfer rates held. Notice that only one object (one image) has been backed up.

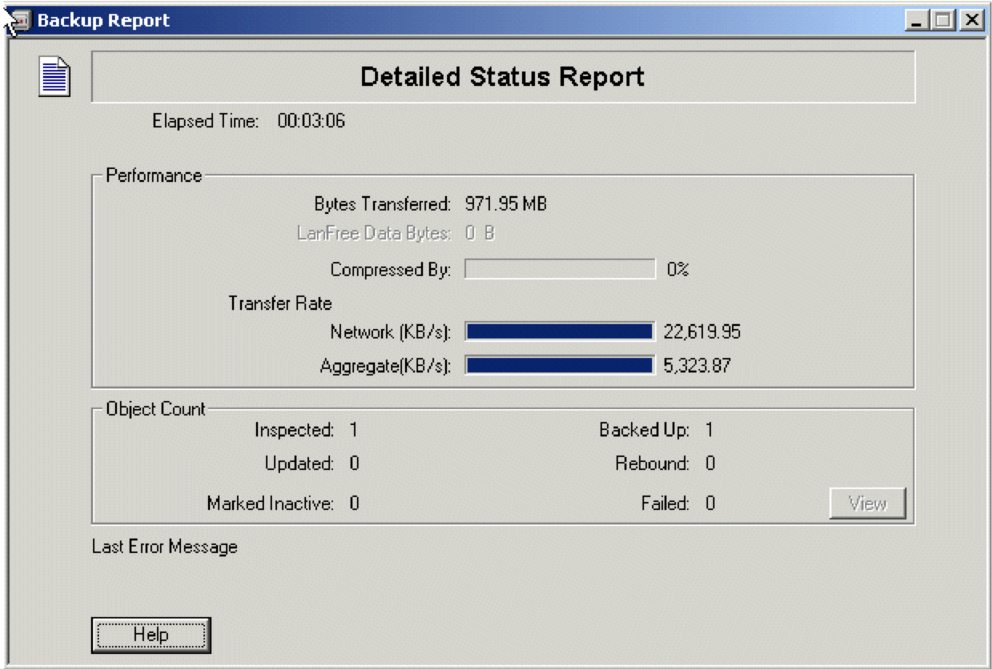

Figure 5-66 Server free backup status report

You can use the Command Line Interface to perform the same image backup. By default, the image backup is a snapshot, but you can request a static backup if you prefer.Here's the kind of display you will see when performing a server free image backup via the CLI. Notice that we indicate on the client that a server background process has been started to control the data movement from the client's SAN attached disk to the server's SAN attached tape device using the SAN Data Gateway datamover functions. Once again we report the size of the entire volume backed up and the number of actual bytes transferred to the tape device.

```
Tivoli Storage Manager
Tivoli Storage Manager
Command Line Backup/Archive Client Interface - Version 5, Release 1, Level 1.0 b3
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

tsm> Backup Image D:
Node Name: CLYDE
Session established with server JAMAICA: Windows
  Server Version 5, Release 1, Level 1.0
  Server date/time: 03/25/2002 11:58:36  Last access: 03/25/2002 11:54:16


Backup Image Function Invoked.

<       78  B> [-   ]

Process initiated on TSM server.

        Start date/time: 03/25/2002 11:59:43
              Operation: Server-Free Image Backup
                   Node: CLYDE
                 Source: \\clyde\d$
    Bytes in operation : 971.82 MB

Volume --> 10,733,957,120 \\clyde\d$ [Sent]
Selective Backup processing of '\\clyde\d$' finished without failure.


Total number of objects inspected:        1
Total number of objects backed up:        1
Total number of objects updated:          0
Total number of objects rebound:          0
Total number of objects deleted:          0
Total number of objects expired:          0
Total number of objects failed:           0
Total number of bytes transferred:   971.95 MB
Data transfer time:                    44.00 sec
Network data transfer rate:        22,619.95 KB/sec
Aggregate data transfer rate:       4,850.40 KB/sec
Objects compressed by:                    0%
Elapsed processing time:           00:03:25
tsm>
```

Figure 5-67 Server free image backup using CLI

This background process that is initiated on the Tivoli Storage Manager server can also be seen in the server's Activity Log of course. A 'Query Actlog' command would show:

```
13:17:34   ANR8337I 3570 volume 080CED mounted in drive MAGDRV1
            (mt1.4.0.3).
13:17:58   ANR0984I Process 6 for BACKUP SERVER-FREE started in the
            BACKGROUND at 13:17:58.
13:17:58   ANR1079I Server-free backup of node CLYDE, file system
            \\clyde\d$, started as process 6 by administrator .
13:20:00   ANR1067I Server Free Backup process 6 completed.
13:20:00   ANR0988I Process 6 for BACKUP SERVER-FREE running in the
            BACKGROUND processed 1,019,156,376 bytes with a
            completion state of SUCCESS at 13:20:00.
```

Figure 5-68 Activity log extract showing server free backup

A 'Query Process' command could also be used to show any server-free processes in progress on the server.

Restoring an image backup server free can also be done from the client GUI or command line interface. With the GUI you simply select the image you wish to restore and click the Restore button as shown in Figure 5-69.

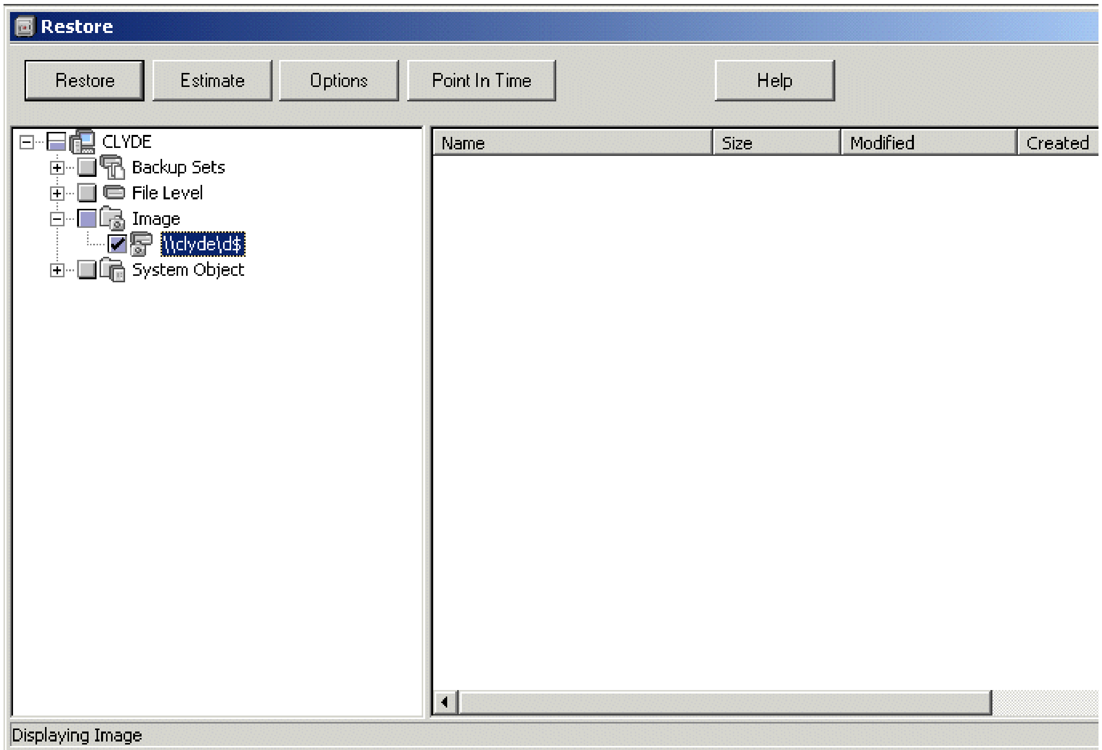

Figure 5-69 Server free restore selection

You will get a warning that the entire volume is about to be overwritten:

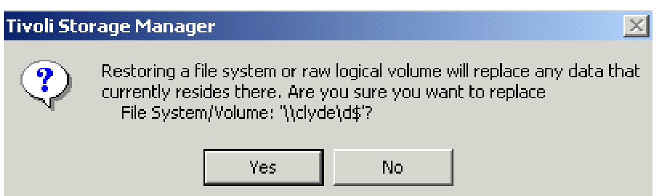

Figure 5-70 Warning on image restore

Once you click ‘Yes’ then the Task List panel is displayed. Notice that this is where you see an indication that the restore is server free.

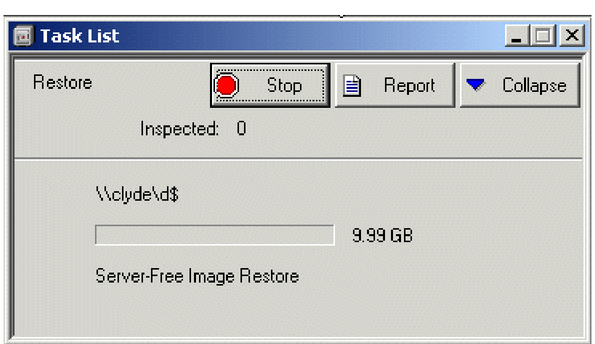

Figure 5-71 Server free image restore

As with the backups, the blue progress bar will not completely fill the space as it is recording only the data transferred (in the occupied blocks) and not the entire volume.

Image restore using the CLI can also be done. The default is to restore the image to its original location. As with the image backup, there is an indication that a background process has been started on the Tivoli Storage Manager server to control data movement from the tape device back to the client's SAN attached disk.

```
Tivoli Storage Manager
Tivoli Storage Manager
Command Line Backup/Archive Client Interface - Version 5, Release 1, Level 1.0 b3
(C) Copyright IBM Corporation 1990, 2002 All Rights Reserved.

tsm> Restore Image D:
Node Name: CLYDE
Session established with server JAMAICA: Windows
  Server Version 5, Release 1, Level 1.0
  Server date/time: 03/27/2002 15:13:49  Last access: 03/27/2002 14:55:10


Restore Image Function Invoked.

***************************** WARNING ********************************
Restoring a file system or raw logical volume will replace any data that
currently resides there. Are you sure you want to replace
    File System/Volume: 'D:'?  (Yes (Y)/No (N)) y
<       78  B> [ - ]

Process initiated on TSM server.

        Start date/time: 03/27/2002 15:14:39
              Operation: Server-Free Image Restore
                   Node: CLYDE
                 Source: \\clyde\d$
            Destination: \\clyde\d$
    Bytes in operation : 971.94 MB

Restoring  10,733,957,120 \\clyde\d$ [Done]

Restore processing finished.

Total number of objects restored:         1
Total number of objects failed:           0
Total number of bytes transferred:   971.82 MB
Data transfer time:                   97.00 sec
Network data transfer rate:        10,259.25 KB/sec
Aggregate data transfer rate:      2,061.09 KB/sec
Elapsed processing time:           00:08:02
tsm> _
```

Figure 5-72 Server free image restore using CLI

As with the backup, this background process can be seen in the Tivoli Storage Manager server Activity Log as shown in Figure 5-73.

```
13:36:29   ANR0984I Process 7 for RESTORE SERVER-FREE started in the
            BACKGROUND at 13:36:29.
13:36:29   ANR1186I Restore of node CLYDE, file system \\clyde\d$,
            started as process 7 by administrator .  A full image for
            this file system will be restored to destination
            Harddisk1.
13:41:17   ANR1067I Server Free Restore process 7 completed.
13:41:17   ANR0988I Process 7 for RESTORE SERVER-FREE running in the
            BACKGROUND processed 1,019,031,552 bytes with a
            completion state of SUCCESS at 13:41:17.
```

Figure 5-73 Activity log extract showing server free restore

Attention: You may notice that the number of bytes processed on the restore is slightly less than the number of bytes processed during the backup. Don't panic! We have not lost any data. The byte count for the backup includes all the bytes moved to the tape drive. This includes the frame headers and the object header. The byte count for the restore only includes the bytes actually moved to the disk so does not contain any headers. The larger the image you backup, the greater is this discrepancy.

5.5.5 Journal based backups on MSCS clusters

Journal Based Backup provides an alternative to traditional Progressive Incremental backup which under certain circumstances may dramatically increase overall backup performance.

The primary difference between Journal Based Backup and Progressive Incremental Backup is the method in which the list of backup candidate objects is derived.

The backup candidate list specifies objects for a particular filesystem which are to be backed up, expired, or updated on the Tivoli Storage Manager server by a Tivoli Storage Manager Backup Client.

Progressive Incremental Backup derives the backup candidate list by building and comparing the list of active previously backed up objects stored on the Tivoli Storage Manager server with the list of objects currently residing on the local filesystem.

The server list is obtained over the network and the local list is obtained by scanning the local filesystem. Objects which exist in the local list but don't exist in the server list are added as backup candidates to the candidate list.

Objects which exist in both lists but differ in some way (attributes, policy, size, etc.) are also added as backup candidates unless only the Tivoli Storage Manager database attributes differ, in which case they are added as attribute update candidates.

Objects which exist in the server list but don't exist in the local list are added as expiration candidates.

Under Journal Based Backup, the Tivoli Storage Manager Backup client obtains the backup candidate list by contacting the Journal Based Backup Daemon.This is a local background process (a service on Windows NT/2000) which manages and maintains a journal database of change activity for each file system being journaled.

Journal Database entries are generated by real time file system change activity and specify objects to backup or expire (attribute update actions are currently not supported).The type of change activity which generate Journal entries is configurable by the user and may consist of any combination of the following:

- Objects created, deleted, or renamed on the file system
- Size changes
- Modification time/date changes
- Access time/date changes
- Attribute changes
- Security (acl) changes

Once a journal entry has been successfully processed, the Tivoli Storage Manager Backup Client notifies the Journal Based Backup Daemon to remove the journal entry from the journal database.

Journal Based backup introduced in Tivoli Storage Manager V4.2 now work on MSCS clusters too. How to setup a backup-archive client in a Microsoft Cluster Server environment please refer to the *Backup-Archive Installation and User's Guide, Appendix G,* GC32-0788.

You must install a Tivoli Storage Manager Journal engine before JBB will work. To do this select the Setup Wizard from the Utilities drop down on the client GUI then check **Help me configure the TSM Journals**, as shown in Figure 5-74.

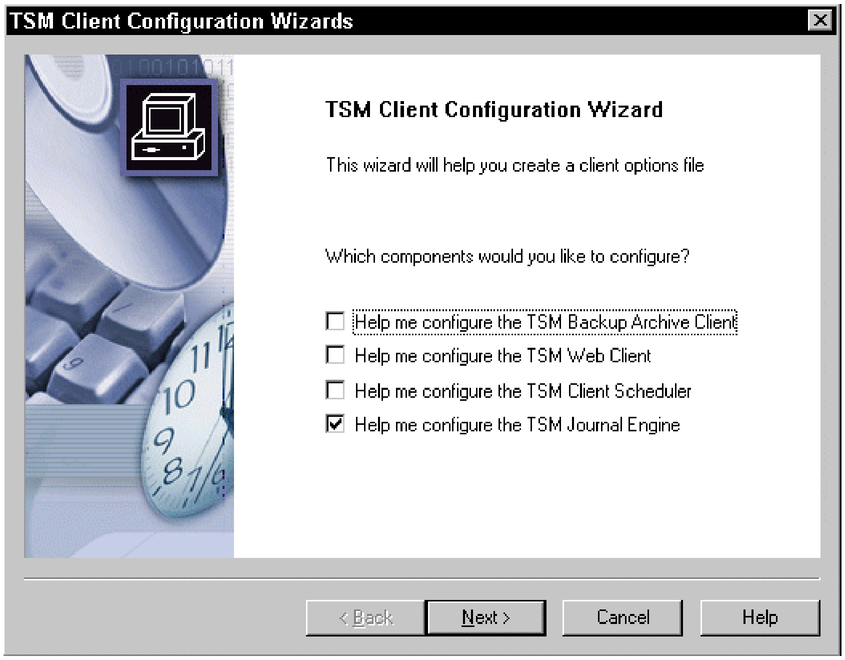

Figure 5-74 Installing the Tivoli Storage Manager journal engine

Cluster exploitation is accomplished by configuring a Journal Engine on each machine in the cluster. It can journal both local and shared resources. Shared resources must be defined with the deferFSMonStart setting enabled so that they may move to different machines in the cluster which currently owns the resource. The location of the journal directory must always be accessible from any machine in the cluster, so the JournalDir setting should define different locations for shared and local resources.

There are some new settings relating to Journal Based Backup on Tivoli Storage Manager V5.1.

The Directory where journal database files are stored and written can now be changed. The default location is the Journal Service install directory. By default this setting applies to all journaled file systems but may be overridden via an override stanza for each journaled file system. If the default value is a fully qualified path (for example c:\tsmjournal) all journal db files will be written to the specified directory. If the default value does not specify a drive letter (for example \tsmjournal) the journal db files for each journal file system will be written to the specified directory on each journal filesystem.

PreserveDBOnExit

A value of 1 indicates that the journaled file system journal database will not be deleted when the journal file system goes offline and will be valid when the journal file system comes back online. This value should be used with caution as any file system change activity which occurs while the journaled file system is offline will not be reflected in the journal database. The PreserveDBOnExit setting may be enabled to allow a shared resource journal to remain valid when a shared resource is moved.

DeferFsMonStart

Defers bringing a journaled file system online until the specified journaled filesystem is valid/available and the specified journal directory can be created/accessed. It will also monitor online journaled filesystems for availability and will bring the filesystem offline if it becomes unavailable.

Resources are checked at the interval specified by the deferRetryInterval setting.

This setting is most commonly used in a cluster environment where shared resources may move to different machines in the cluster.

DeferRetryInterval

The value in seconds journaled filesystems with the deferFsMonStart setting enabled will be checked for availability. The default value is 60 seconds.

logFsErrors

Indicates if errors encountered accessing a journaled filesystem or journal directory should be logged (both to the jbb errorlog and the NT eventlog). Usually used in conjunction with deferFsMonStarts setting to eliminate excessive File System unavailable messages from being written to the logs when bringing a journaled filesystem online is deferred.

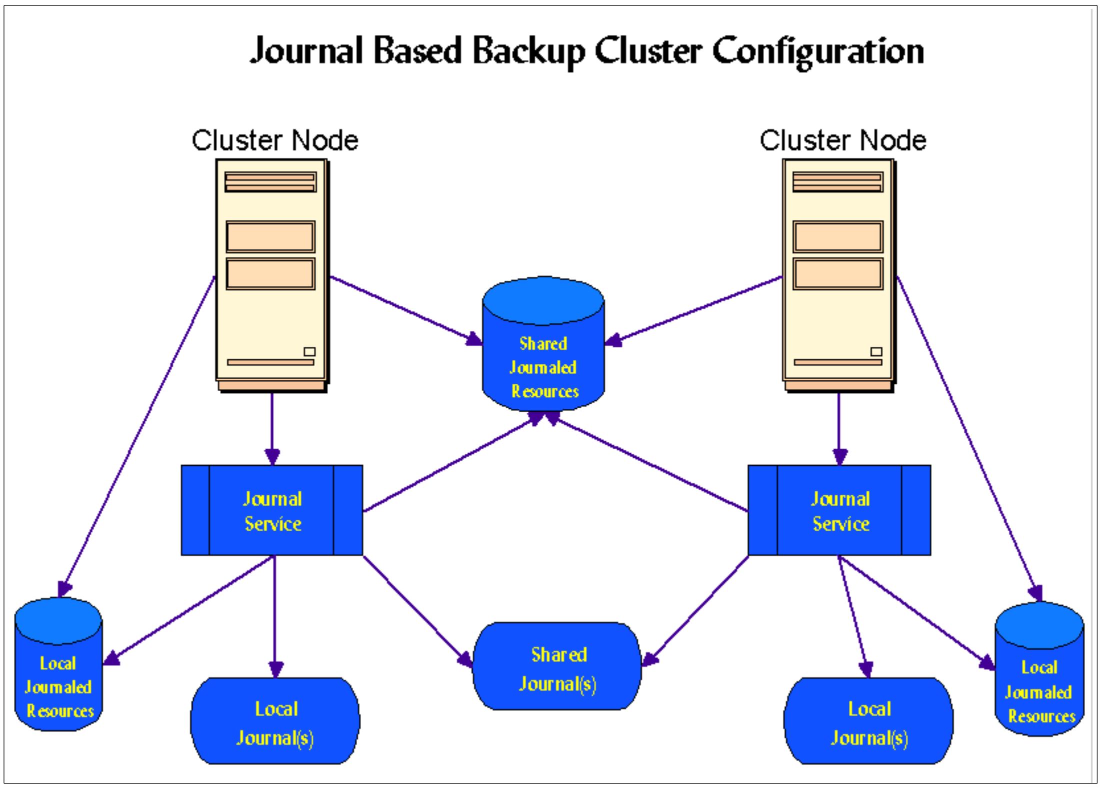

Figure 5-75 Journal based backup diagram

We tested Journal Based Backup in an MSCS cluster configuration in our lab. Two Win/2000 clients (Bonnie and Clyde) shared a D: and E: drive and had a quorum drive (Q:). We installed a Journal Service on both machines, using the wizard, and configured them to monitor both local and shared resources. We then used the Tivoli Storage Manager command line client to do incremental backups of the shared drives and could see that the journal was being used.

Note: You have to do a complete full Incremental backup first before a journaled backup can take place.

5.6 Netware

The IBM Tivoli Storage Manager 5.1 client for Netware introduces some new features and additional support.

5.6.1 Changes in IBM Tivoli Storage Manger Version 5.1 client

Following is a summary of changes since Tivoli Storage Manager Version 5.1:

Netware

Changes in IBM Tivoli Storage Manger version 5.1 client

- Enhanced Web Client Interface
- Support for Long File Names for Backup Sets
- Support for Long File Name Conversions
- Support for files greater than 4 GB on Novell Storage Services (NSS) Volumes
- Multiple session No Query Restore
- Consistent Client Return Codes

Figure 5-76 Changes in IBM Tivoli Storage Manger Version 5.1 client

- **Enhanced Web Client Interface**

  The Web client interface is enhanced to support a JRE 1.3.1 Swing-enabled browser. This type of browser facilitates the use of assistive devices for users with disabilities.

- **Support for Long File Names for Backup Sets**

  Tivoli Storage Manager supports the use of long file names for backup sets on NetWare Version 5.1.0.

- **Support for Long File Name Conversions**

  Tivoli Storage Manager supports long file name conversions without being affected by the memoryefficientbackup option.

- **Support for files greater than 4 GB on Novell Storage Services (NSS) Volumes**

  Tivoli Storage Manager supports files larger than 4 GB for NSS file spaces on NetWare 5.x and NetWare 6.

- **Multiple session No Query Restore**

  The backup-archive clients can now perform multiple restore sessions for no query restore operations, increasing the speed of restores. This is similar to the multiple backup session support. It exploits the mount point available on the server. If data is backed up on multiple tapes, and if the server has multiple mount points available, then the restore starts a session for each tape, up to the number your administrator configures.

- **Consistent Client Return Codes**

  Reliable, consistent, and documented return codes have been added to the command line client and the scheduler. This facilitates automation of client operations via user-written scripts. Administrators can now distinguish between scheduled backups that completed successfully with no skipped files and scheduled backups that completed successfully with one or more skipped files. Also if the preschedulecmd command ends with non-zero return codes, the scheduled event will not run. This ensures that scheduled events will not run if prerequisite commands do not complete successfully.

5.6.2 Enhanced Web client interface

The Web client interface is enhanced to support a JRE 1.3.1 Swing-enabled browser. This type of browser facilitates the use of assistive devices for users with disabilities. The first time you access to Tivoli Storage Manager V5.1 server you will be prompted to download and install a Java (TM) 2 runtime environment plug in.

5.6.3 Support for long file names for backup sets

Tivoli Storage Manager supports the use of long file names for backup sets on NetWare Version 5.1.0. Use the ‘Location’ option to specify where Tivoli Storage Manager searches for the backup set for a query or restore operation. This can be used to locate backup sets on the Tivoli Storage Manager server or on local files.

5.6.4 Support for long file name conversions

Tivoli Storage Manager now supports long file name conversions without being affected by the *memoryefficientbackup* option. If you have file names longer than the DOS name space permits, you can use the long name when you back up the file. If long name space support is removed from a NetWare volume, Tivoli Storage Manager Version 5.1.0 assumes the NTW:LONG file space already exists on the server and ceases any back up or archive operation. If you want to continue backing up using long names, add the long name space support back to the volume in question. You can also use a short name to back up the volume by renaming the file space on the server before attempting a back up.

Macintosh name space is not supported

Tivoli Storage Manager does not support Macintosh name space. If a NetWare volume has only DOS and Mac name spaces, Tivoli Storage Manager uses the DOS name for Macintosh file names. If you use only DOS and Mac names spaces on a NetWare volume, you should add Long name space to that volume. This will ensure that NetWare converts Macintosh file names to the Long name before the file is backed up.

Migration to the long name space

Using the Long name space allows greater flexibility for backing up files whose names are longer than traditional DOS names.

To determine the name spaces that are loaded on a volume, enter **volumes** at the NetWare console. For example: `411:volumes`. Information displays which is similar to the following:

```
Mounted Volumes                    Name Spaces               Flags
    SYS                              DOS, LONG               Cp Sa
    VOL1                             DOS, MAC, LONG          Cp SA
411:
```

Figure 5-77 NetWare mounted volumes display

In this example, server 411 has two volumes. The SYS volume has DOS and long name spaces. VOL1 has DOS, Mac and long name spaces. The Tivoli Storage Manager client inspects these settings when it begins an operation, such as backup. The client then determines the name space used by the backed up volume.

5.6.5 Support for files larger than 4 GB on NSS volumes

Tivoli Storage Manager supports files larger than 4 GB for Novell Storage Services (NSS) file spaces on NetWare 5.x and NetWare 6. The maximum file size for backup, restore, archive, and retrieve operations is now 4,294,963,200 (4GB -4KB). Netware 5.x and Netware 6 clients support a maximum file size greater than 4 GB on NSS volumes.

5.6.6 Multiple session no query restore

The backup-archive clients can now perform multiple restore sessions for no query restore operations, increasing the speed of restores. This is similar to the multiple backup session support. It exploits the mount points available on the server. If data is backed up on multiple tapes, and if the server has multiple mount points available, then the restore starts a session for each tape, up to the number your administrator configures.

When you request a restore, the default is to use a maximum of two sessions, based on how many tapes the requested data is stored on, how many tape drives are available, and the maximum number of mount points allowed for the node. For example, if the data you want to restore is on 5 different tape volumes, the maximum number of mount points is 5 for your node, and *resourceutilization* is set to 3, then 3 sessions will be used for the restore. If you increase the *resourceutilization* setting to 5, then 5 sessions will be used for the restore. There is a 1 to 1 relationship to the number of restore sessions allowed and the *resourceutilization* setting. Multiple restore sessions is only allowed for no query restore operations.

There are potentially undesirable aspects of running multiple sessions. The client could produce multiple accounting records or the server may not start enough concurrent sessions. To avoid this, the servermaxsessions parameter must be reviewed and possibly changed.

5.6.7 Consistent client return codes

Reliable, consistent, and documented return codes have been added to the command line client and the scheduler. This facilitates automation of client operations via user-written scripts. Administrators can now distinguish between scheduled backups that completed successfully with no skipped files and scheduled backups that completed successfully with one or more skipped files. Also if the **preschedulecmd** command ends with non-zero return codes, the scheduled event will not run. This ensures that scheduled events will not run if prerequisite commands do not complete successfully.

If you use Tivoli Storage Manager commands in executables, the return code from a Tivoli Storage Manager executable indicates a successful backup, archive, or restore of specific objects. The explanation of the return codes is explained as follows:

0 The operation completed successfully with no errors or warnings.

4 The operation completed successfully, but some files were not processed.There were no other errors or warnings. This return code is very common.You will see this return code when an operation skips a file because the file is in an exclude list.

8 The operation was mostly successful, but there were warning messages. If your node has a duplicate include/exclude option, a backup will yield message ANS1496W to the error log and return code 8. If a file is skipped because it was changing during a backup attempt, the operation returns code 8.

A macro that is composed of several commands yields a single return code. The return code is the highest return code among all the commands.

For people who don't want a non-zero RC from PRESCHEDULECMD command to prevent a scheduled operation from running you can use the IGNOREPRESCHEDULECMD test flag. This will cause Tivoli Storage Manager to ignore the return code from the PRESCHEDULECMD command.

Those who don't want non-zero RC from POSTSCHEDULECMD command to set the RC to 8 can use the IGNOREPOSTSCHEDULECMD test flag, which will cause Tivoli Storage Manager to ignore the return code from the POSTSCHEDULECMD command.

Those who don't want any of the consistent RC features; but want Tivoli Storage Manager to behave as it had prior to Version 5.1, can use the DISABLECONSISTENTRC test flag, which will cause Tivoli Storage Manager to issue the "old-style" return codes. This test flag implicitly sets the IGNOREPRESCHEDULECMD and IGNOREPOSTSCHEDULECMD test flags, and disables the SHOWALLCONSISTENTRC test flag.

If you want to see return codes for all commands, after each command is issued use the SHOWALLCONSISTENTRC test flag.

The rule of thumb is that the message severity sets the RC.

- ANSnnnnI = RC 0
- ANSnnnnW = RC 8
- ANSnnnnE = RC 12

The exception to the above rule: ANSnnnnW or ANSnnnnE messages that result in skipped files will set the RC to 4. Sometimes it is a matter of personal interpretation whether a particular message constitutes an error.

5.7 Mac OS X

The Apple Mac OS X is a packaging of Apple's Darwin operating system together with different level of system and user interfaces, development libraries and toolkits.

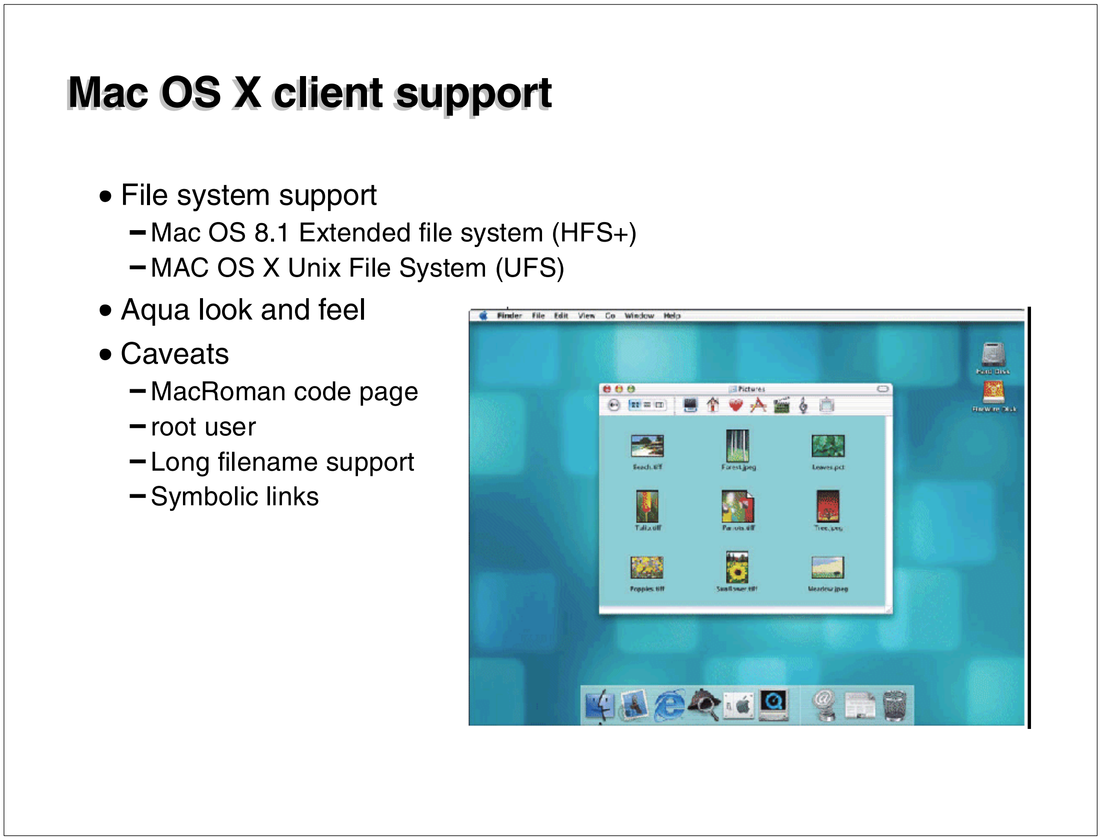

Figure 5-78 Mac OS X client support

The basic parts of Mac OS X are shown in Figure 5-79.

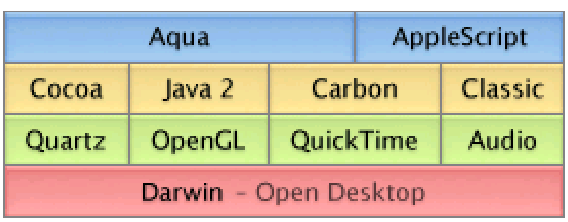

Figure 5-79 Mac OS X's fundamental layers

Mac OS X provides protected memory and preemptive multitasking like other UNIX based operating systems. However the Graphical User Interface (GUI) is quite different from X-Windows and Microsoft Windows based windowing environments - as it always has been.

Figure 5-80 Apple marketing focus of Mac OS X

5.7.1 Darwin

Darwin is the core of Mac OS X and Apple has released the source code to virtually all of the components of Darwin to the developer community.Darwin consists of five major components:

- Mach micro-kernel
- I/O device driver kit
- Virtual File System
- Networking Kernel Extensions
- Berkeley Software Distribution UNIX

Mach micro-kernel

The Mach 3.0 micro-kernel from Carnegie Mellon University, manages processor resources such as CPU usage and memory, handles scheduling, provides memory protection and provides a messaging-centered infrastructure to the rest of the operating system layers. Mach provides Mac OS X with protected memory, preemptive multitasking, virtual memory, and real-time support.

I/O device driver kit

Darwin provides an object-oriented framework, I/O Kit, for the development of device drivers. The I/O Kit provides much of the infrastructure that device drivers require and will allow true plug-and-play, dynamic device management (“hot plugging”) and power management of devices.

Virtual File System

The file system component of Darwin is based on an enhanced Virtual File System (VFS) design, which provides the ability to add in new file systems and enhance those already supported, including HFS, HFS+, UFS, and ISO 9660.

Networking Kernel Extensions

Mac OS X provides an extensible networking system by implementing Network Kernel Extensions (NKEs). Basic UNIX communication protocols and services using the Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) Internet protocols (IP), or more commonly referred to as TCP/IP, supported by Mac OS X are:

- Point to Point Protocol (PPP)
- Printer Access Protocol (PAP)
- Hypertext Transport Protocol (HTTP)
- File Transfer Protocol (FTP)
- Domain Name Services (DNS)
- Service Location Protocol (SLP)

- Dynamic Host Configuration Protocol (DHCP)
- Bootstrap Protocol (BOOTP)
- Lightweight Directory Access Protocol (LDAP)
- Network Time Protocol (NTP)

Berkeley Software Distribution UNIX

Darwin wraps a customized version of Berkeley Software Distribution (BSD) 4.4 around the Mach kernel. Darwin's implementation of BSD provides the process model, basic security policies, and threading support for Mac OS X. It also includes POSIX APIs and provides tools for accessing the file system and networking interfaces.

The Aqua user interface

The most visible part of Mac OS X is its user interface, *Aqua*, which is shown in Figure 5-81.

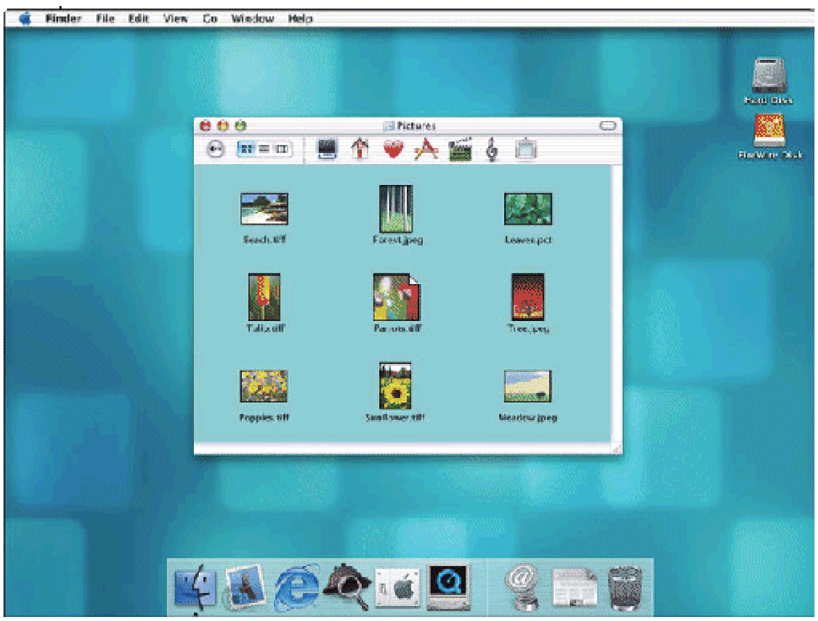

Figure 5-81 MAC OS X desktop

5.7.2 Current limitations and caveats

The Mac OS X TSM 5.1 client supports the Mac OS 8.1 Extended file system (HFS+) and the MAC OS X UNIX File System (UFS) with case sensitivity. The Tivoli Storage Manager 5.1 client has to be executed as the system administrator (root). This will cause the Tivoli Storage Manager 5.1 client to create restored directories with the *root* user as owner.

Other current limitations are that you need to make sure that all files and directories that are to be backed up, only have characters from the *MacRoman* code page, otherwise they will be excluded by the Tivoli Storage Manager 5.1 client. Other limitations are the current lack of long filename support and that symbolic links are not restored properly regarding the ownership and permissions.

Note: When editing the configuration files manually, use the **`SimpleText`** tool or save the file as **`plaintext`** when using the **`TextEdit`** tool.

Should the Tivoli Storage Manager 5.1 client not start due to a configuration error, and the error can not be corrected by the preference editor, either delete the current configuration and create it again, or edit the configuration file with a text editor such as the **TextEdit** tool.

To monitor what the Tivoli Storage Manager 5.1 client scheduler operations, create a file named “TSM Daemon.txt” in the same folder as the daemon is started from. To monitor the authorization tools create a file named “Auth Trace.txt” in the `/tmp` directory.

For more information on the details of the Mac OS X TSM client please refer to the *Tivoli Storage Manager for Macintosh Backup-Archive Installation and User's Guide Version 5 Release 1*, GC32-0787.

The look and feel of the Tivoli Storage Manager Client GUI adheres to the MAC OS X standard.

A

Tivoli Storage Manager SQL

This appendix describes the Tivoli Storage Manager 5.1 Structured Query Language (SQL) interface and how it can be used to monitor the Tivoli Storage Manager server. We will also show some examples using the Windows ODBC client from Lotus Approach and Microsoft Excel.

Structured Query Language

Structured Query Language (SQL) is a relational database language that was first adopted as an industry standard in 1986. Since then, it has been formally adopted as an International Standard by the International Organization for Standardization (ISO) with two major revisions named SQL2 from 1992 and SQL3 from 1999. The SQL standard defines three major classes of entities: objects, data types, and language elements. The **dsmadmc SELECT** command, and the open database connectivity (ODBC) driver, use a subset of the language element part for the **SELECT** query of the SQL2 and SQL3 standards.

For additional information on the Tivoli Storage Manager implementation of SQL please refer to “Using SQL to Query the Tivoli Storage Manager Database” in *Tivoli Storage Manager for* <your platform> *Administrator’s Guide Version 5 Release 1*.

For more information on SQL standards refer to the International Organization for Standardization (ISO), http://www.iso.ch, and American National Standards Institute (ANSI), http://www.ansi.org.

SQL basics

The basic structure of the relational database model is a table, consisting of rows and columns. The data or information are conceptually stored in these tables in rows. The rows can be viewed as records of data. The fields in the rows/records are called columns.

Tables in Tivoli Storage Manager can be accessed by selecting rows that satisfy a given search condition for output. Tables can be manipulated to produce results by operations such as joins on matching columns. However the Tivoli Storage Manager implementation of the **select** SQL query command does not support the following operations:

- UNION
- INTERSECT
- EXCEPT
- Correlated subqueries (returning multiple values)

The select SQL query syntax

The basic syntax description for the **select** SQL query command is:

```
SELECT [ALL | DISTINCT] column1[,column2]
FROM table1[,table2]
[WHERE "conditions"]
[GROUP BY "column-list"]
[HAVING "conditions]
[ORDER BY "column-list" [ASC | DESC] ]
```

Note: A semicolon is *not* allowed to end the **select** SQL query command in Tivoli Storage Manager.

A more thorough syntax description for the **select** SQL query command as it is used by Tivoli Storage Manager is shown in Figure A-1 below.

```
>>-SELECT------------------------------------------------------>

          .-,---------------------------------------------.
          V                                               |
>---------+-value_expression---+---+------------------+--+------>
          '-aggregate_function-'   '-AS--column_name--'

               .-,----------------------.
               V                        |
>----FROM-----table_name--subquery---+---+--------------------+-->
                                         '-WHERE--predicate--'

>-----+-----------------------------------------+-------------->
      '-GROUP BY--+-----------+--.--column_name--'
                  '-table_name-'

>-----+-------------------+------------------------------------>
      '-HAVING--predicate--'

>-----+-----------------------------------------------------+------><
      |                .-,----------------------.            |
      |                V                        |            |
      '-ORDER BY--+----output_column--+------+--+----+--'
                  |                   +-ASC--+       |
                  |                   '-DESC-'       |
                  |  .-,-------------------------.   |
                  |  V                           | |
                  '----positive_integer--+------+--+-'
                                         +-ASC--+
                                         '-DESC-'
```

Figure A-1 Select SQL query syntax for Tivoli Storage Manager

SELECT Specifies the columns to be returned. This is a required parameter. Valid values are:

- **value_expression** Specifies the set of source columns that are to be returned. In the simplest form, this is a list of one or more columns in the table that is specified in the FROM clause. In a more advanced form, it can include DISTINCT, CAST, or CASE expressions or subqueries. DISTINCT will discard the duplicate records for the columns specified after the SELECT statement
- **aggregate_function** Specifies a function that extracts a single value from groups of column names. For example, AVG, COUNT, MAX or SUM.

AS column_name Specifies the text to be used for the column title. This is a required parameter. The default is to display the column name. For value expressions or aggregate functions this is Unnamed[n].

FROM Specifies where in the database the data is to be found. Valid values are:

- **table_name** Specifies one or more source tables from which to extract the query rows and columns. You can find these names by querying `SYSCAT.TABLES`. If you specify two or more tables, you are requesting a JOIN of the tables. You can specify the columns for the JOIN criteria in the WHERE parameter. Or the tables can be joined by matching every row from one table to every row from another table.

subquery Specifies another SELECT command from which the selected columns are to be taken.

WHERE predicate Specifies that only certain rows are displayed based on criteria in the predicate. This is an optional parameter. You can use the AND, OR, and NOT operators to string predicates together.

GROUP BY column_name Specifies groups of rows to be formed if aggregate functions (for example, AVG, COUNT, MAX, SUM) are specified. This is an optional parameter.

HAVING predicate Specifies a condition to be used to filter the extracted values before displaying them. This is an optional parameter.

ORDER BY Specifies how output is to be sorted for display. You can specify column names or numeric positions and whether you want the sort in ascending or descending order. This is an optional parameter. Valid values are:

output_column Specifies order by column. The columns are sorted according to the order in which they are specified.

positive_integer Specifies order by the numeric position of the columns. The columns are sorted according to the order in which they are specified.

ASC Specifies that the columns are sorted in ascending order.

DESC Specifies that the columns are sorted in descending order.

The following comparison operators can be used in a `WHERE` clause:

=	Equal
>	Greater than
<	Less than
>=	Greater than or equal to
<=	Less than or equal to
<>	Not equal to (do not use !=)
LIKE	The LIKE operator allows selection on rows that matches the pattern specified after the operator.

The value that follows a comparison operator in the where clause must be enclosed with single quotes (**'**) if it is non-numeric. With the `LIKE` operator, the percent sign (`%`) can be used as a wild card to match any possible character that might appear before or after the characters specified. For example; '`A%`' matches any character string starting with a capital A. The underscore (_) can be used to match exactly one character.

The select SQL query functions

The Tivoli Storage Manager server `select` SQL query syntax allows the usage of additional functions and aggregate functions which is shown below:

+	Returns the sum of two numbers.
-	Returns the difference between two numbers.
*****	Returns the product of two numbers.
/	Returns the result of dividing one number by another.
%	Returns the remainder of the division of two numbers.
ABS(n)	Returns the absolute value of n.
ABSVAL(n)	Returns the absolute value of n. Same as ABS(n).
BIT_LENGTH(n)	Returns the number of bits used to store n.
CHAR(expression [, n])	Returns a string. expression can be a date expression, string, or number. If specified, n refers to the length of the returned string.
CHAR_LENGTH(string)	Returns the length of string.Same as CHARACTER_LENGTH(string).
CHARACTER_LENGTH(string)	Returns the length of string.Same as CHAR_LENGTH(string).
COALESCE(value1, value2, …)	Returns the first non-NULL value in the list of values.
CONCAT(string1, string2)	Returns the result of appending string2 to the end of string1.
DATE(timestamp)	Returns the date portion of a timestamp value.
DAY(timestamp)	Returns the day of the month of a timestamp value.
DAYNAME(timestamp)	Returns the name of the day from a timestamp value.
DAYOFWEEK(timestamp)	Returns the a number representing the day of the week from a timestamp value. The value returned is 1 for Sunday, 2 for Monday, …, 7 for Saturday.
DAYOFYEAR(timestamp)	Returns the day of the year from a timestamp value.
DAYS(timestamp)	Returns the number of days since January 1, Year 1.
DAYSINMONTH(timestamp)	Returns the number of days in the month from a timestamp value.
DAYSINYEAR(timestamp)	Returns the number of days in the year from a timestamp value.
DEC(string, precision, scale)	Converts a string to a decimal number. The string must conform to the format of a valid decimal number, i.e. '234.65'. Same as DECIMAL(string, precision, scale).
DECIMAL(string, precision, scale)	Converts a string to a decimal number. The string must conform to the format of a valid decimal number, i.e. '234.65'. Same as DEC(string, precision, scale).
DOUBLE(expression)	Returns the result of converting expression to a double-precision number. expression can either be a number or a string. Same as DOUBLE_PRECISION(expression).
DOUBLE_PRECISION(expression)	Returns the result of converting expression to a double-precision number. expression can either be

	a number or a string. Same as DOUBLE(expression).
EXP(n)	Returns the exponential value of n.
FLOAT(expression)	Returns the result of converting expression to a double-precision number. expression can either be a number of a string. Same as DOUBLE (expression).
HEX(expression)	Returns a hexadecimal string that shows how the value is stored in memory.
HOUR(timestamp)	Returns the hour of the time from timestamp.
INT(n | string)	Converts a decimal number n or string to an integer. Same as INTEGER(n).
INTEGER(n | string)	Converts a decimal number n or string to an integer. Same as INT(n).
LCASE(string)	Converts string to all lowercase. Same as LOWER(string).
LEFT(string, n)	Returns the left-most n characters from string.
LENGTH(string)	Returns the length of string.
LOG(n)	Returns the natural logarithm of n in the form of a double.
LOG10(n)	Returns the base 10 logarithm of n in the form of a double.
LOWER(string)	Converts string to all lowercase. Same as LCASE(string).
LTRIM(string)	Removes the leading spaces from string.
MINUTE(timestamp)	Returns the minute of the time from timestamp.
MONTH(timestamp)	Returns the numeric month from timestamp,i.e. January = 1, February = 2, …, December = 12.
MONTHNAME(timestamp)	Returns the name of the month from timestamp.
NULLIF(expression1, expression2)	Returns NULL if both expressions are equal; otherwise it returns expression1.
OCTET_LENGTH(string)	Returns the length of the string in octets (one octet = eight bits).
POSSTR(string1, string2)	Returns the position in string1 where string2 first appears. If string2 does not appear in string1 at all, then the result is 0.
POWER(x, y)	Returns the result of raising x to the yth power. The result is represented as a double value.
QUARTER(timestamp)	Returns the quarter of the year from timestamp.
REAL(n)	Returns the real value of n.
RIGHT(string, n)	Returns the right-most n characters from string.
RTRIM	Removes the trailing spaces from string.
SECOND(timestamp)	Returns the second of the time from timestamp.
SIGN(n)	Returns –1 if n is less than 0, 0 if n is equal to 0, or 1 if n is greater than 0.
SMALLINT(n | string)	Converts a decimal number n or string to a small integer.
SQRT(n)	Returns the square root of n.
SUBSTR(string, p, n)	Returns the part of string from position p for n characters.
TIME(timestamp)	Returns the time from timestamp.
TIMESTAMP(timestamp)	Returns a timestamp value computed from timestamp TIMESTAMP(date, time) or a combination of date and time.

TRIM([option] [optchar] [FROM] string)	Removes leading and/or trailing characters from string. option can be LEADING, TRAILING, or BOTH. optchar can be a single character. The default option is BOTH and the default character is a space. Thus the default is to remove leading and trailing spaces from string.Ifoption and/or optchar are used, then FROM must also be coded.
UCASE(string)	Converts string to all uppercase. Same as UPPER(string).
UPPER(string)	Converts string to all uppercase. Same as UCASE(string).
VARCHAR(timestamp) **VARCHAR(string [, n])**	Converts string to a string, optionally returning a string of length n, or converts timestamp to a string and returns the result.
WEEK(timestamp)	Returns the week of the year from timestamp.
YEAR(timestamp)	Returns the year from timestamp.

System catalog tables

Tivoli Storage Manager provides three system catalog tables to help you find what information is available in the database:

SYSCAT.TABLES	Contains information about all tables that can be queried with the **select** SQL query command.
SYSCAT.COLUMNS	Contains information about the columns in each table.
SYSCAT.ENUMTYPES	Contains information about the valid values for each enumerated type and the order of the values for each type.

Restrictions

Since complicated or lengthy queries can affect Tivoli Storage Manager server performance the usage of the **select** SQL query command have a few restrictions that should be considered beforehand:

- You cannot issue the **select** SQL query command from a server console.
- If an **select** SQL query command requires significant server time or resources, you will be asked to confirm your request.
- Any **select** SQL query command use temporary table storage in the database and require at least one free 4 MB partition, without this partition, temporary table storage space will become exhausted. The more processing that a query requires, the more free space is required. To determine how much temporary table storage space is available in the database, issue the **dsmadmc query db** command as is shown in Example A-1 below.

Example: A-1 dsmamdc query db

```
ANS8000I Server command: 'query db'

Available Assigned   Maximum   Maximum     Page     Total      Used   Pct  Max.
    Space Capacity Extension Reduction     Size    Usable     Pages  Util   Pct
     (MB)     (MB)      (MB)      (MB)  (bytes)     Pages                  Util
--------- -------- --------- --------- ------- --------- --------- ----- -----
      204        4       200         0   4,096     1,024       338  33.0  33.0
```

Check the value in the `Maximum Reduction` field. If this field shows a value of at least 4 MB, you can perform **select** SQL queries, however if the `Maximum Reduction` value is below 4 MB, you might not be able to perform **select** SQL queries. Simple queries like full table scans will usually be possible anyway. In the Example A-1 above, we have `200 MB` that are prepared but not assigned for database usage. The following **dsmadmc** command will take care of that:

```
extend db 200
```

Now when we run the **dsmadmc query db** command again, the output will look similar to the sample in Example A-2 below.

Example: A-2 dsmamdc query db

```
ANS8000I Server command: 'q db'

Available  Assigned    Maximum    Maximum     Page       Total        Used    Pct    Max.
    Space  Capacity  Extension  Reduction     Size      Usable       Pages   Util     Pct
     (MB)      (MB)       (MB)       (MB)  (bytes)       Pages                       Util
---------  --------  ---------  ---------  -------   ---------   ---------  -----   -----
      200       200          0        196    4,096      51,200         514    1.0     1.2
```

The value in the `Maximum Reduction` field is now `196 MB`.

How to use the command line interface

The command line interface is accessed either by using the Web server interface to the Tivoli Storage Manager server or the **dsmadmc** command.

First we will show how the two most common error messages when using the select SQL query interface looks like. The first is the addition of the semicolon to the command line with **dsmadmc** and the second is not using single quotes around the comparison value.

In all the following examples we use the following command to run the select SQL query commands:

```
dsmadmc -id=admin -password=admin
```

Adding a semicolon at the end of the select SQL query statement will produce the error shown in the sample in Example A-3 below.

Example: A-3 Using semicolon to end the select SQL query statement

```
tsm: TSM>select tabname,remarks from syscat.tables where tabname like 'O%';
ANR2910E Invalid SQL token - ';'.

                                                                  |
         .............................................................V
         ct tabname,remarks from syscat.tables where tabname like 'O%' ;

ANS8001I Return code 3.
```

By not using single quotes around the comparison value in the were statement, will produce the error shown in the sample in Example A-4 below.

Example: A-4 Not using single quotes around the comparison value

```
tsm: TSM>select tabname,remarks from syscat.tables where tabname like "O%"
ANR2940E The reference 'O%' is an unknown SQL column name.

                                                             |
```

```
         ............................................................V...
         lect tabname,remarks from syscat.tables where tabname like "O%"

ANS8001I Return code 3.
```

The sample in Example A-5 below, show the result of the same select SQL query statement with the correct syntax.

Example: A-5 select SQL query statement with the correct syntax

```
tsm: TSM>select tabname,remarks from syscat.tables where tabname like 'O%'

TABNAME                 REMARKS
------------------      ------------------
OCCUPANCY               Client storage
                         occupancy
OPTIONS                 Server Options
```

When the select SQL query returns the following message, it means that the query criteria did not find any match. Either it is because there is not match or because the query is spelled wrong or does not perform what was intended. The Example A-6 below, show the message when no match is found.

Example: A-6 No match found for the query

```
ANR2034E SELECT: No match found using this criteria.
ANS8001I Return code 11.
```

If you perform a complex query, try and reduce the complexity in steps until you find a match, to verify that you indeed have constructed the select SQL query as intended. Another useful method is to reverse the selection filtering which is shown in the following two step example shown in Example A-7 below.

Example: A-7 Reversing the selection criteria

```
tsm: TSM>select node_name from nodes where invalid_pw_count <> 0
ANR2034E SELECT: No match found using this criteria.
ANS8001I Return code 11.

tsm: TSM>select node_name from nodes where invalid_pw_count = 0

NODE_NAME
------------------
CRETE
ONE-ON-ONE
...<lines omitted>...
```

To create reports with different column headings than the default, which is the column name, suffix the column name in the select SQL query with as and the new column header string enclosed in double quotes (“). The next sample in Example A-8 below, queries the admins table.

Example: A-8 Renaming the displayed column names

```
tsm: TSM>select admin_name as "dude",system_priv as "can do?" from admins order by
system_priv

dude                    can do?
------------------      ------------------
```

```
CRETE                  No
ONE-ON-ONE             No
TSM                    No
ADMIN                  Yes
RODEN                  Yes
SERVER_CONSOLE         Yes
```

SQL displaymode

To query some display properties for the SQL environment, use the **dsmadmc query sqlsession** command as shown in the next sample in Example A-9 below.

Example: A-9 query sqlsession command

```
tsm: TSM>q sqlses

 Column      Date-Time     Arithmetic     Cursors
Display       Format          Mode        Allowed?
 Format
--------     ---------     ----------     --------
 Narrow         ISO         Truncate         Yes
```

The following dsmadmc set commands can be used to alter the display output from the select SQL query command:

SET SQLDATETIMEFORMAT Controls the formatting of date and time in the display of SQL queries.

SET SQLDISPLAYMODE Controls the column width in the display of SQL queries.

SET SQLMATHMODE Controls how decimal numbers are displayed in SQL queries.

The `Column Display Format` field can either be `Narrow` or `Wide` from Narrow:

Narrow Specifies that the column display width is set to 18. Any wider string is forced onto multiple lines at the client. NARROW is the default.

Wide Specifies that the column display width is set to 250

To change from `Narrow` to `Wide` use the following **dsmadmc** command:

```
set sqldisplaymode wide
```

Example SQL queries

The first example shows how to query the `SYSCAT.TABLES` to find out information about the different tables maintained by Tivoli Storage Manager, in Example A-10.

Example: A-10 Query SYSCAT.TABLES

```
tsm: TSM>select * from SYSCAT.TABLES

TABSCHEMA     TABNAME                      CREATE_TIME     COLCOUNT     INDEX_COLCOUNT     UNIQUE_INDEX     REMARKS
---------     ------------------     ------------------     --------     --------------     ------------     ------------------
ADSM          ACTLOG                                             11                  1            FALSE     Server activity
                                                                                                             log
ADSM          ADMINS                                             17                  1             TRUE     Server administra-
                                                                                                             tors
...<lines omitted>...
```

The next sample, in Example A-11 below, only differs in the report display where we only require the results from the `tabname` and `remarks` columns from the `SYSCAT.TABLES` table.

Example: A-11 Query SYSCAT.TABLES with selective display

```
tsm: TSM>select tabname,remarks from SYSCAT.TABLES
```

```
TABNAME               REMARKS
------------------    ------------------
ACTLOG                Server activity
                       log
ADMINS                Server administra-
                       tors
...<lines omitted>...
```

The following sample in Example A-12 below, use the **dsmadmc select** command to extract all columns from `SYSCAT.COLUMNS` for all tables with additional processing by using SQL functions.

Example: A-12 SQL query to display columns for all tables from syscat.columns with headers

```
SELECT -
CHAR(CONCAT(CONCAT(tab.tabname,'.'),col.colname),35) AS "TABLE_COLUMN",-
CHAR(COALESCE(NULLIF(SUBSTR(col.typename,1,POSSTR(col.typename,'(')-1),''),col.typename),10) AS "TYPE",-
CHAR(col.length,5) AS "#CHR",-
col.remarks AS "DESCRIPTION" -
FROM syscat.columns AS col, syscat.tables AS tab -
WHERE col.tabname = tab.tabname -
ORDER BY table_column
```

Note: Observer the usage of the line continuation character (**-**). The hyphen indicate to the SQL parser that the select statement continues on the next line (after a CR/LF or just LF depending on operating system).

Save the query in a file, in this example we use the name `desc.sql`, and run it with the **dsmadmc macro** command:

```
macro desc.sql
```

The output from the above macro will look similar to the output shown in Example A-13 below.

Example: A-13 Example output from the desc.sql macro

```
TABLE_COLUMN          TYPE          #CHR      DESCRIPTION
------------------    ----------    -----     ------------------
ACTLOG.DATE_TIME      TIMESTAMP     0         Date/Time
ACTLOG.DOMAINNAME     VARCHAR       30        Policy Domain Name
ACTLOG.MESSAGE        VARCHAR       250       Message
ACTLOG.MSGNO          INTEGER       0         Message number
ACTLOG.NODENAME       VARCHAR       64        Node Name
ACTLOG.ORIGINATOR     VARCHAR       20        Originator
ACTLOG.OWNERNAME      VARCHAR       64        Owner Name
ACTLOG.SCHEDNAME      VARCHAR       30        Schedule Name
ACTLOG.SERVERNAME     VARCHAR       64        Server Name
<...lines omitted...>
```

To save the output in a file use the redirection operator and the **-tab** command option to the **dsmadmc** command:

```
dsmadmc -id=admin -password=admin -tab “macro desc.sql > file.out”
```

This would produce a result like what is shown in the following sample in Example A-14 below.

Example: A-14 Example output from the desc.sql macro

```
ACTLOG.DATE_TIME                        TIMESTAMP      0       Date/Time
ACTLOG.DOMAINNAME                       VARCHAR        30      Policy Domain Name
ACTLOG.MESSAGE                          VARCHAR        250     Message
ACTLOG.MSGNO                            INTEGER        0       Message number
ACTLOG.NODENAME                         VARCHAR        64      Node Name
ACTLOG.ORIGINATOR                       VARCHAR        20      Originator
ACTLOG.OWNERNAME                        VARCHAR        64      Owner Name
ACTLOG.SCHEDNAME                        VARCHAR        30      Schedule Name
ACTLOG.SERVERNAME                       VARCHAR        64      Server Name
ACTLOG.SESSID                           INTEGER        0       Sess Number
ACTLOG.SEVERITY                         ENUMERATED     5       Message severity
<...lines omitted...>
```

However since the column headers will not appear in the **-tab** output, you could use the following version of `desc.sql`, shown in Example A-15 below, instead:

Example: A-15 SQL query to display data for all tables from syscat.columns without headers

```
SELECT -
CHAR(CONCAT(CONCAT(t.tabname,'.'),c.colname),35) AS "TC",-
CHAR(COALESCE(NULLIF(SUBSTR(c.typename,1,POSSTR(c.typename,'(')-1),''),c.typename),10),-
CHAR(c.length,5),-
c.remarks -
FROM syscat.columns AS c, syscat.tables AS t -
WHERE c.tabname = t.tabname -
ORDER BY tc
```

It is also possible to add text between columns in the row output by enclosing the text in quotation marks (“), as shown in the next sample in Example A-16 below. However to make it more interesting we have enclosed the SQL output with a header and a footer to create a viewable HTML page, and it works on both UNIX and DOS.

Example: A-16 Creating HTML output with SQL select

```
echo "<html>" > volume_access.html
echo "<head>" >> volume_access.html
echo "<title>Tivoli Storage Manager version 5.1</title>" >> volume_access.html
echo "</head>" >> volume_access.html
echo "<body>" >> volume_access.html
echo "<table>" >> volume_access.html
echo "<td><font size=-1 face="Sans-Serif">Volume</font></td>" >> volume_access.html
echo "<td> </td>" >> volume_access.html
echo "<td><font size=-1 face="Sans-Serif">Access</td>" >> volume_access.html

dsmadmc -id=admin -password=admin -tab "select '<tr><td><font size=-2
face="Sans-Serif">',char(VOLUME_NAME,30),'</font></td><td width=20>&nbs
p;</td><td><font size=-2 face="Sans-Serif">',ACCESS,'</font></td></tr>' from volumes where access like
 'R%' >> volume_access.html"

echo "</table>" >> volume_access.html
echo "</body>" >> volume_access.html
echo "</html>" >> volume_access.html
```

The following is a sample output view by a browser in Figure A-2 below, from running a script with the commands shown in Example A-16.

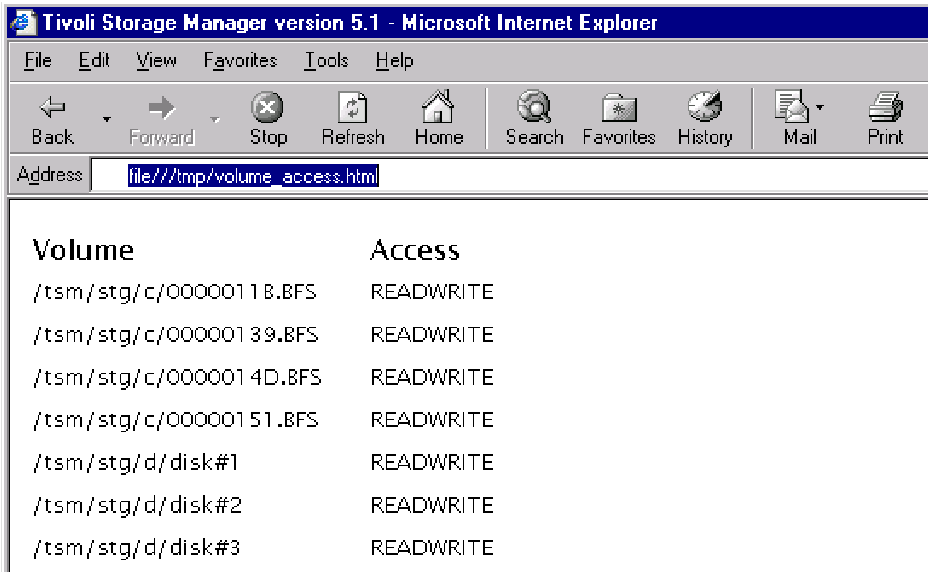

Figure A-2 HTML result

The output sample below in Example A-17 is part of the generated HTML source file.

Example: A-17 Part of the generated HTML source file

```
<<html>
<head>
<title>Tivoli Storage Manager version 5.1</title>
</head>
<body>
<table>
<td><font size=-1 face=Sans-Serif>Volume</font></td>
<td> </td>
<td><font size=-1 face=Sans-Serif>Access</td>
<tr><td><font size=-2 face=Sans-Serif>  /tsm/stg/c/0000011B.BFS         </font></td>
<td width=20> </td>
<td><font size=-2 face=Sans-Serif>  READWRITE       </font></td></tr>
<...lines omitted...>
```

Illustrative examples

The following questions have a corresponding select SQL query appended, they are provided to illustrate the versatility of the query interface.

1. Which client nodes currently locked from server access?

   ```
   select node_name from nodes where locked='YES'
   ```

2. Which administrative clients currently locked from server access?

   ```
   select admin_name from admins where locked='YES'
   ```

3. Which client nodes that has not specified the correct password lately?

   ```
   select node_name from nodes where invalid_pw_count <> 0
   ```

4. Which administrative clients that has not specified the correct password lately?

   ```
   select admin_name from admins where invalid_pw_count <> 0
   ```

5. Which nodes in the `WINDOWS` policy domain are not associated with the daily backup schedule `STANDARD`?

   ```
   select node_name from nodes where domain_name='WINDOWS' and node_name not in -
   ```

```
(select node_name from associations -
where domain_name='WINDOWS' and schedule_name='STANDARD')
```

6. Which administrators have policy authority?

```
select admin_name from admins -
where upper(system_priv) <> 'NO' or upper(policy_priv) <> 'NO'
```

7. What messages of type E (ERROR) or W (WARNING) have been issued in the time period for which activity log records have been maintained?

```
select date_time,msgno,message from actlog where severity='E' or severity='W'
```

8. Which administrative schedules have been defined or altered by administrator ADMIN?

```
select schedule_name from admin_schedules where chg_admin='ADMIN'
```

9. What are the relative administrative schedule priorities?

```
select schedule_name,priority from admin_schedules order by priority
```

10. Which mananagement classes have an archive copy group with a retention period greater than 365 days?

```
select domain_name,set_name,class_name -
from ar_copygroups where retver='NOLIMIT' or cast(retver as integer) > 365
```

11. Which management classes specify more than 5 backup versions?

```
select domain_name,set_name,class_name -
from bu_copygroups where verexists = 'NOLIMIT' or cast(verexists as integer) > 5
```

12. Which client nodes are using the client option set named SECURE?

```
select node_name from nodes where option_set='SECURE'
```

13. How many client nodes are in each policy domain?

```
select domain_name,num_nodes from domains
```

14. How many files have been archived from each node?

```
select node_name,count(*) from archives group by node_name
```

15. Which clients are using space management?

```
select node_name from auditocc where spacemg_mb <> 0
```

16. If the reclamation threshold were to be changed to 50 percent for storage pool TAPE, how many volumes would be reclaimed?

```
select count(*) from volumes -
where stgpool_name='TAPE' and upper(status)='FULL' and pct_utilized < 50
```

17. If the DAILY management class in the STANDARD policy domain is changed or deleted, how many backup files would be affected for each node?

```
select node_name, count(*) as "Files" -
from backups where class_name='DAILY' and -
node_name in (select node_name from nodes where domain_name='STANDARD') -
group by node_name
```

18. For all active client sessions, determine how long have they been connected and their effective throughput in bytes per second.

```
select session_id as "Session", -
client_name as "Client", state as "State", -
current_timestamp-start_time as "Elapsed Time", (-
cast(bytes_sent as decimal(18,0)) / cast((current_timestamp-start_time)seconds as
decimal(18,0))) as "Bytes sent/second", -
(cast(bytes_received as decimal(18,0)) / cast((current_timestamp-start_time)seconds as
decimal(18,0))) as "Bytes received/second" -
from sessions
```

19. How long have the current background processes been running and what is their effective throughput in time and files per second?

```
select process_num as "Number", process, -
current_timestamp-start_time as "Elapsed Time", -
(cast(files_processed as decimal(18,0)) / cast((current_timestamp-start_time)seconds as
ecimal(18,0))) as "Files/second",-
(cast(bytes_processed as decimal(18,0)) / cast((current_timestamp-start_time)seconds as
decimal(18,0))) as "Bytes/second" -
from processes
```

20. How many client nodes are there for each platform type?

```
select platform_name,count(*) as "Number of Nodes" from nodes group by platform_name
```

21. How many filespaces does each client node have, listed in default ascending order?

```
select node_name, count(*) as "number of filespaces" -
from filespaces group by node_name order by 2
```

22. How to display all columns for all tables from `syscat.columns` without headers

```
select char(concat(concat(t.tabname,'.'),c.colname),35) as "TC",-
char(coalesce(nullif(substr(c.typename,1,posstr(c.typename,'(')-1),''),c.typename),10),-
char(c.length,5),-
c.remarks -
from syscat.columns as c, syscat.tables AS t -
where c.tabname = t.tabname -
order by tc
```

23. How to examine which volumes are UNAVAILABLE?

```
select VOLUME_NAME,ACCESS from volumes where access = 'UNAVAILABLE'
```

24. How to examine which volumes have more than three write errors?

```
select VOLUME_NAME,WRITE_ERRORS from volumes where write_errors > 3
```

25. How to examine which volumes read errors?

```
select VOLUME_NAME,READ_ERRORS from volumes where read_errors > 0
```

26. How to examine which volumes have an error state different from `No`?

```
select VOLUME_NAME,ERROR_STATE from volumes where error_state != 'No'
```

27. How to examine which volumes have access different from `READWRITE`?

```
select VOLUME_NAME,ACCESS from volumes where access != 'READWRITE'
```

28. How to examine which volumes have less than ten percent utilization in device classs beginning with the letters `SUN`?

```
select volume_name,pct_utilized,status,access from volumes where pct_utilized < 10 and
devclass_name like 'SUN%'
```

29. How to examine which volumes does not have a access beginning with the letters `READ`?

```
select volume_name,pct_utilized,pct_reclaim,stgpool_name,status,access from volumes
where access not like 'READ%'
```

30. How to list the content of all volumes and display the filesize in MB and order by client node name, volume name and size?

```
select node_name,-
volume_name, -
decimal(file_size/1024/1024,12,2) mb,-
concat(substr(file_name,1,posstr(file_name,' ')-1),-
substr(file_name,posstr(file_name,' ')+1)) -
from contents -
order by node_name,volume_name,mb
```

31.How to find all clients which stores their backup data in the `DISKPOOL` storage pool?

```
select node_name as "CLIENT NODENAME", -
bu_copygroups.destination as "STGPOOL DESTINATION", -
nodes.domain_name as "CLIENT DOMAIN", -
bu_copygroups.domain_name as "COPYGROUP DOMAIN" -
from nodes,bu_copygroups where -
nodes.domain_name = bu_copygroups.domain_name and -
bu_copygroups.destination=upper('diskpool') and -
bu_copygroups.set_name=upper('active') -
order by nodes.domain_name
```

32.How to find all volumes and their status which have data for a specified client?

```
select volumeusage.volume_name,-
volumes.access,-
volumes.error_state,-
volumeusage.stgpool_name  -
from volumeusage,volumes  -
where volumeusage.node_name='ONE-ON-ONE' and volumeusage.volume_name =
volumes.volume_name -
order by volume_name
```

33.How to find all storage pools where a client (DUDE) has stored data?

```
select distinct(STGPOOL_NAME) from OCCUPANCY where node_name='DUDE'
```

ODBC Driver

Tivoli Storage Manager provides an ODBC driver for Windows. The driver supports the ODBC 3.5 specification. Because Tivoli Storage Manager supports only the **`select`** SQL query command, the driver does not conform to any ODBC API or SQL grammar conformance level. After you install this driver, you can use a spreadsheet or database application that supports ODBC to access the database for information, such as Lotus Approach, Lotus 1-2-3 or Microsoft Access.

Installation

The ODBC driver set-up is included in the sever installation package. The server installation program can install the ODBC driver and set the corresponding registry values for the driver and data sources. The following example screen shows the selection to choose for the installation of the client ODBC driver.

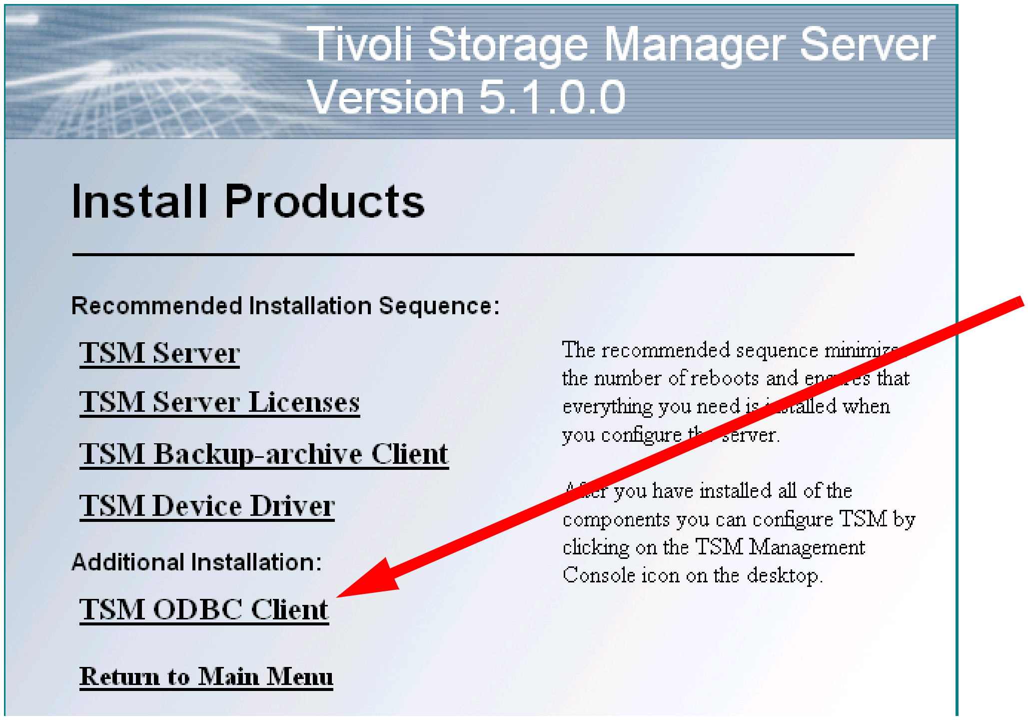

Figure A-3 Client ODBC driver installation screen

Connecting to the server

To open the database through an ODBC application, you must log on to the server (the defined data source). Use the *name* and *password* of a *registered administrator*. After you log on to the server, you can perform query functions provided by the ODBC application to access database information.

Lotus Approach

The following example suite is using Lotus Approach to establish the ODBC connection between a Windows PC and the IBM Tivoli Storage Manager 5.1 server. First you need to create a database, in this example we create a **Blank Database** as is shown in the example below.

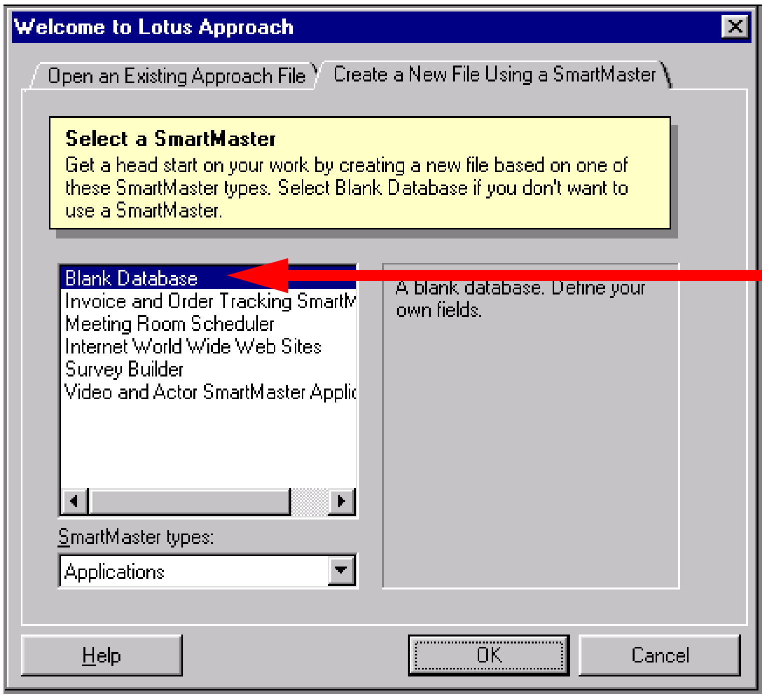

Figure A-4 Create a blank database

Next we select the ODBC source which in this example is **TSM ODBC Drive (*)** as pointed to below and **Create type** as **ODBC Data Sources (*)**.

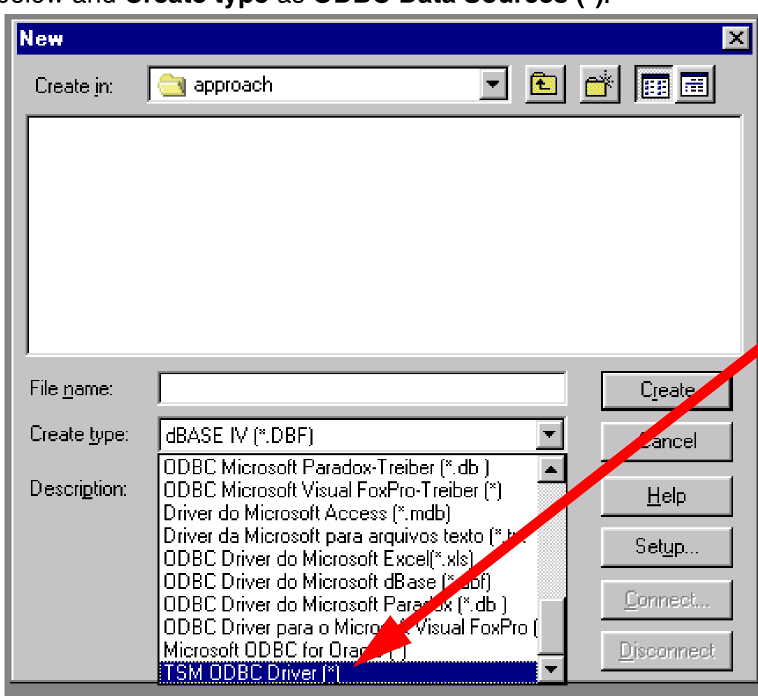

Figure A-5 Selecting ODBC source

Now we get a ODBC logon screen and we are prompted to enter our Tivoli Storage Manager server IP address or DNS hostname, the TCP portnumber by which the ODBC driver will send and receive data to the server, the name and password for the administrator, since it is a security restricted operation to access the Tivoli Storage Manager server through the SQL interface.

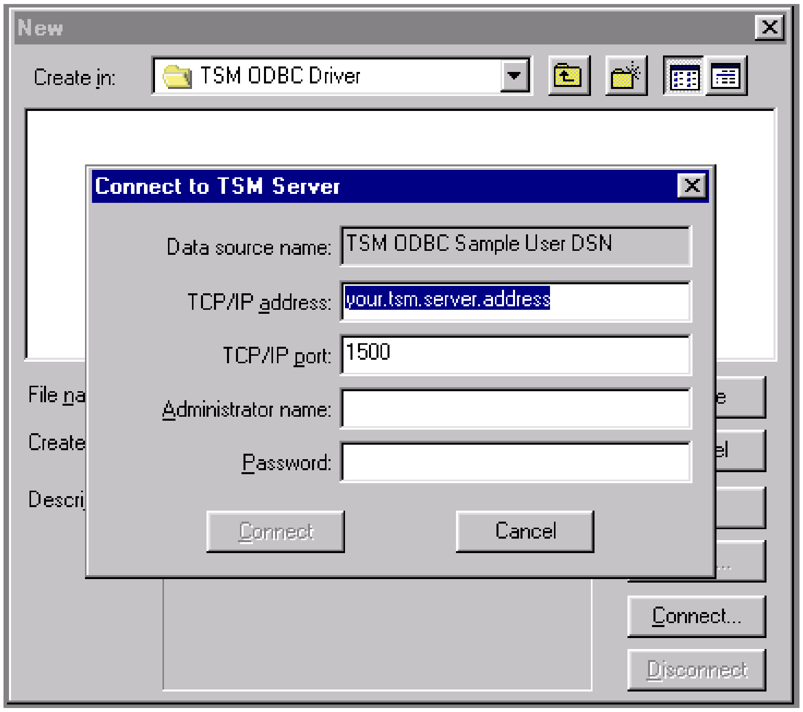

Figure A-6 Default ODBC logon window

In Figure A-7 we have entered our Tivoli Storage Manager servers DNS hostname (`crete`), we use the default TCP portnumber (`1500`) and the administrators name is also the default one (`admin`) and not shown is the password (*****).

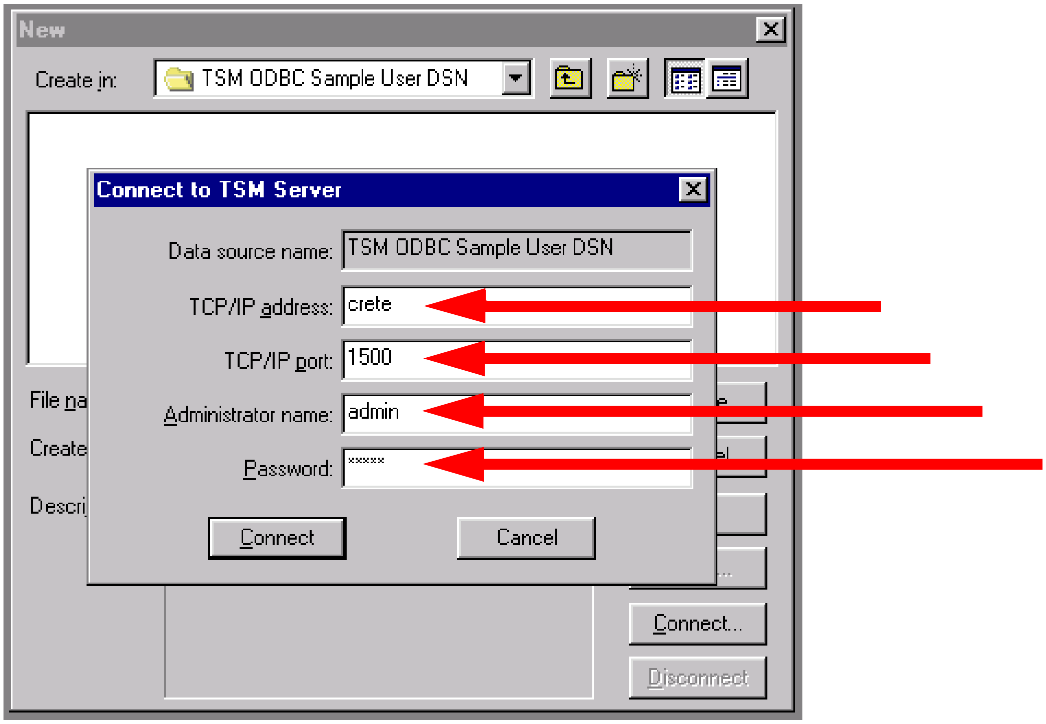

Figure A-7 Customized ODBC logon window

When we have filled in all the values we select **Connect** which bring us to the next screen shown below where we have established the ODBC Tivoli Storage Manager Client/Server connection.

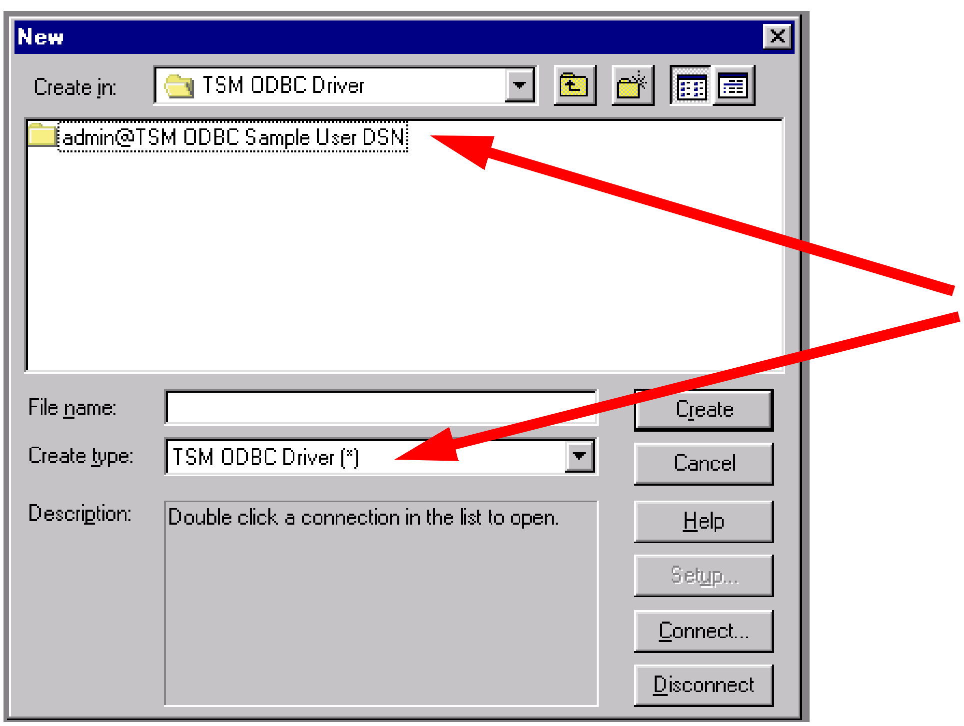

Figure A-8 Established ODBC Tivoli Storage Manager client/server connection

To gain access to the Tivoli Storage Manager servers database, double click on the **admin@TSM ODBC Sample User DSN** entry, and the output will be similar to the one shown in the next example below.

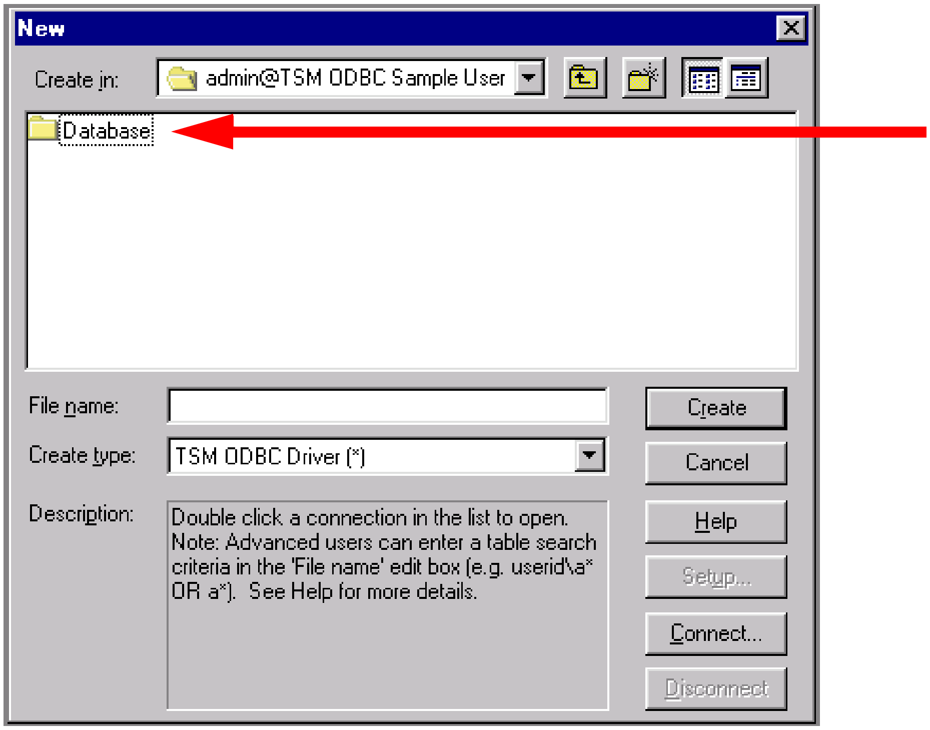

Figure A-9 ODBC acquired information from the Tivoli Storage Manager server

To access different tables in the Tivoli Storage Manager server, double click on the **Database** entry, and the output will be similar to the one shown in the next example below verifying that the ODBC connection is working.

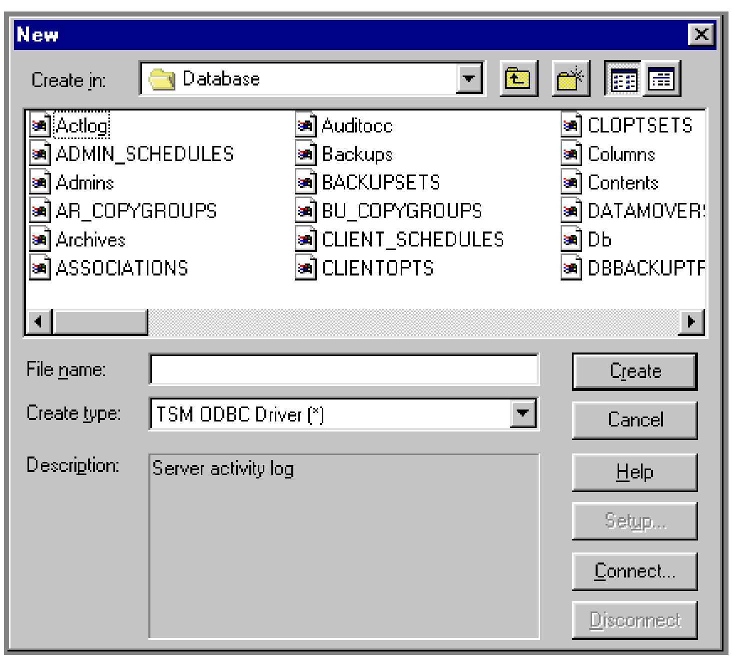

Figure A-10 ODBC acquired tables from the Tivoli Storage Manager server

Microsoft Excel

The following example suite is using Lotus Approach to establish the ODBC connection between a Windows PC and the IBM Tivoli Storage Manager 5.1 server. First you need to create a new `Workbook` or `Datasheet`. Then you create a new database query through the menu **Data->Get External Data->New Database Query**, as is shown in the example below.

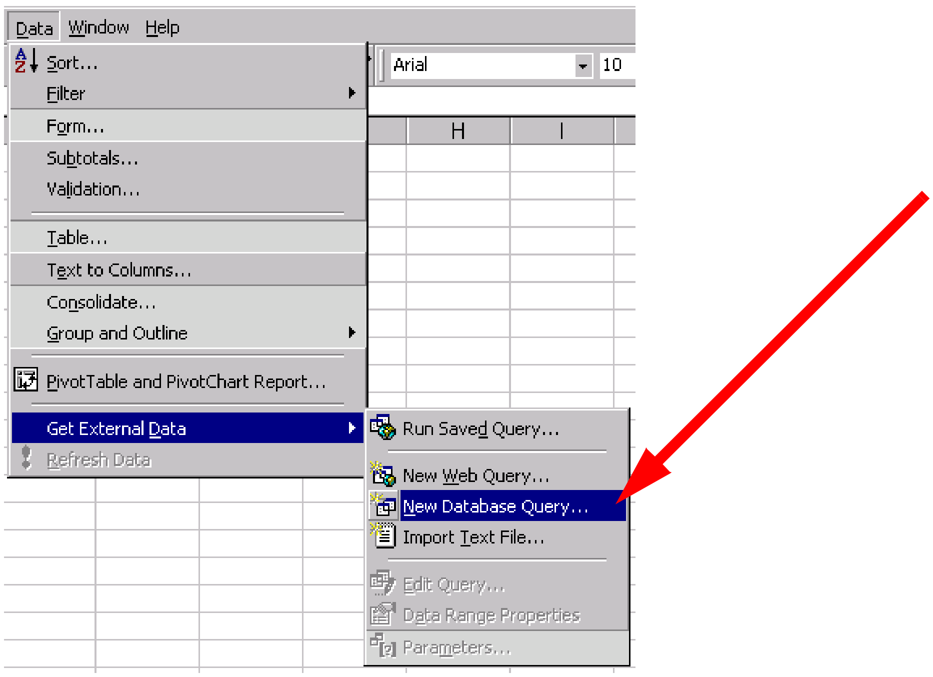

Figure A-11 Creating a new database query

Then select **TSM ODBC Sample User DSN** as **Data Source** in the pop up window that will appear as shown below.

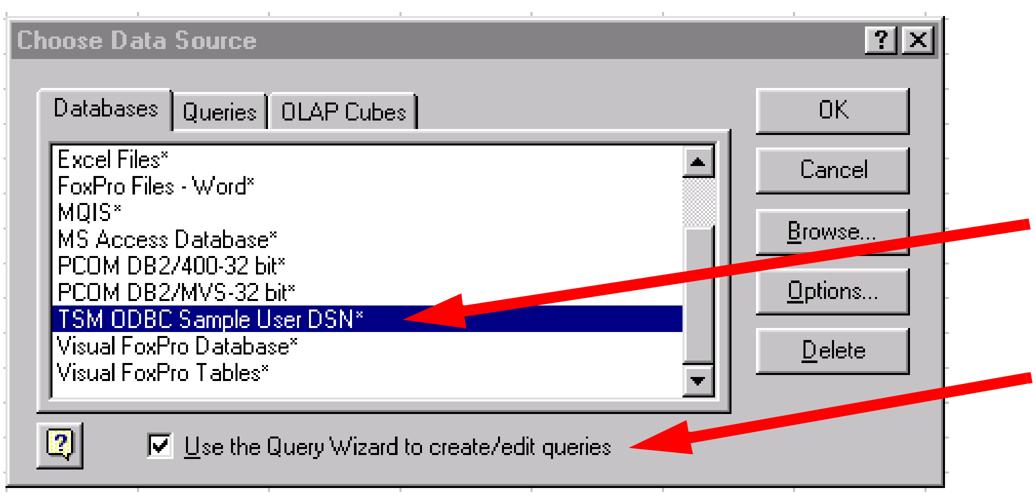

Figure A-12 Choosing data source

Now we get a ODBC logon screen and we are prompted to enter our Tivoli Storage Manager server IP address or DNS hostname, the TCP portnumber by which the ODBC driver will send and receive data to the server, the name and password for the administrator, since it is a security restricted operation to access the Tivoli Storage Manager server through the SQL interface. Please note that for simplicity later on in the dialog chain, we have left the selection box for the **Query Wizard** checked, it will assist in creating queries.

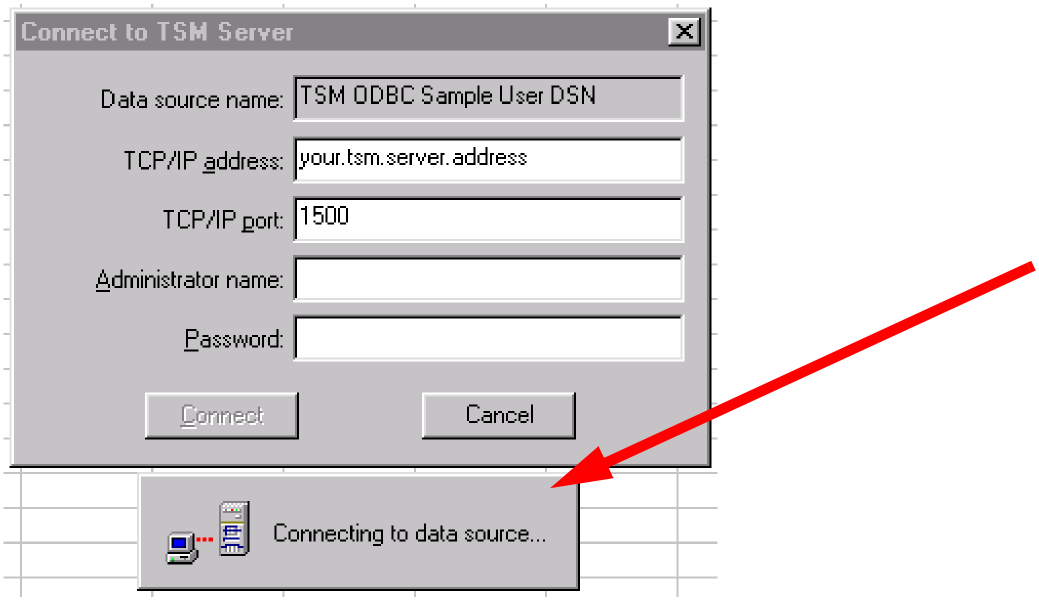

Figure A-13 Default ODBC logon window

In the next screen we have entered our Tivoli Storage Manager servers DNS hostname (`crete`), we use the default TCP portnumber (`1500`) and the administrators name is also the default (`admin`) and the password (*****) is not shown.

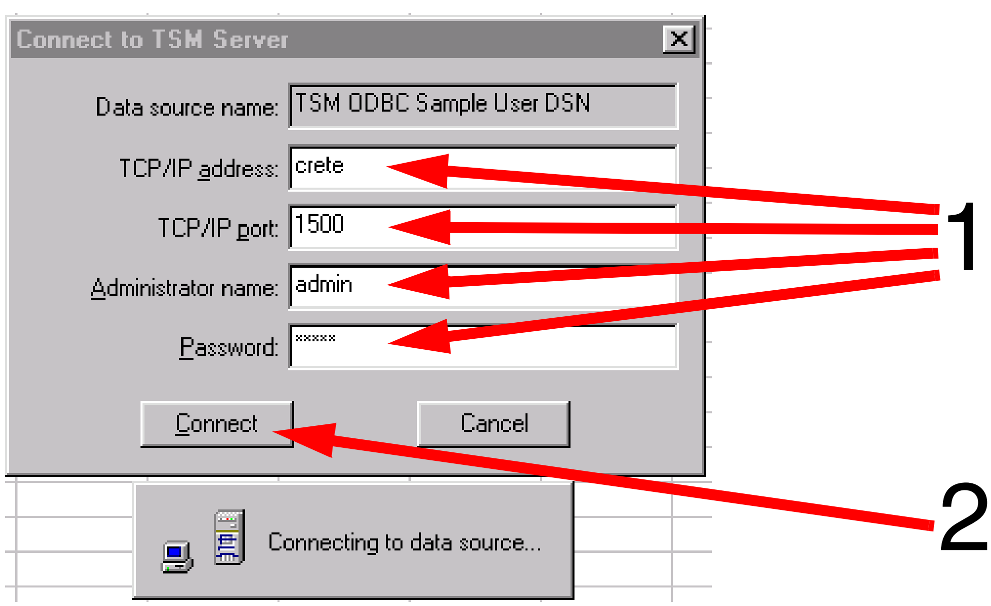

Figure A-14 Customized ODBC logon window

When we have filled in all the values we select **Connect** which bring us to the next screen shown below where we have established the ODBC Tivoli Storage Manager Client/Server connection.

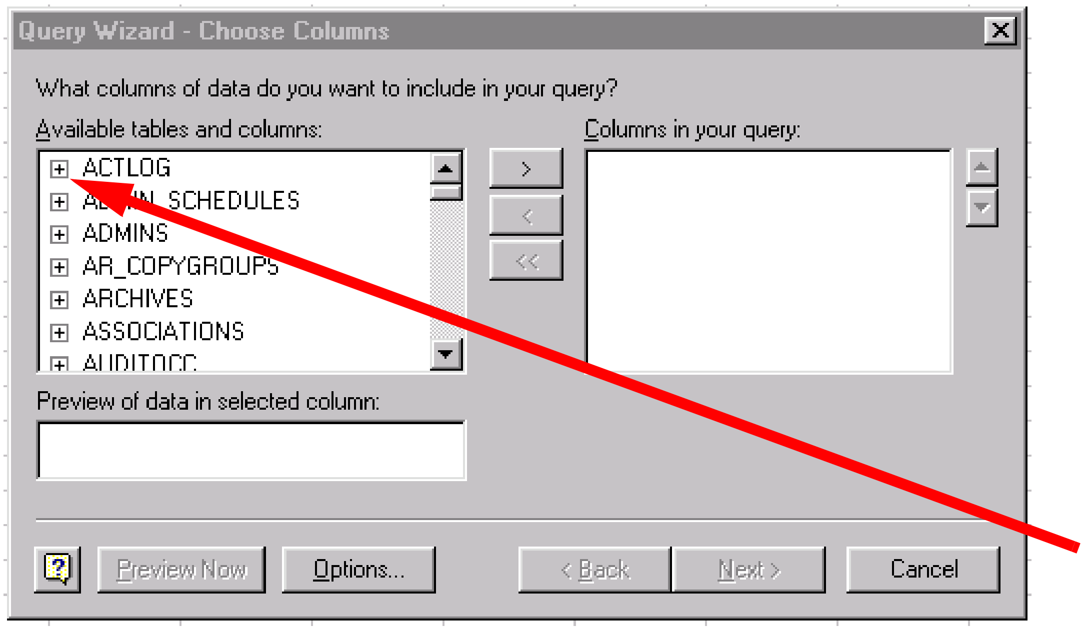

Figure A-15 Table selection

In the popup dialog window above, we can select the desired tables and columns. In the next example we have selected the + box for `ACTLOG` and then selected the `DATE_TIME`, `MSGNO` and `MESSAGE` columns. Then select **Next** and continue with the next dialog in the sequence.

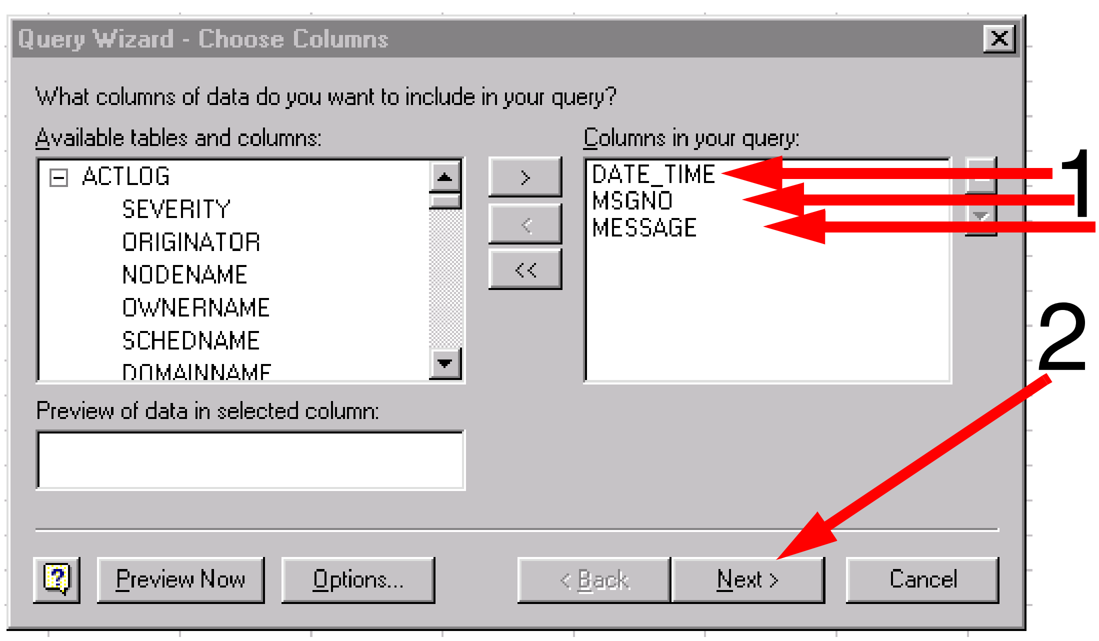

Figure A-16 Column selection

After selecting **Next** the **Filter Data** popup dialog window will appear. Here it is possible to select "filtering" criteria for the query, but we just select **Next** to continue to the next dialog screen.

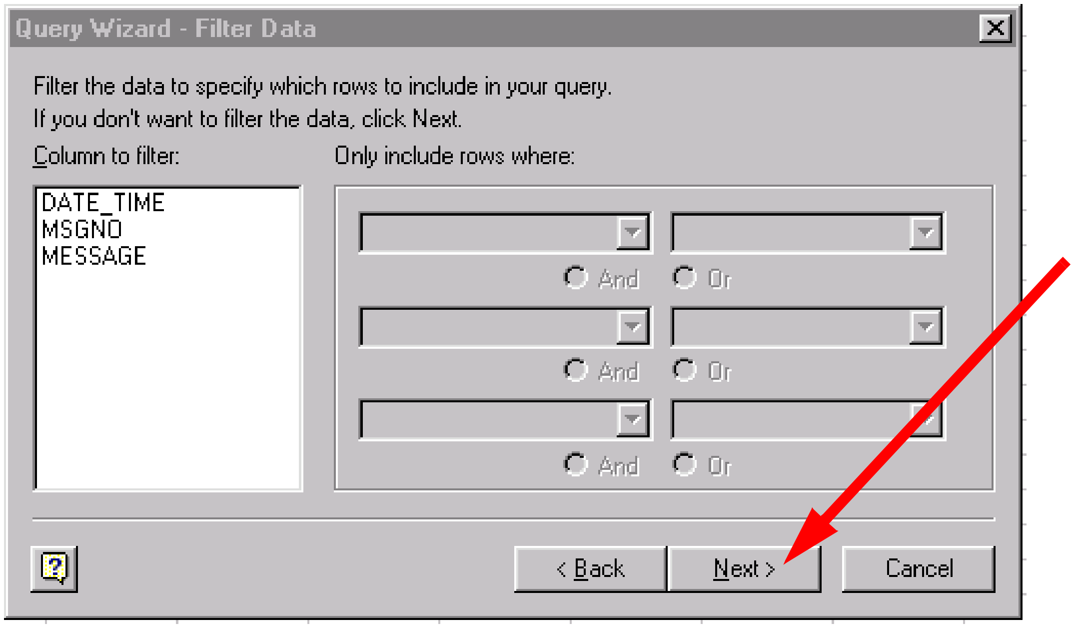

Figure A-17 Filter data

The **Sort Order** popup dialog window allows definition on how the report is sorted, but we just select **Next** to continue to the next dialog screen.

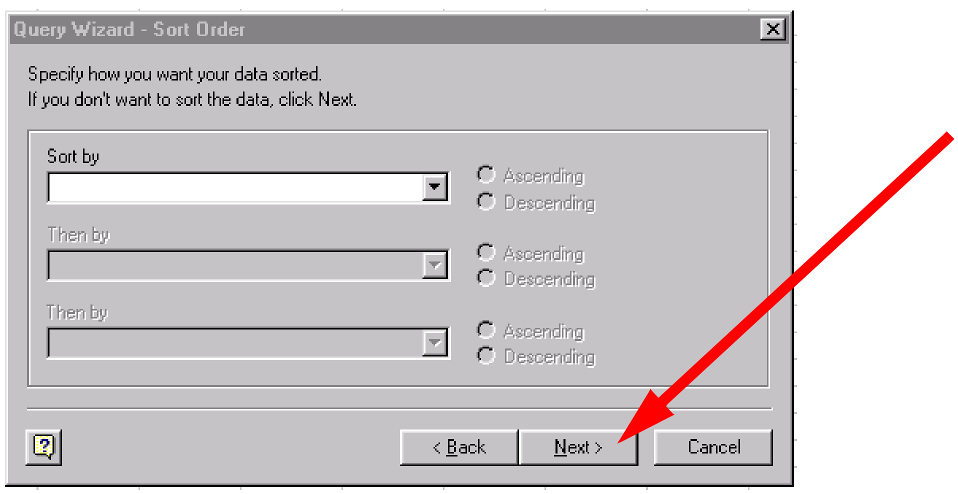

Figure A-18 Sort order

Finally we arrive at the **Finish** popup dialog window and here we can choose how to store the query and the query result. For this example we just select **Finish**.

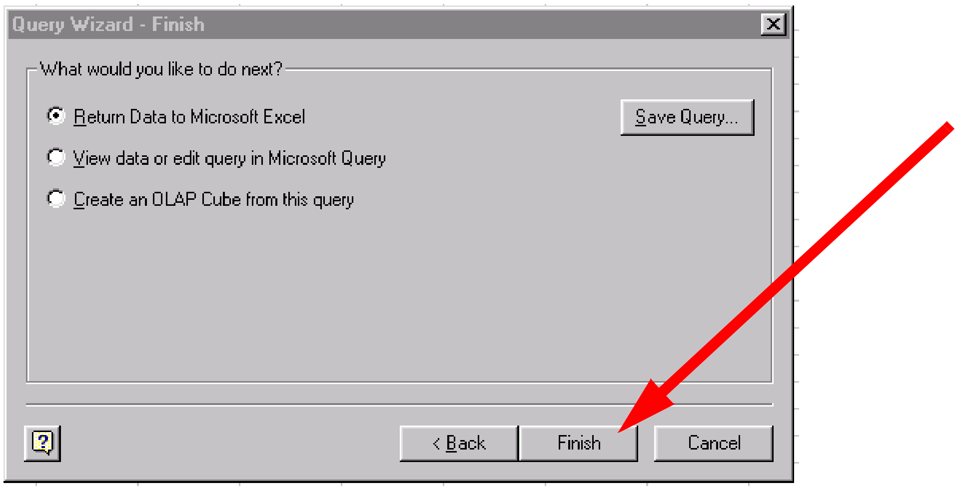

Figure A-19 Finish dialog window

The select SQL query statement generated will look like the following:

```
SELECT ACTLOG.DATE_TIME, ACTLOG.MSGNO, ACTLOG.MESSAGE FROM ACTLOG ACTLOG
```

The result of the query can be quite extensive since we did not apply any filtering but a small sample of the output from the Tivoli Storage Manager server `crete` is shown below.

C7 = ANR2103I Activity log pruning completed: 98 records removed.

	A	B	C
1	**DATE_TIME**	**MSGNO**	**MESSAGE**
2	3/11/02 0:00	2562	ANR2562I Automatic event record deletion started.
3	3/11/02 0:00	2565	ANR2565I 0 schedules for immediate client actions have been deleted.
4	3/11/02 0:00	2563	ANR2563I Removing event records dated prior to 03/01/02 00:00:00.
5	3/11/02 0:00	2102	ANR2102I Activity log pruning started: removing entries prior to 03/10/02 00:00:00.
6	3/11/02 0:00	2564	ANR2564I Automatic event record deletion ended - 0 records deleted.
7	3/11/02 0:00	2103	ANR2103I Activity log pruning completed: 98 records removed.

Figure A-20 Sample result from the Excel ODBC SQL query

B

Microcode update of an IBM LTO Ultrium tape drive

This appendix contains the procedures for updating microcode:

Updating the microcode using NTUtil program

The following procedure describes how to update the drive microcode in the 3580, 3581, and 3583 products, over the SCSI bus using NTUtil on a Windows based Tivoli Storage Manager server. For downloading the latest Microcode use:

http://www.storage.ibm.com/hardsoft/tape/index.html

1. If you have not already done so, install the IBM Ultrium Device Driver. For further information how to install the IBM Ultrium Device Driver, please refer to the Redbook *Using IBM LTO Ultrium with Open Systems,* SG-6502-00.
2. If you are using a software application that utilizes a different device driver, you will need to stop that device driver. In case of Tivoli Storage Manager the LTO library is under control of Tivoli Storage Manager and the drives are under control of the Windows 2000 IBM Ultrium Device Driver. Use the Tivoli Storage Manager Management Console to stop the Tivoli Storage Manager Device Driver.
3. Open a DOS window, go to the directory where IBM Ultrium Device Driver is installed (e.g. C:\Program Files\IBM Corporation\IBM Ultrium Device Driver\5.0.2.4) and start NTUtil by typing **ntutil** at the command prompt. The following screen should appear:

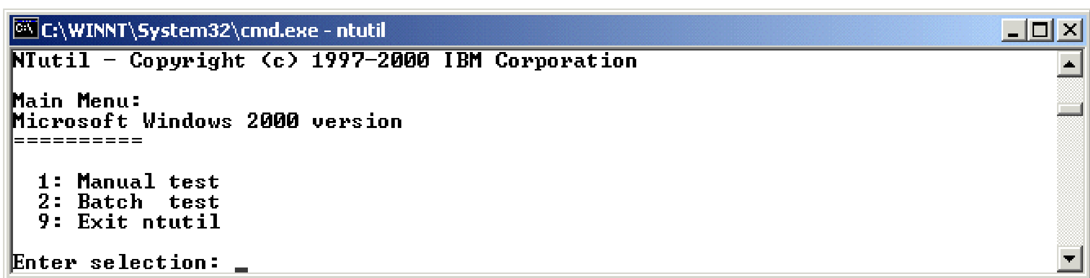

Figure B-1 NTUtil window

4. Type 1 and press Enter, the following NTUtil main menu should appear:

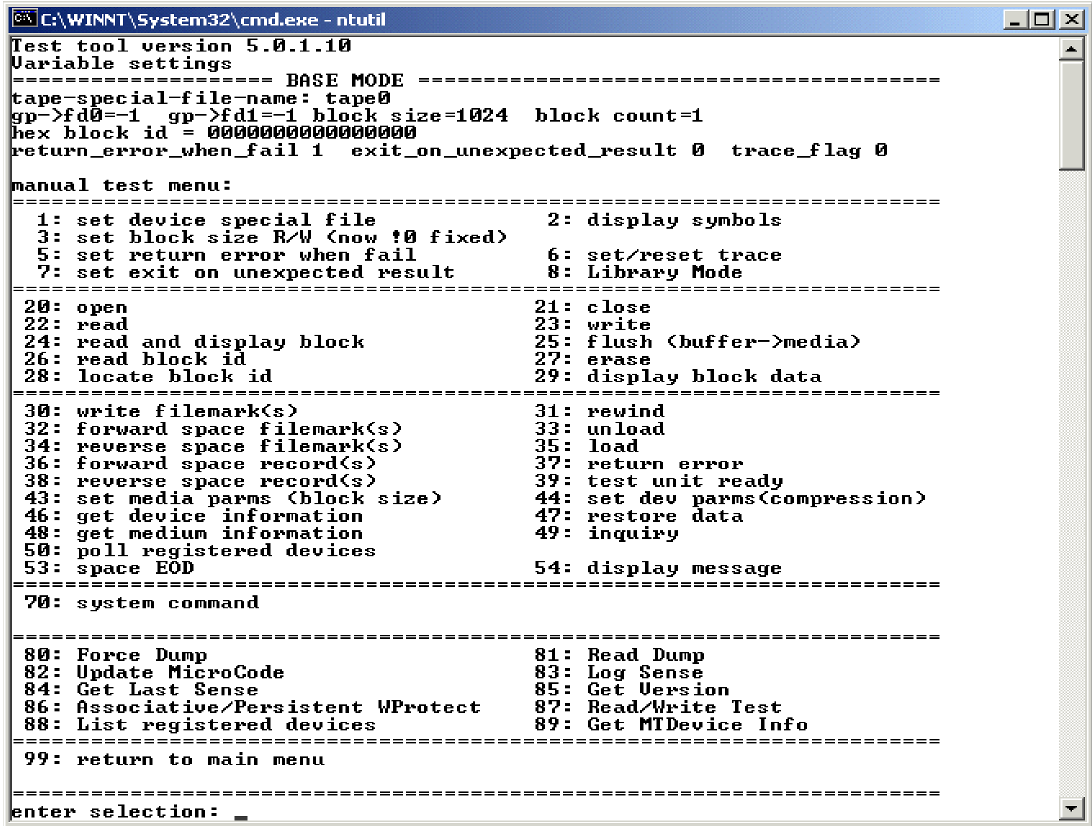

Figure B-2 NTUtil main window

If your window does not come up with a scroll bar that allows you to move up and down the menu items in NTUtil, right click in **Title Bar** portion of the window (top of the window, usually dark blue in color), then left click on **Properties** to bring up the "Command Prompt" Properties window. Click on the **Layout** tab and adjust the **Window Size Width** and **Height** numbers until scroll bars appear.

5. Ensure that the firmware image file to be downloaded to the drive has the correct naming convention: x.fmr to xxxxxxxx.fmr (1 to 8 characters with an extension of "fmr"), and that the image file resides in the same directory that NTUtil is started from.
6. Enter option **20** to open the drive, autoloader or library. Use open mode **1** for read/write operation.

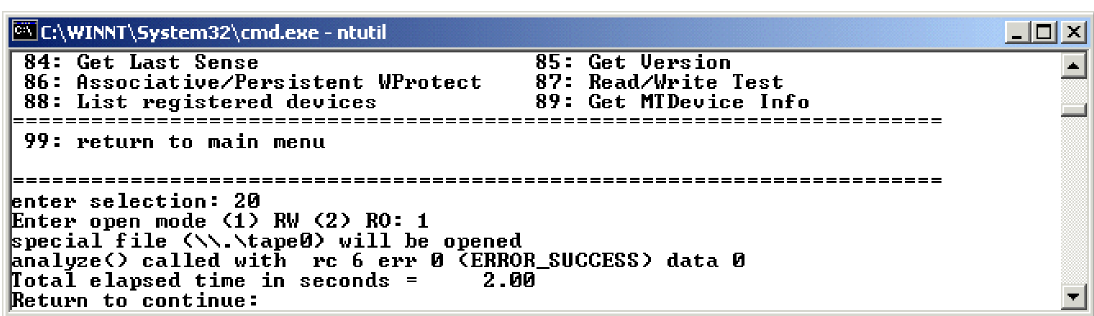

Figure B-3 Output of the open command

7. Now the media changer is in state open. Press **Return** to go back to the main window.

8. To select the drive you want to update, enter option **1**, press **Return** and enter the device name of the drive you want to update.

```
C:\WINNT\System32\cmd.exe - ntutil
 82: Update MicroCode                        83: Log Sense
 84: Get Last Sense                          85: Get Version
                                             87: Read/Write Test
 88: List registered devices
========================================================================
 99: return to main menu

========================================================================
enter selection: 1
Enter device special file (tape path) name or return for no change: mt0.2.0.3
Total elapsed time in seconds =     6.00
Return to continue: _
```

Figure B-4 Drive selection

9. Press **Return** to go back to the main window. Enter option **82** to update the microcode at the NTUtil selection prompt. The following screen will appear prompting you to enter the microcode image name. Only enter the characters and not the "fmr" file extension

```
C:\WINNT\System32\cmd.exe - ntutil
========================================================================
 80: Force Dump                              81: Read Dump
 82: Update MicroCode                        83: Log Sense
 84: Get Last Sense                          85: Get Version
                                             87: Read/Write Test
 88: List registered devices
========================================================================
 99: return to main menu

========================================================================
enter selection: 82
Enter microcode image name (1 to 8 characters), or return to skip: 22U0
```

Figure B-5 Microcode update

10. After pressing the Enter or Return key, the following screen will display and the status light located on the bezel or faceplate of the drive will flash amber. As the NTUtil screen instructs, do not touch the drive until the firmware load has completed.

```
C:\WINNT\System32\cmd.exe - ntutil
======================================================================
enter selection: 82
Enter microcode image name (1 to 8 characters), or return to skip: 22U0
Image size = 1048576
Transferring 65536 image bytes, total bytes = 65536
Transferring 65536 image bytes, total bytes = 131072
Transferring 65536 image bytes, total bytes = 196608
Transferring 65536 image bytes, total bytes = 262144
Transferring 65536 image bytes, total bytes = 327680
Transferring 65536 image bytes, total bytes = 393216
Transferring 65536 image bytes, total bytes = 458752
Transferring 65536 image bytes, total bytes = 524288
Transferring 65536 image bytes, total bytes = 589824
Transferring 65536 image bytes, total bytes = 655360
Transferring 65536 image bytes, total bytes = 720896
Transferring 65536 image bytes, total bytes = 786432
Transferring 65536 image bytes, total bytes = 851968
Transferring 65536 image bytes, total bytes = 917504
Transferring 65536 image bytes, total bytes = 983040
Transferring 65536 image bytes, total bytes = 1048576
======================================================================
Waiting for tape device to write flash-prom
Do not touch drive until complete (Avg 45 - 50 sec)
Microcode updated, standby for the drive to initialize
```

Figure B-6 Microcode update status

11. After the microcode load completes successfully, the drive Status Light will be solid green and the following message on the NTUtil screen will appear.

```
C:\WINNT\System32\cmd.exe - ntutil
Transferring 65536 image bytes, total bytes = 1048576
======================================================================
Waiting for tape device to write flash-prom
Do not touch drive until complete (Avg 45 - 50 sec)
Microcode updated, standby for the drive to initialize

analyze() called with  rc 1 err 0 (ERROR_SUCCESS) data 0
Total elapsed time in seconds =      7.00
Return to continue: _
```

Figure B-7 Microcode update completed

12. If you have a IBM Ultrium 3583 library with more than one drive installed you have to repeat steps 8 to 12 for each drive.

13. Now you should exit the NTUtil program and restart the Ultrium drive, autoloader or library.

Updating the microcode using Ultrium Tape Library Specialist

The 3583 LTO Ultrium tape library with an RMU (Remote Management Unit) provides remote access to the library control panel and other functions via the network. You can use this rather then NTutil to upgrade the drive microcode levels if you prefer. We noticed that the file transfer using the Web based Ultrium Tape Library Specialist took considerably longer than using NTUtil. We therefore recommend you use NTUtil.

For downloading the latest Microcode use:
`www.ssddom02.storage.ibm.com/tape/lto/tape_drive/firmware.html`

You can read more details on using the Ultrium Tape Library Specialist in the redbook *Using IBM LTO Ultrium with Open Systems, SG24-6502-00*. Here we will concentrate only on upgrading the drive microcode in order to support Server-Free Data Movement.

1. Using the 3583 operator panel buttons select **More —> Setup —> Library —>RMU**. The following panel is displayed which identifies the IP address of the library. In our case it was 9.1.38.170.

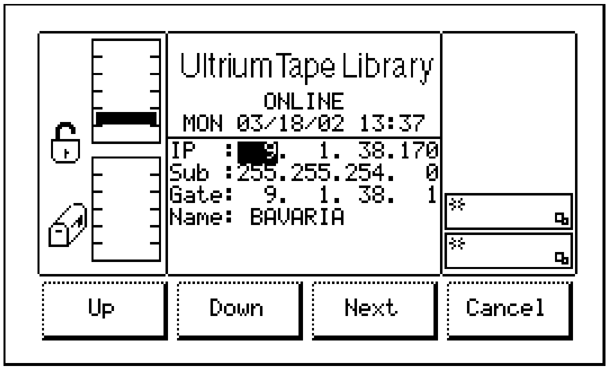

Figure B-8   LTO RMU panel

2. You can now access the Ultrium Tape Library Specialist from your favorite Web browser by entering the just identified IP address like this: http://9.1.38.170. A login screen like this will appear:

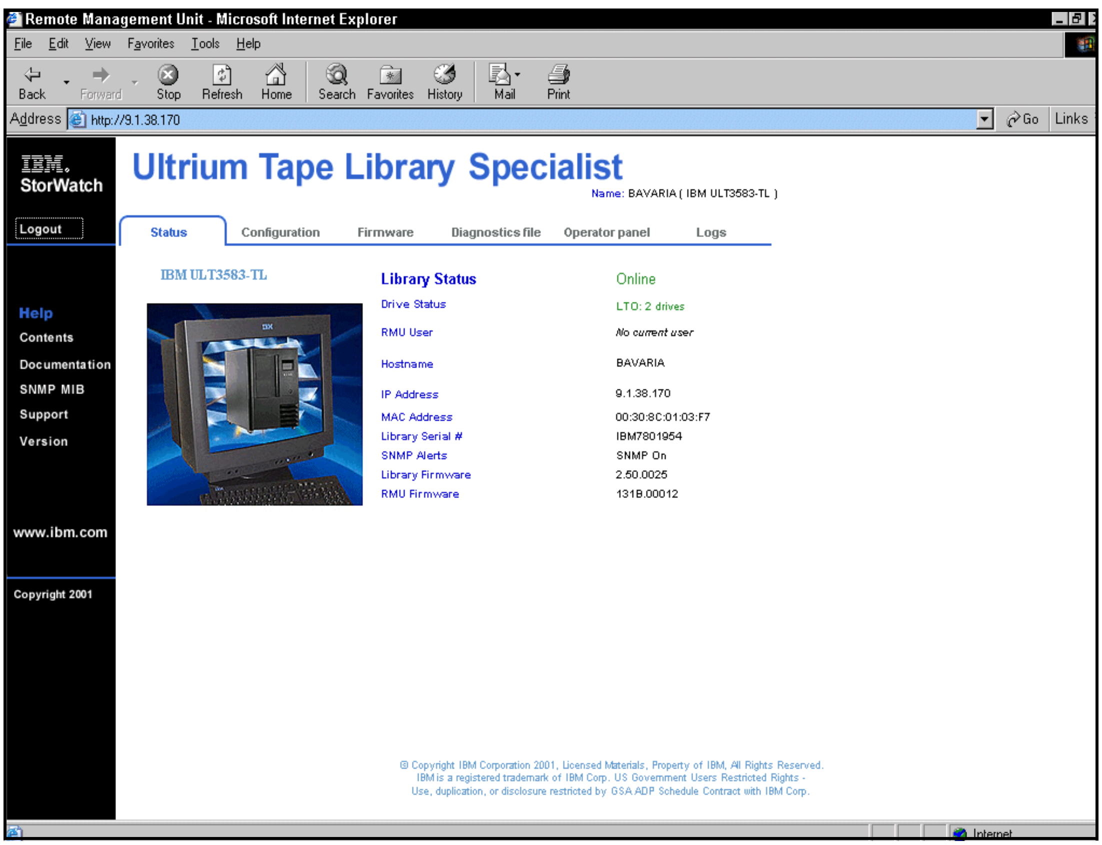

Figure B-9   LTO 3583 Library Specialist login panel

3. Any action on this panel will prompt for an administrator ID and password. They are originally set to ‘admin’ and ‘secure’ but you can change them once logged in. Now you can select the Firmware tab and from here choose the drive you want to upgrade the microcode on. You can browse for the ‘*.fmr’ file in which the microcode is shipped. In our case we were upgrading the drives to firmware 22u0 so browsed for the ‘22u0.fmr’ file which we had downloaded from the Ultium driver Web page earlier.

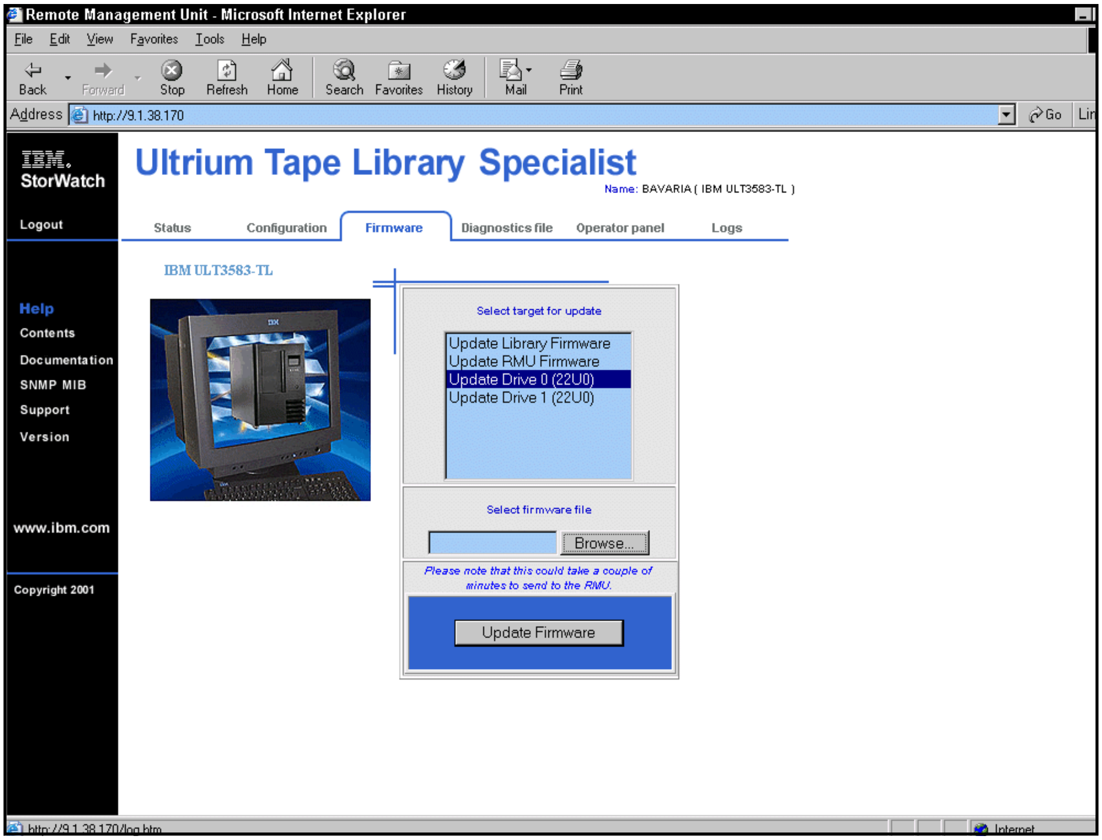

Figure B-10 Upgrading Ultrium drive firmware using Library Specialist

4. Once you click on the ‘Update Firmware’ button the drive microcode will be downloaded and installed onto the chosen drive. The following screen is displayed while the microcode download is taking place:

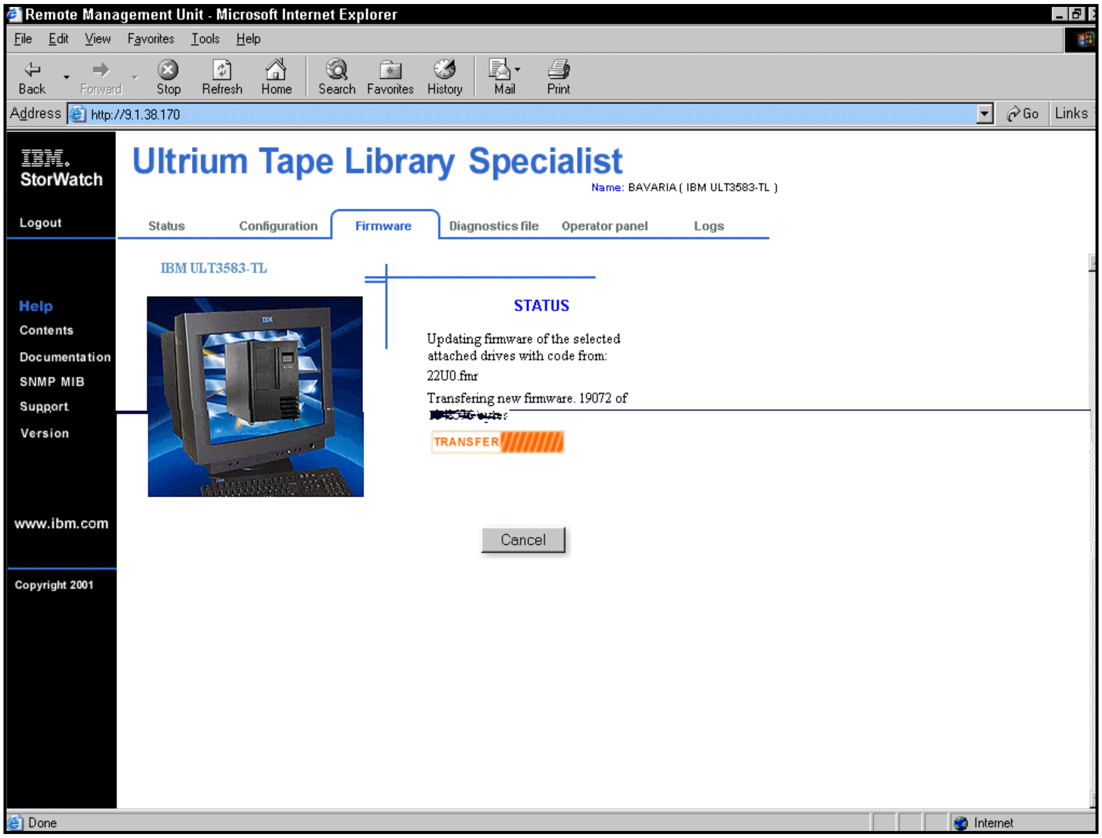

Figure B-11 Drive microcode transfer

5. If you have a IBM Ultrium 3583 tape library with more than one drive installed you have to repeat steps 3 to 5 for each drive.
6. Now you should exit the Web browser and restart the IBM Ultrium 3583 tape library.

Glossary

Agent. A software entity that runs on endpoints and provides management capability for other hardware or software. An agent has the ability to spawn other processes. An example is the Tivoli Storage Manager Client Acceptor for Windows which can spawn the Tivoli Storage Manager Client Scheduler.

Allocated storage. The space that is allocated to volumes, but not assigned.

Allocation. The entire process of obtaining a volume and unit of external storage, and setting aside space on that storage for a data set.

Arbitrated loop. A Fibre Channel interconnection technology that allows up to 126 participating node ports and one participating fabric port to communicate. See also Fibre Channel Arbitrated Loop and loop topology.

Bandwidth. A measure of the data transfer rate of a transmission channel.

Bootstrap Protocol (BOOTP). A protocol that allows a client to find both its Internet Protocol (IP) address and the name of a file from a network server.

Client. A function that requests services from a server, and makes them available to the user. A term used in an environment to identify a machine that uses the resources of the network.

Client authentication. The verification of a client in secure communications where the identity of a server or browser (client) with whom you wish to communicate is discovered. A sender's authenticity is demonstrated by the digital certificate issued to the sender.

Client-server relationship. Any process that provides resources to other processes on a network is a server. Any process that employs these resources is a client. A machine can run client and server processes at the same time.

Cluster. A loosely-coupled collection of independent systems (nodes) organized into a network for the purpose of sharing resources and communicating with each other.

Console. A user interface to a server.

Device driver. A program that enables a computer to communicate with a specific device, for example, a disk drive.

Domain Name System (DNS). A server program that supplies name-to-address conversion by mapping domain names to IP addresses. The domain name server allows users to request services of a computer by using a symbolic name, which is easier to remember than an IP address.

Dynamic Host Configuration Protocol (DHCP). An application-layer protocol that allows a machine on the network, the client, to get an IP address and other configuration parameters from the server.

Event. In the Tivoli environment, any significant change in the state of a system resource, network resource, or network application. An event can be generated for a problem, for the resolution of a problem, or for the successful completion of a task.

Fabric. The Fibre Channel employs a fabric to connect devices. A fabric can be as simple as a single cable connecting two devices. The term is often used to describe a more complex network utilizing hubs, switches, and gateways.

Fiber optic. The medium and the technology associated with the transmission of information along a glass or plastic wire or fiber.

Fibre Channel. A technology for transmitting data between computer devices at a data rate of up to 1 Gb. It is especially suited for connecting computer servers to shared storage devices and for interconnecting storage controllers and drives.

Fibre Channel Arbitrated Loop. A reference to the FC-AL standard, a shared gigabit media for up to 127 nodes, one of which can be attached to a switch fabric. See also arbitrated loop and loop topology.

Fibre Channel standard. An ANSI standard for a computer peripheral interface. The I/O interface defines a protocol for communication over a serial interface that configures attached units to a communication fabric.

File system. An individual file system on a host. This is the smallest unit that can monitor and extend. Policy values defined at this level override those that might be defined at higher levels.

Gateway. In the SAN environment, a gateway connects two or more different remote SANs with each other. A gateway can also be a server on which a gateway component runs.

Hardware zoning. Hardware zoning is based on physical ports. The members of a zone are physical ports on the fabric switch. It can be implemented in the following configurations: one to one (1:1), one to many (1:N), and many to many (N:N).

High Availability (HA) clusters. A system typically made up of two or more robust machines that mirror each other. Usually half the nodes are actively working, while the other half are quietly waiting to take over in case of a failure. HA cluster design varies according to its function.

Host. Any system that has at least one Internet address associated with it. A host with multiple network interfaces can have multiple Internet addresses associated with it.

Host bus adapter (HBA). A Fibre Channel HBA connection that allows a workstation to attach to the SAN network.

Hub. A Fibre Channel device that connects up to 126 nodes into a logical loop. All connected nodes share the bandwidth of this one logical loop. Hubs automatically recognize an active node and insert the node into the loop. A node that fails or is powered off is automatically removed from the loop.

Java. A programming language that enables application developers to create object-oriented programs that are very secure, portable across different machine and operating system platforms, and dynamic enough to allow expandability.

Java Virtual Machine (JVM). The execution environment within which Java programs run. The Java virtual machine is described by the Java Machine Specification which is published by Sun Microsystems.

JBOD. Just a Bunch Of Disks.

Logical unit number (LUN). The LUNs are provided by the storage devices attached to the SAN. This number provides you with a volume identifier that is unique among all storage servers. The LUN is synonymous with a physical disk drive or a SCSI device. For disk subsystems such as the IBM Enterprise Storage Server, a LUN is a logical disk drive. This is a unit of storage on the SAN which is available for assignment or unassignment to a host server.

LUN assignment criteria. The combination of a set of LUN types, a minimum size, and a maximum size used for selecting a LUN for automatic assignment.

LUN masking. This allows or blocks access to the storage devices on the SAN. Intelligent disk subsystems like the IBM Enterprise Storage Server provide this kind of masking.

Management Information Base (MIB). A logical database residing in the managed system which defines a set of MIB objects. A MIB is considered a logical database because actual data is not stored in it, but rather provides a view of the data that can be accessed on a managed system.

Mirroring. The creation of a mirror image of data to be reserved in the event of disk failure.

N_Port node port. A Fibre Channel-defined hardware entity at the end of a link which provides the mechanisms necessary to transport information units to or from another node.

Network File System (NFS). A distributed file system that allows users to access files and directories located on remote computers and to treat those files and directories as if they were local.

Network topology. A physical arrangement of nodes and interconnecting communications links in networks based on application requirements and geographical distribution of users.

NL_Port node loop port. A node port that supports arbitrated loop devices.

Open system. A system whose characteristics comply with standards made available throughout the industry, and therefore can be connected to other systems that comply with the same standards.

Point-to-point topology. A PtoP topology consists of a single connection between two nodes. All the bandwidth is dedicated for these two nodes.

Port. An end point for communication between applications, generally referring to a logical connection. A port provides queues for sending and receiving data. Each port has a port number for identification. When the port number is combined with an Internet address, it is called a socket address.

Port zoning. In Fibre Channel environments, port zoning is the grouping together of multiple ports to form a virtual private storage network. Ports that are members of a group or zone can communicate with each other but are isolated from ports in other zones. See also LUN masking and subsystem masking.

Protocol. The set of rules governing the operation of functional units of a communication system if communication is to take place. Protocols can determine low-level details of machine-to-machine interfaces, such as the order in which bits from a byte are sent. They can also determine high-level exchanges between application programs, such as file transfer.

Redundant Array of Independent Disks (RAID). A set of physical disks that act as a single physical volume and use parity checking to protect against disk failure.

SCSI Small Computer System Interface. An ANSI standard for a logical interface to computer peripherals and for a computer peripheral interface. The interface utilizes a SCSI logical protocol over an I/O interface that configures attached targets and initiators in a multi-drop bus topology.

Serial Storage Architecture (SSA). An expanded storage adapter for multi-processor data sharing in UNIX-based computing, allowing disk connection in a high-speed loop.

Simple Network Management Protocol (SNMP). A protocol designed to give a user the capability to remotely manage a computer network by polling and setting terminal values and monitoring network events.

Simple Network Management Protocol (SNMP). A protocol governing network management and the monitoring of network devices and their functions.

Small Computer System Interface (SCSI). An adapter supporting the attachment of various direct-access storage devices.

Software zoning. Is implemented within the Simple Name Server (SNS) running inside the fabric switch. When using software zoning, the members of the zone can be defined with: node WWN, port WWN, or physical port number. Usually the zoning software also allows you to create symbolic names for the zone members and for the zones themselves.

Storage area network (SAN). A managed, high-speed network that enables any-to-any interconnection of heterogeneous servers and storage systems.

Subsystem masking. The support provided by intelligent disk storage subsystems like the Enterprise Storage Server. See also LUN masking and port zoning.

Switch. A component with multiple entry and exit points or ports that provide dynamic connection between any two of these points.

Switch topology. A switch allows multiple concurrent connections between nodes. There can be two types of switches, circuit switches and frame switches. Circuit switches establish a dedicated connection between two nodes. Frame switches route frames between nodes and establish the connection only when needed. A switch can handle all protocols.

Topology. An interconnection scheme that allows multiple Fibre Channel ports to communicate. For example, point-to-point, arbitrated loop, and switched fabric are all Fibre Channel topologies.

Transmission Control Protocol (TCP). A reliable, full duplex, connection-oriented, end-to-end transport protocol running on of IP.

Trivial File Transfer Protocol (TFTP). A set of conventions for transferring files between hosts using minimal protocol.

Zoning. In Fibre Channel environments, zoning allows for finer segmentation of the switched fabric. Zoning can be used to instigate a barrier between different environments. Ports that are members of a zone can communicate with each other but are isolated from ports in other zones. Zoning can be implemented in two ways: hardware zoning and software zoning.

Abbreviations and acronyms

ABI	Application Binary Interface
ACSLS	Automated Cartridge System Library Software
ADSM	Adstar Distributed Storage Manager
ADSM	ADSTAR Distributed Storage Manager
AFS	Andrew File System
AIT	Advanced Intelligent Tape
AIX	Advanced Interactive Executive
AL	Arbitrated Loop
ANSI	American National Standards Institute
API	Application Programming Interface
ATL	Automated Tape Library
B/A	Backup/Archive
BOOTP	Bootstrap Protocol
BSD	Berkeley Software Distribution
CD-ROM	Compact Disc Read-Only Memory
CIFS	Common Internet File System
CLI	Command Line Interface
CPU	Central Processing Unit
CRC	Cyclic Redundancy Check
CRC	Cyclic Redundancy Check
DASD	Direct Access Storage Device
DAT	Digital Audio Tape
DB	Database
DFS	Distributed File System
DHCP	Dynamic Host Configuration Protocol
DLT	Digital Linear Tape
DLT	Digital Linear Tape
DNS	Domain Name Services
DRM	Disaster Recovery Manager
DVD	Digital Versatile Disc
E_Port	Expansion Port
ESCON	Enterprise System Connection
F_Node	Fabric Node
F_Port	Fabric Port
FAT	File Access Table
FC	Fibre Channel
FCA	Fibre Channel Association
FC-AL	Fibre Channel Arbitrated Loop
FCLC	Fibre Channel Loop Association
FCP	Fibre Channel Protocol
FC-PLDA	Fibre Channel Private Loop Direct Attach
FC-SW	Fibre Channel Switch Fabric
FICON	Fiber Connection
FL_Port	Fabric Loop Port
FLOGI	Fabric Login
FS	File System
FSID	File Space Identifier
FSP	Fibre Channel Service Protocol
FTP	File Transfer Protocol
G_Port	Generic Port
Gb	Gigabit
GB	Gigabyte
GBIC	GigaBit Interface Converter
GLM	Gigabit Link Module
GUI	Graphical User Interface
HACMP	High Availability Cluster Multi Processing
HACMP/ES	High Availability Cluster Multi Processing/Enhanced Scalability
HBA	Host Bus Adapter
HFS	Extended file system
HSM	Hierarchical Storage Manager
HTTP	Hypertext Transport Protocol
I/O	Input/Output
IBM	International Business Machines Corporation
IFS	Integrated File System
ILE	Integrated Language Environment
IP	Internet protocols
ISO	International Standards Organization
ITSO	International Technical Support Organization
JFS	Journaled File System
JRE	Java Runtime Environment
Kb	Kilobit
KB	Kilobyte
L_Port	Loop Port

LAN	Local Area Network
LDAP	Lightweight Directory Access Protocol
LE	Logical Extents
Linux	A Unix like operating system
LP	Logical Partition
LTO	Linear Tape-Open
LUN	Logical Unit Number
LV	Logical Volume
LVD	Low Voltage Differential
LVM	Logical Volume Manager
LVSA	Logical Volume Snapshot Agent
Mb	Megabit
MB	Megabyte
MIA	Media Interface Adapter
MMF	Multi-Mode Fiber
N_Port	Node Port
NAS	Network-Attached Storage
NDMP	Network Data Management Protocol
NFS	Network File System
NKE	Network Kernel Extensions
NMS	Network Management System
NSS	Novell Storage Services
NT	NoT a Unix like operating system
NTFS	NT file system
NTP	Network Time Protocol
PAP	Printer Access Protocol
PASE	Portable Application Solutions Environment
PE	Physical Extents
PLDA	Private Loop Direct Attached
PP	Physical Partition
PPP	Point to Point Protocol
PTF	Program Temporary Fix
PV	Physical Volume
QFS	Quick File System
QIC	Quarter Inch Cartridge
RAID	Redundant Array of Inexpensive or Independent Disks
RSS	Remote Storage Server
SAN	Storage Area Network
SCSI	Small Computer Interface Standard
SDG	SAN Data Gateway
SLIC	System Licensed Internal Code
SLP	Service Location Protocol
SMC	SCSI Medium Changer
SMF	Single-Mode Fiber
SNIA	Storage Networking Industry Association
SNMP	Simple Network Management Protocol
SP	Super Parallel
SQL	Structured Query Language
SSA	Serial Storage Architecture
STP	Shielded Twisted Pair
T_Port	Trunk Port
Tb	Terrabit
TB	Terabyte
TCP	Transmission Control Protocol
TDP	Tivoli Data Protection
TIMI	Technology Independent Machine Interface
TSM	Tivoli Storage Manager
TXN	Transaction
UDP	User Datagram Protocol
UFS	Unix File System
Unix	Good operating system
UTF	Unicode Transformation Format
UTP	Unshielded Twisted Pair
VFS	Virtual File System
VMM	Virtual Memory Manager
VPN	Virtual Private Network
VPS	Virtual Private SAN
WAN	Wide Area Network
WORM	Write-Once Read-Many
WWN	World Wide Name
WWNN	World Wide Node Name
WWPN	World Wide Port Name

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

IBM Redbooks

For information on ordering these publications, see "How to get IBM Redbooks" on page 341.

- *A Practical Guide to Tivoli SANergy,* SG24-6146
- *Using Tivoli Storage Manager in a SAN Environment,* SG24-6132

Other resources

These publications are also relevant as further information sources:

- *IBM Tivoli Storage Manager for AIX Administrator's Reference V5.1,* GC32-0769
- *IBM Tivoli Storage Manager for Windows Administrator's Guide V5.1*, GC32-0782
- *IBM Tivoli Storage Manager for Windows Administrator's Reference V5.1*, GC32-0783
- *IBM Tivoli Storage Manager for Windows Quick Start V5.1*, GC32-0784
- *IBM Tivoli Storage Manager Using the Application Program Interface V5.1*, GC32-0793
- *IBM Tivoli Storage Manager for Sun Solaris Managed Systems for SAN Storage Agent User's Guide V5.1*, GC32-0781
- *IBM Storage Area Network Data Gateway Installation and User's Guide,* SC26-7304
- *IBM Tivoli Storage Manager for UNIX Backup-Archive Clients Installation and User's Guide V5.1*, GC32-0789
- *IBM Tivoli Space Manager for UNIX Using Hierarchial Storage Management Clients*, GC32-0794

Referenced Web sites

These Web sites are also relevant as further information sources:

- Tivoli Storage Manager Web site

  http://www.tivoli.com/tsm
- Tivoli SANergy Web site

  http://www.tivoli.com/products/index/sanergy
- Tivoli Product License Information

  http://www.tivoli.com/products/licensing
- IBM Tivoli Storage Manager Enterprise Edition

  http://www.tivoli.com/products/index/storage_mgr/storage_mgr_concepts.html
- Tivoli Customer Support

  http://www.tivoli.com/support

- Tivoli Storage and SAN Product Family Technical Support
  http://www.tivoli.com/support/storage_mgr/tivolimain.html
- Tivoli Storage Manager Products Technical Support
  http://www.tivoli.com/support/storage_mgr/requirements.html
- Sun Microsystems QFS Web site
  http://www.sun.com/storage/software/storage_mgmt/qfs/index.html
- Sunsolve Online
  http://sunsolve.sun.com
- Java 2 Platform Standard Edition V1.3.1
  http://java.sun.com/j2se/1.3/
- IBM Content Manager CommonStore Web site
  http://www-4.ibm.com/software/data/commonstore
- IBM.com ftp Software Server
  ftp://ftp.software.ibm.com/storage/devdrvr
- Performance Management Guide
  http://publibn.boulder.ibm.com/doc_link/en_US/a_doc_lib/aixbman/prftungd/prftungd.htm
- iSeries Information Center
  http://publib.boulder.ibm.com/pubs/html/as400
- Integrated Application Servers
  http://www.as400.ibm.com
- White Papers for AS/400 and iSeries
  http://www.ibm.com/servers/eserver/iseries/whpapr/ifs.html
- IBM Storage Media Product Selector
  http://www.storage.ibm.com/media/products.html
- Tape and Optical Storage
  http://www.storage.ibm.com/hardsoft/tape/index.html
- Kernel Extensions and Device Support Programming Concepts
  http://publibn.boulder.ibm.com/doc_link/en_US/a_doc_lib/aixprggd/kernextc/ls_devconfig_subr.htm
- IBM HP-UX Tape and Medium Changer Device Driver (ATDD) — Readme file
  ftp://ftp.software.ibm.com/storage/devdrvr/HPUX/README
- IBM Developer Kit for AIX, Java Technology Edition
  http://www-106.ibm.com/developerworks/java/jdk/aix/index.html
- ACM Portal
  http://portal.acm.org
- QLogic Technical Support
  http://www.qlogic.com/support/home_support.asp
- SDK and RTE 1.3 for hp-ux 11.0 and 11i PA-RISC downloads
  http://www.hp.com/products1/unix/java/java2/sdkrte1_3/downloads/index.html
- Java Runtime Plugin, v 1.3.1 for SGI IRIX® — MR Release
  http://www.sgi.com/developers/devtools/languages/javaplugin131.html

- International Organization for Standardization

  http://www.iso.ch

- American National Standards Organization

  http://www.ansi.org

How to get IBM Redbooks

You can order hardcopy Redbooks, as well as view, download, or search for Redbooks at the following Web site:

ibm.com/redbooks

You can also download additional materials (code samples or diskette/CD-ROM images) from that site.

IBM Redbooks collections

Redbooks are also available on CD-ROMs. Click the CD-ROMs button on the Redbooks Web site for information about all the CD-ROMs offered, as well as updates and formats.

Index

Numerics

A

B

C

D

E

F

G

H

I

J

Q

R

S

T

U

IBM Tivoli Storage Manager Version 5.1: Technical Guide

(0.5” spine)
0.475”<->0.873”
250 <-> 459 pages

IBM Tivoli Storage Manager Version 5.1 Technical Guide

Redbooks

Server-free backup and restore for Windows 2000

Online image backup for Linux and Windows 2000

Performance enhancements and HACMP support

IBM Tivoli Storage Manager, the backbone product of the Tivoli Storage Management product set, is an enterprise-wide network storage management solution. This IBM Redbook presents an overview of IBM Tivoli Storage Manager Version 5.1 and provides updates on the Tivoli Storage Management product set. It gives a detailed description of each of the new functions of Tivoli Storage Manager.

IBM Tivoli Storage Manager Version 5.1 was introduced in March 2002 and is the successor to Tivoli Version 4.2 and 3.7 and Tivoli ADSM Version 3.1. It provides many exciting new functions such as server-free backup and online image backup for Windows 2000.

This book is essential reading for customers, consultants, IBM Business Partners, IBM employees, and Tivoli staff who are familiar with Tivoli Storage Manager Version 3.7 and 4.2 and who want to understand the new functionality of Version 5.1.

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

SG24-6554-00 ISBN 0738425192