Tivoli

Tivoli Storage Manager Version 3.7.3 & 4.1: Technical Guide

Exploiting SAN environments for LAN-free backup of application data

Complete data protection solution for mobile systems

Windows 2000 exploitation and support

Roland Leins
German Agudelo
Jawad Butt
Massimo Canella
Angela Hopkins

Redbooks

ibm.com/redbooks

International Technical Support Organization

SG24-6110-00

Tivoli Storage Manager Version 3.7.3 & 4.1: Technical Guide

September 2000

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix B, “Special notices” on page 331.

First Edition (September 2000)

This edition applies to Version 3, Release 7 and Version 4, Release 1 of Tivoli Storage Manager, Program Numbers 5697-TSM and 5698-TSM for use with the Microsoft Windows NT and Windows 2000, IBM AIX, Sun Solaris, HP-UX, and IBM OS/400, to Version 2, Release 1 of Tivoli SANergy File Sharing, Program Number 5698-SFS for use with Microsoft Windows NT and Windows 2000 and Sun Solaris, and to Version 2, Release 1 of Tivoli Decision Support, Program Number 5698-TDS for use with MIcrosoft Windows NT.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. 471F 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

Preface

Tivoli Storage Manager, the backbone product of the Tivoli Storage Management product set, is an enterprise-wide network storage management solution. Tivoli Storage Manager Version 4.1 was introduced in July 2000 and is the successor of Tivoli ADSM Version 3.1 and Tivoli Storage Manager Version 3.7. It provides many new functions such as support for mobile systems backup, LAN-free backup and restore of data using a Storage Area Network (SAN) environment, and Windows 2000 exploitation and support.

This IBM Redbook presents an overview of Tivoli Storage Manager Version 3.7.3 and Version 4.1. The book provides updates on the Tivoli Storage Management product set, and it gives a detailed description of each of the new functions of Tivoli Storage Manager. The book also discusses in detail the Tivoli Storage Manager Windows 2000 exploitation and support, and introduces Tivoli SANergy File Sharing and Tivoli Decision Support for Storage Management Analysis, two other members of the Tivoli Storage Management product set.

This book is intended for customers, consultants, IBM Business Partners, IBM employees, and Tivoli staff who are familiar with ADSM Version 3.1 and Tivoli Storage Manager Version 3.7, and who need to understand what is new in Tivoli Storage Manager Version 3.7.3 and 4.1.

It particularly applies to the following versions: Version 3, Release 7 and Version 4, Release 1 of Tivoli Storage Manager, Program Numbers 5697-TSM and 5698-TSM for use with the Microsoft Windows NT and Windows 2000, IBM AIX, Sun Solaris, HP-UX, and IBM OS/400, to Version 2, Release 1 of Tivoli SANergy File Sharing, Program Number 5698-SFS for use with Microsoft Windows NT and Windows 2000 and Sun Solaris, and to Version 2, Release 1 of Tivoli Decision Support, Program Number 5698-TDS for use with MIcrosoft Windows NT.

The team that wrote this IBM Redbook

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

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

German Agudelo is an IT Specialist in the ATSC in San Jose. He has 22 years of experience in marketing and sales. He holds a degree in Computer Science from Queens College in New York. His areas of expertise include storage sub-systems hardware and software, particular TSM server and clients.

Jawad Butt is a Solutions Architect at GE Capital IT Solutions, an IBM Business Partner in Canada. He has 6 years of experience in designing and implementing UNIX solutions, primarily in the IBM RS/6000 and SP systems arena. Within GE, his responsibilities include consulting and project management in the area of multi-vendor storage solutions, including the frenetic world of SANs. Jawad holds a degree in Electrical Engineering from the University of Engineering and Technology, Lahore, Pakistan.

Massimo Canella is an Advisory IT Specialist from IBM Italy. He has 25 years of IBM experience and 4 years in ADSM/TSM proposal and implementation services on UNIX, NT, and AS/400 platforms. His areas of expertise include VM, VSE, OS/2, Windows NT, AIX, and Networking.

Angela Hopkins is a data storage specialist and works for Storm in the UK. She has 6 years of experience in the field of distributed data and storage management, working with IBM ADSM and Tivoli Storage Manager. Her areas of expertise include Tivoli Storage Manager implementation and consultancy services on most of the supported platforms as well as Tivoli Storage Manager education.

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

Charlotte Brooks
International Technical Support Organization, San Jose Center

Karen Dutch
Tivoli Systems, San Jose

Tricia Jiang
Tivoli Systems, San Jose

Richard Harrison
Tivoli Systems, Austin

Jim Smith
Tivoli Systems, San Jose

Mike Dile
Tivoli Systems, San Jose

Rob Edwards
Tivoli Systems, San Jose

Avishai Hochberg
Tivoli Systems, San Jose

Jon Viksne
Tivoli Systems, San Jose

Colin Dawson
Tivoli Systems, Tucson

Barry Fruchtman
Tivoli Systems, Tucson

Mike Kaczmarski
Tivoli Systems, Tucson

William Scheid
Tivoli Systems, Tucson

Glen Hattrup
Tivoli Systems, Tucson

Comments welcome

Your comments are important to us!

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

- Fax the evaluation form found in “IBM Redbooks review” on page 347 to the fax number shown on the form.
- Use the online evaluation form found at **ibm.com**/redbooks
- Send your comments in an Internet note to redbook@us.ibm.com

Chapter 1. Tivoli Storage Manager Version 4.1 overview

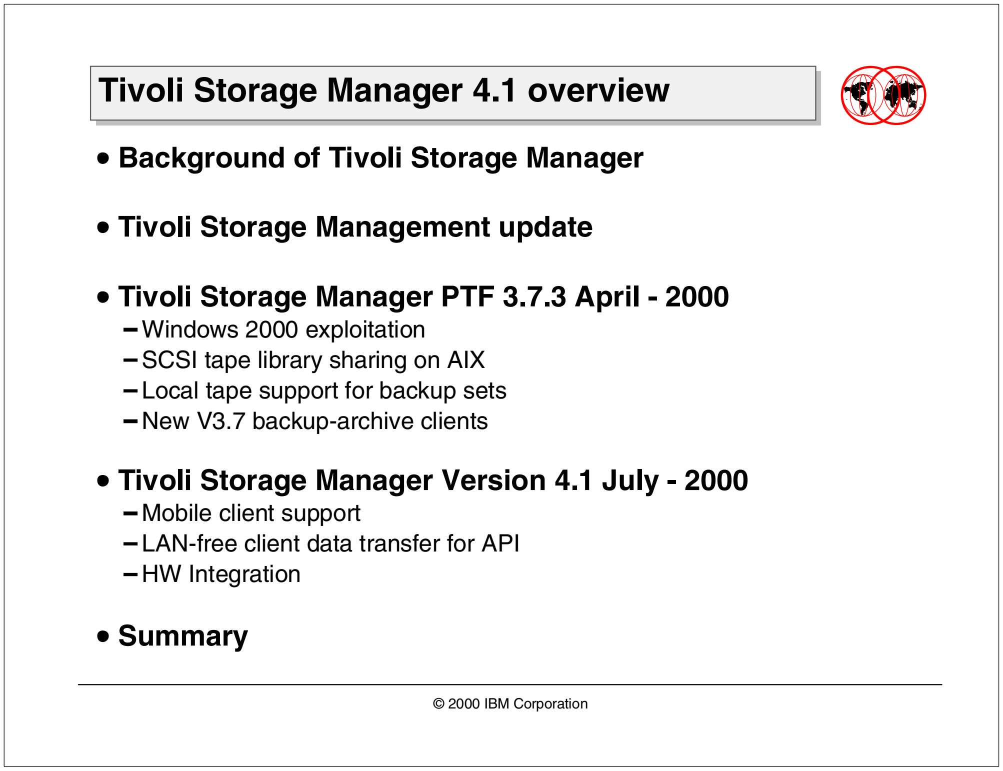

Tivoli Storage Manager 4.1 overview

- **Background of Tivoli Storage Manager**
- **Tivoli Storage Management update**
- **Tivoli Storage Manager PTF 3.7.3 April - 2000**
  - Windows 2000 exploitation
  - SCSI tape library sharing on AIX
  - Local tape support for backup sets
  - New V3.7 backup-archive clients
- **Tivoli Storage Manager Version 4.1 July - 2000**
  - Mobile client support
  - LAN-free client data transfer for API
  - HW Integration
- **Summary**

This chapter provides an overview of Tivoli Storage Manager Version 3.7.3 and 4.1. Topics covered are:

- Background information explaining the strategy of Tivoli Storage Manager
- Update of the Tivoli Storage Management product set
- Tivoli Storage Manager Version 3.7.3 and 4.1:
  - Windows 2000 support and exploitation
  - New features and functions, including the exploitation of SAN technology
  - New features to support the special backup and recovery requirements of mobile systems
  - Special features to support backup and recovery of data stored on advanced disk subsystems
  - Client and server usability enhancements
- Tivoli Storage Manager Version 3.7 summary

1.1 Background

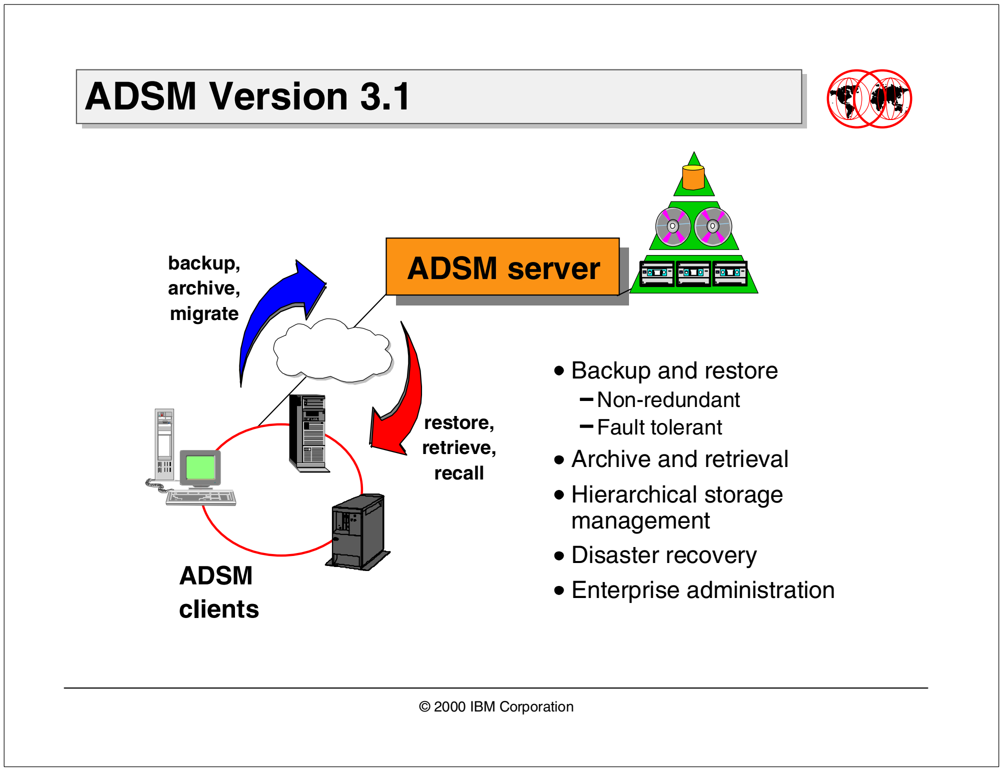

ADSM Version 3.1

- Backup and restore
  - Non-redundant
  - Fault tolerant
- Archive and retrieval
- Hierarchical storage management
- Disaster recovery
- Enterprise administration

© 2000 IBM Corporation

1.1.1 ADSM Version 3.1

ADSTAR Distributed Storage Manager (ADSM) Version 3.1 was made available in 1997 and is a solution for distributed data management in an enterprise network environment. ADSM supports a wide variety of platforms for both small and large systems, and delivers many data management functions, for example, data backup, data archival, space management, and disaster recovery — all integrated in one software solution. With its strong enterprise administration features, it is well suited as a total solution for complex enterprise environments.

ADSM was rebranded to Tivoli ADSM in February 1999 and is now the base for Tivoli Storage Manager Version 3.7. It is the kernel of the Tivoli Storage Management product set.

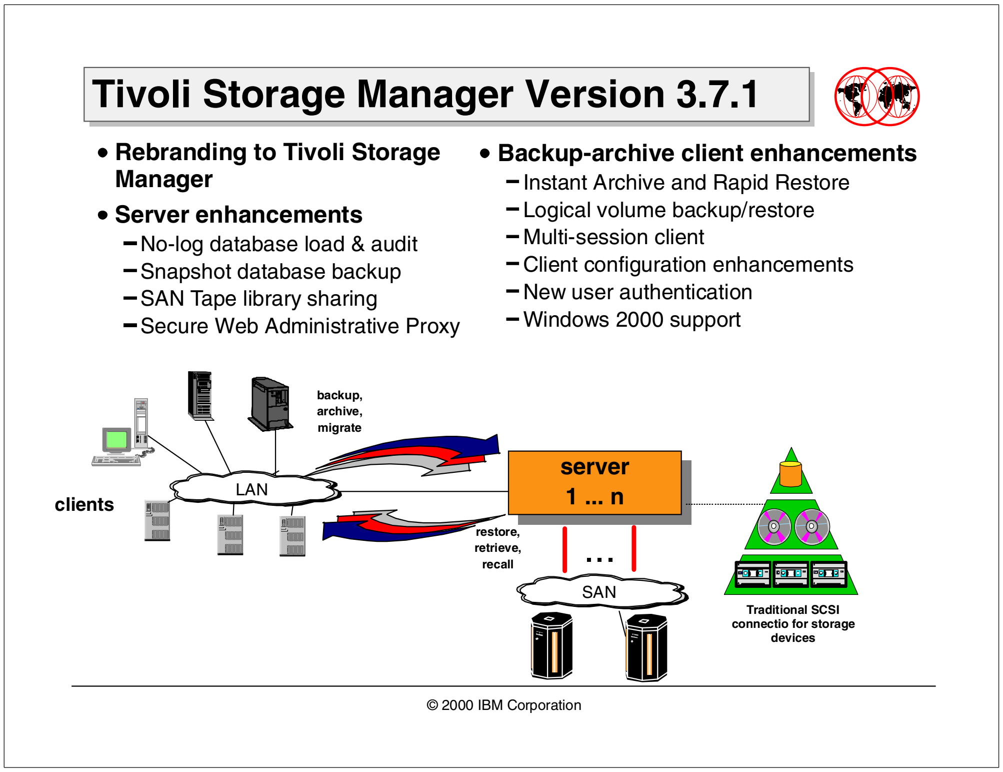

Tivoli Storage Manager Version 3.7.1

- Rebranding to Tivoli Storage Manager
- Server enhancements
  - No-log database load & audit
  - Snapshot database backup
  - SAN Tape library sharing
  - Secure Web Administrative Proxy
- Backup-archive client enhancements
  - Instant Archive and Rapid Restore
  - Logical volume backup/restore
  - Multi-session client
  - Client configuration enhancements
  - New user authentication
  - Windows 2000 support

© 2000 IBM Corporation

1.1.2 Tivoli Storage Manager Version 3.7.1

The first version of Tivoli Storage Manager, Version 3.7.1, and the successor of ADSM V3.1, was announced in September 1999. Besides the rebranding of the entire product with Version 3.7.1, several enhancements for both server and client, as well as complementary products, have introduced a new age for backup restore solutions.

Feature improvements for server and client respond to the necessity of speeding up backup and restore operations. Tivoli Storage Manager started to exploit the new SAN technology with SAN library sharing, thus giving to multiple servers the ability to share an automated library in a high performance Storage Area Network (SAN) configuration.

Besides numerous usability enhancements both on client and server, the first client features to support the new Microsoft operating system Windows 2000 have been introduced.

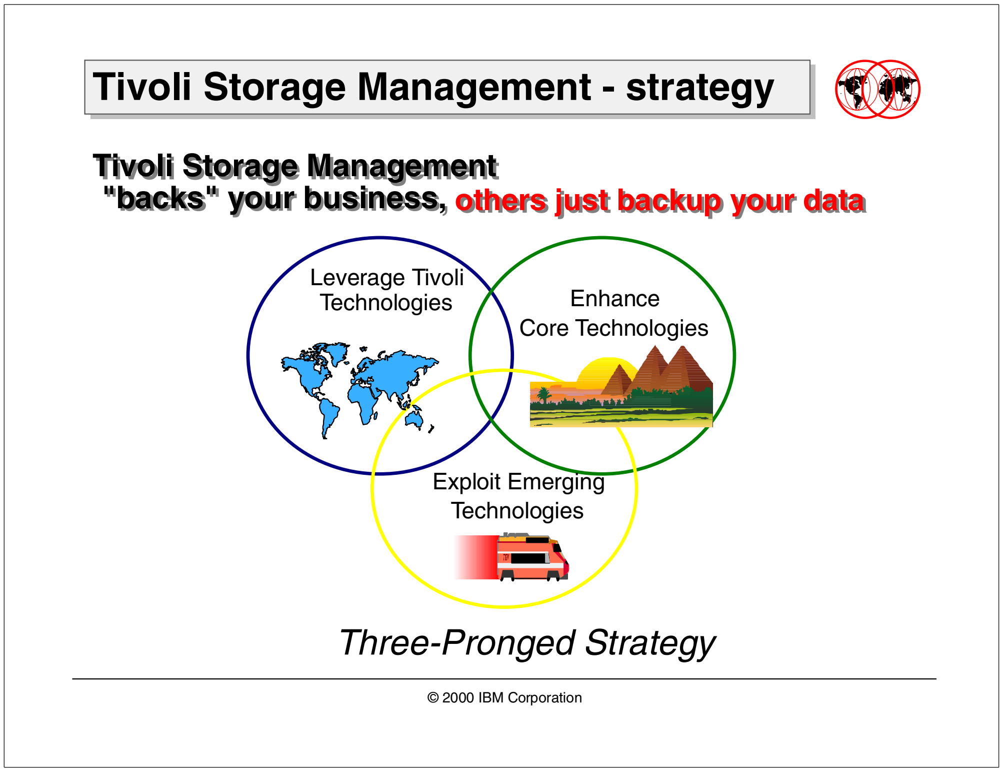

Tivoli Storage Management - strategy

Tivoli Storage Management
"backs" your business, others just backup your data

Leverage Tivoli Technologies

Enhance Core Technologies

Exploit Emerging Technologies

Three-Pronged Strategy

1.1.3 Tivoli Storage Management strategy

The Tivoli Storage Management strategy focuses on the delivery of a complete business solution, rather than trying to solve only the data backup and recovery problems. It is based on the strengths of ADSM Version 3.1, and addresses the emerging trends in the network environment by:

- Leveraging Tivoli technologies to integrate storage management seamlessly into the overall business management concept
- Enhancing core technologies to satisfy the needs of complex environments and to address customer requirements
- Exploiting new emerging technologies like the SAN technology

1.2 Tivoli Storage Management product set update

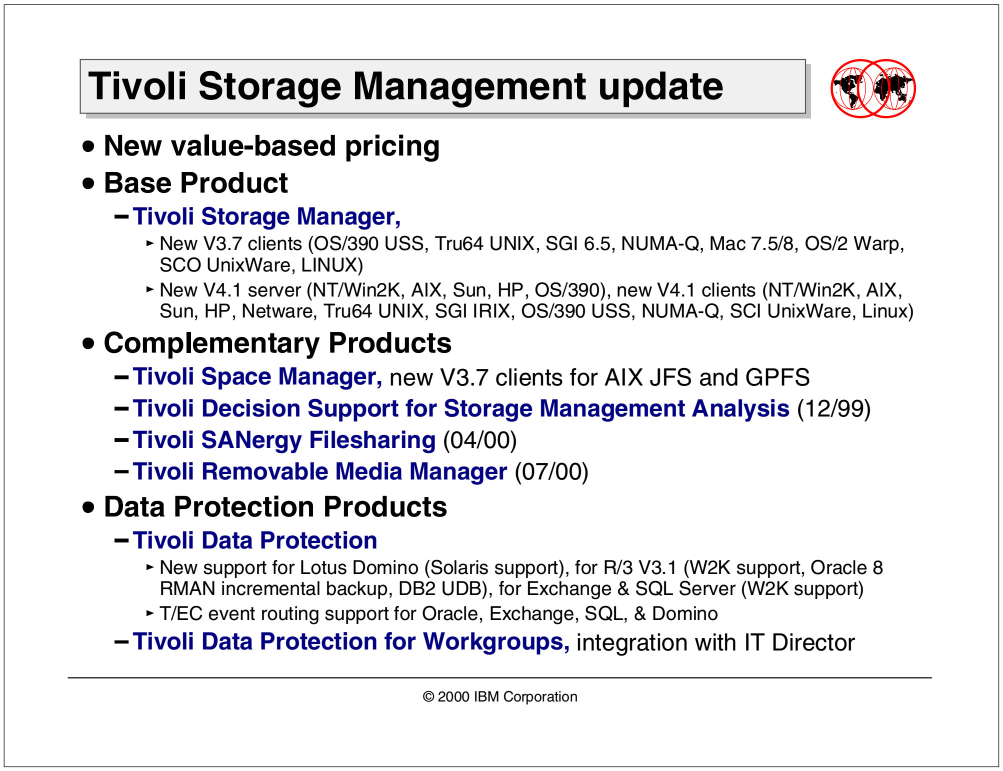

Tivoli Storage Management update

- **New value-based pricing**
- **Base Product**
  - **Tivoli Storage Manager,**
    - New V3.7 clients (OS/390 USS, Tru64 UNIX, SGI 6.5, NUMA-Q, Mac 7.5/8, OS/2 Warp, SCO UnixWare, LINUX)
    - New V4.1 server (NT/Win2K, AIX, Sun, HP, OS/390), new V4.1 clients (NT/Win2K, AIX, Sun, HP, Netware, Tru64 UNIX, SGI IRIX, OS/390 USS, NUMA-Q, SCI UnixWare, Linux)
- **Complementary Products**
  - **Tivoli Space Manager,** new V3.7 clients for AIX JFS and GPFS
  - **Tivoli Decision Support for Storage Management Analysis** (12/99)
  - **Tivoli SANergy Filesharing** (04/00)
  - **Tivoli Removable Media Manager** (07/00)
- **Data Protection Products**
  - **Tivoli Data Protection**
    - New support for Lotus Domino (Solaris support), for R/3 V3.1 (W2K support, Oracle 8 RMAN incremental backup, DB2 UDB), for Exchange & SQL Server (W2K support)
    - T/EC event routing support for Oracle, Exchange, SQL, & Domino
  - **Tivoli Data Protection for Workgroups,** integration with IT Director

The *Tivoli Storage Management product set* covers the various aspects of storage management within an enterprise environment. In general, the product set can be structured in three parts: The *Tivoli Storage Manager*, which provides the storage management backbone; the *complementary products*, which add special storage management functions (like disaster recovery management or hierarchical storage management) in conjunction with Tivoli Storage Manager; and the *Tivoli Data Protection group of products*, which integrate application data management into the Tivoli Storage Manager storage management environment.

Besides the introduction of the new value-based pricing model for all Tivoli Storage Management products, the following product updates have occurred since the announcement of Tivoli Storage Manager Version 3.7.1. in September 1999:

Tivoli Storage Manager:

- New V3.7 clients (OS/390 USS, Tru64 UNIX, SGI 6.5, NUMA-Q, Mac 7.5/8, OS/2 Warp, SCO UnixWare, Linux) in April 2000
- New V4.1 server (NT/Win2K, AIX, Sun, HP, OS/390), new V4.1 clients (NT/Win2K, AIX, Sun, HP, Netware, Tru64 UNIX, SGI IRIX, OS/390 USS, NUMA-Q, SCI UnixWare, Linux) in July 2000

Complementary products:

- New Tivoli Space Manager V3.7 client provides support for the AIX JFS and AIX GPFS file system
- Tivoli Decision Support for Storage Management Analysis (TDSfSMA) provides a business-relevant reporting and decision support solution for Tivoli Storage Manager covering event and performance analysis.
- Tivoli SANergy File Sharing provides a true data sharing solution for Storage Area Network environments and allows Tivoli Storage Manager to implement LAN-free and server-free backup solutions of data using Storage Area Networks.
- Tivoli Removable Media Manager (TRMM) is a new tape library and media management program which allows multiple applications running in a distributed environment to share tape library and media resources.

Tivoli Data Protection for Application Clients:

- This function is used to integrate backup and recovery of specific application data, such as Oracle, Lotus Domino, and others, into the Tivoli Storage Manager solution.
- A new feature provided is TDP for R/3, the successor of the BackInt product.
- There is also new platform support available for TDP for Lotus Domino, MS Exchange, and SQL, and for enhanced T/EC event routing capabilities.
- There are new message codes for several TDP products to identify them as event sources.

Tivoli Data Protection for Workgroups:

- This function, which is a bare-metal disaster recovery solution for Microsoft NT, has been integrated into Tivoli IT DIrector.

1.3 Tivoli Storage Manager 3.7.3 and 4.1

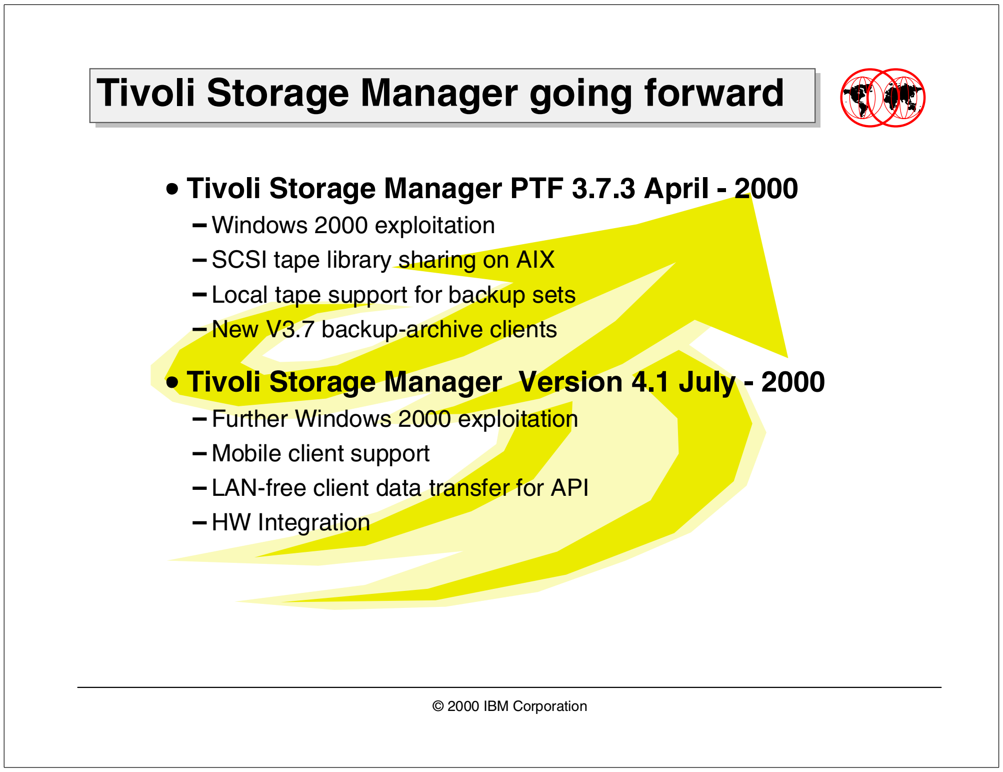

Tivoli Storage Manager going forward

- **Tivoli Storage Manager PTF 3.7.3 April - 2000**
  - Windows 2000 exploitation
  - SCSI tape library sharing on AIX
  - Local tape support for backup sets
  - New V3.7 backup-archive clients
- **Tivoli Storage Manager Version 4.1 July - 2000**
  - Further Windows 2000 exploitation
  - Mobile client support
  - LAN-free client data transfer for API
  - HW Integration

Since the announcement of Tivoli Storage Manager in September 1999, there have been two major new versions of Tivoli Storage Manager introduced, Version 3.7.3 in April 2000 and Version 4.1 in July 2000. The following section gives an introduction to the major new functions and features of both versions, which are:

- Windows 2000 exploitation and support of client and server
- New AIX and Windows 2000 platform support tape library sharing in a SAN environment for SCSI connected libraries
- Client tape support for backup set restore on the local client system without interaction with the Tivoli Storage manager server
- New V3.7 backup-archive clients, including the newly supported Linux platform
- Backup-archive client and server enhancements supporting the special backup requirements of mobile systems
- LAN-free client data transfer for the API client in a SAN environment
- Tivoli Storage Manager hardware integration to implement server-free backup solutions for enhanced disk subsystems like IBM ESS and EMC Symmetrix

.

Windows 2000 exploitation

- **Support and exploitation of Windows 2000 architectural enhancements**

- **Windows 2000 Server Exploitation**
  - Microsoft Cluster Server integration
  - Active Directory support
  - Performance enhancements
  - Multiple Server instances
  - Microsoft Management Console
  - Removable Storage Manager
  - Microsoft System Installer

- **Windows 2000 Client Support**
  - Back up Windows 2000 System State items
  - Back up Windows 2000 File System items

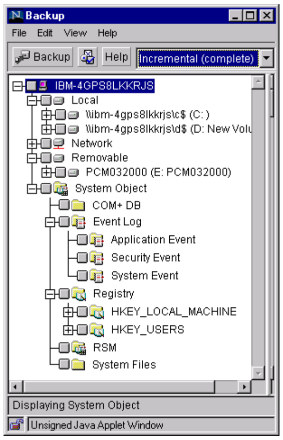

© 2000 IBM Corporation

1.3.1 Windows 2000 exploitation

Microsoft Windows 2000, the next major release of the Windows NT operating system, introduces a number of architectural enhancements that require special attention from system administrators responsible for protecting data. Tivoli Storage Manager V3.7 and 4.1 exploit the new features of Windows 2000 and provide comprehensive data protection with ease of management.

Tivoli Storage Manager server exploits many of the new architectural enhancements introduced with the Windows 2000 Server Family of Products. Some new features have also been introduced for the Tivoli Storage Manager server to take advantage of the Windows 2000 server environments. The Tivoli Storage Manager server introduces support enhancements for:

- Microsoft Cluster Server
- Active Directory
- Terminal Services
- Removable Storage Manager
- Microsoft Management Console
- Microsoft Windows Installer

In addition, new server features include:

- Performance tuning wizard
- Support of multiple server instances

The Tivoli Storage Manager client has the ability to back up and process some new Windows 2000 System Objects including:

- Active Directory
- MSCS cluster databases
- Removable Storage Manager database
- COM+ Class Registration database
- System Volume
- Certificate Server DB
- System File Protection

Tivoli Storage Manager client also exploits the new Windows 2000 File System items including:

- Disk Quotas
- Object (file) Encryption
- Reparse Points utilized by Volume Mount Points
- Multiple Data Streams
- Change Journal
- Sparse File Support
- Distributed Link-Tracking
- File System Replication
- Distributed File Systems (DFS)

.

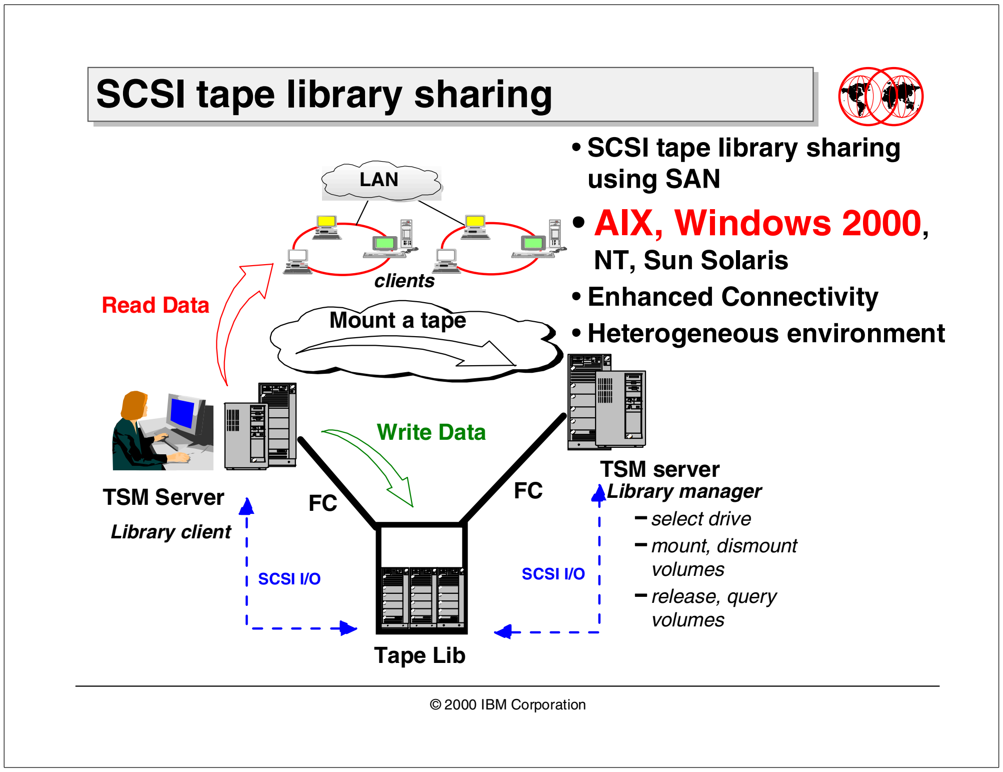

SCSI tape library sharing

- SCSI tape library sharing using SAN
- AIX, Windows 2000, NT, Sun Solaris
- Enhanced Connectivity
- Heterogeneous environment

1.3.2 SCSI tape library sharing

With Tivoli Storage Manager server 3.7.1 and above, multiple Tivoli Storage Manager servers can dynamically share library volume and tape drive resources of a physical tape library connected to a SAN using the Tivoli Storage Manager server SCSI tape library sharing feature. The hosts can thus maintain high speed connections to the same devices through the SAN fabric. Applications which immediately benefit from this include backup and restore. The effect is pronounced for environments with large amounts of data to back up over shrinking windows and constrained LAN bandwidth.

AIX and Windows 2000 have been added as a new Tivoli Storage Manager server platform that supports the SCSI tape library sharing solution. SCSI tape library sharing for SCSI libraries is now supported across the following platforms:

- IBM AIX
- Windows 2000
- Windows NT
- Sun Solaris

It is also supported in heterogeneous server environments.

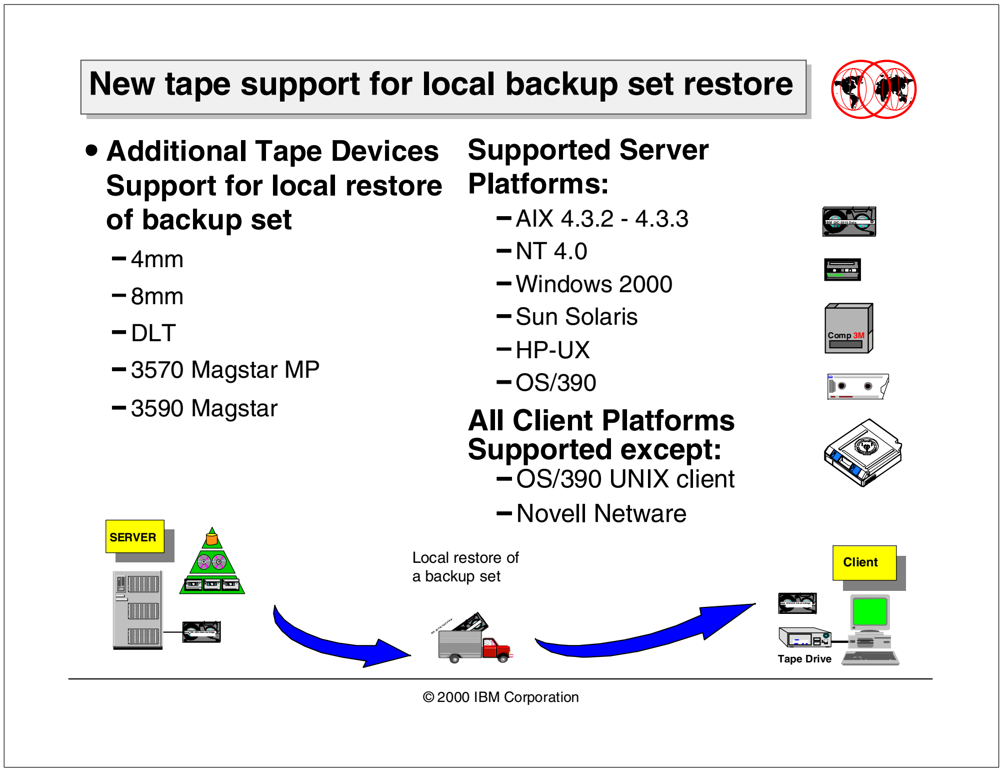

New tape support for local backup set restore

- **Additional Tape Devices Support for local restore of backup set**
  - 4mm
  - 8mm
  - DLT
  - 3570 Magstar MP
  - 3590 Magstar

Supported Server Platforms:
- AIX 4.3.2 - 4.3.3
- NT 4.0
- Windows 2000
- Sun Solaris
- HP-UX
- OS/390

All Client Platforms Supported except:
- OS/390 UNIX client
- Novell Netware

© 2000 IBM Corporation

1.3.3 New tape support for local backup set restore

The local restore of a backup set, also called *Rapid Recovery*, gives a client the ability to restore all or parts of filespaces previously backed up to TSM server without transferring the data over the LAN. Backup sets generated on tape media provide an efficient solution to move large amounts of data.

In Tivoli Storage Manager 3.7.1, only the AIX client platform using an 8mm tape device was supported to restore a server-generated backup set on tape media. New client code, starting with Version 3.7.2 for all platform clients except OS/390, extends the ability to restore backup sets generated onto tape units common to the TSM server. In addition, more media devices are now supported.

These devices are: 8mm, 4mm, DLT, 3570, and 3590. The same device must be available on both the server and client machine to ensure that a backup set generated on the server can be restored on the client machine. The client machine will use the native driver for the tape drive.

Linux client support

- **Interfaces and supported distributions**
  - GUI, Web and command line
  - SuSE 6.3, RedHat 6.1, Caldera 2.3, TurboLinux 6.0
- **Installation and setup**
  - *Installing the Clients* V3R7.2 and README file coming with the client code
  - Executable in `/opt/tivoli/tsm/`
- **Features and limitations**
  - Client configuration wizard
  - Backup set restore from FILE and Tape
  - Backup set restore only from FILE using Web Client
  - Image backup not supported

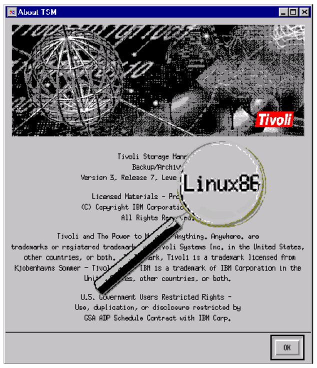

Tivoli Storage Mana
Backup/Archiv
Version 3, Release 7, Leve

Licensed Materials - Pro
(C) Copyright IBM Corporatio
All Rights Re

Tivoli and The Power to M Anything. Anywhere. are
trademarks or registered trademark ivoli Systems Inc. in the United States,
other countries, or both. ark, Tivoli is a trademark licensed from
Kjobenhavns Sommer - Tivol IBM is a trademark of IBM Corporation in the
Unit es, other countries, or both.

U.S. Government Users Restricted Rights -
Use, duplication, or disclosure restricted by
GSA ADP Schedule Contract with IBM Corp.

1.3.4 Linux client support

Tivoli Storage Manager Version 3.7.2 introduces a fully functional and fully supported backup-archive client for the Linux operating system. The Linux client is supported for the following distributions:

- SuSE 6.3
- RedHat 6.1
- Caldera 2.3
- TurboLinux 6.0

All interfaces are available:

- Backup-archive client (command-line and GUI)
- Administrative client (command line)
- Web backup-archive client

API and API multi-thread are also available for applications that will use them, for example, the Web client.

The only communication method supported is TCP/IP. Any supported network adapter can be used.

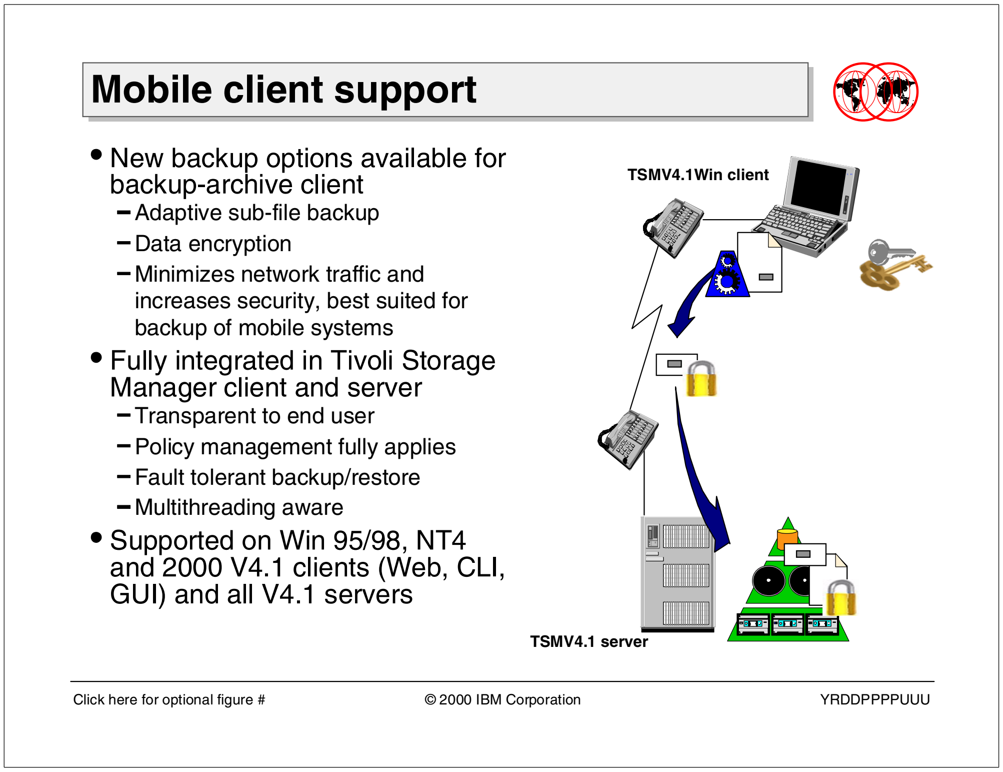

Mobile client support

- New backup options available for backup-archive client
  - Adaptive sub-file backup
  - Data encryption
  - Minimizes network traffic and increases security, best suited for backup of mobile systems
- Fully integrated in Tivoli Storage Manager client and server
  - Transparent to end user
  - Policy management fully applies
  - Fault tolerant backup/restore
  - Multithreading aware
- Supported on Win 95/98, NT4 and 2000 V4.1 clients (Web, CLI, GUI) and all V4.1 servers

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

1.3.5 Mobile client support

Mobile and remote computers have limited access to the infrastructure that services the rest of the company. This limited access both increases the criticality of storage management services and limits the applicability of traditional methods and policies. Typically connected to the backup server over a modem, these machines can only send small amounts of data in a reasonable amount of time. Critical data needs to be backed up to protect the remote users.

1.3.5.1 Adaptive sub-file backup

Tivoli Storage Manager offers a new option to resolve the limited bandwidth problem, which is the new *adaptive sub-file backup* features. The new feature introduces a new level of granularity for backup objects, which is on a sub-file level, either on block or byte bases. A client backup operation using this feature will be able to send only the changed bytes or blocks of a file to the server, instead of the entire file. The feature is integrated into the Tivoli Storage Managers Version 4.1 Windows 32-bit clients, and all Version 4.1 servers.

1.3.5.2 Data encryption support

In order to improve the security of backup of data over unsecured telephone lines, the Version 4.1 Tivoli Storage Manager Windows clients implement an encryption function. This allows for data to be encrypted before it is sent to the Tivoli Storage Manager server. Not only does this help secure backed up data during transmission, but it also means that the data stored on the Tivoli Storage Manager server is encrypted and would thus be unreadable by any malicious administrators.

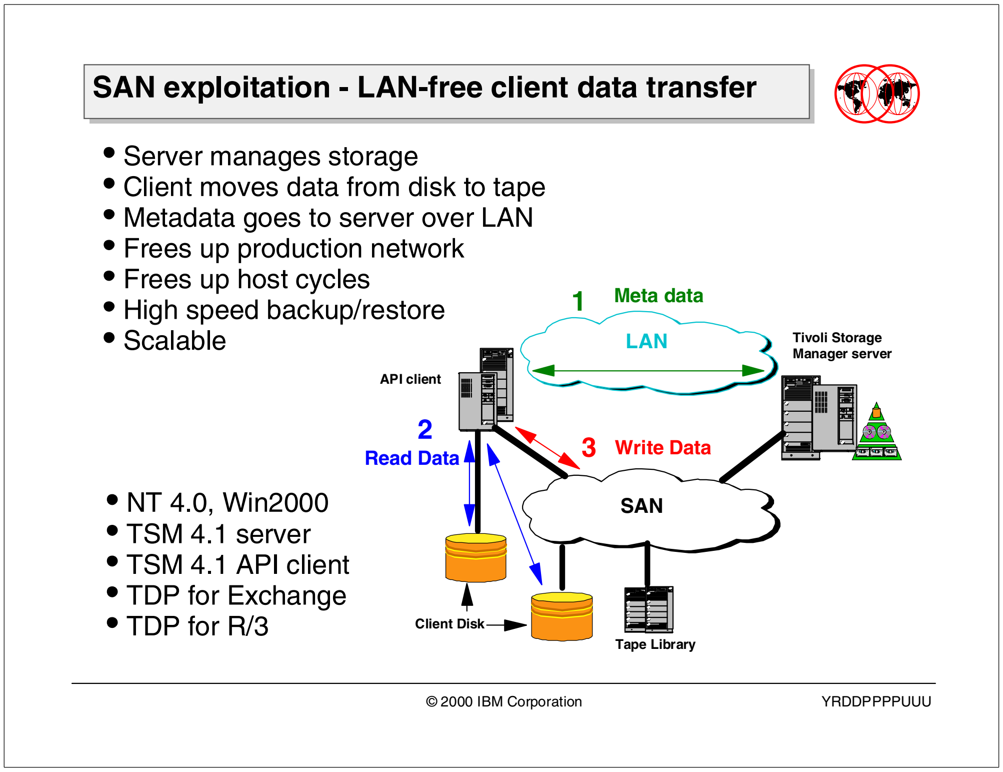

SAN exploitation - LAN-free client data transfer

- Server manages storage
- Client moves data from disk to tape
- Metadata goes to server over LAN
- Frees up production network
- Frees up host cycles
- High speed backup/restore
- Scalable

- NT 4.0, Win2000
- TSM 4.1 server
- TSM 4.1 API client
- TDP for Exchange
- TDP for R/3

© 2000 IBM Corporation YRDDPPPPUUU

1.3.6 SAN exploitation - LAN-free client data transfer

The advent of SAN technology provides an alternative path for data movement between the Tivoli Storage Manager client and the server. Shared storage resources are accessible to both the client and the server through the storage area network. It is now possible for the client to write to tape storage managed by the server. Data movement is off-loaded from the LAN and from the server processor for greater scalability.

LAN-free client data transfer is a new feature of Tivoli Storage Manager V4.1 that allows a Tivoli Data Protection for Application Client to transfer data directly to a SAN attached tape device, that is known to both the client and server. The TSM API client in conjunction with the enhanced TSM Server and a new TSM Storage Agent have been enhanced with the ability to write directly to server-owned tape storage media in a format that is consistent with that written by the server today. TDP for Exchange and TDP for R/3 are the only supported applications in the initial release.

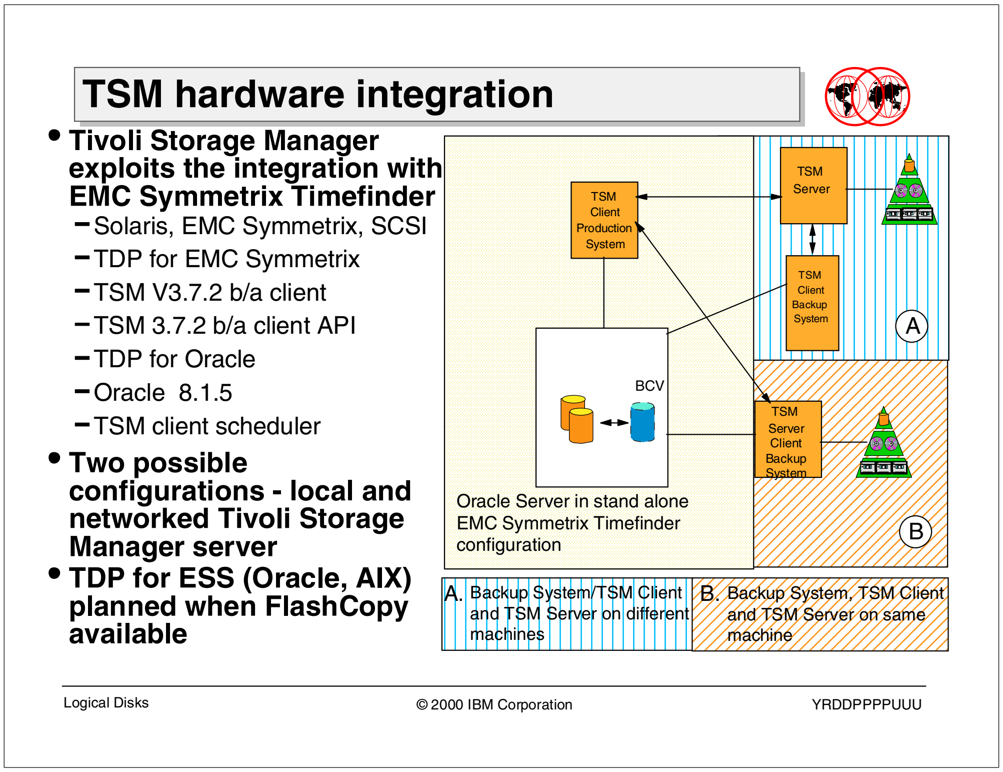

TSM hardware integration

- Tivoli Storage Manager exploits the integration with EMC Symmetrix Timefinder
  - Solaris, EMC Symmetrix, SCSI
  - TDP for EMC Symmetrix
  - TSM V3.7.2 b/a client
  - TSM 3.7.2 b/a client API
  - TDP for Oracle
  - Oracle 8.1.5
  - TSM client scheduler
- Two possible configurations - local and networked Tivoli Storage Manager server
- TDP for ESS (Oracle, AIX) planned when FlashCopy available

TSM Client Production System

TSM Server

TSM Client Backup System

A

BCV

TSM Server Client Backup System

B

Oracle Server in stand alone EMC Symmetrix Timefinder configuration

A. Backup System/TSM Client and TSM Server on different machines

B. Backup System, TSM Client and TSM Server on same machine

Logical Disks © 2000 IBM Corporation YRDDPPPPUUU

1.3.7 TSM hardware integration

As databases grow in size, customers are looking for faster backup and restore solutions. The backup operation for databases with large amounts of data may take a long time to perform. If a local area network is used to back up data from the database server to the Tivoli Storage Manager server, the network often becomes the main bottleneck.

Tivoli Storage Manager Version 3.7.3 and above have been enhanced to provide automated exploitation of the EMC TimeFinder copy function for the Oracle database server. The backup process of an Oracle database on a Sun Solaris machine with an EMC Symmetrix storage subsystem can now be efficiently performed on a backup copy of the actual production data.

A new backup server system is required, which is SCSI-connected to the storage subsystem. The backup system hosts the Tivoli Storage Manager client application (TDP for Oracle). To minimize data movement of the backup data on the network, a more powerful and integrated solution can be implemented by having the Tivoli Storage Manager server also running on the backup system.

A new Tivoli Data Protection for EMC Symmetrix has been developed and is available with a sample command script. TDP for EMC Symmetrix can be used for Oracle databases or for SAP R/3 using an Oracle database to do offline or online backups of the full database. The current release will support the Veritas File System and Universal File System.

As soon as FlashCopy for heterogeneous systems becomes available, the same concept will be used to provide a TDP for IBM ESS product for a similar Oracle backup solution on AIX using the IBM ESS disk storage system.

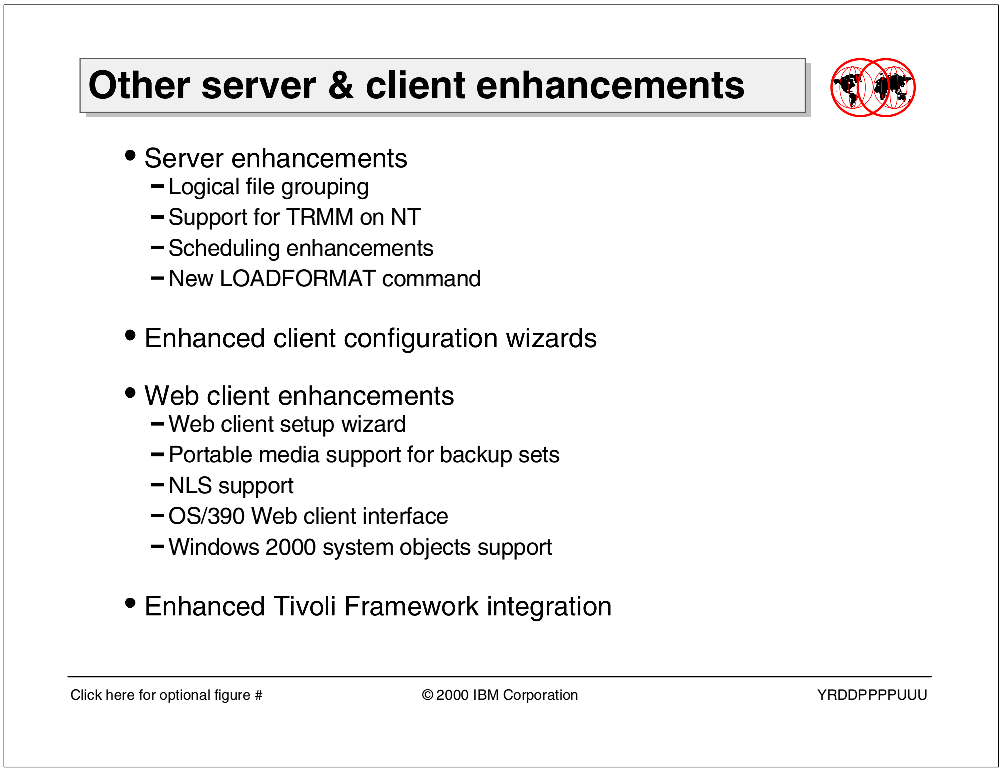

Other server & client enhancements

- Server enhancements
  - Logical file grouping
  - Support for TRMM on NT
  - Scheduling enhancements
  - New LOADFORMAT command
- Enhanced client configuration wizards
- Web client enhancements
  - Web client setup wizard
  - Portable media support for backup sets
  - NLS support
  - OS/390 Web client interface
  - Windows 2000 system objects support
- Enhanced Tivoli Framework integration

Click here for optional figure #

YRDDPPPPUUU

1.3.8 Other server and client enhancements

Tivoli Storage Manager Version 3.7.2 and Version 4.1 introduce the following new features and functions for client and server:

- The concept of Logical File Grouping on the Tivoli Storage Manager server to support the adaptive sub-file backup feature and the backup of complex Windows 2000 system objects
- Support of Tivoli Removable Media Manager on Windows NT as an external library manager
- Server scripting and scheduling enhancements to support the new hardware integration feature and to introduce event driven scheduling
- A new DSMSERV LOADFORMAT command for server database maintenance
- Enhanced and new client configuration wizards to improve the configuration and usage of the backup-archive client, Web client and the client scheduler
- Numerous enhancements of the Web client interface as portable media support, NLS support, Windows 2000 system object support and others
- Enhanced Tivoli Framework integration

1.4 Tivoli Storage Management summary

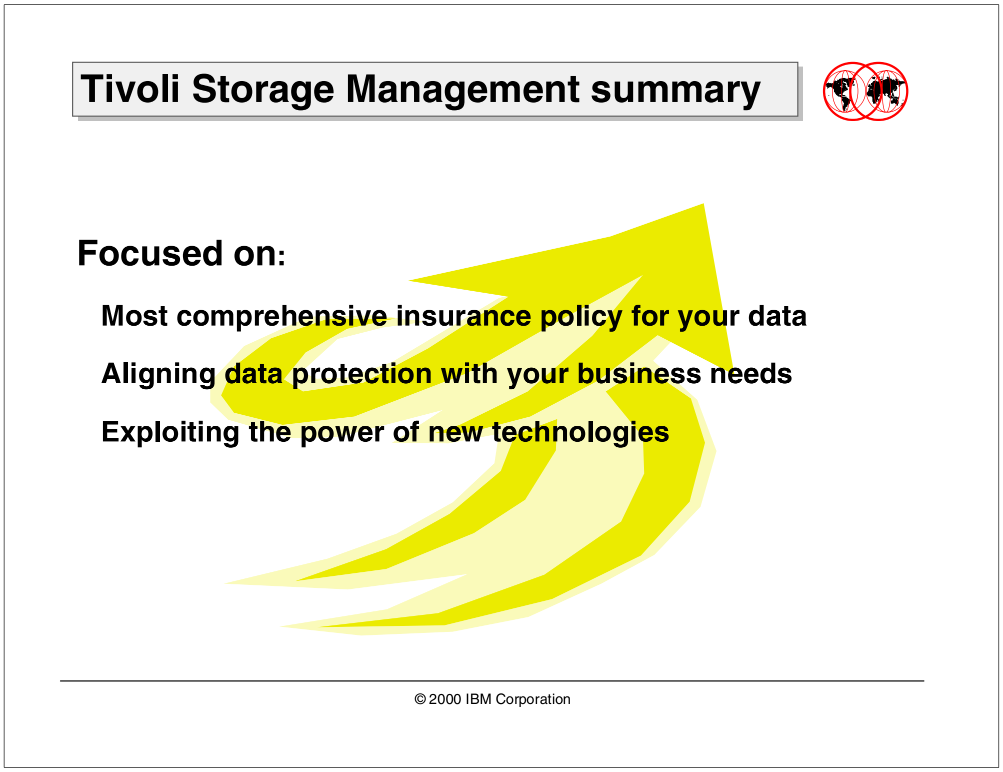

Tivoli Storage Management summary

Focused on:

Most comprehensive insurance policy for your data

Aligning data protection with your business needs

Exploiting the power of new technologies

The Tivoli Storage Management Solution is focused on the most comprehensive insurance policy for the customer’s business data. It aligns data protection with customer business needs and exploits the power of new technologies.

Chapter 2. Backup-archive client enhancements

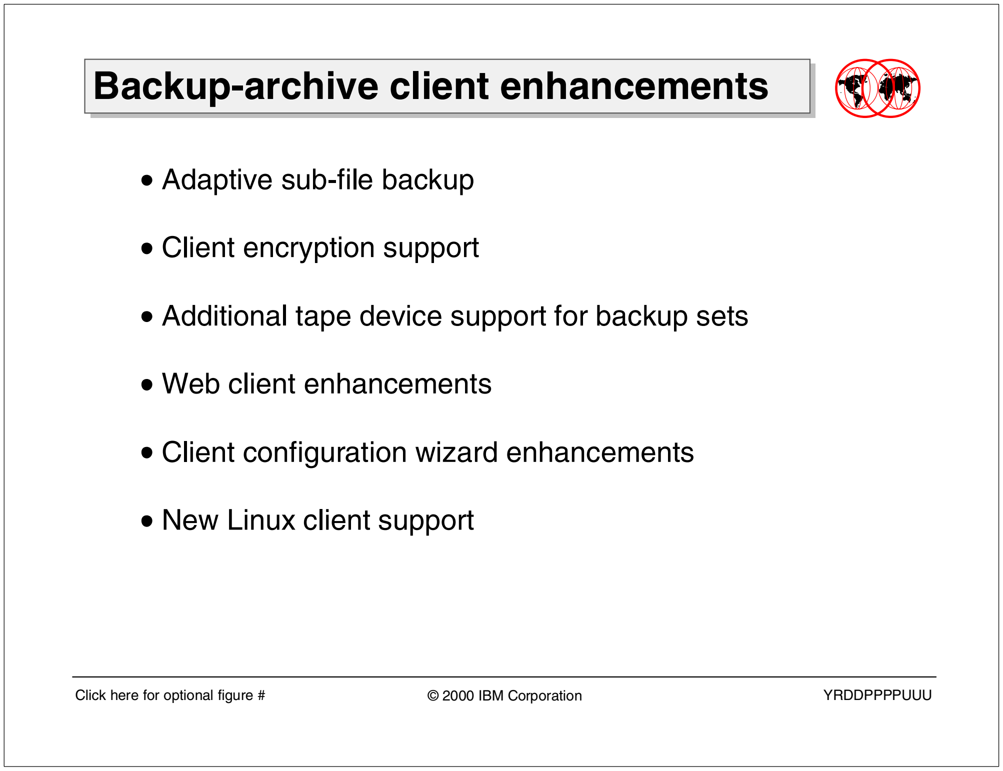

Backup-archive client enhancements

- Adaptive sub-file backup
- Client encryption support
- Additional tape device support for backup sets
- Web client enhancements
- Client configuration wizard enhancements
- New Linux client support

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

This chapter covers major backup-archive client enhancements, delivered with Tivoli Storage Manager Version 3.7.2 and Version 4.1.0. These enhancements are:

- Adaptive sub-file backup, a new client backup technology which reduces transferred data objects to only the modified portions of a physical file
- Client encryption support, a new client feature which allows for encrypting data before sending it to the Tivoli Storage Manager server
- Additional tape device support for restoring backup sets LAN-free on the client system
- New Web client enhancements, including NLS support, new supported system objects, support of portable media, and a new configuration wizard
- Client configuration wizard enhancements
- An introduction to the newly supported backup-archive client for the popular Linux platform

2.1 Adaptive sub-file backup

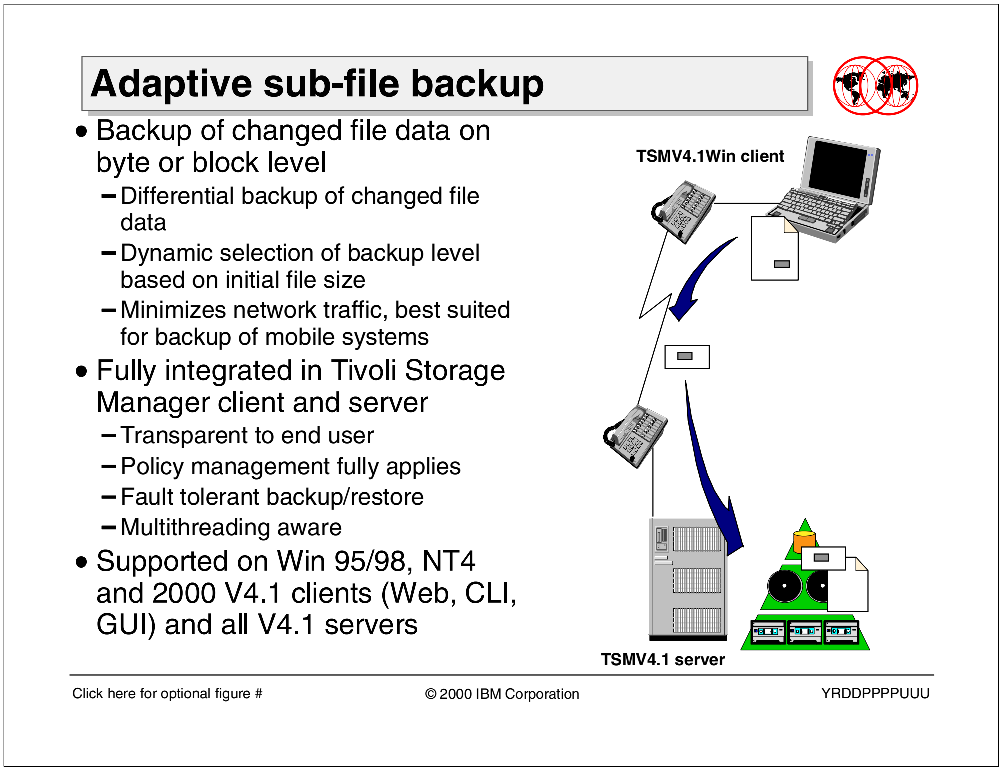

Adaptive sub-file backup

- Backup of changed file data on byte or block level
  - Differential backup of changed file data
  - Dynamic selection of backup level based on initial file size
  - Minimizes network traffic, best suited for backup of mobile systems
- Fully integrated in Tivoli Storage Manager client and server
  - Transparent to end user
  - Policy management fully applies
  - Fault tolerant backup/restore
  - Multithreading aware
- Supported on Win 95/98, NT4 and 2000 V4.1 clients (Web, CLI, GUI) and all V4.1 servers

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

2.1.1 Overview

As the number of laptop PCs approaches 20 percent of the PC install base, many central support organizations will need to provide storage management services for their mobile and remote workers. Mobile and remote computers have limited access to the infrastructure that services the rest of the company. Some limitations include being attached to the corporate network with limited bandwidth, limited connect time, and limited or no assistance to perform the backup.

This limited access both increases the criticality of storage management services and limits the applicability of traditional methods and policies. Tivoli Storage Manager is helping to resolve these problems, with its new *adaptive sub-file backup* feature, which reduces the amount of data transferred while backing up changed files.

The new features enable the backup-archive client (Web client, command line and GUI) to back up only the changed portion of a file, either on byte level or on block level, instead of transferring the whole file to the server every time. The changed file portion is backed up as a differential backup relative to the last complete backup of the file (*base* or *reference file*) and it is called *delta file*. All the changes of the file, since the last complete backup of the file, are included in this delta file. In the case of a restore, this allows for the restore of the whole file by restoring only two *sub-file components,* one delta file and the last complete backup of the whole file, the base file.

The decision as to whether the differential is based on byte level or on block level will be done at the backup of the base file and depends on the size of the file. Sub-file backup technology is not used for very small files (less than 1 KB in size) and for very large files (bigger than 2GB). If the delta file size exceeds 60 percent of the base file at the last sub-file backup, a new base file will be transferred.

The adaptive sub-file backup as well as the restore of a file consisting of a base file and the delta file is completely transparent to the end user. All the needed separations or reconstructions of the file data happen under-the-covers of the backup-archive client. Also, all the other Tivoli Storage Manager features such as policy management or fault tolerant backup and restore still fully apply. Adaptive sub-file backup is used for incremental as well as for selective backup. It is aware of multithreading and will work together with client data compression and encryption.

This new feature is available for Tivoli Storage Manager Version 4 Windows 95/98, Windows NT4, and Windows 2000 clients, and for all Version 4 servers.

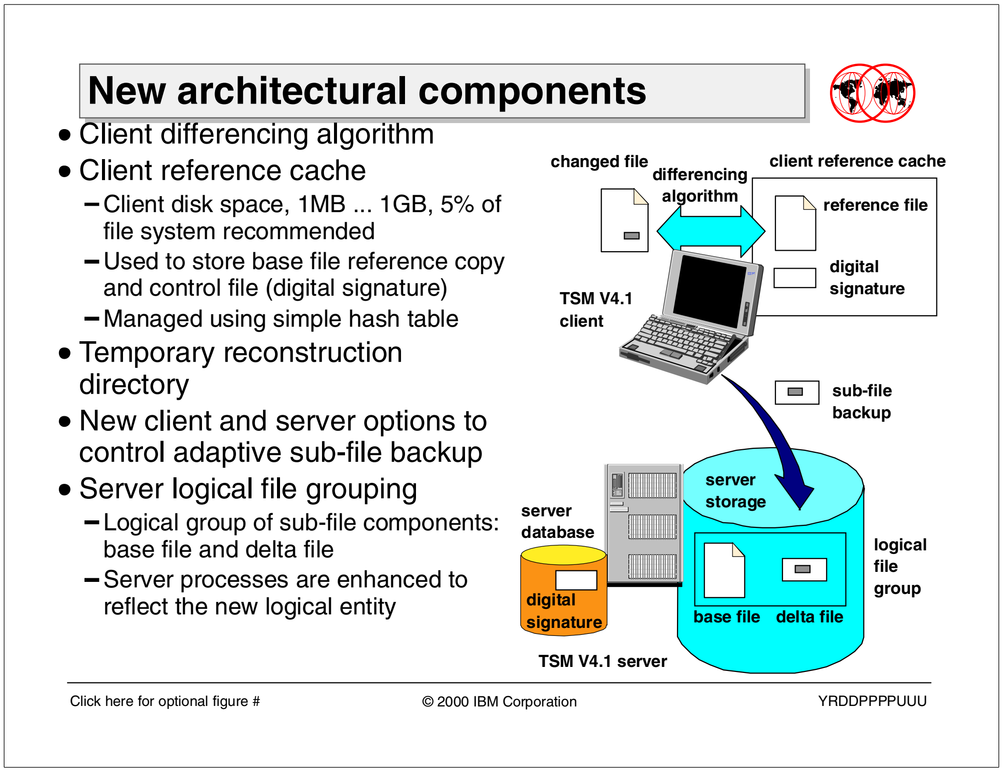

New architectural components

- Client differencing algorithm
- Client reference cache
  - Client disk space, 1MB ... 1GB, 5% of file system recommended
  - Used to store base file reference copy and control file (digital signature)
  - Managed using simple hash table
- Temporary reconstruction directory
- New client and server options to control adaptive sub-file backup
- Server logical file grouping
  - Logical group of sub-file components: base file and delta file
  - Server processes are enhanced to reflect the new logical entity

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

2.1.2 New architectural components

In order to implement adaptive sub-file backup with y the Tivoli Storage Manager client and server, the external interface remained unchanged. All the architectural changes do not affect the way the end user initiates backups and restores, nor does it introduce a new backup or restore option. The only externals are additional client and server options to enable and configure adaptive sub-file backup for the whole client.

However, under-the-covers, numerous architectural changes are implemented on the client and server to accommodate the new function. The client sub-file backup introduces:

- The *differencing algorithm,* a patented technology, which finds the file changes between two versions of the same file, either on byte or block level.
- The *client reference cache* stores reference information, which is used to determine file changes and controls the consistency of this information.
- The *temporary reconstruction directory*, which is used to store the base file and delta file and to reconstruct the file version from both components.

On the server, Version 4.1 introduces:

- The *logical file grouping*, which allows the grouping and managing of all delta files and the base file of a file as a logical entity. It also ensures that, according to the implemented file management policies, every file version can be reconstructed. More information on logical file grouping on the server can be found in Section 3.1, “Logical file grouping server support” on page 82.

2.1.2.1 Client reference cache

The client reference cache is a directory in the client file system, which is used to store reference information needed to determine file changes. It also contains a small database used to control the consistency of this information. The maximum size, which can range between 1MB and 1GB, and the location of this cache directory, is controlled by client options. The cache is created the first time a file is backed up while sub-file backup is activated.

When a file is backed up for the first time in the cache, two objects are stored:

- A copy of the base-file component, called the *reference file*
- A 256 bit file checksum, called the *digital signature* of the reference file

The reference file is either a byte-by-byte copy of the original file at the time of the base file backup or a sequence of 256 bit block signatures computed from the original file at the time of the base file backup. The blocksize used to compute the signatures depends on the size of the original file and will be chosen as an increment of 4K to achieve a reference file size below 64K. The decision as to what type of reference file is stored depends on whether byte-level or block-level sub-file backup was chosen by the client.

The digital signature of the reference file is also created at the time of the first backup of a file under adaptive sub-file backup and represents a 256 bit checksum of the whole file. The digital signature is also sent to the server each time when a base file or delta file is transferred and will be stored in the server database. It is used to ensure the consistency between the reference file (stored in the cache) and the base file (stored in the server storage repository).

The client reference cache is managed by a simple hash table, called the *cache control database*. Each reference file stored in client cache has an entry in that database, which hashes the file names, holds additional information, such as whether byte- or block-level differencing is used, and the last compression ratio; and it chains the entries in a least-recently-used chain. If the cache fills up, reference files are purged from the cache according to a least-recently-used algorithm. This triggers a new base file backup at the next backup of the file.

2.1.2.2 Differencing algorithm

The differencing algorithm is implemented in the backup-archive client in order to detect differences between the reference file (stored in the client reference cache) and a changed version of the original file. Next it creates a delta file, which is then backed up to the server instead of the complete version of the file. The process can be described using the following formula:

$V_{reference} + V_{version}$ ->Delta($V_{reference}$, $V_{version}$)

Reconstruction, the inverse operation, requires the base file and the sub-file component to rebuild the version

V_{base} + Delta($V_{reference}$, $V_{version}$) -> $V_{version}$

2.1.2.3 Temporary reconstruction directory

During a restore, the file components are restored to a temporary client directory using special naming conventions. This is important to keep track of the restore in case the restore session is interrupted.

As soon as both the base file and the delta file are completely present in the directory, the reconstruction will take place in the same directory to create the file version, which the user attempted to restore. As a final step, the file is copied to the original location, and the base and delta files in the temporary directory are deleted.

The common location of that temporary directory is `\~tsmtemp\` of that file system, to which the restore was directed. The directory will be created automatically and all the files within that directory will be cleaned up automatically after completion of the restore. The user should not try to alter or delete entries in that directory.

Note

If there are no restartable restores for any node on that client OR any active restore operations going on, then the user can remove any entries in this temporary directory to avoid over-committing the file system. However, this should only be done if both of the above-mentioned conditions are true.

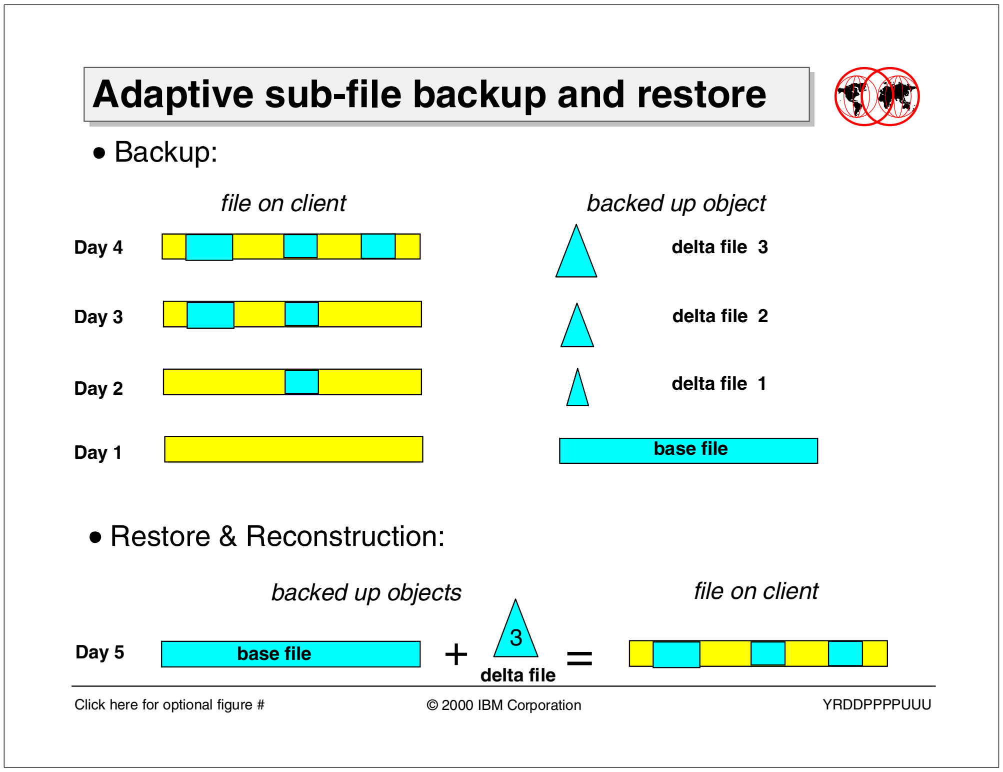

Adaptive sub-file backup and restore

- Backup:

file on client

backed up object

Day 4

delta file 3

Day 3

delta file 2

Day 2

delta file 1

Day 1

base file

- Restore & Reconstruction:

backed up objects

file on client

Day 5

base file + 3 delta file =

Click here for optional figure #

© 2000 IBM Corporation

YRDDPPPPUUU

2.1.3 Sub-file backup and restore

When a user attempts to back up a file having adaptive sub-file backup enabled, basically, two different kinds of backups can occur. The user has no influence on this; the backup-archive client itself decides based on built-in rules, which of the following backup steps will occur:

- Back up the file as a base file component.
- Back up the file as a delta file component.

The first step is always to take a backup of the entire file. This backup is considered to be the base file backup. Subsequent backups use delta file backups. Delta file backups take a delta of what has changed relative to and based on the reference file. The delta file backup can occur up to 20 times before a new base file backup occurs again. However, if the changes to the file are too numerous, and at the last delta file backup, a compression of at least 40 percent could not be achieved, then a new base file backup will be performed.

The base file and delta file represent versions of the same file. To restore a file version that has been backed up using adaptive sub-file technology, Tivoli Storage Manager does not create incremental delta chains from all the backed up sub-file components. This means that in the case of 3 delta files backed up on 3 consecutive days, in order to recreate the latest version of the file the backup-archive client does not have to restore 3 delta files. Rather, only the base file and the delta file of the last day are restored. The delta file of the last day contains all changes that delta file 2 and delta file 1 contain.

The differencing is performed only on the file data, and not on any other control data associated with the file, for example, ACL and DACL. During the restore operation, the control data of the delta file is preserved, as it represents the most current version of the control data.

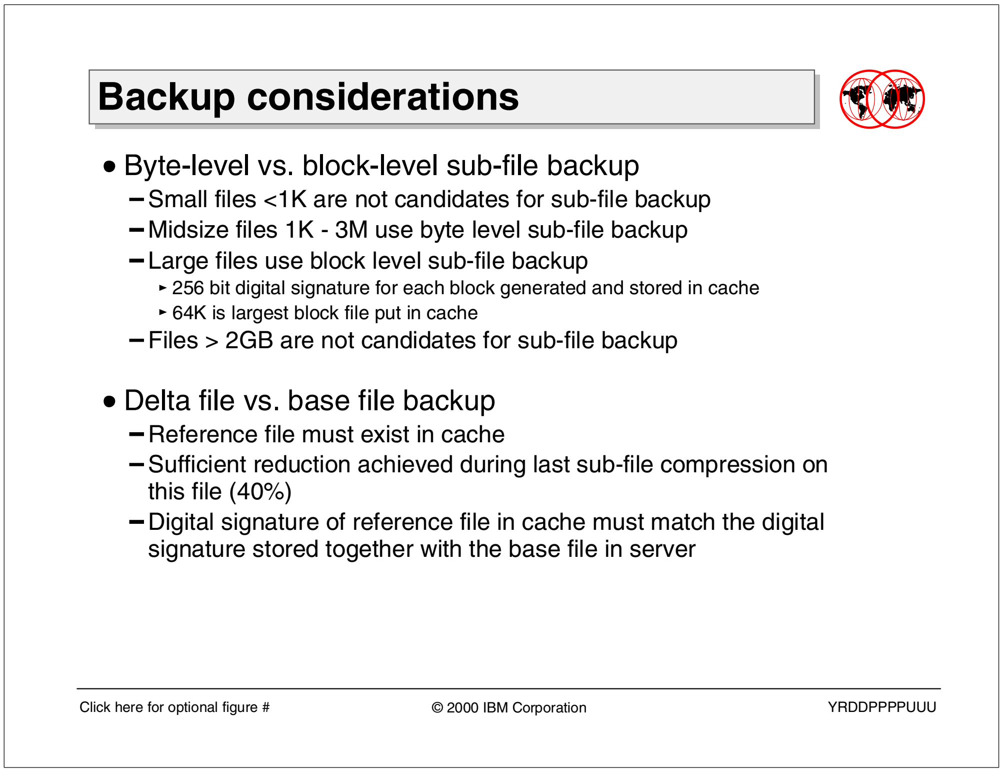

Backup considerations

- Byte-level vs. block-level sub-file backup
  - Small files <1K are not candidates for sub-file backup
  - Midsize files 1K - 3M use byte level sub-file backup
  - Large files use block level sub-file backup
    - 256 bit digital signature for each block generated and stored in cache
    - 64K is largest block file put in cache
  - Files > 2GB are not candidates for sub-file backup
- Delta file vs. base file backup
  - Reference file must exist in cache
  - Sufficient reduction achieved during last sub-file compression on this file (40%)
  - Digital signature of reference file in cache must match the digital signature stored together with the base file in server

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

2.1.3.1 Backup consideration

Tivoli Storage Manager combines several methods to create a sophisticated strategy for backing up files with the least amount of bandwidth and computer overhead needed. The delta file is either based on byte level differential between the reference file and the version file, or on block level changes dependent on the initial size of the file when a base file backup was performed. The following rules apply:

- Small files <1 KB are not candidates for the adaptive sub-file backup technology, and are always backed up complete. No reference file will be created in the client reference cache.
- Midsize files 1 KB to 3 MB are candidates for sub-file backup. The delta file backup will be performed using byte-level differentials.
- Large files between 3 MB and 2 GB are candidates for sub-file backup. The delta file backup will be performed using block-level differentials.
- Huge files larger than 2 GB are not candidates for the adaptive sub-file backup technology, and are always backed up complete. No reference file will be created in the client reference cache.

The requirements for sub-file backup to occur are:

- The file being backed up must be present in the client cache with a reference file and a digital signature.
- Sufficient reduction must have been achieved during the last sub-file compression on this file.
- The digital signature of the reference file in cache must match the digital signature stored together with the base file in the server.

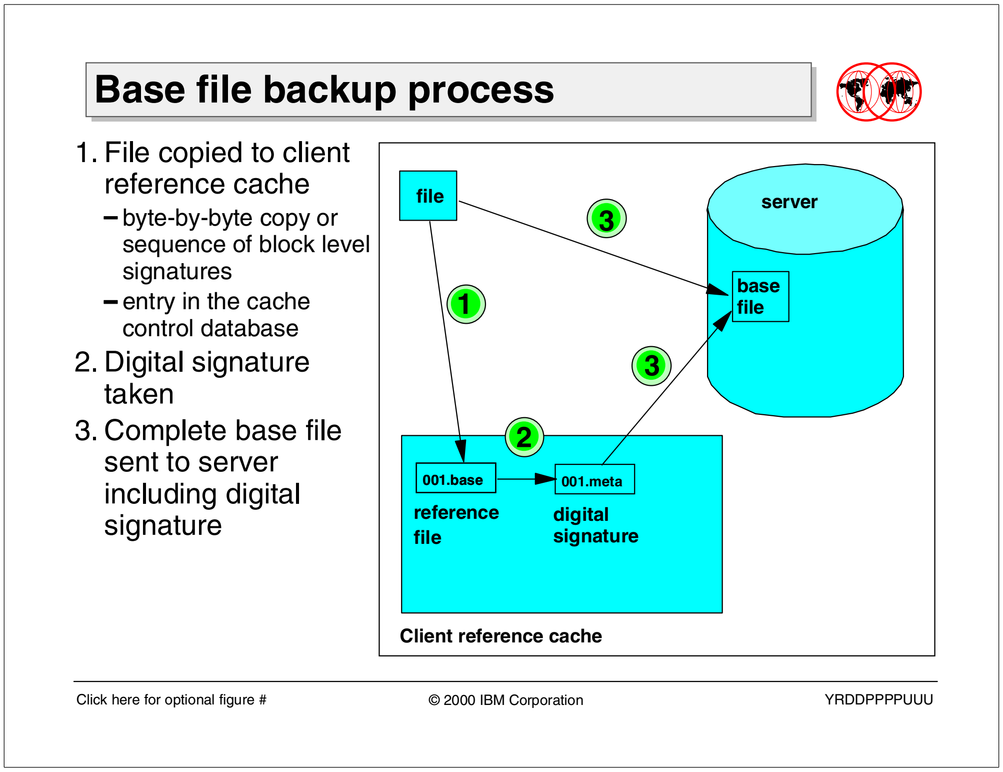

Base file backup process

1. File copied to client reference cache
   - byte-by-byte copy or sequence of block level signatures
   - entry in the cache control database
2. Digital signature taken
3. Complete base file sent to server including digital signature

Click here for optional figure #

© 2000 IBM Corporation

YRDDPPPPUUU

2.1.3.2 Base file backup process

Based on the previous discussion, the first step of the sub-file backup process is a base file backup. As soon as sub-file backup is enabled and a file eligible for sub-file backup is backed up by a user using the incremental or the selective backup function, the backup-archive client starts the following process:

1. The file is copied to the client reference cache to create a reference file there. Depending on the size of the file, this is either a byte-by-byte copy or a sequence of block signatures. In the cache control database an entry will be added, which hashes the real file name to the unique name of the reference file and contains additional control information.
2. The digital signature of the file is taken and also stored in the client reference cache.
3. The original file is backed up to the server storage, including the created digital signature, which is stored in the server database.

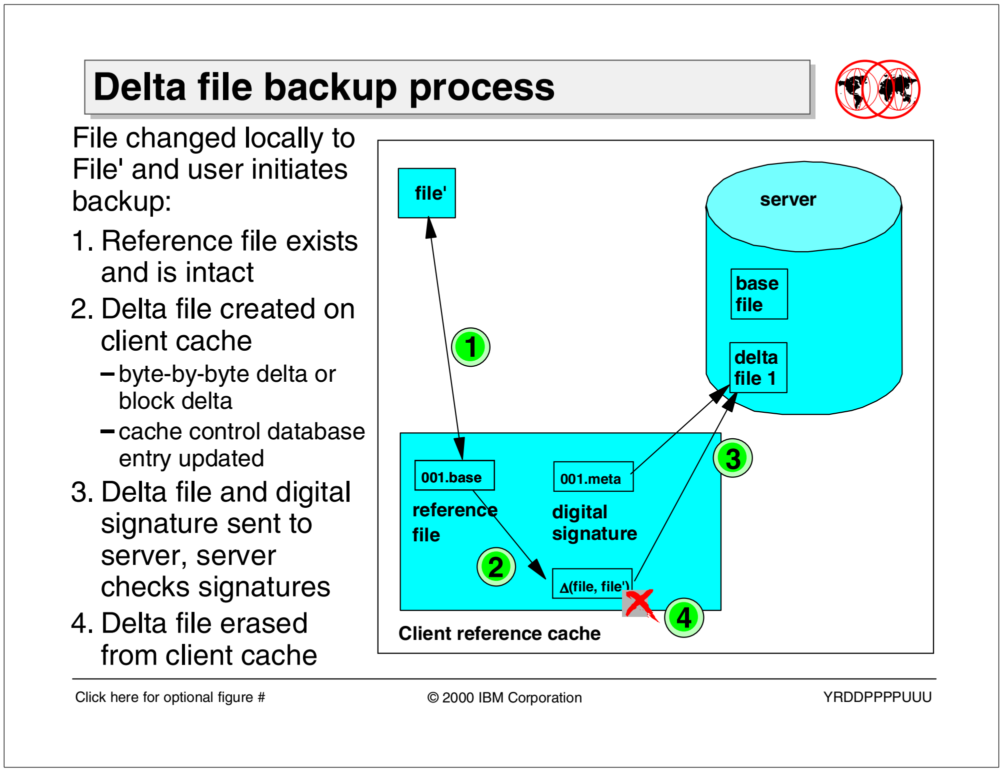

Delta file backup process

File changed locally to File' and user initiates backup:

1. Reference file exists and is intact
2. Delta file created on client cache
   - byte-by-byte delta or block delta
   - cache control database entry updated
3. Delta file and digital signature sent to server, server checks signatures
4. Delta file erased from client cache

Click here for optional figure #

© 2000 IBM Corporation

YRDDPPPPUUU

2.1.3.3 Delta file backup process

After the same file which was backed up in Section 2.1.3.2, “Base file backup process” on page 28 as a base file has changed and the user initiates another incremental or selective backup on that file, the following process is performed by the backup-archive client:

1. The client checks for the existence of a reference file in the client reference cache. Then the client makes sure that all the other conditions for a sub-file backup are fulfilled: the reference file is valid; the cache control database entry is consistent; and the compression achieved from the last sub-file backup is higher than 40 percent.
2. Then the client creates the delta file in the client cache, using its patented differencing algorithm. Depending on whether byte-level or block-level differencing is used, the delta file will contain information about the changed bytes or changed blocks of the current version of the file relative to the reference file (or base file). The differencing is not performed on non-file data, like the ACL data. So the delta file will also contain all the non-file data (ACLs) of the current version of the file.
3. The created delta file and the digital signature stored in the client cache are sent to the server. The server compares the received digital signature with the stored digital signature from base file backup and if they match, the server will accept the delta file. If both signatures do not match, a transaction failure will occur, and a new base file backup will be triggered.
4. The delta file is erased from the client reference cache.

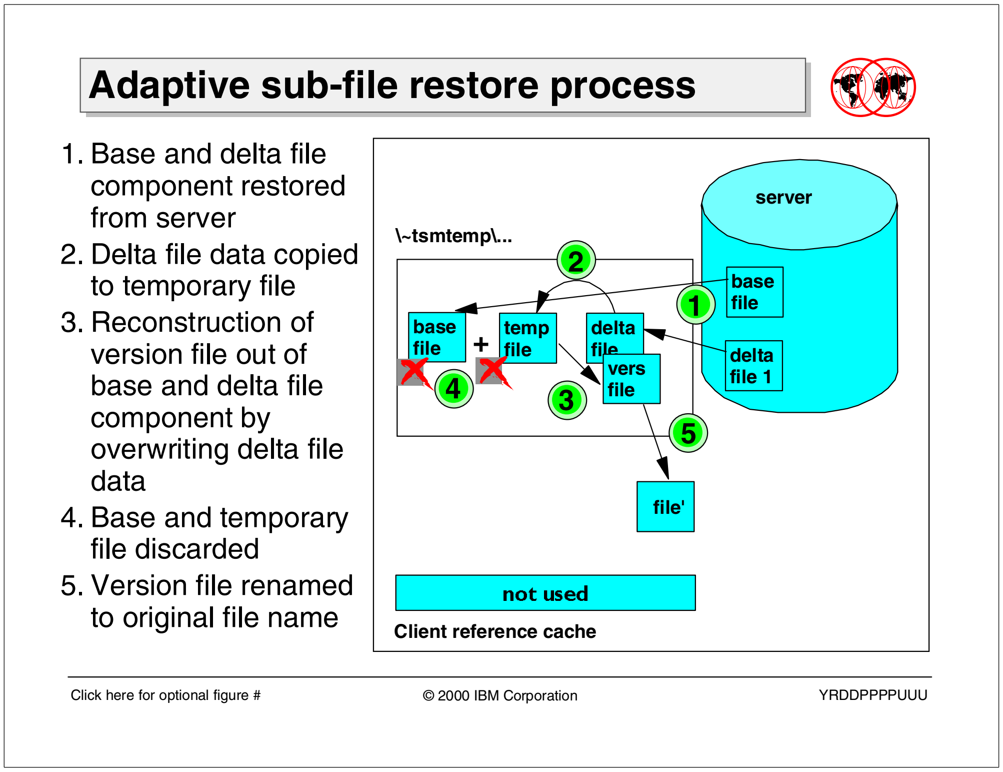

Adaptive sub-file restore process

1. Base and delta file component restored from server
2. Delta file data copied to temporary file
3. Reconstruction of version file out of base and delta file component by overwriting delta file data
4. Base and temporary file discarded
5. Version file renamed to original file name

Click here for optional figure #

© 2000 IBM Corporation

YRDDPPPPUUU

2.1.3.4 Sub-file restore process

The server versions of the file are represented either through the base file or the delta file component. To restore a file version the end user initiates a restore of one of those components, under-the-covers, however, the client and the server are performing the following steps:

1. The corresponding base and delta file component of the file version are restored from the server to a temporary reconstruction directory under `\~tsmtemp` on the logical drive to which the file is attempting to be restored. The file components are restored under a temporary name in order to allow reconstruction of them before moving the version file to the original location and to keep track of the restore status as well as the relationship between the components.
2. Delta file data is copied to a *temporary rebuild file* in order to allow reconstruction of the file by overwriting the restored delta file data so it can utilize the ACL data of that delta file version.
3. The base file and the delta file component stored by the temporary reconstruction file are reconstructed into the version file by overwriting the data part of the restored delta file. The ACLs of this component (representing the ACL of the restored version of the file) remain intact.
4. The base file and the temporary reconstruction file are deleted from the reconstruction directory.
5. The version file is renamed back to its original name.

File components can be restored in any order; the components do not have to be restored sequentially.

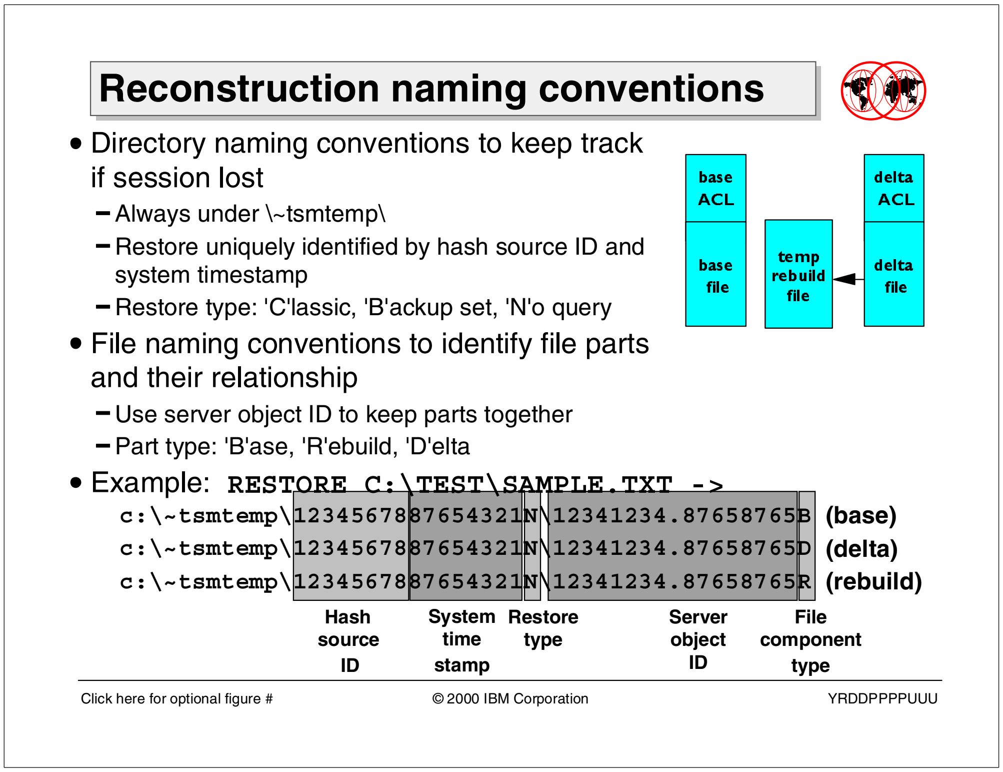

Reconstruction naming conventions

- Directory naming conventions to keep track if session lost
  - Always under \~tsmtemp\
  - Restore uniquely identified by hash source ID and system timestamp
  - Restore type: 'C'lassic, 'B'ackup set, 'N'o query
- File naming conventions to identify file parts and their relationship
  - Use server object ID to keep parts together
  - Part type: 'B'ase, 'R'ebuild, 'D'elta
- Example: `RESTORE C:\TEST\SAMPLE.TXT ->`

```
c:\~tsmtemp\1234567887654321N\12341234.87658765B (base)
c:\~tsmtemp\1234567887654321N\12341234.87658765D (delta)
c:\~tsmtemp\1234567887654321N\12341234.87658765R (rebuild)
```

Click here for optional figure #

© 2000 IBM Corporation

YRDDPPPPUUU

2.1.3.5 Naming conventions during version file reconstruction

As shown in Section 2.1.3.4, “Sub-file restore process” on page 30 during the restore process the sub-file components are stored in a temporary reconstruction directory. Within this directory the components are stored in temporary directories under temporary names following specific naming conventions. Once the file is reconstructed, the final file version is renamed to its original restore path name.

Temporary directory naming conventions are necessary to be able to keep track of restores, if the session is lost. The directory conventions identify that a sub-file component restore occurs, they uniquely identify the restore request and mark the general type of restore. The directory is a subdirectory of the \~tsmtemp\ in the same file system as restore. To uniquely identify the restore, the hash source ID and a timestamp are used. The type of restore is marked by a trailing capital letter: C for classic, B for backup set, and N for no-query restore.

Temporary file naming conventions are used until the file has been reconstructed to identify the relationship between the file parts and to mark the part type. The relationship between the file parts is maintained using the object-ID of the base file from the server, the file component type is coded by a trailing capital letter: B for base file component, R for temporary rebuild, and D for delta file component.

The following example illustrates these conventions. Assume that a user issues the following restore command `RESTORE "C:\TEST\*"`. The client creates the `\~tsmtemp\` directory on C: drive and hashes `"C:\TEST\*"` to 0x12345678. From the system, it gets a timestamp 0x87654321. Because this is a no-query restore, the trailing capital letter of the temporary directory name is N. The complete restore path would be `C:\~tsmtemp\1234567887654321N\` for storing the sub-file components. The path is stored in the server for restart.

For example, assume that c:\test\sample.txt is the first file which is restored, consisting of base file and delta file components. The object-ID on the server is 0x1234123487658765. Following the file naming conventions, the base file is named C:\~tsmtemp\1234567887654321N\ 12341234.87658765B, the restored delta file is named C:\~tsmtemp\1234567887654321N\ 12341234.87658765D and the temporary rebuild file is called C:\~tsmtemp\1234567887654321N\ 12341234.87658765R.

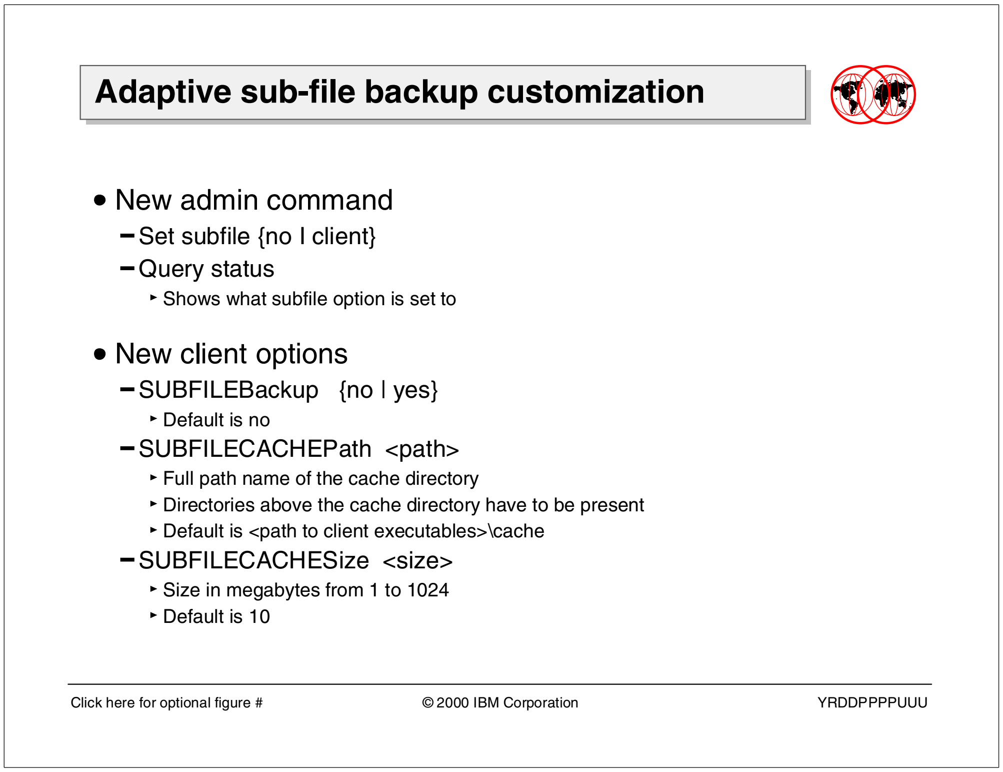

Adaptive sub-file backup customization

- New admin command
  - Set subfile {no I client}
  - Query status
    - Shows what subfile option is set to
- New client options
  - SUBFILEBackup {no I yes}
    - Default is no
  - SUBFILECACHEPath <path>
    - Full path name of the cache directory
    - Directories above the cache directory have to be present
    - Default is <path to client executables>\cache
  - SUBFILECACHESize <size>
    - Size in megabytes from 1 to 1024
    - Default is 10

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

2.1.4 Adaptive sub-file backup configuration

To use the adaptive sub-file backup feature, new options are introduced on client and server.

2.1.4.1 Controlling sub-file backup on server

To enable sub-file backup on the server the option SUBFILE needs to be set to CLIENT. This can be done by issuing the `SET SUBFILE` administrative command.

```
tsm: 78-BB053_SERVER1>set subfile client
ANR2288I Subfile set to CLIENT.
```

The default of this option is NO, which disables sub-file backup for all clients. The administrator can use the `QUERY STATUS` administrative command to find out what the option is set to.

2.1.4.2 Customizing sub-file backup on client

The SUBFILEBACKUP, SUBFILECACHEPATH, and SUBFILECACHESIZE are new client options added to the backup-archive client in order to customize the sub-file backup behavior of the client. The options can be set in the client options file `dsm.opt`.

Setting the option SUBFILEBACKUP to YES enables sub-file backup on the client side. From this point, **all** files backed up to the server using incremental or selective backup will be transferred using the adaptive sub-file backup technology. The default value of this option is NO.

The option SUBFILECACHEPATH defines the path to the client reference directory. If the first file is backed up using the adaptive sub-file backup technology, the backup-archive client creates this directory as well as the cache control database, and stores the reference file copy and the digital signature of the file there. The parent directory of this directory has to exist. If SUBFILECACHEPATH is not specified, the Tivoli Storage Manager client will create a directory `\cache` under the directory where the Tivoli Storage Manager client executables reside.

The maximum size of the client reference cache is controlled by the SUBFILECACHESIZE option. The size value specifies the size of the maximum size of the cache in megabytes. The range of values is 1 to 1024 MB. The default value is 10. If the cache fills up to the specified size, reference files are purged from the cache using a least-recently-used algorithm. If a file is backed up again, a new base file backup will be performed.

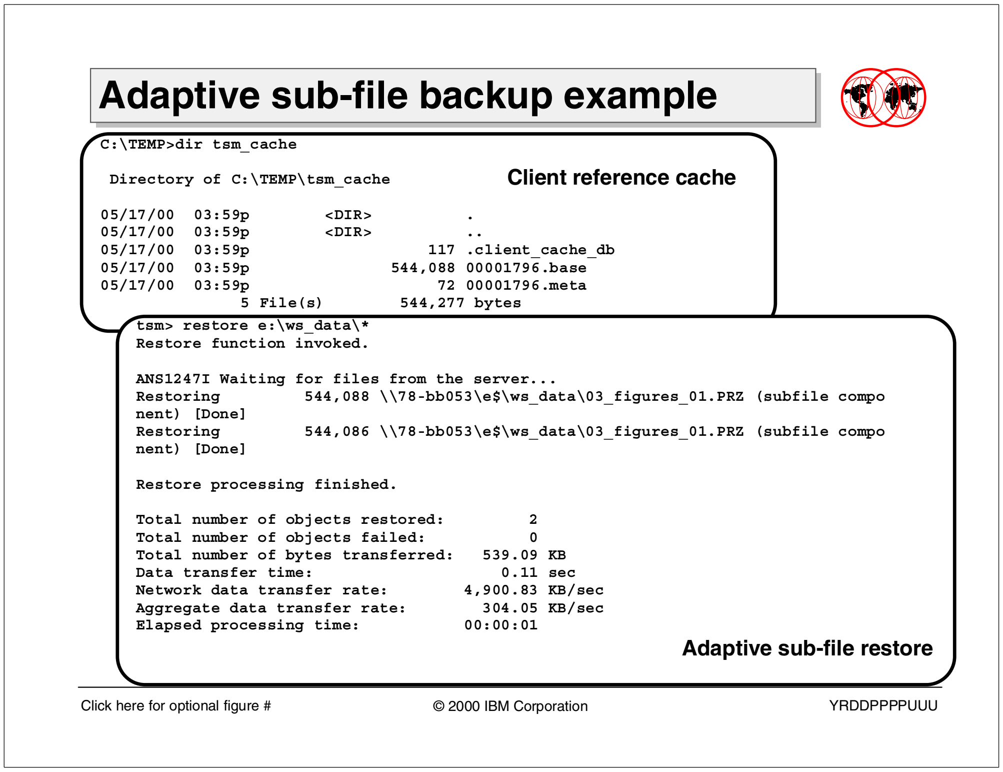

Adaptive sub-file backup example

```
C:\TEMP>dir tsm_cache

 Directory of C:\TEMP\tsm_cache

05/17/00  03:59p        <DIR>          .
05/17/00  03:59p        <DIR>          ..
05/17/00  03:59p                   117 .client_cache_db
05/17/00  03:59p               544,088 00001796.base
05/17/00  03:59p                    72 00001796.meta
               5 File(s)        544,277 bytes
```

Client reference cache

```
tsm> restore e:\ws_data\*
Restore function invoked.

ANS1247I Waiting for files from the server...
Restoring         544,088 \\78-bb053\e$\ws_data\03_figures_01.PRZ (subfile compo
nent) [Done]
Restoring         544,086 \\78-bb053\e$\ws_data\03_figures_01.PRZ (subfile compo
nent) [Done]

Restore processing finished.

Total number of objects restored:         2
Total number of objects failed:           0
Total number of bytes transferred:   539.09 KB
Data transfer time:                    0.11 sec
Network data transfer rate:        4,900.83 KB/sec
Aggregate data transfer rate:        304.05 KB/sec
Elapsed processing time:           00:00:01
```

Adaptive sub-file restore

Click here for optional figure #

© 2000 IBM Corporation

YRDDPPPPUUU

2.1.5 Adaptive sub-file backup example

In this example, a user backs up a 531 KB data file using selective backup with adaptive sub-file backup enabled. As shown in Figure 1, the selective backup command indicates that 531 KB of data is transferred to the server.

```
tsm> sel e:\ws_data\*
Selective Backup function invoked.

Directory-->                   0 \\78-bb053\e$\ws_data [Sent]
Normal File-->           544,088 \\78-bb053\e$\ws_data\03_figures_01.PRZ
 [Sent]
Selective Backup processing of '\\78-bb053\e$\ws_data\*' finished without failure.

Total number of objects inspected:       2
Total number of objects backed up:       2
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:   531.42 KB
Data transfer time:                    0.18 sec
Network data transfer rate:        2,952.34 KB/sec
Aggregate data transfer rate:         75.27 KB/sec
Objects compressed by:                   0%
Elapsed processing time:           00:00:07
```

Figure 1. Selective backup #1 using command line — 531 KB — base file transferred

Under-the-covers, the client creates the client reference cache and stores the reference file copy and the digital signature in the cache. In our example, the cache directory is `C:\temp\tsm_cache`. Figure 2 shows the content of the client reference cache after the selective backup transferred the base file to the server.

```
C:\TEMP>dir tsm_cache
 Volume in drive C has no label.
 Volume Serial Number is 36B0-570D

 Directory of C:\TEMP\tsm_cache

05/17/00  03:59p        <DIR>          .
05/17/00  03:59p        <DIR>          ..
05/17/00  03:59p                   117 .client_cache_db
05/17/00  03:59p               544,088 00001796.base
05/17/00  03:59p                    72 00001796.meta
               5 File(s)        544,277 bytes
                            326,415,872 bytes free
```

Figure 2. Content of the client reference cache

The file `0001796.base` is the reference file; the file `00001796.meta` is the digital signature. Both files are crucial for the correct function of the adaptive sub-file backup and may not to be altered or deleted under any circumstances. The file `.client_cache_db` is the client reference cache database, which contains a hash table and other information about the cache contents.

In our example now, the data file has been changed and backed up a second time as a new version, using selective backup as shown in Figure 3.

```
tsm> sel e:\ws_data\*
Selective Backup function invoked.

Directory-->                   0 \\78-bb053\e$\ws_data [Sent]
Normal File-->           544,086 \\78-bb053\e$\ws_data\03_figures_01.PRZ
[Sent]
Selective Backup processing of '\\78-bb053\e$\ws_data\*' finished without failure.


Total number of objects inspected:        2
Total number of objects backed up:        2
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:     7.66 KB
Data transfer time:                    0.00 sec
Network data transfer rate:            0.00 KB/sec
Aggregate data transfer rate:          1.09 KB/sec
Objects compressed by:                    0%
Elapsed processing time:           00:00:07
```

Figure 3. Selective backup #2 using command line — 7.66 KB — delta file transferred

In this backup, only 7.66 KB of data have been transferred. This is the size of the delta file.

Next, the original file was deleted and restored with the restore command from the command line client as shown in Figure 4.

```
tsm> restore e:\ws_data\*
Restore function invoked.

ANS1247I Waiting for files from the server...
Restoring         544,088 \\78-bb053\e$\ws_data\03_figures_01.PRZ (subfile compo
nent) [Done]
Restoring         544,086 \\78-bb053\e$\ws_data\03_figures_01.PRZ (subfile compo
nent) [Done]

Restore processing finished.

Total number of objects restored:         2
Total number of objects failed:           0
Total number of bytes transferred:   539.09 KB
Data transfer time:                    0.11 sec
Network data transfer rate:        4,900.83 KB/sec
Aggregate data transfer rate:        304.05 KB/sec
Elapsed processing time:           00:00:01
```

Figure 4. File restored — 539 KB — base file and delta file components transferred

The client shows that two sub-file components (base file and delta file) of the file are transferred from the server to the client. The client shows that, for both components, the actual size of the base file at the time of backup of each component. However, the total number of bytes transferred shows the correct amount of data (539 KB) transferred from server to the client.

Note

The reference file stored in the client cache is never used for the restore.

The components are temporarily restored to the reconstruction directory shown in Figure 5 (after the base file component was restored), the version file is created and renamed to its original location.

```
E:\~tsmtemp>dir 050086dd39232620N
 Volume in drive E has no label.
 Volume Serial Number is 2243-1D03

 Directory of E:\~tsmtemp\050086dd39232620N

05/17/00  04:07p        <DIR>          .
05/17/00  04:07p        <DIR>          ..
05/17/00  01:41p               544,088 0.2050B
               3 File(s)        544,088 bytes
                            273,317,888 bytes free
```

Figure 5. Temporary reconstruction directory and the base file component

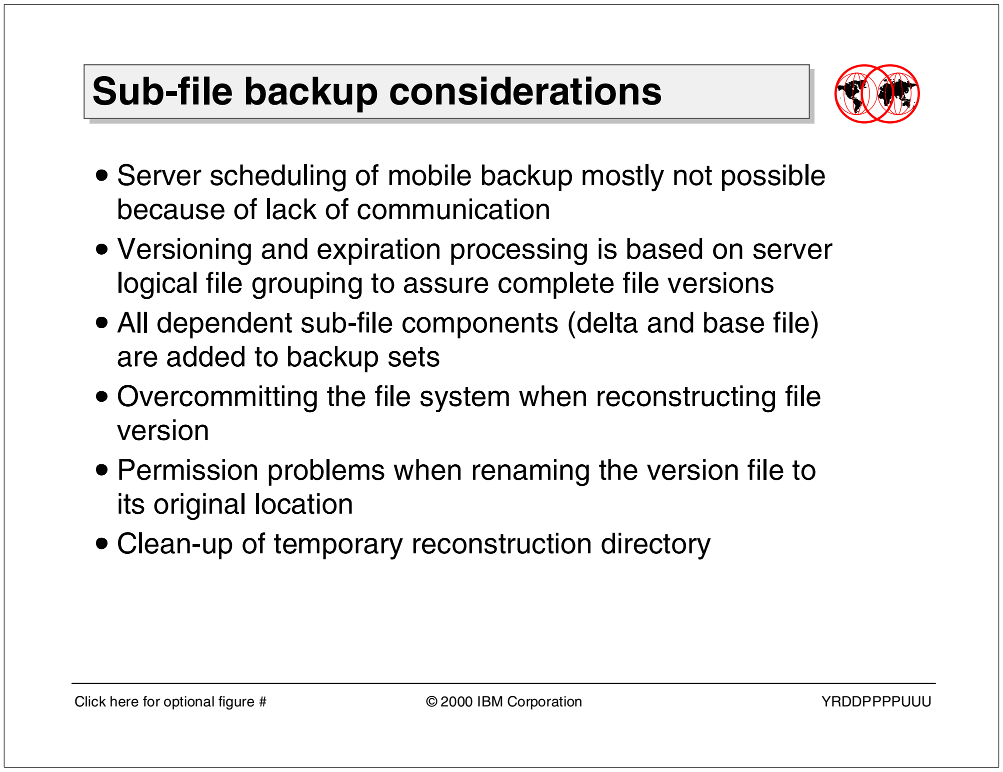

Sub-file backup considerations

- Server scheduling of mobile backup mostly not possible because of lack of communication
- Versioning and expiration processing is based on server logical file grouping to assure complete file versions
- All dependent sub-file components (delta and base file) are added to backup sets
- Overcommitting the file system when reconstructing file version
- Permission problems when renaming the version file to its original location
- Clean-up of temporary reconstruction directory

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

2.1.6 Adaptive sub-file backup considerations

When using adaptive sub-file backup technology with mobile system backup, the following should be considered.

2.1.6.1 Scheduling of mobile backup

Backups of mobile systems can be scheduled using the TSM scheduling wizard. However, because mobile system cannot be reached from the server due to the lack of communication, various other methods should be considered.

2.1.6.2 Versioning and expiration process

Because delta files are useless without the corresponding base file, the server processes the expiration of base files differently. Using the new server feature, logical file grouping, the server recognizes a base file as eligible for expiration, but does not delete the file until all its dependent delta files have expired. For more information, see Section 3.1, “Logical file grouping server support” on page 82.

2.1.6.3 Adding sub-file components to backup sets

When a delta file component is added to a backup set, the server also includes its corresponding base file with the backup set. If the base file and dependent delta files are stored on separate volumes when a backup set is created, additional volume mounts may be required to create the backup set.

2.1.6.4 Handling of non-file data

Non-file data, for example, alternate file streams or ACL information, are not processed by adaptive sub-file backup. Non-file data is restored directly from the delta file component, because there is no way to copy non-file data from one place to another on the file system.

2.1.6.5 Restore limitations

When restoring adaptive sub-file backup files to the client system, multiple file copies are stored in the temporary reconstruction directory. In extreme cases this leads to 2.7 times over-committing of the file system size. If the file system fills up, the restore will stop. Reducing the number of files restored at once is thus recommended.

Another limitation occurs when the client runs under a different user than the user who originally created the file version which the client is attempting to reconstruct. In this case, it can cause permission problems when renaming a file to its original name.

If the restore session stops for any unexpected reason, the temporary files will remain in the temporary directory. If the user does not restart the restore session, the client will not clean up this directory, and the file system can fill up.

Note

If there are no restartable restores for any node on that client OR any active restore operations going on, then the user can remove any entries in this temporary directory to avoid over-committing the file system. However, this should only be done if both of the above-mentioned conditions are true.

2.2 Client encryption support

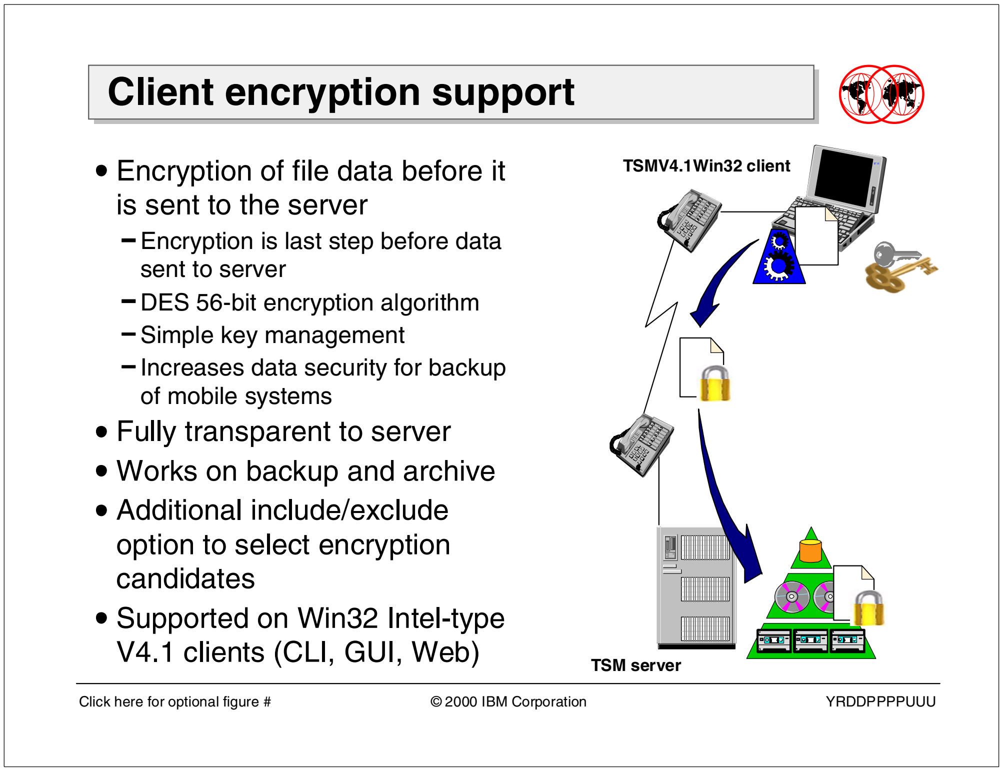

Client encryption support

- Encryption of file data before it is sent to the server
  - Encryption is last step before data sent to server
  - DES 56-bit encryption algorithm
  - Simple key management
  - Increases data security for backup of mobile systems
- Fully transparent to server
- Works on backup and archive
- Additional include/exclude option to select encryption candidates
- Supported on Win32 Intel-type V4.1 clients (CLI, GUI, Web)

Click here for optional figure #   © 2000 IBM Corporation   YRDDPPPPUUU

2.2.1 Overview

Files being backed up using a modem connection are exposed to the security hazards of public telephone lines. In order to make the backed up data more secure, the Tivoli Storage Manager client now implements an encryption function, which allows for encrypting data before it is sent to the Tivoli Storage Manager server. Not only does this help secure backed up data during transmission, but it also means that the data stored on the Tivoli Storage Manager server is encrypted and thus is unreadable by any malicious administrators.

The new function uses a standard 56-bit DES routine to provide the encryption. It allows the user to choose what files are subject to encryption via include/exclude processing. The encryption uses a very simple, key management system, which means the user must either remember the *encryption key password* during restore, or store it locally on the client system within the registry. The encryption processing is done as the last task on the client system before the data is sent to the server; other client operations like compression happen before encryption is done.

In the initial implementation, encryption support is only implemented on Windows 32-bit Intel clients. Encryption and adaptive sub-file backup can be used together or used separately. There is no inter-dependency. Encryption works for backup as well as for archive.

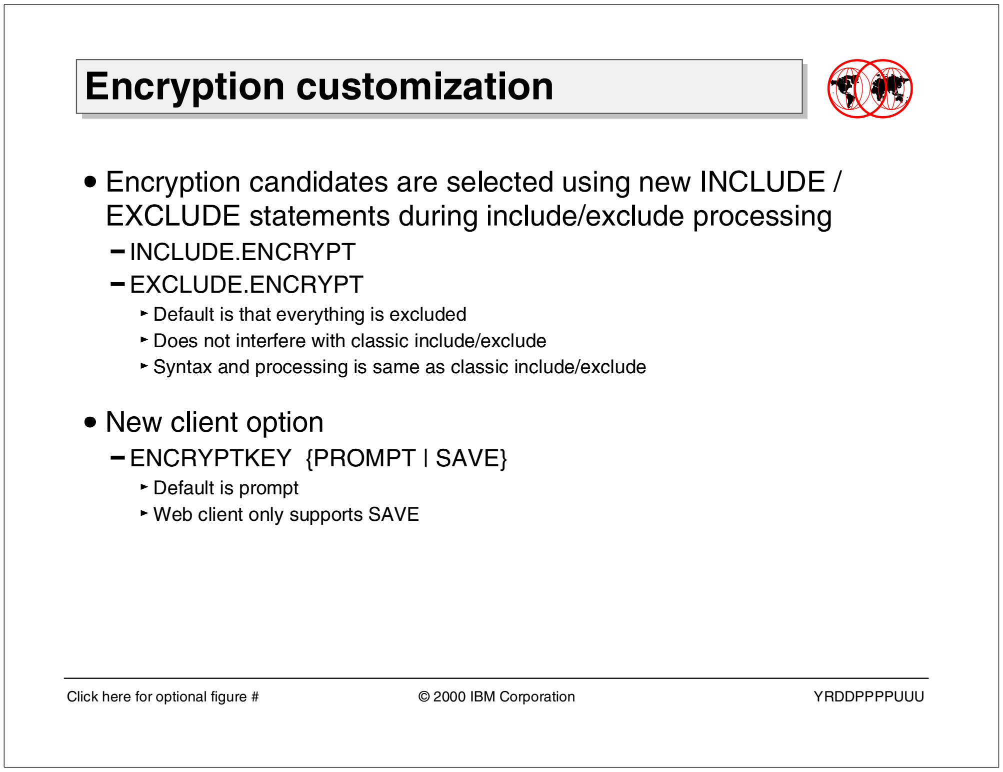

Encryption customization

- Encryption candidates are selected using new INCLUDE / EXCLUDE statements during include/exclude processing
  - INCLUDE.ENCRYPT
  - EXCLUDE.ENCRYPT
    - Default is that everything is excluded
    - Does not interfere with classic include/exclude
    - Syntax and processing is same as classic include/exclude
- New client option
  - ENCRYPTKEY {PROMPT | SAVE}
    - Default is prompt
    - Web client only supports SAVE

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

2.2.2 Encryption customizing

The complete encryption customizing is done on the client side using new include/exclude statements and client options. There is no need to enable encryption on the server.

2.2.2.1 Selection of encryption candidates

The selection of encryption candidates is done during the include/exclude processing using the two new include/exclude statements INCLUDE.ENCRYPT and EXCLUDE.ENCRYPT. Using these new statements, the user can select files and directories which have to be encrypted during the backup and archive processing. The syntax of these statements and the processing order is identical to the classical file include/exclude statements. The only difference is that by default, all client files are excluded from encryption. The user has to explicitly enable encryption for files by entering include statements in the client option file. The statements shown in Figure 6 would enable encryption for all files stored in the `c:\system32\myfiles` directory except temporary files.

```
EXCLUDE.ENCRYPT c:\system32\myfiles\*.tmp
INCLUDE.ENCRYPT c:\system32\myfiles\*
```

Figure 6. INCLUDE.ENCRYPT - EXCLUDE.ENCRYPT example

There is no dependency or interference between the classical file include/exclude statements and the encryption include/exclude statements. However, the file include/exclude statements will be used first to determine what files are eligible for backup. If a file is excluded from backup, it makes no sense to include it for encryption, because it will not backed up in the first place.

2.2.2.2 Saving the encryption key password

The new option ENCRYPTKEY allows the user to control whether the encryption key password has to be entered for every backup or restore, or whether it is be saved in the registry. The default value is PROMPT, which means the user is always prompted for an encryption key password. Setting the option to SAVE will trigger the client to store the password in the client registry. As soon the password is stored, it will be used for all backups and restores.

Note

The Web client only supports the SAVE mode, which means that the password has to be stored on the client system. It will not remotely prompt for an encryption key password.

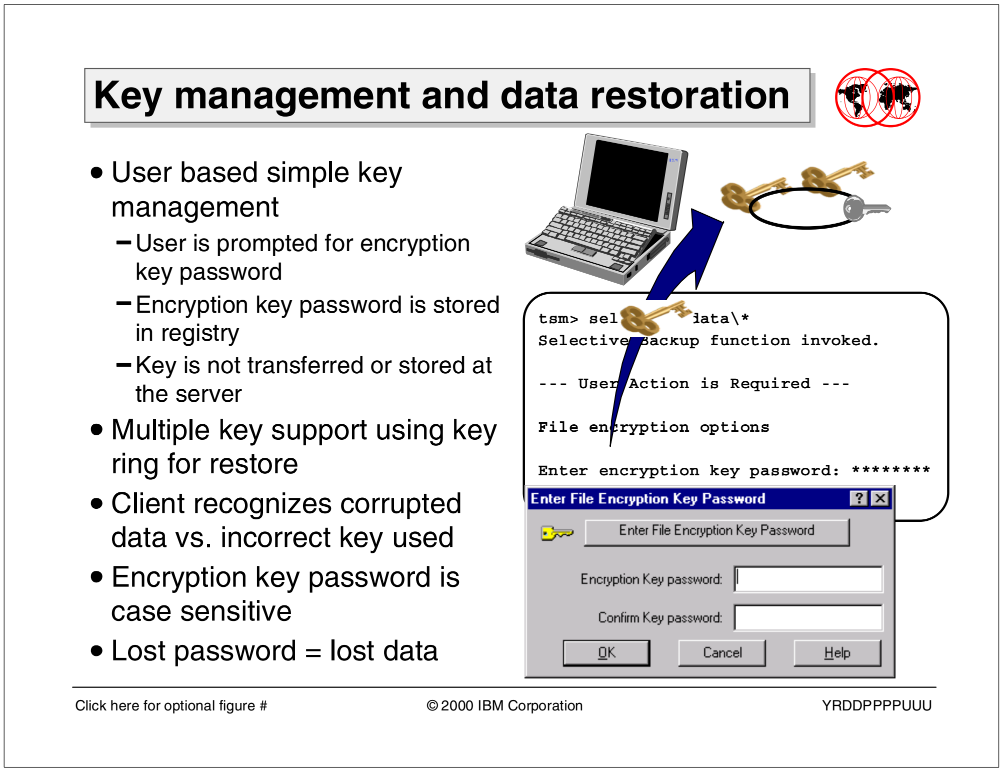

Key management and data restoration

- User based simple key management
  - User is prompted for encryption key password
  - Encryption key password is stored in registry
  - Key is not transferred or stored at the server
- Multiple key support using key ring for restore
- Client recognizes corrupted data vs. incorrect key used
- Encryption key password is case sensitive
- Lost password = lost data

```
tsm> sel          data\*
Selective Backup function invoked.

--- User Action is Required ---

File encryption options

Enter encryption key password: ********
```

Enter File Encryption Key Password
Enter File Encryption Key Password
Encryption Key password:
Confirm Key password:
OK
Cancel
Help

Click here for optional figure #   © 2000 IBM Corporation   YRDDPPPPUUU

2.2.3 Key management and data restoration

Tivoli Storage Manager client encryption support uses a very simple key management method. The key is only used at the client, it is not transferred or stored at the server. The encryption key password is either provided every time a file is encrypted or decrypted via a prompt, or it is stored in the client's registry.

Multiple keys can be used, but only the key entered when the ENCRYPTKEY client option was set to SAVE is stored in the registry. As shown in Figure 7 during restores, the user will be prompted for the proper decryption key password. The client will create the key from it and keep the key on an internal *key-ring cache* as long as the client program is running. Files are only restored if the matching key is found on the key-ring, or if the user can supply a proper decryption key password after being prompted.

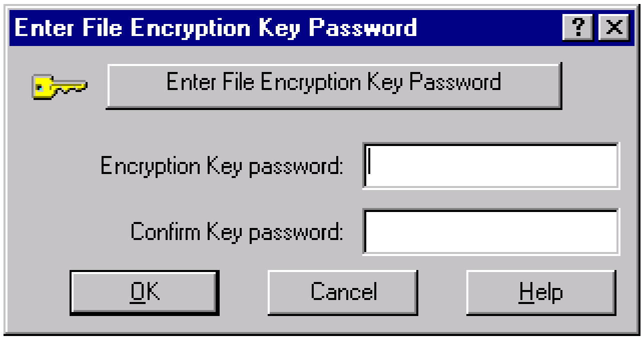

Enter File Encryption Key Password
?
Enter File Encryption Key Password
Encryption Key password:
Confirm Key password:
OK
Cancel
Help

Figure 7. Prompt for encryption key password

The encryption function is able to differentiate between invalid data and data that was decrypted with an incorrect user key. Information stored in the file stream on the server indicates that encryption was used and which type. The key-ring on restore caches potential user keys and uses the information in the file stream to determine if the encryption key password supplied by user is incorrect.

Unlike the Tivoli Storage Manager user password, the encryption key password is case-sensitive. If the password is lost or forgotten, the encrypted data cannot be decrypted, which means that the data is lost.

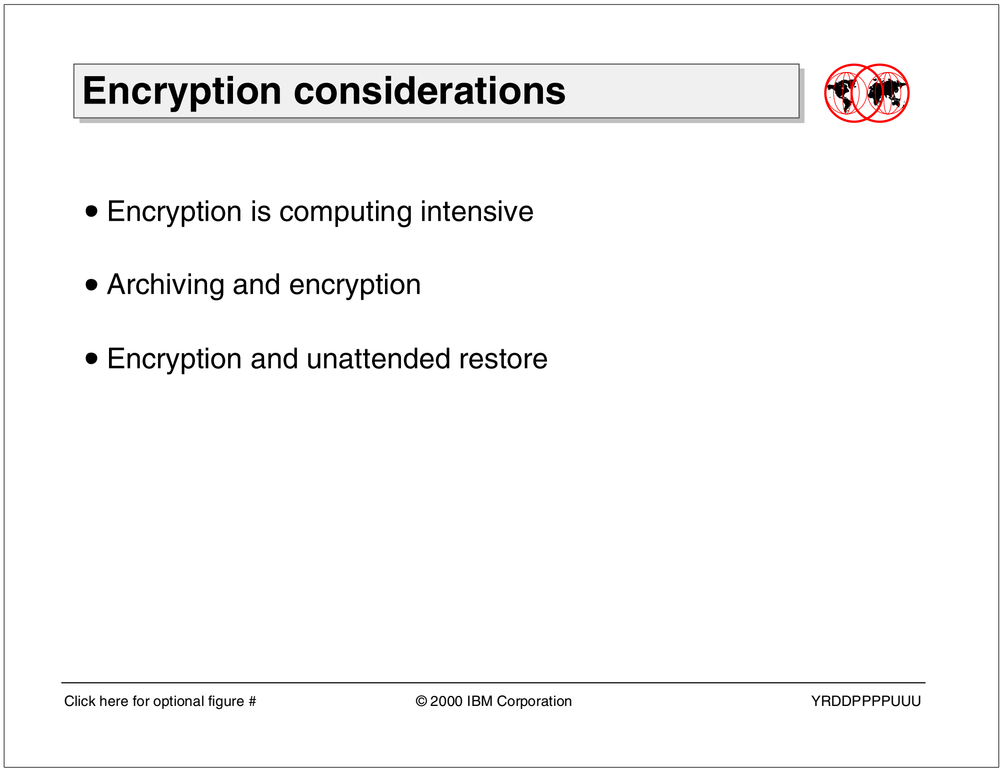

Encryption considerations

- Encryption is computing intensive
- Archiving and encryption
- Encryption and unattended restore

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

2.2.4 Encryption considerations

When using encryption while backing up data, the following should be taken into consideration.

2.2.4.1 Encryption is computing intensive

Because of the nature of the encryption process, it will add a lot of additional computing requirements on the client CPU. The user should very carefully decide what data items really need to be encrypted and control them using the include/exclude statements. The administrator has the option to overwrite client selections using client options sets.

2.2.4.2 Encryption and archiving

Especially in the case of long-term archiving of data, problems can be occur with the simple key management scheme. If data is archived using encryption, organizational rules have to make sure that the encryption key password remains available for retrieve. Cyclic password changes or no external password management can cause a situation where data cannot be successfully retrieved.

2.2.4.3 Encryption and unattended restore

When restoring data using the scheduling function, it can happen that the files are not restored, because the needed encryption key password is not stored in the registry (PROMPT mode), or the restored files are encrypted with a different key than the stored encryption key password.

2.3 Client Configuration Wizard

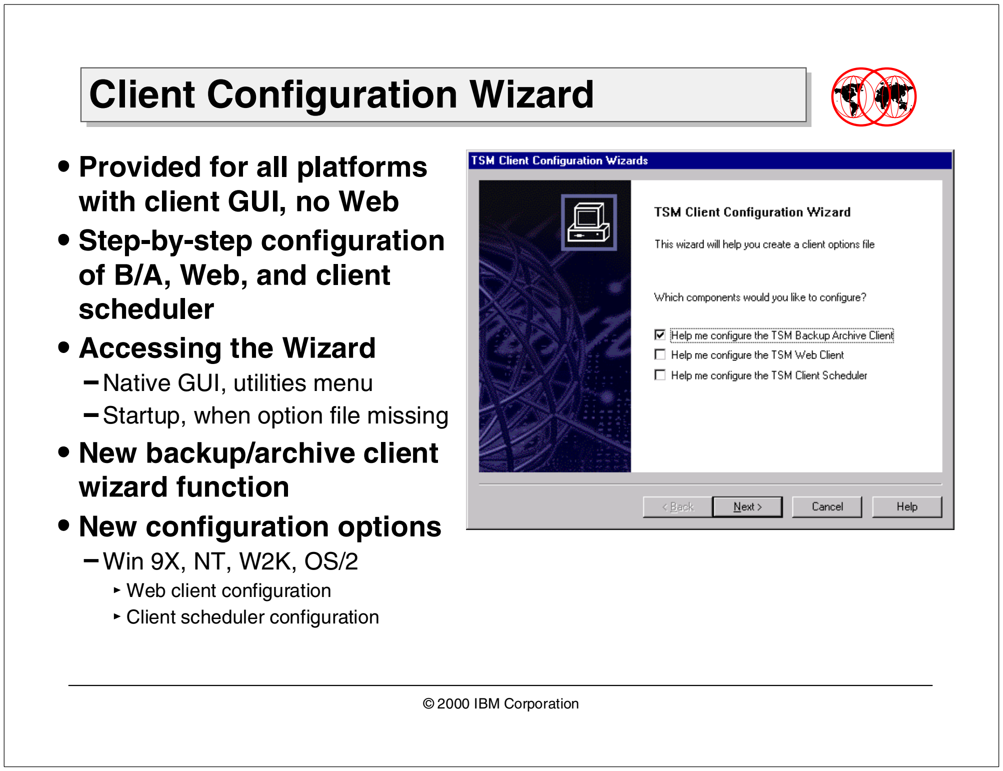

Client Configuration Wizard

- **Provided for all platforms with client GUI, no Web**
- **Step-by-step configuration of B/A, Web, and client scheduler**
- **Accessing the Wizard**
  - Native GUI, utilities menu
  - Startup, when option file missing
- **New backup/archive client wizard function**
- **New configuration options**
  - Win 9X, NT, W2K, OS/2
    - Web client configuration
    - Client scheduler configuration

TSM Client Configuration Wizards

TSM Client Configuration Wizard

This wizard will help you create a client options file

Which components would you like to configure?

☑ Help me configure the TSM Backup Archive Client
☐ Help me configure the TSM Web Client
☐ Help me configure the TSM Client Scheduler

© 2000 IBM Corporation

The wizard function available since TSM client 3.7.1 has the ability to help users in performing the configuration of a TSM client. The new 3.7.2 TSM client code improves the wizard with new functions:

- New backup-archive client configuration wizard functions
- New Web client configuration wizard (Win9X, NT, W2K and OS/2)
- New client scheduler configuration wizard (Win9X, NT, W2K and OS/2)

The client services can be installed, updated, and removed using the easy-to-use graphical wizard panels. These new wizards are available for those platforms with client GUI, and some of them are available for only specific client platforms at this time.

The wizards are started automatically at the first startup of the backup-archive client’s GUI, when no option file is present. The wizards can also be manually started by clicking on the Utilities menu in the Client GUI main panel.

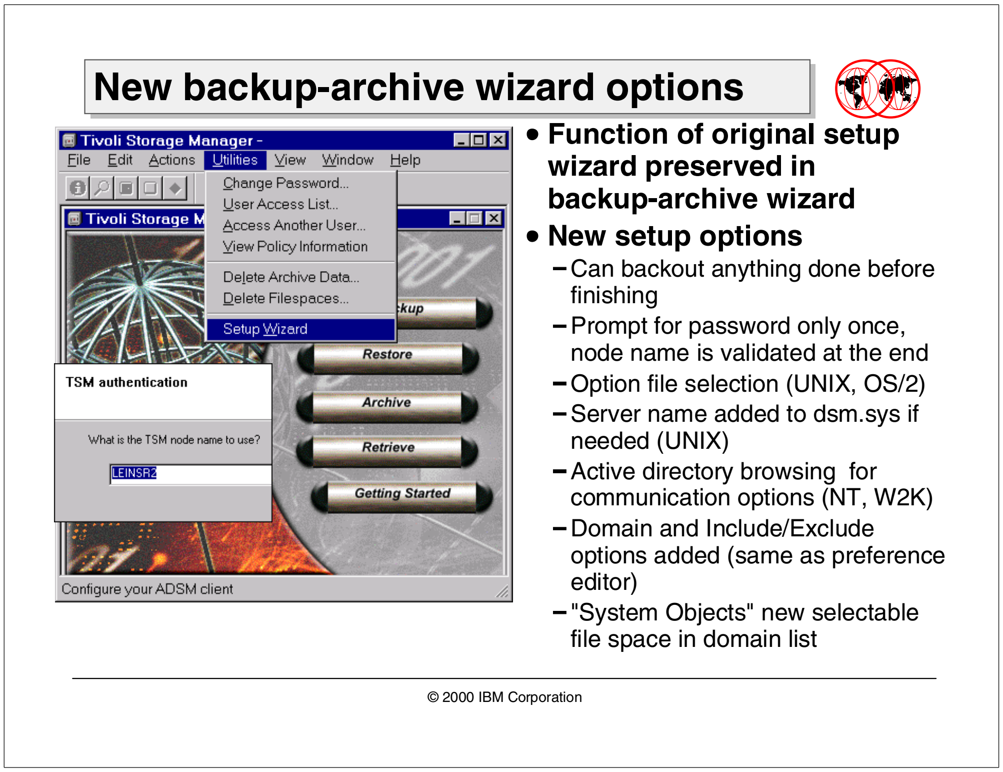

New backup-archive wizard options

TSM authentication

What is the TSM node name to use? LEINSR2

- **Function of original setup wizard preserved in backup-archive wizard**
- **New setup options**
  - Can backout anything done before finishing
  - Prompt for password only once, node name is validated at the end
  - Option file selection (UNIX, OS/2)
  - Server name added to dsm.sys if needed (UNIX)
  - Active directory browsing for communication options (NT, W2K)
  - Domain and Include/Exclude options added (same as preference editor)
  - "System Objects" new selectable file space in domain list

2.3.1 Backup-archive client configuration wizard overview

The backup-archive client configuration wizard is the successor of the V3.7.1 initial client configuration wizard. All functions present in this wizard are preserved in the new backup-archive client configuration wizard.

When the Client GUI is launched for the first time and does not find an option file (dsm.sys for UNIX clients, dsm.opt for all other clients), a new one is created, and the setup wizard guides users through the configuration process.

The user is asked to perform the initial task to create, update, or import the client option file (Figure 8 shows the three options).

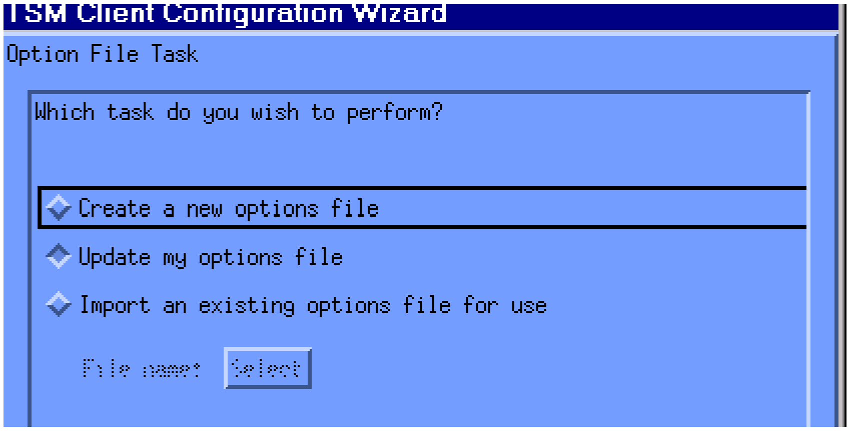

TSM Client Configuration Wizard

Option File Task

Which task do you wish to perform?

◇ Create a new options file

◇ Update my options file

◇ Import an existing options file for use

File name: Select

Figure 8. Option file specification

The wizard assigns a default node name depending on the platform. For example it can be the machine name for NT and W2K, or the host name for UNIX clients. Users can accept or change the node name. The validation is done at the end of the configuration. The node name must be a registered node, or the server must allow open registration, otherwise you will receive an error message.

Next, the communication method and server address (based on communication method) are configured. The domain list and include/exclude list editing is the same as for the preference editor. This feature now gives users the ability to determine which disks or file systems are to be considered as domain resources, and it defines local policy for include/exclude files and directories. Default values are offered.

Finally, a prompt is given to contact the TSM server. The `userid` and `password` must be provided. The option `passwordacces=generate` is automatically set into the option file.

2.3.2 New function of the backup-archive client configuration wizard

Additional functions of the new wizard are:

- For UNIX only: the server name will be added to dsm.sys if necessary.
- The user is no longer prompted for a password unless necessary. It is validated against the server near the end of the interaction.
- Active Directory browsing is a specific NT and W2K capability. The server name address browse button is enabled only if the system can find a suitable Windows 2000 domain (see Figure 9).

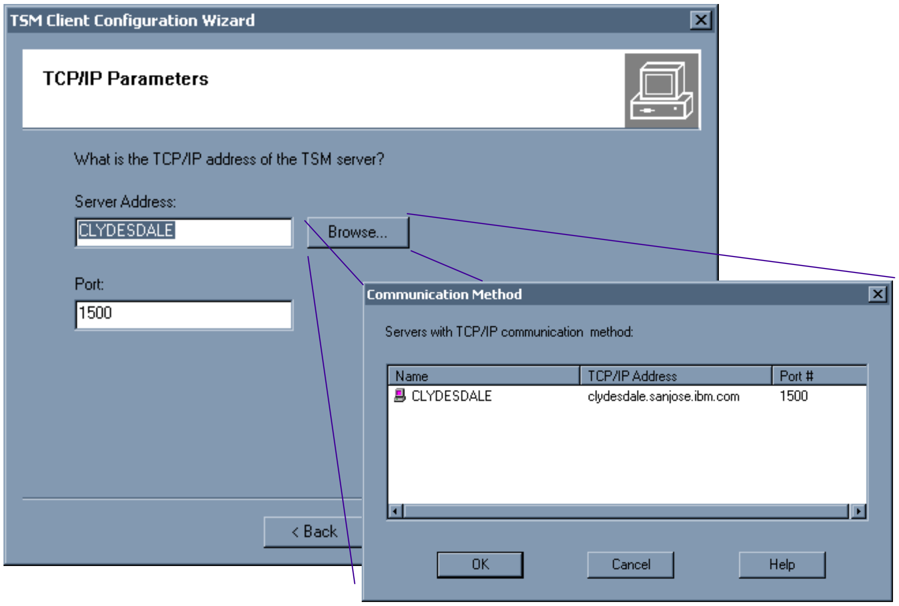

TSM Client Configuration Wizard

TCP/IP Parameters

What is the TCP/IP address of the TSM server?

Server Address: CLYDESDALE

Port: 1500

Communication Method

Servers with TCP/IP communication method:

Name	TCP/IP Address	Port #
CLYDESDALE	clydesdale.sanjose.ibm.com	1500

Figure 9. Active Directory Browsing

- In the domain list there is a new entry SYSTEM OBJECTS for NT and W2K clients. By default, SYSTEM OBJECTS are included (see Figure 10).

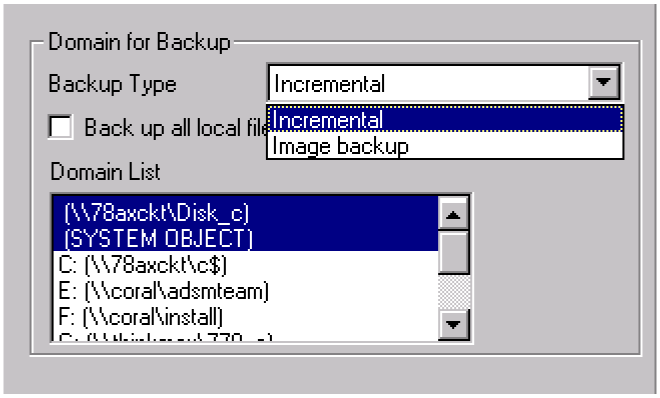

Domain for Backup

Backup Type: Incremental

□ Back up all local fil

Domain List

(\\78axckt\Disk_c)
(SYSTEM OBJECT)
C: (\\78axckt\c$)
E: (\\coral\adsmteam)
F: (\\coral\install)

Figure 10. New entry in domain list

- All changes take effect when the user hits the Finish button, and any mistakes will display an error message: `Data was written to the option file, however there was at least one error that must be fixed before we can connect to a server`. At any time the user can cancel the configuration without changing the option files.

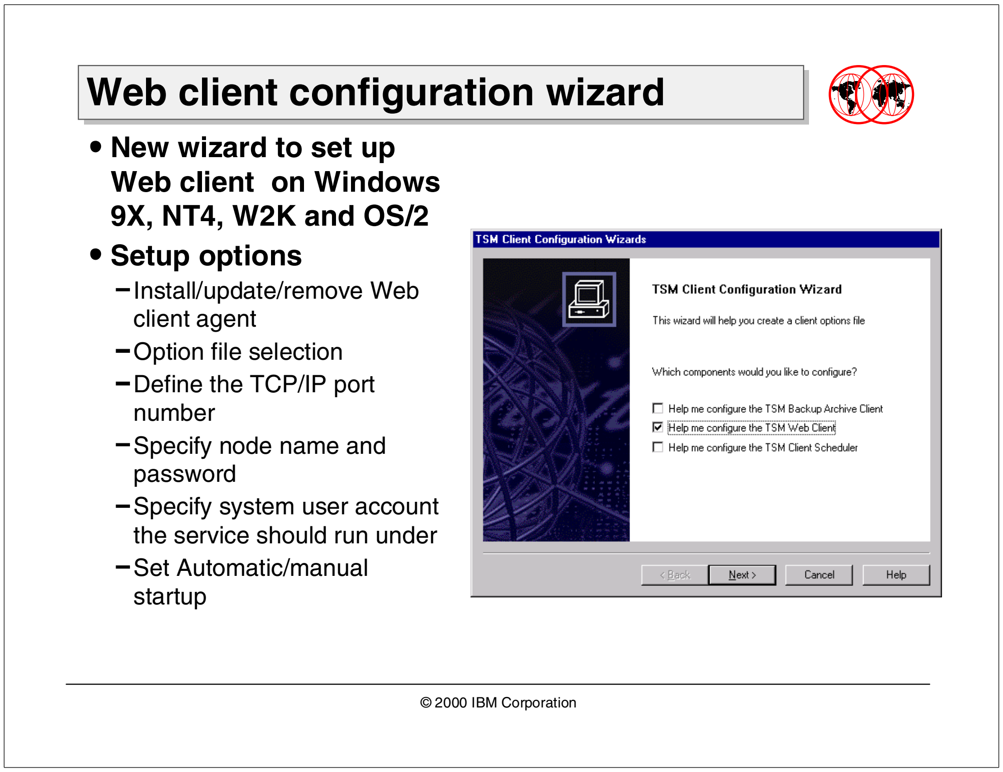

Web client configuration wizard

- New wizard to set up Web client on Windows 9X, NT4, W2K and OS/2
- Setup options
  - Install/update/remove Web client agent
  - Option file selection
  - Define the TCP/IP port number
  - Specify node name and password
  - Specify system user account the service should run under
  - Set Automatic/manual startup

TSM Client Configuration Wizards

TSM Client Configuration Wizard

This wizard will help you create a client options file

Which components would you like to configure?

☐ Help me configure the TSM Backup Archive Client
☑ Help me configure the TSM Web Client
☐ Help me configure the TSM Client Scheduler

© 2000 IBM Corporation

2.3.3 Web client configuration wizard

The Web client is a Java-based remote control client that has the same functionality as the native GUI with a few limitations. On Windows and OS/2 clients, you can define and update the Web client using the *Web client configuration wizard*.

On Windows NT or Windows 2000, the following configuration steps are available:

1. The wizard is simply launched from the setup wizard launch window by selecting the "Help me configure the TSM Web Client" option. The launch window either starts automatically at the first start of the backup-archive client GUI, or can be started from the Utilities menu of the backup-archive client GUI.
2. The next step allows you to install, update, or remove a Web client agent (see Figure 11).
3. You can specify whether the Web Client service should be started automatically when Windows boots, or whether it has to be manually started (default setting).

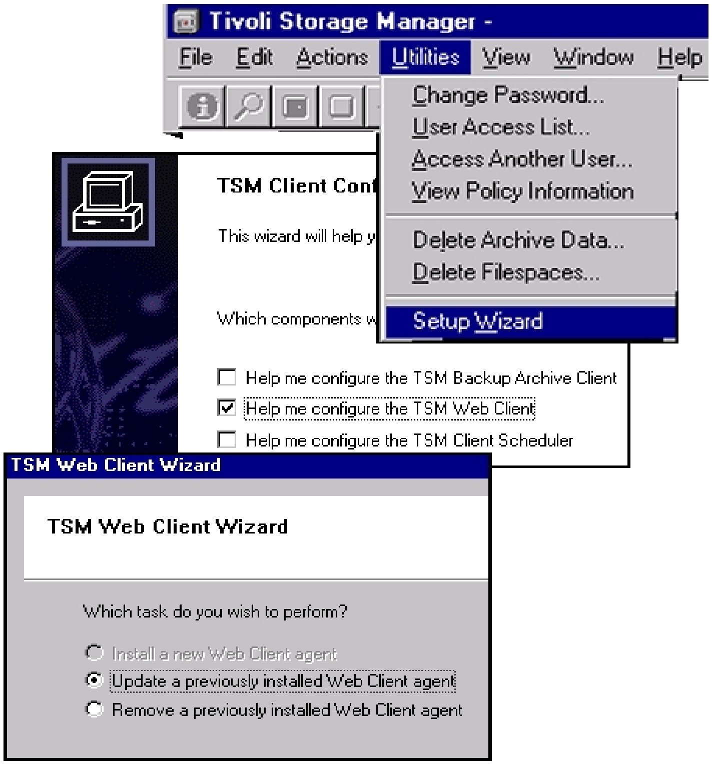

TSM Client Conf

This wizard will help y

Which components w

□ Help me configure the TSM Backup Archive Client

☑ Help me configure the TSM Web Client

□ Help me configure the TSM Client Scheduler

TSM Web Client Wizard

TSM Web Client Wizard

Which task do you wish to perform?

○ Install a new Web Client agent

● Update a previously installed Web Client agent

○ Remove a previously installed Web Client agent

Figure 11. Defining a Web Client

4. Defining the option file to be used by the Web client is the next step. You can accept the default option file already used by the backup-archive client GUI, or browse the disks to look for another previously created file.
5. A port number for TCP/P communications must be defined. 1581 is the default port. Two other buttons are provided to grant or revoke the access privilege to the Web client.
6. A node name and password must be provided. This node must be already registered in Tivoli Storage Manager server, otherwise you get an error message (even if the configuration process says: “Web Client successfully installed”).
7. In the last panel you must provide an operating system user account under which the Web client agent should be executed. The default user is the system account user. In the same panel you can specify how the service should be started.

The corresponding panels with setup steps 4 to 7 are shown in Figure 12.

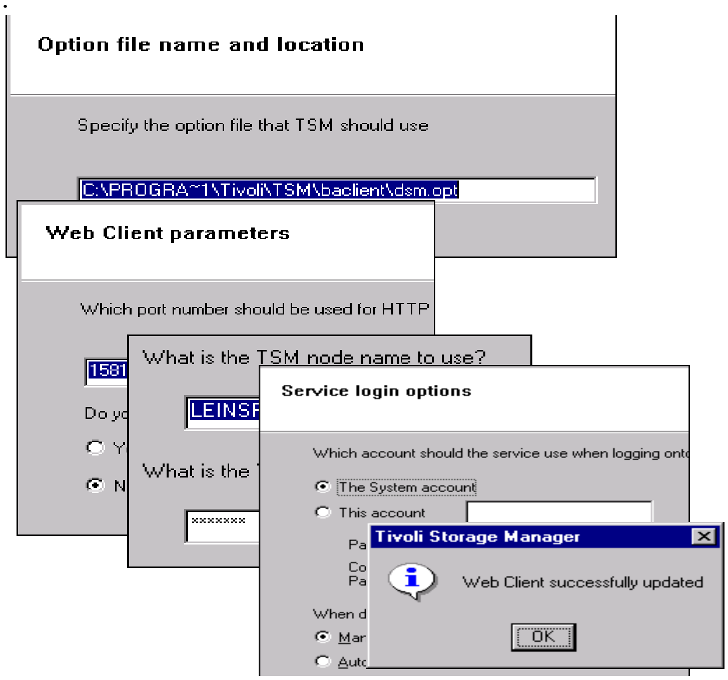

Option file name and location

Specify the option file that TSM should use

C:\PROGRA~1\Tivoli\TSM\baclient\dsm.opt

Web Client parameters

Which port number should be used for HTTP

1581

Do yo

○ Y

● N

What is the TSM node name to use?

LEINSF

What is the T

****

Service login options

Which account should the service use when logging ont

● The System account

○ This account

Pa

Co
Pa

When d

● Man

○ Auto

Tivoli Storage Manager

Web Client successfully updated

OK

Figure 12. Completing the configuration of a Web Client

Once defined, a Web Client can be updated or removed using the same Wizard panels.

The `dsmcutil.exe` command is still available in the `~TSM_directory\baclient` and can also be used to manage the Web Client services.

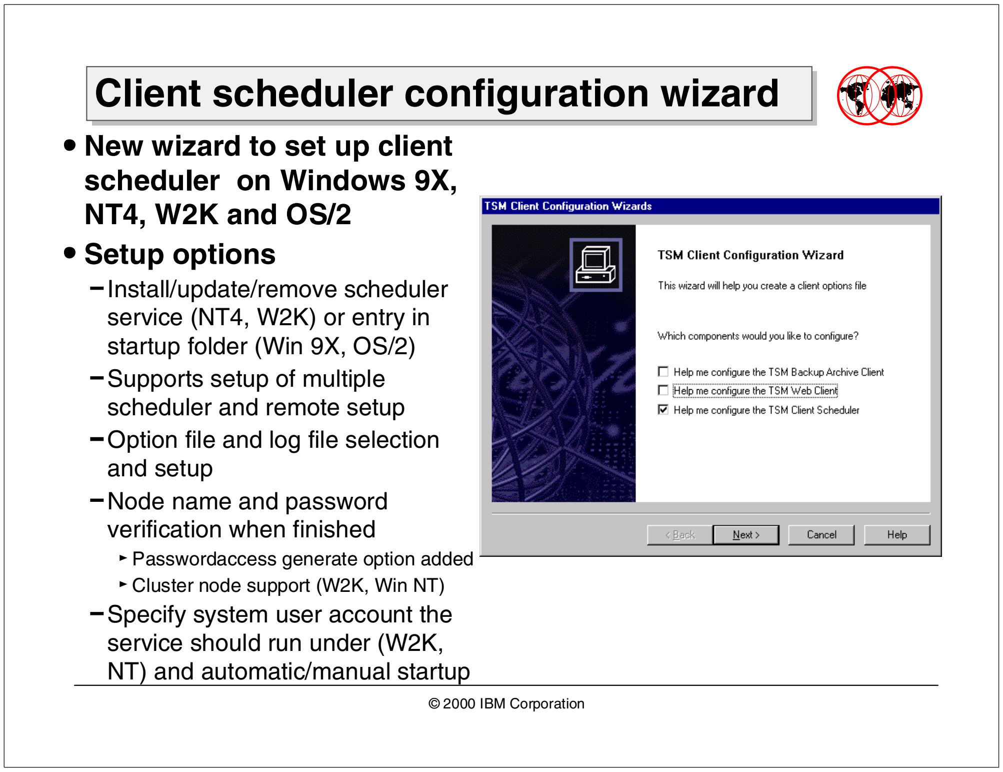

Client scheduler configuration wizard

- **New wizard to set up client scheduler on Windows 9X, NT4, W2K and OS/2**
- **Setup options**
  - Install/update/remove scheduler service (NT4, W2K) or entry in startup folder (Win 9X, OS/2)
  - Supports setup of multiple scheduler and remote setup
  - Option file and log file selection and setup
  - Node name and password verification when finished
    - Passwordaccess generate option added
    - Cluster node support (W2K, Win NT)
  - Specify system user account the service should run under (W2K, NT) and automatic/manual startup

TSM Client Configuration Wizards

TSM Client Configuration Wizard

This wizard will help you create a client options file

Which components would you like to configure?

☐ Help me configure the TSM Backup Archive Client

☐ Help me configure the TSM Web Client

☑ Help me configure the TSM Client Scheduler

© 2000 IBM Corporation

2.3.4 Client scheduler configuration wizard

The *Client scheduler configuration wizard* for Windows NT, Windows 2000, Windows 95/98, and OS/2, is available to help users in defining the scheduler services.

2.3.4.1 Install a scheduler service

Using this wizard it is possible to install, update and remove a scheduler service. Some platforms allow more than one scheduler. A unique name has to be used to define the scheduler service (see Figure 13). If more than one is required, then different names have to be used on the name page on the install panel.

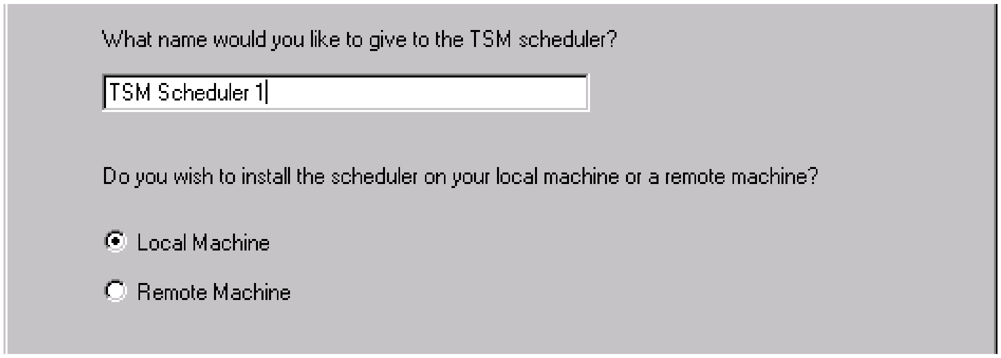

What name would you like to give to the TSM scheduler?

TSM Scheduler 1

Do you wish to install the scheduler on your local machine or a remote machine?

● Local Machine

○ Remote Machine

Figure 13. Scheduler name definition

2.3.4.2 Defining a scheduler service for a remote machine

It is also possible to define a scheduler service for a remote machine. A panel with a browser push-button gives the ability to navigate on the network to find the requested machine (see Figure 14).

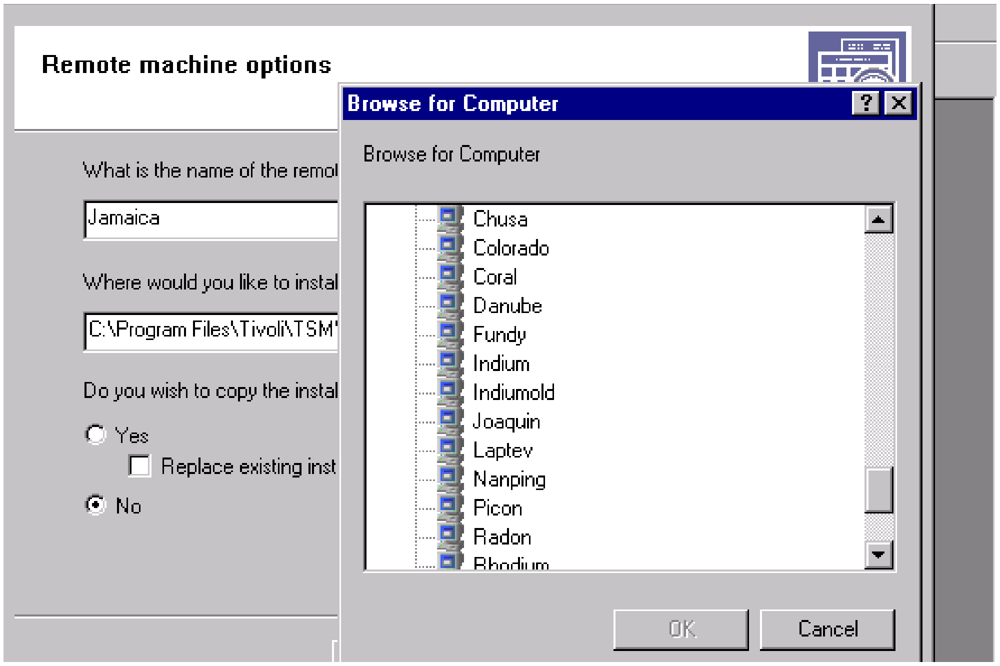

Remote machine options

What is the name of the remot

Jamaica

Where would you like to instal

C:\Program Files\Tivoli\TSM\

Do you wish to copy the instal

○ Yes

□ Replace existing inst

● No

Figure 14. Remote machine option

If more than one client scheduler is required, then different names have to be used on the name page on install panel. Defining multiple schedulers gives the ability to use different client option files. Thus, users can save data with different node names, and can include/exclude specifications (management class binding) or user account authority.

2.3.4.3 Update or remove a scheduler service

The update and remove options give the list of the defined schedulers and allow for the update or removal of the scheduler service (see Figure 15).

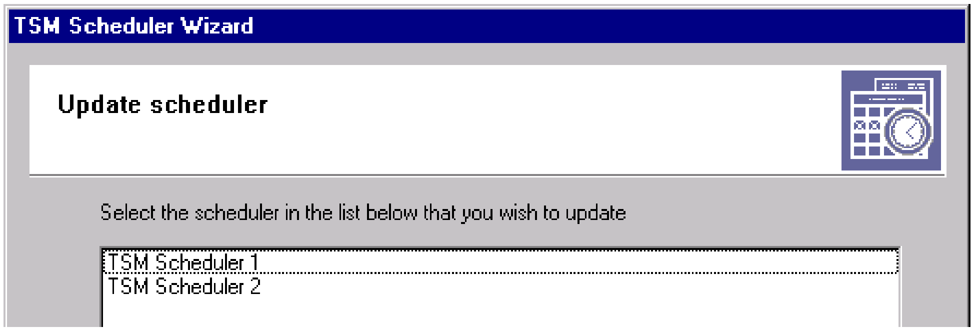

TSM Scheduler Wizard

Update scheduler

Select the scheduler in the list below that you wish to update

TSM Scheduler 1
TSM Scheduler 2

Figure 15. Selecting a scheduler service

2.3.4.4 Specifying a user account

The scheduler service can be registered using the default system account or a specific account. It may be helpful to join a scheduler to a particular user with specific application authority (see Figure 16).

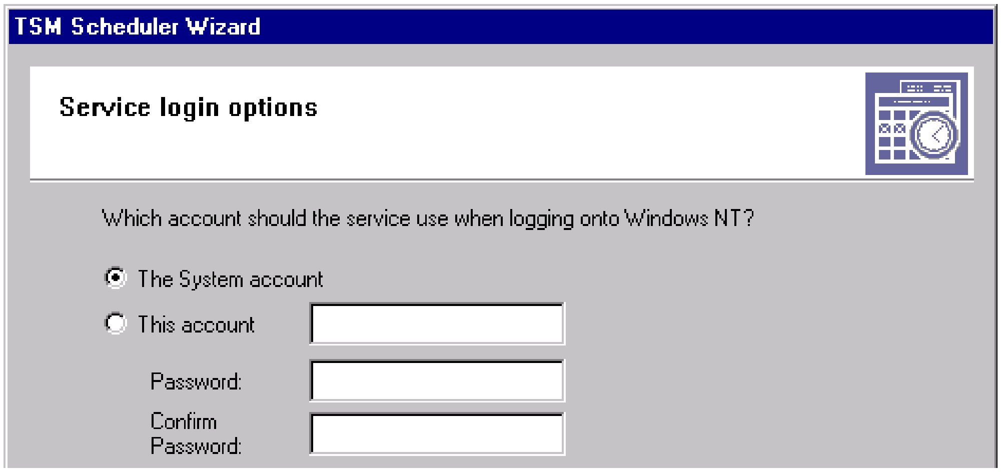

TSM Scheduler Wizard

Service login options

Which account should the service use when logging onto Windows NT?

● The System account

○ This account

Password:

Confirm Password:

Figure 16. User account definition

2.3.4.5 Cluster node detecting

If defining a scheduler service on an MS Cluster Node, the wizard detects this special configuration and helps the user to fill in the proper definitions.

2.3.4.6 Specifying log files

A service logging page is provided to specify name and location for the log files. (see Figure 17).

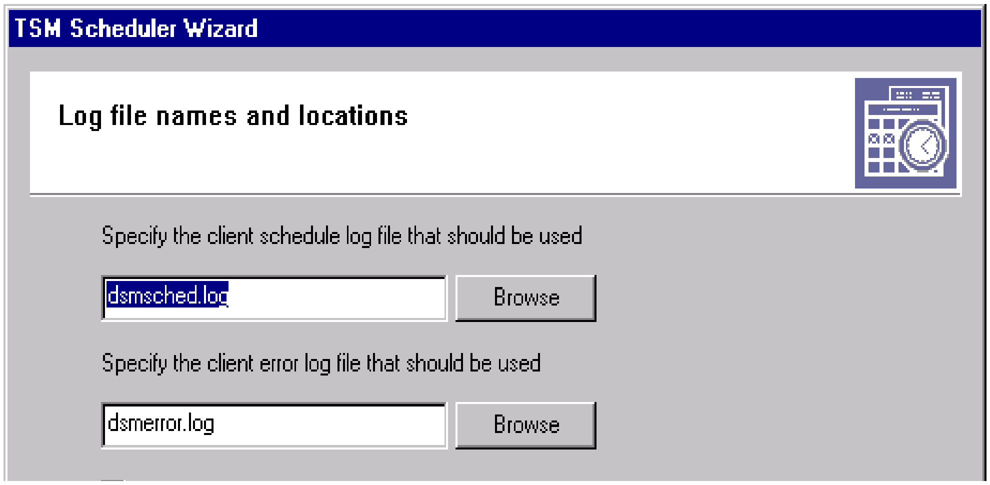

TSM Scheduler Wizard

Log file names and locations

Specify the client schedule log file that should be used

dsmsched.log

Specify the client error log file that should be used

dsmerror.log

Figure 17. Specifying the log files

The program **dsmcutil** is still available and can be used in a command line DOS prompt to install, update, or remove any scheduler service.

2.4 Additional tape device support for backup sets

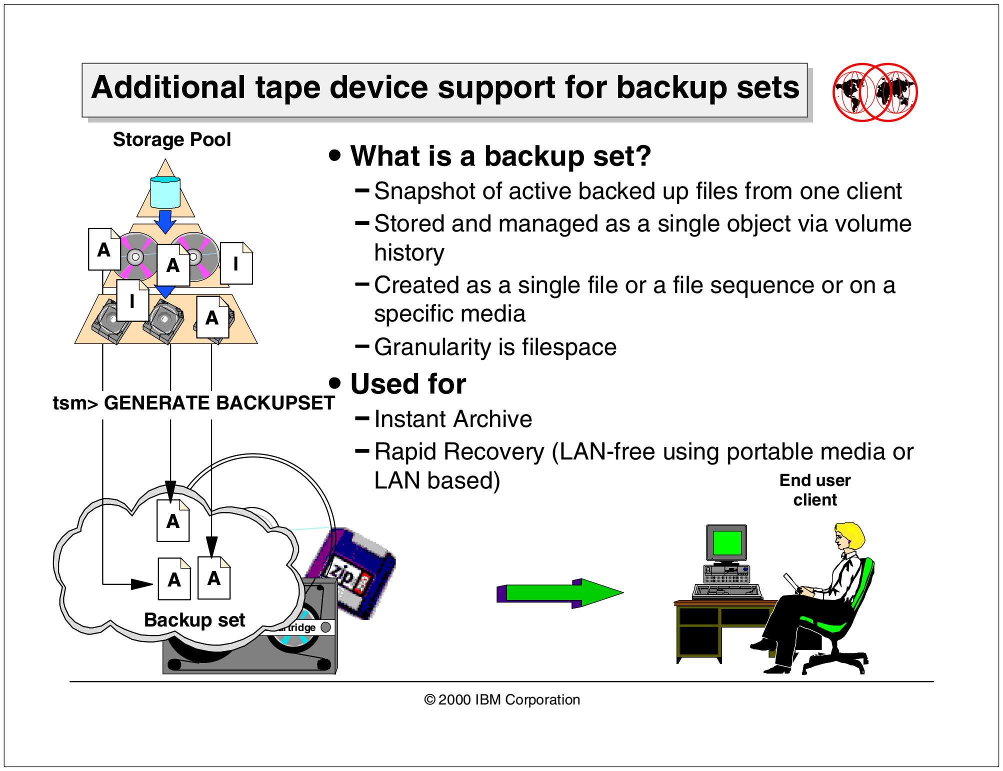

Additional tape device support for backup sets

- **What is a backup set?**
  - Snapshot of active backed up files from one client
  - Stored and managed as a single object via volume history
  - Created as a single file or a file sequence or on a specific media
  - Granularity is filespace
- **Used for**
  - Instant Archive
  - Rapid Recovery (LAN-free using portable media or LAN based)

© 2000 IBM Corporation

With Tivoli Storage Manager Version 3.7.1, the feature of backup sets was introduced. With Version 3.7.2 of the backup-archive clients, new tape device support for local backup set restore was introduced.

The following section will give a short introduction to backup sets and discuss the new tape device support for local backup set restore on the Tivoli Storage Manager backup-archive client (GUI, CLI and Web).

2.4.1 What is a backup set?

Tivoli Storage Manager is able to create a snapshot or an independent copy of a client’s latest backup from the server storage pools onto sequential media and to manage it as a single object. This object is called *a backup set*.

The client machine is not involved in the generation process. All backed up client files are already in the server storage pools, so no data is sent across the network. The backup set may be interactively created by an administrator or it can be scheduled as a server process.

A backup set may contain all the file spaces of a node or just one of them. So, the filespace is the smallest backup set possible. By default Tivoli Storage Manager will include all the filespaces of a node.

A backup set can be intended as a snapshot of the whole active backed up data or part of it. Backup sets are created and stored on the TSM server. Their volumes are tracked in the volume history file and will expire as defined by the retention option.

2.4.2 Usage of backup sets

The two main usages of backup sets are for:

- Instant archive
- Rapid recovery

2.4.2.1 Instant archive

Administrators can generate multiple copies of backup sets that correspond to some point-in-time backup image of a node and can be retained for various time periods. This satisfies the requirement to retain long term storage of periodic backups without having data sent over the network again.

Instant archive is the term used to describe these images. Every backup set contains all data and meta-data required to restore the client files and directories, but the server will consider it a single object. It means that individual files are not indexed and tracked in the TSM database. An instant archive can be done occasionally, or scheduled on regular base. In the example above, multiple backup set generation is performed for different periods of retention. Every image is a frozen copy of the client machine and can be utilized to recover a client or to recreate certain desired conditions as well.

2.4.2.2 Rapid recovery

A backup set can be used to recover an entire client system at the most current state without utilizing the server's database. Because the backup set is a self-describing image of the client, it can be used independently to recover the client system. Generally, two methods of recovery can be distinguished:

- LAN-based restore
- LAN-free restore

LAN-based restore

The restore of a backup set, unlike a normal data restore, is performed as a no-query restore operation, without confirmation of verbs. Using the backup-archive client, the backup set can be queried and selected to be restored from the server. A response is required for the destination path, then data is sent over the network.

LAN-free restore of backup sets

The local option will give the ability to manage the backup set using locally attached devices and specific media, called *portable media*. The portable media can be any kind of media created on compatible devices for client and server machines, for example CD-ROM, JAZ, ZIP, or tape. The restore process will be as fast as allowed by the device. A client-server session has to be established only if using the client backup-archive GUI. With the CLI, the restore can be performed in a stand-alone machine even in the absence of a network connection.

The server and client machine can be different platforms; so a backup set generated on an AIX server can be moved and locally restored in a Windows NT client, and vice versa. The major requirements are that server and client must use the same tape device, and the formatting of the media needs to be compatible.

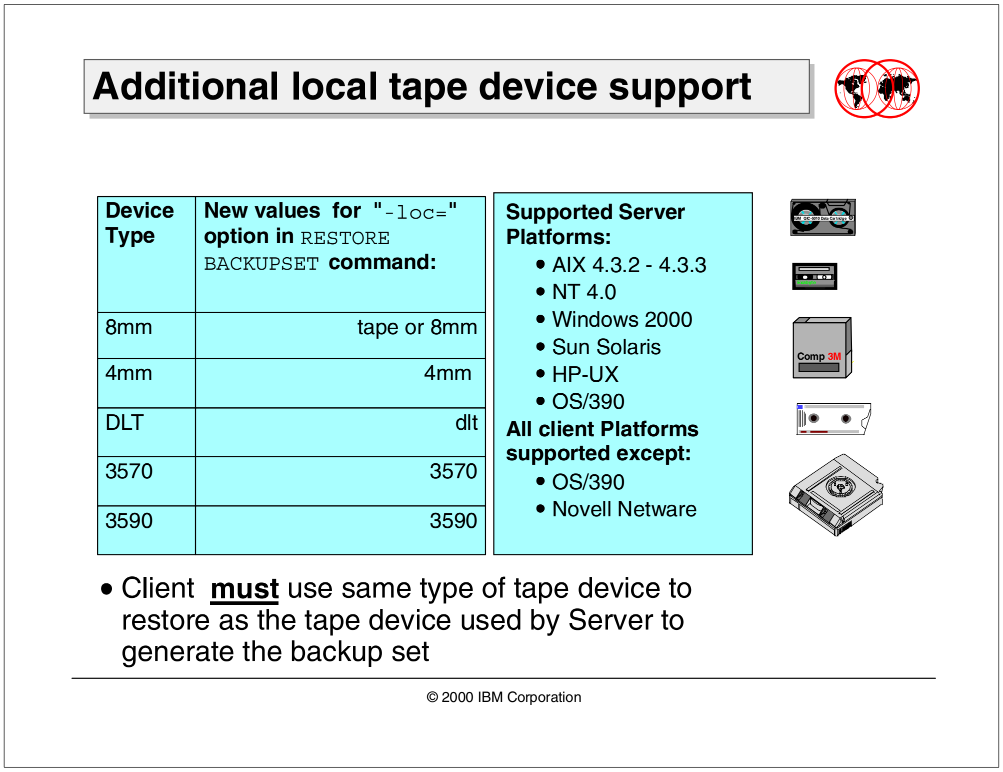

Additional local tape device support

Device Type	New values for "-loc=" option in `RESTORE BACKUPSET` command:
8mm	tape or 8mm
4mm	4mm
DLT	dlt
3570	3570
3590	3590

Supported Server Platforms:

- AIX 4.3.2 - 4.3.3
- NT 4.0
- Windows 2000
- Sun Solaris
- HP-UX
- OS/390

All client Platforms supported except:

- OS/390
- Novell Netware

- Client **must** use same type of tape device to restore as the tape device used by Server to generate the backup set

2.4.3 Additional tape device support

In Tivoli Storage Manager 3.7.1, only the AIX client platform provided tape support, and that was only for 8mm tapes. The 3.7.3 TSM client extends tape support to most client platforms (except S/390 and Novel Netware), and adds additional tape devices to AIX.

The newly supported device types are: 8mm, 4mm, DLT, 3570 and 3590. Such devices must be available on both server and client machines to ensure that a backup set generated on the server can be queried and restored on the client machine. The device class parameter used in the generate command, now can be any of the tape device classes defined on the server machine. However, the same device with the same characteristics must be available on the client machine in order to be able to restore it locally.

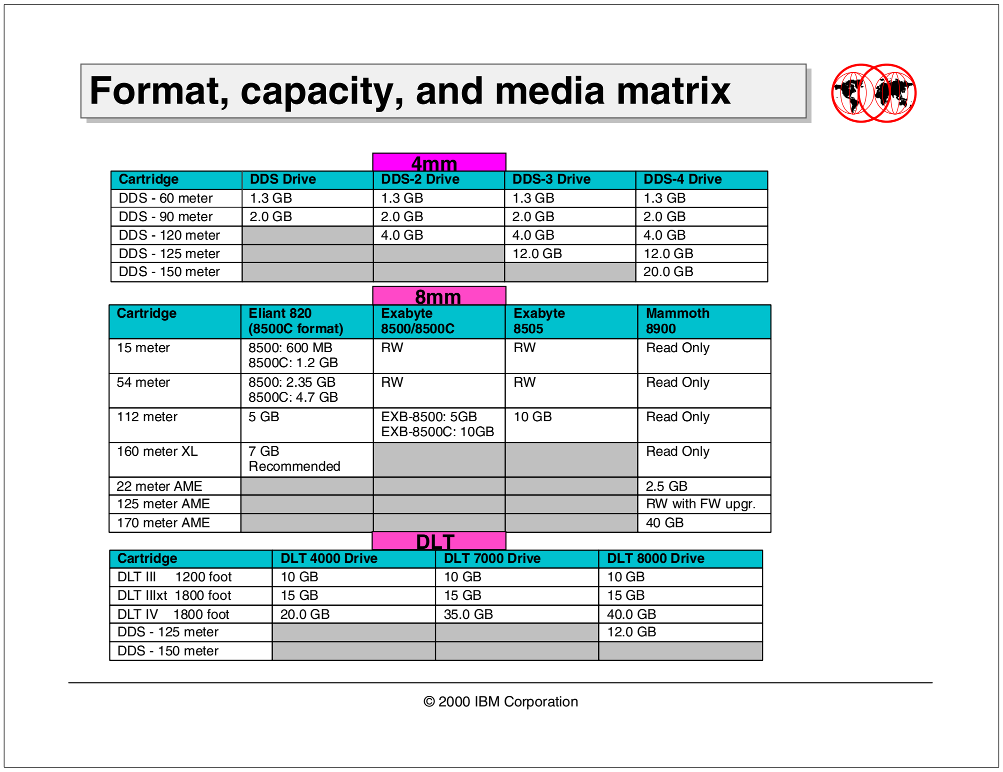

Format, capacity, and media matrix

4mm

Cartridge	DDS Drive	DDS-2 Drive	DDS-3 Drive	DDS-4 Drive
DDS - 60 meter	1.3 GB	1.3 GB	1.3 GB	1.3 GB
DDS - 90 meter	2.0 GB	2.0 GB	2.0 GB	2.0 GB
DDS - 120 meter		4.0 GB	4.0 GB	4.0 GB
DDS - 125 meter			12.0 GB	12.0 GB
DDS - 150 meter				20.0 GB

8mm

Cartridge	Eliant 820 (8500C format)	Exabyte 8500/8500C	Exabyte 8505	Mammoth 8900
15 meter	8500: 600 MB 8500C: 1.2 GB	RW	RW	Read Only
54 meter	8500: 2.35 GB 8500C: 4.7 GB	RW	RW	Read Only
112 meter	5 GB	EXB-8500: 5GB EXB-8500C: 10GB	10 GB	Read Only
160 meter XL	7 GB Recommended			Read Only
22 meter AME				2.5 GB
125 meter AME				RW with FW upgr.
170 meter AME				40 GB

DLT

Cartridge	DLT 4000 Drive	DLT 7000 Drive	DLT 8000 Drive
DLT III 1200 foot	10 GB	10 GB	10 GB
DLT IIIxt 1800 foot	15 GB	15 GB	15 GB
DLT IV 1800 foot	20.0 GB	35.0 GB	40.0 GB
DDS - 125 meter			12.0 GB
DDS - 150 meter			

2.4.3.1 Format, capacity and media matrix

The client machine will use its own native driver for the tape drive. It is important to consider that a device can support different media formats, but not all formats can be used with the same device.

The above table shows how to consider the relationship between the tape and the media format. The default value using the `DEFINE DEVCLASS` command for the `format` parameter is DRIVE and may not be appropriate for the tape device available on the client site (see Section 2.4.4, “Generating a backup set on tape” on page 60).

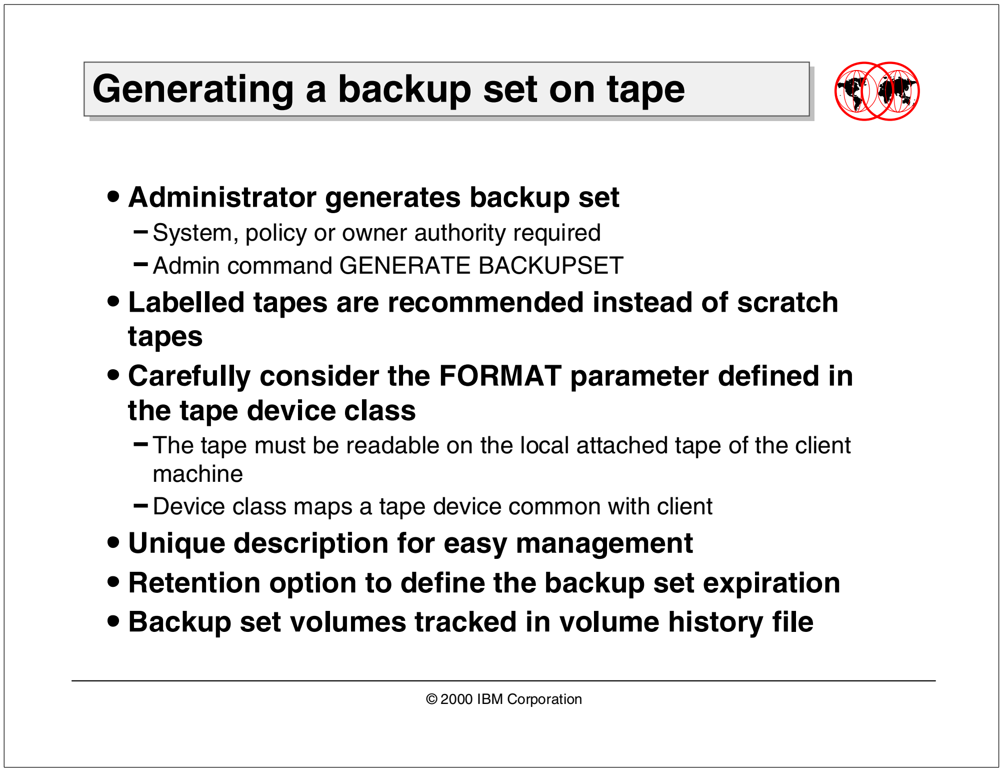

Generating a backup set on tape

- **Administrator generates backup set**
  - System, policy or owner authority required
  - Admin command GENERATE BACKUPSET
- **Labelled tapes are recommended instead of scratch tapes**
- **Carefully consider the FORMAT parameter defined in the tape device class**
  - The tape must be readable on the local attached tape of the client machine
  - Device class maps a tape device common with client
- **Unique description for easy management**
- **Retention option to define the backup set expiration**
- **Backup set volumes tracked in volume history file**

2.4.4 Generating a backup set on tape

A backup set may be generated using the FILE device class or a tape device class. They are all considered as sequential media. A FILE device class can be used to generate a backup set. The resulting file or files may then be copied to removable media such as CD-ROM, JAZ, or ZIP devices. Removable files is also supported, so backup set can be directly generated on JAZ, ZIP, or floppy disk media. These portable media may be useful especially for client machines with limited a quantity of data or without a local tape. However, for larger amounts of data, tape media should be considered.

To generate a backup set, an administrator with system authority or an administrator with restricted policy or owner authority is required.

A tape unit common with the client machine needs to be defined on the server. The *device class* describes the format to be used to write the backup set on the media. When defining the device class, the `format` parameter must be carefully considered. The default value for `format` is DRIVE. It specifies that Tivoli Storage Manager selects the highest format that can be supported by the drive unit on which a volume is mounted. We recommend an explicit value for the `format` parameter, compatible with the tape format supported on the local attached drive of the client machine (see Section 2.4.3.1, “Format, capacity and media matrix” on page 59).

Tivoli Storage Manager server will, by default, use scratch tapes to store the backup set. If using devices without a barcode feature such as 4mm or 8mm, we suggest a hand-written description label, that will be helpful to match the tapes.

Any tape used for generating a backup set has to be labeled as a TSM tape. The dsmlabel command, the wizard option "media labeling", or the administrative Web interface, may be used to label a tape.

For Windows NT or Windows 2000, this could be:

```
C:\Program Files\TIVOLI\TSM\utils> dsmlabel -drive=mt3.0.0.0
```

For the AIX server, this could be:

```
root > dsmlabel -drive=/dev/rmt0
```

A backup set is created by the server using a command like this:

```
tsm: 78-BB053_SERVER1> generate backupset JAMAICA JAMSET_M \
cont> \\jamaica\c$ devclass=C8MM volume=JAMSET1 \
cont> descripton='Jamaica HD C' retention=30
```

Here, `JAMAICA` is the node name and `JAMSET_M` is the backup set name. A tape volume labeled `JAMSET1` will be used. It is possible to specify more than one tape just by inserting the volume names separated by a comma. If volume names are not used, then scratch tapes will be used.

The `retention` parameter will provide the duration of the package. At the expiration date the backup set media will become scratch (if tape) or deleted (if file). The volume history also will be updated, erasing entries for backup set volumes.

The node for which the backup set is to be created must be a registered node, and an incremental backup must be already completed.

The `GENERATE BACKUPSET` command runs as a background process, and if stopped, the target media will not contain a complete backup set. So it will be necessary to delete the backup set and restart from the beginning.

Managing backup sets

- QUERY BACKUPSET
- UPDATE BACKUPSET
- DELETE BACKUPSET
- QUERY BACKUPSETCONTENT
- QUERY VOLHISTORY TYPE=BACKUPSET

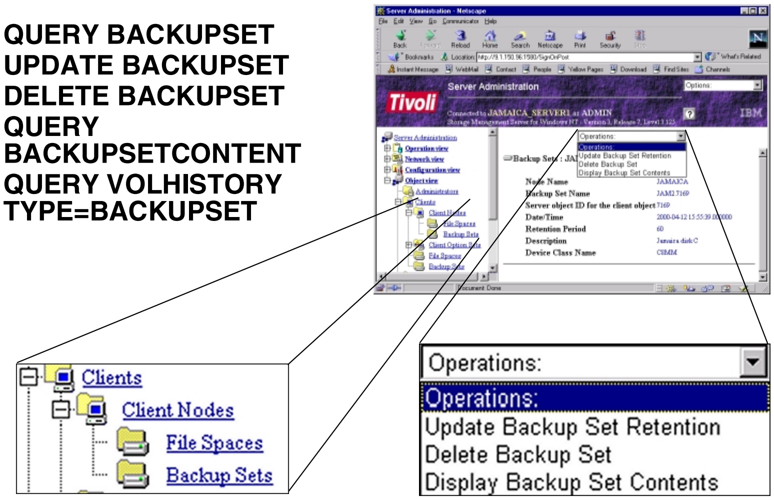

Node Name	JAMAICA
Backup Set Name	JAM2.7169
Server object ID for the client object	7169
Date/Time	2000-04-12 15:55:39.000000
Retention Period	60
Description	Jamaica disk C
Device Class Name	C8MM

© 2000 IBM Corporation

2.4.5 Managing backup sets

Administrators can manage backup sets using the following commands:

```
QUERY BACKUPSET

UPDATE BACKUPSET

QUERY BACKUPSETCONTENT

DELETE BACKUPSET

QUERY VOLHISTORY TYPE=BACKUPSET
```

As shown in the following screen, the QUERY BACKUPSET command displays information about backup set, for example node, description, creation date and retention period.

```
tsm> query backupset jamaica
Node Name: JAMAICA
  Backup Set Name: JAMSET_M.117211
        Date/Time: 04/11/2000 16:30:09
 Retention Period: 60
Device Class Name: C8M
      Description: Jamaica disk C
```

The retention period value can be changed using the UPDATE BACKUPSET command.

It is also possible to display information about the content of a backup set. The command is QUERY BACKUPSETCONTENT.

The deletion of a backup set can be done using the administrative `DELETE BACKUPSET` command. The volume history maintains information about the backup set that can be managed using `QUERY/UPDATE/DELETE VOLHISTORY` commands.

The WEB client interface (see Figure 18) as well may be used to issue the above commands.

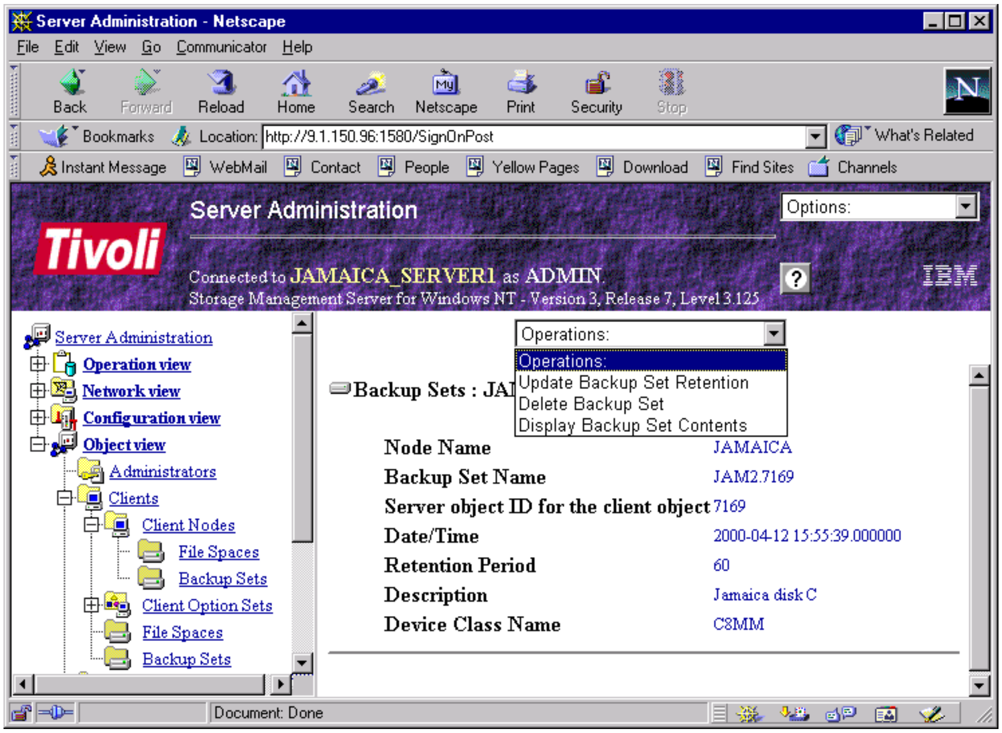
Server Administration - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://9.1.150.96:1580/SignOnPost What's Related
Instant Message WebMail Contact People Yellow Pages Download Find Sites Channels
Server Administration
Options:
Tivoli
IBM
Server Administration
Operation view
Network view
Configuration view
Object view
Administrators
Clients
Client Nodes
File Spaces
Backup Sets
Client Option Sets
File Spaces
Backup Sets
Operations:
Operations:
Update Backup Set Retention
Delete Backup Set
Display Backup Set Contents
Document: Done

Connected to JAMAICA_SERVER1 as ADMIN.
Storage Management Server for Windows NT - Version 3, Release 7, Level 3.125

Backup Sets : JA

Node Name	JAMAICA
Backup Set Name	JAM2.7169
Server object ID for the client object	7169
Date/Time	2000-04-12 15:55:39.000000
Retention Period	60
Description	Jamaica disk C
Device Class Name	C8MM

Figure 18. Managing backup sets using WEB client

A backup set generated on one server can be moved to another server if they are using the same device. The `DEFINE BACKUPSET` command makes it available to the other server. This improved flexibility in moving backup sets to different servers allows the user the ability to restore their data from a server other than one on which the backup set was created.

The client information about backup sets may be displayed using the `QUERY BACKUPSET` client command as shown in the following screen:

```
tsm> q backupset

   Backup Set Name  Generation Date     Retention  Description
   ---------------- ------------------- ---------- -----------------------------
JAMSET_M.117211   04/11/2000 16:30:09 60         Backup Set for Jamaica Disk C
```

The information is in different format, but quite similar. However, only on the server output is the device class displayed.

Figure 19 shows the client GUI for a restore request selecting the Server option:

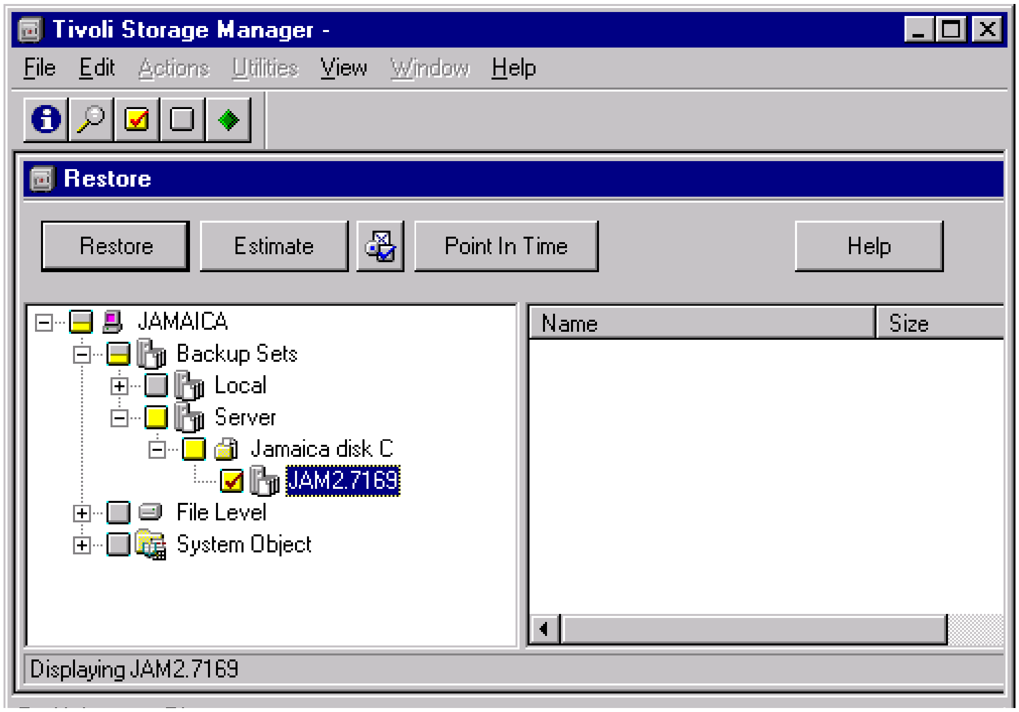

Figure 19. Server based restore

On the expanded tree are shown “Jamaica disk C:” as the description, and “JAM2.7169” as backup set name. The unique number “7169” is a suffix automatically appended to “JAM2” at backup set generation time.

If using a manual tape drive such as 4mm or 8mm to create backup set, it is recommended to apply a hand-written label to tapes, with a description, backup set name, and volume name, for easier management.

The `QUERY VOLHISTORY T=BACKUPSET` command will show information like this:

```
 Date/Time: 04/12/00 11:09:34
     Volume Type: BACKUPSET
   Backup Series:
Backup Operation:
      Volume Seq: 1
    Device Class: C8MM
     Volume Name: JAMAICA
 Volume Location:
         Command: GENERATE BACKUPSET JAM2 JAMAICA \\jamaica\c$ DEVCLASS=C8MM
                   RETENTION=60 VOLUMENAMES=JAMAICA SCRATCH=YES
                   DESCRIPTION="Jamaica Disk C" WAIT=NO
```

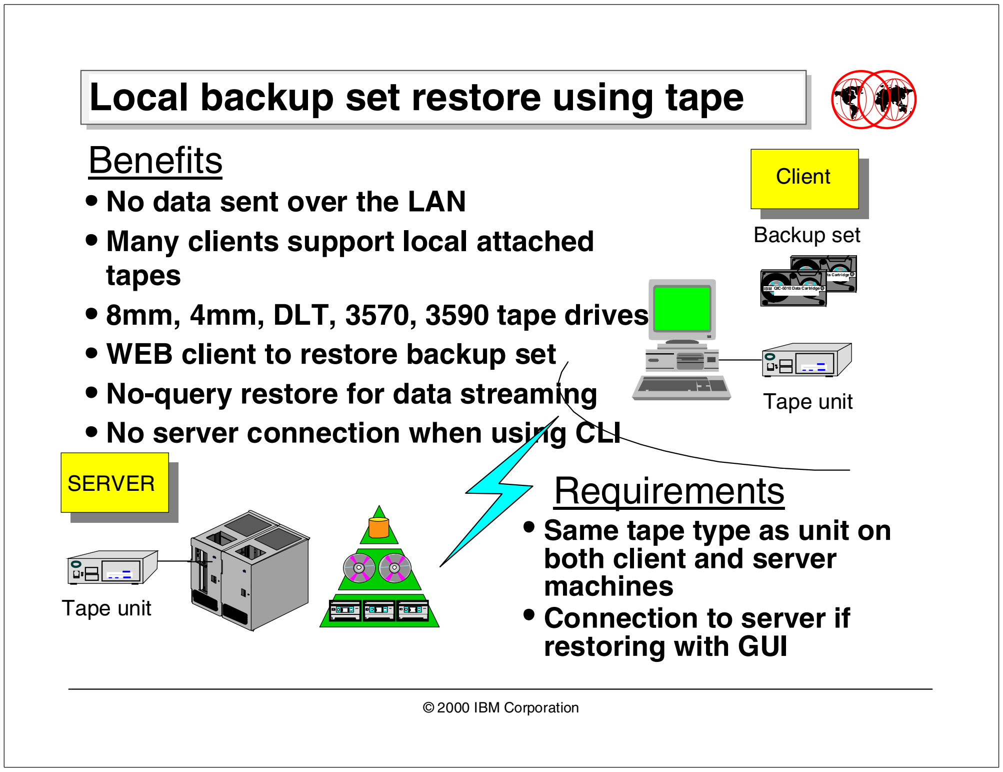

Local backup set restore using tape

Benefits

- **No data sent over the LAN**
- **Many clients support local attached tapes**
- **8mm, 4mm, DLT, 3570, 3590 tape drives**
- **WEB client to restore backup set**
- **No-query restore for data streaming**
- **No server connection when using CLI**

Requirements

- **Same tape type as unit on both client and server machines**
- **Connection to server if restoring with GUI**

2.4.6 Local backup set restore using tape

The restore of a backup set, using a locally attached tape unit, gives the ability to a client to restore all or part of filespaces previously backed up to a TSM server in a LAN-free configuration.

The major benefits of tape usage are:

1. All data is locally available on tape media; this avoids sending data across the network.
2. All V3.7.2 clients, except OS/390 UNIX clients and Novell Netware, support local attached tapes.
3. These devices are: 8mm, 4mm, DLT, 3570 MP and 3590 tape drives.
4. UNIX and Windows NT/2000 clients are able to restore a backup set using the Web client interface. At the moment, the Linux Web client allows you to restore backup sets only from a `FILE` device class.
5. Data is streamed as a no-query restore operation, without confirmation of verbs.
6. If using the Command Line Interface the complete restore operation can be performed in a stand-alone machine without the network connection to the Tivoli Storage Manager server.

These requirements are to be considered:

1. The tape unit **must** be the same type on both client and server, and the tape formatting **must** be compatible for both tape units.
2. If using GUI interface, the client needs to contact the TSM server on the LAN.

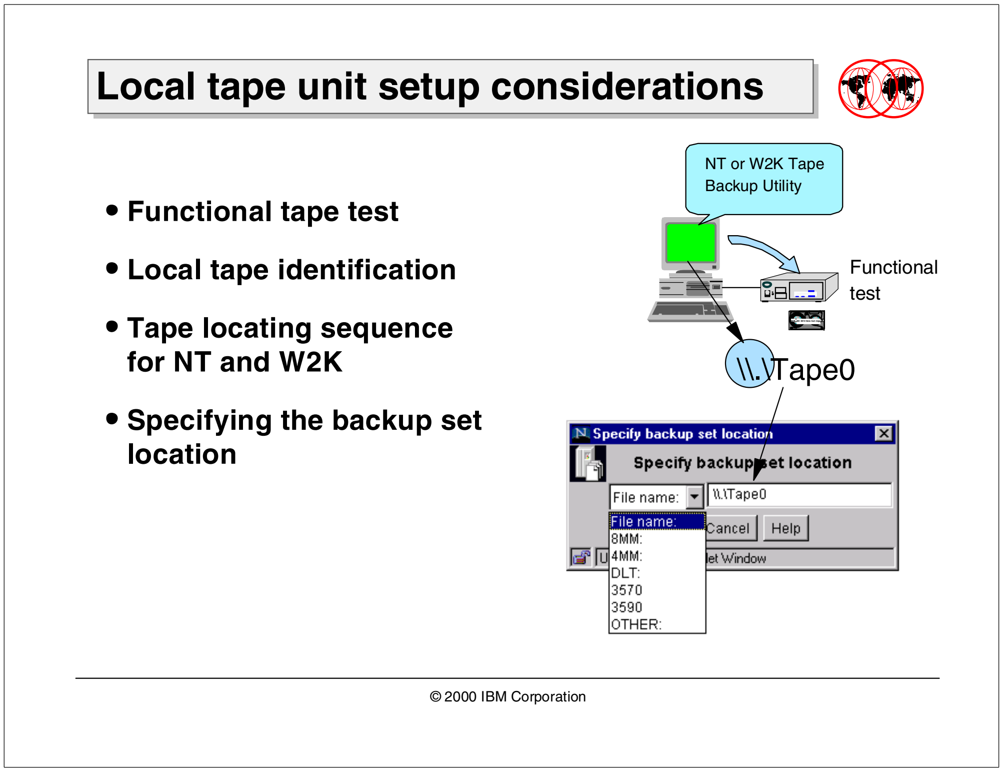

Local tape unit setup considerations

- Functional tape test
- Local tape identification
- Tape locating sequence for NT and W2K
- Specifying the backup set location

© 2000 IBM Corporation

2.4.6.1 Local tape unit setup considerations

Using tape devices may be a fast recovery solution, but the following setup requirements must be considered.

Functional tape test

A tape unit must be defined and operative on the client machine. The specific device driver must be used. It may be helpful to test the unit starting a tape backup operation using the Operating System native utility.

Local tape identification

UNIX based systems refer to the tape unit by a special device name such `/dev/rmt0`, so this information is usually very simple to find. Other operating systems such as Windows NT and Windows 2000 (W2K) *map* the tape using a special definition. It is normally identified as "\\.\TapeX" where "\\.\" is a specific NT/W2K prefix and "X" will be 0 for the first tape device, 2 for the second, etc.

In a Windows NT machine, this number can be retrieved clicking on the Tape Devices icon in the Control Panel, then double-clicking on the tape for properties. The name of the tape unit (for example, `Tape0`), however the "\\.\" is not shown, and users must add these special characters before the name itself.

In Windows 2000, the full name of the tape can be located as a removable storage device. Since the location of such a name is not so intuitive, especially the very first time, we offer a method to identify the proper tape name (see Section 2.4.6.2, "Tape locating sequence for Windows 2000" on page 67).

2.4.6.2 Tape locating sequence for Windows 2000

This sequence is done by a series of clicks on the Windows objects, which should be very familiar to the Windows user:

1. Open the “Computer Management” application following the menu path **Start->Programs->Administrative Tools->Computer Management**.
2. In the “Computer Management” application, expand the menu tree **Removable Storage->Physical Locations->Tape folder name**, as shown in Figure 20.
3. Click on **Drives**.
4. On the right portion of the window, right-click on your tape device and the “Device Name” value (\\.\Tape0) is finally shown (Figure 20).

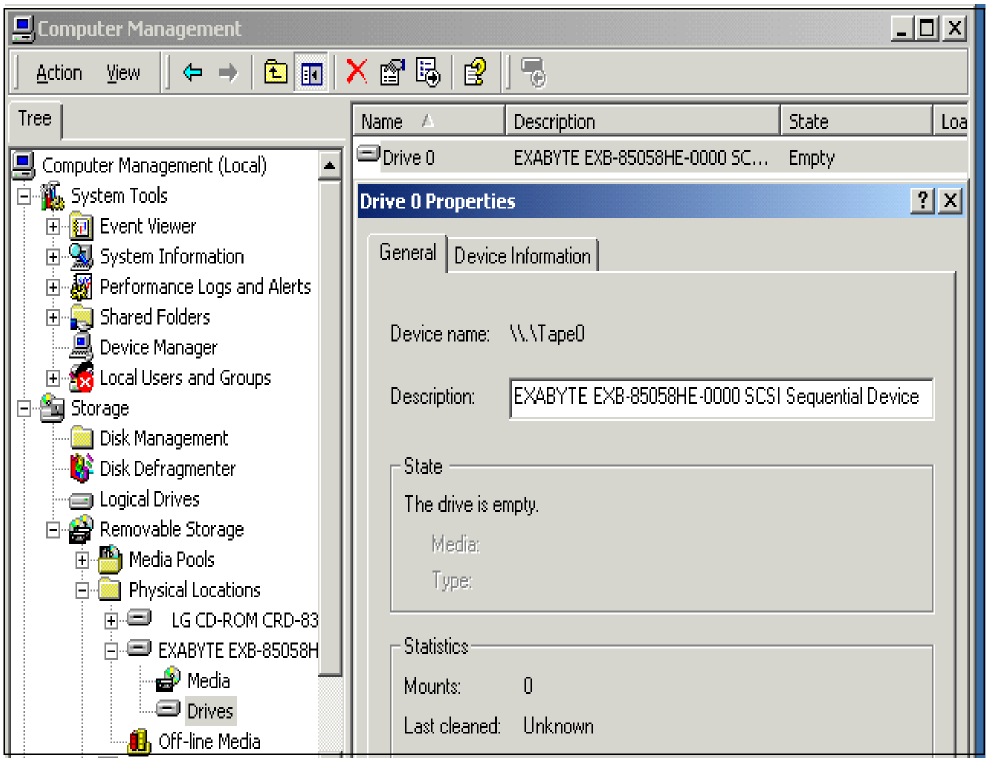

Name	Description	State	Loa
Drive 0	EXABYTE EXB-85058HE-0000 SC...	Empty	

Device name: \\.\Tape0

Description: EXABYTE EXB-85058HE-0000 SCSI Sequential Device

The drive is empty.

Media:

Type:

Mounts: 0

Last cleaned: Unknown

Figure 20. Tape locating sequence

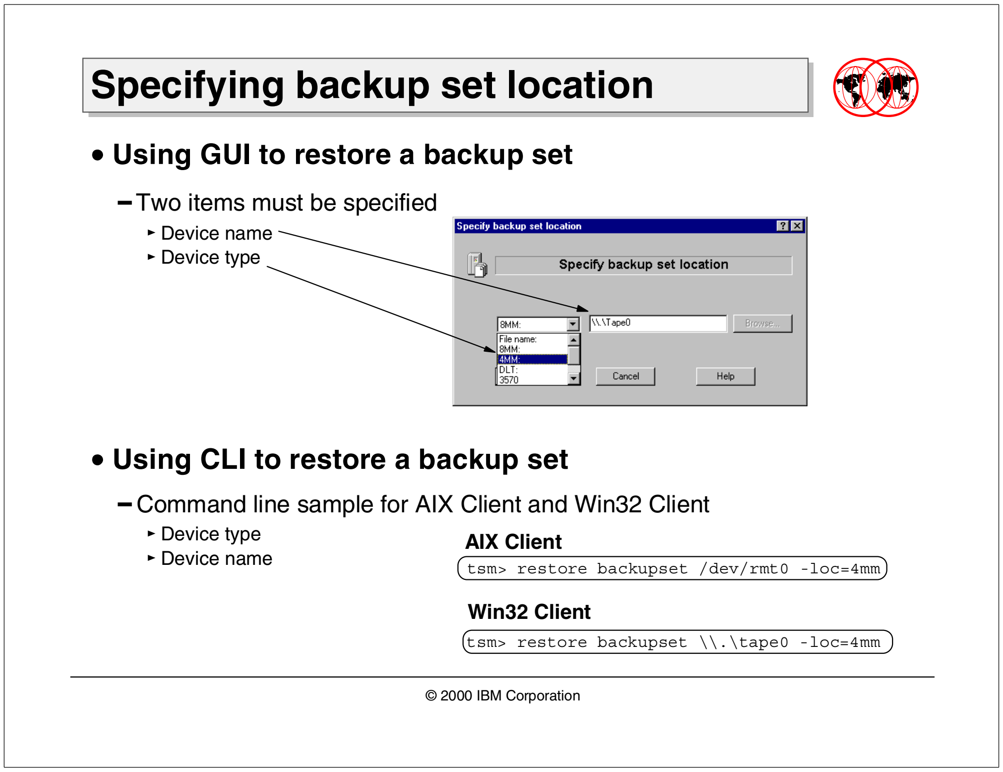

Specifying backup set location

- **Using GUI to restore a backup set**
  - Two items must be specified
    - Device name
    - Device type

8MM:

\\.\Tape0

- **Using CLI to restore a backup set**
  - Command line sample for AIX Client and Win32 Client
    - Device type
    - Device name

AIX Client

```
tsm> restore backupset /dev/rmt0 -loc=4mm
```

Win32 Client

```
tsm> restore backupset \\.\tape0 -loc=4mm
```

2.4.6.3 Specifying backup set location

If using the Tivoli Storage Manager client GUI for a local backup set restore, a server authentication has to be performed. Userid and password are needed to qualify the client.

Then, by clicking on “local option”, a prompt will be displayed. In the left field of the window, the media type must be entered. In the right field, the tape identifier or tape name must be entered. This structure of the name is defined by the operating system and can be determined as shown in Section 2.4.6.1, “Local tape unit setup considerations” on page 66.

If the tape contains a backup set, its name and description will be displayed on the GUI. Now it is possible to start the restore of the backup set. The destination path can be different than the original path, and the option for a new restore path is shown.

The backup set can be restored as well by using the client CLI. In this case, the server connection is not needed. The client can query the local tape and perform a backup set restore using the specific commands.

Many TSM clients can also perform operations on backup sets using the Web client interface; this improves the ability to manage a restore process using any other machine with a Web browser and proper authority on the network.

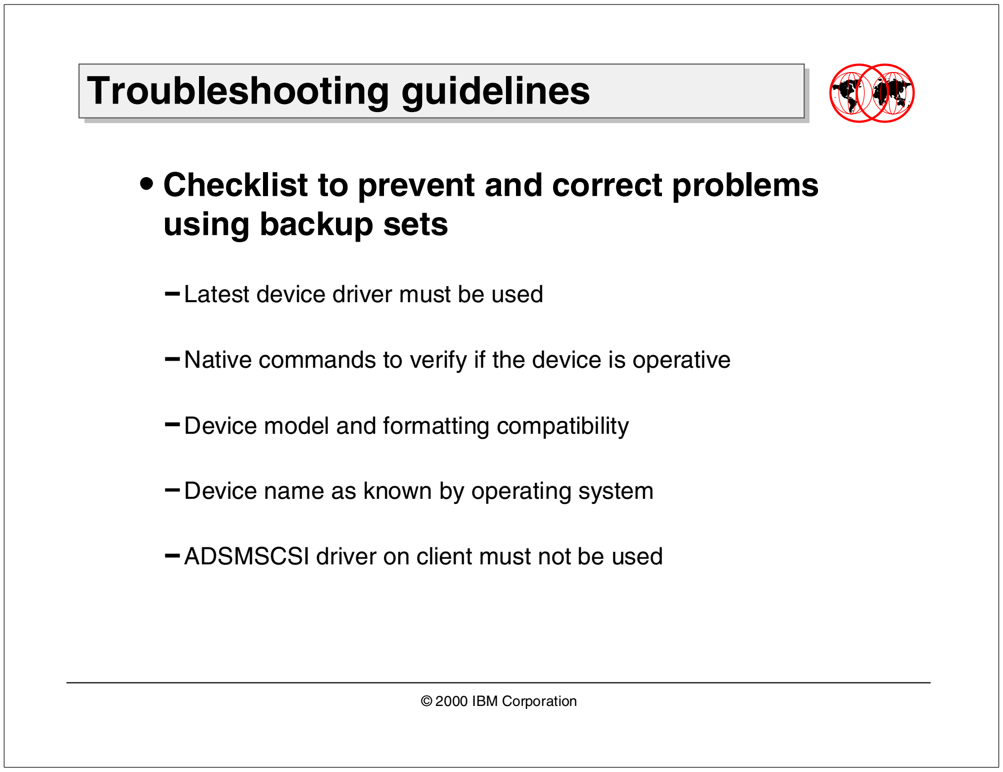

Troubleshooting guidelines

- **Checklist to prevent and correct problems using backup sets**
  - Latest device driver must be used
  - Native commands to verify if the device is operative
  - Device model and formatting compatibility
  - Device name as known by operating system
  - ADSMSCSI driver on client must not be used

2.4.7 Troubleshooting guidelines

Devices should be carefully configured on both server and client machines to ensure that a backup set generated on the server can be queried and restored on the client machine. The following hints and considerations may be helpful to solve problems during a local backup set restore operation:

1. Does the machine have the latest device driver? This is very important, because an old device may not work as expected, and the result can be inconsistent behavior or complete failure using the device. The latest device driver can be obtained from the device manufacturer's Web site.
2. Is the device driver loaded? Can the operating system use the device? On NT clients you can use the backup utility in NT administrative tools or the tar command in AIX to test if the operating system can communicate with the tape drive.
3. The tape drive's model: This information is necessary to determine capacity and format of the tape (see 2.4.3.1, "Format, capacity and media matrix" on page 59). The format used to generate the backup set must be considered. If the device class has DRIVE as a format option (default value), the server will try to use the highest possible format for the attached device (for example, the highest format for a DLT 8000 is 40 GB using a DLT type IV). The device that will be used on the client machine must be able to read that format type.
4. In the left field of the local restore panel, which comes up using the GUI or the WEB Client, the device name must be entered as known by the operating system. For example, `/dev/rmt0` can be a device for a UNIX client; NT will map the tape unit as `\\.\Tape0`. In the right field, the drive type must be selected (8mm, 4mm, DLT, 3570, 3590 or other).

5. If restoring a backup set on a Windows NT machine that is both a TSM server and TSM client, then the adsmscsi driver used for the tape unit must be stopped. Also check that adsmscsi is not automatically loaded at boot time. After a reboot of the operating system, the native driver must be loaded to allow the correct selection by the client in the backup set restore panel. The backup set restore will not work with the adsmscsi device driver.

2.5 Web client enhancements

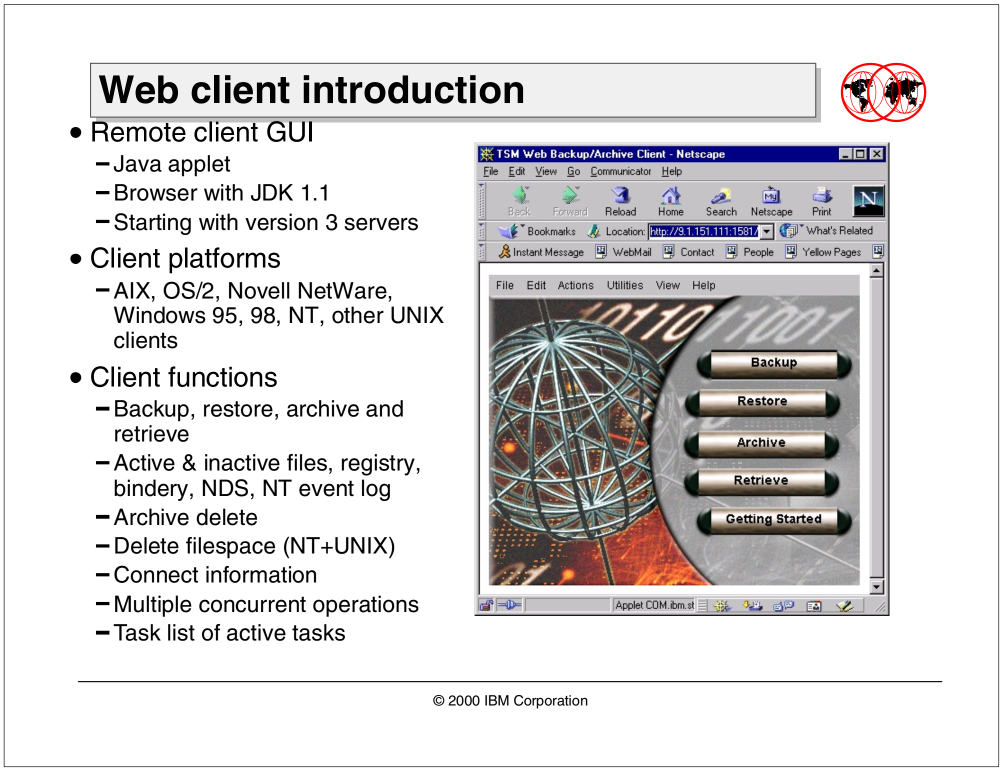

Web client introduction

- Remote client GUI
  - Java applet
  - Browser with JDK 1.1
  - Starting with version 3 servers
- Client platforms
  - AIX, OS/2, Novell NetWare, Windows 95, 98, NT, other UNIX clients
- Client functions
  - Backup, restore, archive and retrieve
  - Active & inactive files, registry, bindery, NDS, NT event log
  - Archive delete
  - Delete filespace (NT+UNIX)
  - Connect information
  - Multiple concurrent operations
  - Task list of active tasks

TSM Web Backup/Archive Client - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print
Bookmarks Location: http://9.1.151.111:1581/
What's Related
Instant Message WebMail Contact People Yellow Pages
File Edit Actions Utilities View Help
Backup
Restore
Archive
Retrieve
Getting Started
Applet COM.ibm.st

© 2000 IBM Corporation

2.5.1 Web client introduction

The Tivoli Storage Manager Web backup-archive client is a Java client interface that enables authorized users to perform remote operations on a backup-archive client system.

2.5.1.1 Remote client GUI

The Web client is a Java applet that provides a remote client GUI for Tivoli Storage Manager backup-archive clients. Authorized users can access the Web client remotely using a Web browser. The graphic is the initial client hub window displayed in the Web browser when the Web client is accessed.

The Web backup-archive client (Web client) permits an authorized administrator, help desk person, or end user to perform backup, restore, archive, and retrieve services on any machine or platform that supports a Java 1.1.6-capable browser, such as Netscape Navigator 4.06 or higher, or Microsoft Internet Explorer 4.01 or higher. If the browser does not have the correct level of Java support, the client hub window is displayed as a grey box.

The Web client is supported only on Tivoli Storage Manager Version 3 servers.

2.5.1.2 Client platforms

The Web client is available since PTF6 of ADSM. The Tivoli Storage Manager Client 3.7.2 enhances the functionality.

The following client platforms are supported:

- AIX
- OS/2 Warp
- Novell NetWare
- Windows 95, Windows 98, and Windows NT 4.0
- Digital UNIX 4
- HP-UX 11.0
- SGI IRIX 6.5
- Sun Solaris 2.6 and 2.7

The Web client is a subcomponent of the backup-archive client package. The native backup-archive clients can be installed with the Web client interface. The Web client interface is not a separately installed package as was the case with the previous Webshell client, which is replaced by the Web client.

2.5.1.3 Client functions

The Web client supports most of the Version 3 backup-archive client functions. It can be used to perform backup, restore, archive, and retrieve operations. For restore operations, both active and inactive files can be restored, and the point-in-time restore function is supported. The functions not currently supported include cross client restores, and the file find and search functions.

The Windows NT Web client can back up and restore the NT registry and event log. The NetWare Web client can back up and restore the NetWare 3.x bindery and NetWare 4.x Directory Services (NDS) and provides, for the first time, a NetWare client GUI.

The Web client is another client interface. It can be used only to perform backup-archive client operations. If used from a Web browser on a remote workstation, no local access to backup or archive data is provided on that remote workstation. Data cannot be restored or retrieved locally; it can only be restored to or retrieved to the client workstation that owns the data. A Web browser can also be used locally on the client workstation to invoke the Web client as an alternative interface to the backup-archive client command line interface or GUI.

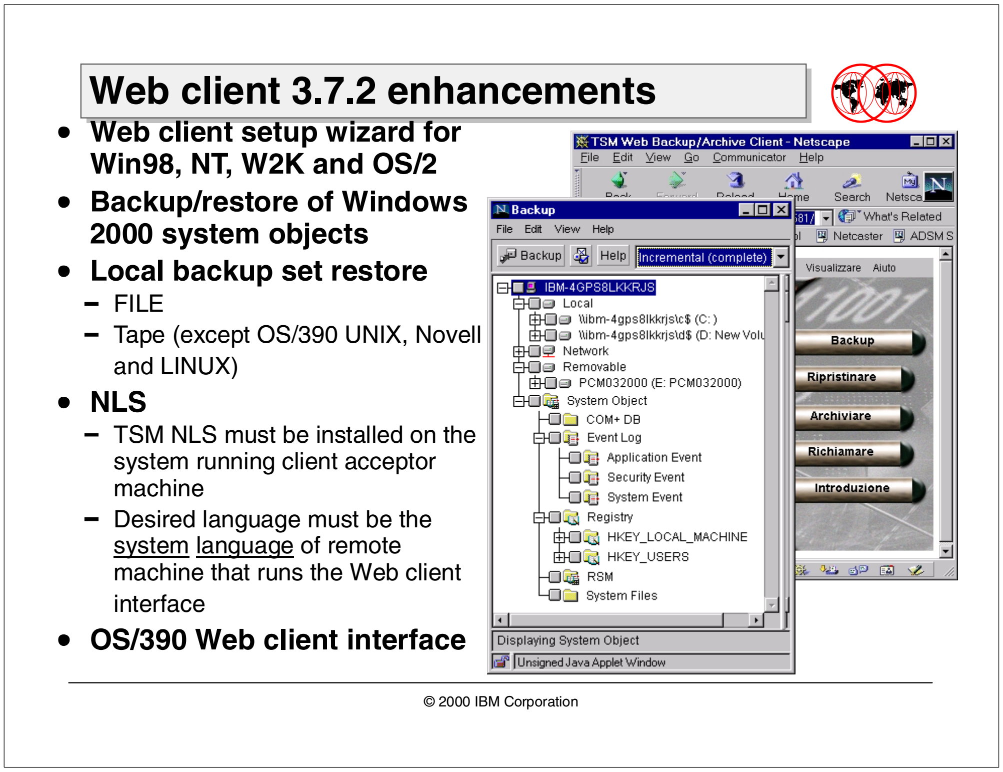

Web client 3.7.2 enhancements

- **Web client setup wizard for Win98, NT, W2K and OS/2**
- **Backup/restore of Windows 2000 system objects**
- **Local backup set restore**
  - FILE
  - Tape (except OS/390 UNIX, Novell and LINUX)
- **NLS**
  - TSM NLS must be installed on the system running client acceptor machine
  - Desired language must be the system language of remote machine that runs the Web client interface
- **OS/390 Web client interface**

© 2000 IBM Corporation

2.5.2 Web client 3.7.2 enhancements

Web client enhancements to improve functionality of the Java-based remote Web client interface covering the following items:

- Backup/restore of Windows 2000 system objects
- Backup set restore from FILE and local attached tape
- National language support
- OS/390 Web client interface

For the Web client configuration wizard check Section 2.3.3, “Web client configuration wizard” on page 50.

2.5.2.1 Backup and restore of Windows 2000 system objects

Figure 21 shows the backup and restore panels prompted in a Web client interface for a Windows 2000 advanced server.

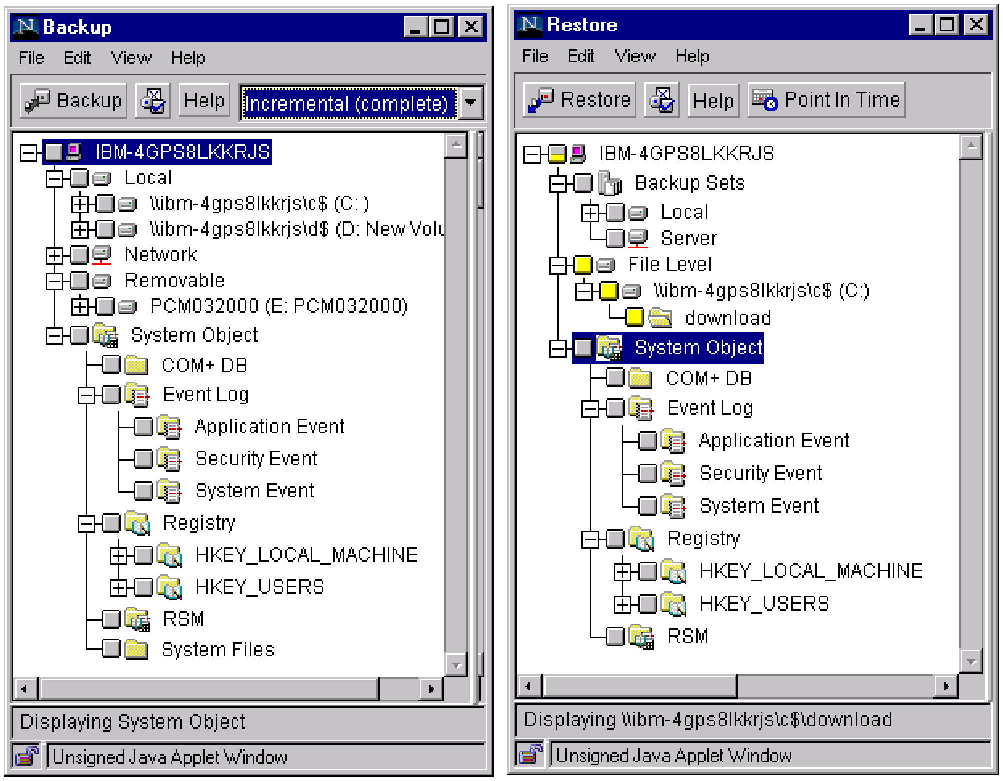

Figure 21. GUI backup and restore panels for Windows 2000

2.5.2.2 Backup set restore using Web client

Backup set restore can be done same way as using the backup-archive client GUI. Only browsing is not allowed by Web interface. In the window “Specify backup set location” it is only possible to specify the complete path of the backup set volume. The volume must be locally available at the client machine location.

All local attached tape devices that are supported for both client and server machines are available in most of Web clients. All clients can perform a “FILE” backup set restore from any supported media (for example, CD-ROM, JAZ, ZIP, or external SCSI disk).

2.5.2.3 National Language Support

National Language Support (NLS) is now available for the Web client interface. Users can access a remote client and have an NLS Web Client screen in the same format as defined in their operating system.

The desired client language **must be installed** in the system that starts the client acceptor. In the backup-archive client executable path (for example under Windows in `C:\Program Files\Tivoli\TSM\baclient`), there must be a folder “it_IT”, after the NLS installation was successful. Locally, the user may edit the option file (`dsm.opt`) to set the language (for example `LANG=ITALIAN`), or may use the English default language as well: it does not affect the remote Web client.

The remote system **must define the language as system default language**.

The same rule must be observed for UNIX clients. They also need to install the TSM national language, which can be used or not for local GUI. A directory “it_IT” must exist in the path `usr/Tivoli/tsm/client/ba/bin` after the installation of NLS.

UNIX clients must set the desired language as a system option in order to be able to have it using the Web client interface to contact a remote client.

To define the language for the NT Operating System, double-click on **Regional Settings** in the “Control Panel”, then select the language and click on **Set as system default locale**. The language definition for Windows 95 and 98 is identical; for Windows 2000, after clicking on **Regional Settings**, click on **General** to modify the language. Figure 22 shows the selection panel for Windows NT.

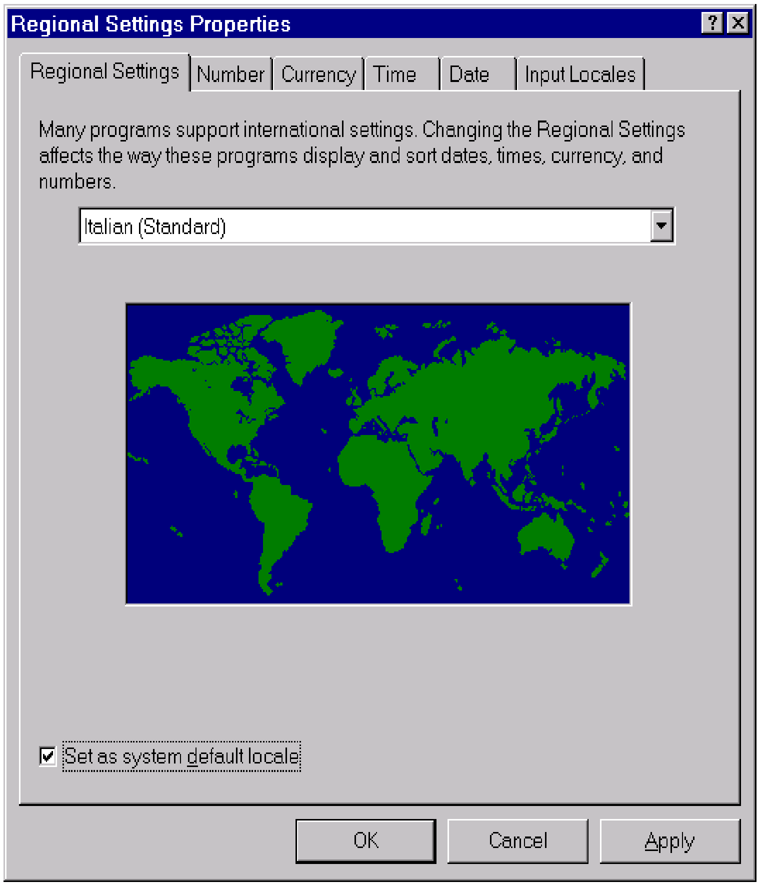

Many programs support international settings. Changing the Regional Settings affects the way these programs display and sort dates, times, currency, and numbers.

Italian (Standard)

☑ Set as system default locale

Figure 22. Setting NLS in Windows NT operating system

2.5.2.4 OS/390 Web client interface

The Web client interface is now available for the OS/390 UNIX System Services client. It looks the same as an AIX Web client, and the agent is started running the executable file `dsmcad` in `/usr/tivoli/tsm/client/ba/bin`. The agent can be manually started, or it can be launched automatically by inserting an entry in `/etc/inittab`.

2.6 Linux client support

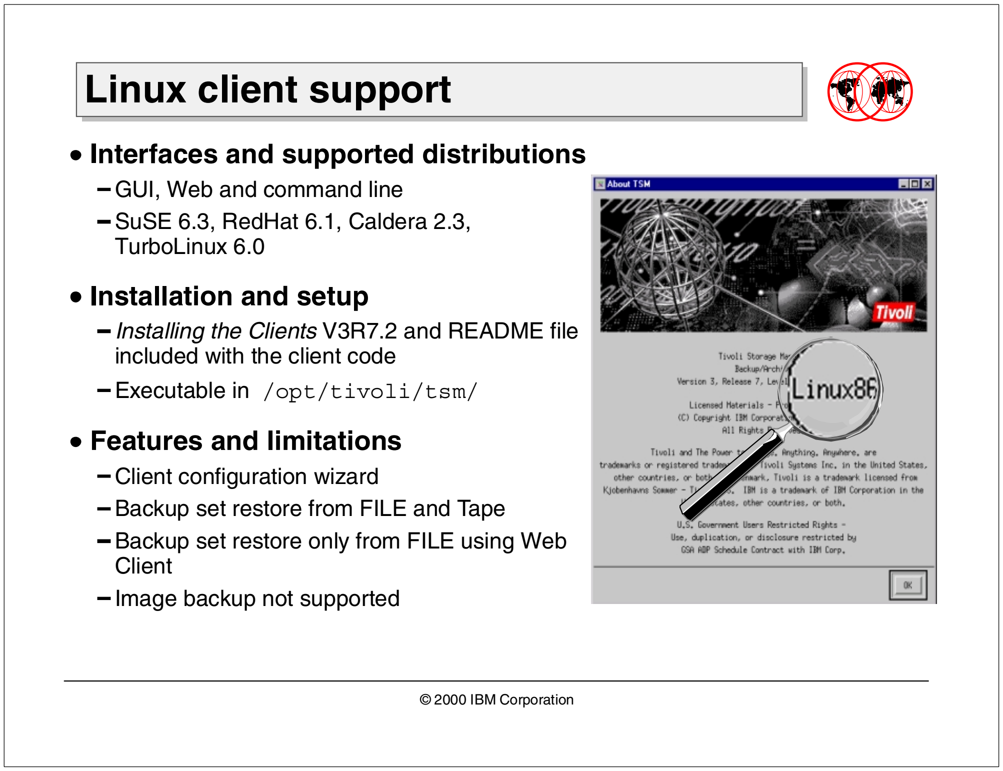

Linux client support

- **Interfaces and supported distributions**
  - GUI, Web and command line
  - SuSE 6.3, RedHat 6.1, Caldera 2.3, TurboLinux 6.0
- **Installation and setup**
  - *Installing the Clients* V3R7.2 and README file included with the client code
  - Executable in `/opt/tivoli/tsm/`
- **Features and limitations**
  - Client configuration wizard
  - Backup set restore from FILE and Tape
  - Backup set restore only from FILE using Web Client
  - Image backup not supported

About TSM

Tivoli Storage Ma
Backup/Arch
Version 3, Release 7, Level
Linux86

Licensed Materials - P
(C) Copyright IBM Corporat
All Rights

Tivoli and The Power t , Anything, Anywhere, are
trademarks or registered trade Tivoli Systems Inc. in the United States,
other countries, or both enmark, Tivoli is a trademark licensed from
Kjobenhavns Sommer - Ti . IBM is a trademark of IBM Corporation in the
states, other countries, or both.

U.S. Government Users Restricted Rights -
Use, duplication, or disclosure restricted by
GSA ADP Schedule Contract with IBM Corp.

© 2000 IBM Corporation

This section contains client environment information, TSM client components, and hardware and software requirements for the Linux platform.

2.6.1 Interfaces and supported environments

The Linux client is supported for the following distributions:

- SuSE 6.3
- RedHat 6.1
- Caldera 2.3
- TurboLinux 6.0

All interfaces are available:

- Backup-archive client (command-line and GUI)
- Administrative client (command line)
- Web backup-archive client

API and API Multi-Thread are also available for applications that will use them, such as the Web client, for example.

The only communication method supported is TCP/IP. Any supported network adapter can be used.

The backup-arrive GUI requires:

- X Window System X11R6
- Linux kernel 2.2.5 or later
- glibc 2.1.x
- C++ shared libraries (libstdc++2.9.x), and thread libraries (libthread.so)

2.6.2 Installation and setup

The latest Linux client code is downloadable at the following Tivoli site:
`http://www.tivoli.com/support/storage_mgr/clients.html#linux`

There are five files in the package:

- `README.WEBCLI`
- `README.API`
- `README`
- `TIVsm-API.i386.rpm`
- `TIVsm-BA.i386.rpm`

Although the installation instructions for Linux are in the manual *Tivoli Storage Manager, Installing the clients*, SH26-4119, we recommend that you look carefully at the README file for the latest information. The installation steps are very simple to perform:

1. Log in as root user and mount the CD-ROM to /cdrom and change directory to: `/cdrom/tsmcli/linux` or, if downloaded, change directory on the path used for the operation.
2. Enter the following command to install the backup-archive client (command-line and GUI), the administrative client (command-line), and the Web client:

   ```
   rpm -i TIVsm-BA.i386.rpm
   ```

3. Enter the following command to install the API:

   ```
   rpm -i TIVsm-API.i386.rpm
   ```

4. Answer Yes (y) to all questions during installation.

The client code is installed as follows:

- The TSM backup-archive and Web client files are installed in the directory: `/opt/tivoli/tsm/client/ba/bin`
- The TSM administrative client (command-line) is installed in the directory: `/opt/tivoli/tsm/client/admin/bin`
- The TSM API files are installed in the directory: `/opt/tivoli/tsm/client/api/bin`
- The TSM documentation files are installed in the directories: `/opt/tivoli/tsm/client/books/html` and `/opt/tivoli/tsm/client/books/pdf`

2.6.2.1 Starting the Web client

In the directory `/opt/tivoli/tsm/client/ba/bin` is the executable file `dsmcad`. Before running this command, the environment variables DSM_DIR, DSM_CONFIG and DSM_LOG must be set.

It may be helpful to create a script with the export command to be executed in the same shell. Then it will be possible to run from there the `dsmcad` program. By default the port number is always 1581, but it can be changed in the `dsm.opt` file or using the preference setting option in the GUI as shown in Figure 23.

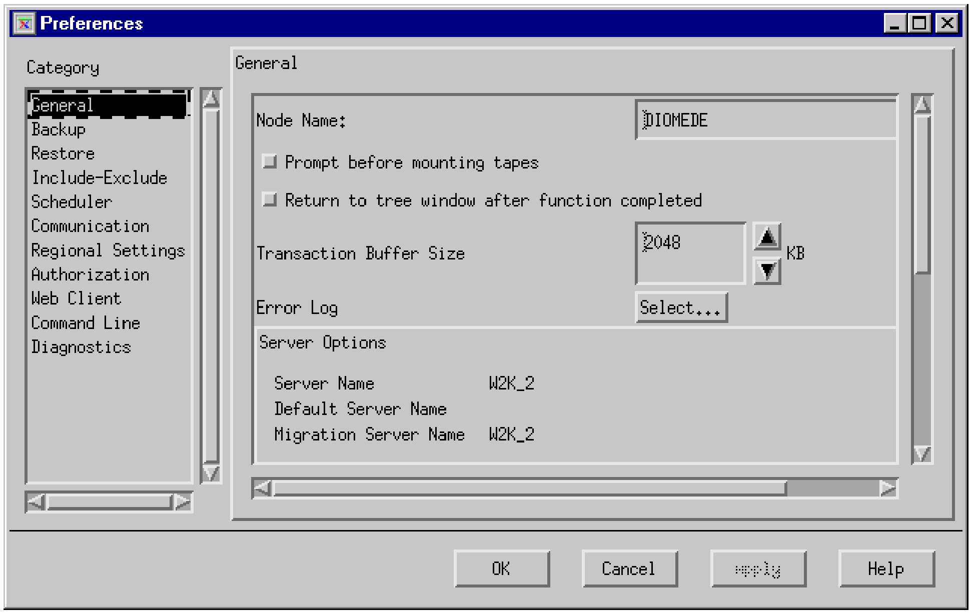

Node Name: DIOMEDE

□ Prompt before mounting tapes

□ Return to tree window after function completed

Transaction Buffer Size 2048 KB

Error Log

Server Options

Server Name W2K_2
Default Server Name
Migration Server Name W2K_2

Figure 23. Customizing of preferences in a Linux client GUI

Once started, any user is allowed to obtain the login Web client panel at the URL: `HTTP://Linux_host_IP_address:1581` (Figure 24).

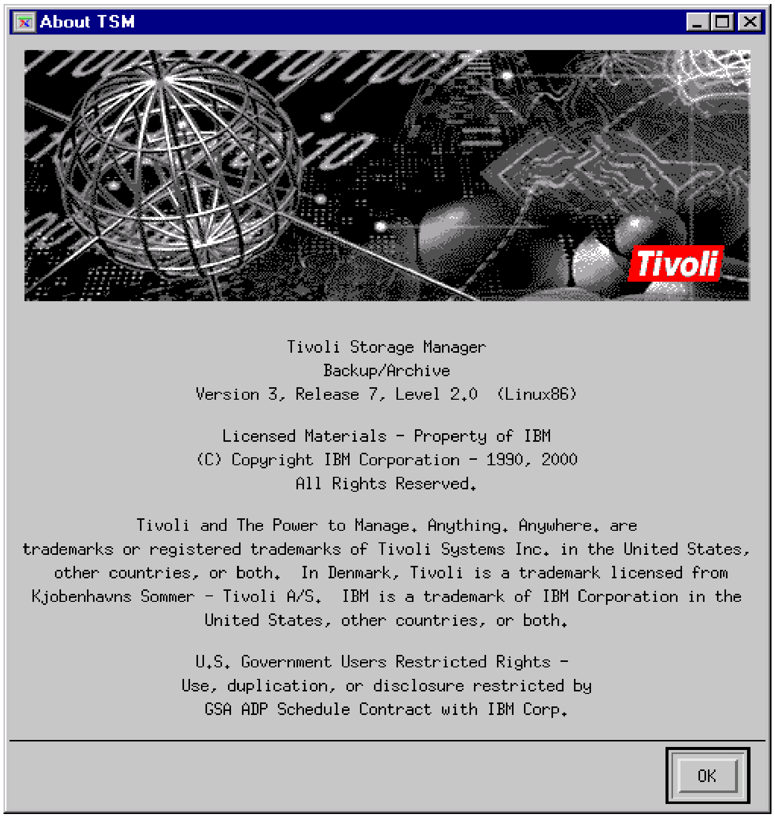

Tivoli Storage Manager
Backup/Archive
Version 3, Release 7, Level 2.0 (Linux86)

Licensed Materials - Property of IBM
(C) Copyright IBM Corporation - 1990, 2000
All Rights Reserved.

Tivoli and The Power to Manage. Anything. Anywhere. are trademarks or registered trademarks of Tivoli Systems Inc. in the United States, other countries, or both. In Denmark, Tivoli is a trademark licensed from Kjobenhavns Sommer - Tivoli A/S. IBM is a trademark of IBM Corporation in the United States, other countries, or both.

U.S. Government Users Restricted Rights -
Use, duplication, or disclosure restricted by
GSA ADP Schedule Contract with IBM Corp.

Figure 24. Linux Web client panel

The connection information shows the platform identification for the Linux client (Figure 25).

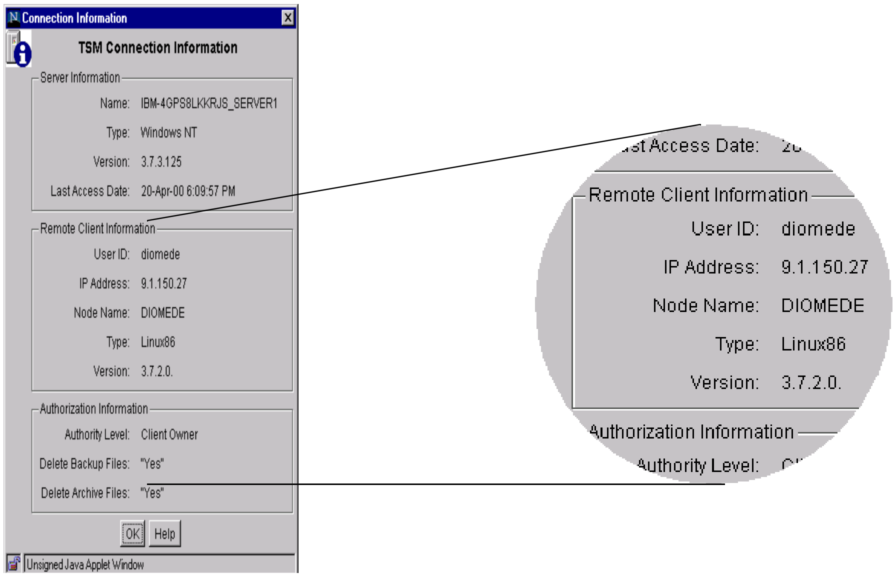

TSM Connection Information

Server Information

Name: IBM-4GPS8LKKRJS_SERVER1
Type: Windows NT
Version: 3.7.3.125
Last Access Date: 20-Apr-00 6:09:57 PM

Remote Client Information

User ID: diomede
IP Address: 9.1.150.27
Node Name: DIOMEDE
Type: Linux86
Version: 3.7.2.0.

Authorization Information

Authority Level: Client Owner
Delete Backup Files: "Yes"
Delete Archive Files: "Yes"

Figure 25. Connection information in a Linux Web client GUI

2.6.3 Features and limitations

The current implementation of the backup-archive client for Linux operating systems has the following features and limitations.

2.6.3.1 Client Configuration Wizard

A Client Configuration Wizard is provided in the Utilities pull-down menu in the GUI interface. Using the wizard, the user is prompted for “Help to configure a TSM Backup Archive Client”*;* then a series of panels will be displayed to complete the customizing.

These panels are:

1. Create/Update/Import an option file
2. Specify a TSM Server name
3. Specify the communication method
4. Specify the TCP/IP address of the TSM Server
5. Specify the TSM node name to use

2.6.3.2 Backup set restore from FILE and tape

The Linux client can perform the restore of a backup set from FILE device class or using a local attached SCSI tape device as other UNIX clients can do.

2.6.3.3 Backup set restore from FILE only using Web client

At the current time the restore of a backup set from a local attached tape is not supported using the Web client. Only restore from FILE (Figure 26) will be allowed using this interface.

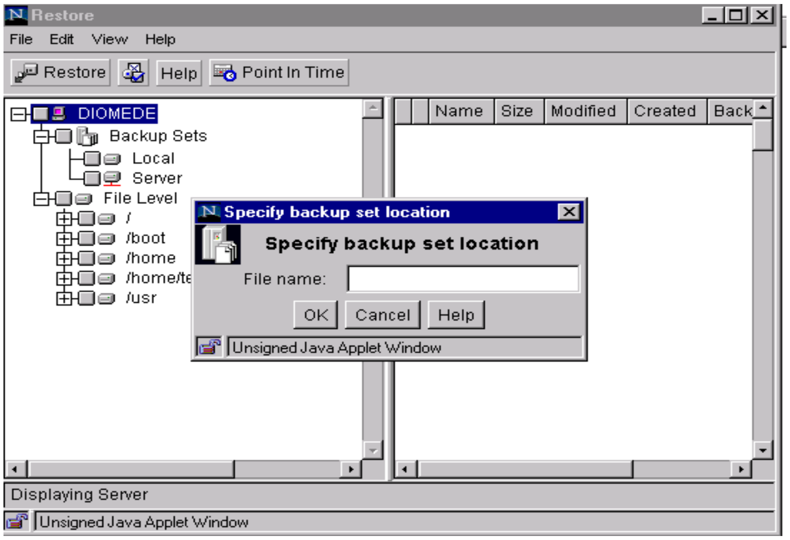

Figure 26. Restore backup set using Linux Web client GUI

2.6.3.4 Image backup function not supported

The Linux client presently does not support the image backup function.

Chapter 3. Server enhancements

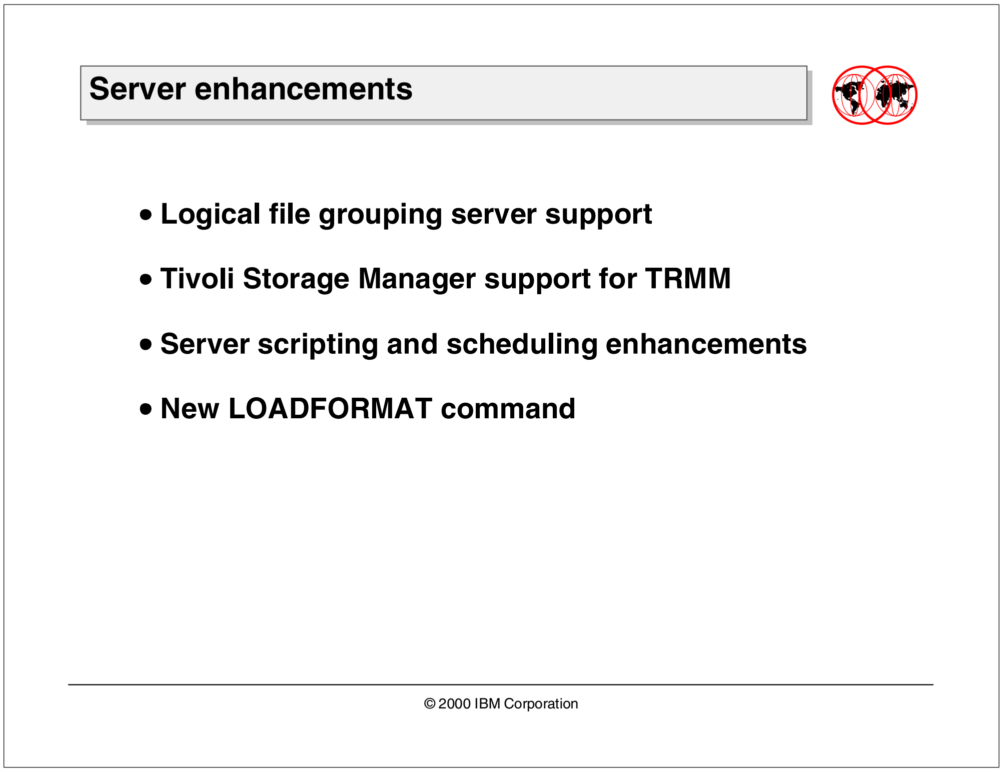

Server enhancements

- **Logical file grouping server support**
- **Tivoli Storage Manager support for TRMM**
- **Server scripting and scheduling enhancements**
- **New LOADFORMAT command**

This chapter covers the major Tivoli Storage Manager server enhancements:

- Server support for logical file grouping
- Tivoli Storage Manager support for Tivoli Removable Media Manager (TRMM)
- Server scripting and scheduling enhancements
- New LOADFORMAT command

3.1 Logical file grouping server support

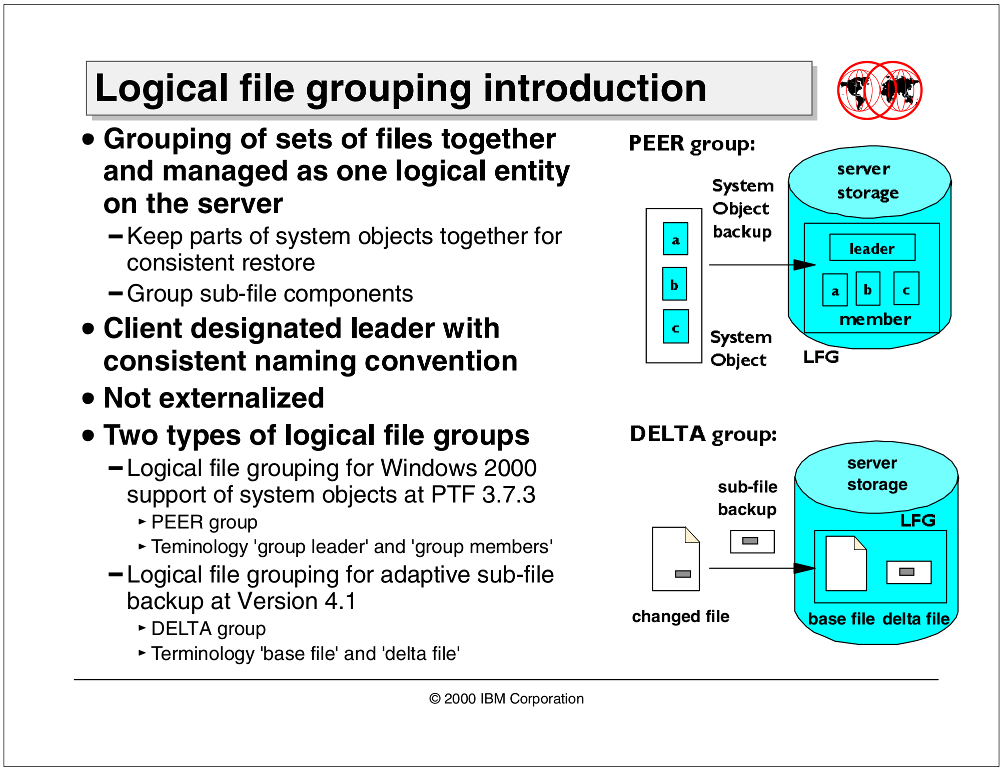

Logical file grouping introduction

- **Grouping of sets of files together and managed as one logical entity on the server**
  - Keep parts of system objects together for consistent restore
  - Group sub-file components
- **Client designated leader with consistent naming convention**
- **Not externalized**
- **Two types of logical file groups**
  - Logical file grouping for Windows 2000 support of system objects at PTF 3.7.3
    - PEER group
    - Teminology 'group leader' and 'group members'
  - Logical file grouping for adaptive sub-file backup at Version 4.1
    - DELTA group
    - Terminology 'base file' and 'delta file'

© 2000 IBM Corporation

Logical file grouping support was initially introduced in the Tivoli Storage Manager server V3.7.3 for system object backup support in a Windows NT and Windows 2000 environment. This support was further enhanced in Tivoli Storage Manager V4.1 to support the adaptive sub-file backup for mobile clients.

Logical file grouping provides a method of grouping sets of related files together by the backup-archive client. The grouping allows the client to “relate” a number of files together in a group, which is managed as a single entity. These related files will be managed together as a single object on the server. However, this is all invisible to the Tivoli Storage Manager users or administrators.

3.1.1 Logical file grouping introduction

The objective of logical file grouping is to give a means to group and manage multiple files as a single object. The functionality is controlled by the client, but there are changes to the Tivoli Storage Manager server to enable management of this file groups. All changes are not externalized. Logical file grouping is only available for backup objects.

3.1.1.1 Logical file grouping for Windows 2000 system objects

Many Windows 2000 system objects, for example, the Active Directory, are data structures which consist of numerous files. These files have to be backed up and restored at one time, and the data management application has to ensure that all components are always available at once. To implement this requirement, Tivoli Storage Manager introduced logical file grouping in Version 3.7.3.

Logical file grouping for Windows 2000 system objects introduces two new terms: *group leader*, who manages all the objects in a group; and *group members*, who are the objects in the group and represent the files of a system object. The group leader represents the entire system object and gets all the versioning and expiration information assigned.

A logical file group consisting of group leader and group members is also called a *PEER group*.

3.1.1.2 Logical file grouping for adaptive sub-file backup

In Tivoli Storage Manager V4.1, logical file grouping is used to implement a relationship between of sub-file components (base file and delta files) needed for the adaptive sub-file backup feature. Unlike the previously-discussed implementation in this case all components of the group get versioning and expiration information assigned.

A logical file group consisting of base file and delta files is also called a *DELTA group*.

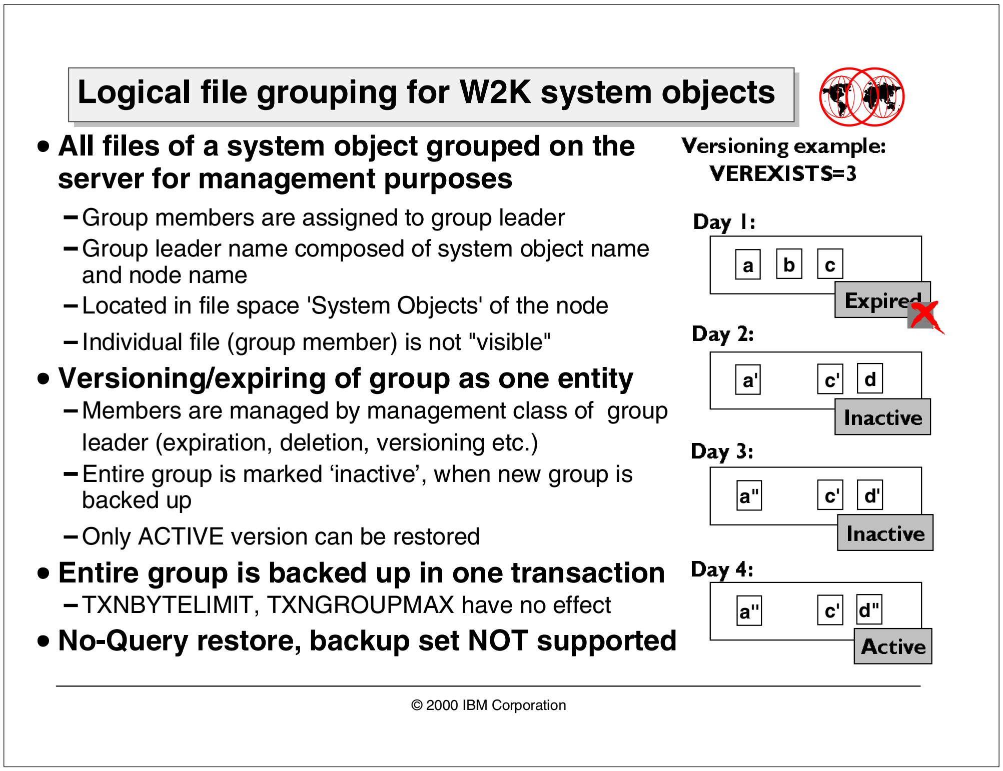

Logical file grouping for W2K system objects

- All files of a system object grouped on the server for management purposes
  - Group members are assigned to group leader
  - Group leader name composed of system object name and node name
  - Located in file space 'System Objects' of the node
  - Individual file (group member) is not "visible"
- Versioning/expiring of group as one entity
  - Members are managed by management class of group leader (expiration, deletion, versioning etc.)
  - Entire group is marked ‘inactive’, when new group is backed up
  - Only ACTIVE version can be restored
- Entire group is backed up in one transaction
  - TXNBYTELIMIT, TXNGROUPMAX have no effect
- No-Query restore, backup set NOT supported

© 2000 IBM Corporation

3.1.2 Characteristics of logical file grouping for W2K system objects

Logical file grouping for Windows 2000 system objects at the Tivoli Storage Manager server has the following characteristics:

1. When a logical group is sent to the server, the previous active version of the group, as well as all the members, are marked “inactive”.
2. A set of files is grouped on the server for management purposes.
3. A group leader is defined with a name composed from the name of the system object and the node name of the sending client. This is the same name each time; it is added to the Tivoli Storage Manager database.
4. Logical file grouping allows for a set of files to belong to a logical group, and they are not valid as individual files outside of the group.
5. All files in the group are located in the same filespace called “system objects”.
6. All members of the group are managed by the management class the “group leader” is bound to.

Deletion of a group such as *delete filespace* and *expiration* processing is managed by the *group leader*. If a group leader is deleted (delete filespace or delete volume), all group members are deleted. With expiry processing, no *group members* are expired until the *group leader* is expired.

All group members are assigned a group leader; and if a member is deleted, for example, through the *delete volume* command in Tivoli Storage Manager server, it will cause all members of the group to be deleted. Expiration processing again will only delete a group if the group leader is eligible for expiration.

All files in the group are backed up as one transaction. TXNBYTELIMIT and TXNGROUPMAX have no effect on this processing. The backup of system objects can be a load on the Tivoli Storage Manager recovery log, as there can be thousands of objects, for example, in the Active Directory.)

At this time, the “No Query Restore” function and backup set generation is not supported for Windows NT and Windows 2000 system objects. Neither can an *inactive* version of a System Object be restored, only the *active* version.

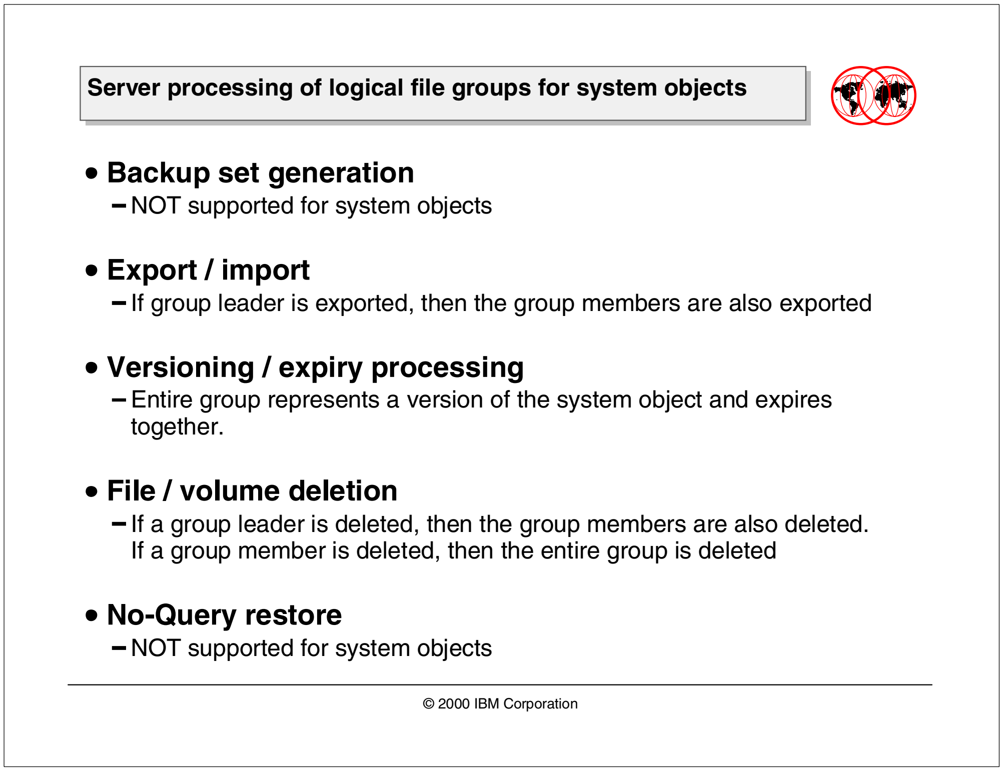

Server processing of logical file groups for system objects

- **Backup set generation**
  - NOT supported for system objects
- **Export / import**
  - If group leader is exported, then the group members are also exported
- **Versioning / expiry processing**
  - Entire group represents a version of the system object and expires together.
- **File / volume deletion**
  - If a group leader is deleted, then the group members are also deleted. If a group member is deleted, then the entire group is deleted
- **No-Query restore**
  - NOT supported for system objects

3.1.3 Function support for logical file groups for system objects

The Tivoli Storage Manager server supports the following commands with logical file grouping of system objects. There is no change to the Tivoli Storage Manager commands.

- Backup sets are not supported for system objects
- Export/import

  Export and import of client nodes is supported. If a group is eligible, then the entire group is exported. The import must take place on a supported Tivoli Storage Manager Server. If it is determined by the client that a group member is missing, then the client issues an error.
- Expiry processing

  The group members of a system object do not expire. The group leader does and acts on the expiry. Therefore, once the group leader is expired and deleted, then the group members associated with this group leader also expire. Expiration processing will only delete a group if the group leader is eligible for expiration.
- File deletion

  If a group leader is deleted when issuing the commands *delete filespace* or *delete volume,* then all the group members are also deleted. If a group member is deleted, it will delete all members of the group.
- No-query restore — not supported
- Incremental paradigm — not supported

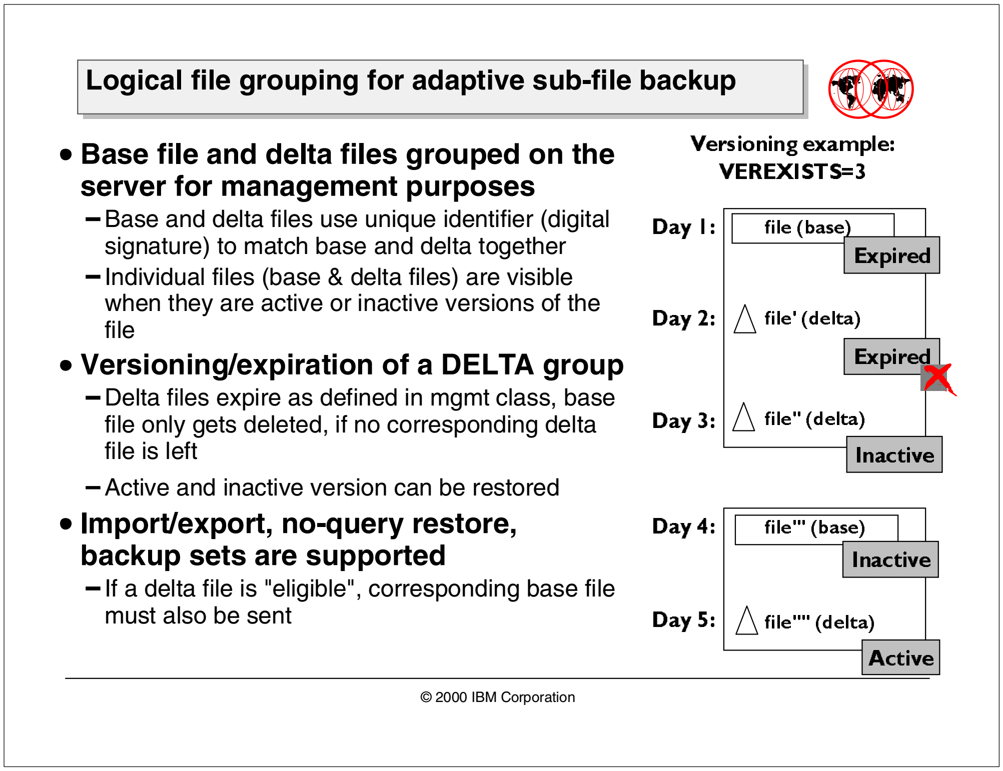

Logical file grouping for adaptive sub-file backup

- **Base file and delta files grouped on the server for management purposes**
  - Base and delta files use unique identifier (digital signature) to match base and delta together
  - Individual files (base & delta files) are visible when they are active or inactive versions of the file
- **Versioning/expiration of a DELTA group**
  - Delta files expire as defined in mgmt class, base file only gets deleted, if no corresponding delta file is left
  - Active and inactive version can be restored
- **Import/export, no-query restore, backup sets are supported**
  - If a delta file is "eligible", corresponding base file must also be sent

Versioning example:
VEREXISTS=3

Day 1: file (base) Expired

Day 2: △ file' (delta) Expired

Day 3: △ file'' (delta) Inactive

Day 4: file''' (base) Inactive

Day 5: △ file'''' (delta) Active

© 2000 IBM Corporation

3.1.4 Characteristics of logical file grouping for adaptive sub-file backup

Adaptive sub-file backup functionality introduced with Tivoli Storage Manager V4.1 allows for mobile users to back up files (the base file), and thereafter, only the pieces of a file that have changed (the delta files).

The Tivoli Storage Manager server manages these files with logical file grouping. The digital signature that is transferred with every sub-file backup from the client to the server matches the different components. Unlike PEER groups, the individual components of these DELTA groups are visible as long as they represent active or inactive versions of the original file. Expiration and versioning information is bound to both delta files and base files.

Expiration processing will expire a *delta* file when it is marked for expiry and will not wait for the group leader to expire. When the *base file* is expired, it will only be deleted if there are no dependent *delta* files. For file deletion, if a *base file* is deleted, then all the *deltas* are also deleted, but if the delete is issued for a *delta* file, then only that *delta* file is removed.

Logical file grouping for adaptive sub-file backup supports “no-query restore”, import/export and backup set generation, and allows for the restore of active and inactive versions.

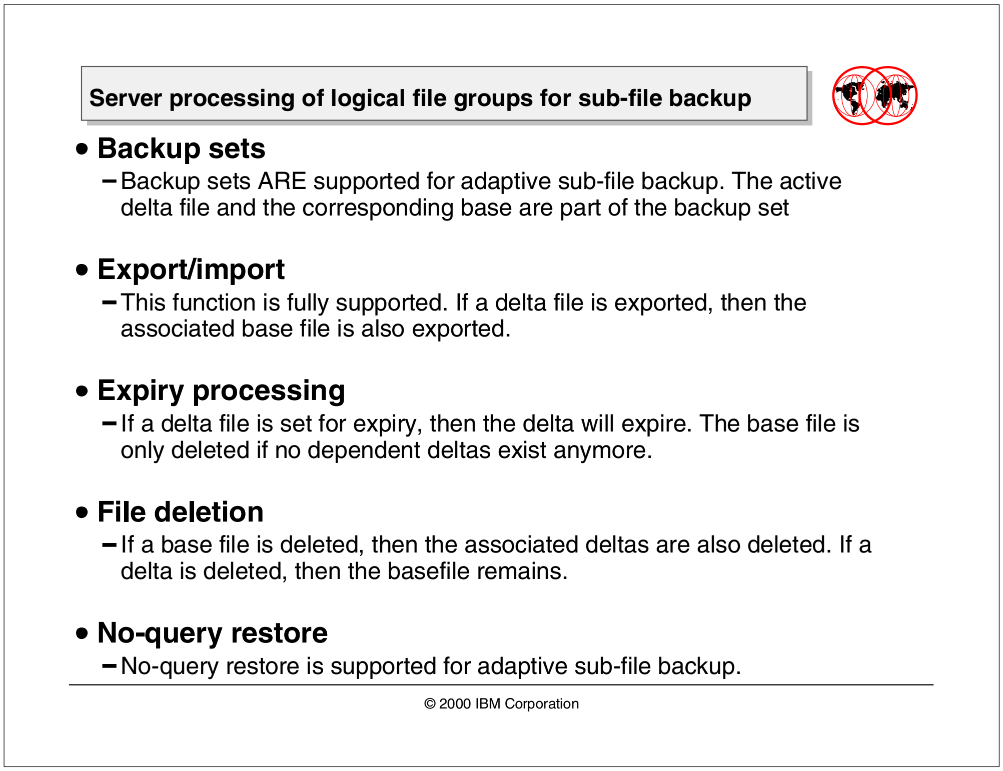

Server processing of logical file groups for sub-file backup

- **Backup sets**
  - Backup sets ARE supported for adaptive sub-file backup. The active delta file and the corresponding base are part of the backup set
- **Export/import**
  - This function is fully supported. If a delta file is exported, then the associated base file is also exported.
- **Expiry processing**
  - If a delta file is set for expiry, then the delta will expire. The base file is only deleted if no dependent deltas exist anymore.
- **File deletion**
  - If a base file is deleted, then the associated deltas are also deleted. If a delta is deleted, then the basefile remains.
- **No-query restore**
  - No-query restore is supported for adaptive sub-file backup.

3.1.5 Function support of logical file groups for adaptive sub-file backup

The Tivoli Storage Manager server supports the following commands with logical file grouping for adaptive sub-file backup.

- Backup sets

  The backup set must contain the base file and the delta file associated with the file version. If a delta file is eligible for the backup set, then the base file will be copied to the backup set.

- Export/import

  Export and Import of client nodes is supported. If a delta file is eligible for export, the base file also must be exported. The import must take place on a Tivoli Storage Manager server supporting adaptive sub-file backup.

- Expiry processing

  Deltas can be expired as needed; they are not dependent on the base file being marked for expiration. The base file is only deleted if there are no deltas associated with the base file.

- File deletion

  If the base file is deleted, then all deltas are deleted; however, if a delta file is deleted, then the base file is not deleted.

- No-query restore

  If a delta file is eligible for a no-query restore, then the corresponding base file is also available and must be sent.

3.2 Tivoli Storage Manager support for TRMM

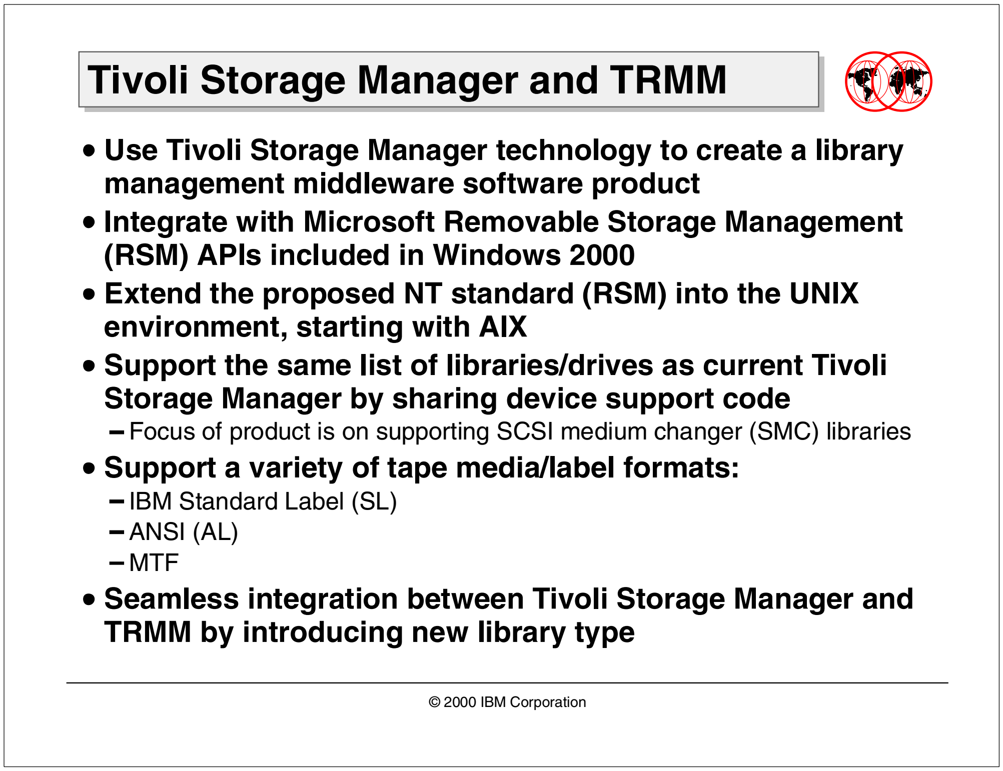

Tivoli Storage Manager and TRMM

- **Use Tivoli Storage Manager technology to create a library management middleware software product**
- **Integrate with Microsoft Removable Storage Management (RSM) APIs included in Windows 2000**
- **Extend the proposed NT standard (RSM) into the UNIX environment, starting with AIX**
- **Support the same list of libraries/drives as current Tivoli Storage Manager by sharing device support code**
  - Focus of product is on supporting SCSI medium changer (SMC) libraries
- **Support a variety of tape media/label formats:**
  - IBM Standard Label (SL)
  - ANSI (AL)
  - MTF
- **Seamless integration between Tivoli Storage Manager and TRMM by introducing new library type**

The major functions for the Tivoli Removable Media Manager (TRMM) is to provide tape library management, tape drive management, and media management, all in an open systems environment. For Version 1 of TRMM, the tape library drive sharing will be among different applications on a single system. Tivoli Storage Manager server needs to be at Version 4.1.

TRMM library management is used to manage and control removable media libraries and to facilitate the sharing of removable media libraries between applications running concurrently on a single Windows NT system. TRMM library management includes:

- The ability to enter and eject removable media from an automated tape libraries and autochangers.
- The ability to mount/dismount removable media onto tape drives contained in an automated tape library or autochanger.
- The ability to inventory the contents of a removable media library
- The ability to audit the contents of removable media libraries to ensure that the removable media cartridges that the software currently believes reside in the library are physically present.

Media Management is the ability to manage, control, and coordinate the use of removable media. For TRMM, this includes:

- The ability to label, identify, and classify media based on readable external labels, machine readable external labels (for example, bar codes), and machine recorded and readable (magnetically, optically) on-media identifier or labels.
- The ability to provide for media pooling, that is, the ability to create logical groupings of physical pieces of removable media to be used for various purposes.
- The ability to track the life cycle of removable media and to recommend replacement of media based on various criteria such as usage or age.

Currently the TRMM V1 support is for Windows/NT4 only and does not support stand-alone drives. TRMM has been developed alongside Microsoft Removable Storage Management (RSM) APIs and Highground Systems, which was formally called NT Media Services (NTMS). RSM is embedded into the Windows 2000 operating system as an operating system service; TRMM intends to move this standard into the UNIX world, starting with AIX.

The object of this section is to give you a brief description of TRMM and how it interfaces with the Tivoli Storage Manager.

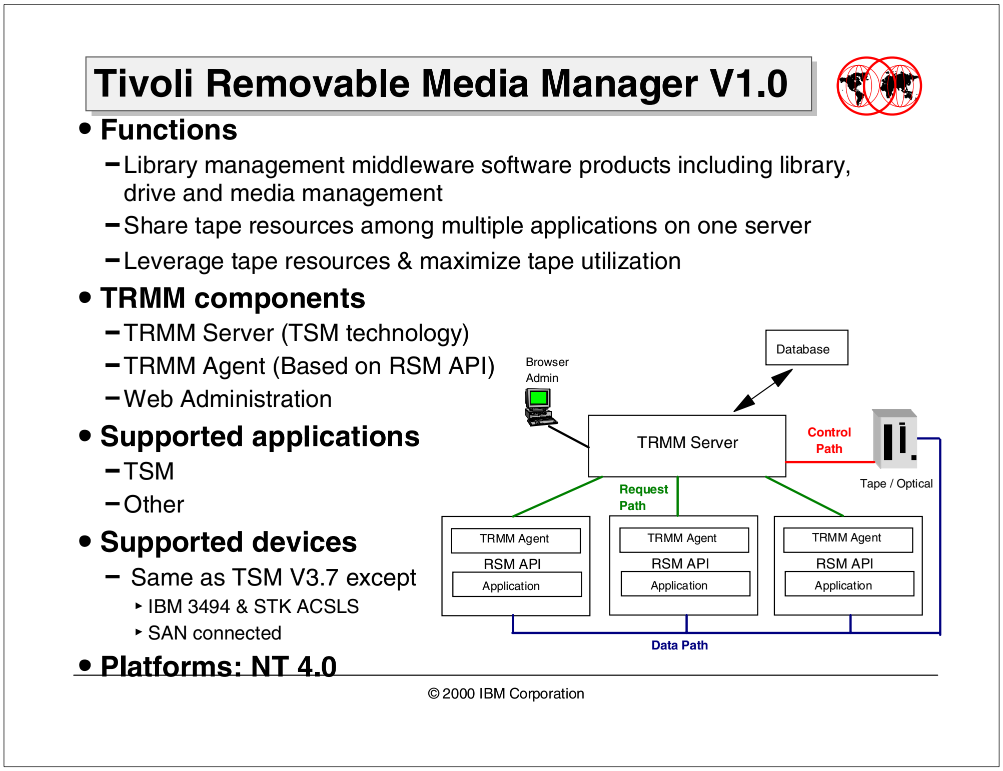

Tivoli Removable Media Manager V1.0

- **Functions**
  - Library management middleware software products including library, drive and media management
  - Share tape resources among multiple applications on one server
  - Leverage tape resources & maximize tape utilization
- **TRMM components**
  - TRMM Server (TSM technology)
  - TRMM Agent (Based on RSM API)
  - Web Administration
- **Supported applications**
  - TSM
  - Other
- **Supported devices**
  - Same as TSM V3.7 except
    - IBM 3494 & STK ACSLS
    - SAN connected
- **Platforms: NT 4.0**

3.2.1 Tivoli Removable Media Manager Version 1.0

Tivoli Removable Media Manager (TRMM) is a software product that provides services for the management, sharing, and control of removable media libraries, drives and media in the Microsoft Windows NT operating system environments. Tivoli Removable Media Manager allows multiple applications to dynamically share removable media resources (drives, media, and automated libraries). This results in more efficient usage of these resources and lower total cost of ownership. It implements the industry standard Removable Storage Management (RSM) API that was created by HighGround Systems and is included in the Microsoft Windows 2000 operating system. Please see Section 5.6, “Removable Storage Manager support” on page 167 for more details on RSM.

3.2.1.1 Function

The functions provided by the TRMM fall into three broad categories:

- Library management and sharing
- Drive management and sharing
- Media management and sharing

For purposes of TRMM, library management is defined as the ability to manage and control removable media libraries and to facilitate the sharing of removable media libraries between applications running concurrently on a single Windows NT system. Drive management and sharing is defined as the ability to manage, control, and coordinate the allocation and use of drives (tape, optical, CDROM, DVD, and so on) between applications. Media management is defined as the ability to manage, control, and coordinate the use of removable media.

3.2.1.2 Components

The major components of TRMM are described below.

TRMM Server — The TRMM server is built on TSM server technology. The server manages the removable media libraries and devices; controls the allocation and deallocation of drives to applications; controls the mounting, demounting, and pooling of volumes; and maintains a database to store any required removable storage information needed.

Database — The TRMM database contains all of the information about the objects tracked by TRMM. This includes physical and logical storage objects, including:

- Removable Media Libraries
- Drives
- Physical Media
- Partitions, Logical Media
- Media Pools And States
- Library Requests
- Operators Requests
- Administrators

The TRMM database is the mechanism used by TRMM to track media related system components, their interrelationships, and their status. The database is a relational database that contains a set of tables corresponding to the objects tracked by TRMM. Information in the TRMM database can be queried via a variety of human and programming interfaces.

TRMM provides utilities to back up and recover the database. The database has an associated log file which is used for logging transactions (inserts, updates, and deletes) performed on the database tables.

TRMM API Agent — The TRMM agent provides an application programming interface (API) that can be invoked by other software products written by independent software vendors (ISVs) and is used to request removable media services from the TRMM Server. The TRMM agent API is equivalent to the NT Media Services (NTMS) or Removable Storage Management (RSM) API provided by Microsoft in the Windows 2000. It is a group of functions that you can use to build a media management application or incorporate the use of removable media into existing applications. The functions include:

- Mount and dismount media
- Clean drives
- Inject and eject media
- Inventory libraries
- Enable and disable libraries, drives, and media
- Access information in the TRMM database

The TRMM agent communicates requests that it receives from application programs for removable media services via the API to the TRMM Server. The TRMM server returns the results of the request to the TRMM agent, who in turn, returns the results of the request/operation to the software application that invoked the API.

TRMM Administrative Web Interface — TRMM provides an administrator interface to manage the TRMM server and the removable media resources being managed/controlled by the TRMM server (Figure 27). This administrator interface is a Web-browser administrator interface similar in nature to the TSM Web administration interface. This interface allows a storage administrator to administer, manage, and query the removable media resources (libraries, drives, and media) under TRMM management, and to manage the TRMM server.

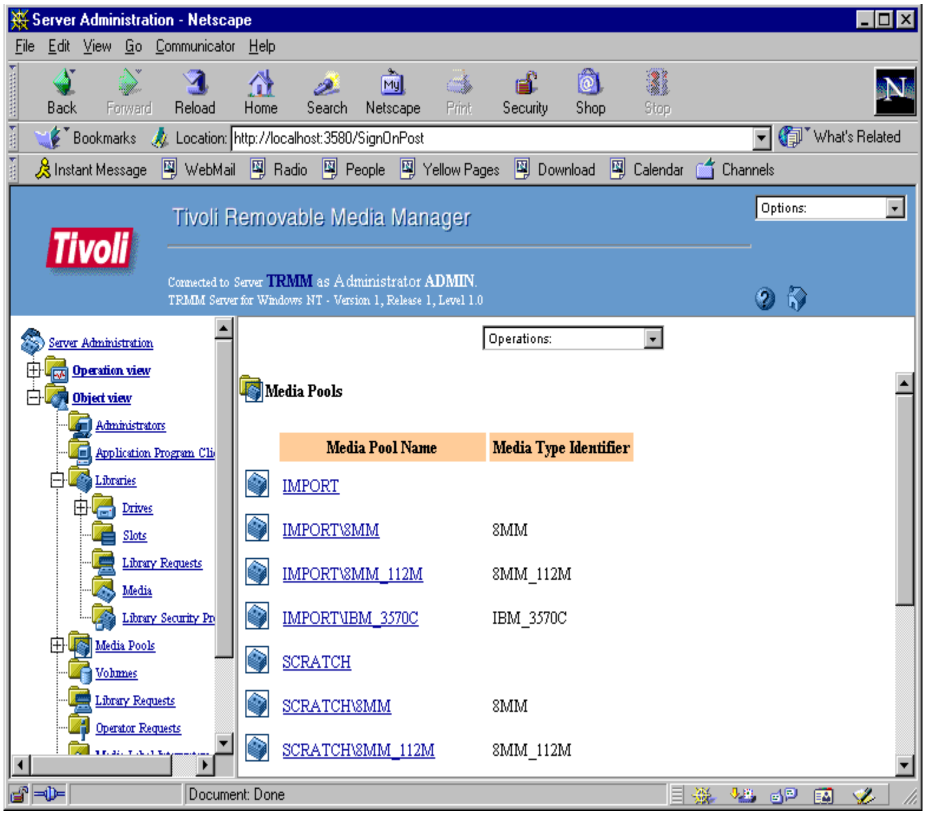
Server Administration - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Shop Stop
Bookmarks Location: http://localhost:3580/SignOnPost
What's Related
Instant Message WebMail Radio People Yellow Pages Download Calendar Channels
Tivoli
Tivoli Removable Media Manager
Options:
Server Administration
Operation view
Object view
Administrators
Application Program Cli
Libraries
Drives
Slots
Library Requests
Media
Library Security Pr
Media Pools
Volumes
Library Requests
Operator Requests
Operations:
Document: Done

Connected to Server TRMM as Administrator ADMIN.
TRMM Server for Windows NT - Version 1, Release 1, Level 1.0

Media Pools

Media Pool Name	Media Type Identifier
IMPORT	
IMPORT\8MM	8MM
IMPORT\8MM_112M	8MM_112M
IMPORT\IBM_3570C	IBM_3570C
SCRATCH	
SCRATCH\8MM	8MM
SCRATCH\8MM_112M	8MM_112M

Figure 27. TRMM Web Administration interface

3.2.1.3 Applications

An application like TSM can utilize TRMM for removable media management functions. TRMM would allow TSM to coexist with another application on the same system that uses the same media resources, such as tape libraries, drives, and volumes.

3.2.1.4 Supported devices and platforms

TRMM initially supports all TSM V3.7.1 supported devices, except IBM 3494 and ACSLS-controlled ATL and SAN-connected tape devices.

In the current implementation, the TRMM server and the TRMM agent reside on the same machine and provide single system sharing of removable media resources. At the time of writing of this book, the only supported platform for TRMM is NT 4.0.

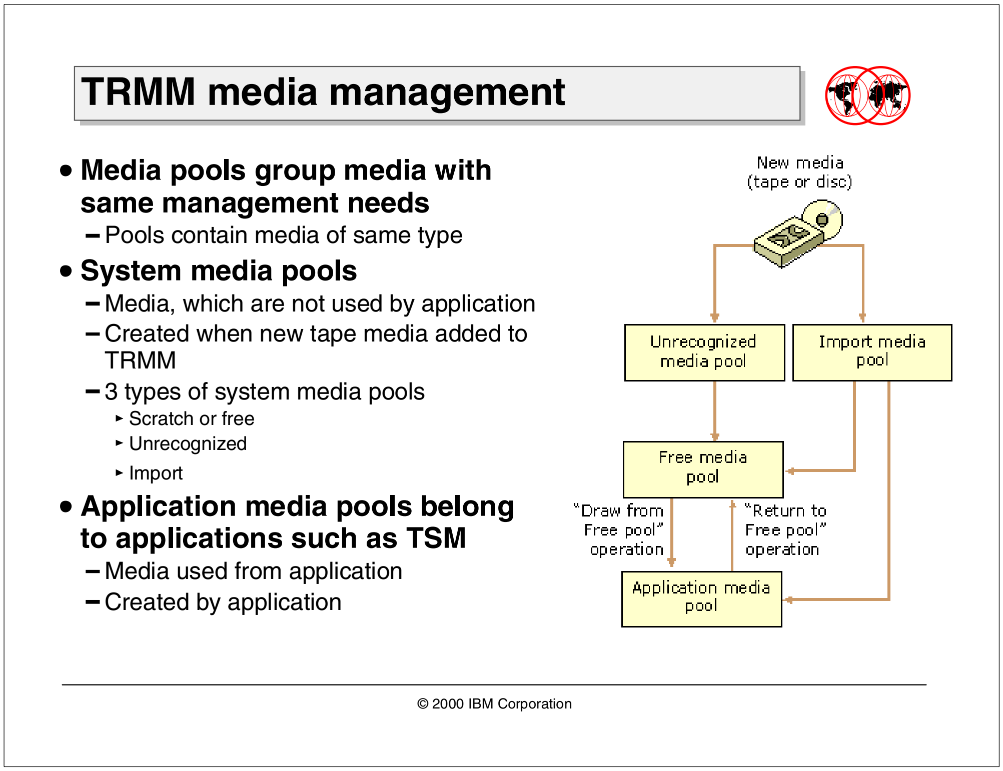

TRMM media management

- **Media pools group media with same management needs**
  - Pools contain media of same type
- **System media pools**
  - Media, which are not used by application
  - Created when new tape media added to TRMM
  - 3 types of system media pools
    - Scratch or free
    - Unrecognized
    - Import
- **Application media pools belong to applications such as TSM**
  - Media used from application
  - Created by application

New media (tape or disc)

Unrecognized media pool

Import media pool

Free media pool

"Draw from Free pool" operation

"Return to Free pool" operation

Application media pool

© 2000 IBM Corporation

3.2.2 TRMM media management

All media managed within TRMM belong to a *media pool*. Media pools are logical collections of media that have the same media management policies. They are used by applications, in this case Tivoli Storage Manager, to control access to media within a TRMM system. A media pool can only contain one type of media. For each type of media, for example, 8MM and 3570 in the same system, there will be media pools allocated to each different media type.

The purpose of media pools is to provide a mechanism for the logical grouping of removable media. Policies and properties can be defined to logical groups. Media groups allows for the media in the same library to have different policies and properties, also allowing a media policy or property to span libraries.

There are two types of media pools within TRMM: *system media pools* and *application media pools*. The system media pools include *scratch* or *free, unrecognized,* and *import* pools. TRMM creates one scratch or free, one unrecognized and one import media pool for each media type. When TRMM is installed and libraries and drives are defined to TRMM, the media pools are allocated automatically.

Application media pools are created by application programs using the TRMM API or via the Web Interface. The application media pool is used by applications to determine which media can be accessed by which applications, and to set policies for that media. Media in that pool is owned by the application.

Scratch or free media pools contain media that are not allocated by an application; TRMM creates one scratch pool for each media type. Applications (Tivoli Storage Manager) can be configured so they can draw from the scratch media pool. An application can be configured to return media to the scratch media pool.

Unrecognized media pools contain blank media (media without labels) and media that have on-media labels which are not known by TRMM. TRMM administrators have to label media in that pool or install the label recognition extension for TRMM.

TRMM puts media in an import pool when it recognizes the format of the on-media labels, but it does not recognize the label-id (volser) of the media. This means the label-id is not cataloged in the TRMM database. Media in an import pool could be media from another physical location, which is introduced to TRMM. TRMM know what the media is and recognizes the on-media labels, but does not have information about that specific piece of media in the TRMM database. The administrator will move that media depending on whether data are stored on it or not, to the free media pool or to an application media pool.

Through the API, physical media are referenced by physical media identifiers (PMIDs). TRMM applications use PMIDs to perform operations normally associated with media such as eject, enable, disable, and delete. Tivoli Storage Manager will use, for example, the media type identifier **IBM_3570C**.

When a Tivoli Storage Manager is defined to TRMM, then an application pool is created called “Tivoli Storage Manager”, as well as a server instance pool.

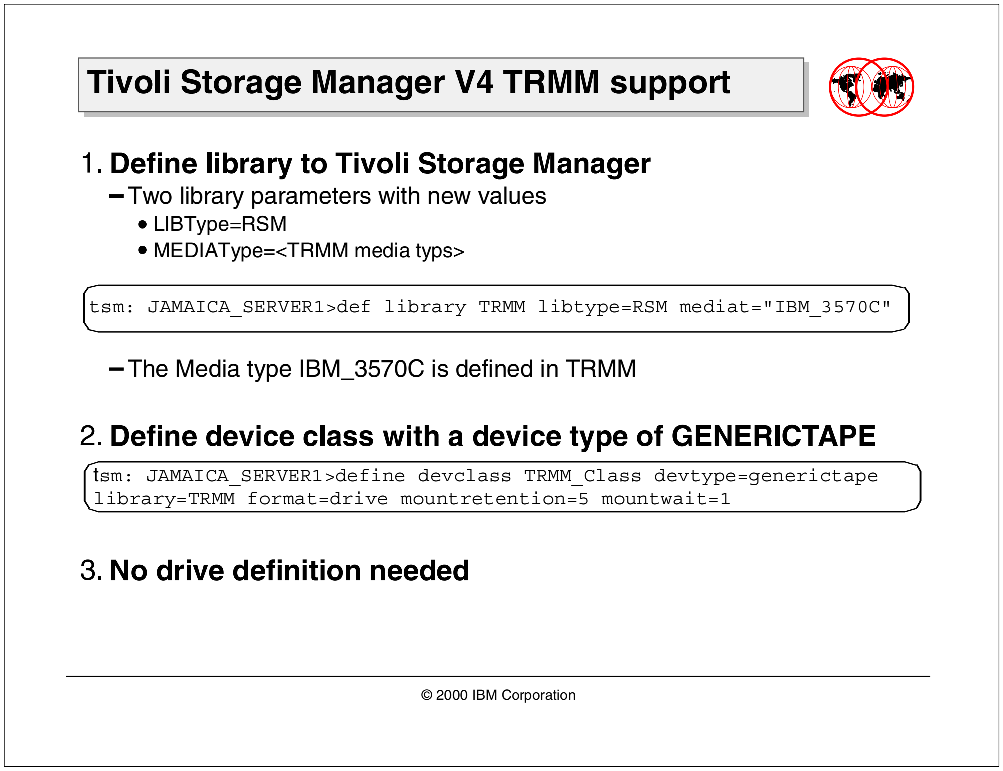

Tivoli Storage Manager V4 TRMM support

1. **Define library to Tivoli Storage Manager**
   - Two library parameters with new values
     - LIBType=RSM
     - MEDIAType=<TRMM media typs>

```
tsm: JAMAICA_SERVER1>def library TRMM libtype=RSM mediat="IBM_3570C"
```

   - The Media type IBM_3570C is defined in TRMM

2. **Define device class with a device type of GENERICTAPE**

```
tsm: JAMAICA_SERVER1>define devclass TRMM_Class devtype=generictape
library=TRMM format=drive mountretention=5 mountwait=1
```

3. **No drive definition needed**

3.2.3 Tivoli Storage Manager V4.1 TRMM support

The Tivoli Storage Manager Server for NT has to be at the V4.1. level. For Tivoli Storage Manager to use TRMM, the following Tivoli Storage Manager definitions must be made:

1. A library must first be defined to Tivoli Storage Manager. The library definition for Tivoli Storage Manager uses a new library type, RSM (see the following example using the `DEFINE LIBRARY` command).

```
tsm: JAMAICA_SERVER1>def library TRMM libtype=RSM mediat=”IBM_3570C”
```

The `MEDIAType` parameter used is a TRMM media type.

When a Tivoli Storage Manager administrator creates and configures a library of the library type RSM, Tivoli Storage Manager triggers TRMM to create a top-level media pool named “Tivoli Storage Manager” and a server instance specific pool.

Under this pool Tivoli Storage Manager will create sub-pools that are media type specific. The first pool will be associated with the library and the second pool will be a media specific import pool. This is detailed in Section 3.2.4, “TRMM and Tivoli Storage Manager Architecture” on page 98. A TRMM media pool is analogous to a directory or folder. The names of the volumes in the pool are listed in the folder. The volumes contain TSM data, and TRMM retains information that maps physical media to devices.

A library of `library` type can be deleted, but it cannot be updated. Changing the media type of an existing library is not allowed.

2. A Tivoli Storage Manager device class associated with this library must be defined using a device type GENERICTAPE as shown in the following screen.

```
tsm: JAMAICA_SERVER1>define devclass TRMM_Class devtype=generictape \
cont>library=TRMM format=drive mountretention=5 mountwait=1 mountl=2
```

   As with external libraries, the mount retention and the mount wait can be adjusted to a more reasonable level than 60 to allow more sharing of TRMM device resources. The `MOUNTLimit` parameter needs to be equal to the number of drives within the library.

3. Tape drives are not defined to Tivoli Storage Manager when using TRMM.

The library type of RSM is much like a library type of EXTERNAL in that it does not manage the volumes as libvols, but depends on TRMM to manage the volumes. As such, the functions of Checkout and Checkin are not supported, but the commands Move Media, Move Drmedia, Query Drmedia, Query Media, and Mount with a Location are supported.

It is strongly recommended that Tivoli Storage Manager administrators utilize a TRMM library only for new devices introduced to the system. Direct media migration from an existing Tivoli Storage Manager library to a new TRMM library is only possible if the Tivoli Storage Manager library drives were previously defined with the GENERICTAPE type. For other formats, please see the administrative command `Move Data` for further details.

Tivoli Storage Manager will still have volume and / or volhist inventory entries for storage pool and database backup volumes. The query volume commands in Tivoli Storage Manager will still work as before.

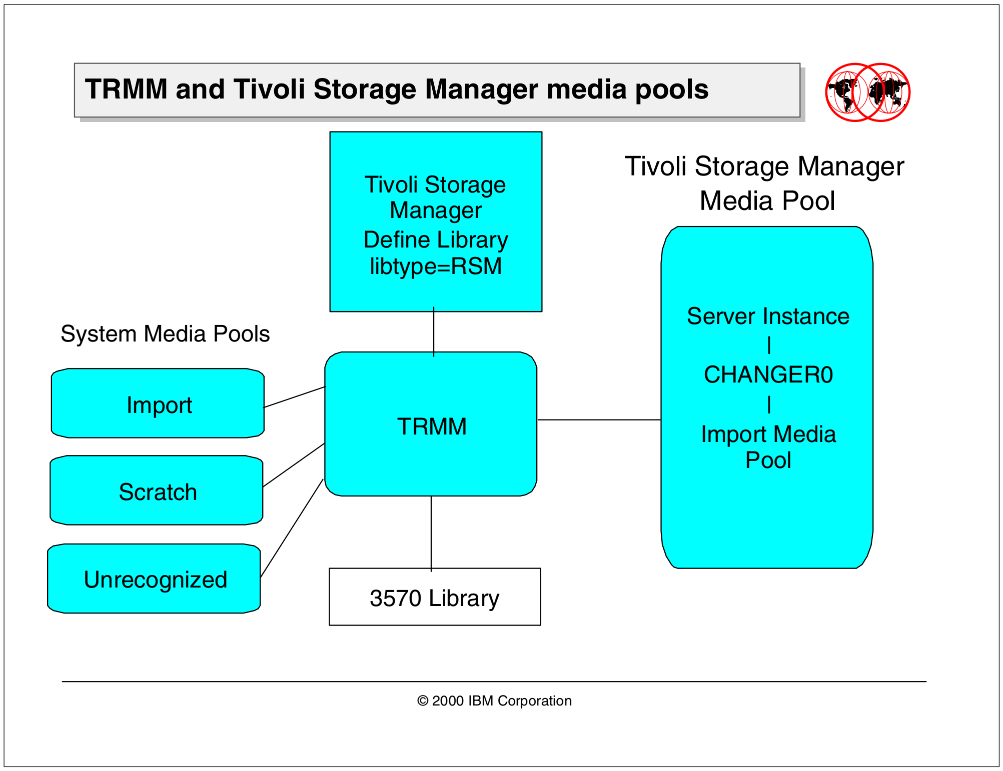

TRMM and Tivoli Storage Manager media pools

Tivoli Storage Manager
Define Library
libtype=RSM

Tivoli Storage Manager
Media Pool

System Media Pools

Import

Scratch

Unrecognized

TRMM

Server Instance
|
CHANGER0
|
Import Media
Pool

3570 Library

© 2000 IBM Corporation

3.2.4 TRMM and Tivoli Storage Manager Architecture

When Tivoli Storage Manager administrators configure a library with a library type of RSM as described in Section 3.2.3, “Tivoli Storage Manager V4.1 TRMM support” on page 96, then TRMM creates a top-level media pool named “Tivoli Storage Manager” (the application media pool as described in Section 3.2.2, “TRMM media management” on page 94) and a server instance specific pool. If there were two Tivoli Storage Manager servers running on the same system then TRMM would create a second server instance specific pool under the Tivoli Storage Manager pool.

Under the Tivoli Storage Manager server instance pool, TRMM will create two other sub-pools. One is an Import pool and the other is the Changer pool. The Import pool is unique for a server instance and represents labeled TSM media which are not assigned to specific storage pools or otherwise used in a library. The Changer pool represents all media under control of an autochanger.

All pools have media specific sub-pools. If, for example, the TIvoli Storage Manager server is using a 3570 and 8mm tape library, multiple Import sub-pools would be created. An instance-specific media type Import pool can be described as a repository for all media of a single media type in one or more libraries that is under the use of one specific Tivoli Storage Manager server instance.

The Changer pool corresponds to a specific library defined in Tivoli Storage Manager and can span one or multiple physical libraries. All volumes in that library that belong to the Tivoli Storage Manager server reside in the Changer pool. These volumes include both onsite and offsite volumes as well as database backup or storage pool volumes.

Volumes are moved to Changer pools either from the systems scratch or free pool in the case of a Tivoli Storage Manager scratch mount request, or from the instance Import pool in the case of a Tivoli Storage Manager private volume request. Volumes will stay in the Changer pool until Tivoli Storage Manager expires or deletes them. The expired or deleted volumes automatically go from the Changer pool to the Scratch pool for the media type.

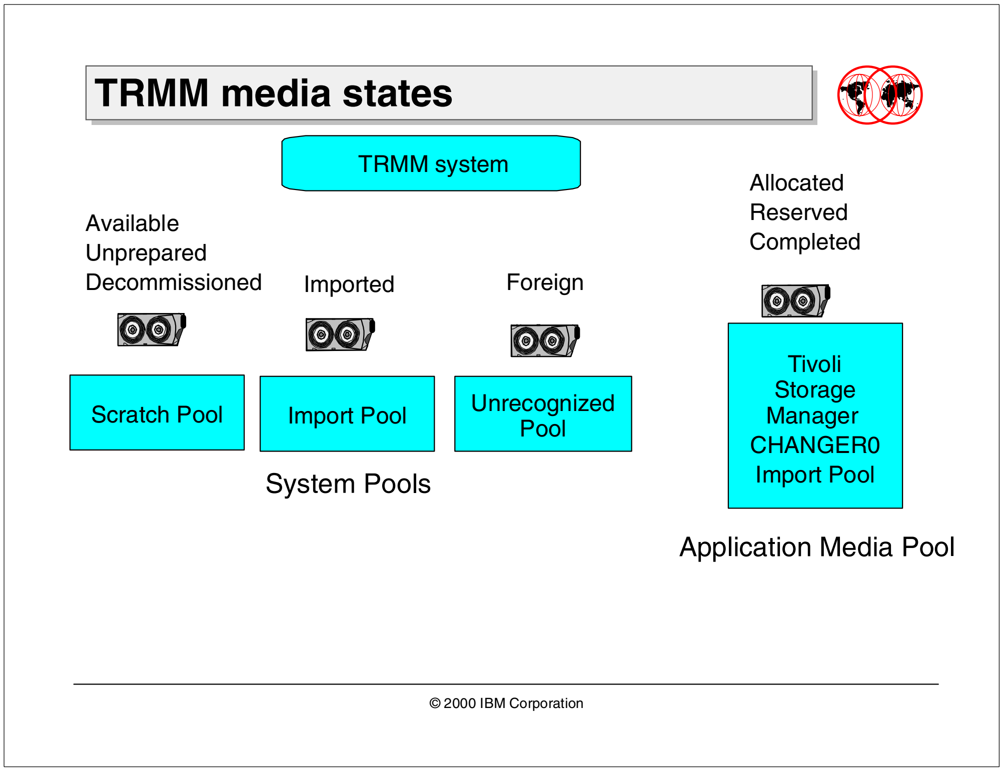

TRMM media states

TRMM system

Available
Unprepared
Decommissioned

Imported

Foreign

Allocated
Reserved
Completed

Scratch Pool

Import Pool

Unrecognized Pool

System Pools

Tivoli Storage Manager CHANGER0 Import Pool

Application Media Pool

© 2000 IBM Corporation

3.2.5 TRMM media states

Media states in TRMM identify the condition of each piece of media in the system and defines which applications have access to it. When media is introduced to TRMM, TRMM looks at the on-media labels and bar codes to determine the states of the media. This is known as classification. There are nine media states in TRMM:

- Allocated — Specifies that the media is in use and assigned to an application media pool, for example, Tivoli Storage Manager.
- Available — Specifies that the media is ready for use by applications.
- Completed — Specifies that the media is in use and assigned to application media pools, but can no longer be used for write operations.
- Decommissioned — Specifies that the media cannot be used again.
- Unrecognized — Specifies that the media label types and label IDs are not recognized by TRMM, or the media is new and does not contain any on-media labels.
- Imported — Specifies that the media label types are recognized, but label IDs are not recognized.
- Incompatible — Specifies that the media is not compatible with the library or drive in which they were classified.
- Reserved — Specifies that the media are held by an application.
- Unprepared — Specifies that the media are in a scratch media pool, but have not been assigned scratch labels by TRMM.

Media in the Import state must be changed to Available or Allocated by the application (for example, Tivoli Storage Manager), which means that Tivoli Storage Manager can either assign the media to a system scratch pool or to the application pool.

Media in the Foreign state can be moved to an Available or Allocated state either by an application, or manually by the TRMM administrator.

Media in the Allocated state can be changed to the Available, Completed, or Reserved state. If the application returns the media with an Allocated state to an Available state, the media returns to scratch. If the application changes the state of the media from Allocated to Completed or Reserved, the media remains in the Application media pool.

3.3 Server scripting and scheduling enhancements

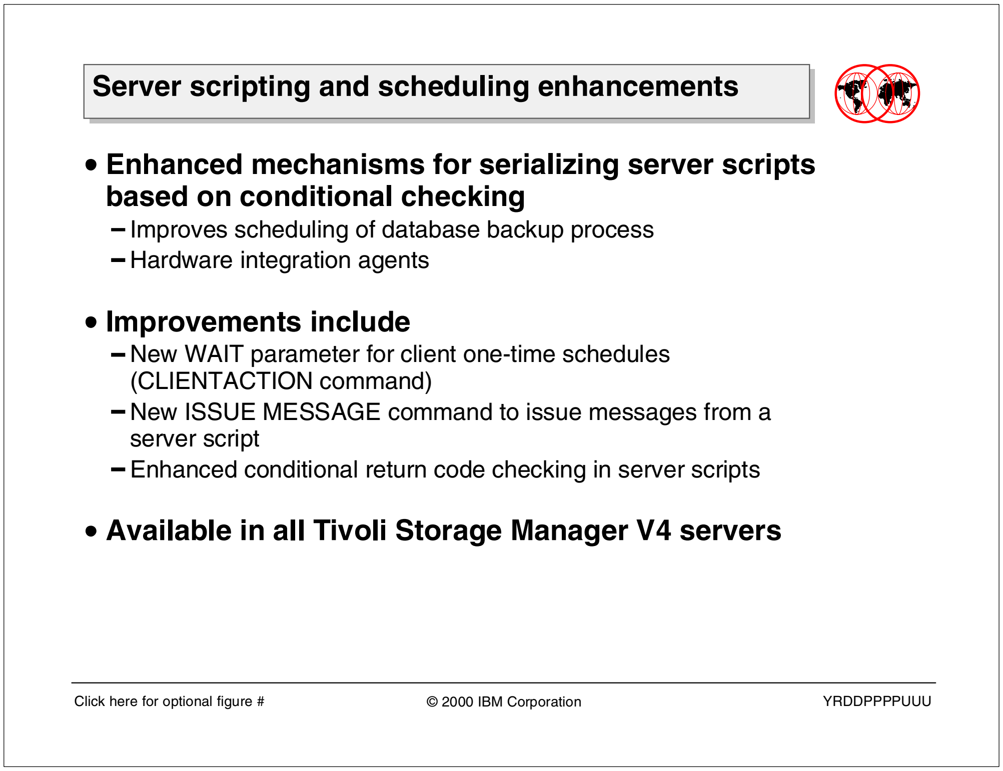

Server scripting and scheduling enhancements

- **Enhanced mechanisms for serializing server scripts based on conditional checking**
  - Improves scheduling of database backup process
  - Hardware integration agents
- **Improvements include**
  - New WAIT parameter for client one-time schedules (CLIENTACTION command)
  - New ISSUE MESSAGE command to issue messages from a server script
  - Enhanced conditional return code checking in server scripts
- **Available in all Tivoli Storage Manager V4 servers**

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

The Tivoli Storage Manager scripting and scheduling enhancements are available in V4.1 of the server. They were initially developed to support IBM Enterprise Storage Server (ESS) Flashcopy or EMC Symmetrix BCVs backup of mirrored volumes. However, these new enhancements can benefit the execution of other scripts in order to introduce event driven scheduling.

For example, the automation of database backup scripts could benefit from these enhanced commands. The simplest way of scheduling an automated backup is to write a short script to perform the backup. Then, use Tivoli Storage Manager central scheduler to define a schedule to issue a command to execute the script.

Before, when running this type of scheduled script, there was no ability to let actions complete (such as quiescing a database) before backup could commence. With these new enhancements, a schedule can now be defined to “wait” on another event (*event driven scheduling*). User-defined messages can be issued from the script, to show the condition on completion of the script or macro.

The enhancements include:

- A new `WAIT` parameter for client one-time schedules.
- A new `ISSUE MESSAGE` command to create user-defined messages in a server script
- Enhanced conditional return code checking in server scripts

The objective of this section is to describe the new scripting and scheduling enhancements and to give an example of how these could be used.

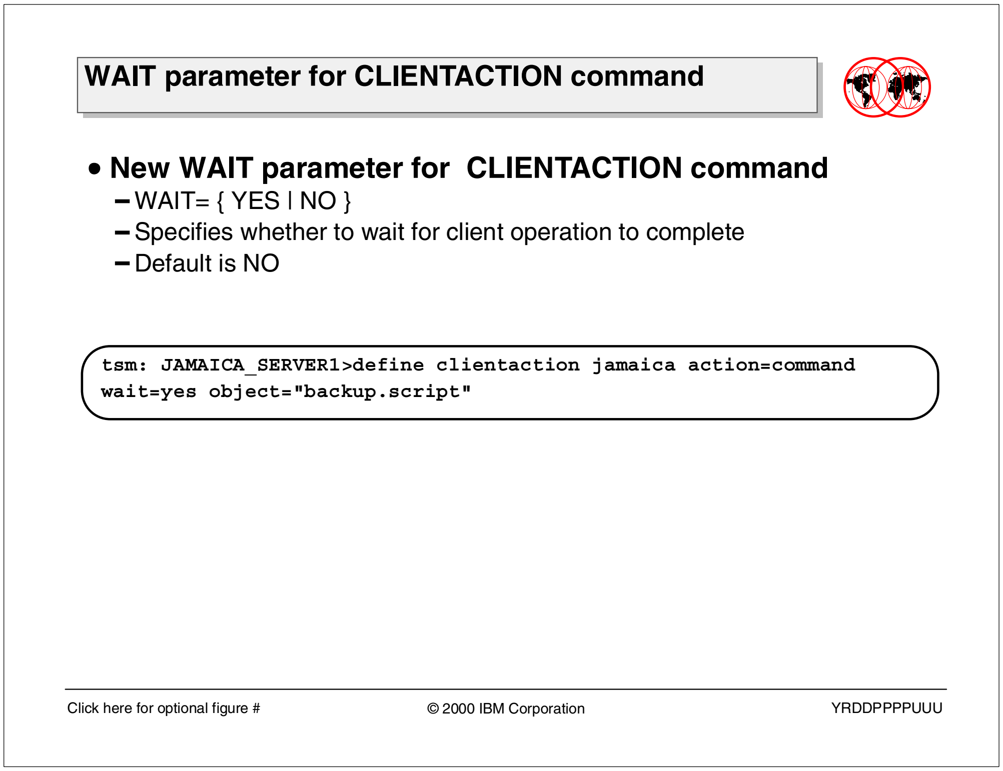

WAIT parameter for CLIENTACTION command

- **New WAIT parameter for CLIENTACTION command**
  - WAIT= { YES | NO }
  - Specifies whether to wait for client operation to complete
  - Default is NO

```
tsm: JAMAICA_SERVER1>define clientaction jamaica action=command
wait=yes object="backup.script"
```

Click here for optional figure #   © 2000 IBM Corporation   YRDDPPPPUUU

3.3.1 WAIT parameter for CLIENTACTION command

This specifies whether to wait for a scheduled client operation to complete. This allows for the CLIENTACTION command to complete before processing the next step. This is a synchronous operation and provides the serialization in a script.

You cannot issue the DEFINE CLIENTACTION command with WAIT=YES from the server console. However, from the server console, you can:

- Specify WAIT=YES with DEFINE CLIENTACTION as the command line of a DEFINE SCRIPT command.
- Specify WAIT=YES with DEFINE CLIENTACTION as the command line of a file whose contents will be read into the script that is defined by a DEFINE SCRIPT command.

```
tsm: JAMAICA_SERVER1>define clientaction jamaica action=command \
cont>wait=yes object="backup.script"
```

.

Note

If you specify the DEFINE CLIENTACTION command with WAIT=YES in a macro, the immediate schedules defined by the command will not roll-back if the macro does not complete successfully.

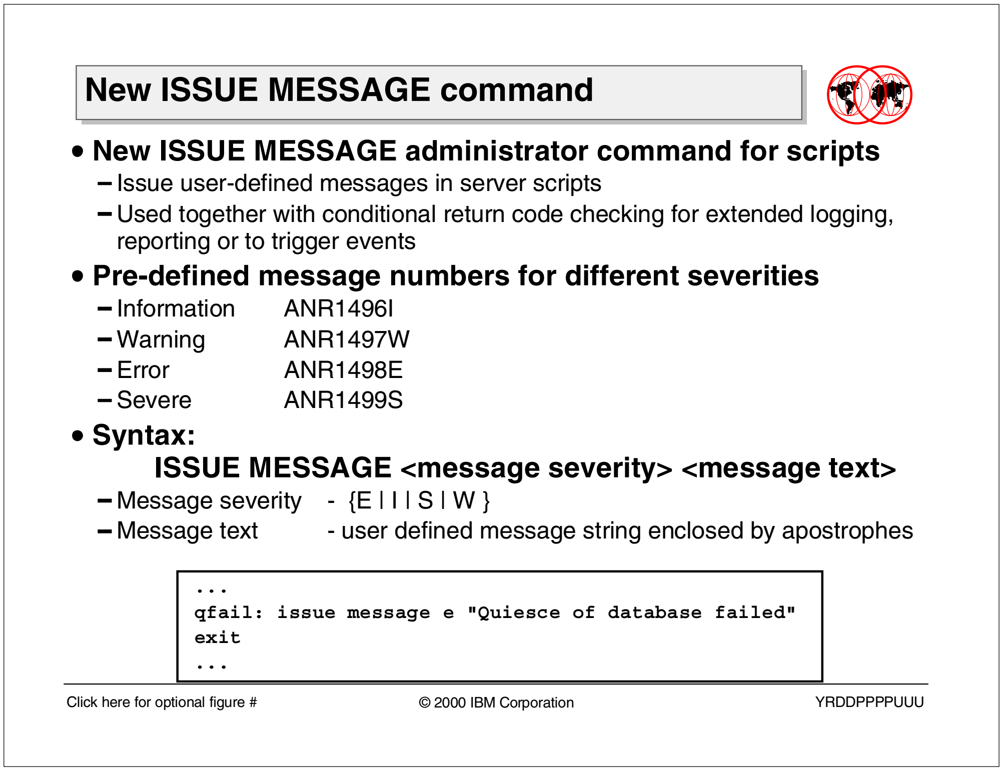

New ISSUE MESSAGE command

- **New ISSUE MESSAGE administrator command for scripts**
  - Issue user-defined messages in server scripts
  - Used together with conditional return code checking for extended logging, reporting or to trigger events
- **Pre-defined message numbers for different severities**

– Information	ANR1496I
– Warning	ANR1497W
– Error	ANR1498E
– Severe	ANR1499S

- **Syntax:**

  ISSUE MESSAGE <message severity> <message text>

– Message severity	- {E | I | S | W }
– Message text	- user defined message string enclosed by apostrophes

```
...
qfail: issue message e "Quiesce of database failed"
exit
...
```

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

3.3.2 New ISSUE MESSAGE command

The objective of this new parameter is to allow for user-defined messages to be issued from a server script to determine, if there is a problem with the command in the script. This allows for the server script to indicate the outcome of an action. Message numbers are predefined for specific severity levels.

The new Tivoli Storage Manager scripting command is:

- `ISSUE MESSAGE message_severity message_text`

  Where:

  message_severity specifies the severity of the message. The message severity indicators are:

  - E = Error. ANR1498E is displayed in the message text.
  - I = Information. ANR1496I is displayed in the message text.
  - S = Severe. ANR1499S is displayed in the message text.
  - W = Warning. ANR1497W is displayed in the message text.

  Where:

  message_text specifies the description of the message.

The `ISSUE MESSAGE` command can only be used in conjunction with the return code function and can only be used in a script. It can be used in a server script to report about the result of a quiesce database operation. For example, a script results in a non-zero return code. Use the `ISSUE MESSAGE` command to determine the severity of the return code and the message description.

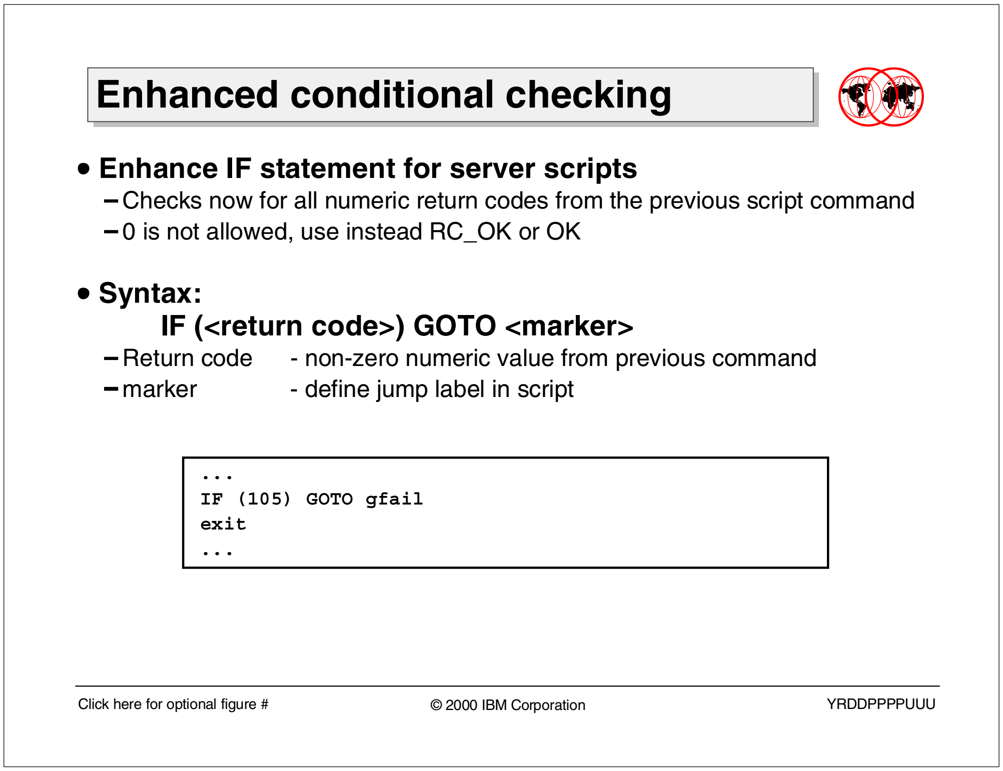

Enhanced conditional checking

- **Enhance IF statement for server scripts**
  - Checks now for all numeric return codes from the previous script command
  - 0 is not allowed, use instead RC_OK or OK

- **Syntax:**

  IF (<return code>) GOTO <marker>

  - Return code - non-zero numeric value from previous command
  - marker - define jump label in script

```
...
IF (105) GOTO gfail
exit
...
```

Click here for optional figure #

© 2000 IBM Corporation

YRDDPPPPUUU

3.3.3 Enhanced conditional checking

In previous versions of Tivoli Storage Manager it was possible to write server scripts that included return codes. As each command is processed in a script, the return code can be used to determine how processing of the script should proceed based on the return code value. The return codes can be one of three severities: OK, WARNING, ERROR.

Tivoli Storage Manager scripts use the symbolic return code for processing, not the numeric value. The numeric values are displayed by the administrative client when a command is run.

The Tivoli Storage Manager conditional ‘if’ enhancement can now evaluate a single numeric return code in a script statement.

The syntax is:

```
if(<rc>)
```

`rc` = is a single numeric value

The rc zero (0) cannot be coded, but instead the “RC_OK” or “OK” can be used.

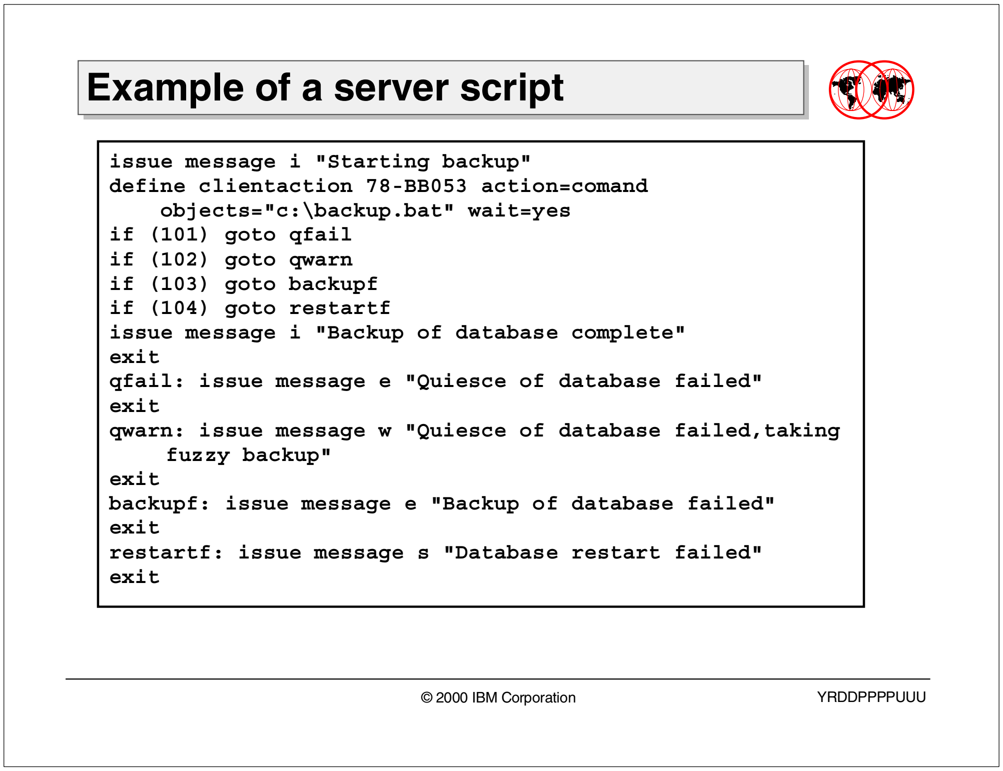

Example of a server script

```
issue message i "Starting backup"
define clientaction 78-BB053 action=comand
    objects="c:\backup.bat" wait=yes
if (101) goto qfail
if (102) goto qwarn
if (103) goto backupf
if (104) goto restartf
issue message i "Backup of database complete"
exit
qfail: issue message e "Quiesce of database failed"
exit
qwarn: issue message w "Quiesce of database failed,taking
     fuzzy backup"
exit
backupf: issue message e "Backup of database failed"
exit
restartf: issue message s "Database restart failed"
exit
```

© 2000 IBM Corporation YRDDPPPPUUU

3.3.4 Example of a server script

The above is an example of how the new enhancements could be used.

In the example we are issuing a message stating that we are starting a backup. The script then goes on to define a client action in which we run the script `backupscript` but with a `wait=yes`. The wait coded in this `define clientaction` command will mean the script will stop until the `backupscript` has finished.

The return codes ‘101’, ‘102’, ‘103’ and ‘104’ are coded in the `backupscript`. This shows the `if` conditional checking. If the `backupscript` should issue one of the return codes contained in the script, then the script will exit, and will issue a message depending on the return code.

The above use of these enhancements in a server script will be beneficial to database backups that are required to be shut down before backing up the database and subsequently starting it up again.

3.4 New DSMSERV LOADFORMAT command

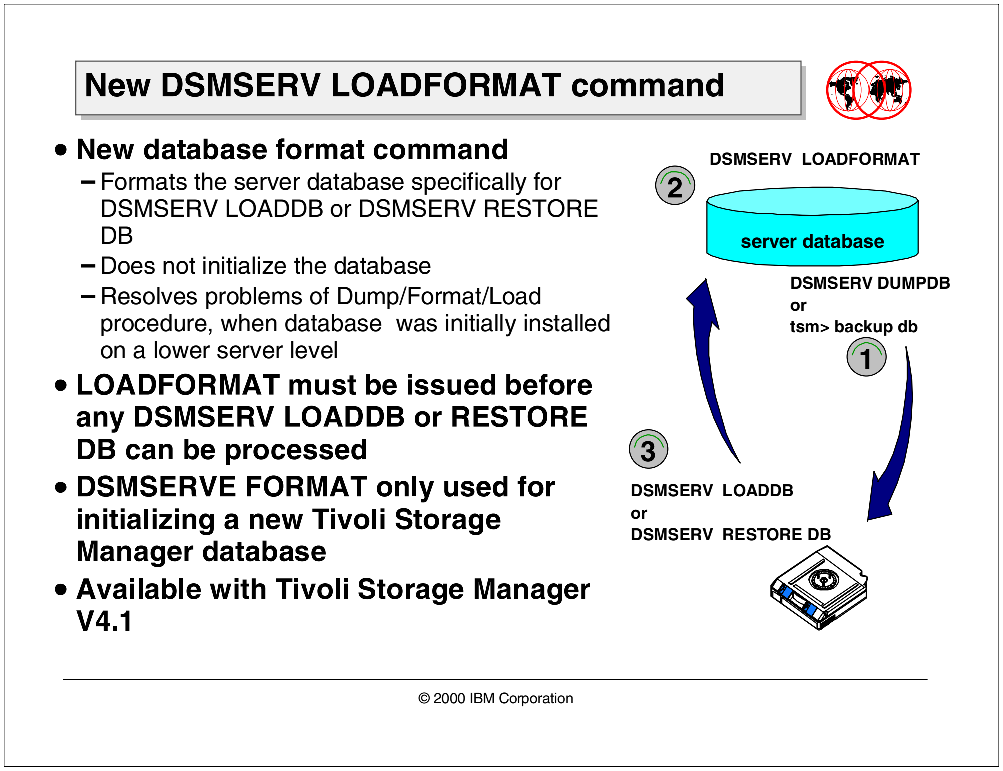

New DSMSERV LOADFORMAT command

- **New database format command**
  - Formats the server database specifically for DSMSERV LOADDB or DSMSERV RESTORE DB
  - Does not initialize the database
  - Resolves problems of Dump/Format/Load procedure, when database was initially installed on a lower server level
- **LOADFORMAT must be issued before any DSMSERV LOADDB or RESTORE DB can be processed**
- **DSMSERVE FORMAT only used for initializing a new Tivoli Storage Manager database**
- **Available with Tivoli Storage Manager V4.1**

© 2000 IBM Corporation

Tivoli Storage Manager V4.1 introduces a new DSMSERV LOADFORMAT command. In the following section, we will detail when this command is be used.

3.4.1 DSMSERV LOADFORMAT description

The new DSMSERV LOADFORMAT command is replacing the DSMSERV FORMAT command when used in conjunction with DSMSERV LOADDB or DSMSERV RESTORE DB.

The Dump/Format/Load procedure can create problems when the Tivoli Storage Manager database was initially installed at a lower version level (for example, ADSM V2.1, V3.1 or V3.7), and when new defaults have been added. If new default values are added or the default values are changed in a later version of the Tivoli Storage Manager code, there can be a problem when DSMSERV LOADDB or DSMSERV RESTORE DB is run on a Tivoli Storage Manager server that has been constantly upgraded with the DSMSERV UGRADEDB when migrating to new Tivoli Storage Manager versions.

The DSMSERV FORMAT will create and format the database. It also initializes default values, and DSMSERV UPGRADEDB might under some circumstances not create the new values which may come with the new code versions.

3.4.2 DSMSERV LOADFORMAT operation

The DSMSERV LOADFORMAT command is required for formatting a Tivoli Storage Manager database before the DSMSERV LOADDB command. It can also be used for formatting a Tivoli Storage Manager database before the DSMSERV RESTORE DB command. It is important to remember that the DSMSERV LOADFORMAT command cannot be used for formatting a Tivoli Storage Manager database for a new server; the DSMSERV FORMAT command must be used.

The DSMSERV LOADDB command will only be allowed after the DSMSERV LOADFORMAT has been run. After the DSMSERV LOADFORMAT command has been issued, a DSMSERV LOADDB or DSMSERV RESTORE is required before the Tivoli Storage Manager server will start.

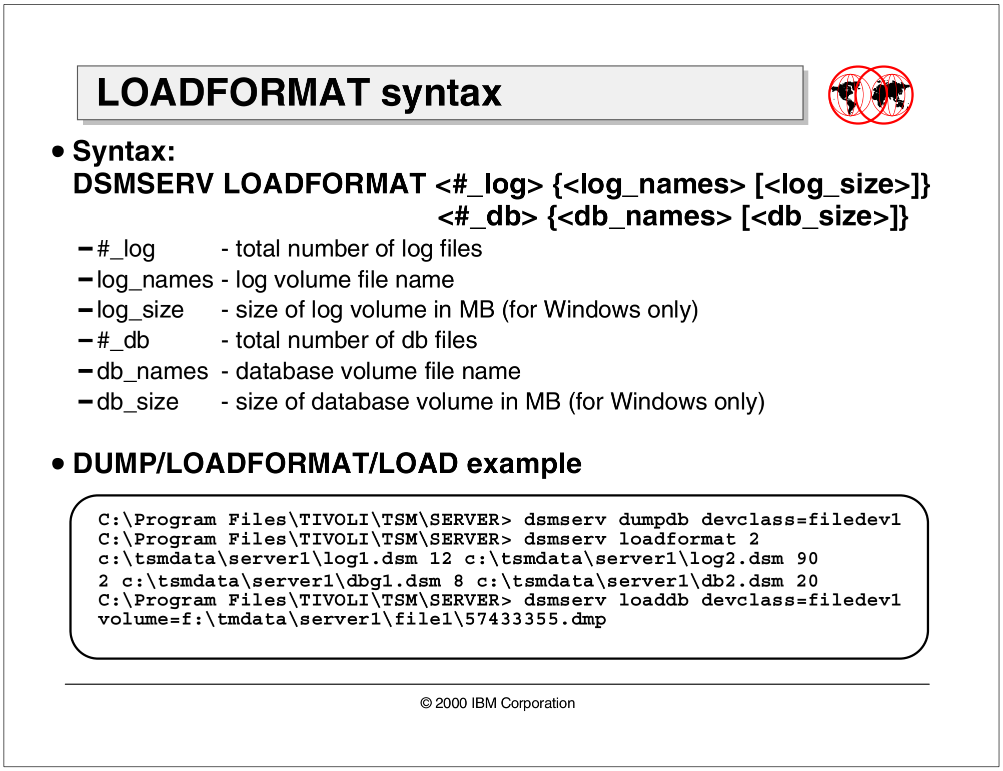

LOADFORMAT syntax

- **Syntax:**
  DSMSERV LOADFORMAT <#_log> {<log_names> [<log_size>]}
  <#_db> {<db_names> [<db_size>]}
  - #_log - total number of log files
  - log_names - log volume file name
  - log_size - size of log volume in MB (for Windows only)
  - #_db - total number of db files
  - db_names - database volume file name
  - db_size - size of database volume in MB (for Windows only)
- **DUMP/LOADFORMAT/LOAD example**

```
C:\Program Files\TIVOLI\TSM\SERVER> dsmserv dumpdb devclass=filedev1
C:\Program Files\TIVOLI\TSM\SERVER> dsmserv loadformat 2
c:\tsmdata\server1\log1.dsm 12 c:\tsmdata\server1\log2.dsm 90
2 c:\tsmdata\server1\dbg1.dsm 8 c:\tsmdata\server1\db2.dsm 20
C:\Program Files\TIVOLI\TSM\SERVER> dsmserv loaddb devclass=filedev1
volume=f:\tmdata\server1\file1\57433355.dmp
```

3.4.3 LOADFORMAT syntax

This section will show an example of the new DSMSERV LOADFORMAT command in a demonstration of when this command should be used.

3.4.3.1 DSMSERV DUMPDB

If there is a catastrophic error in the Tivoli Storage Manager database, then the database uses the DSMSERV DUMPDB salvage utility. To give an example of the new DSMSERV LOADFORMAT facility, we did a DSMSERV DUMPDB.

```
C:\Program Files\TIVOLI\TSM\SERVER>dsmserv dumpdb devclass=filedev1
```

3.4.3.2 DSMSERV LOADFORMAT syntax

Before any Tivoli Storage Manager DSMSERV LOADDB operation can be issued by an administrator, the new DSMSERV LOADFORMAT command has to be issued.

The DSMSERV LOADFORMAT command syntax is:

```
C:\Program Files\TIVOLI\TSM\SERVER>dsmserv loadformat 2
c:\tsmdata\server1\log1.dsm 12 c:\tsmdata\server1\log2.dsm 90 2
c:\tsmdata\server1\db1.dsm 8 c:\tsmdata\server1\db2.dsm 20
```

Tivoli Storage Manager is now ready for the DSMSERV LOADDB command.

3.4.3.3 DSMSERV LOADDB

Once the DSMSERV LOADFORMAT command has completed successfully, the DSMSERV LOADDB command can be issued.

The syntax for the command is:

```
C:\Program Files\TIVOLI\TSM\SERVER>dsmserv loaddb devclass=filedev1
volume=f:\tmdata\server1\file1\57433355.dmp
```

Watch the results of the DSMSERV LOADDB command, as sometimes an audit of the Tivoli Storage Manager database is required. In our case, an audit was required.

Chapter 4. Tivoli integration enhancements

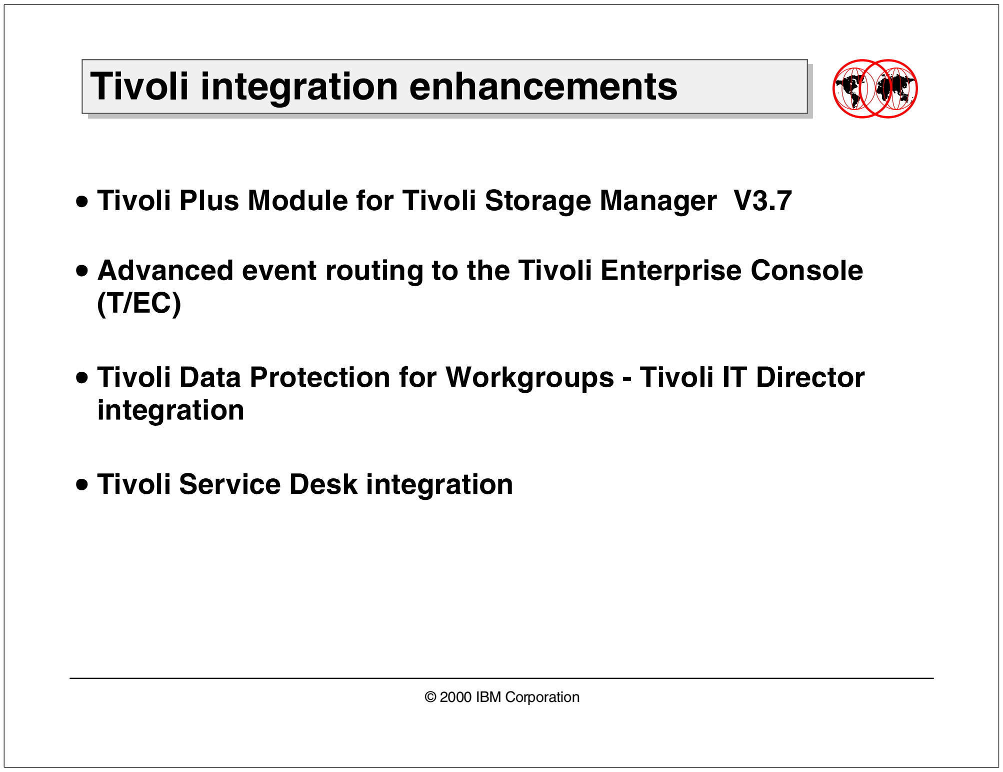

Tivoli integration enhancements

- Tivoli Plus Module for Tivoli Storage Manager V3.7
- Advanced event routing to the Tivoli Enterprise Console (T/EC)
- Tivoli Data Protection for Workgroups - Tivoli IT Director integration
- Tivoli Service Desk integration

© 2000 IBM Corporation

This chapter covers the major Tivoli Storage Manager administration enhancements:

- Enhancements to the Tivoli Plus Module for Tivoli Storage Manager V3.7
- Enhanced event routing to the Tivoli Enterprise Console (T/EC) using new unique event classes for the Tivoli Storage Manager server, client, API and Tivoli Data Protection applications. This does not include Tivoli Data Protection for Workgroups.
- Tivoli IT Director integration with Data Protection for Workgroups.
- Enhanced Tivoli Service Desk integration.

4.1 Tivoli Plus Module for Tivoli Storage Manager V3.7

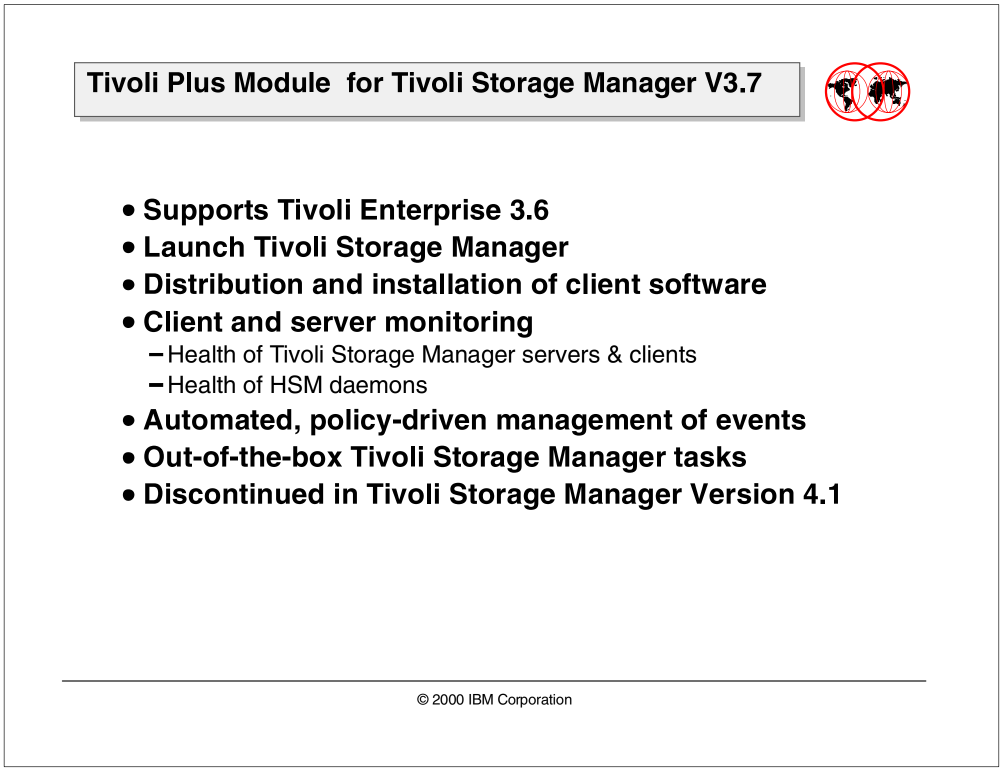

Tivoli Plus Module for Tivoli Storage Manager V3.7

- Supports Tivoli Enterprise 3.6
- Launch Tivoli Storage Manager
- Distribution and installation of client software
- Client and server monitoring
  - Health of Tivoli Storage Manager servers & clients
  - Health of HSM daemons
- Automated, policy-driven management of events
- Out-of-the-box Tivoli Storage Manager tasks
- Discontinued in Tivoli Storage Manager Version 4.1

© 2000 IBM Corporation

The Tivoli Plus Module provides integration of the Tivoli Storage Manager software into the Tivoli Enterprise environment.

The Tivoli Plus Module for Tivoli Storage Manager V3.7 is a new module. Older versions of the Tivoli Plus Module must be removed and then the new module can be installed. This version of the Tivoli Plus Module will only work with Tivoli Storage Manager V3.7 and not with any earlier versions of ADSM / Tivoli Storage Manager.

The Tivoli Plus module provides central management of the Tivoli Storage Manager application across a multi-platform network. Including:

- Icons for launching the Tivoli Storage Manager application
- Subscription lists for clients and servers
- Monitors for Tivoli Management Environment 10 Distributed Monitoring

TME 10 Distributed Monitoring must be installed and configured before the Tivoli Plus Module is installed. TME 10 Enterprise Console (T/EC) must be installed to use T/EC functions with the Tivoli Plus Module. If the Tivoli Plus Module is installed before any of these TME 10 applications, the Plus Module will need to be re-installed before any of these features are operational.

The Tivoli Plus Module V3.7 uses a *log file adapter* and reads messages from the Tivoli log file which then parses and builds the Tivoli event descriptors. These are then forwarded to the T/EC. The log file adapter only recognizes a few hundred of the thousands of Tivoli Storage Manager messages. Tivoli Storage Manager now provides better integration into the TME 10 environment than the Tivoli Plus Module provides.

4.2 Enhanced T/EC event routing

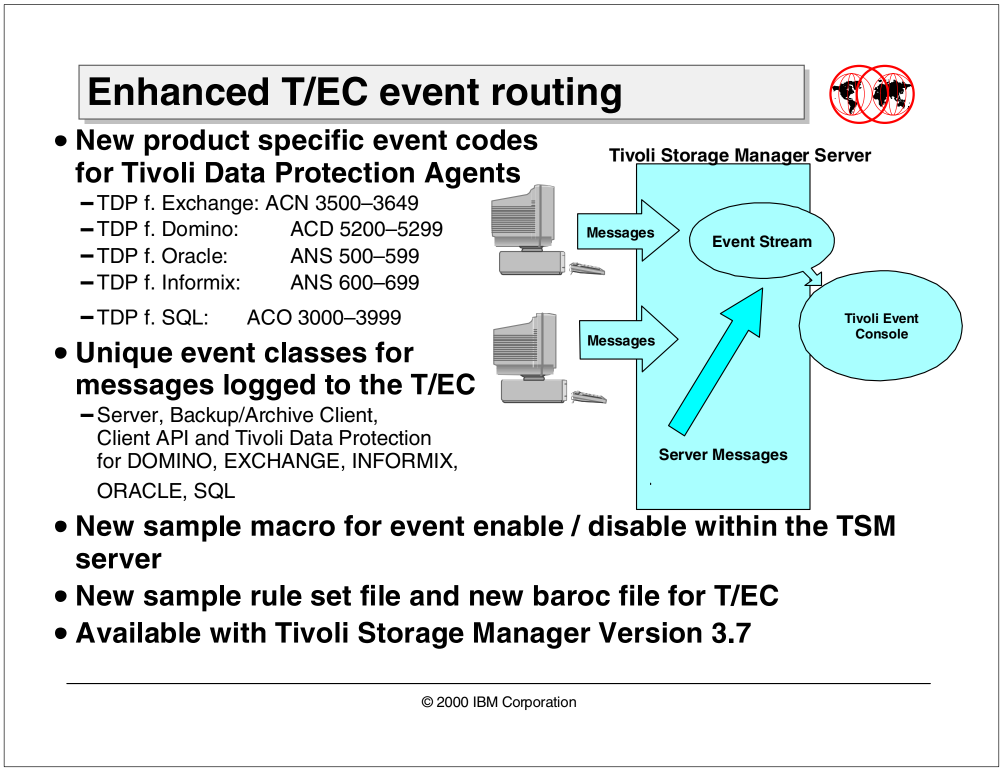

Enhanced T/EC event routing

- **New product specific event codes for Tivoli Data Protection Agents**
  - TDP f. Exchange: ACN 3500–3649
  - TDP f. Domino: ACD 5200–5299
  - TDP f. Oracle: ANS 500–599
  - TDP f. Informix: ANS 600–699
  - TDP f. SQL: ACO 3000–3999
- **Unique event classes for messages logged to the T/EC**
  - Server, Backup/Archive Client, Client API and Tivoli Data Protection for DOMINO, EXCHANGE, INFORMIX, ORACLE, SQL
- **New sample macro for event enable / disable within the TSM server**
- **New sample rule set file and new baroc file for T/EC**
- **Available with Tivoli Storage Manager Version 3.7**

T/EC *event adapters* are the interfaces through which the T/EC receives events. Event adapters are typically unique to a type of event. For example, the T/EC provides event adapters for NetView/6000, HP OpenView, logfiles, and others. These event adapters can only be used to handle events from those sources. Other applications can create their own event adapters, which was done with Tivoli Storage Manager Version 3.7

An event adapter has two alternate methods for sending events to the T/EC. A "secure" connection makes use of the object request broker technology within the Tivoli framework. This method requires that both the originating system and the T/EC are Tivoli-managed nodes. The second method is a direct TCP/IP socket connection. This is termed an “unsecure” connection by Tivoli, as the originating system does not have to be a managed node. Tivoli Storage Manager Version 3 uses this direct TCP/IP socket connection, as it provides greater flexibility, supports all Tivoli Storage Manager server platforms, and removes the requirement that the Tivoli Storage Manager server must be a Tivoli-managed node.

The Tivoli Enterprise Console (T/EC) is where messages from various programs are collected and acted upon, depending on a set of rules. This may include an automated task or paging an operator. Tivoli Storage Manager sends messages to the T/EC via the log file adapter, using Tivoli Plus Module or the Tivoli Event Receiver. The log file adapter only recognizes a few hundred of the Tivoli Storage Manager messages and could only send these to the T/EC. The Tivoli Event Receiver can send ALL Tivoli Storage Manager Server messages to the T/EC.

Only a subset of the Tivoli Storage Manager Client and TDP messages can be sent to the T/EC. The Tivoli Storage Manager Server has many thousands of messages, and not all of these need to be sent to the T/EC. A sample macro file has been provided with the Tivoli Storage Manager V3.7.3 to provide some guidance regarding which messages to enable and disable.

The Tivoli Storage Manager server can be configured to log exception messages, as events, to the T/EC. These are the type of messages a system administrator should be informed of, and these will vary with each enterprise configuration. Tivoli Storage Manager communicates with the T/EC through TCP/IP and sends pre-defined messages up to the T/EC.

4.2.1 New functions and features

Tivoli Storage Manager Version 3.7.3 enhanced TE/C support features combine:

- New product specific event codes for Tivoli Data Protection agents
- Unique event classes for messages logged to the T/EC
- A new sample rule-set file and a new baroc file for T/EC, and a new sample macro for event enable / disable events within the TSM server

4.2.1.1 New T/EC event prefix ranges for TDP applications

The Tivoli Data Protection Agents now have unique event classes for messages that can be forwarded to the T/EC.

The Prefix Ranges are as follows:

- TDP for Exchange ACN 3500–3649
- TDP for Domino ACD 5200–5299
- TDP for Oracle ANS 500–599
- TDP for Informix ANS 600–699
- TDP for SQL ACO 3000–3999

Note

The application client must have enhanced T/EC support enabled in order to route the above messages to the T/EC.

Because of the number of messages, you should not enable all messages from a node to be logged to the Tivoli Storage Manager Server.

4.2.1.2 New unique T/EC event classes

There are new event classes for the Tivoli Storage Manager server, the Backup/archive client, the Tivoli Storage Manager API Client and the Tivoli Data Protection for Domino, MS Exchange, Informix, Oracle and SQL products.

The new unique event classes provide faster filtering for the T/EC and easier rule generation. This enhancement also provides the ability to log native error codes from the Tivoli Data Protection agents to the T/EC, via the Tivoli Storage Manager Server.

Only error messages of ANE4990 through to ANE4994 (inclusive) will be passed from the Tivoli Storage Manager API to the Tivoli Storage Manager server. The Tivoli Storage Manager server will send the ANE4990 - ANE4994 messages to the T/EC for API applications that are not Tivoli Data Protection agents.

For the Backup/Archive client messages between the range of ANE4900 through to ANE4989 (inclusive), there will be a separate event class generated for each of these messages. This provides the same functionality / uniqueness as with server messages.

The Tivoli Data Protection clients now have unique event classes, within specific ranges, for messages that will be forwarded to the T/EC via the Tivoli Storage Manager server.

The following examples show the message formats:

TSM_SERVER_ANR####_platform (Server)

TSM_CLIENT_ANE#### (Backup-archive client)

TSM_APPL_ANE#### (API client)

TSM_TDP_agent_XXX#### (TDP Agents)

Where:

represents the four digit message number.
XXX represent 3-digit application type code

The platform flag will only appear for platform specific messages. The platform refers to platform specific messages (S390, AIX, AS400, HPUX, WINNT,OS2, SUN). The agent refers to the Tivoli Data Protection Agent type (DOMINO, EXCHANGE, INFORMIX, ORACLE, SQL) respectively

4.2.1.3 New sample rule-set file, baroc file, and event macro

Also provided is a new sample *rule-set file* called `ibmtsm.rls` and a new macro for recommended Enable / Disable commands called `ibmtsm.macro`. There is a new *baroc file* `ibmtsm.baroc`; this is a new developed file and must replace any older versions of the file if the T/EC is to function correctly. The baroc file defines all potential messages to the T/EC.

The new sample rule-set file supplies rules to the T/EC to inform the T/EC on how it should react if it should receive a specific message, a sequence of messages, or a certain quantity of messages from the TSM server.

The new macro file is a set of enable and disable commands. It represents a subset of the Tivoli Storage Manager messages as a recommendation for which messages should be enabled for routing to the T/EC. This will provide guidance for a Tivoli Storage Manager administrator when deciding what messages should be enabled. This file has to be run as a macro from the Tivoli Storage Manager Administrative client. The macro is shipped with the Tivoli Storage Manager V3.7.3 server code.

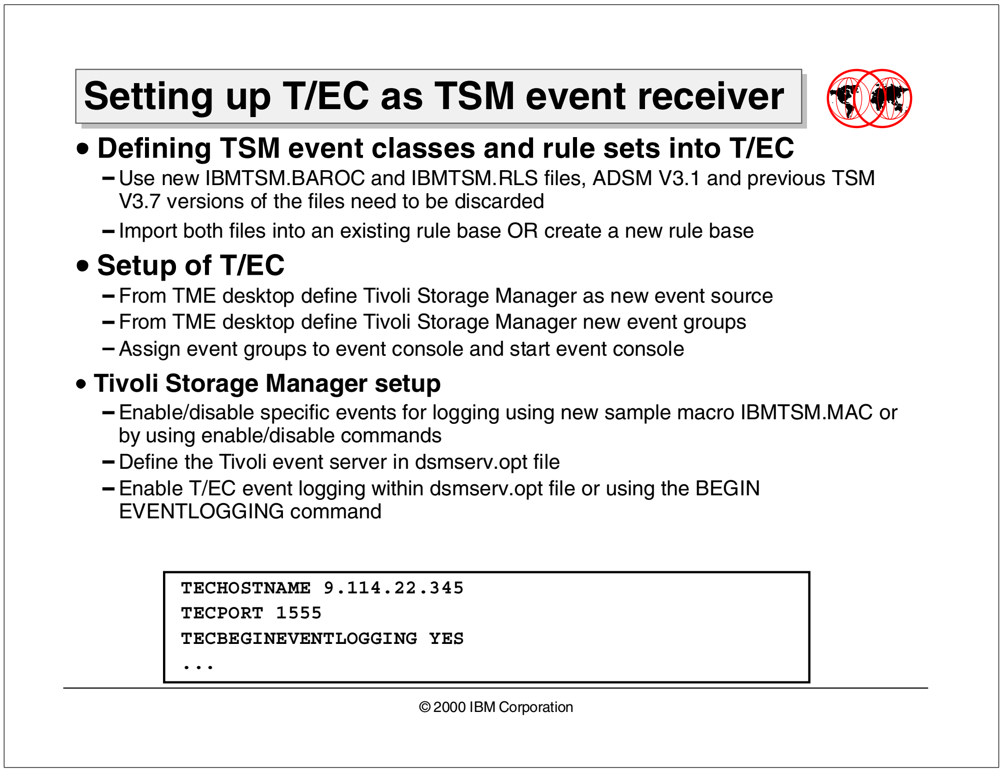

Setting up T/EC as TSM event receiver

- **Defining TSM event classes and rule sets into T/EC**
  - Use new IBMTSM.BAROC and IBMTSM.RLS files, ADSM V3.1 and previous TSM V3.7 versions of the files need to be discarded
  - Import both files into an existing rule base OR create a new rule base
- **Setup of T/EC**
  - From TME desktop define Tivoli Storage Manager as new event source
  - From TME desktop define Tivoli Storage Manager new event groups
  - Assign event groups to event console and start event console
- **Tivoli Storage Manager setup**
  - Enable/disable specific events for logging using new sample macro IBMTSM.MAC or by using enable/disable commands
  - Define the Tivoli event server in dsmserv.opt file
  - Enable T/EC event logging within dsmserv.opt file or using the BEGIN EVENTLOGGING command

```
TECHOSTNAME 9.114.22.345
TECPORT 1555
TECBEGINEVENTLOGGING YES
...
```

4.2.2 Setting up Tivoli as an event receiver

The T/EC *event server* must be customized to receive and process Tivoli Storage Manager events. Customization involves defining the Tivoli Storage Manager event classes to the T/EC. These event classes are provided in a T/EC baroc file, `IBMTSM.BAROC`, which comes with the Tivoli Storage Manager server. This is a completely new file. The `IBMADSM.BAROC` distributed with ADSM should be replaced with the new baroc file.

If you have migrated from a Tivoli Storage Manager Version 3.7.1 or ADSM Version 3.1 and have an existing *rule base*, then you must either replace the event classes by exchanging the baroc files or create a new rule base.

4.2.2.1 Using the new ibmtsm.baroc file

Before the events are displayed on a T/EC, you must import the `ibmtsm.baroc` into an existing rule base or create a new rule base and then import the `ibmtsm.baroc` file.

Adding the ibmtsm.baroc into an existing rule base:

1. From the TME desktop, click on the **RulesBase** icon to display the pop-up menu and select **Import** and specify the location of the `ibmtsm.baroc` file.
2. Select the **Compile** pop-up menu.
3. Select the **Load** pop-up menu and **Load, but activate only when server restarts** from the resulting dialog.
4. Shut down the event server and restart it.

Creating a new rule base and loading the ibmtsm.baroc file:

1. From the TME desktop, open the “Event Server Rules Bases” window, by double-clicking on the Event Server icon.
2. Select the **Create->RuleBase** menu or copy the contents of an existing rule base by selecting the **Copy** pop-up menu from the rule base to be copied.
3. Click on the RuleBase icon to display the pop-up menu.
4. Select **Import** and specify the location of the `ibmtsm.baroc` file.
5. Select the **Compile** pop-up menu.
6. Select the **Load** pop-up menu and **Load, but activate only when server restarts** from the resulting dialog.
7. Shut down the event server and restart it.

4.2.2.2 TME setup

1. From the TME desktop, select **Source** from the Event Server pop-up menu. From the resulting dialog, define a new source whose name is Tivoli Storage Manager.
2. From the TME desktop, select **Event Groups** from the EventServer pop-up menu. From the resulting dialog, define a new event group for TSM and a filter that includes the event classes IBMTSMSERVER_EVENT and IBMTSMCLIENT_EVENT.
3. From the event console icon, select the **Assign Event Group** pop-up menu item and assign the new event group to the event console.
4. Double-click on the event console icon to start the configured event console.

4.2.2.3 Tivoli Storage Manager setup

1. Enable events for logging to the Tivoli receiver by running the sample macro `ibmtsm.macro` or using enable/disable administrator commands. For example, to enable all severe and error server events, enter:

```
tsm: JAMAICA_SERVER1>enable events tivoli severe,error
```

2. In the server options file *(*`dsmserv.opt`*)*, specify the location of the host on which the Tivoli server is running. For example, to specify a Tivoli server at the IP address 9.114.22.345:1555, enter the following:

   ```
   techostname 9.114.22.345
   ```

   ```
   tecport 1555
   ```

3. Begin event logging for the Tivoli receiver. This can be done one of two ways:

   Begin event logging automatically at server startup by specifying the following server option:

   ```
   tecbegineventlogging yes
   ```

   s enter the following Tivoli Storage Manager administrator command:

```
tsm: JAMAICA_SERVER1>begin eventlogging tivoli
```

4.3 Tivoli Data Protection for Workgroups — IT Director integration

TDPfW - Tivoli IT Director integration

- IT Director interacts with TDPfW through the IT Management Console (Management Console)

- Integration includes
  - Software distribution
  - Event management
  - Task automation & scheduling
  - Remote control

- Supported for Tivoli IT Director V2 and TDPfW V1.1.1

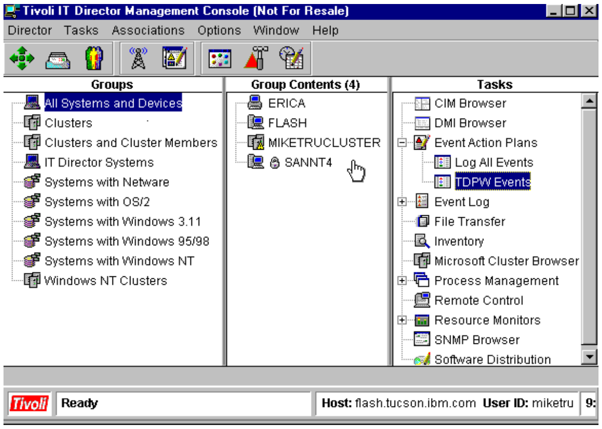

Tivoli Data Protection for Workgroups (TDPfW) provides a storage management solution NT data protection. With TDPfW you can recover individual files, full volumes, or even a complete Windows NT server in case of a disaster.

Tivoli IT Director Version 2 (IT Director) is a management product for system administrators in small to medium growth companies. IT Director is focused on Windows and Intel-centric PC and LAN networks and addresses real system management problems.

The integration with TDPfW is achieved by extending the existing tasks of IT Director.

4.3.1 Integrating TDPfW with IT Director

The IT Director administrator interacts with TDPfW through the IT Management Console (Management Console). The Management Console enables you to perform all administrative tasks. Some examples of these are:

- Process management
- Software distribution
- Event management

Areas of integration between TDPfW and IT Director are:

- Software distribution

  A TDPfW software distribution package template is incorporated into the IT Director Management Console which will enable IT Director administrators to create a TDPfW software distribution package.

- Event management

  TDPfW maintains a history file that contains messages, including the start and stop time, of jobs processed by TDPfW. These messages can also be copied to the Windows NT Application Log. IT Director's event management task supports the displaying of Windows NT events. Using these functions together will enable TDPfW to generate events which will appear in the IT DIrector event log.

- TDPfW task automation and scheduling

  TDPfW includes RepCmd, a command line interface application which processes parameter files to perform most TDPfW operations. IT Director enables an administrator to manage individual processes on remote systems. The process management task enables an administrator to start, stop, and monitor applications and processes. Using these functions together, IT Director tasks can be created to perform many TDPfW operations. Once a process manager task has been created to start a TDPfW operation, that task can be scheduled within IT Director.

- Remote control

  Remote control enables an administrator to remotely control the desktop of a native managed system. This enables the sending of keystrokes and mouse commands to the remote system and displaying the remote the system's desktop on the IT Director Management Console. An administrator can access a managed system with TDPfW installed and run TDPfW remotely.

Software distribution

- **TDPfW provides IT Director package template files**
- **Setup of TDPfW software distribution**
  - Copy tdpw1.1.1.tpl, .html and .gif files into templates directory
  - Start Software Distribution from the management console
  - Start TDPfW wizard from the template folder and follow the steps to build the package
  - TDPfW package is added to the software distribution task pane of the management console
  - Distribute software by drag & drop

Software Distribution Manager window

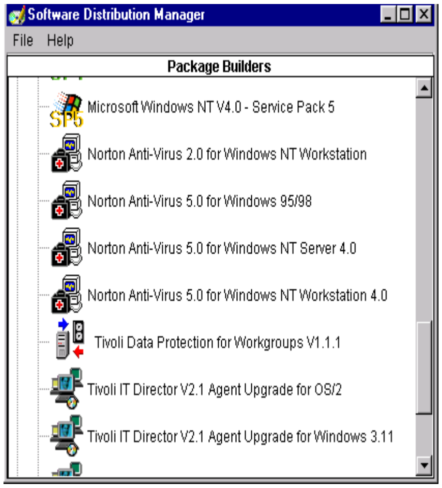

Software Distribution Manager

Package Builders

Microsoft Windows NT V4.0 - Service Pack 5
Norton Anti-Virus 2.0 for Windows NT Workstation
Norton Anti-Virus 5.0 for Windows 95/98
Norton Anti-Virus 5.0 for Windows NT Server 4.0
Norton Anti-Virus 5.0 for Windows NT Workstation 4.0
Tivoli Data Protection for Workgroups V1.1.1
Tivoli IT Director V2.1 Agent Upgrade for OS/2
Tivoli IT Director V2.1 Agent Upgrade for Windows 3.11

© 2000 IBM Corporation

4.3.2 Software distribution

IT Director Administrators can bundle, distribute, and install applications on their managed systems. IT Director supports package template files that are shortcuts to building distribution packages for common applications. Each template maps to a specific version of an application.

TDPfW provides a package template file that will act as a shortcut to building distribution packages for TDPfW. Also, an HTML file and GIF file are supplied. The HTML file contains the instructions that guide an administrator in preparing a product for package recreation. The GIF file contains the icon that is displayed by IT Director next to the name of the template.

To install the files provided with TDPfW, copy them into the `TivoliWG\Data\Templates` directory on the IT Director server.The IT Director Management Console are updated the next time the software Distribution window is opened.

Once the files have been successfully copied into IT Director, the Software Distribution Manager window of the IT Director console will contain a TDPfW icon the next time it is opened.

1. Start the IT Director Management Console.
2. Double-click on **Software Distribution** in the **Tasks** pane. The **Software Distribution Manager** dialog is displayed.
3. Double-click on the **Templates** folder.
4. Double-click on the **Tivoli Data Protection** icon. The installation wizard is started and the first panel is displayed.
5. Follow the steps to build a distribution package.
6. After the package is built, the TDPfW package is added as a child of **Software Distribution** in the **Tasks** pane of the **Management Console**.
7. Now an administrator can distribute the TDPfW package to a managed system by selecting the package and dragging it to the system.

IT Director event management

- TDPfW messages go to the IT Director via the NT event Log
- Set up TDPfW to log messages to NT event log and create IT Director action plan to forward those events to IT Director server

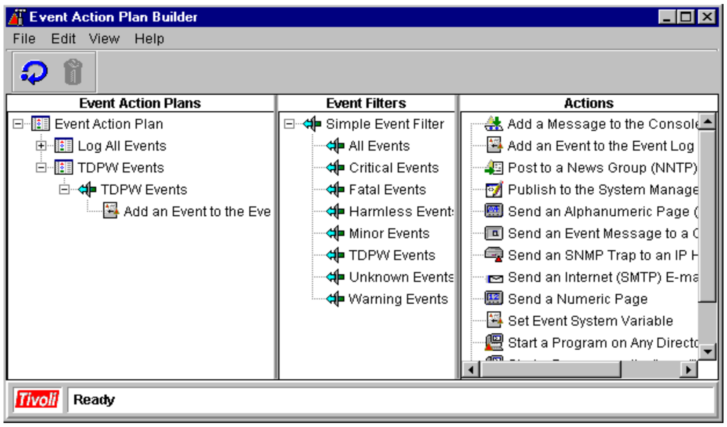
Event Action Plan Builder
File
Edit
View
Help
Event Action Plans
Event Action Plan
Log All Events
TDPW Events
TDPW Events
Add an Event to the Eve
Event Filters
Simple Event Filter
All Events
Critical Events
Fatal Events
Harmless Event:
Minor Events
TDPW Events
Unknown Events
Warning Events
Actions
Add a Message to the Console
Add an Event to the Event Log
Post to a News Group (NNTP)
Publish to the System Manage
Send an Alphanumeric Page (
Send an Event Message to a C
Send an SNMP Trap to an IP H
Send an Internet (SMTP) E-ma
Send a Numeric Page
Set Event System Variable
Start a Program on Any Directo
Tivoli
Ready

4.3.3 Event management

Through the IT Director event management task, an administrator can view a log of events that have occurred for a managed system or group of systems, and can create event action plans to associate an event with a desired action. The Windows NT event log is supported in the IT Director Log. The TDPfW messages can be sent to the Windows NT event log.

TDPfW maintains a *history file* that contains messages. An administrator can choose to have TDPfW copy the history file entries to the WIndows NT application log. When the TDPfW messages are written to the Windows NT event log, RepSrv is the source or the name of the software that logged the event.

Two steps must be performed to enable TDPfW events in the IT Director event log.

1. The TDPfW settings must be changed so that the TDPfW history file entries are also written to the Windows NT Application log.
2. At the IT Director server, an Event Action Plan must be created which will cause all TDPfW events written to the Window NT Application log to be forwarded to the IT Director server.

Setting TDPfW preferences

- Start the TDPfW administrator
- Select the File menu, select the General tab
- Check write history to application event log

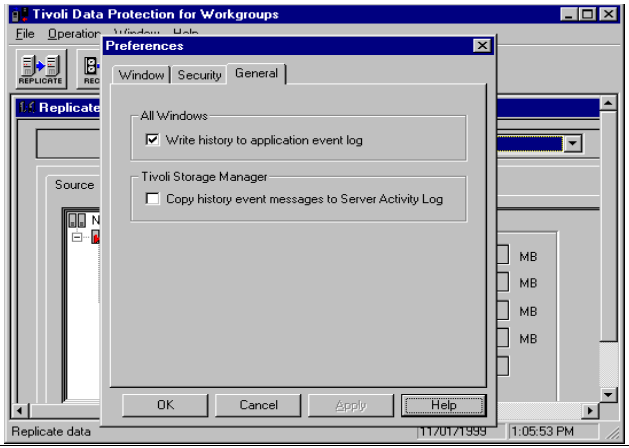

All Windows

☑ Write history to application event log

Tivoli Storage Manager

☐ Copy history event messages to Server Activity Log

Replicate data 11/01/1999 1:05:53 PM

4.3.3.1 Setting TDPfW preference

To set TDPfW preferences, do the following:

1. Start the TDPfW Administrator.
2. Select the File menu. The Preferences window is displayed.
3. Select the General tab. The general preferences are displayed.
4. Check Write history to application event log.

Creating an IT Director event action plan

- Create event action plan
- TDPfW events written to NT Event Log get forwarded to IT Director

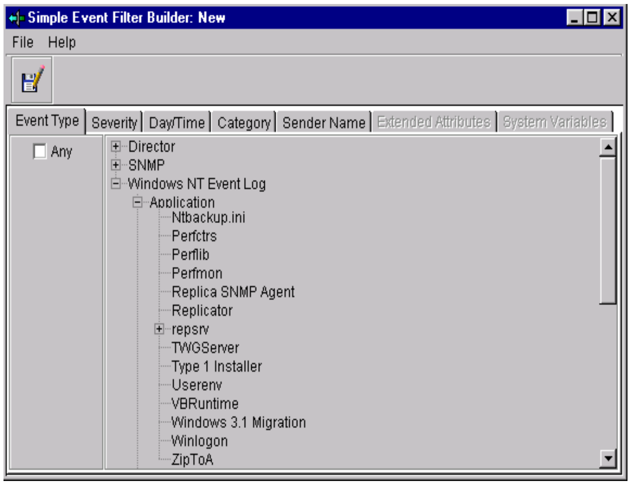

© 2000 IBM Corporation

4.3.3.2 Creating an IT Director event action plan

1. Start the IT Director Management Console.
2. Click the **Event Action Plan** icon.
3. In the **Event Filters** pane, right-click.
   - – Click New — Simple Event Filter.
   - – Create a TDPfW event filter:
     - – On the **Event Type** tab, uncheck the **Any** box.
     - – On the **Event Type** tab, select **Windows NT Event LOg.Apllication.RepSrv**, where **RepSrv** is the source of the event on the managed system.
   - – On the remaining tabs, leave the **Any** box checked.
   - – Name the filter and save it.
4. In the **Event Action Plans** pane, right-click. The menu is displayed.
   - – Click **New -- Event Action Plan**. IT Director displays the **Create Event Action Plan** window.
   - – Create TDPfW action plan.
5. In the **Event Filters** pane, select the TDPfW filter created in step 3.

6. In the **Event Action Plans** pane, select the TDPfW Action Plan created in step 4:
   - Right-click. The menu is displayed.
   - Click **Add Event Filter**
   - The filter selected in Step 5 is added to the plan selected in Step 6.
7. In the **Actions** pane, select the action **Add an Event to the Event Log**.
8. In the Event Action Plans pane, select the TDPfW Filter just added in Step 6:
   - Right-click. The menu is displayed.
   - Select **Add Action**.
   - The action selected in Step 7 is added to the filter selected in Step 8.
9. Associate this Event Action Plan with the system(s) running TDPfW.

TDPfW events from a managed system

- TDPfW messages at the IT Director event log

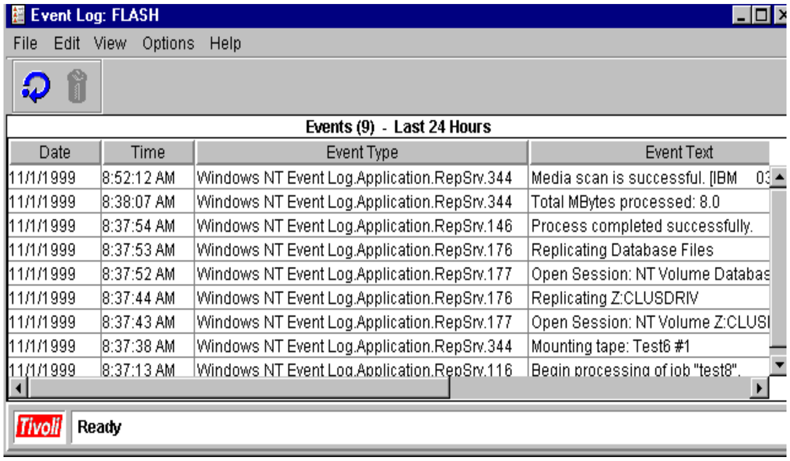

Events (9) - Last 24 Hours

Date	Time	Event Type	Event Text
11/1/1999	8:52:12 AM	Windows NT Event Log.Application.RepSrv.344	Media scan is successful. [IBM 03
11/1/1999	8:38:07 AM	Windows NT Event Log.Application.RepSrv.344	Total MBytes processed: 8.0
11/1/1999	8:37:54 AM	Windows NT Event Log.Application.RepSrv.146	Process completed successfully.
11/1/1999	8:37:53 AM	Windows NT Event Log.Application.RepSrv.176	Replicating Database Files
11/1/1999	8:37:52 AM	Windows NT Event Log.Application.RepSrv.177	Open Session: NT Volume Databas
11/1/1999	8:37:44 AM	Windows NT Event Log.Application.RepSrv.176	Replicating Z:CLUSDRIV
11/1/1999	8:37:43 AM	Windows NT Event Log.Application.RepSrv.177	Open Session: NT Volume Z:CLUSI
11/1/1999	8:37:38 AM	Windows NT Event Log.Application.RepSrv.344	Mounting tape: Test6 #1
11/1/1999	8:37:13 AM	Windows NT Event Log.Application.RepSrv.116	Begin processing of job "test8".

Ready

4.3.3.3 TDPfW messages from a managed system

The results are that the IT Director event Log now contains TDPfW events from a managed system.

The TDPfW events (messages) are now logged centrally.

TDPfW task automation & scheduling

- **TDPfW's command line interface (RepCmd)**
  - Can be used in .bat and .cmd files to start TDPfW operation
  - Command file can be started as IT Director process tasks
  - Process task can be scheduled
- **Setup of TDPfW automation**
  - Create TDPfW command file and parameter file and save both files on managed system
  - Create IT Director process task by adding the path of the command file in the command field
  - Apply the task to a managed system
  - Create a new schedule job to start the task automaticaly

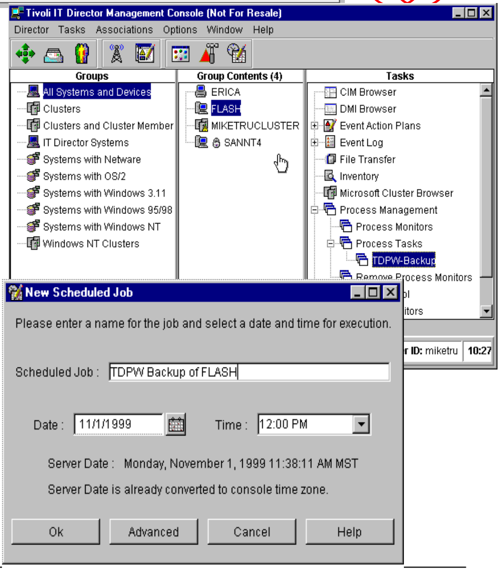

Please enter a name for the job and select a date and time for execution.

Scheduled Job : TDPW Backup of FLASH

Date : 11/1/1999

Time : 12:00 PM

Server Date : Monday, November 1, 1999 11:38:11 AM MST

Server Date is already converted to console time zone.

© 2000 IBM Corporation

4.3.4 TDPfW task automation and scheduling

TDPfW has a command line interface, RepCmd, which processes parameter files to perform virtually all TDPfW operations from the command line. The RepCmd can be integrated into automated command files such as **bat**. files.

An IT Director *process task* can have as its command a TDPfW command line file containing TDPfW commands. The IT Director can schedule these tasks.

Note

In the RepCmd command statement, do not use the parameters -el, -ep, -al, and -ap. They are not required and could cause problems when running from a service.

TDPfW jobs submitted through IT Director run as standard jobs and these can be stopped or put on hold through the standard TDPfW status window. Backups created via IT Director are entered into the TDPfW tape databases and can be used for file or volume recovery. The tapes produced by backups started by IT Director are the same as those started directly through TDPfW.

4.3.4.1 TDPfW task setup

1. Prepare the TDPfW commands.

   The type of tape device attached to TDPfW determines the content of the command files to be used. The following example shows command files for standard devices, not autochanger devices.

a. Create a *parameter file* (.RCF) that contains the parameters for the job to be scheduled. A file names `BACKSERV.RCF` could contain the lines as shown in Figure 28:

```
REPLICA_CMD=FULL_BACKUP
JOB_NAME=Through IT Director
```

Figure 28. IT Director parameter file example

The commands folder (`C:\Program Files\Tivoli\TDPfW\tdphostCommands`) contains `FullBack.RCF` and `VolBack.RCF`. These files contain the lines needed to back up an entire system, or just the C: drive.

b. Create a command file (.CMD) that contains the `RepCmd` command to run the .RCF file. For example, the `BACKSERV.CMD` file may contain:

```
C:
cd\Prgram Files\Tivoli\TDPW\tdphost
repcmd -i “c:\backups\backserv.rcf” -oa “c:\backups\backout.txt” -t
```

c. Store both files on the TDPfW system on which the job will run. For easier management, it is probably best to keep them all in one folder.

2. Create an IT Director Process Task.

a. Start the IT Director Management Console.

b. In the **Tasks** pane, expand the **Process Management** icon to access the subtasks.

c. Right-click on **Process Tasks**, then select **Open** from the menu. The **Process Task** dialog is displayed.

d. Enter the filepath and command syntax in the **Command** field.

e. Select **File ===> Save As**. The **Save As** dialog is displayed.

f. Enter the name to assign to this process, then select **OK**.

g. The TDPfW task can now be applied to a managed system to execute the command.

4.3.4.2 TDPfW task sheduling

The task scheduler feature of IT Director enables an administrator to schedule any task. Once a task has been defined to IT Director, an administrator can schedule these tasks. To schedule a TDPfW task from the IT Director Management Console window, drag-and-drop the task to be executed onto a managed system.

When a TDPfW task is selected to be performed, it can either be performed immediately, or it can be scheduled for a later date. If this is selected, the **New Scheduled Job** dialog will prompt the administrator for basic scheduling information.

Note

When installing the IT Director client on a remote system, remote control needs to be enabled. Logging in to a remote system through remote control requires that the **Require user authorization for screen access** setting is disabled on the remote screen.

4.4 Tivoli Service Desk integration

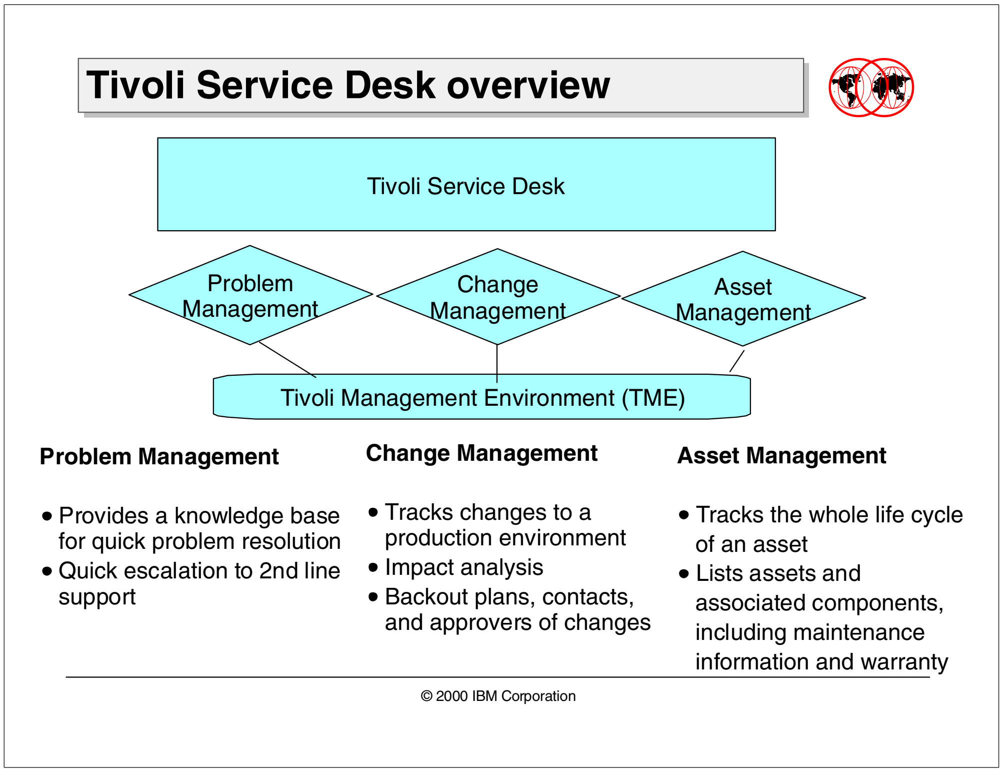

Tivoli Service Desk overview

Problem Management

- Provides a knowledge base for quick problem resolution
- Quick escalation to 2nd line support

Change Management

- Tracks changes to a production environment
- Impact analysis
- Backout plans, contacts, and approvers of changes

Asset Management

- Tracks the whole life cycle of an asset
- Lists assets and associated components, including maintenance information and warranty

There are a number of applications that comprise the Tivoli Service Desk (TSD). These applications include Tivoli Problem Management, Tivoli Change Management, Tivoli Asset Management, and Tivoli Decision Support. This chapter will briefly explain each application and show how Tivoli Storage Manager integrates with TSD.

4.4.1 Tivoli Service Desk overview

TSD is a suite of Problem Management, Change Management, and Asset Management applications. These applications are designed to work together so that you can manage the complete life cycle of problem tracking, change requests, and asset tracking activities in the enterprise information technology (IT) infrastructure.

4.4.1.1 Problem Management module

The Tivoli Service Desk V 6.0 Problem Management (TPM) module provides a multi-support organizational structure. This structure effectively applies Problem Management solutions for IT-related problems for more than one organization. Tivoli Service Desk allows first-level support personnel to easily look up similar problems and find solutions through a search of the knowledge base. If a request is not resolved while in contact with the customer or end user, the problem can be assigned to second level support personnel.

4.4.1.2 Change Management module

The Tivoli Service Desk V 6.0 Change Management (TCM) module is used to track all changes to the production environment that might affect business performance or operations. Impact analysis and potential problems with a change can be detected before the change is physically executed. Fallback procedures, contacts, and approvers are defined for each individual change request. A change cannot be executed without all the proper approvals in place. Changes can be tracked, monitored, and measured to ensure minimal disruption to the production environment. The Tivoli Service Desk Change Management module supports the business processes and infrastructure of an enterprise.

4.4.1.3 Asset Management module

The Tivoli Service Desk V 6.0 Asset Management (TAM) module is designed to track the whole life cycle of an asset. It provides a detailed list of assets and associated components, including maintenance information, cost centers, and warranty information.

4.4.1.4 TSD integration with the Tivoli environment

Tivoli Service Desk integration with the Tivoli Management Environment (TME):

- Facilitates effective management and execution of Help Desk tasks by allowing access to the distributed TME Inventory, Remote Control, and User Administration applications from the Tivoli Service Desk desktop.
- Allows TME applications to monitor Tivoli Service Desk as a main business application through monitoring of Tivoli Service Desk internal application tasks.

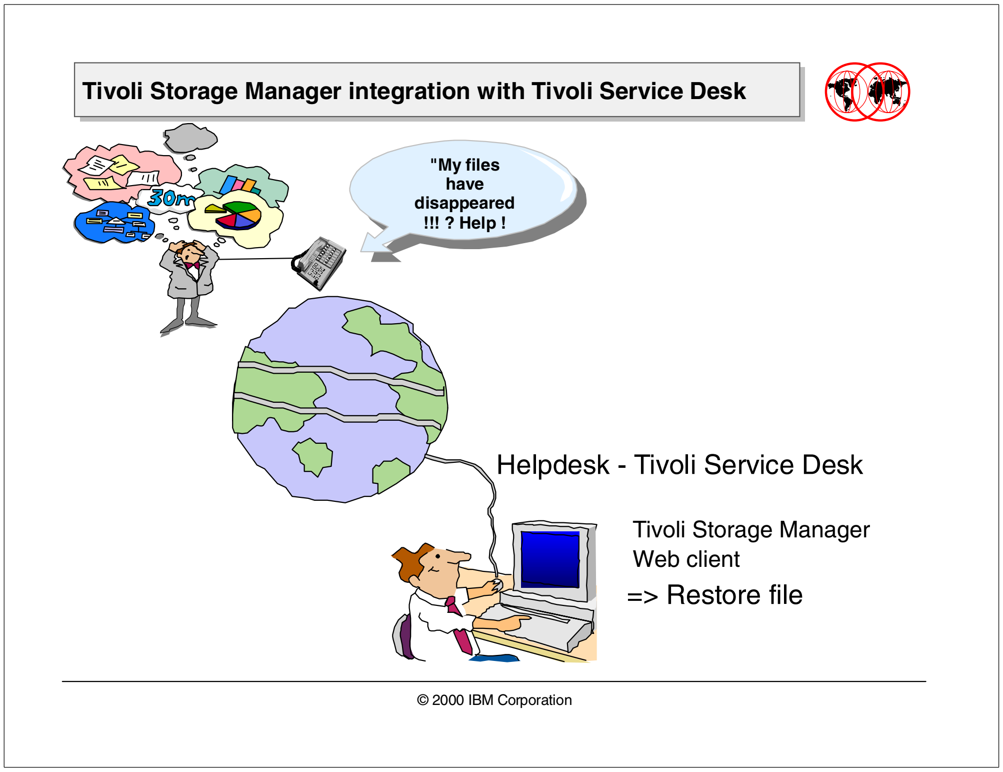

Tivoli Storage Manager integration with Tivoli Service Desk

Helpdesk - Tivoli Service Desk

Tivoli Storage Manager
Web client

=> Restore file

© 2000 IBM Corporation

4.4.2 Tivoli Storage Manager integration with Tivoli Service Desk

Tivoli Storage Manager integration was included with the Tivoli Service Desk V6.0 release.

The integration works with the Tivoli Storage Manager Web Client Interface. The TSD Integrations Action Group has been amended to include a new menu item for the Tivoli Storage Manager Web client. When this menu is selected in TSD, it brings up a prompt for a machine name or IP address where the administrator needs to perform an action. The TSD administrator can perform any TSM Web Client action as necessary; for example, a file restore.

Chapter 5. Windows 2000 and TSM server exploitation

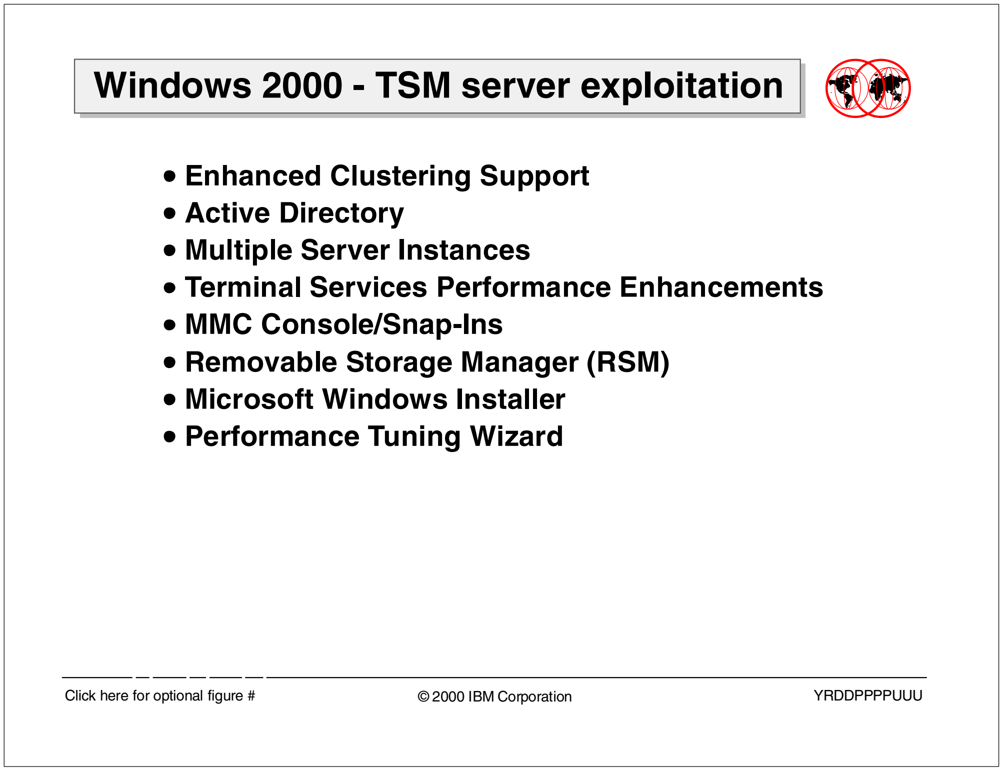

Windows 2000 - TSM server exploitation

- Enhanced Clustering Support
- Active Directory
- Multiple Server Instances
- Terminal Services Performance Enhancements
- MMC Console/Snap-Ins
- Removable Storage Manager (RSM)
- Microsoft Windows Installer
- Performance Tuning Wizard

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

Microsoft Windows 2000, the latest major release of the Windows NT operating system, introduces a number of architectural enhancements that require special attention from system administrators. This chapter discusses some of the new features within Windows 2000. It gives an overview of how Tivoli Storage Manager (TSM) server integrates with and leverages these features to provide a solution of choice.

The topics include:

- Enhanced clustering support
- Active directory
- Multiple server instances
- Terminal services performance enhancements
- MMC console/snap-ins
- Removable Storage Manager
- Microsoft Windows Installer
- Performance Tuning Wizard

5.1 Microsoft Cluster server

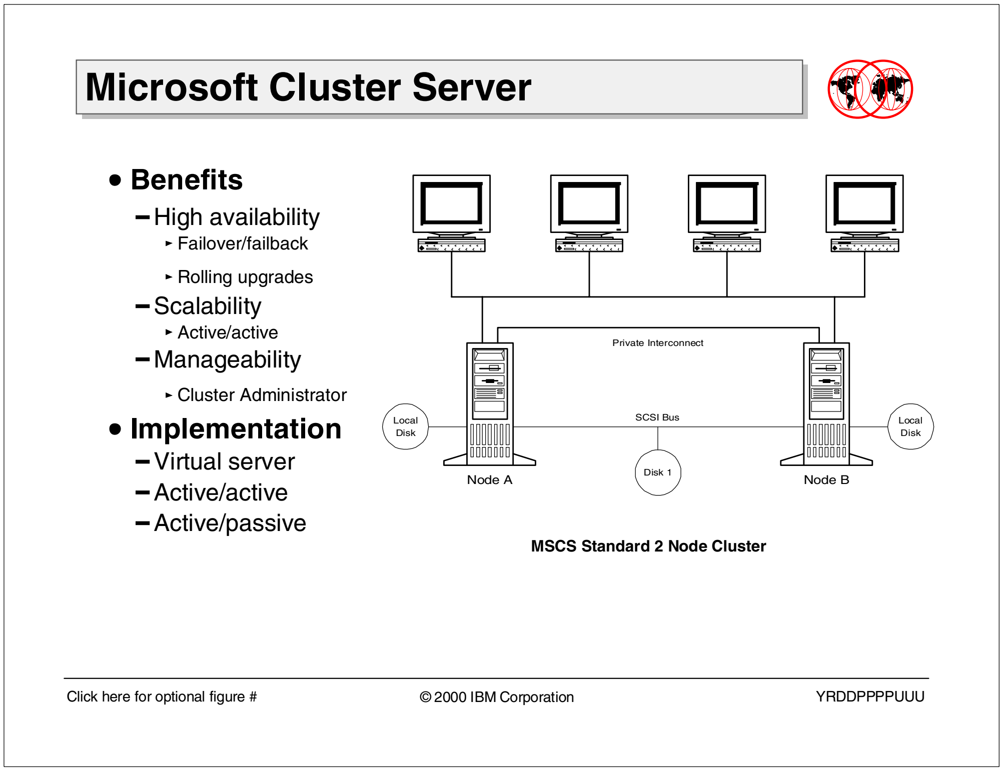

Microsoft Cluster Server

- **Benefits**
  - High availability
    - Failover/failback
    - Rolling upgrades
  - Scalability
    - Active/active
  - Manageability
    - Cluster Administrator
- **Implementation**
  - Virtual server
  - Active/active
  - Active/passive

MSCS Standard 2 Node Cluster

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

Microsoft Cluster Server (MSCS) provides high availability, scalability, and manageability for resources and applications by clustering multiple servers running Windows 2000 Advanced Server or Windows 2000 Datacenter Server. It preserves client access to applications and resources during failures and planned outages. If a cluster server, resource, or application is unavailable due to failure or maintenance, resources and applications will migrate to an available server within the cluster.

MSCS achieves high availability by:

- **Failover/Failback**. In the event of a failure, *failover* occurs. Ownership of resources, such as disk drives and IP addresses, is automatically transferred from a failed server to a surviving server. The part of the workload that is capable of restarting on the surviving server is then restarted. For example, a print queue may be restarted, but an SQL query that was executing when the failure occurred may not restart. *Failback* is the ability to automatically rebalance the workload in a cluster when a failed server comes back online.
- **Rolling Upgrades**. Administrators can move all the workload onto one server and maintain or upgrade the unloaded server. Once the maintenance or upgrade is completed and tested, the server is brought back online, and it automatically rejoins the cluster. This process can be repeated when maintaining or upgrading the other server in the cluster.

MSCS achieves scalability by:

- **Active/Active Clustering**. All servers run their own workload. This means every computer in the cluster is available to do real work, and each computer in the cluster is also available to recover the resources and workload of any other computer in the cluster. There is no need to have a wasted, idle server waiting for a failure.

MSCS achieves manageability by:

- **Cluster management applications**. You use cluster management applications, such as Cluster Administrator, to configure, control, and monitor clusters. You can install Cluster Administrator on any computer running Windows 2000. You can also administer a Windows 2000 server cluster remotely from a computer running Windows NT 4.0 Service Pack 3 or later using the Windows NT Server, Enterprise Edition 4.0, Cluster Administrator tool. Cluster Administrator allows you to manage cluster objects, establish groups, initiate failover, handle maintenance, and monitor cluster activity through a graphical console.

5.1.1 Implementation

A server cluster consists of two or more independent Microsoft Windows NT Server/Windows 2000 Server systems. The computer systems that make up the collection are known as *cluster nodes*. Nodes have the following characteristics:

- Every node has access to all cluster configuration data.
- Every node communicates with the other nodes in the cluster through one or more physically independent networks (sometimes referred to as interconnects). Network adapters, referred to in server clusters as network interfaces, attach nodes to networks.
- Every node in the cluster is aware when another system joins or leaves the cluster.
- Every node in the cluster is aware of the resources that are running locally as well as the resources that are running on the other cluster nodes.
- All nodes in the cluster are grouped under a common name, the cluster name, which is used for accessing and managing the cluster.

When a node starts, it searches for Active nodes on the networks designated for internal communication. If it finds an Active node, it attempts to join the node's cluster. If it cannot find an existing cluster, it attempts to form a cluster by taking control of the *quorum resource*. The quorum resource stores the most current version of the *Cluster Database*, which contains cluster configuration and state data. A server cluster maintains a consistent, updated copy of the Cluster Database on all Active nodes.

A node can host physical or logical units, referred to as *resources*. Administrators organize these cluster resources into functional units called *groups* and assign these groups to individual nodes. If a node fails, the server cluster transfers the groups that were being hosted by the node to other nodes in the cluster. This transfer process is called failover. The reverse process, failback, occurs when the failed node becomes active again and the groups that were failed over to other nodes are transferred back to the original node.

Nodes, resources, and groups are three kinds of *cluster objects*. The others are networks, network interfaces, and resource types. All server cluster objects are associated with a set of properties, which are data values that describe an object's identity and behavior in the cluster. Administrators manage cluster objects by manipulating their properties, typically through a cluster management application such as Cluster Administrator.

Developers use the Server Cluster API to create cluster-aware applications, cluster management applications, and custom resource types.

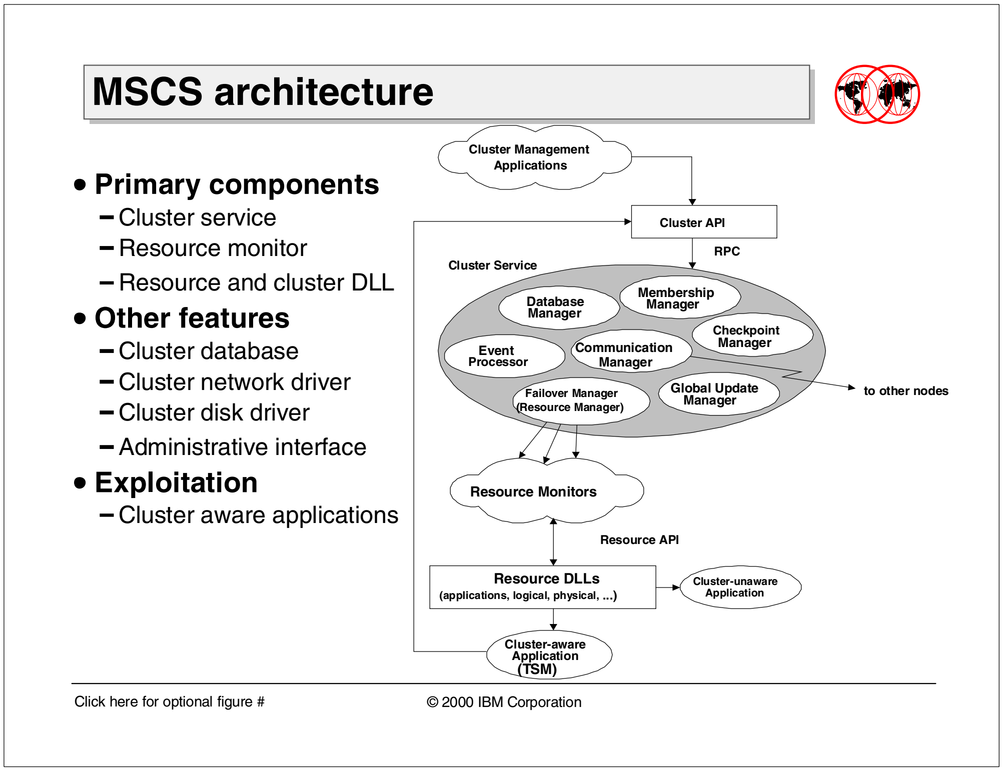

MSCS architecture

- **Primary components**
  - Cluster service
  - Resource monitor
  - Resource and cluster DLL
- **Other features**
  - Cluster database
  - Cluster network driver
  - Cluster disk driver
  - Administrative interface
- **Exploitation**
  - Cluster aware applications

Click here for optional figure #

© 2000 IBM Corporation

5.1.2 MSCS architecture

Microsoft Cluster Server is comprised of three key components: Cluster Service, Resource Monitors, and Resource DLLs. Other software components that are implemented in the server include the Cluster Database, Cluster Network, Disk Drivers, and Administrative Interfaces.

5.1.2.1 The Cluster Service

The *Cluster Service* (which is composed of the Event Processor, the Failover Manager/Resource Manager, the Global Update Manager, and so forth) is the core component of MSCS and runs as a high-priority system service. The Cluster Service controls cluster activities and performs such tasks as coordinating event notification, facilitating communication between cluster components, handling failover operations, and managing the configuration. Each cluster node runs its own Cluster Service. The Cluster Service is fundamental to a node's participation in a server cluster. A single instance of the Cluster Service, however, is specific to a single node; and therefore cluster-wide operations depend on all instances of the Cluster Service communicating and working together.

5.1.2.2 The Resource Monitor

A *Resource Monitor* provides a communication, monitoring, and processing layer between the Cluster Service and one or more resources. A Resource Monitor always runs in a process separate from the Cluster Service. If a resource fails, the Resource Monitor isolates the Cluster Service from the effects. If the Cluster Service fails, the Resource Monitor allows its resources to shut down gracefully.

The Cluster Service uses the Resource Monitor to communicate with the resource DLLs. The DLL handles all communication with the resource, thus shielding the Cluster Service from resources that misbehave or stop functioning. To work with a resource, a Resource Monitor loads the resource DLL responsible for that resource type into its process. When the Cluster Service requests an operation on a resource, the Resource Monitor routes the request to the appropriate entry point function of the resource DLL responsible for the resource.

The Resource Monitor performs default processing for some resource operations; it stores synchronized state data, allowing the Cluster Service and resource DLLs to operate asynchronously, checking and updating resource status as needed. A Resource Monitor periodically checks the operational status of all of its resources. Multiple copies of a Resource Monitor can be running on a single node, thereby providing a means by which unpredictable resources can be isolated from other resources.

5.1.2.3 The Resource DLL

Resource DLLs are responsible for carrying out most operations on cluster resources. A Resource DLL contains the resource-specific code necessary to provide high availability for instances of one or more resource types. It exposes this code through a standard interface consisting of a set of entry point functions. It is registered with the Cluster Service to associate one or more resource type names with the name of the DLL. It is always loaded into a Resource Monitor's process when in use.

When the Cluster Service needs to perform an operation on a resource, it sends the request to the Resource Monitor assigned to the resource. If the Resource Monitor does not have a DLL in its process that can handle that type of resource, it uses the registration information to load the DLL associated with the resource type. It then passes the Cluster Service's request to one of the DLL's entry point functions. The Resource DLL handles the details of the operation so as to meet the specific needs of the resource.

This architecture separates the Cluster Service from the implementation details of specific resource types, allowing the Cluster Service to manage file shares, physical disks, IP addresses, and all other types of resources through a standard interface. It also allows for the creation of new resource types. You can define your own resource types to provide customized support for cluster-unaware applications, enhanced support for cluster-aware applications, or specialized support for new kinds of devices.

5.1.2.4 Cluster Database

The *Cluster Database* is a set of keys, sometimes referred to as the “cluster hive”, under HKEY_LOCAL_MACHINE in the Microsoft Windows NT/Windows 2000 registry. It contains information about all physical and logical elements in a cluster, including a listing of cluster objects, their properties, and configuration data. Each node stores a continuously updated copy of the Cluster Database. The Cluster Services on each node cooperate to maintain a consistent, updated image by participating in global updates and periodic consistency checks. The quorum resource also maintains an updated copy of the Cluster Database.

5.1.2.5 Cluster Network Driver

Each node in the cluster runs an instance of the *Cluster Network Driver*, Clusnet.sys. The Cluster Network Driver is responsible for monitoring the status of all network paths between nodes. It does message routing and detects communication failure. Each node's Cluster Network Driver exchanges periodic messages, called heartbeats, with the Cluster Network Driver on the other Active nodes. If a node fails to respond to a heartbeat message, the Cluster Network Driver on the node that detects the failure notifies the Cluster Service, which initiates failover. Note that failure to respond to a heartbeat message can be caused by node failure, network interface failure, or network failure.

5.1.2.6 Cluster Disk Driver

Each node in the cluster runs an instance of the *Cluster Disk Driver*, `Clusdisk.sys`. The Cluster Disk Driver is responsible for ensuring that only one node at a time can access a cluster disk.

5.1.2.7 Administration Interfaces

Server clusters provide two cluster management applications for administering cluster operations: *Cluster Administrator* and `Cluster.exe`.

Any Microsoft Windows NT/Windows 2000 system, regardless of whether it is a cluster node, can install Cluster Administrator. Cluster Administrator allows administrators to manage cluster objects, establish groups, initiate failover, handle maintenance, and monitor cluster activity through a convenient graphical interface. Third-party developers providing their own resource types can write an extension that enables Cluster Administrator to manage their custom types.

`Cluster.exe` is a command-line interface for administering server clusters.

5.1.2.8 Cluster-aware applications

Applications that implement their own Resource DLLs to communicate with the Cluster Service and that use the Cluster API to request and update cluster information are defined as *cluster-aware applications*. Applications and services that do not use the Cluster or Resource APIs and cluster control code functions are unaware of clustering and have no knowledge that MSCS is running. These cluster-unaware applications are generally managed as generic applications or services.Both cluster-aware and cluster-unaware applications run on a cluster node and can be managed as cluster resources. However, only cluster-aware applications can take advantage of features offered by Cluster Server through the Cluster API.

MSCS includes two tools that perform basic cluster management: these tools are Cluster Administrator (`CluAdmin.exe`) and a command line management tool (`Cluster.exe`). Cluster-aware applications also implement Cluster Administrator extension DLLs, which contain implementations of interfaces from the Cluster Administrator extension API. A Cluster Administrator extension DLL allows an application to be configured into the Cluster Administrator tool (`CluAdmin.exe`). Implementing custom resource and Cluster Administrator extension DLLs allows for specialized management of the application and its related resources, and enables the system administrator to install and configure the application more easily.

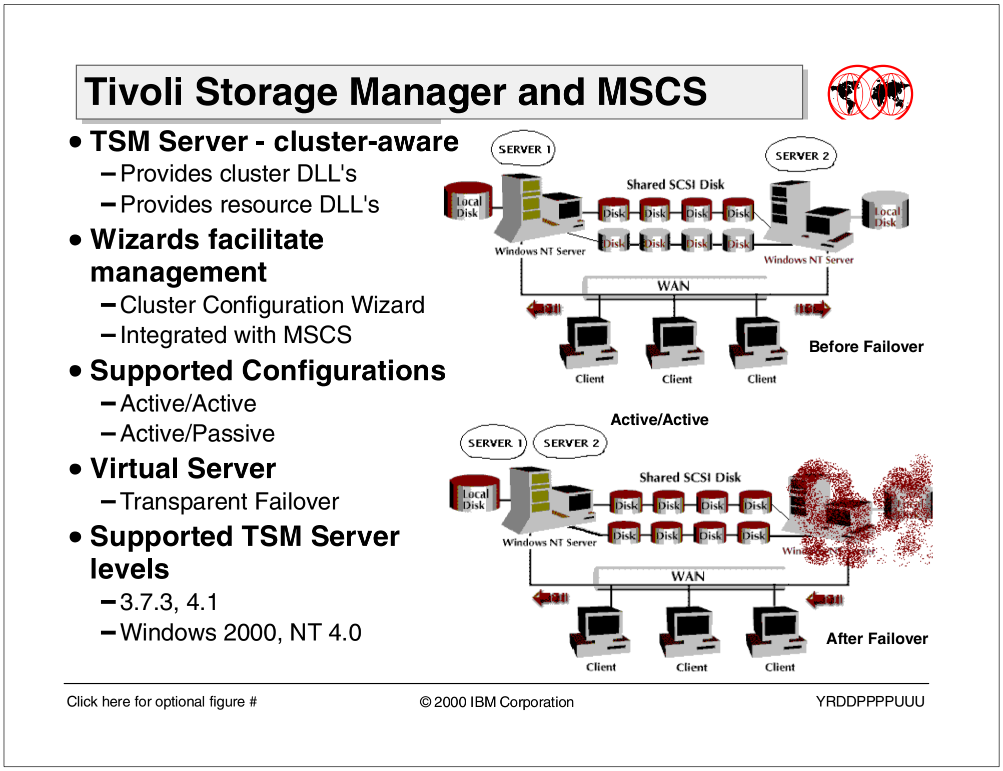

Tivoli Storage Manager and MSCS

- **TSM Server - cluster-aware**
  - Provides cluster DLL's
  - Provides resource DLL's
- **Wizards facilitate management**
  - Cluster Configuration Wizard
  - Integrated with MSCS
- **Supported Configurations**
  - Active/Active
  - Active/Passive
- **Virtual Server**
  - Transparent Failover
- **Supported TSM Server levels**
  - 3.7.3, 4.1
  - Windows 2000, NT 4.0

Click here for optional figure #

© 2000 IBM Corporation

YRDDPPPPUUU

5.1.3 TSM exploitation of MSCS

TSM is a cluster-aware application that fully integrates with and takes advantage of MSCS clustering and administrative capabilities.

5.1.3.1 Failover support

TSM uses MSCS failover to offer two configurations:

- **Active/Passive configuration** — For the Active/Passive configuration, you create one instance of a TSM server that can run on either node. The other node serves as an online (hot) backup. In active/passive configurations, the server has one database, recovery log, and one set of storage pool volumes. These volumes are created on a shared disk accessible by both nodes. TSM code is installed on the local disk of each node.
- **Active/Active configuration** — For the Active/Active configuration, the cluster runs two independent instances of a TSM server. Although the instances typically run on separate nodes, one node can run both instances. (All the server utilities and wizards have been updated to support multiple server instances on the same node). Even if both instances are running on the same physical server, users believe they are accessing separate servers.

  In Active/Active configurations, each virtual server has a separate database, recovery log, and a separate set of storage pool volumes because virtual servers cannot share data, nor can they share the database or recovery logs. MSCS requires each TSM server instance to have a private set of disk resources. Although nodes can share disk resources, only one node can actively control a disk at a time. TSM code is installed on each server's local disk.

5.1.3.2 Cluster configuration wizard

Both the Tivoli Storage Manager client and server are cluster aware and take full advantage of Windows 2000 clustering capabilities. TSM enhances integration and manageability by providing Resource and Cluster Administrator DLLs.

The installation of Tivoli Storage Manager on a Cluster Server has been simplified with the help of wizards that have been added to the product. The Server Utilities now includes a *cluster configuration wizard* and is now fully integrated with the Microsoft Cluster Server Administrator. This wizard can be run on different cluster nodes to help ensure that TSM is properly configured to take advantage of high availability environments.

5.1.3.3 Virtual servers

MSCS lets you place TSM server cluster resources into a virtual server. A *virtual server* is an MSCS cluster group that looks like a Windows server. The virtual server has a network name, an IP address, one or more physical disks, and a service. A TSM server can be one of the virtual services provided by an MSCS virtual server.

The virtual server name is independent of the name of the physical node on which the virtual server runs. The virtual server name and address migrate from node to node with the virtual server. Clients connect to a TSM Server using the virtual server name, rather than the Windows server name. The virtual server name is implemented as a cluster network name resource and maps to a primary or backup node, depending on where the virtual server currently resides. Any client that uses WINS or directory services to locate servers can automatically track the virtual server as it moves between nodes. Automatically tracking the virtual server does not require client modification or reconfiguration.

As mentioned above, each virtual server has its own disk as part of a cluster resource group. Therefore, they cannot share data. Each TSM server that has been implemented as a virtual server has its database, recovery log, and set of storage pool volumes on a separate disk owned by that virtual server.

Since the server’s location is transparent to client applications, this affords TSM the maximum ease of failover and failback, while minimizing the impact on the TSM clients.

Note

Currently, MSCS only supports an "IP Address" as a resource. This means that any TSM Server running on a cluster must limit its supported communication methods to just TCP/IP. Any client NOT using TCP/IP as a communication method will not be able to reach the virtual server if it should failover to the other cluster node

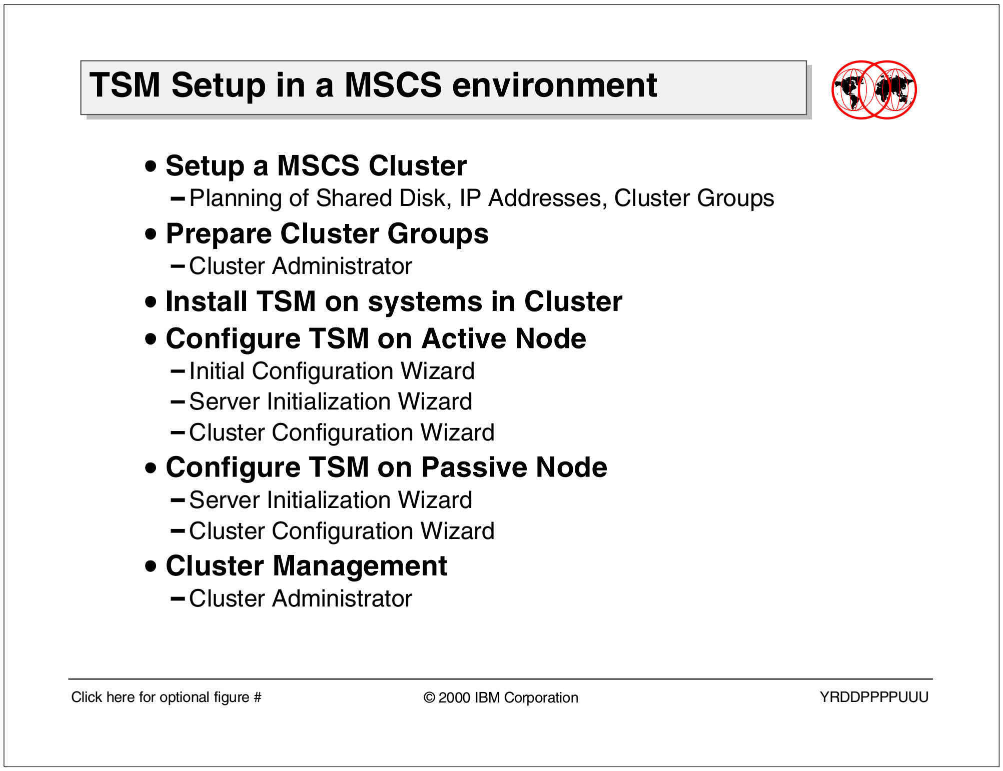

TSM Setup in a MSCS environment

- **Setup a MSCS Cluster**
  - Planning of Shared Disk, IP Addresses, Cluster Groups
- **Prepare Cluster Groups**
  - Cluster Administrator
- **Install TSM on systems in Cluster**
- **Configure TSM on Active Node**
  - Initial Configuration Wizard
  - Server Initialization Wizard
  - Cluster Configuration Wizard
- **Configure TSM on Passive Node**
  - Server Initialization Wizard
  - Cluster Configuration Wizard
- **Cluster Management**
  - Cluster Administrator

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

5.1.4 TSM setup in an Active/Passive MSCS environment

To enable TSM for high availability, perform the following steps:

Step 1 — Set up an MSCS cluster

Setting up an MSCS cluster requires a lot of planning up-front. The design considerations include the type of configuration to use, disk resources to be dedicated to the application, and IP addresses for primary and failover adapters as well as cluster resource groups. There are planning worksheets that can be utilized. See your MSCS documentation and the Windows and Cluster Administrator online help for the details of setting up and configuring an MSCS cluster. The basic steps include these:

1. Install the Windows operating system on all machines in the cluster.
2. Install and configure the Microsoft Cluster Server.
3. Verify that each node and the shared disk in the cluster is operational.

Step 2 — Prepare cluster groups

Prepare one or more cluster groups. Each TSM server instance requires a cluster resource group. The group must contain only disk resources. You can create a new group and move disk resources to it. You can choose to rename an existing resource group that contains only disk resources. Use the Cluster Administrator program on the computer that owns the shared disk resource. This will be the Active node. From the Cluster Administrator, create the group.

1. Ensure that the groups have been given distinctive names, for example, TSM Srv1 and TSM Srv 2. **Do not change the group names after this step.** A name change could cause a corrupt configuration.

2. Ensure that all nodes in the cluster are online. The minimum requirement is two nodes.
3. Ensure that the group is online and owned by the node where you will be installing the initial server instance. Move groups as needed. All disk resources in the group must also be online. The system restart will cause ownership of groups to change, you must move the group back to the node where TSM will be configured.
4. Test the cluster to ensure proper operation of each node and the shared disk.

Step 3 — Install TSM on machines in the cluster

On every machine in the cluster, install TSM.

1. Install TSM to a local disk. It is recommended to use the same local-disk drive letter for each machine in the cluster.
2. Restart the system after TSM installation completes. After the system comes up, ensure that the TSM group and disk resources are again online and owned by the node at which TSM is being configured.

Step 4 — Configure TSM on the Active node

We are considering an Active/Passive configuration. You must configure one TSM server on the initial, or active, node; and one TSM server on the Passive node. You must perform configuration on the Active node first. If you try to perform configuration on the Passive node first, it fails.

From the TSM Server Utilities, do the following:

1. Start the Initial Configuration Wizard, run through the steps to set up the TSM environment. During the Server Initialization, the Windows Cluster is detected. Click Yes to configure TSM in a Windows Cluster and click Next as shown in Figure 29. The Select Cluster group dialog displays. Select the group defined in Step 2 (TSM Srv 1) and continue.

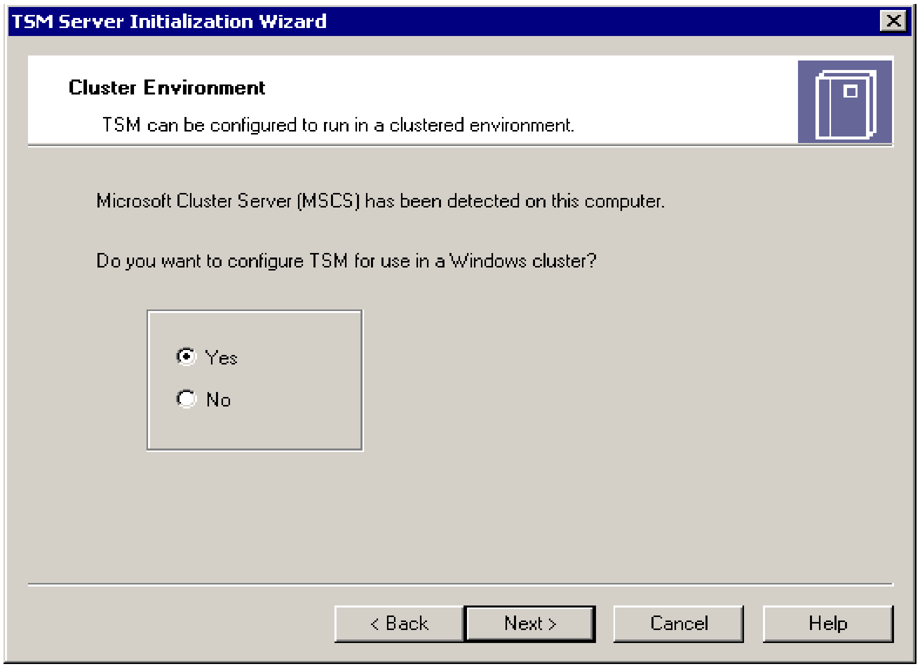

TSM Server Initialization Wizard

Cluster Environment

TSM can be configured to run in a clustered environment.

Microsoft Cluster Server (MSCS) has been detected on this computer.

Do you want to configure TSM for use in a Windows cluster?

● Yes

○ No

Figure 29. TSM Server Initialization Wizard

2. After the Server Initialization Wizard completes, the TSM Cluster Configuration Wizard begins. The Cluster Configuration Wizard displays only if Windows detects that a server cluster environment exists. The wizard asks you to select a group. Select the same group (as shown in Figure 30) you picked in step 1, which is the default.

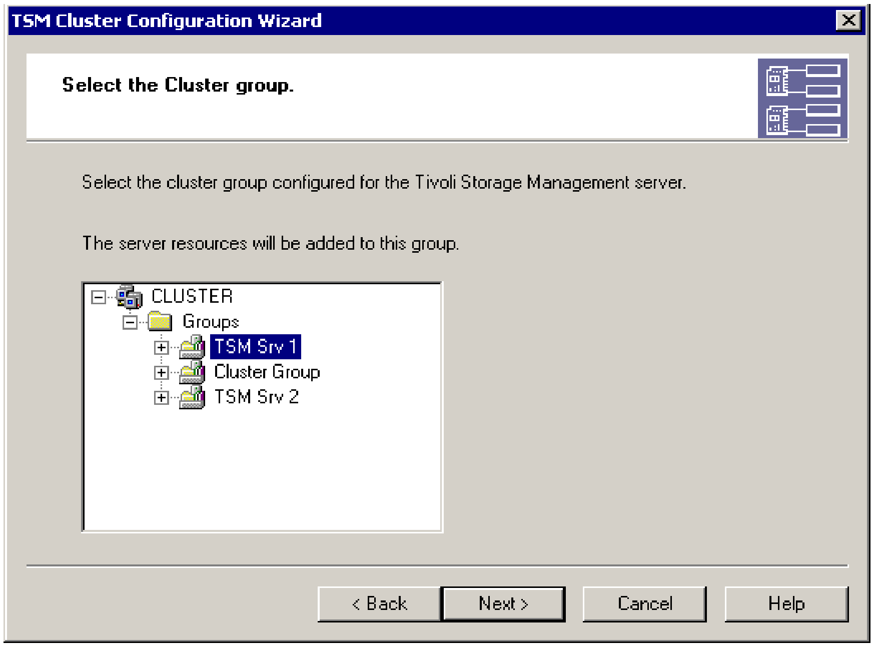

TSM Cluster Configuration Wizard

Select the Cluster group.

Select the cluster group configured for the Tivoli Storage Management server.

The server resources will be added to this group.

Figure 30. TSM Cluster Configuration Wizard — Select Cluster Group

3. Next you will be asked to specify an IP address, subnet mask, and the network to be used. Enter the information from your planning worksheet as shown in Figure 31.

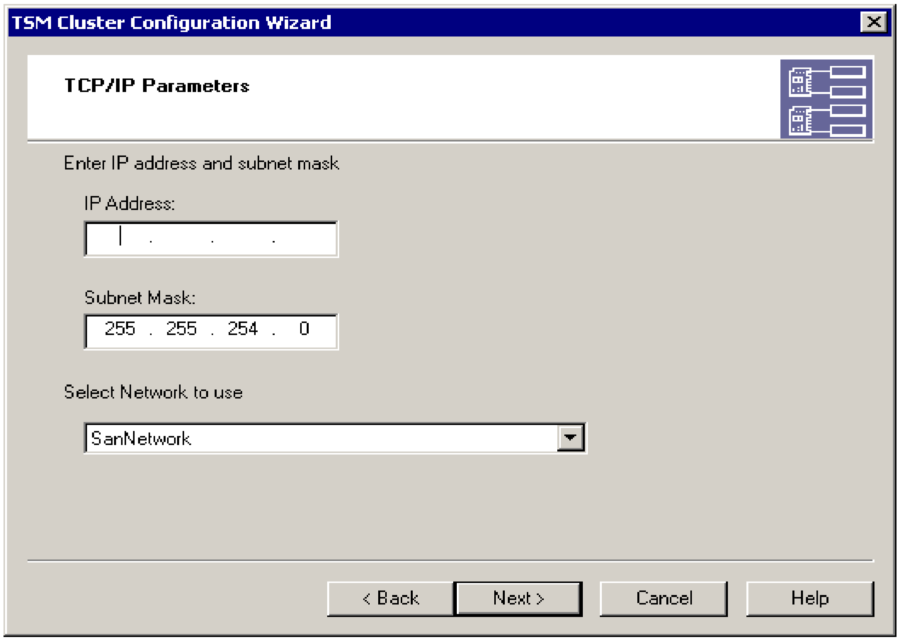

TSM Cluster Configuration Wizard

TCP/IP Parameters

Enter IP address and subnet mask

IP Address: . . .

Subnet Mask: 255 . 255 . 254 . 0

Select Network to use

SanNetwork

Figure 31. TSM Cluster Configuration Wizard — TCP/IP Parameters

4. Next the Cluster Configuration Wizard prompts you to enter the cluster network name. This is the name of the virtual server (as discussed in Section 5.1.3.3, “Virtual servers” on page 143 and see Figure 32).

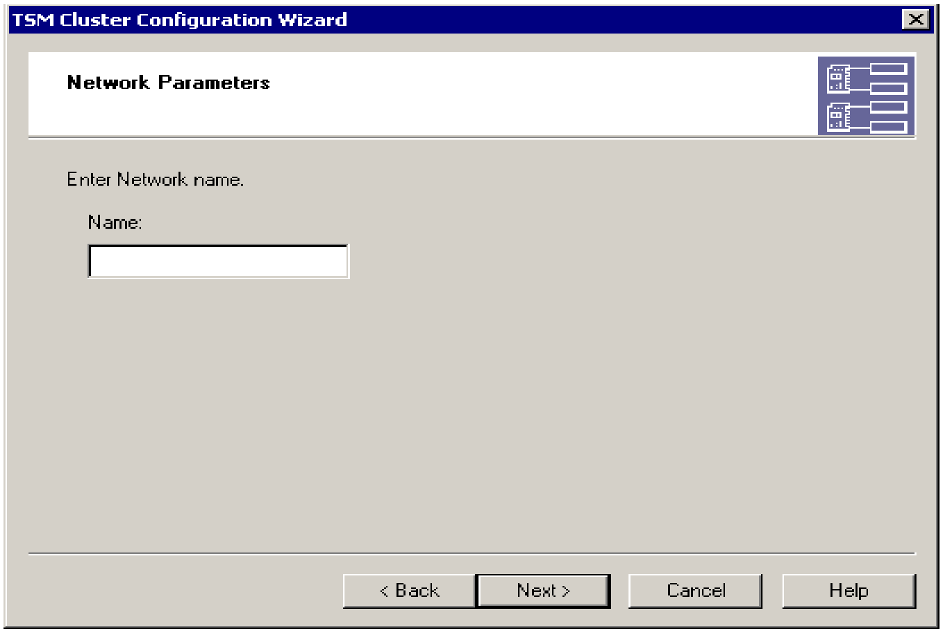

TSM Cluster Configuration Wizard

Network Parameters

Enter Network name.

Name:

Figure 32. TSM Cluster Configuration Wizard — Cluster Network Name

5. The Cluster Configuration Wizard completes, showing the configured setup (see Figure 33).

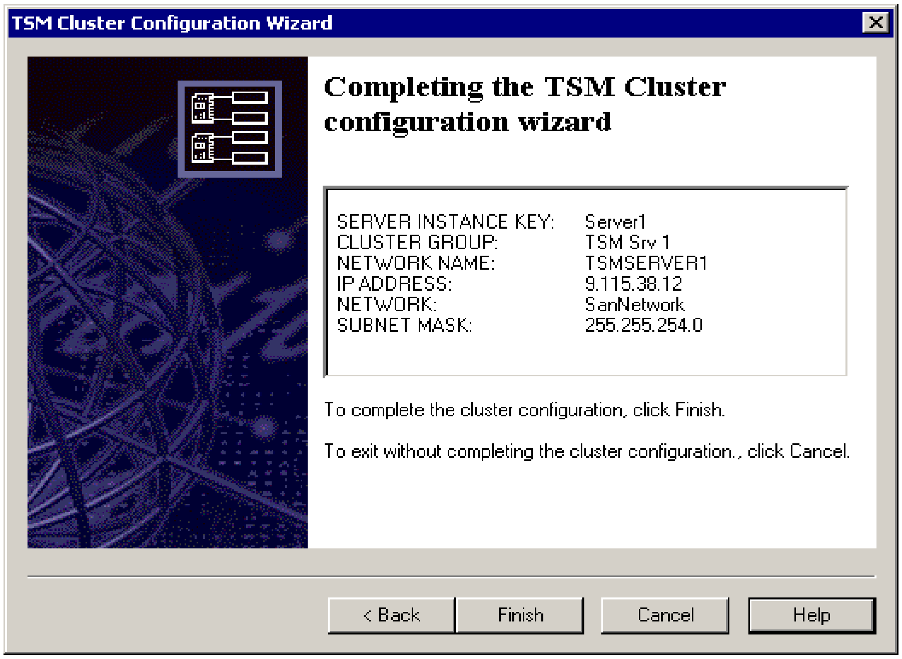

TSM Cluster Configuration Wizard

Completing the TSM Cluster configuration wizard

```
SERVER INSTANCE KEY:     Server1
CLUSTER GROUP:           TSM Srv 1
NETWORK NAME:            TSMSERVER1
IP ADDRESS:              9.115.38.12
NETWORK:                 SanNetwork
SUBNET MASK:             255.255.254.0
```

To complete the cluster configuration, click Finish.

To exit without completing the cluster configuration., click Cancel.

Figure 33. TSM Cluster Configuration Wizard — Cluster Definitions

6. Continue with the initial configuration. After completing the initial configuration, stop the server instance from the TSM Server Utilities. **Important:** Check to make sure the server instance is not running before proceeding to the next step.

Step 5 — Configure TSM on the Passive node

After you complete an active installation on one node, you can configure TSM on a Passive node.

1. From the Cluster Administrator, move the resource group (TSM Srv1) to the Passive node.
2. From the TSM Server Utilities, cancel the Initial Configuration wizard. Initial Configuration of the server instance to be used here was already performed during the server Installation on the Active node.
3. Select and start the TSM Server Initialization wizard. During the Server Initialization processing, the Windows Cluster is detected. The wizard asks if you want to use the cluster for this server instance. Click Yes and then Next.
4. Select and start the Cluster Configuration wizard. Select the group where this server instance will reside. It must match the group (TSM Srv 1) selected during the configuration of the Active node.
5. This time you will not be asked to provide an IP Address, Subnet mask, and Network name. The values provided during Active node configuration will be used.
6. After the Cluster Configuration Wizard completes, the Active-Passive cluster is configured, and the server instance is stopped. The group is then moved back to the Active server and the server instance is started.

5.1.5 TSM setup in an Active/Active MSCS configuration

To create an Active/Active configuration, you need to create an additional virtual server with a unique name. Creating additional server instances is a two-phase process:

5.1.5.1 Phase 1

From the Cluster Administrator

Verify that all nodes in the cluster are online. The minimum requirement is two nodes.

Verify that the group is online and owned by the node where you will be installing the server instance. Move groups as needed. All disk resources in the group must also be online.

From the Server Utilities

Start the Initial Configuration Wizard.

Select Configure a New Server. Make all other selections based on your environment.

During server initialization your Windows cluster is detected. Click Yes to configure TSM in a Windows Cluster.

Click the group where this server instance will reside.

During the Cluster Configuration, select the same group as in the previous step and enter the requested information from your planning worksheet.

After the Cluster Configuration Wizard completes, continue with the TSM initial configuration. Attention: After initial configuration is completed, stop the server instance from the Server Utilities.

5.1.5.2 Phase 2

From the Cluster Administrator

Verify that all nodes in the cluster are online. The minimum requirement is two nodes.

Verify that the group is online and owned by the node where you will be installing the initial server instance. Move groups as needed. All disk resources in the group must also be online.

From the Server Utilities

Cancel the TSM Initial Configuration wizard.

Select and start the TSM Server Initialization wizard. During server initialization the Windows Cluster is detected. Click Yes to configure TSM in a Windows Cluster.

Click the group where this server instance will reside. (This is the same group that you selected in Phase 1).

Note: Remember this group name; you will need to select the same group during the Cluster Configuration Wizard.

Select and start the Cluster Configuration wizard.

During the Cluster Configuration, select the same group as specified in the third step and verify the displayed information from your planning worksheet.

Ensure that the correct server instance is selected on the Logon Information dialog.

Once the Cluster Configuration Wizard completes, TSM is in an Active/Active cluster configuration.

5.1.6 Cluster management

The Microsoft Cluster Administrator is the preferred tool for managing TSM virtual servers. It is available through the Administrative Tools program group. The Cluster Administrator main window displays a detailed view of a virtual server configuration, including the physical Windows servers that make up the cluster and their resources, network connections, and status.

Use the Cluster Administrator to view the components of a virtual server configuration and to start, stop, or fail back a virtual server that has failed over. If you use the TSM Server Utilities or Windows NT Service Control Manager to stop a virtual TSM server, Clustering Service treats the shutdown as a failure and restarts the server on the cluster’s secondary node.

Note

When using tape drives in a clustered environment, the cluster administrator has to consider different configuration alternatives to support the failover of tape devices with minimal manual intervention.

Valid solutions are the usage of twin-tailed libraries or the connection of the library using a SAN. Also the usage of large disk pools (to cache backups) and a manual reconfiguration of the library can be considered.

5.2 Active Directory exploitation

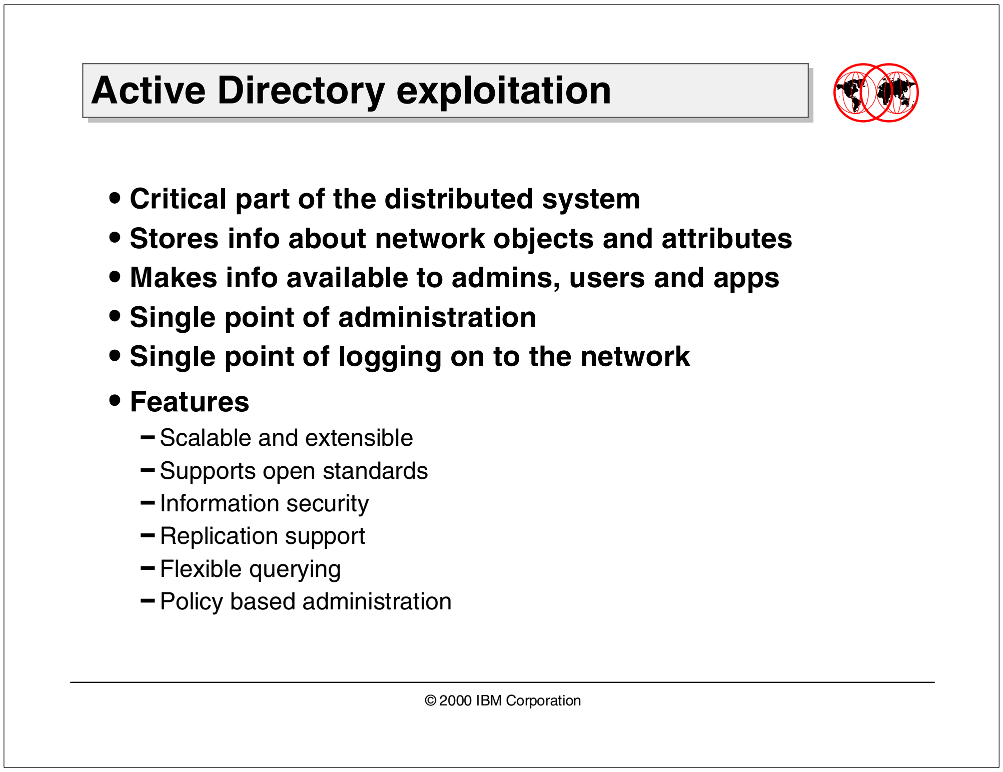

Active Directory exploitation

- **Critical part of the distributed system**
- **Stores info about network objects and attributes**
- **Makes info available to admins, users and apps**
- **Single point of administration**
- **Single point of logging on to the network**
- **Features**
  - Scalable and extensible
  - Supports open standards
  - Information security
  - Replication support
  - Flexible querying
  - Policy based administration

Active Directory is the directory service for Windows 2000 Server. It stores information about objects on the network and makes this information easy for administrators and users to find and use. Active Directory directory service uses a structured data store as the basis for a logical, hierarchical organization of directory information.

Active Directory provides a single point of administration for all published resources, including files, peripheral devices, host connections, databases, Web access, users, and other arbitrary objects and services. It provides a single point for logging on to the network.

TCP/IP is the mandatory networking protocol for the Windows 2000 Active Directory. Windows 2000 Server installs TCP/IP as its default protocol.

Active Directory is integrated with both Internet and intranet environments. It works with tools that users already use such as Web browsers. Scales from a small business with a few hundred objects on a few servers to the largest enterprise with millions of objects on thousands of servers.

Security is integrated with Active Directory through logon authentication and access control to objects in the directory. With a single network logon, administrators can manage directory data and organization throughout their network, and authorized network users can access resources anywhere on the network. Policy-based administration eases the management of even the most complex network.

5.2.1 Features of Active Directory

Active Directory includes the following features:

Support for open standards. Active directory supports open standards to facilitate cross-platform directory services, including support for standard protocols such as Lightweight Directory Access Protocol (LDAP), X500 and MAPI-RPC. Backward compatibility with previous version of NT and novell Netware environments.

Integration with DNS. Support is provided for standard name formats to ensure ease of use and migration. The Internet namespace standard is the Domain Name System (DNS), which translates names to numeric TCP/IP addresses. Active Directory leverages DNS as the service that is used to locate objects.

Flexible querying. Users and administrators can use the Search command on the Start menu, the My Network Places icon on the desktop, or the Active Directory Users and Computers snap-in to quickly find an object on the network using object properties. For example, you can find a user by first name, last name, e-mail name, office location, or other properties of that person's user account. Finding information is optimized by use of the global catalog.

Extensibility. Active Directory is extensible, which means that administrators can add new classes of objects to the schema and can add new attributes to existing classes of objects. The schema contains a definition of each object class, and each object class’s attributes, that can be stored in the directory.

Policy-based administration. Group Policies are configuration settings applied to computers or users as they are initialized. All Group Policy settings are contained in Group Policy Objects (GPOs) applied to Active Directory sites, domains, or organizational units. GPO settings determine access to directory objects and domain resources, what domain resources (such as applications) are available to users, and how these domain resources are configured for use.

Scalability. Active Directory includes one or more domains, each with one or more domain controllers, enabling you to scale the directory to meet any network requirements. Multiple domains can be combined into a domain tree and multiple domain trees can be combined into a forest. Management is through a hierarchical domain structure.

Open application programming interfaces (APIs). These can be used to query the directory and modify the records in the directory. Supported API’s include LDAP API, ADSI (Active Directory Service Interfaces) and MAPI.

Information security. Management of user authentication and access control, both fully integrated with Active Directory, are key security features in the Windows 2000 operating system. Active Directory centralizes authentication. Access control can be defined not only on each object in the directory, but also on each property of each object.

Information replication. Active Directory uses multi master replication, which lets you update the directory at any domain controller. Deploying multiple domain controllers in one domain provides fault tolerance and load balancing. If one domain controller within a domain slows, stops, or fails, other domain controllers within the same domain can provide necessary directory access, since they contain the same directory data.

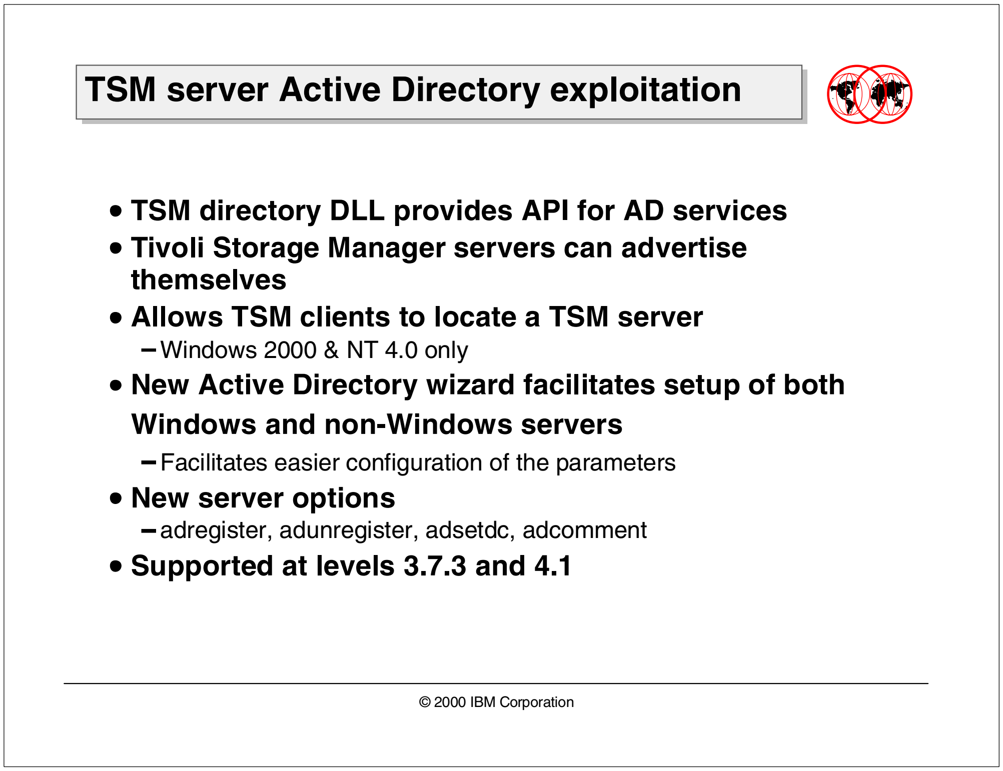

TSM server Active Directory exploitation

- **TSM directory DLL provides API for AD services**
- **Tivoli Storage Manager servers can advertise themselves**
- **Allows TSM clients to locate a TSM server**
  - Windows 2000 & NT 4.0 only
- **New Active Directory wizard facilitates setup of both Windows and non-Windows servers**
  - Facilitates easier configuration of the parameters
- **New server options**
  - adregister, adunregister, adsetdc, adcomment
- **Supported at levels 3.7.3 and 4.1**

5.2.2 Tivoli Storage Manager server Active Directory exploitation

Starting with Version 3.7.3 of Tivoli Storage Manager server on all platforms, server information including server name, list of registered clients, and communication parameters, can now be published in Active Directory (AD). The TSM Directory DLL provides an API that allows TSM to store and retrieve information from a directory service.

A new Active Directory Configuration Wizard has been added to the Server Utilities. An administrator can add, remove, or edit TSM server entries with this wizard. This provides a method for administrators to publish information about non-Windows TSM servers. The TSM backup-archive client configuration wizard includes the capability to browse server information in the AD. The client can use the information to determine which server to connect to and what communication protocol to use. Also, several new server options now support Active Directory.

At level 3.7.3 and above, three new server options have been added:

- ADREGISTER – Allows the TSM server to register with AD at startup. The valid choices are Yes or No. In order to support publishing of server information, this box should be checked.
- ADUNREGISTER — Allows the TSM server to unregister with AD at TSM server shutdown. The valid choices are Yes or No. This is useful so that the servers that are down are not published for clients as valid backup servers.
- ADSETDC — Sets the domain controller name on which AD resides.

An additional option has been added at level 4.1:

- ADCOMMENT — Any additional comments for client reference, and so on.

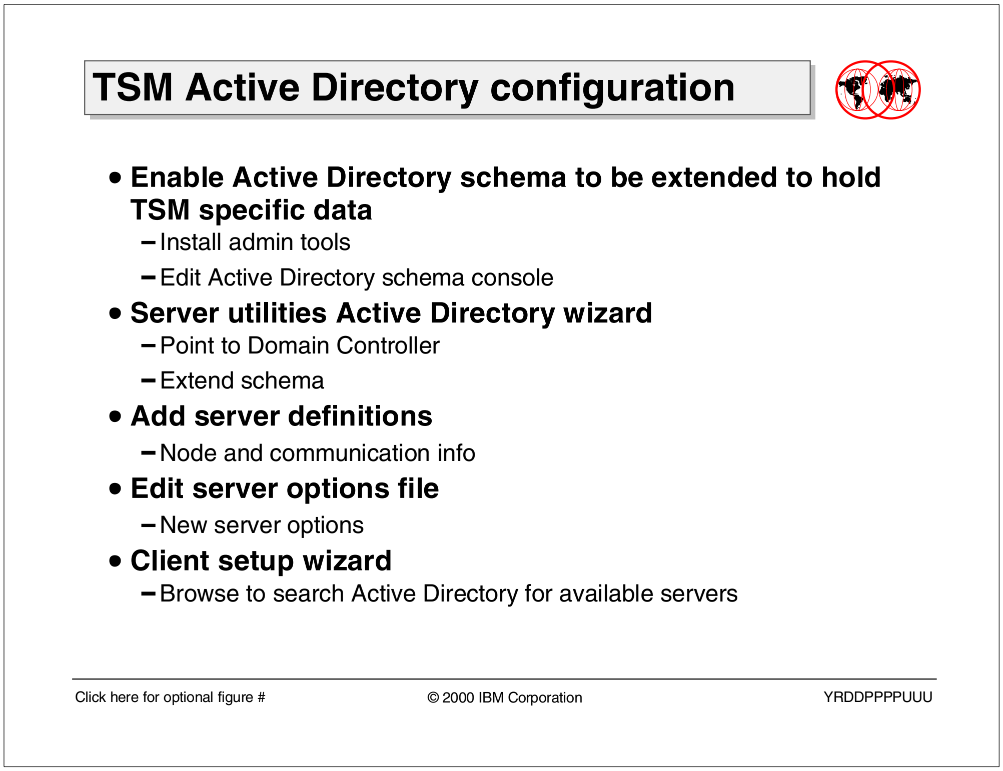

TSM Active Directory configuration

- **Enable Active Directory schema to be extended to hold TSM specific data**
  - Install admin tools
  - Edit Active Directory schema console
- **Server utilities Active Directory wizard**
  - Point to Domain Controller
  - Extend schema
- **Add server definitions**
  - Node and communication info
- **Edit server options file**
  - New server options
- **Client setup wizard**
  - Browse to search Active Directory for available servers

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

5.2.3 TSM Active Directory configuration

Active Directory should be properly set up and configured. The Admin tools for Windows 2000 should be installed on the domain controller. A console for *Active Directory schema* should be available. (For detailed information on all these activities, please see the online help utilities for Windows 2000).

The Active Directory schema is effectively a big data structure. It must be expanded to hold TSM specific data.

1. The first task is to enable the schema to be updated:
   a. Open the console for Active Directory schema.
   b. Right-click on Active directory schema and select Operations Master.
   c. Check the box which allows the schema to be modified on this domain controller.
   d. Right-click on Active Directory schema and select Permissions.
   e. Update an administrator account with Full Control permissions for the schema.
2. The second task is to extend the schema for TSM:
   a. Start the TSM Active Directory Configuration Wizard from TSM server utilities.
   b. Specify the proper domain controller name as shown in Figure 34.
   c. Click Next to configure/extend the schema.

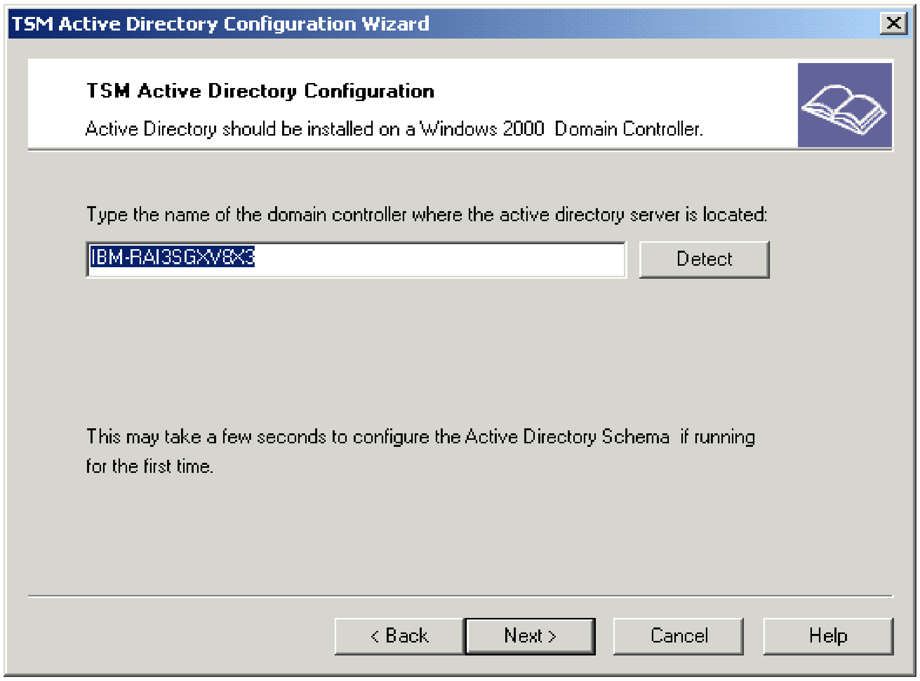

TSM Active Directory Configuration Wizard

TSM Active Directory Configuration

Active Directory should be installed on a Windows 2000 Domain Controller.

Type the name of the domain controller where the active directory server is located:

IBM-RAI3SGXV8X3

This may take a few seconds to configure the Active Directory Schema if running for the first time.

Figure 34. Schema extension

3. The third task is to define servers so they can be published in the Active Directory:
   a. The Active Directory configuration wizard takes you to the server definition screen as shown in Figure 35.
   b. Add all Windows and non-Windows servers, including list of registered nodes and communication information.

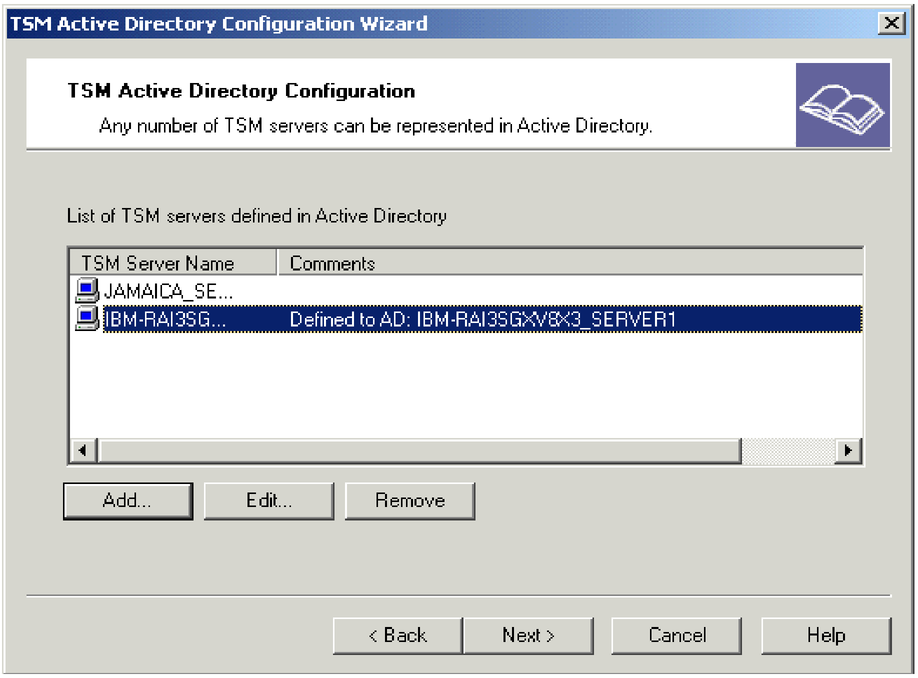

TSM Active Directory Configuration Wizard

TSM Active Directory Configuration

Any number of TSM servers can be represented in Active Directory.

List of TSM servers defined in Active Directory

TSM Server Name	Comments
JAMAICA_SE...	
IBM-RAI3SG...	Defined to AD: IBM-RAI3SGXV8X3_SERVER1

Figure 35. Server definition

4. The final task in preparing the TSM server is to update the server options:
    a. Open the Server Options File Editor in the server utilities and you will get the screen shown in Figure 36.
    b. Check the feature to register with Active directory on TSM server startup
    c. Check the feature to unregister with Active Directory on TSM server shutdown
    d. Select the proper domain controller, either by explicitly specifying its name or using the automatic detection feature.

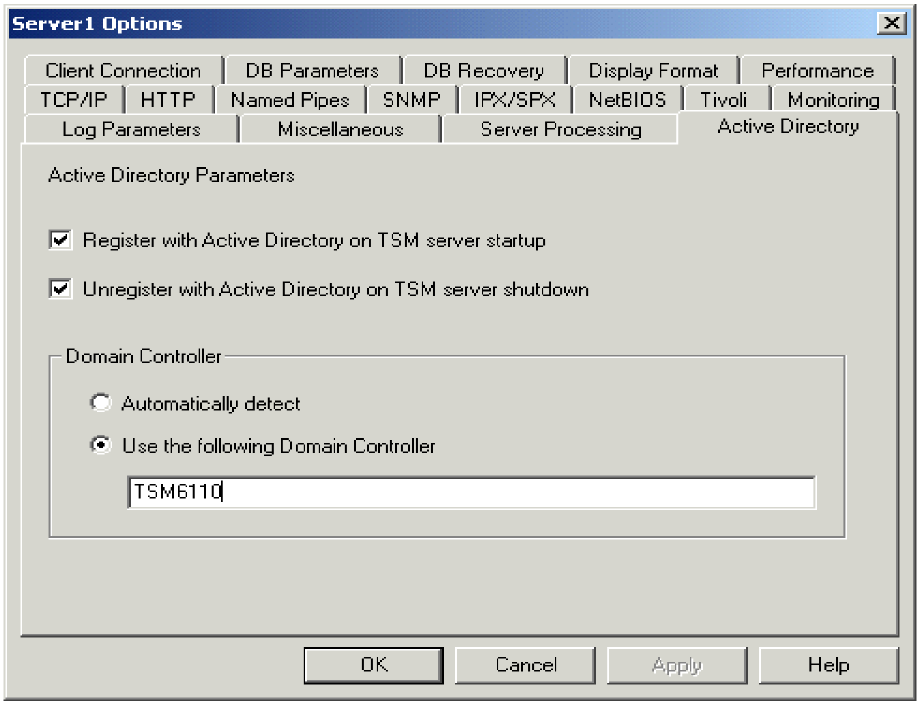

Active Directory Parameters

☑ Register with Active Directory on TSM server startup

☑ Unregister with Active Directory on TSM server shutdown

Domain Controller

○ Automatically detect

● Use the following Domain Controller

TSM6110

Figure 36. Tivoli Storage Manager Server options

5. To see the client side of things:
    a. Start the Backup Archive client GUI.
    b. Go to Utilities and start the Setup Wizard.
    c. At the stage where TCP/IP parameters of the TSM server have to be specified, a Browse option is available (compare with screen in the background of Figure 37).
    d. Clicking on Browse will bring up all information on all the TSM servers (as shown in the front window in Figure 37), that have registered this node and support the requested communication method.

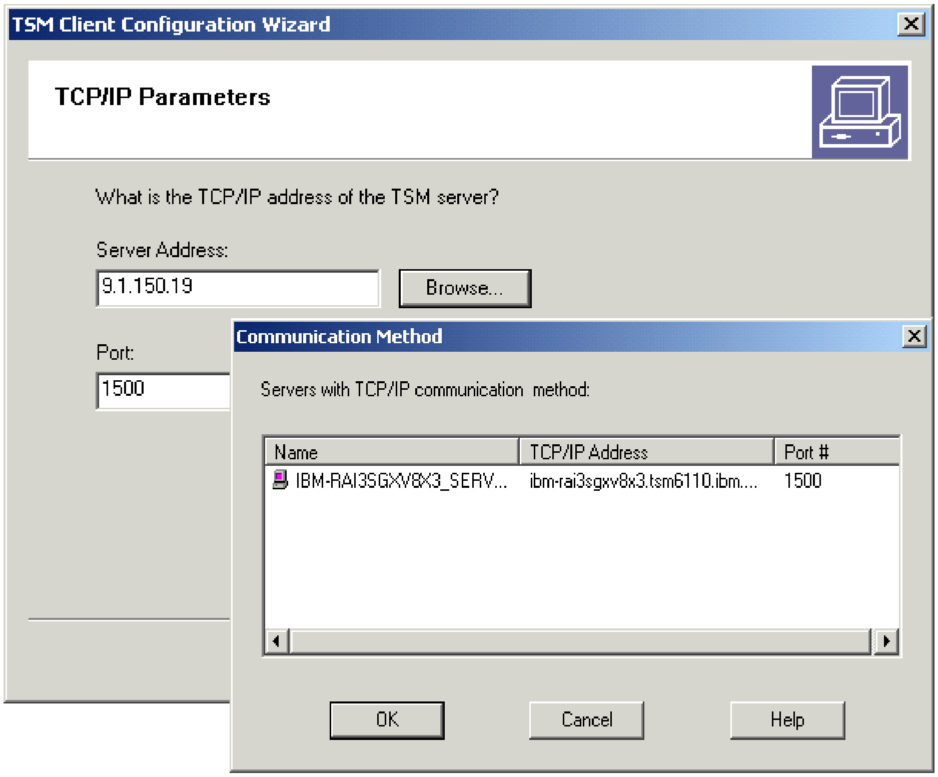

TCP/IP Parameters

What is the TCP/IP address of the TSM server?

Server Address: 9.1.150.19

Port: 1500

Servers with TCP/IP communication method:

Name	TCP/IP Address	Port #
IBM-RAI3SGXV8X3_SERV...	ibm-rai3sgxv8x3.tsm6110.ibm....	1500

Figure 37. Client query for published TSM servers

Note

Active Directory is supported on Windows 2000 and NT 4.0 clients only. It is not supported for Windows 95/98 clients.

5.3 TSM support for multiple server instances

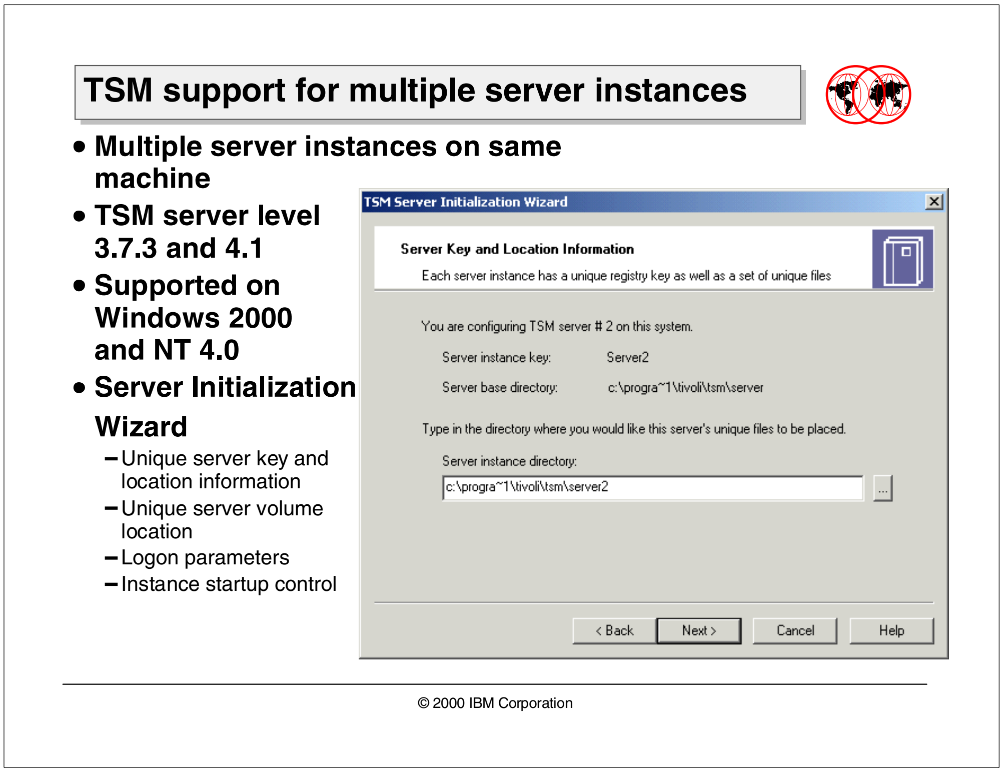

TSM support for multiple server instances

- **Multiple server instances on same machine**
- **TSM server level 3.7.3 and 4.1**
- **Supported on Windows 2000 and NT 4.0**
- **Server Initialization Wizard**
  - Unique server key and location information
  - Unique server volume location
  - Logon parameters
  - Instance startup control

TSM Server Initialization Wizard

Server Key and Location Information

Each server instance has a unique registry key as well as a set of unique files

You are configuring TSM server # 2 on this system.

Server instance key: Server2

Server base directory: c:\progra~1\tivoli\tsm\server

Type in the directory where you would like this server's unique files to be placed.

Server instance directory: c:\progra~1\tivoli\tsm\server2

© 2000 IBM Corporation

Starting with Version 3.7.3, Tivoli Storage Manager supports the user with easy setup of multiple instances of the server on a single Windows 2000 or NT 4.0 machine.

To support this functionality, server initialization has moved out of the setup program and into a *Server Initialization wizard* accessible through server utilities. The wizard prompts the user to type in the name of a directory where the server's unique files will be placed, as shown in Figure 38.

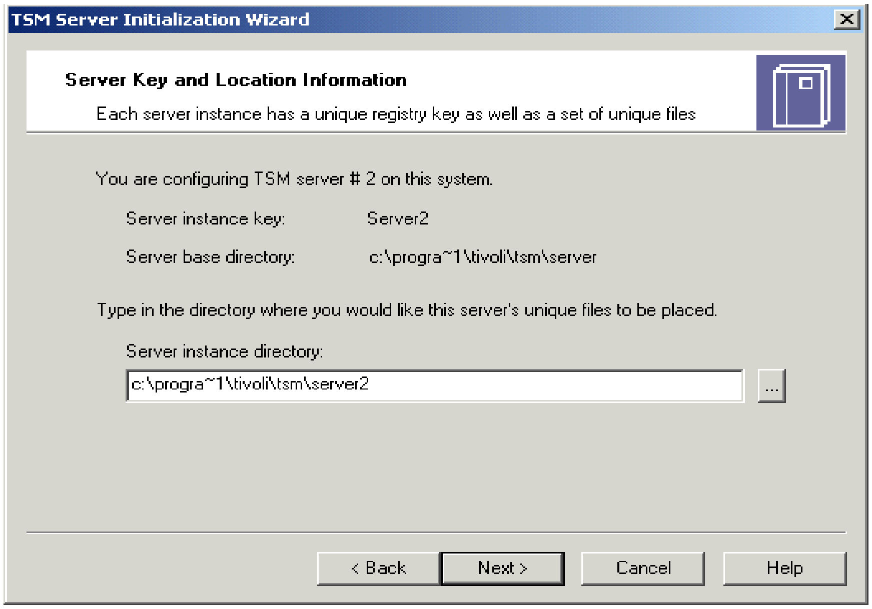

Server Key and Location Information

Each server instance has a unique registry key as well as a set of unique files

You are configuring TSM server # 2 on this system.

Server instance key: Server2

Server base directory: c:\progra~1\tivoli\tsm\server

Type in the directory where you would like this server's unique files to be placed.

Server instance directory:

c:\progra~1\tivoli\tsm\server2

Figure 38. Server instance location

The wizard then prompts the user to type in the name of a directory for database, recovery log, and storage pool volumes. Options are also provided at this stage to automatically create additional database and recovery log volumes as needed. It is recommended that the user select these options (see Figure 39).

The wizard also takes clustering and performance data into account during initialization.

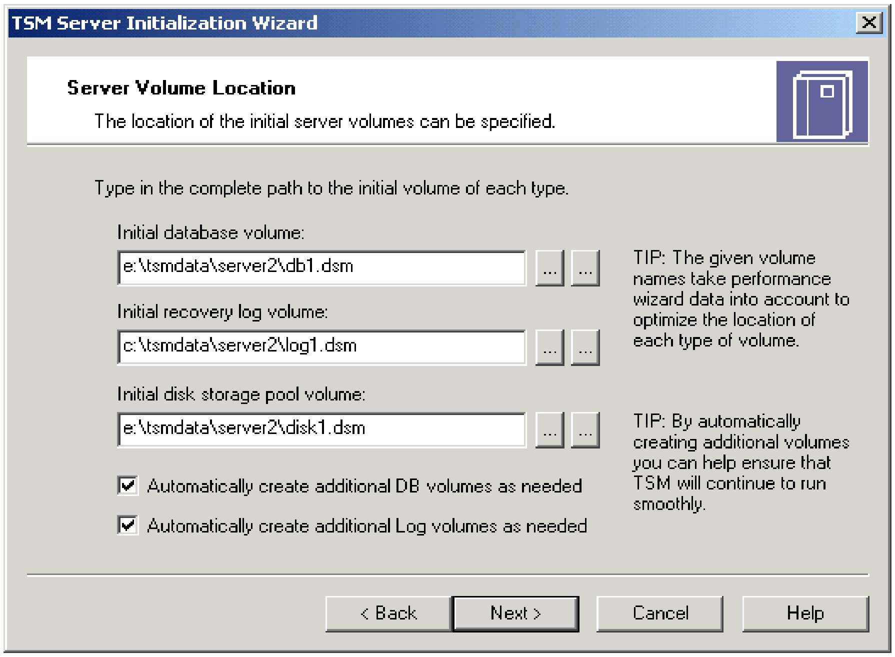

Server Volume Location

The location of the initial server volumes can be specified.

Type in the complete path to the initial volume of each type.

Initial database volume:

e:\tsmdata\server2\db1.dsm

Initial recovery log volume:

c:\tsmdata\server2\log1.dsm

Initial disk storage pool volume:

e:\tsmdata\server2\disk1.dsm

☑ Automatically create additional DB volumes as needed

☑ Automatically create additional Log volumes as needed

TIP: The given volume names take performance wizard data into account to optimize the location of each type of volume.

TIP: By automatically creating additional volumes you can help ensure that TSM will continue to run smoothly.

Figure 39. Volume locations

As a final step, the wizard inquires about the server service logon parameters, which include account information for managing the particular server instance, as shown in Figure 40.

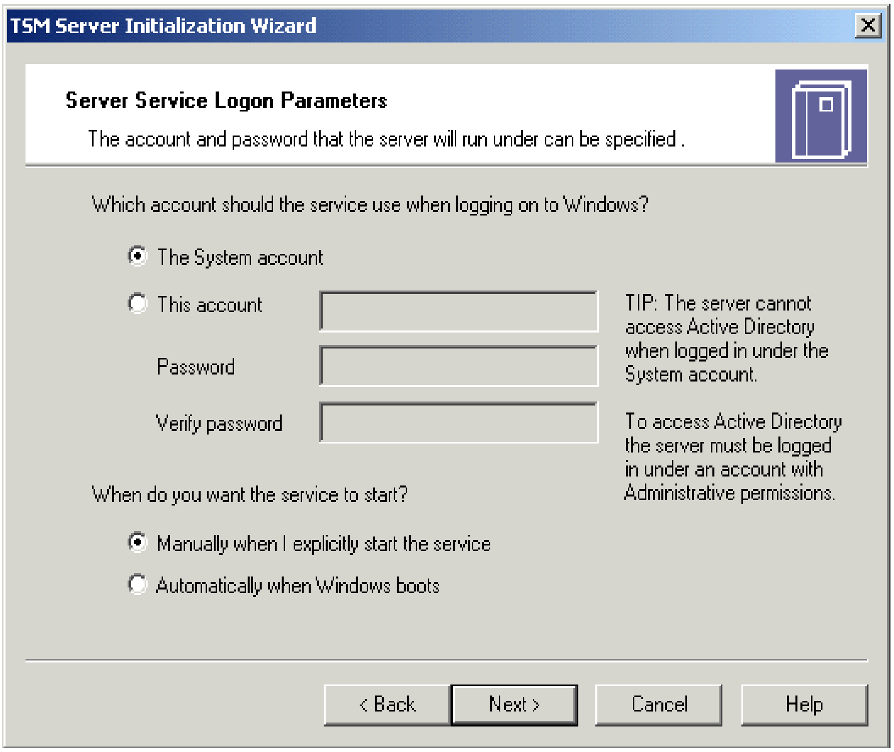

TSM Server Initialization Wizard

Server Service Logon Parameters

The account and password that the server will run under can be specified .

Which account should the service use when logging on to Windows?

● The System account

○ This account

Password

Verify password

TIP: The server cannot access Active Directory when logged in under the System account.

To access Active Directory the server must be logged in under an account with Administrative permissions.

When do you want the service to start?

● Manually when I explicitly start the service

○ Automatically when Windows boots

Figure 40. Startup Options

Startup options are also provided by the wizard. The user has the option to either explicitly start the instance manually or have it start up automatically when Windows boots.

All server utility wizards have been updated to take instance information into account.

5.4 Microsoft Management Console

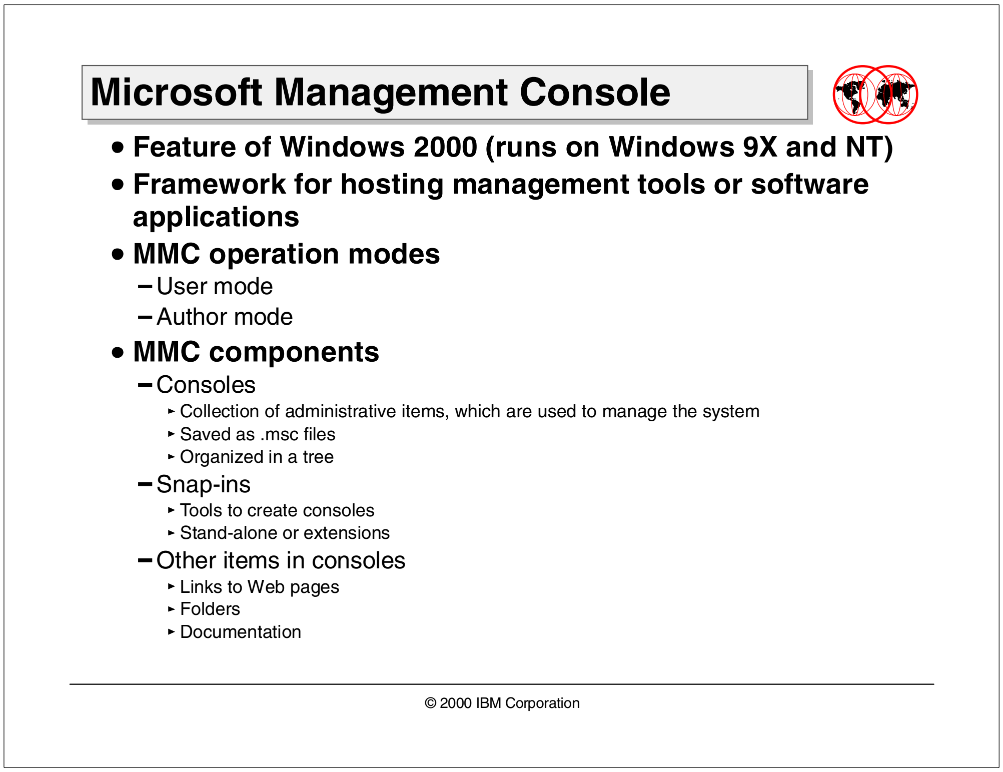

Microsoft Management Console

- **Feature of Windows 2000 (runs on Windows 9X and NT)**
- **Framework for hosting management tools or software applications**
- **MMC operation modes**
  - User mode
  - Author mode
- **MMC components**
  - Consoles
    - Collection of administrative items, which are used to manage the system
    - Saved as .msc files
    - Organized in a tree
  - Snap-ins
    - Tools to create consoles
    - Stand-alone or extensions
  - Other items in consoles
    - Links to Web pages
    - Folders
    - Documentation

© 2000 IBM Corporation

Microsoft Management Console (MMC) is used to create, save, and open administrative tools (called MMC consoles) that manage the hardware, software, and network components of your Windows system.

MMC is a feature of the Windows 2000 operating system, but you can also run MMC on Windows NT, Windows 95, and Windows 98 operating systems. In addition, MMC is a feature of several software applications designed to run on Windows.

MMC does not perform administrative functions, but hosts tools that do. The primary type of tool you can add to a console is called a snap-in. Other items that you can add include ActiveX controls, links to Web pages, folders, taskpad views, and tasks.

There are two general ways that you can use MMC:

User mode — Working with existing MMC consoles to administer a system.

Author mode — Creating new consoles or modifying existing MMC consoles.

5.4.1 Consoles and snap-ins

In basic terms, a console is a collection of administrative items. A snap-in is a primary tool used to create a console. Other items in the console may include wizards, Web links, documentation, monitor controls, and so on. The primary purpose of a console is to manage the specific aspect of the system for which it was designed.

5.4.1.1 Consoles

A new MMC *console* consists of a window divided into two panes. The left pane contains two tabs: the Tree tab and the Favorites tab. The right pane contains the details pane. The Tree tab, also called the console tree, shows the items that are available in a given console. The details pane shows information about, and functions pertaining to, these items. As you click different items in the console tree, the information in the details pane changes. The details pane can display many types of information, including Web pages, graphics, charts, tables, and columns. Each console has its own menus and toolbar, separate from those of the main MMC window, that help a user perform tasks. The operating system you are using may already have preconfigured and saved console files available on the Programs menu or in the Administrative Tools folder in Control Panel.

Consoles are highly customizable and are also persistent. In general, users have the ability to add to the console or update it to include all their preferred snap-ins. When a console is saved, it is saved as a .msc file. This file contains all the persistent information for each contained snap-in. These custom consoles can be sent by E-mail, shared in a network folder or posted on the Web. They can also be assigned to users, groups or computers with Group Policy settings. As an example, Figure 41 shows the new TSM MMC.

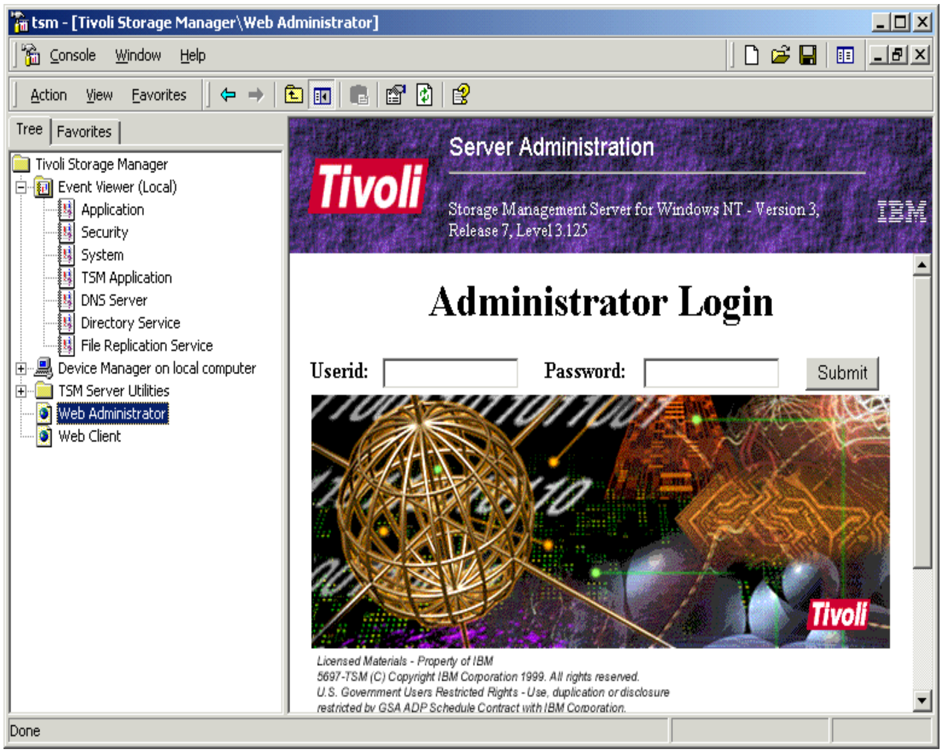

Storage Management Server for Windows NT - Version 3, Release 7, Level 3.125

Administrator Login

Licensed Materials - Property of IBM
5697-TSM (C) Copyright IBM Corporation 1999. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corporation.

Figure 41. TSM MMC console

5.4.1.2 Snap-ins

A *snap-in* is the basic component of an MMC console. Snap-ins always reside in a console; they do not run by themselves. When you install a component that has a snap-in associated with it on a computer running Windows, the snap-in is available to anyone creating a console on that computer (unless restricted by a user policy).

MMC supports two types of snap-ins: stand-alone snap-ins and extension snap-ins. You can add a stand-alone snap-in, usually just called a snap-in, to a console tree without adding another item first. An extension snap-in, usually called an extension, is always added to a stand-alone or extension snap-in that is already on the console tree. When extensions are enabled for a snap-in, they operate on the objects controlled by the snap-in, such as a computer, printer, modem, or other device. When you add a snap-in or extension to a console, it may appear as a new item in the console tree, or it may add context menu items, additional toolbars, additional property pages, or wizards to a snap-in already installed in the console.

You can add a single snap-in or multiple snap-ins and other items to a console. In addition, you can add multiple instances of a particular snap-in to the same console to administer different computers or to repair a damaged console. Each time you add a new instance of a snap-in to a console, any variables for the snap-in are set at default values until you configure the snap-in. For instance, if you configure a particular snap-in to manage a remote computer and then add a second instance of the snap-in, the second instance will not automatically be configured to manage the remote computer. In general, you can only add snap-ins that are installed on the computer you are using to author a console. However, in Windows 2000, if your computer is part of a domain, you can use MMC to download any snap-ins that are not locally installed, but that are available in the Active Directory directory service.

Within MMC you can add and remove snap-ins through the parent window Console menu with the Add/Remove snap-in command as shown in Figure 42.

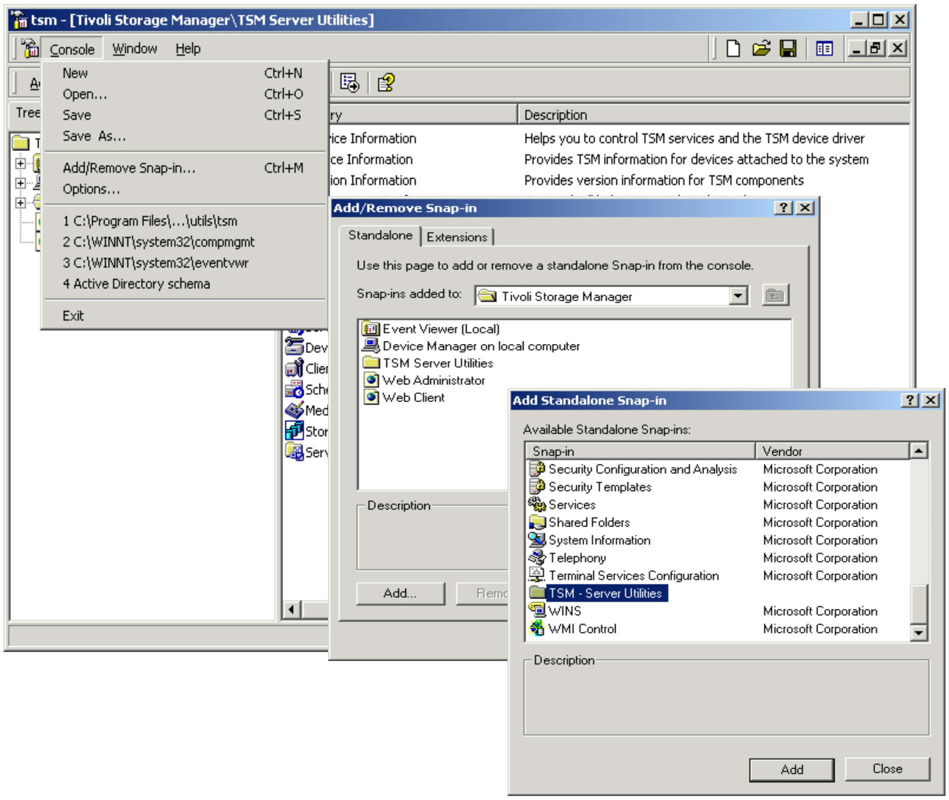
tsm - [Tivoli Storage Manager\TSM Server Utilities]
Console
Window
Help
New Ctrl+N
Open... Ctrl+O
Save Ctrl+S
Save As...
Add/Remove Snap-in... Ctrl+M
Options...
1 C:\Program Files\...\utils\tsm
2 C:\WINNT\system32\compmgmt
3 C:\WINNT\system32\eventvwr
4 Active Directory schema
Exit
Tree
Description
Dev
Clier
Sche
Med
Stor
Serv
Add/Remove Snap-in
Standalone
Extensions
Add...
Remo
Add Standalone Snap-in
Add
Close

ry	Description
ice Information	Helps you to control TSM services and the TSM device driver
ce Information	Provides TSM information for devices attached to the system
ion Information	Provides version information for TSM components

Use this page to add or remove a standalone Snap-in from the console.

Snap-ins added to: Tivoli Storage Manager

Event Viewer (Local)
Device Manager on local computer
TSM Server Utilities
Web Administrator
Web Client

Description

Available Standalone Snap-ins:

Snap-in	Vendor
Security Configuration and Analysis	Microsoft Corporation
Security Templates	Microsoft Corporation
Services	Microsoft Corporation
Shared Folders	Microsoft Corporation
System Information	Microsoft Corporation
Telephony	Microsoft Corporation
Terminal Services Configuration	Microsoft Corporation
TSM - Server Utilities	
WINS	Microsoft Corporation
WMI Control	Microsoft Corporation

Description

Figure 42. Snap-ins

TSM management console/snap-ins

- **TSM management console**
  - Shipped with TSM as "tsm.msc"
  - TSM snap-ins
    - Web client
    - Web administrator
    - Server utilities
  - System snap-ins
    - Event viewer
    - Device manager
- **Original server utilities GUI still used**
  - Needed if MMC 1.2 not installed

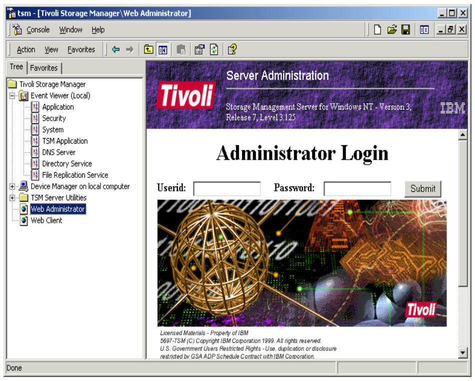

Storage Management Server for Windows NT - Version 3, Release 7, Level 3.125

Administrator Login

Licensed Materials - Property of IBM
5697-TSM (C) Copyright IBM Corporation 1999. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corporation.

5.4.2 Tivoli Storage Manager exploitation

A TSM management console is shipped with the product and includes TSM and system related snap-ins (as already shown in Figure 41 on page 162). These include:

- Server utilities
- Web administrator
- Web client
- Device manager
- Event viewer

The console can be activated by going into `Start`, then selecting `Programs`, then selecting `Tivoli Storage Manager`, and finally selecting `MMC Admin Console`. This executes the `tsm.msc` file in the background. All files with an extension of.msc are recognized as console files and will launch the application `mmc`. The location of this file is:

```
C:\Program Files\Tivoli\TSM\Utils\tsm.msc
```

The original server utilities GUI can still be built and used independently of the snap-in. This may be required if the user does not have MMC 1.2 installed. These can be activated by going into `Start`, then selecting `Programs`, then selecting `Tivoli Storage Manager`, and finally selecting `Server Utilities`. This runs the command `adsm` in the background. The location of the executable file is:

```
C:\Program Files\Tivoli\TSM\Utils\adsm.exe
```

5.5 Terminal Services performance enhancements

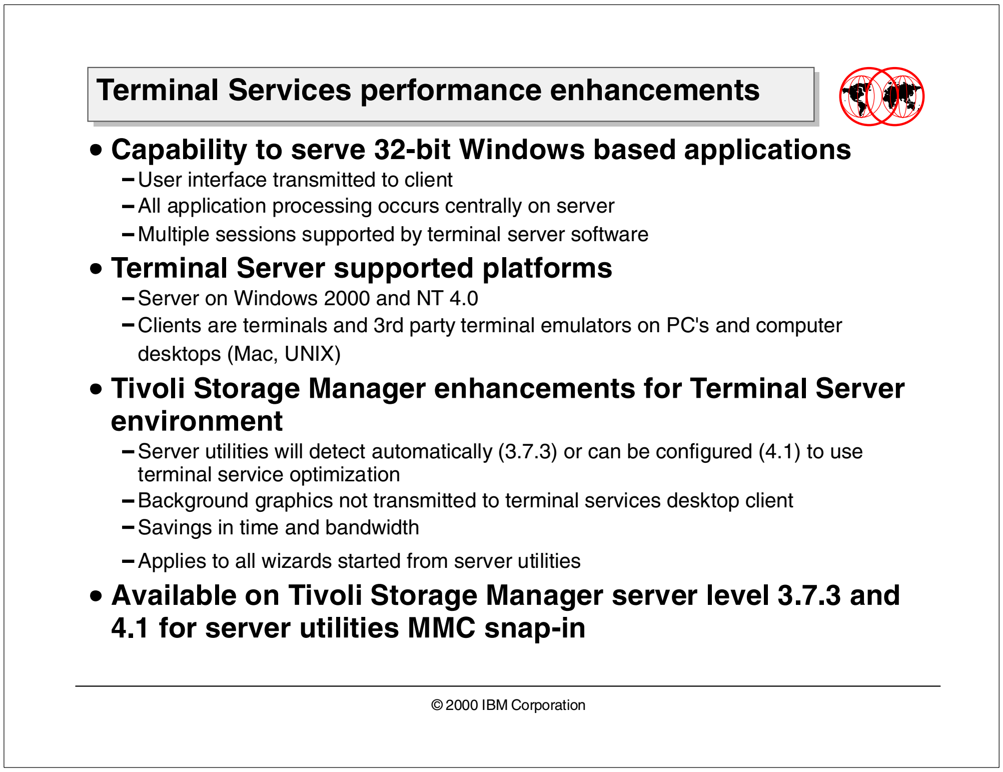

Terminal Services performance enhancements

- **Capability to serve 32-bit Windows based applications**
  - User interface transmitted to client
  - All application processing occurs centrally on server
  - Multiple sessions supported by terminal server software
- **Terminal Server supported platforms**
  - Server on Windows 2000 and NT 4.0
  - Clients are terminals and 3rd party terminal emulators on PC's and computer desktops (Mac, UNIX)
- **Tivoli Storage Manager enhancements for Terminal Server environment**
  - Server utilities will detect automatically (3.7.3) or can be configured (4.1) to use terminal service optimization
  - Background graphics not transmitted to terminal services desktop client
  - Savings in time and bandwidth
  - Applies to all wizards started from server utilities
- **Available on Tivoli Storage Manager server level 3.7.3 and 4.1 for server utilities MMC snap-in**

© 2000 IBM Corporation

5.5.1 Microsoft Terminal Services overview

Terminal Services provides remote access to a server desktop through "thin client" software, serving as a terminal emulator. Terminal Services transmits only the user interface of the program to the client. The client then returns keyboard and mouse clicks back to be processed by the server. With Terminal Services, all program execution, data processing and data storage occur on the server, centralizing the deployment of programs.

Each user logs on and sees only their individual session, which is managed transparently by the server operating system and is independent of any other client session.

Client software can run on a number of client hardware devices, including computers and Windows-based terminals. Other devices, such as Macintosh computers or UNIX-based workstations, can also connect to a terminal server with additional third party software.

Terminal Services can be deployed on the server in either application server or remote administration mode.

As an application server, Terminal Services provides an effective and reliable way to distribute Windows-based programs with a network server. In application server mode, Terminal Services delivers the Windows desktop and the most current Windows based applications to computers that might not normally be able to run Windows.

When used for remote administration, Terminal Services provides remote access for administering your server from virtually anywhere on your network.

5.5.2 Tivoli Storage Manager server enhancements

Splash screens and other graphical information that displays while an application is running works well in local video environments because of the speed of delivery. When that same splash screen is transmitted to a Terminal Server desktop client over the network, the transmission consumes extra network bandwidth and forces the user to wait when using the application. This is especially true when Terminal Server Clients are run over low bandwidth connections such as phone lines.

The Tivoli Storage Manager server utilities currently provide graphical splash and about screens, and they paint background bitmaps on all wizard pages in accordance with the Wizard97 style guidelines.

Enhancements have been made that optimize the performance of the server utilities when running under a Terminal Server Client.

The server utilities are being re-hosted as a Microsoft Management Console snap-in. Snap-in conventions provide a standardized About box with no graphics, and do not display a splash screen, as shown in Figure 43.

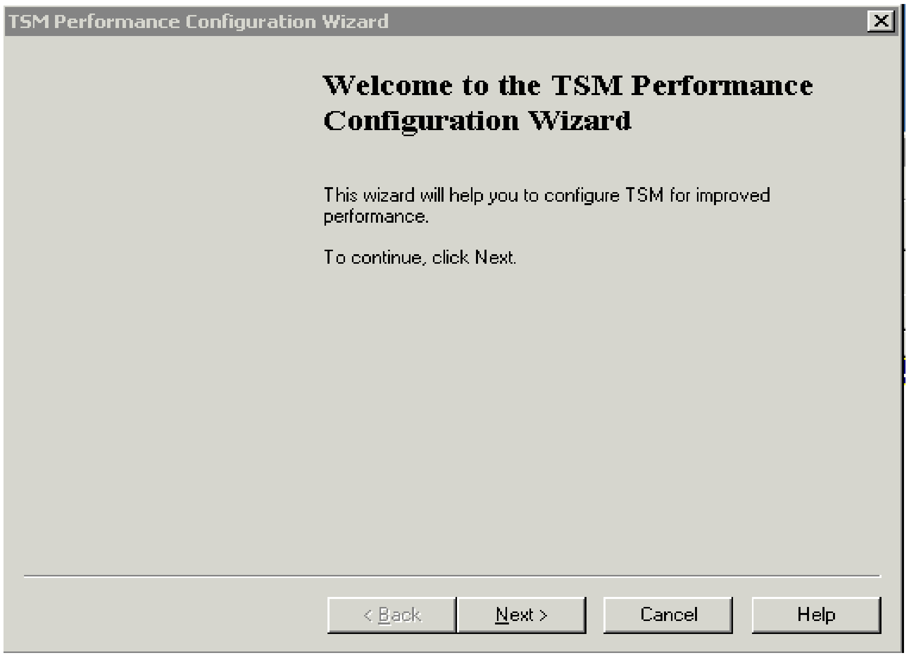

Welcome to the TSM Performance Configuration Wizard

This wizard will help you to configure TSM for improved performance.

To continue, click Next.

Figure 43. No background graphics for the wizards

Tivoli Storage Manager server utilities will automatically detect when they are being run in a Terminal Server session. In a Terminal Server session, they will not display the splash screen or any background wizard graphics, thus significantly improving interactive performance.

5.6 Removable Storage Manager support

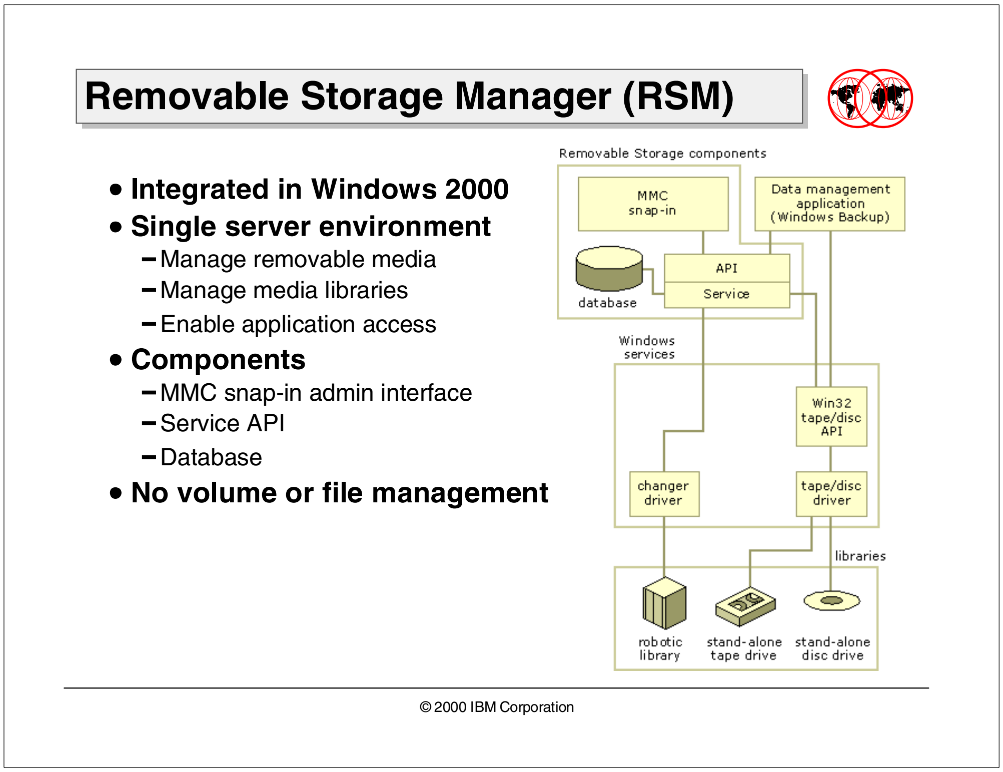

Removable Storage Manager (RSM)

- Integrated in Windows 2000
- Single server environment
  - Manage removable media
  - Manage media libraries
  - Enable application access
- Components
  - MMC snap-in admin interface
  - Service API
  - Database
- No volume or file management

© 2000 IBM Corporation

Removable Storage Manager (RSM) is a service included with Microsoft Windows 2000 that facilitates communication among applications, robotic changers, and media libraries. It enables multiple applications to share local robotic media libraries and tape or disk drives, and manage removable media within a single-server system.

Removable Storage makes it easy for you to track your Removable Storage media (tapes and optical discs) and to manage the hardware libraries, such as changers and jukeboxes, that contain them. The Removable Storage Manager labels, catalogs, and tracks media; controls library drives, slots, and doors; and provides drive cleaning operations.

Removable Storage works in concert with your data-management programs, such as Backup. You use data-management programs to manage the actual data stored on the media.

Removable Storage makes it possible for multiple programs to share the same storage media resources, which can reduce your costs. Removable Storage organizes all the media in your libraries into different media pools. Removable Storage also moves media between media pools in order to provide the amount of data storage your applications require.

Removable Storage does not provide volume management, such as for media siding or striping. Also, Removable Storage does not provide file management, such as for data backup or disk-extender operations. These services are done by data-management applications such as Tivoli Storage Manager.

You must run all your data-management programs on the same computer that connects to your library. Removable Storage does not support multiple data-management programs running on different computers that are connected to the same library.

Note

RSM needs the NTFS 5 structures and therefore is not available on Windows 95/98, Windows NT 4, and Windows 2000 Professional.

5.6.1 RSM components

Removable Storage Manager is comprised of a Windows 2000 service with an API, a database and an administrative interface implemented as a Microsoft Management Console (MMC) snap-in.

5.6.1.1 RSM API

The RSM API is the interface that applications use to access the services and media in an RSM system. It is a dynamic-link library (DLL) that contains a group of functions you can use to build your data management application. These functions perform the following operations:

- Mount and dismount media
- Clean drives
- Insert and eject media
- Inventory libraries
- Enable and disable libraries, drives, and media.
- Access information in the RSM database.

When you develop your media management application using the RSM API, you gain the following benefits:

- Multiple applications can share the same library, drive, and media resources.
- Code written to the RSM API for a specific library can execute on a wide range of libraries, both those that exist now and those that are introduced in the future.
- A single computer can track multiple types of media.
- A single computer can track media inside or outside a media library unit.

5.6.1.2 RSM database

The RSM database contains information about the objects in an RSM system. RSM maintains the data in the database and updates it whenever the administrator, an application, or a device makes a change in an RSM system. For example, if the administrator wants to mount a piece of media in an online library, RSM mounts the media and updates the database to reflect the change. The database includes the following information:

- Library, drive, and media configuration and state
- Media pool configuration and contents
- Library work list
- Operator requests

5.6.1.3 RSM MMC snap-in

The RSM MMC snap-in is an administrative interface to RSM that is a snap-in for the Microsoft Management Console (MMC) for Windows 2000. Administrators can use this snap-in to insert and eject media, perform inventories, mount and dismount media, clean drives, and check status information. RSM objects can be physical objects (such as libraries, drives, and media) or groups of media (called media pools).

Most users should never be required to use the snap-in, because each RSM client application takes care of routine tasks. For example, a simple backup application should not require a user to go to the snap-in to place media in a drive. In general, administrative support should only be required if a system is running more than one RSM client application, backup and a document management system, for example, or if the system has a sophisticated robotic library that requires administrative management. The MMC snap-in can help administrators perform the following tasks:

- Complete or refuse operator requests
- Perform point-and-click operations, such as inserting media, ejecting media, drive cleaning, dismounting media, and moving media from one pool to another
- Cancel work queue items
- Enable and disable drives and libraries

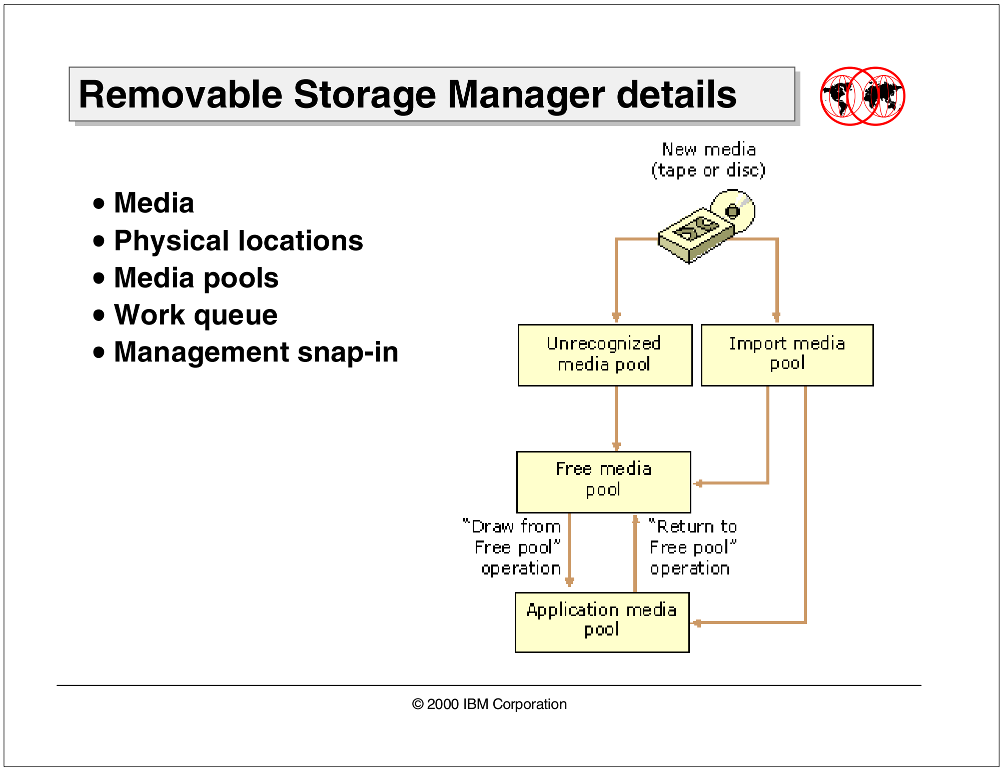

Removable Storage Manager details

- Media
- Physical locations
- Media pools
- Work queue
- Management snap-in

New media
(tape or disc)

Unrecognized media pool

Import media pool

Free media pool

"Draw from Free pool" operation

"Return to Free pool" operation

Application media pool

5.6.2 Removable Storage Manager details

Removable Storage Manager can be described in terms of five basic concepts: media, physical locations, media pools, work queue items, and operator requests. The last four of these are the top level nodes in the RSM snap-in. The first, media, is the most fundamental concept and affects all the others.

5.6.2.1 Media

Units of *media* store information. Each unit of media, or medium, also referred to as a "cartridge", is of a certain type, such as 8mm tape, magnetic disk, optical disk, or CD-ROM. Removable Storage Manager represents media physically and logically. Physical media are the tangible media that are inserted and removed from libraries and mounted in drives.

When an application needs to access data on a cartridge, Removable Storage Manager generates a logical identifier (ID) that allows the application to request the data on that cartridge. Since access to the data occurs only through that ID, Removable Storage Manager can manage the physical location of the data. For example, if the original cartridge begins to fail, Removable Storage Manager can move the data to a new cartridge without having to notify the application. Physical media and the sides they contain are tracked in Removable Storage Manager and can take on various states as they are entered into a library and used.

Removable Storage Manager provides the infrastructure necessary to share media among various applications. Removable Storage Manager ensures that all media are used in such a way as to preserve the data they contain. Removable Storage Manager accomplishes this by identifying and verifying the identity of each cartridge.

5.6.2.2 Media identifiers

Removable Storage uses two identification methods to track media and perform inventories:

- On-media identifiers
- Bar codes

On-media identifiers — *On-media identifiers* are electronically recorded on media the first time a tape or disc is inserted into a library. The next time the medium is inserted into the library, Removable Storage uses the on-media identifier to recognize and track the tape or disc.

An on-media identifier has two parts: a label type and a label ID. The label type identifies the specific format used to record information on the tape or disc, such as Microsoft Tape Format (MTF). The label ID is a unique identifier for the tape or disc in the Removable Storage system. If Removable Storage Manager does not recognize the label type, it puts the medium in an unrecognized media pool. If Removable Storage recognizes the label type but does not recognize the label ID, it puts the media in an import media pool. If Removable Storage recognizes the label type and the label ID, it updates the Removable Storage database to reflect that the medium is online and available for use.

> **Note**
>
> Removable Storage Manager does not record on-media identifiers for read-only or write-once optical media, such as CD-ROM, DVD-ROM, CD-R, DVD-R, or WORM discs. Instead, Removable Storage uses the volume and serial number information that is already associated with the disc.

Bar codes — If your library supports *bar codes*, Removable Storage Manager can use them to identify media. Media with bar codes also have on-media identifiers, and Removable Storage can use either one. Using bar codes to track media is generally much faster because each tape or disc does not have to be mounted to read the on-media identifier.

5.6.2.3 Media states

The *media state* identifies the operational status of each tape or disc in a Removable Storage system and therefore determines what operations are permitted. When media are inserted into a library, Removable Storage uses the on-media identifier or bar code to determine its state.

There are two sets of media states that govern media movement and usage by Removable Storage:

Physical states — The *physical states* identify operations performed that involve media location or movement. There are five possible physical states:

- Idle — Media is currently in a storage slot in a robotic library or is shelved offline.
- In-use — Media is in the process of being moved.
- Loaded — Media is mounted in a drive, and data is available to be read or written.

- Mounted — Media is mounted in a drive, but data is not yet available to be read or written.
- Unloaded — Media has been dismounted and is ready to be removed from a drive.

Side states — A tape or disc side is where actual information is stored. *Side states* identify operations performed that involve media usage rather than location or movement. There are nine possible side states:

- Allocated — Side is reserved for use by an application. The side is not available to any other application.
- Available — Side is currently available for use by an application.
- Completed — Side is in use, but cannot be used for write operations. The side is full.
- Decommissioned — Side is no longer available for use, as it has reached its allocation maximum.
- Imported — Side's label type is recognized, but its label ID is not recognized by RSM.
- Incompatible — The media type of the side is not compatible with the library. The tape or disc should be removed from the library.
- Reserved — Side is unavailable for allocation except by a single application. This applies only to two-sided media where one side has already been allocated.
- Unprepared — Side has been placed in a free media pool, but does not yet have a free media label. This is a temporary state.
- Unrecognized — Side's label type and label ID are not recognized by RSM.

There is a relationship between side states and media pools:

- Import media pools can only hold media sides in the imported state.
- Unrecognized media pools can only hold media sides that are in the unrecognized state.
- Free media pools can only hold media sides that are in the available (and transitional unprepared) state.

5.6.2.4 Physical locations

Removable Storage Manager manages two classes of *physical locations*:

- Libraries
- Offline media

Libraries include both media and the means to read and write them. The *offline media* physical location is a special holder for media that is cataloged by Removable Storage, but does not reside in a library. The set of all physical locations includes the libraries and the offline media physical location.

5.6.2.5 Libraries

In its simplest form, a *library* is composed of data storage media and a means to read and write the media. A CD-ROM drive with a disc inserted is a simple library with one drive, no slots, and no transport. A more complex example of a library is a robotic-based tape library, which holds several (up to several thousand) tapes, has one or more tape drives, and has a mechanical means to move tapes into and out of the drives. Specifically, libraries are composed of the following:

Slots — are storage locations in the library. A stand-alone drive library has no slots. Sometimes slots are organized into collections of slots called magazines. Magazines are usually removable.

Drives — A drive is a device that can read or write to a cartridge. A library has at least one drive.

Transports — A transport is the robotic device that moves a cartridge from its slot to a drive and back again. Robotic libraries usually have only one transport.

Bar code readers — A label with a bar code is attached to the outside of a cartridge that is both human readable and computer readable. Libraries that hold media with bar codes attached might have a bar code reader.

Doors — A door is a means to gain unconstrained access to the media in a library. Doors on smaller libraries might not always look exactly like doors, but they still provide the same functionality. For example, some small changers have all their slots in a magazine.

Insert/Eject Ports — While doors provide unconstrained access to the media in a library, an insert/eject (IE) port provides very controlled access. When an administrator adds media to a library through an IE port, the media are not placed directly into a slot. Instead, the media are placed in the IE port and then the library uses the transport to move the media from the IE port to a slot.

5.6.2.6 Types of libraries

Robotic — A robotic library can have any of the components described above. At a minimum, a robotic library has some media, slots to hold the media, one or more drives, a transport, and either a door or an IE port. Robotic libraries are sometimes referred to as changers or jukeboxes. No human intervention is required to place a cartridge in a library in one of its drives.

Stand-alone — A stand-alone drive library or stand-alone drive is a special kind of library. A cartridge must be manually placed in a drive. The CD-ROM drive found on most desktop computers is a stand-alone drive library. Removable Storage Manager treats stand-alone drives as libraries with one drive and an IE port. Removable Storage Manager models these drives as libraries to make its program interfaces simpler. For this reason, an application need not know whether a cartridge is mounted by a transport or a human.

5.6.2.7 Offline media physical locations

The offline media physical location is where Removable Storage Manager lists media that are not in a library. The physical location of media in an online library is the library in which it resides. Media that are not in any of these libraries, such as archived backup tapes on a shelf, are offline media and reside in the offline media physical location.

When a user or administrator moves an offline cartridge into a library, Removable Storage Manager changes its physical location to be the library into which it was placed. When a cartridge is taken out of an online library, Removable Storage Manager notes that it now resides in the offline media physical location.

5.6.2.8 Media pools

A *media pool* is a collection of tapes or discs to which the same management properties apply. All media in a Removable Storage system belong to a media pool, and each media pool holds only one type of media. Data management programs use media pools to gain access to specific tapes or discs within a library.

Using media pools, you can define properties that apply to a set of media. This is useful because Removable Storage allows multiple programs to share the same media within a single library. A library can include media from different media pools, each with different properties. A single media pool can span multiple libraries. You can also create hierarchies of media pools, or media pools that contain other media pools. For example, you can create a media pool for each specific media type required by a program, and then create another media pool that contains this collection of media pools. Media pools can contain either media or other media pools, but not both.

A Removable Storage system provides two classes of media pools: system and application.

5.6.2.9 System media pools

These include free media pools, unrecognized media pools, and import media pools. Removable Storage creates one free, one unrecognized, and one import media pool for each media type in your Removable Storage system. The system media pools are used to hold media that are not currently being used by an application. The different media pools are defined as follows:

Unrecognized media pools — These contain blank (new) media and media that Removable Storage Manager does not recognize. You should immediately move a new tape or disc from an unrecognized media pool to a free media pool so that the tape or disc can be used by applications, or remove it from the library. Unrecognized media are automatically deleted from the Removable Storage database when they are ejected from a library.

Import media pools — These contain media that Removable Storage recognizes in the database but that have not been used before in a particular Removable Storage Manager system. For example, media in an import media pool could be media from one office location that are introduced into a Removable Storage Manager system at another office location. You can move media from import media pools to free media pools or application media pools for reuse at the new location.

Free media pools — These contain media that are not currently in use by applications and do not contain useful data. Media in free media pools are available for use by applications. Application media pools can be configured to automatically draw media from free media pools when there are not sufficient media available in a particular application media pool. If this configuration is not implemented, you must manually move media from a free media pool when needed.

5.6.2.10 Application media pools

These are created by data management applications (and by you). Application media pools determine what media can be accessed by which applications. Media in an application media pool are controlled by that application or by an administrator. An application can use more than one media pool, and more than one application can share a single media pool. For example, Tivoli Storage Manager might use one media pool (tape storage pool) for backups, and another media pool for archives each containing a different media type.

There can be any number of application media pools in a Removable Storage Manager system. Media that are currently reserved for use by an application, called allocated media, cannot be moved between media pools. Allocation controls how the media are used by applications.

5.6.2.11 Work queue

Removable Storage Manager enables multiple applications to share a bulk storage device.

RSM implements and manages a *work queue*. All requests are sent to the work queue. The requests may be varied in terms of the nature of work and may come from same or different applications. The work queue lists all library requests, or work items, whether initiated by an application or by Removable Storage.

As an example TSM asks RSM to mount a tape in an ATL DLT library, an HSM application asks RSM to dismount a tape in the same library, and an administrator wants to perform a library inventory — all at the same time. The request to mount a tape in a library results in a mount work item listed in the work queue.

Each work item can exist in one of the following states:

- Completed — Work item has been completed successfully.
- Failed — Removable Storage was unable to complete the work item.
- In Process — Removable Storage is actively operating on the work item.
- Queued — This state represents the time between when the request is issued and when Removable Storage starts examining the request.
- Waiting — One or more resources that are needed to satisfy the request are currently in use.

5.6.2.12 Operator requests

Sometimes, even with robotic libraries, manual assistance is needed to complete a request or perform maintenance or support. If an application requests that a cartridge in the offline media physical location be mounted, the cartridge must be manually entered into an online library, generating a request to the operator to enter the cartridge. The requests are viewed from the MMC snap-in for RSM. An *operator request* is a message that asks you to perform a specific task. Operator requests can be issued by Removable Storage or by a program that is aware of Removable Storage, such as Backup.

Removable Storage can generate an operator request when any of the following occur:

- An application initiates a mount request for a tape or disc that is offline.
- There are no available media online. That is, a program requires media, and none exist in the appropriate application media pool or in a free media pool.
- A library has failed and requires servicing.
- A drive needs cleaning, and no usable cleaner cartridges are available in the library.

You can respond to operator requests by either completing or refusing the operator request. When you refuse an operator request, Removable Storage notifies the application that generated the request. By default, Removable Storage stores an operator request in the work queue for one hour after you respond to the request.

- Completed — You have indicated, or Removable Storage has detected, that the operator request was completed.
- Refused — You have indicated that the operator request will not be performed.
- Submitted — Removable Storage or an application has submitted an operator request for action.

You can change how operator requests are displayed on your computer, delete operator requests, cancel any pending operation listed in the work queue, and change the order of mount requests in the work queue.

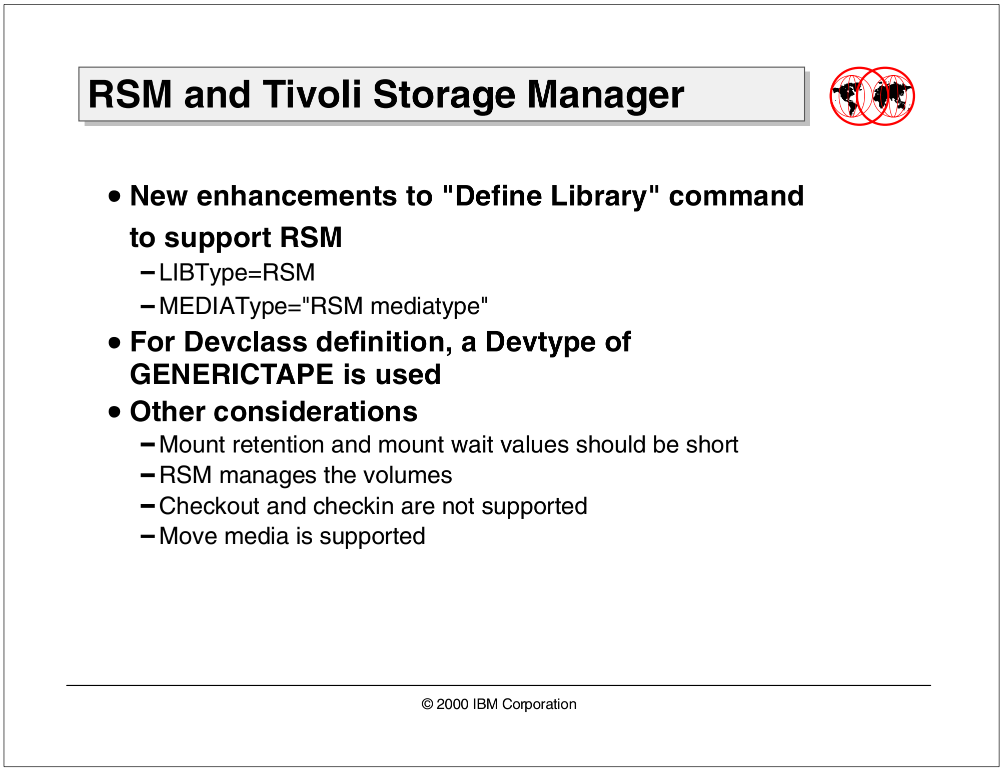

RSM and Tivoli Storage Manager

- **New enhancements to "Define Library" command to support RSM**
  - LIBType=RSM
  - MEDIAType="RSM mediatype"
- **For Devclass definition, a Devtype of GENERICTAPE is used**
- **Other considerations**
  - Mount retention and mount wait values should be short
  - RSM manages the volumes
  - Checkout and checkin are not supported
  - Move media is supported

© 2000 IBM Corporation

5.6.3 TSM setup considerations for RSM support

In order to support the definition of libraries that are being managed by RSM on a Windows 2000 system, a new library type has been created, known as RSM.

A new keyword has been added to the `LIBTYPE` parameter of the TSM `DEFINE LIBRARY` command. When a library is defined to TSM as being managed by RSM (via the `LIBTYPE=RSM` keyword on the `DEFINE LIBRARY` command), then TSM will invoke RSM API's for creating RSM media pools, mounting/dismounting media in the library, ejecting media from the library, and so on.

At this stage, you can only specify `LIBTYPE=RSM` on a `DEFINE LIBRARY` command on a Windows 2000 system.

A new parameter `MEDIATYPE` has also been added to the `DEFINE LIBRARY` command. It specifies the Windows NT media type. The Windows NT System Monitor displays these under media pools.

Consider an example where a standalone 4mm tape drive under the control of RSM is to be used with TSM. The Device Configuration Wizard under TSM Server Utilities would issue the following sequence of commands:

```
define library hp4mm libtype=RSM mediatype="4mm DDS”
```

The `library` name is `hp4mm`. This is a unique name within TSM.

The `LIBtype` is `RSM`, indicating that RSM is being used for all media services.

The `mediatype` is `"4mm DDS"`. Note that the media types contain embedded blanks, thus making it necessary to surround them with quotation marks. At the present time, the media type is restricted to tapes, because the RSM system assumes that opticals contain a file system.

Note

Libraries with a LIBtype of RSM can be deleted but not updated. Changing the media type of an existing library is not allowed.

The device class associated with this library must be GENERICTAPE as shown in the following `DEFINE DEVCLASS` command:

```
define devclass mydevclass devtype=generictape library=hp4mm format=drive
mountretention=5 mountwait=10 mountlimit=1
```

The mountlimit for the define devclass command should be equal to the number of drives in the physical library.

The RSM Physical Volume Repository (PVR) is much like the `EXTERNAL` PVR's in that it does not manage the volumes as libvols but depends on RSM to manage the volumes. The functions of Checkout and Checkin are not supported since RSM handles media eject and inject. Media can be moved between different application or system pools. Therefore, the commands Move Media, Move Drmedia, Query Drmedia, Query Media, and Mount with a location are supported.

Next the storage pool hierarchies are defined. RSM references these storage pools as application storage pools. These are described in more detail in the next section. The commands listed below will be part of a site specific configuration of the storage pool hierarchy.

```
define stgpool genpool1 genclass1 maxscratch=500
update stgpool backuppool next=genpool1
```

All the commands mentioned above can be executed both from the command line or using the Device Configuration Wizard from Server Utilities. The Wizard has been updated to include RSM support.

Note

It is strongly recommended that administrators utilize an RSM library only for new devices introduced to the system. Direct media migration from an existing TSM library to a new RSM library is only possible if the TSM library drives were previously defined with the `GENERICTAPE` type. For other formats, please see the `Move Data` command in *Tivoli Storage Manager for Windows Administrator's Reference*,GC35-0381 for further details.

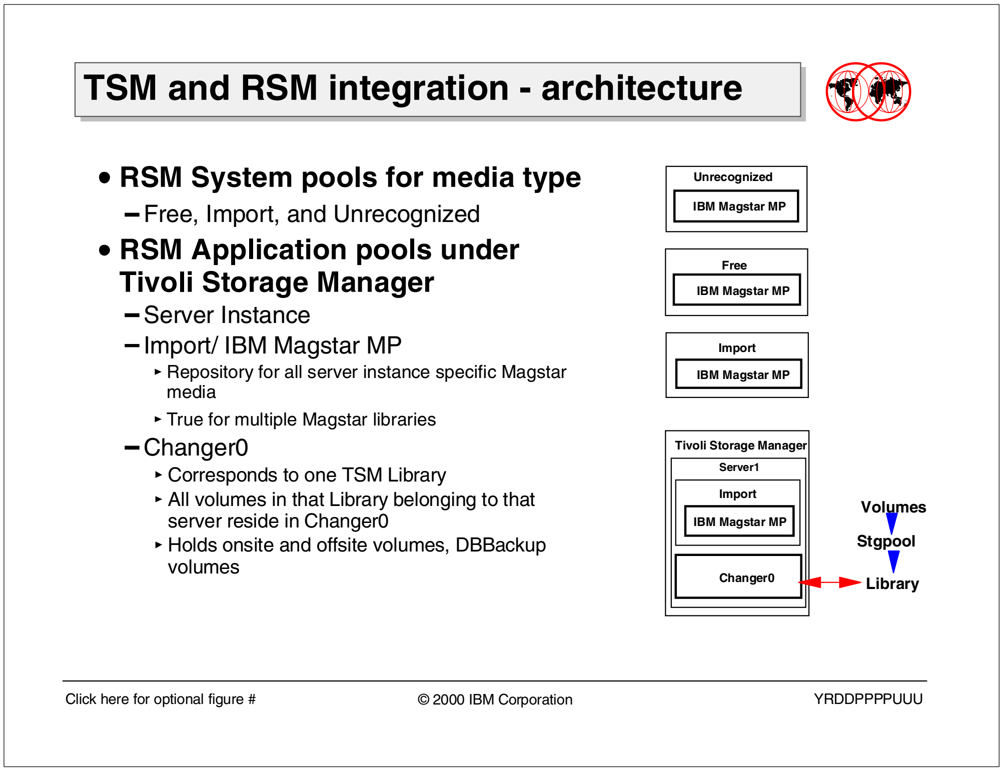

TSM and RSM integration - architecture

- **RSM System pools for media type**
  - Free, Import, and Unrecognized
- **RSM Application pools under Tivoli Storage Manager**
  - Server Instance
  - Import/ IBM Magstar MP
    - Repository for all server instance specific Magstar media
    - True for multiple Magstar libraries
  - Changer0
    - Corresponds to one TSM Library
    - All volumes in that Library belonging to that server reside in Changer0
    - Holds onsite and offsite volumes, DBBackup volumes

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

5.6.4 TSM and RSM integration — architecture

TSM integrates completely with RSM. Once a library is set up to work under RSM, specific pools for the media type are created by RSM under the Free, Import, and Unrecognized pools. As soon as TSM is set up, an RSM application pool named Tivoli Storage Manager is created which contains an Import pool and a Changer pool.

The *Import pool* is a holding directory for specific media type pools. In this case the Import pool contains a specific pool for holding Magstar media named IBM Magstar MP. If there was another Magstar library attached to the same system its contents would also end up in the same media specific instance Import pool. If the TSM server was also using an 8mm tape library, a separate 8mm instance specific Import pool (8mm MP) would be created. An instance specific media type Import pool can be described as a repository for all media of a single media type in one or more libraries that is under the use of one specific TSM server instance.

The *Changer pool* corresponds to a specific TSM library. All volumes in a specific Library that belong to that particular server reside in the Changer pool. These volumes include both on site and off site volumes as well as DBBackup volumes. If the server was using another library whether using same or different type of media, there would be a corresponding Changer1 pool under the server instance pool. Volumes are moved to Changer pools either from the Free pool in which case they are TSM scratch volumes or from the instance Import pool in which case they are TSM private volumes. Volumes will stay in the Changer pool until they are expired and TSM no longer needs them. Then, the expired volumes automatically go from Changer pool to the Free pool for the media type.

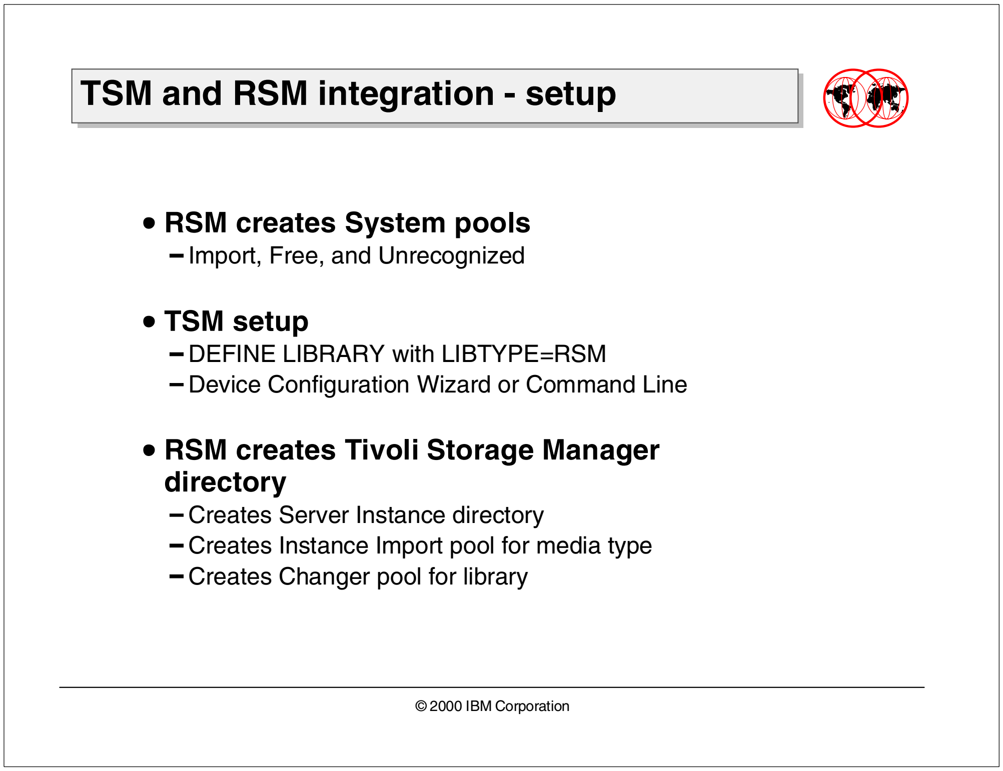

TSM and RSM integration - setup

- **RSM creates System pools**
  - Import, Free, and Unrecognized
- **TSM setup**
  - DEFINE LIBRARY with LIBTYPE=RSM
  - Device Configuration Wizard or Command Line
- **RSM creates Tivoli Storage Manager directory**
  - Creates Server Instance directory
  - Creates Instance Import pool for media type
  - Creates Changer pool for library

© 2000 IBM Corporation

5.6.5 TSM and RSM integration — setup

Administrators set up RSM to manage tapes when they configure a library under the control of RSM. In our example, we have an IBM Magstar 3570 library that is directly attached to a Windows 2000 system. To define a Magstar 3570 library under Windows 2000, you need to go to `ftp://index.storsys.ibm.com` to get the latest drivers. The procedure for installation is documented in the `readme` file that comes with the driver. Once you have the driver set up and configured and you boot the Windows 2000 system, RSM configures the device and creates media pools for the Magstar media type under each of its system pools. These media pools are repositories for holding all Magstar type media.

The first time the `DEFINE LIBRARY` command is issued with `LIBTYPE=RSM` (either using Device Configuration Wizard from Server Utilities or using command line), RSM creates a top level pool named Tivoli Storage Manager. This is an anchor directory for holding other pools, and does not contain any media itself. Under the Tivoli Storage Manager pool, a server instance pool is created. For a system with multiple TSM server instances running there would be multiple high level server instance pools under Tivoli Storage Manager pool.

Under the “Server1” pool, RSM creates two other pools. One is the Import pool and the other is the Changer pool. The instance Import pool categorizes media of specific types. In our case, under the instance Import pool a pool is created which represents all Magstar media. The Changer pool represents an autochanger. (Please see previous section for descriptions). Another Magstar library, therefore, would end up as Changer1 under RSM.

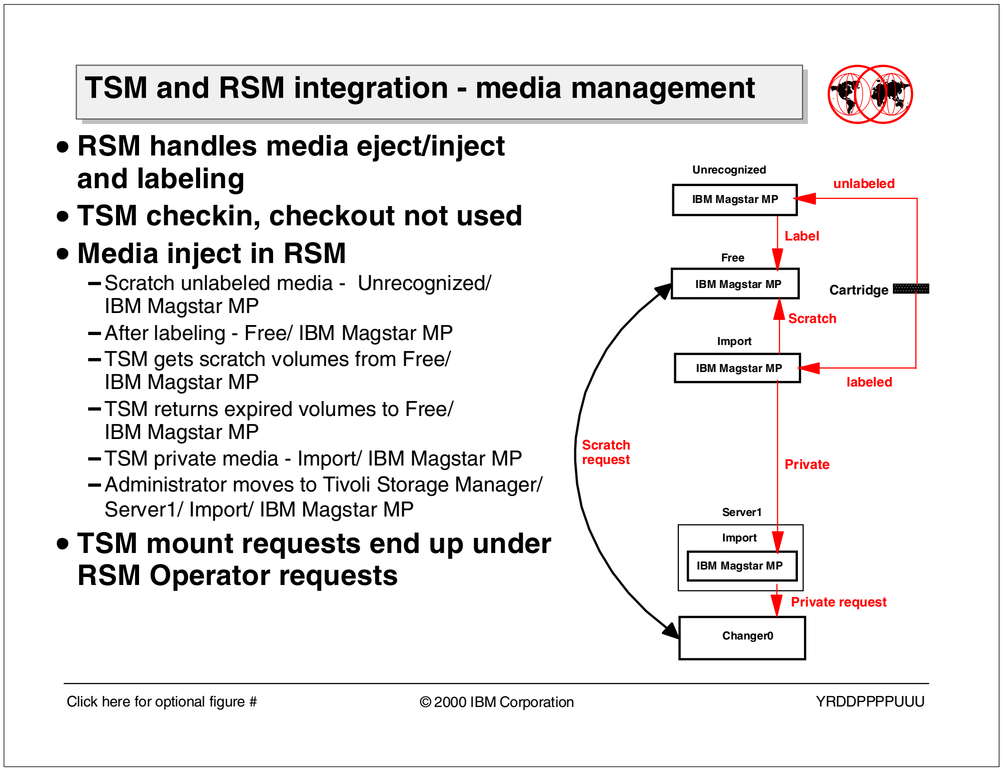

TSM and RSM integration - media management

- RSM handles media eject/inject and labeling
- TSM checkin, checkout not used
- Media inject in RSM
  - Scratch unlabeled media - Unrecognized/ IBM Magstar MP
  - After labeling - Free/ IBM Magstar MP
  - TSM gets scratch volumes from Free/ IBM Magstar MP
  - TSM returns expired volumes to Free/ IBM Magstar MP
  - TSM private media - Import/ IBM Magstar MP
  - Administrator moves to Tivoli Storage Manager/ Server1/ Import/ IBM Magstar MP
- TSM mount requests end up under RSM Operator requests

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

5.6.6 TSM and RSM integration — media management

RSM handles all media inject and eject functions in the library, as well as the labeling of that media. The TSM commands, checkin and checkout, are not used.

When scratch media is injected into an RSM system for the first time, it is sent to the media type Unrecognized pool. From this pool the RSM administrator moves it to the Free pool for the media type where RSM labels it. If TSM needs a scratch volume it can request it and automatically get it from Free pool for the specific media type. Any expired volumes are sent back from the TSM Changer pool to the Free pool for the media type. The volumes in Free pool are accessible to all applications as valid scratch volumes.

When media is injected into an RSM system that has already been used and labeled with TSM, it is sent to the Import pool for the media type. From there the RSM administrator must identify the volumes belonging to TSM and manually move them (drag and drop) from this pool to the Tivoli Storage Manager instance specific Import pool. This operation should not be delayed, since RSM does not prevent other applications from requesting an overwrite of a volume that TSM is using.

When TSM requests a volume, the volume is sent from the instance Import pool for the media type to the Changer pool. The volume stays in the Changer pool until expired after which it is sent back to the Free pool for the media type.

Another point to note is that any media mount requests made by TSM appear in the Operator Requests folder for RSM.

5.7 New Windows 2000 device configuration

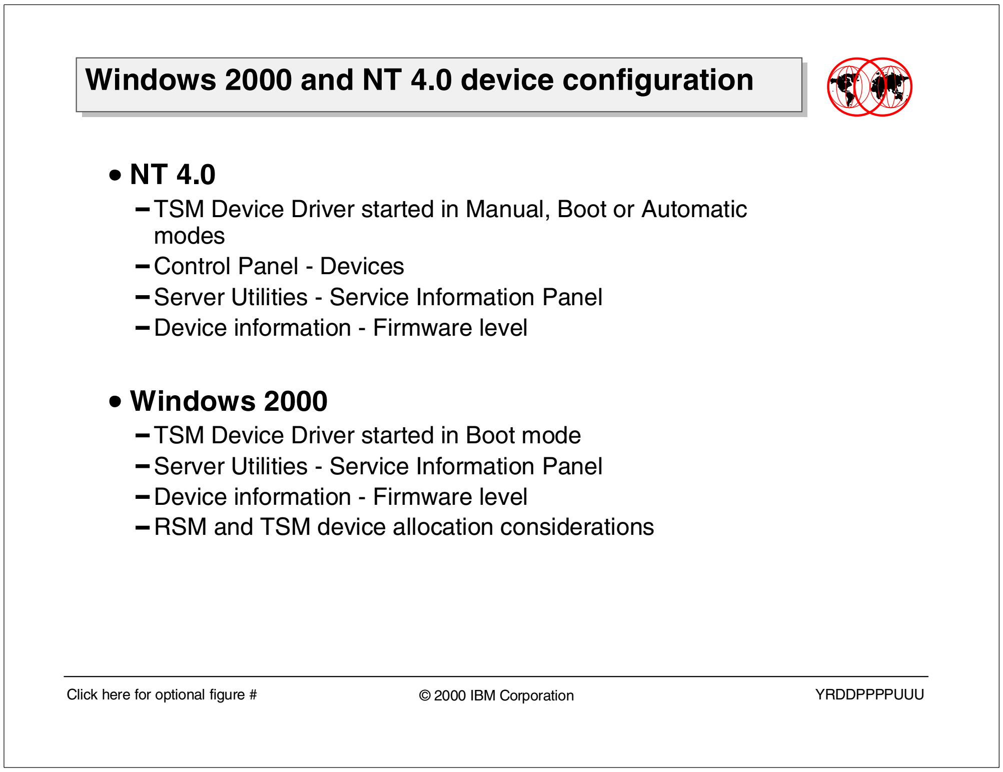

Windows 2000 and NT 4.0 device configuration

- **NT 4.0**
  - TSM Device Driver started in Manual, Boot or Automatic modes
  - Control Panel - Devices
  - Server Utilities - Service Information Panel
  - Device information - Firmware level
- **Windows 2000**
  - TSM Device Driver started in Boot mode
  - Server Utilities - Service Information Panel
  - Device information - Firmware level
  - RSM and TSM device allocation considerations

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

All the devices that TSM supports for NT 4.0 continue to be supported on Windows 2000. However, the way a device is configured is different.

5.7.1 Windows NT 4.0 device configuration

When you are running on a Windows NT 4.0 system, you configure devices in one of the three ways:

- Manual startup option: This is the default. TSM Device Driver can be started from Control Panel — Devices, or from the TSM Server Utilities Service Information panel. Right-click on TSM device driver line gives the option to enable ADSMSCSI.
- Automatic: If the startup option is changed from Manual to Automatic, the system automatically starts TSM device driver. The startup option can be changed from TSM Service Utilities Service Information by right-clicking on TSM device driver.
- Boot option: From TSM Server Utilities Service Information, right-click on TSM device driver. Selecting Enable Optical Device Support box under Options enables the boot option.

Reboot is required after changing the option from Manual start to Automatic or Boot start.

5.7.2 Windows 2000 device configuration

In order to get a device configured for TSM in Windows 2000, it is necessary that the TSM device driver be started with the boot option.

In NT 4.0 you can also start ADSMSCSI from Control Panel — Devices icon. This option is not available on Windows 2000 anymore. You can enable ADSMSCSI and start it only from the TSM Server Utilities in the Windows 2000 environment.

Open the TSM Server Utilities Service Information panel, right-click on the TSM Device Driver and select Options, as shown in Figure 44.

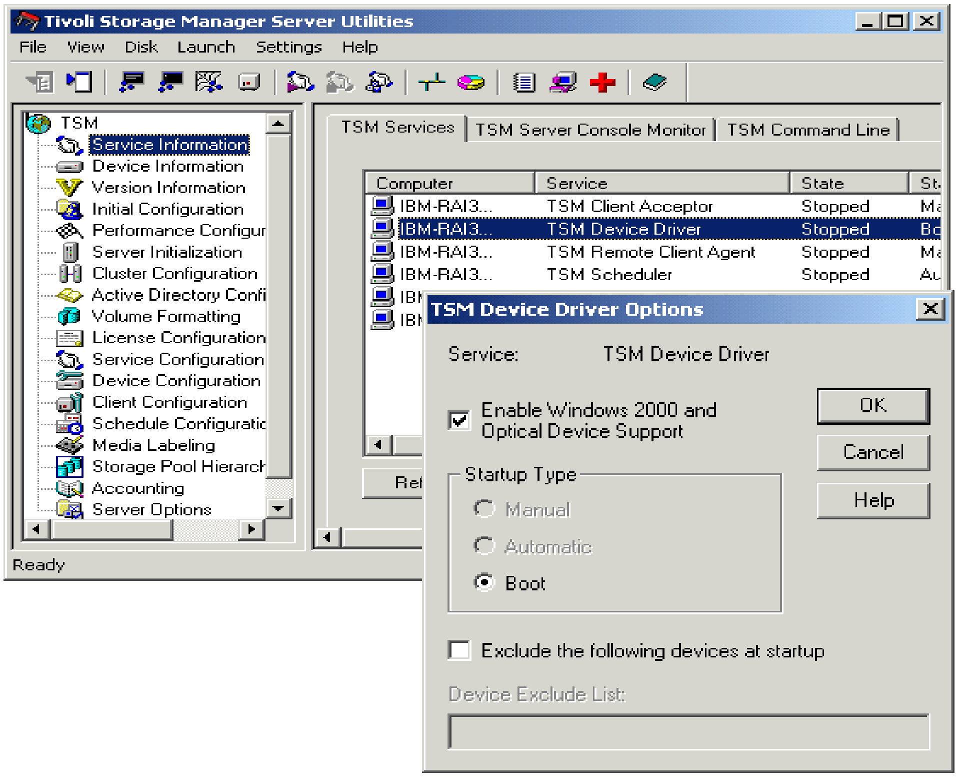

Computer	Service	State	St
IBM-RAI3...	TSM Client Acceptor	Stopped	Ma
IBM-RAI3...	TSM Device Driver	Stopped	Bo
IBM-RAI3...	TSM Remote Client Agent	Stopped	Ma
IBM-RAI3...	TSM Scheduler	Stopped	Au
IBN			
IBN			

TSM Device Driver Options

Service: TSM Device Driver

☑ Enable Windows 2000 and Optical Device Support

Startup Type

○ Manual

○ Automatic

● Boot

□ Exclude the following devices at startup

Device Exclude List:

Figure 44. TSM Device Driver options

Check the box for Enable Windows 2000 and Optical Device Support. This immediately tells the server that the it is being run under Windows 2000. The startup option is changed to boot; and the remaining options (manual and automatic) are excluded from selection.

The reason for this is that Windows 2000 implements a plug-and-play architecture. Tivoli Storage Manager has to take control of the devices at boot time; otherwise, they come under the control of Windows 2000. If you want to reserve any device for other applications, check the box "Exclude the following devices at startup" and specify the devices in Device Exclude List space. The naming conventions used in the exclude list include 4 digits to reference each device, for example, 1.0.0.2.

A reboot is required. The expected behavior after reboot is that you will get all the TSM device names in Device Information in TSM Server Utilities.

5.7.3 Device firmware level information

The device information now also includes firmware levels as shown in Figure 45. This is supported for both NT 4.0 and Windows 2000.

When a drive is recognized by the operating system, Device Manager will list it under Tape Drives or under Disk Drives for optical drivers. However if there is no Windows 2000 hardware driver installed in the system for the attached drives, they don't get recognized. In that case the attached drives will be listed under Other Devices with a yellow question mark as an unknown device.

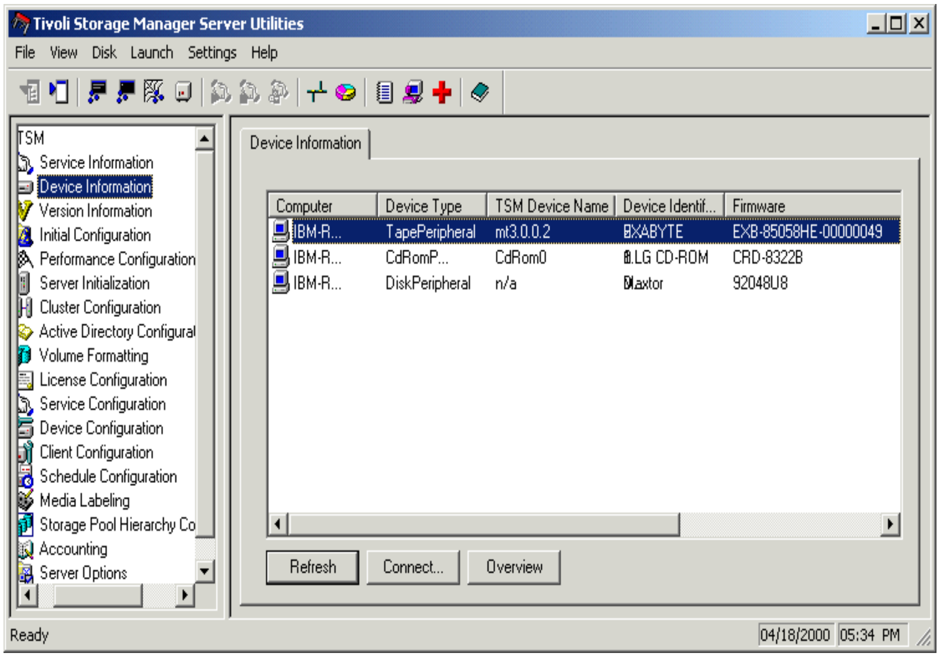

Computer	Device Type	TSM Device Name	Device Identif...	Firmware
IBM-R...	TapePeripheral	mt3.0.0.2	EXABYTE	EXB-85058HE-00000049
IBM-R...	CdRomP...	CdRom0	LG CD-ROM	CRD-8322B
IBM-R...	DiskPeripheral	n/a	Maxtor	92048U8

Figure 45. Server Utilities — Device Information

The tape drives configured by TSM and bound to the adsmscsi.sys driver are identified by TSM device names. Device Manager lists the tapes under Tape Drives or Other Devices.

The tape drives configured by Windows 2000 and bound to native drivers are identified by NT device names. Device Manager always lists the tapes under Tape Drives.

5.7.4 TSM and RSM device allocation issues

RSM runs on Windows 2000 and only uses native device drivers. It does not use ADSMSCSI.

By default, after installation, the TSM device driver (ADSMSCSI) on Windows 2000 is set so it acquires all the devices it can at system startup. In order to configure RSM, it is necessary to either disable ADSMSCSI’s Windows 2000 support or add the RSM devices to ADSMSCSI’s exclude list.

A reboot is required after this step. Once the system comes up, RSM can be set up using Device Configuration Wizard in the server utilities.

If the exclude option is not specified, then TSM gets all the devices.

5.8 Microsoft Windows Installer support

Microsoft Windows Installer

- **Installation program with software management capabilities**
  - Components are Windows Installer Service and Installation Package Files
- **Integrated part of Windows 2000, supported via service pack for Windows 9X and NT**
- **Supported on Tivoli Storage Manager V3.7.3 & V4.1**
- **Features**
  - Application advertisement
  - Efficient installation tracking
  - On-demand installation
  - Clean uninstall
  - Upgrade detection
  - Installshield wizard

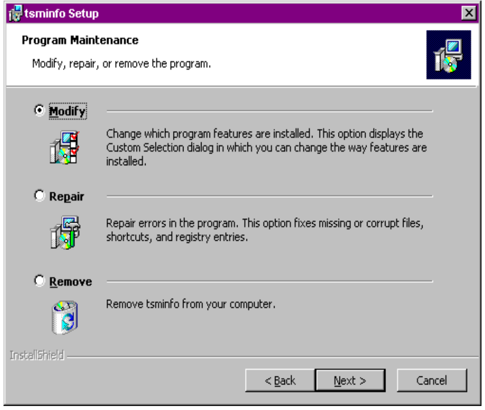

tsminfo Setup

Program Maintenance

Modify, repair, or remove the program.

● **Modify**

Change which program features are installed. This option displays the Custom Selection dialog in which you can change the way features are installed.

○ **Repair**

Repair errors in the program. This option fixes missing or corrupt files, shortcuts, and registry entries.

○ **Remove**

Remove tsminfo from your computer.

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

Microsoft Windows Installer is a component of the Windows 2000 operating system and that simplifies the application installation process.

Windows Installer manages the installation and removal of applications by applying a set of centrally defined setup rules during the installation process. These setup rules define the installation and configuration of the installed application. In addition, this service is used to modify, repair, or remove an existing application.

Windows Installer is not only an installation program; it is also an extensible software management system. Windows Installer manages the installation, addition, and deletion of software components, monitors file resiliency, and maintains basic disaster recovery by way of rollbacks. Additionally, Windows Installer supports installing and running software from multiple sources, and can be customized to install custom applications.

5.8.1 Technology

The Windows Installer technology consists of the Windows Installer service for the Windows operating systems and the package (.msi) file format used to hold information regarding the application setup and installations.

Windows Installer uses the information contained within a package file to install the application.

5.8.1.1 Installer service

Windows Installer service is an operating system service that allows the operating system to manage the installation process. The installer program (Msiexec.exe) program is a component of the Windows Installer. This program uses a dynamic link library, Msi.dll, to read the package files (.msi), apply transforms (.mst), and incorporate command-line options. The installer performs all installation-related tasks.

When Windows Installer is installed on a computer, the file association capabilities of the operating system are modified to recognize the.msi file type. When a file with the.msi extension is double-clicked, the operating system associates the msi file with Windows Installer and runs the Msiexec.exe application.

5.8.1.2 Installation package file

Each package (.msi) file contains multiple interrelated tables that compose a relational type database that stores all the instructions and data required to install (and uninstall) the program across many installation scenarios. For example, a package file could contain instructions for installing an application when a prior version of the application is already installed. The package file could also contain instructions for installing the software on a computer where that application has never been present.

The Windows Installer database tables in the package file reflect the general layout of the entire group of applications, including the available features, components, the relationships between features and components and necessary registry settings. Because the database is relational, changes made to one table are propagated automatically throughout the database. This is a very efficient process for introducing consistent changes into the installation process that simplifies customizing a large application or group of applications.

Another component associated with the package files is the Windows Installer transform (.mst) file. The installation process can be manipulated by applying transforms to the installation database. A transform makes changes to elements of the database. The transform files modify the installation package file at installation time, and can therefore dynamically affect the installation behavior. For example, Windows Installer can use a transform file to change the language in the user interface of an application.

5.8.2 Features

With Windows Installer and the .msi package file format, software installation, and removal has become more reliable and resilient while providing a larger set of installation options. Windows Installer has the ability to perform the following tasks:

Restores original computer state upon installation failure — Windows Installer keeps track of all changes made to the system during the application installation process. If the installation fails, Windows Installer can restore, or roll back, the system to its initial state. This includes restoration of deleted or overwritten files, registry keys, and other resources.

Helps prevent certain forms of inter-application conflicts — Windows Installer enforces installation rules that help to prevent conflicts with shared resources between existing applications. Such conflicts can be caused when an install operation makes updates to a dynamic link library (.dll) shared by an existing application, or when an operation deletes a dynamic link library shared by another application.

Reliably removes existing programs — Windows Installer can reliably uninstall any program it previously installed. Since Windows Installer manages applications at the component level and oversees resources other than files, it ensures that no resource is installed or removed unless the component that owns it is either installed or removed. It removes all the associated registry entries and application files, except for those shared by other installed software. You can uninstall an application at any time after a successful installation. (Removal should not be confused with rollback, which restores a computer to its initial state when an installation failure has occurred).

Diagnoses and repairs corrupted applications — An application can query Windows Installer to determine whether an installed application has missing or corrupted files. If any are detected, Windows Installer repairs the application by recopying only those files found to be missing or corrupted.

Supports on-demand installation of application features — Windows Installer can be instructed to initially install a minimal subset of an application. Later, additional components can be automatically installed the first time the user accesses features that require those components. This is known as advertising. Similarly, Windows Installer can also be configured to purge components that go unused in an application.

Supports unattended application installation — Installation packages can be configured to require no installation process interaction from the user. During the installation process, Windows Installer can query the computer for desktop attributes, including determining whether any applications were previously installed by Windows Installer.

5.8.3 Tivoli Storage Manager exploitation

TSM's integration with the Windows Installer service provides better overall efficiency for the install program.

Installshield — TSM uses the Installshield software for system installation. It is an entirely new release that has been designed to work with Microsoft Windows Installer.

Wide variety of installable services — The following services can be installed using the Microsoft Windows Installer.

- TSM Server
- TSM Device Driver
- TSM Scheduler
- TSM Client Acceptor
- TSM Remote Client Agent

The Installer Service creates and exports DLL functions to perform TSM server initialization and configuration tasks, such as these:

- The `dsmserv` runfile task
- The `dsmserv` initialize task
- The `dsmfmt` task
- The service install task

Clean uninstall — Because Windows Installer service can accurately track what a given component has installed and when that component can be removed, applications like TSM installed using the Windows Installer service can be uninstalled much more cleanly. No resource is installed or removed unless the component that owns it is either installed or removed.

Upgrade detection — In case the Tivoli Storage Manager server is being upgraded to a newer version, the installation program ensures that the existing location is the default location and the existing program folder is the default program folder for the new server code. In addition, the program would look at server configuration information and ensure that the existing database and log sets are reused and new database and log volumes are not generated. Lastly, a lookup on the version info would tell the server whether an upgrade DB is required.

Application advertisement and on-demand install — Applications that use the Windows Installer service management API do not require setup to run again if it is decided later that additional functionality is required. If, while using TSM, the administrator requests functionality that was not previously installed, TSM can call the Windows Installer service to install the necessary feature(s) on the administrator's behalf without the need to re-run setup. In essence, all features of TSM are available even if they are not installed. This notion of availability in the absence of the installed files is known as *advertisement*. It is shown in Figure 46.

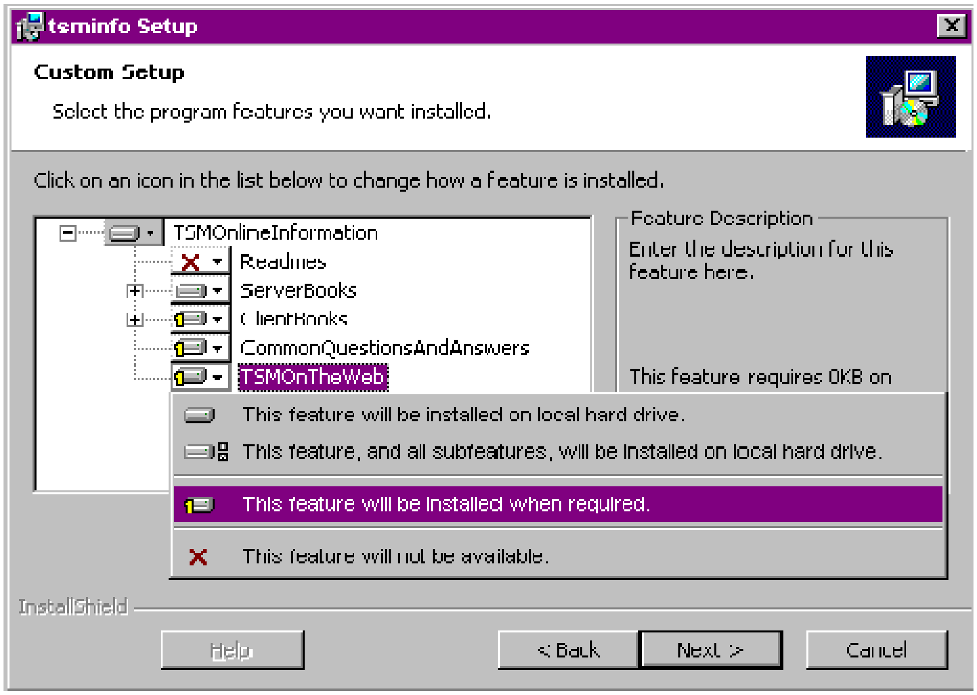

tsminfo Setup

Custom Setup

Select the program features you want installed.

Click on an icon in the list below to change how a feature is installed.

Feature Description

Enter the description for this feature here.

This feature requires 0KB on

This feature will be installed on local hard drive.

This feature, and all subfeatures, will be installed on local hard drive.

This feature will be installed when required.

This feature will not be available.

Figure 46. On-Demand install

Installation Tracking — Tivoli Storage Manager is described in the form of a single Windows Installer package file. The package file is a database format that has been optimized for installation performance, and describes, among other things, the relationships between features, components, and resources for a given product. This makes it easier for the TSM administrator to accurately track all the installed components.

5.9 Performance Tuning Wizard enhancements

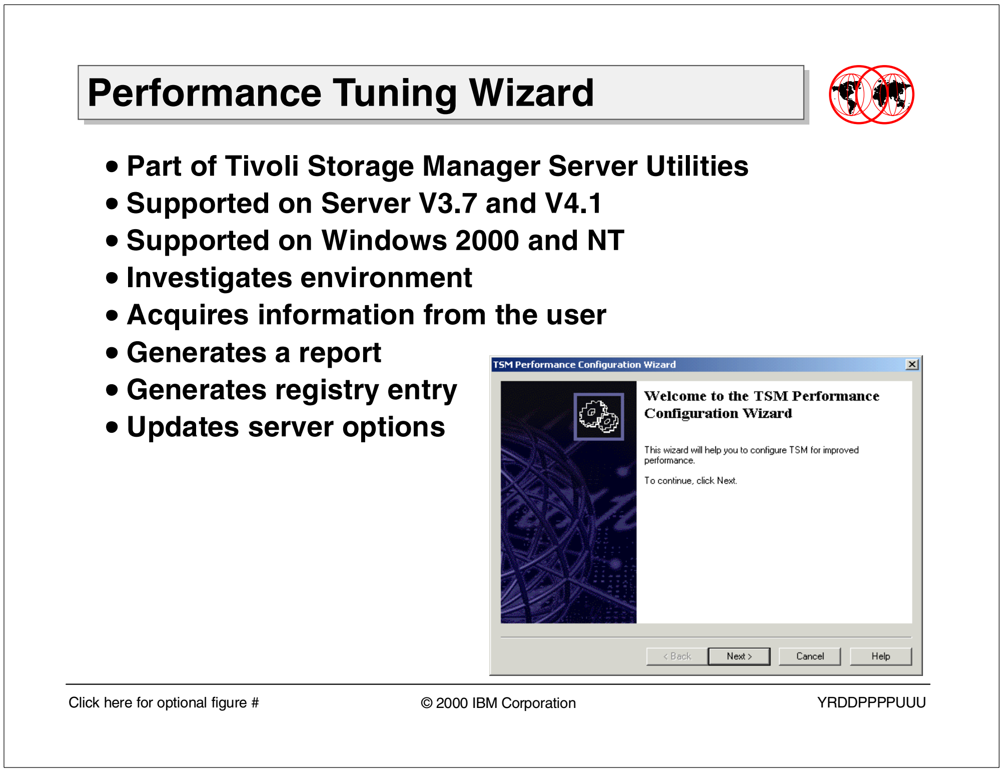

Performance Tuning Wizard

- Part of Tivoli Storage Manager Server Utilities
- Supported on Server V3.7 and V4.1
- Supported on Windows 2000 and NT
- Investigates environment
- Acquires information from the user
- Generates a report
- Generates registry entry
- Updates server options

TSM Performance Configuration Wizard

Welcome to the TSM Performance Configuration Wizard

This wizard will help you to configure TSM for improved performance.

To continue, click Next.

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

The TSM server utilities on Windows 2000 and NT include a performance tuning wizard under the listing Performance Configuration. If engaged, the wizard acquires information from the user, investigates the server environment and provides options for taking certain actions. The information obtained and actions performed are all server instance specific.

As part of acquiring information from the user, the wizard inquires about the number of supported clients and the typical size of stored files (Figure 47).

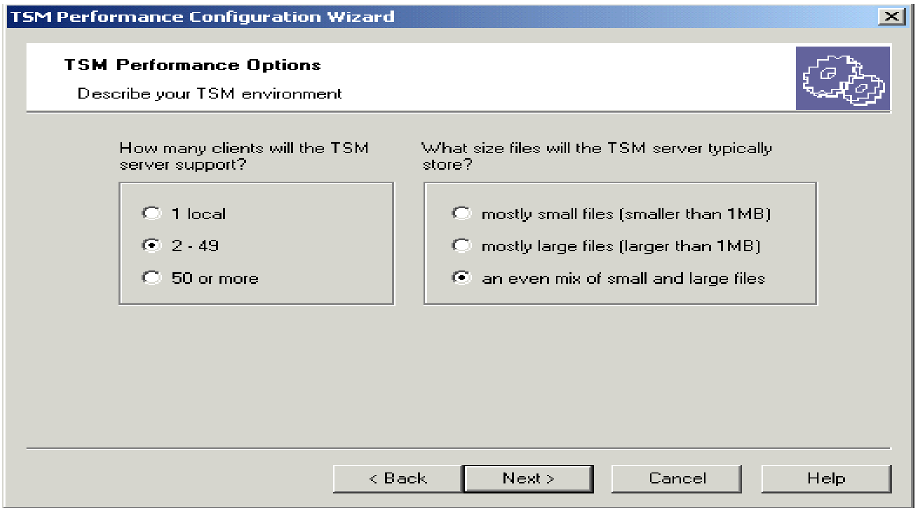

TSM Performance Configuration Wizard

TSM Performance Options

Describe your TSM environment

How many clients will the TSM server support?

○ 1 local
● 2 - 49
○ 50 or more

What size files will the TSM server typically store?

○ mostly small files (smaller than 1MB)
○ mostly large files (larger than 1MB)
● an even mix of small and large files

Figure 47. Acquiring system information from the user

Next, the wizard starts investigating the environment. Part of this is measuring disk performance.

The wizard runs a set of sequential and random access tests on local hard drives. These tests model the way TSM uses the drives for its disk, database and recovery log volumes. TSM uses the results of these tests to determine the best location for each of these type of volumes (disk, db, recovery log). It ranks the drives and allocates space in that order as shown in Figure 48.

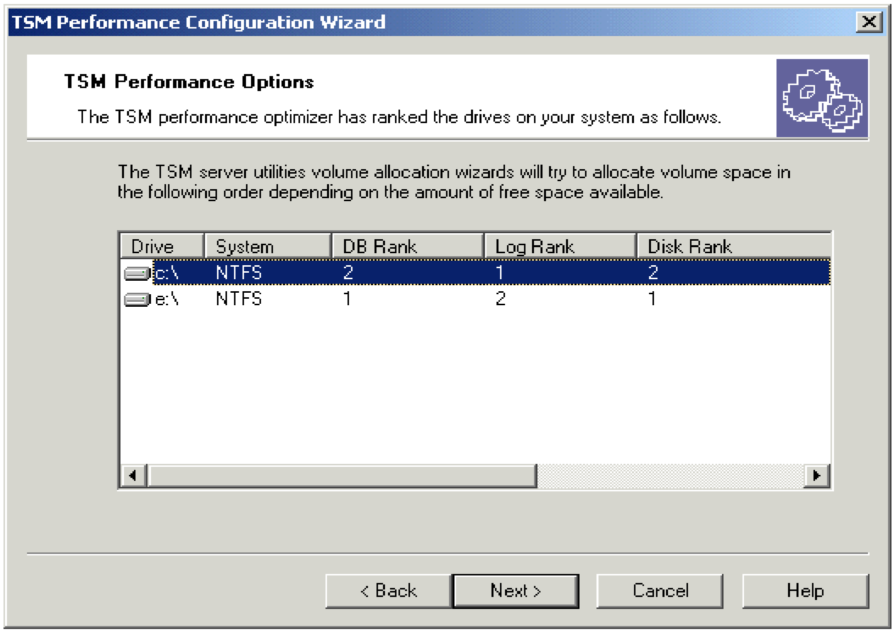

TSM Performance Configuration Wizard

TSM Performance Options

The TSM performance optimizer has ranked the drives on your system as follows.

The TSM server utilities volume allocation wizards will try to allocate volume space in the following order depending on the amount of free space available.

Drive	System	DB Rank	Log Rank	Disk Rank
c:\	NTFS	2	1	2
e:\	NTFS	1	2	1

Figure 48. Ranking available drives for space allocation

The performance wizard generates three pieces of output. The user can select either, all or none of the options that generate the output as shown in Figure 49 on page 192.

- A report (located in `C:\Program Files\Tivoli\TSM\Utils\adsmperf.log`) that provides recommendations on client options, server options, and volume locations.
- A registry entry that stores preferred drives (previously ranked) for TSM volumes.

  The utilities will refer to this entry, if it exists, when deciding on default locations for new volumes.

  When a new volume is created, the utilities will scan the list in order and use the first drive that has enough space for the volume.

  The performance wizard is a step in the initial configuration wizard and will help to ensure volumes are allocated on the right volumes from the start.

  It would be possible for the `dsmfmt` utility to use the information to help decide where to find space to extend volumes as well.
- The performance wizard will make changes to the server options file.

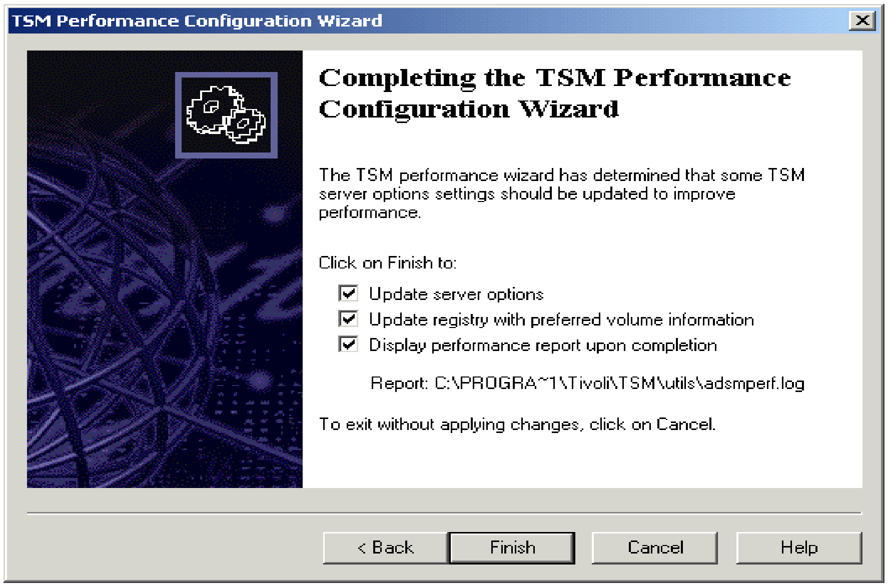

Completing the TSM Performance Configuration Wizard

The TSM performance wizard has determined that some TSM server options settings should be updated to improve performance.

Click on Finish to:

- ☑ Update server options
- ☑ Update registry with preferred volume information
- ☑ Display performance report upon completion

Report: C:\PROGRA~1\Tivoli\TSM\utils\adsmperf.log

To exit without applying changes, click on Cancel.

Figure 49. Output options

A sample report is shown below:

```
Detected 1 processor(s)
Detected 603492352 bytes physical memory
Detected TCP/IP protocol is installed.
Detected NetBIOS over NetBeui protocol is installed.
Found local fixed logical disk c:\
(90.44MB free)
Found local fixed logical disk e:\
(3905.32MB free)

Disk Performance Results (Smaller timing values are better)

      DB Rank        Log Rank       Disk Rank      File
Drive Pos. Raw       Pos. Raw       Pos. Raw       System   Comp. Seq. 4k  Seq. 256k  Rnd. 4k
---------------------------------------------------------------------------------------------
----
c:\    2    1048766  1    1048767   2    1048679   NTFS     Yes   8.8      96.6       9.9
e:\    1    1048768  2    1048767   1    1048710   NTFS     Yes   8.6      65.4       7.3

Timings in ms per I/O request

Drive  SeqW4k    SeqR4k    SeqW256k  SeqR256k  RndW4k    RndR4k
---------------------------------------------------------------------------------------------
----
c:\    8.8       0.0       95.0      1.6       9.8       0.0
e:\    8.6       0.0       63.8      1.6       7.3       0.0

Server is expected to support an even mix of small and large files

Server is expected to support 2-49 clients
```

```
================================================================================
RECOMMENDATIONS BASED ON INFORMATION OBTAINED
================================================================================

Preferred database volume drive list: e:\,c:\ was stored in the registry for future use
Preferred log volume drive list: c:\,e:\ was stored in the registry for future use
Preferred disk volume drive list: e:\,c:\ was stored in the registry for future use


TSM SERVER OPTIONS FILE SETTINGS
txngroupmax     256
uselargebuffers yes
COMMmethod      NAMEDPIPE
NAMEdpipename   \\.\pipe\adsmpipe
TCPWindowsize   63
COMMmethod      TCPIP
TCPPort         1500
TCPNODELAY      Yes
bufpoolsize     73728

Updated server options file c:\progra~1\tivoli\tsm\server1\dsmserv.opt
Changes will take effect once the server is restarted.

TSM SERVER MACHINE SETTINGS
The TSM server performance can be enhanced by adding one or more CPUs to this system.
The TSM server performance may be enhanced by removing any unused protocols on this computer.

LOCAL TSM CLIENT OPTIONS FILE SETTINGS
tcpwindowsize    63
tcpbufsize       32
txnbytelimit     25600
largecommbuffers no
commmethod       namedpipe
compression      no

NETWORK TSM CLIENTS OPTIONS FILE SETTINGS
tcpwindowsize    63
tcpbufsize       32
txnbytelimit     25600
largecommbuffers no
commmethod       tcpip
compression      no

----------------------------------------------------------------------------------------
TSM Server Performance Wizard log file closed
----------------------------------------------------------------------------------------
```

Chapter 6. TSM client Windows 2000 support

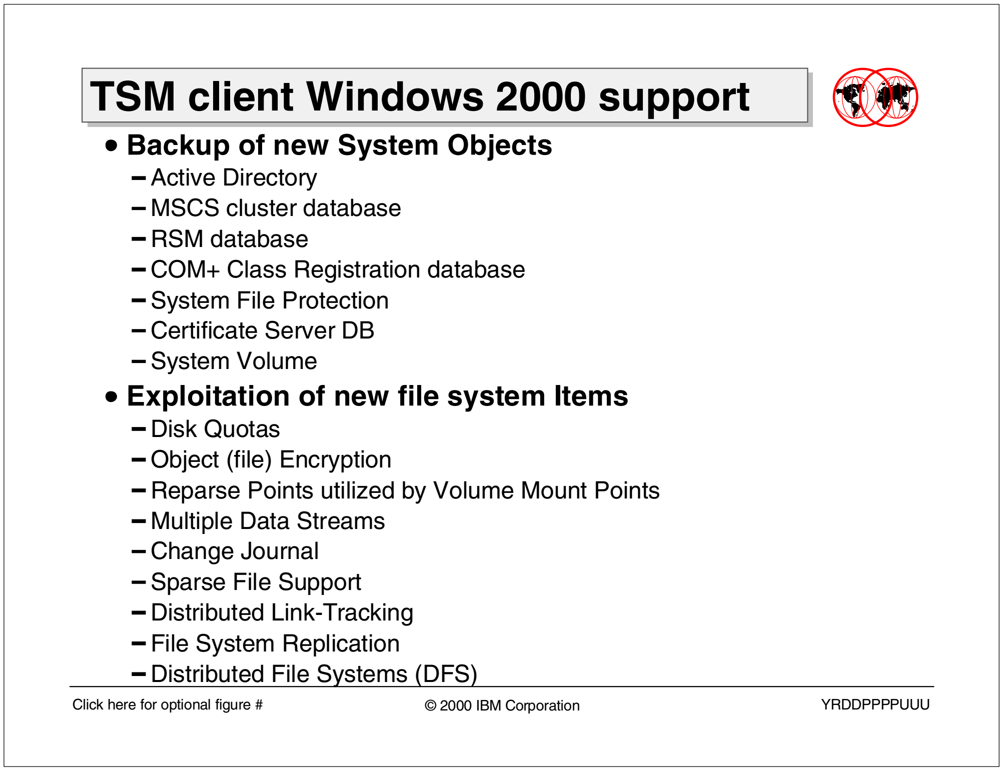

TSM client Windows 2000 support

- **Backup of new System Objects**
  - Active Directory
  - MSCS cluster database
  - RSM database
  - COM+ Class Registration database
  - System File Protection
  - Certificate Server DB
  - System Volume
- **Exploitation of new file system Items**
  - Disk Quotas
  - Object (file) Encryption
  - Reparse Points utilized by Volume Mount Points
  - Multiple Data Streams
  - Change Journal
  - Sparse File Support
  - Distributed Link-Tracking
  - File System Replication
  - Distributed File Systems (DFS)

Click here for optional figure #

© 2000 IBM Corporation

YRDDPPPPUUU

This section discusses some of the new features within Windows 2000. It gives an overview of how Tivoli Storage Manager client supports:

New System Objects, including:

- Active Directory
- MSCS cluster database
- Removable Storage Manager database
- COM+ Class Registration database
- System File Protection
- Certificate Server database
- System Volume

New file system items, including:

- Disk Quotas
- Object (file) Encryption
- Reparse Points utilized by volume mount points
- Multiple Data Streams
- Change Journal
- Sparse File Support
- Distributed Link-Tracking
- File System Replication
- Distributed File Systems (DFS)

The previous Version 3.7.1 of the Tivoli Storage Manager client on Windows 2000 systems already had the ability to back up and restore the Registry, Event Logs and the Removable Storage Manager database. It also supported the new file system features: Disk Quotas, Multiple Data Streams, and Sparse Files. That is why these features will not be covered in this Chapter. For more information, refer to the TSM product documentation or to *Tivoli Storage Manager V3.7: Technical Guide,* SG24-5477.

6.1 TSM support for new System Objects

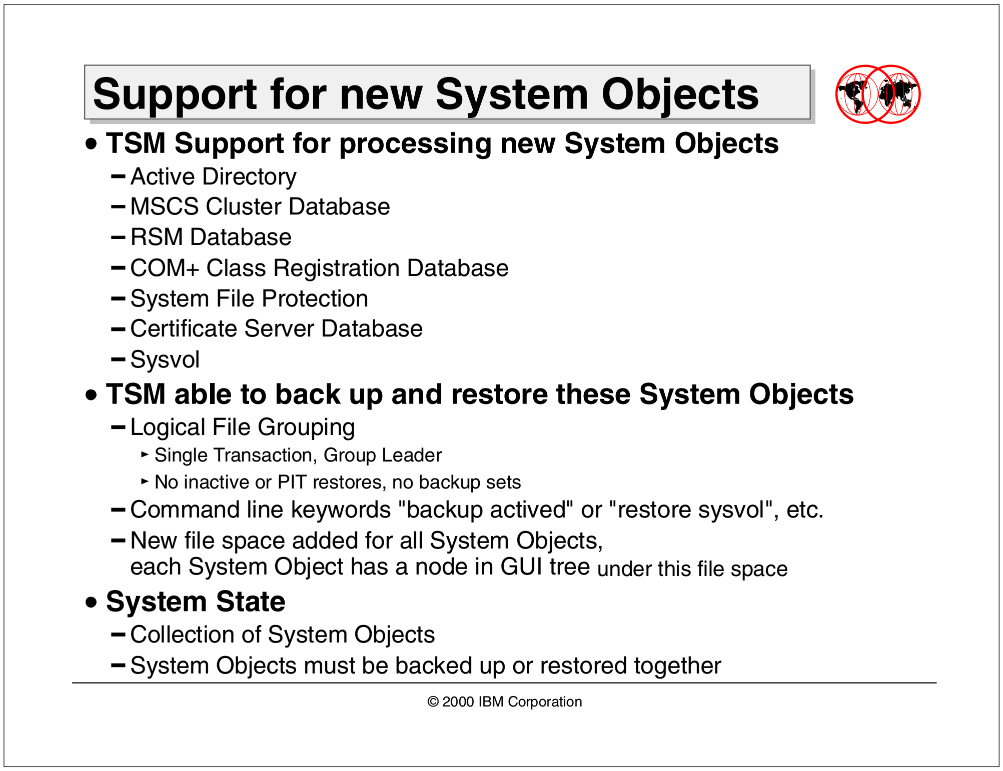

Support for new System Objects

- **TSM Support for processing new System Objects**
  - Active Directory
  - MSCS Cluster Database
  - RSM Database
  - COM+ Class Registration Database
  - System File Protection
  - Certificate Server Database
  - Sysvol
- **TSM able to back up and restore these System Objects**
  - Logical File Grouping
    - Single Transaction, Group Leader
    - No inactive or PIT restores, no backup sets
  - Command line keywords "backup actived" or "restore sysvol", etc.
  - New file space added for all System Objects, each System Object has a node in GUI tree under this file space
- **System State**
  - Collection of System Objects
  - System Objects must be backed up or restored together

© 2000 IBM Corporation

A System Object is a collection of files and/or databases which represent a logical entity and help the system achieve a consistent state. The System Objects can be part of larger, distributed entity known as System State.

6.1.1 TSM Client Support for new System Objects

The previous Version 3.7.1 of Tivoli Storage Manager client on Windows 2000 systems had the already ability to back up and restore the Registry, Event Logs, and the Removable Storage Manager database. For more information, refer to the TSM product documentation or to *Tivoli Storage Manager V3.7: Technical Guide,* SG24-5477. In addition, the new client Versions 3.7.2 and 4.1 support the backup and restore of the other System Objects, including:

- Active Directory directory service
- Cluster service information
- COM+ Class Registration database
- System files (boot files)
- Certificate Services database
- SYSVOL directory

TSM backup/restore commands have been extended to support these new System Objects. The backup/restore operation can be initiated using the command line or the GUI. To support GUI backup/restore, for each client node, a new file space “System Objects” has been added on the server.

For each System Object, there is a node in the GUI tree under this new file space, as shown in Figure 50. The command syntax for all System Objects is the same. Each System Object has a command-line keyword, for example, `backup Sysvol` or `restore actived`.

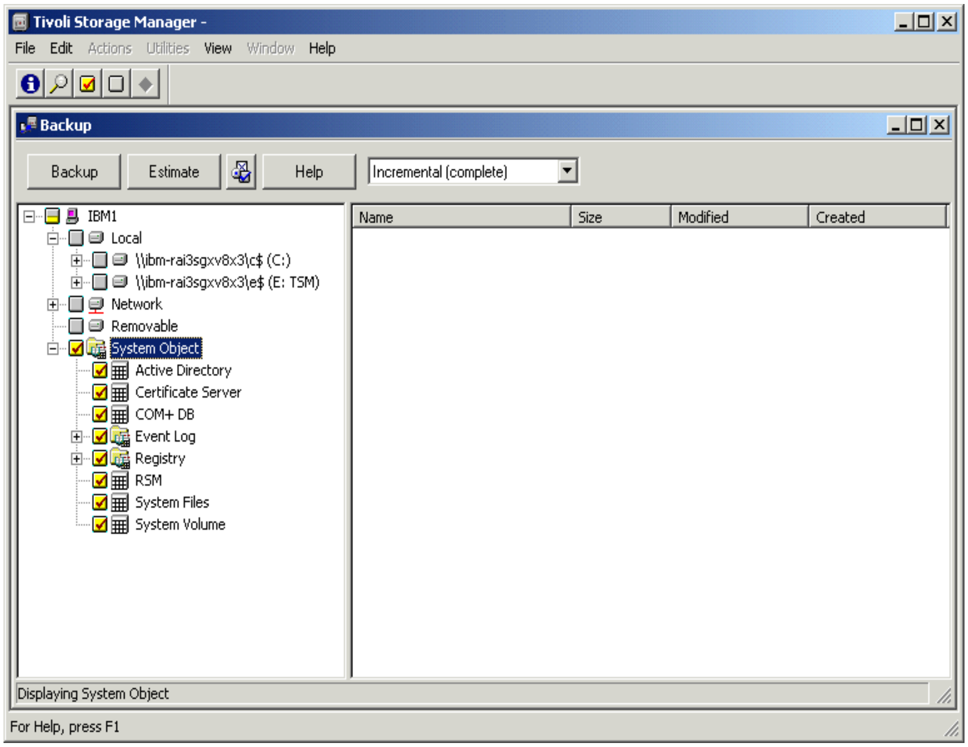

Displaying System Object

For Help, press F1

Figure 50. BA Client GUI Tree — System Objects file space

6.1.2 Logical file grouping

Each System Object can contain several files and/or databases. This can potentially create a situation where a set of files in the System Object changes between backup and consequently extra versions for those set of files are expired off at the server. This would cause an inconsistent restore situation.

Tivoli Storage Manager implements logical file grouping (for more details see Section 3.1, “Logical file grouping server support” on page 82) to take care of all these aspects. The server groups files as a logical entity at client's request. The files are transacted together in a group and are represented by a client designated group leader. During restore of the System Object, the client asks for the name of the group leader. The server replies with a matching Object ID. The client then requests restore of group members based on Object ID of the group leader.

In order to support the backup of System Objects, the traditional transaction limits are overridden for logical file groups. The transactions can be thousands of files depending on the size of the System Object.

There are some limitations as well. Only classic restores are supported for System Objects. Point-in-time or inactive restores are not supported. Also, the System Objects cannot be included as part of backup sets.

6.1.3 System State

System State consists of components that must be processed together to achieve a consistent state of the system. The components are System Objects. The System State consists of the following System Objects:

- System File Protection files
- Active Directory (domain controller)
- SysVol (domain controller)
- Certificate Server
- Cluster database (clusters only)
- Registry
- COM+ database

You can only back up and restore the System State data on a local computer. When you choose to back up or restore the System State data, all of the System State data that is relevant to your computer is backed up or restored; you cannot choose to back up or restore individual components. This is due to dependencies among the System State components. The System State can be backed up in any order, but for restores, it is recommended to restore boot files first and commit the system hive of the registry last.

Note

Certain Windows NT and Windows 2000 System Objects must be backed up or restored together in order to create a consistent System State.

The System Objects include:

- Active directory (domain controller only)
- Certificate server database
- Cluster database (cluster node only)
- COM+ database
- Registry
- System and boot files
- System volume

Other System Objects can be backed up or restored individually. These include the following:

- Event Logs
- Removable Storage Management Database (RSM)
- Replicated File Systems

6.2 Active Directory support

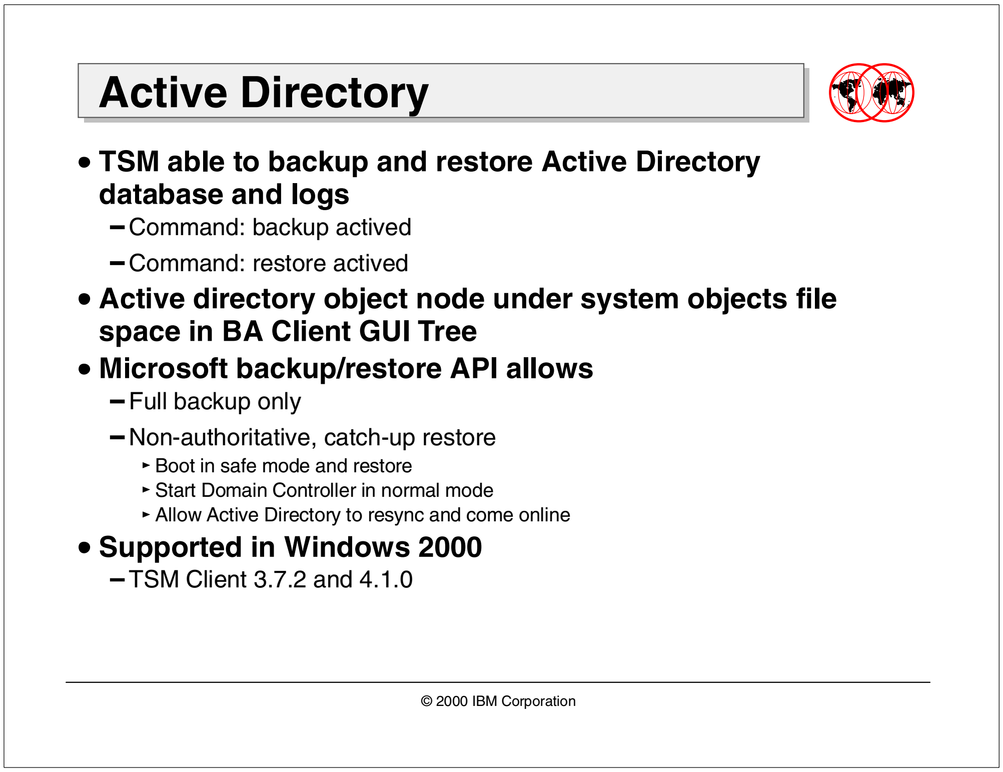

Active Directory

- **TSM able to backup and restore Active Directory database and logs**
  - Command: backup actived
  - Command: restore actived
- **Active directory object node under system objects file space in BA Client GUI Tree**
- **Microsoft backup/restore API allows**
  - Full backup only
  - Non-authoritative, catch-up restore
    - Boot in safe mode and restore
    - Start Domain Controller in normal mode
    - Allow Active Directory to resync and come online
- **Supported in Windows 2000**
  - TSM Client 3.7.2 and 4.1.0

Active Directory is a critical part of the Windows 2000 distributed system environment. It stores information about network System Objects and their attributes and makes it available to administrators, users, and applications. (Please see Section 5.2, “Active Directory exploitation” on page 151 for details).

6.2.1 Backups

The Active Directory is backed up online. Microsoft’s backup/restore API only allows full backups of the Active Directory. TSM supports full backup of the Active Directory, which backs up the database and its associated transaction logs. After the database and logs are backed up, the logs are deleted.

The backup can be initiated either from the GUI or command line. The `Active Directory` node under the `System Objects` file space in the GUI tree for the BA client represents the Active Directory database and logs and can be marked for backups. The command line command is `BACKUP ACTIVED`.

Note:

The `BACKUP ACTIVED` command only backs up the Active Directory database and logs and does not back up any additional files that may be required to properly restore the Active Directory.

For information on System Objects that must be backed up together, please see the Note on page 199.

6.2.2 Restores

The `RESTORE ACTIVED type=catchup online=no` command restores a Windows 2000 Active Directory database and associated logs from a TSM server. Restore can also be initiated from the Active Directory node under the System Objects file space in the GUI tree.

In order to restore the Active Directory, the Windows 2000 Server must have an Active Directory installed, and the Active Directory must be offline. To begin the restore, the system is rebooted in safe mode without the Active Directory. While the directory is offline, all user validation occurs using the Security Accounts Manager (SAM) in the registry.

After a restore in safe mode, the domain controller is started in normal mode. When the directory service starts, the domain controller performs its normal consistency check, and the restored directory becomes online.

Restored pieces are reconciled and allowed only to "catch-up" to the rest of the enterprise. If restored onto a machine which is participating in an Active Directory enterprise, the restored section never becomes the master in the enterprise, and all of the other Active Directory pieces never synchronize with what was restored.

Note

When restoring the Active Directory, the Registry and Sysvol must also be restored.

6.3 MSCS cluster database support

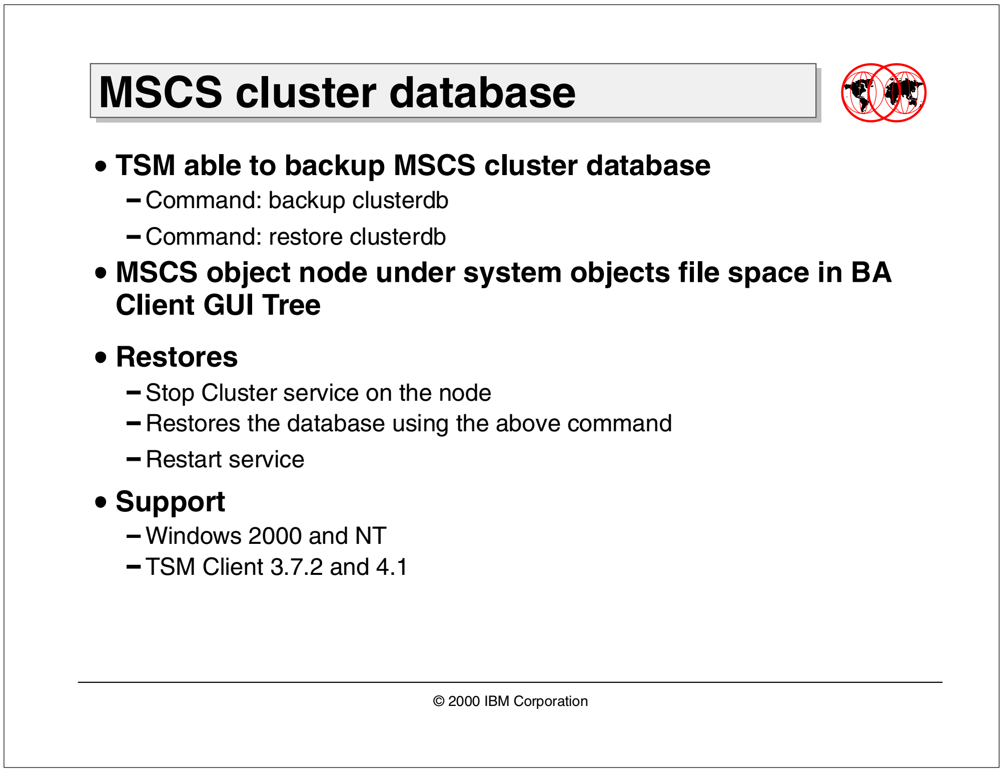

MSCS cluster database

- **TSM able to backup MSCS cluster database**
  - Command: backup clusterdb
  - Command: restore clusterdb
- **MSCS object node under system objects file space in BA Client GUI Tree**
- **Restores**
  - Stop Cluster service on the node
  - Restores the database using the above command
  - Restart service
- **Support**
  - Windows 2000 and NT
  - TSM Client 3.7.2 and 4.1

Microsoft Cluster Server (MSCS) lets you join two Windows servers, or nodes, through a shared disk subsystem so that the nodes share data to provide high resource availability. The combination of these servers is a cluster. (For detailed information on MSCS, including TSM integration, please see Section 5.1, “Microsoft Cluster server” on page 136).

6.3.1 Backup

The cluster database contains information about all physical and logical elements in a cluster. It is a set of keys under `HKEY_LOCAL_MACHINE` in the registry. The cluster database contains state data that is replicated among nodes to ensure that all nodes have a consistent configuration. This replication is achieved via log files that are stored on the quorum drive. Tivoli Storage Manager backs up this information using standard interfaces designed to ensure restore consistency.

The backup can be initiated either from the GUI or command line. The `MSCS` node under `System Objects` file space in the GUI tree for the BA client represents the MSCS cluster database and can be selected for backups. The command line option for the same is the `BACKUP CLUSTERDB` command.

6.3.2 Restore

The restore of the cluster database can be accomplished using the GUI option or command line. The command for restoring an MSCS cluster database is `RESTORE CLUSTERDB`. The cluster database on a node can only be written to and committed to registry when that node is not part of the cluster, or the cluster is down. To restore, TSM stops the running node as part of the cluster. It then restores the cluster database on that node and restarts the service.

6.4 COM+ class registration database support

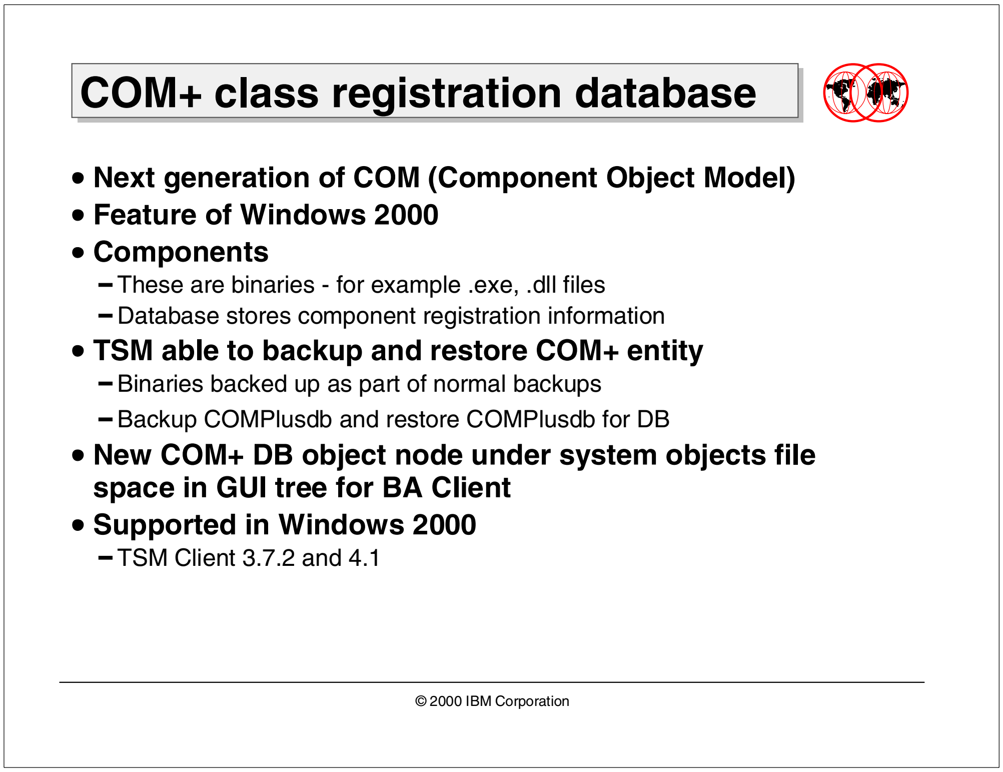

COM+ class registration database

- **Next generation of COM (Component Object Model)**
- **Feature of Windows 2000**
- **Components**
  - These are binaries - for example .exe, .dll files
  - Database stores component registration information
- **TSM able to backup and restore COM+ entity**
  - Binaries backed up as part of normal backups
  - Backup COMPlusdb and restore COMPlusdb for DB
- **New COM+ DB object node under system objects file space in GUI tree for BA Client**
- **Supported in Windows 2000**
  - TSM Client 3.7.2 and 4.1

COM (Component Object Model) is a binary standard for writing component software in a distributed system environment. It is a programming tool that lets application developers write application code that can notify client-side application code when a particular event has happened. Windows 2000 implements COM+ which is a next generation of the component model. It is the integration of Microsoft Transaction Server (MTS) and Microsoft Message Queuing (MSMQ) with the core operating system. It offers a registry independent database for storing component registration information.

The Windows 2000 COM+ entity consists of two elements, the COM+ binaries (EXEs and DLLs) and the COM+ database.

TSM is able to back up/restore the Windows 2000 COM+ entity. The COM+ binaries are backed up/restored through normal backup/restore operations. The COM+ database is backed up/restored via defined API's and is implemented as a TSM System Object. The backup/restore commands have been extended to support this new System Object.

The backup/restore can be initiated either from the GUI or command line. The `COM+ DB` node under `System Objects` file space in the GUI tree for the BA client represents the COM+ Class Registration database and can be selected for backups. The command line options for the same are the `BACKUP COMPLUSDB` and the `RESTORE COMPLUSDB` commands.

6.5 System files

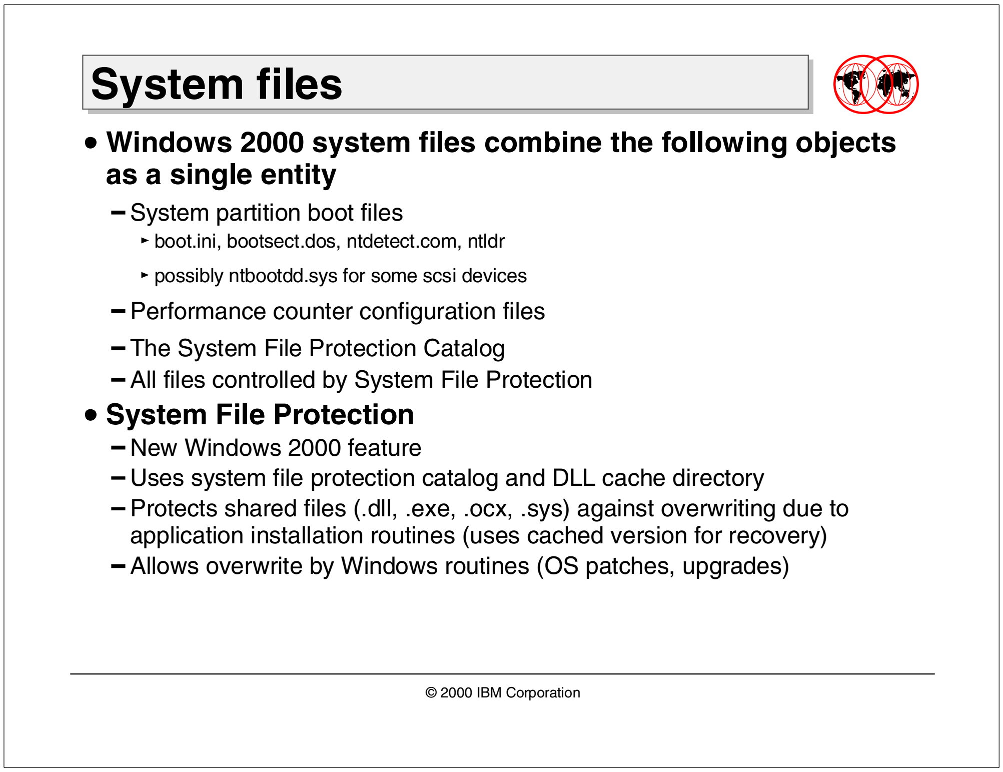

System files

- **Windows 2000 system files combine the following objects as a single entity**
  - System partition boot files
    - boot.ini, bootsect.dos, ntdetect.com, ntldr
    - possibly ntbootdd.sys for some scsi devices
  - Performance counter configuration files
  - The System File Protection Catalog
  - All files controlled by System File Protection
- **System File Protection**
  - New Windows 2000 feature
  - Uses system file protection catalog and DLL cache directory
  - Protects shared files (.dll, .exe, .ocx, .sys) against overwriting due to application installation routines (uses cached version for recovery)
  - Allows overwrite by Windows routines (OS patches, upgrades)

Windows 2000 System Files are one component of the Windows System State entity. This component consists of:

1. **System Partition boot files**
   - `boot.ini`
   - `bootsect.dos`
   - `ntdetect.com`
   - `ntldr`
   - `ntbootdd.sys` (possibly, for some SCSI devices)
2. **All files protected by System File Protection**
   - Files installed by Windows 2000 with the extensions .dll, .exe, .ocx, and .sys.
   - True Type fonts Micross.ttf, Tahoma.ttf, Tahomabd.ttf
3. **All files in the System File Protection service catalog directory**
   - `System32\CatRoot\{F750E6C3-38EE-11D1-85E5-00C04FC295EE}`
4. **Performance counter configuration files**
   - `System32\perf?00?.dat`
   - `System32\perf?00?.bak`

In these file names, “?” represents a single character.

6.5.1 Windows 2000 System File Protection (SFP)

A common problem in the history of the Windows operating systems has been the ability for shared system files to be overwritten by application installations. After these changes were made, the user would often experience unpredictable performance results, ranging from application errors to operating system crashes. This problem affects several types of files, most commonly dynamic link libraries (DLLs) and executable files (EXEs).

In Windows 2000, a new feature, called System File Protection (SFP) has been introduced that prevents the replacement of certain monitored system files. By preventing the replacement of essential system files, file version mismatches can be avoided.

System File Protection provides protection for system files by using a background mechanism that runs inside `WINLOGON.EXE` on a Windows 2000-based system. At the end of GUI-mode setup, SFP runs a scan of all protected files to ensure they have not been modified by applications installed by unattended installation. After successful startup, the System File Protection service performs a check of all catalog (.cat) files used to track correct file versions. If any catalog files are missing or corrupted, System File Protection renames the affected catalog file and retrieves a cached version of that file from the `dllcache` directory. If a cached copy of the catalog file is not available in `dllcache`, SFP requests the appropriate media (that is, Windows 2000 Service Pack, Windows 2000 Hotfix, and so on) to retrieve a clean copy of the catalog file.

System File Protection is triggered when it receives a directory change notification on a file in a protected directory. Once this notification is received, SFP determines which file was changed. If the file is protected, SFP will look up the file signature in a catalog file to determine if the new file is the correct Microsoft version. If it is not, then the system either replaces the file from a cache directory or asks for the original installation media.

The only installation mechanisms that will not trigger the System File Protection functionality are:

- Windows 2000 Service Pack installation (`UPDATE.EXE`).
- Hotfix distributions installed using `HOTFIX.EXE`.
- Operating system upgrade (`WINNT32.EXE`).

Replacement of system files that are on the protected list by any other mechanism will result in the System File Protection mechanism of the Windows 2000 operating system replacing the new system file with a copy of the protected version that is stored in the `dllcache` directory.

If the `dllcache` directory is overwritten or corrupted, its contents may be rebuilt by using System File Checker. It performs a scan of all protected files and replaces any invalid versions, in both `dllcache` and directory locations. The command line option is `C:\> SFC.EXE /SCANONCE`.

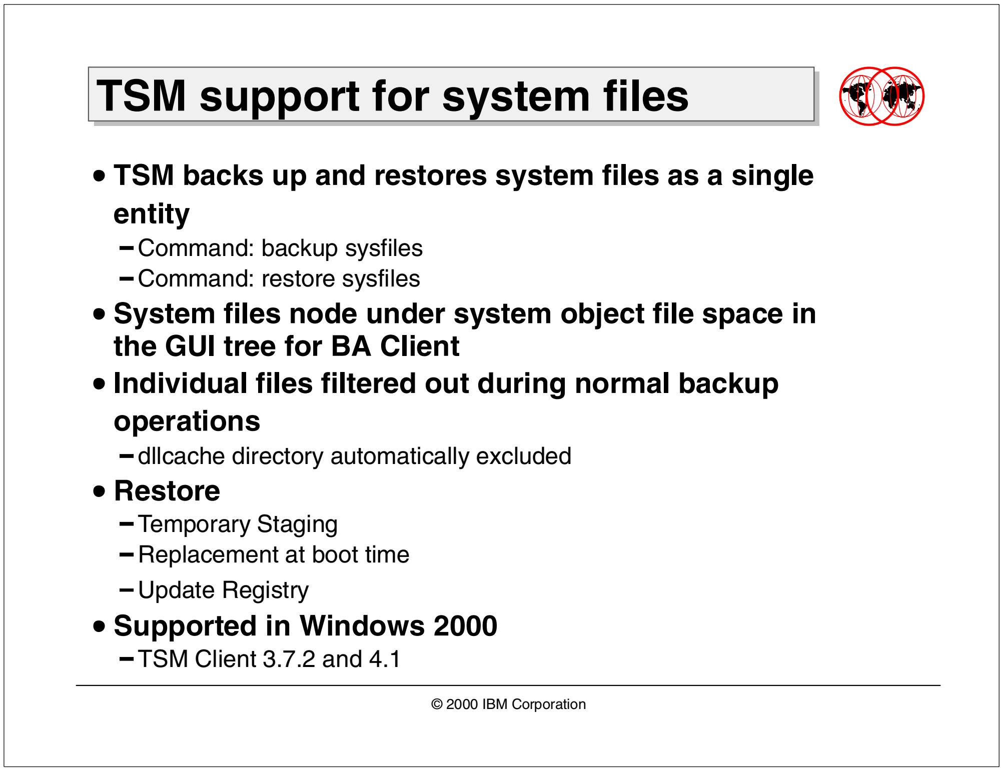

TSM support for system files

- **TSM backs up and restores system files as a single entity**
  - Command: backup sysfiles
  - Command: restore sysfiles
- **System files node under system object file space in the GUI tree for BA Client**
- **Individual files filtered out during normal backup operations**
  - dllcache directory automatically excluded
- **Restore**
  - Temporary Staging
  - Replacement at boot time
  - Update Registry
- **Supported in Windows 2000**
  - TSM Client 3.7.2 and 4.1

6.5.2 TSM support for system files

The boot files are treated as a single entity implemented as a TSM System Object and are backed up (and restored) as a whole. The System Object can be backed up either using the command line or GUI. The backup/restore commands have been extended to support the backup of these files. The commands are `BACKUP SYSFILES` and `RESTORE SYSFILES`.

In addition a `System Files` System Object has been added under the `System Objects` file space in the GUI tree for the BA client to support the same.

Individual files within this new System Object are excluded from normal backups. A mechanism has also been employed to ensure that the files in the `dllcache` directory are automatically excluded from backups.

The restores are somewhat different. The files in this System Object are protected, so they cannot be restored directly. Therefore TSM restores the files to a temporary location and marks them to be replaced at boot time. At boot time these files are copied over to their original location via an API function. The registry is also modified to take this into account.

6.6 Certificate server database support

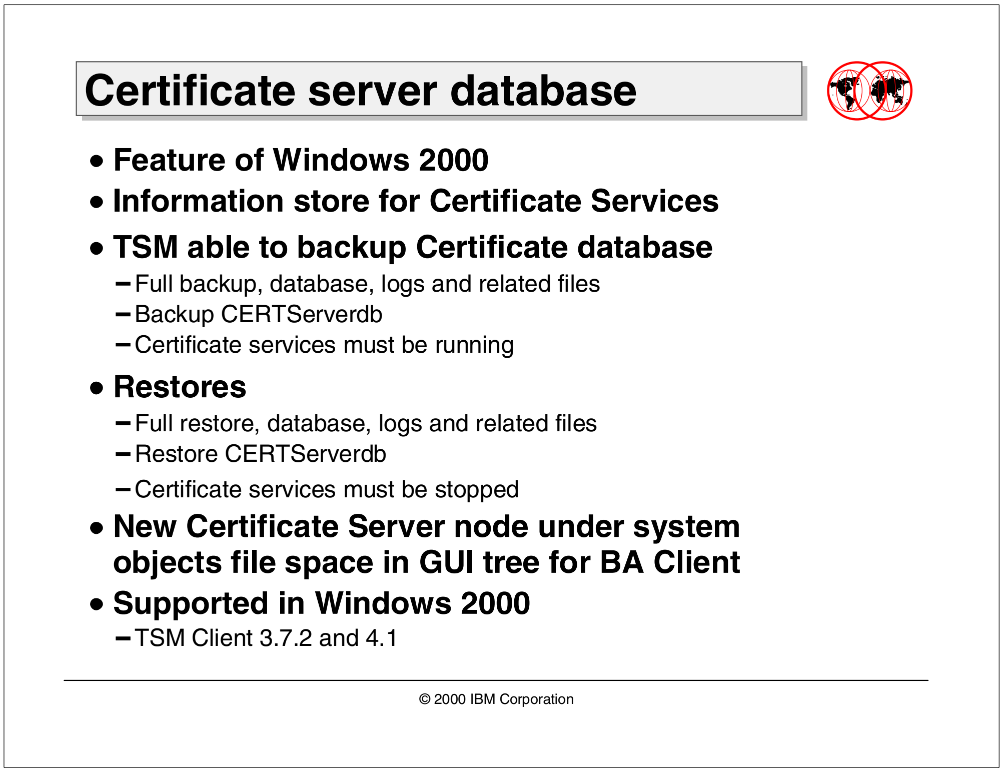

Certificate server database

- **Feature of Windows 2000**
- **Information store for Certificate Services**
- **TSM able to backup Certificate database**
  - Full backup, database, logs and related files
  - Backup CERTServerdb
  - Certificate services must be running
- **Restores**
  - Full restore, database, logs and related files
  - Restore CERTServerdb
  - Certificate services must be stopped
- **New Certificate Server node under system objects file space in GUI tree for BA Client**
- **Supported in Windows 2000**
  - TSM Client 3.7.2 and 4.1

Microsoft Windows 2000 Certificate Services lets you create a certification authority (CA) for managing the Windows 2000 public key infrastructure (PKI). A certification authority (CA) issues certificates that affirm the identity and other attributes of the certificate subject to other entities. PKI refers to a system of digital certificates (also called public key certificates) and CAs that verify and authenticate the validity of each party involved in an electronic transaction (secure E-mail, smart cards, and so on) allowing the secure exchange of information on open networks, such as the Internet, extranets, and intranets.

Certificate server service of Windows 2000 provides X.509 certificates for clients. The Certificate database is a repository for that information.

TSM is able to back up and restore the Windows 2000 Certificate Server database. A new System Object has been added to TSM. The backup and restore commands have been extended for this new System Object. The System Object can be backed up either using the command line or GUI. The commands are `BACKUP CERTSERVERDB` and `RESTORE CERTSERVERDB`. A new `Certificate Server` node has been added under the `System Objects` file space in the GUI tree for the Backup/Archive client to support the same.

Only full backup/restore is supported, which does a backup/restore of the database, logs, and associated files. The certificate services must be running before the backup is initiated, and must be stopped before doing a restore.

6.7 System Volume (Sysvol) support

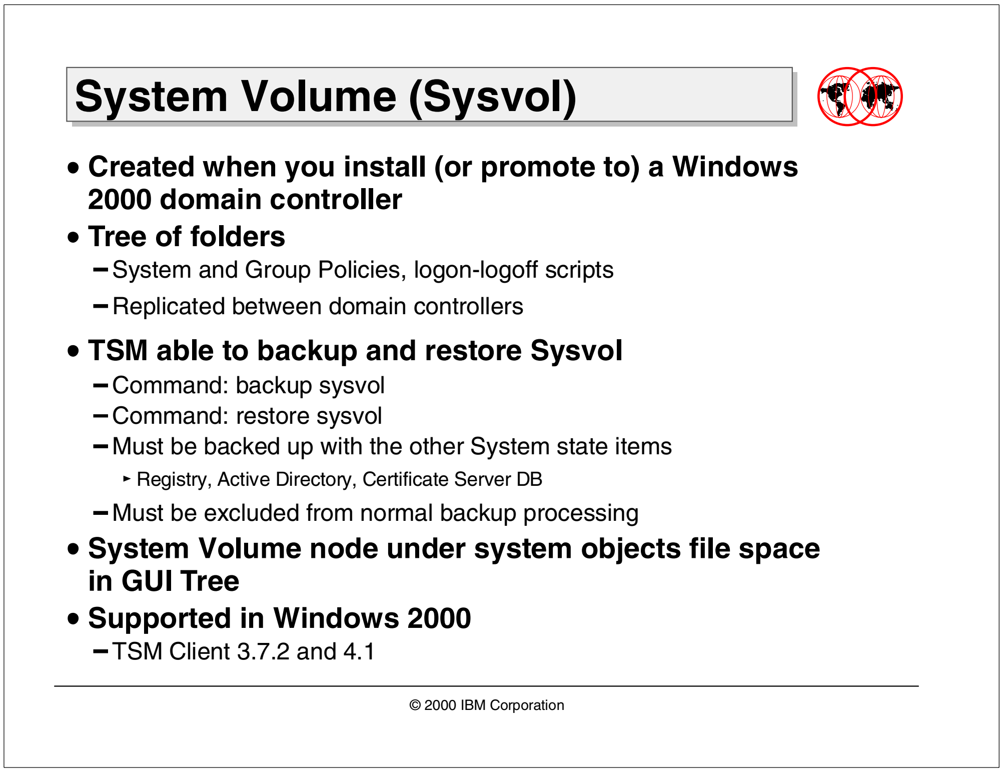

System Volume (Sysvol)

- **Created when you install (or promote to) a Windows 2000 domain controller**
- **Tree of folders**
  - System and Group Policies, logon-logoff scripts
  - Replicated between domain controllers
- **TSM able to backup and restore Sysvol**
  - Command: backup sysvol
  - Command: restore sysvol
  - Must be backed up with the other System state items
    - Registry, Active Directory, Certificate Server DB
  - Must be excluded from normal backup processing
- **System Volume node under system objects file space in GUI Tree**
- **Supported in Windows 2000**
  - TSM Client 3.7.2 and 4.1

© 2000 IBM Corporation

The Windows 2000 System Volume, or Sysvol, is built during the creation of a domain controller. This is the data which has been replicated from the Active Directory and which is not part of the directory tree database. It is a tree of folders containing files that need to be available and synchronized between domain controllers in a domain or forest. The folders include:

- SYSVOL share
- NETLOGON share
- Windows 95, Windows 98, and Windows NT 4.0 system policies
- Windows 2000 Group Policy settings
- User logon and logoff scripts

For example, the default folder structure contains the following folders for policies or scripts used by network clients:

```
\\Winnt\Sysvol\Sysvol\domain_name\Policies
```

```
\\Winnt\Sysvol\Sysvol\domain_name\Scripts
```

When you add, remove, or modify the contents of the Sysvol folder on a domain controller, those changes are replicated (by Active Directory) to the Sysvol folders on all other domain controllers in the domain.

TSM has the ability to back up and restore the information in Sysvol. A new Sysvol System Object has been added to TSM. The backup and restore commands have been extended for this new System Object. The System Object can be backed up either using the command line or GUI. The commands are `BACKUP SYSVOL` and `RESTORE SYSVOL`.

A new `System Volume` node has been added under the `System Objects` file space in the GUI tree for the Backup/Archive client to support the same.

Note:

The data in the Sysvol should be excluded from normal backup processing and must be backed up in conjunction with other System State data. The Sysvol has interdependencies with the Registry, the Active Directory database, and the Certificate Server database.

For details on backing up System State System Objects please see Part 6.1, “TSM support for new System Objects” on page 197.

6.8 Encrypting File System (EFS)

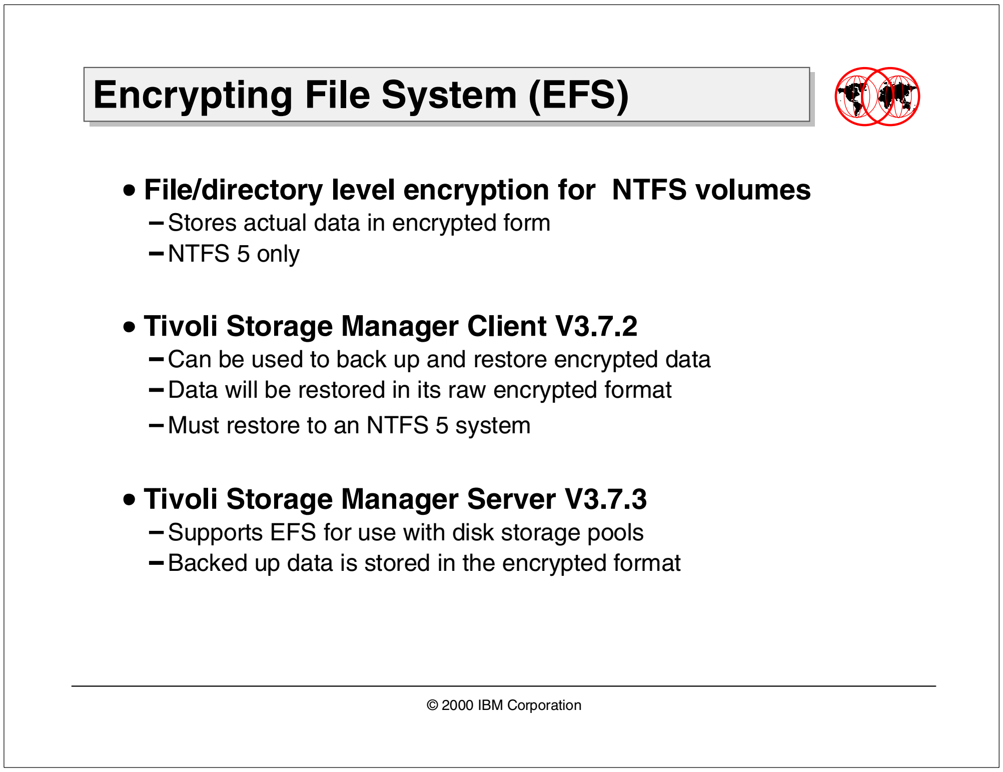

Encrypting File System (EFS)

- **File/directory level encryption for NTFS volumes**
  - Stores actual data in encrypted form
  - NTFS 5 only
- **Tivoli Storage Manager Client V3.7.2**
  - Can be used to back up and restore encrypted data
  - Data will be restored in its raw encrypted format
  - Must restore to an NTFS 5 system
- **Tivoli Storage Manager Server V3.7.3**
  - Supports EFS for use with disk storage pools
  - Backed up data is stored in the encrypted format

File/Directory level encryption is implemented in NTFS for enhanced security in NTFS volumes. It is possible for an unauthorized user to remove a physical storage volume, and mount it on another Windows 2000 system. Once mounted by another system, the administrator of this system can take ownership of all the data, bypassing NTFS security. Laptops are particularly vulnerable to this scenario. Windows 2000 Encrypting File System (EFS) solves this problem by storing actual data in encrypted form, thus providing security when the storage media is removed from a Windows 2000 based system.

The Tivoli Storage Manager client supports EFS and can be used to back up and restore encrypted data, the data will be restored in its raw encrypted format. The Tivoli Storage Manager server supports EFS for use with disk storage pools only and the backed up data is stored in the encrypted format. Tivoli Storage Manager does not try to decrypt the files; the data is stored on the server in a secure format.

EFS is only available in NTFS 5, Windows 2000. EFS backed using the Tivoli Storage Manager client can only be restored to an NTFS 5, Windows 2000 system.

When EFS data is backed up or restored, it is either cached in memory or, if larger than 64K, it is cached to a temporary file. By default, the temporary file path is the path specified by the TMP environment variable, if TEMP is not defined.

6.9 Reparse point support

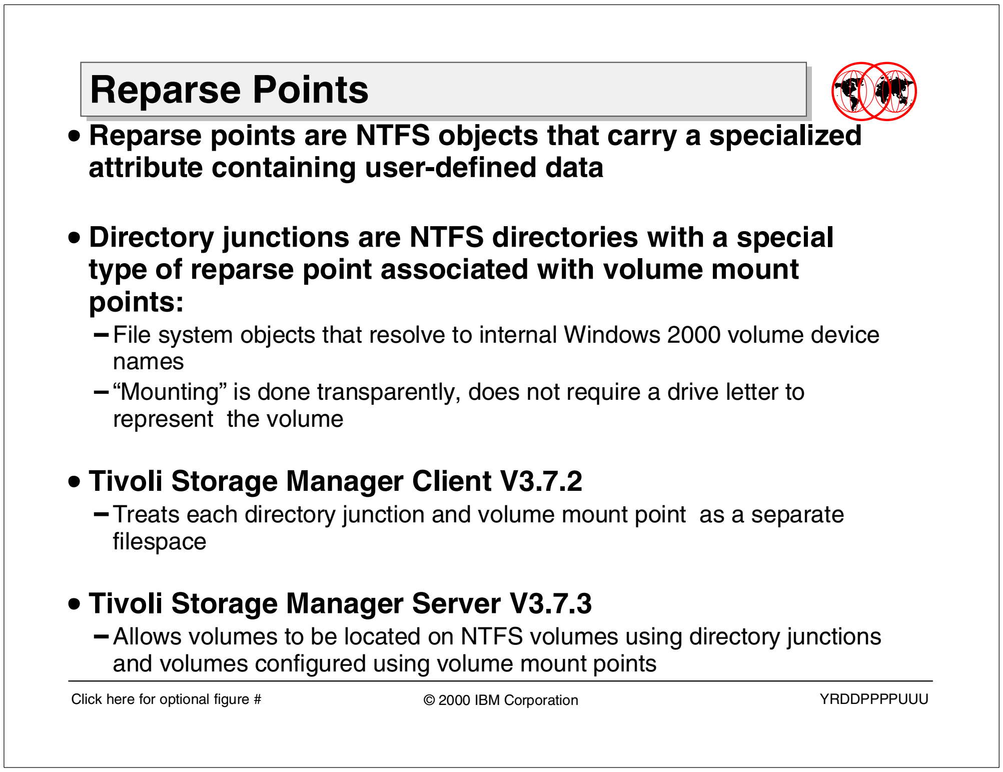

Reparse Points

- **Reparse points are NTFS objects that carry a specialized attribute containing user-defined data**
- **Directory junctions are NTFS directories with a special type of reparse point associated with volume mount points:**
  - File system objects that resolve to internal Windows 2000 volume device names
  - “Mounting” is done transparently, does not require a drive letter to represent the volume
- **Tivoli Storage Manager Client V3.7.2**
  - Treats each directory junction and volume mount point as a separate filespace
- **Tivoli Storage Manager Server V3.7.3**
  - Allows volumes to be located on NTFS volumes using directory junctions and volumes configured using volume mount points

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

Reparse points are NTFS System Objects that carry a specialized attribute containing user-defined data. Reparse points are used to extend functionality in the I/O subsystem. One reparse point is allowed per file or directory. Remote storage and NTFS directory junction points are examples of functionality based on reparse points.

Note that with volume mount points and Windows 2000 directory junction points, the file system namespace is no longer guaranteed to be a directed acyclic graph. It is possible to define reparse points that create cycles in the namespace.

The Tivoli Storage Manager backup client supports NTFS reparse points and handles cycle-related issues correctly. In general, when a reparse point is encountered the Tivoli Storage Manager client will backup the meta-data representing the reparse point.

6.9.1 NTFS directory junctions

An NTFS directory junction is an NTFS directory with a special type of reparse point associated with it. An NTFS directory junction can be mapped to any local or remote target directory. The Tivoli Storage Manager client can be used to back up directory junctions. The Tivoli Storage Manager client treats each directory junction as a separate file space. The Tivoli Storage Manager Server supports directory junctions by allowing Tivoli Storage Manager volumes to be located on NTFS volumes using directory junctions. This support is available in the Tivoli Storage Manager client V3.7.2 and V3.7.3 for the server.

6.9.2 Volume mount points

Volume mount points are file System Objects that resolve to internal Windows 2000 volume device names in a robust, persistent manner. Placing a volume mount point on an NTFS directory causes the storage subsystem to resolve the directory to a specified local volume. The “mounting” is done transparently and does not require a drive letter to represent the volume. A Windows 2000 mount point always resolves to the root directory of the desired volume.

The Tivoli Storage Manager client handles volume mount points just like it does for directory junctions. Again the Tivoli Storage Manager client treats each directory junction as a separate file space. The Tivoli Storage Manager server supports volume mount points by allowing Tivoli Storage Manager volumes to be located on NTFS volumes configured using volume mount points.

Tivoli Storage Manager Server support needs to be at V3.7.3; the Backup/Archive client support needs to be at level V3.7.2.

6.10 Link tracking and Object IDs

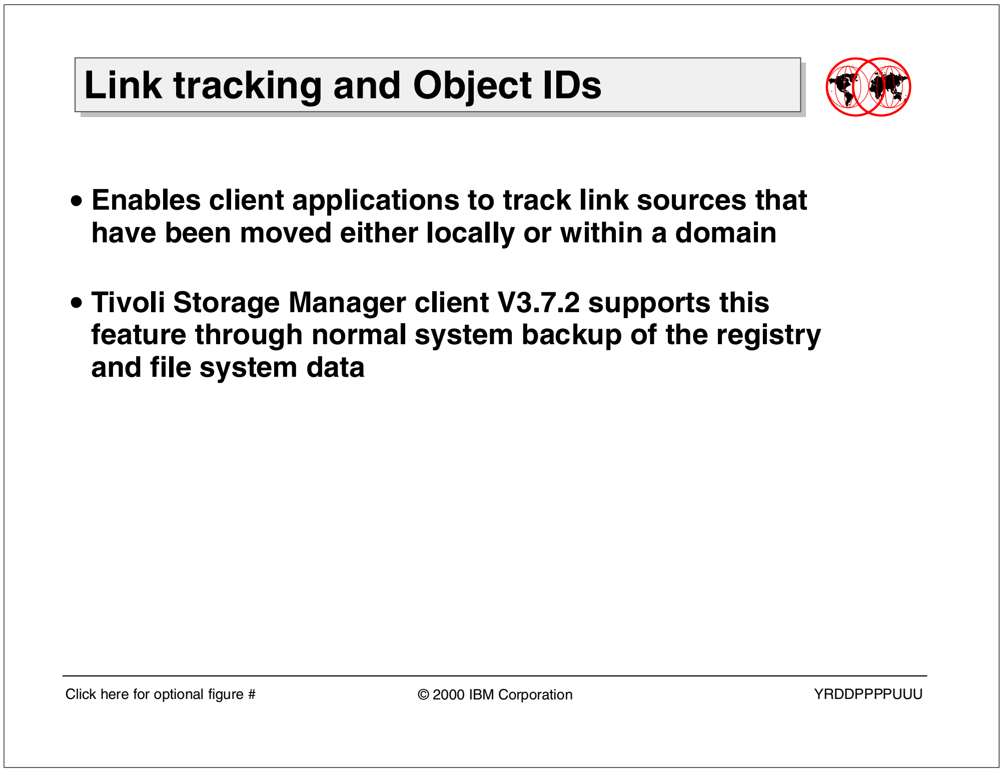

Link tracking and Object IDs

- **Enables client applications to track link sources that have been moved either locally or within a domain**
- **Tivoli Storage Manager client V3.7.2 supports this feature through normal system backup of the registry and file system data**

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

Link tracking and Object IDs (OIDs) are new features with NTFS. Windows 2000 provides a link tracking service that enables client applications to track link sources that have been moved either locally or within a domain. As a result, clients that subscribe to the link tracking service can maintain the integrity of their references because the System Object’s references can be moved transparently. Files in NTFS can be references by a unique Object ID. Link tracking stores a file’s OID as part of its tracking information.

The Tivoli Storage Manager client V3.7.2 supports this feature through normal system backup of the registry and file system data.

6.11 File Replication Service (FRS) support

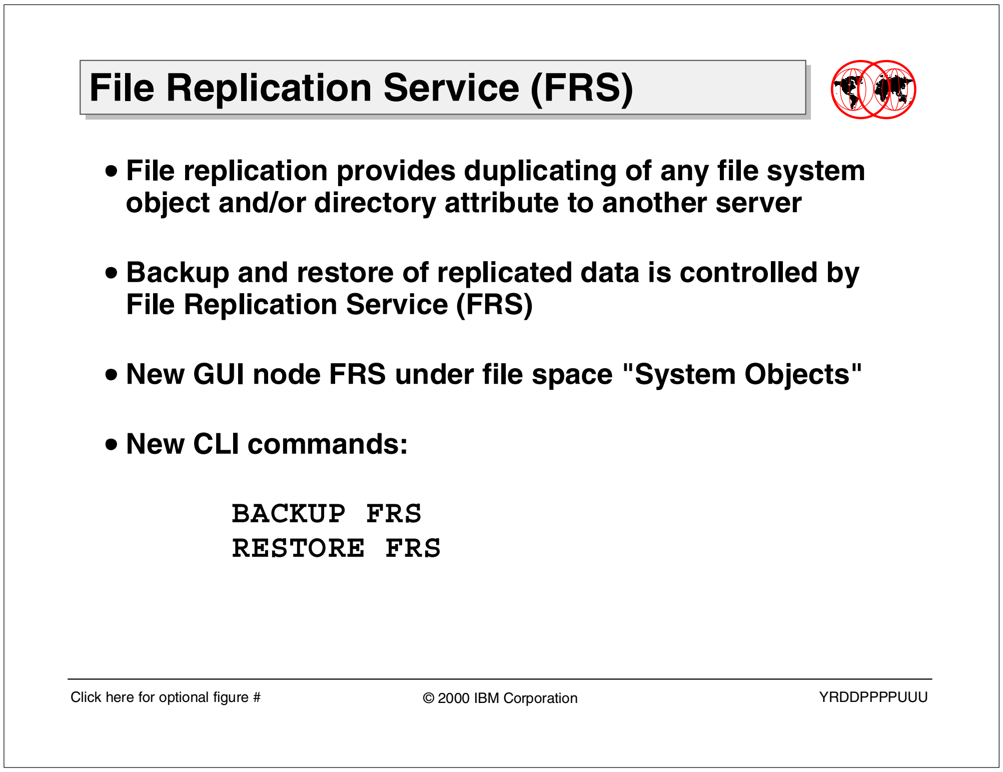

File Replication Service (FRS)

- **File replication provides duplicating of any file system object and/or directory attribute to another server**
- **Backup and restore of replicated data is controlled by File Replication Service (FRS)**
- **New GUI node FRS under file space "System Objects"**
- **New CLI commands:**

```
BACKUP FRS
RESTORE FRS
```

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

Windows 2000 introduces full two-way file replication for NTFS. File replication provides a mechanism for duplicating any file System Object and/or directory attribute to another server. Windows 2000 replication is based on unique sequence numbers used in the Change Journal.

The replication service stores state in both the directory services database and a local Jet database.

The backup and restore of replicated data is controlled by the Windows 2000 File Replication Service (FRS). FRS is employed by Active Directory and DFS to replicate critical files for performance and availability.

Tivoli Storage Manager provides backup and non-authoritative restore of data under the control of FRS. In the GUI, this data is backed up by selecting the FRS System Object node in the GUI tree. On the command line, use the commands:

- `BACKUP FRS`
- `RESTORE FRS`

Tivoli Storage Manager obtains a list of directories from the FRS that are replicated, and then backs up all files and directories under this list of directories. FRS files are grouped as one System Object at the Tivoli Storage Manager Server, and there is no file level granularity available for backup or restore; only the System Object can be restored.

A new term with the FRS is non-authoritative restore, which means that Tivoli Storage Manager does not cause FRS to force-replicate the restored files to FRS partner systems. Tivoli Storage Manager client does not currently handle the Jet DB associated with the FRS.

6.12 Distributed File System (DFS) support

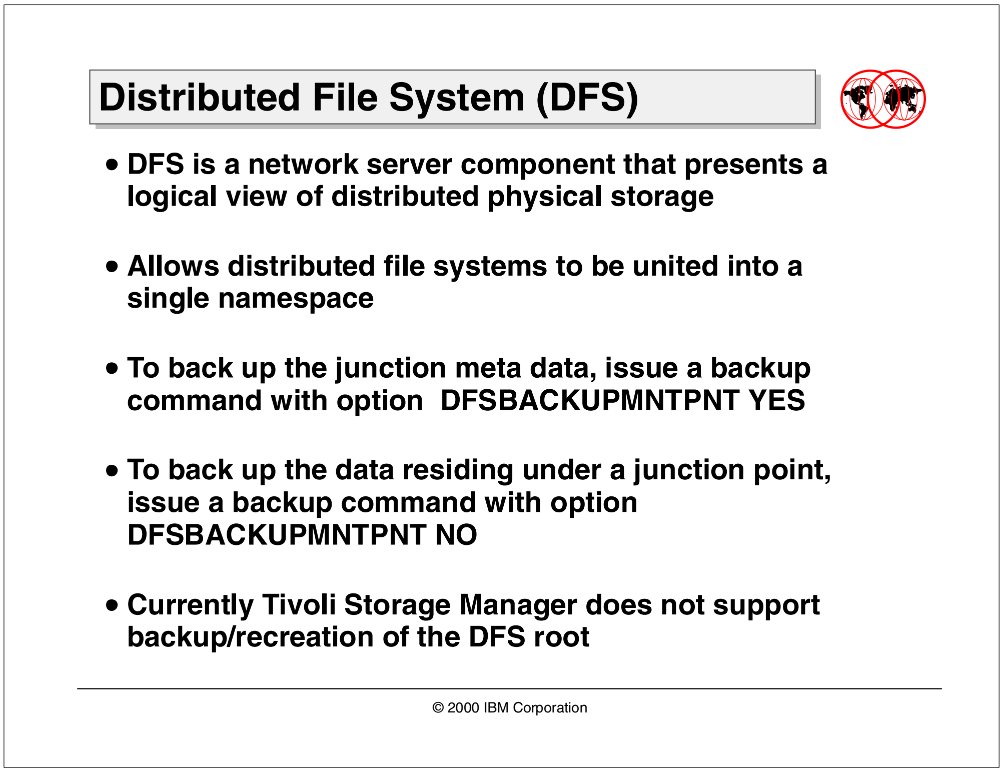

Distributed File System (DFS)

- DFS is a network server component that presents a logical view of distributed physical storage
- Allows distributed file systems to be united into a single namespace
- To back up the junction meta data, issue a backup command with option DFSBACKUPMNTPNT YES
- To back up the data residing under a junction point, issue a backup command with option DFSBACKUPMNTPNT NO
- Currently Tivoli Storage Manager does not support backup/recreation of the DFS root

Distributed File System (DFS) for Windows 2000 is a network server component that presents a logical view of distributed physical storage. DFS allows distributed file systems to be united into a single namespace. Its usage includes higher data availability, load balancing, name transparency, and flexible volume administration.

DFS is an extension of the Common Internet File System (CIFS). DFS extends the logical view of storage to include storage across the entire enterprise network. It presents a logical view of distributed physical storage and allows distributed file systems to be united into a single namespace.

A network can host many individual DFS volumes, and these volumes are accessed using a standard UNC name.

To back up the junction meta data, you have to perform a backup with the option DFSBACKUPMNTPNT YES. You back up the data residing under a junction point by performing a backup with DFSBACKUPMNTPNT NO. Currently Tivoli Storage Manager is not able to back up or recreate the DFS root.

6.13 Change Journal

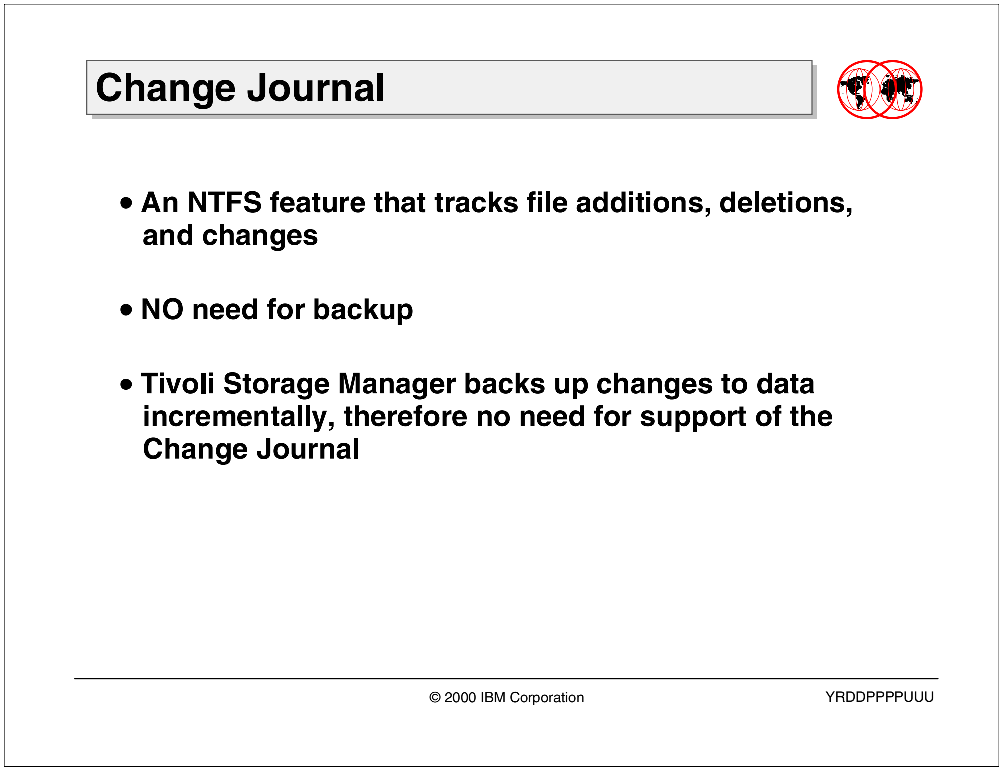

Change Journal

- **An NTFS feature that tracks file additions, deletions, and changes**
- **NO need for backup**
- **Tivoli Storage Manager backs up changes to data incrementally, therefore no need for support of the Change Journal**

© 2000 IBM Corporation YRDDPPPPUUU

The Change Journal is an NTFS feature that tracks file additions, deletions, and changes. While backing up the changed files, there is NO need or requirement for backing up the Change Journal.

Tivoli Storage Manager does not support the backup of the Change Journal. Because Tivoli Storage Manager can back up changes to data incrementally, there does not appear to be a need for supporting a backup of this journal.

Chapter 7. Tivoli SANergy File Sharing

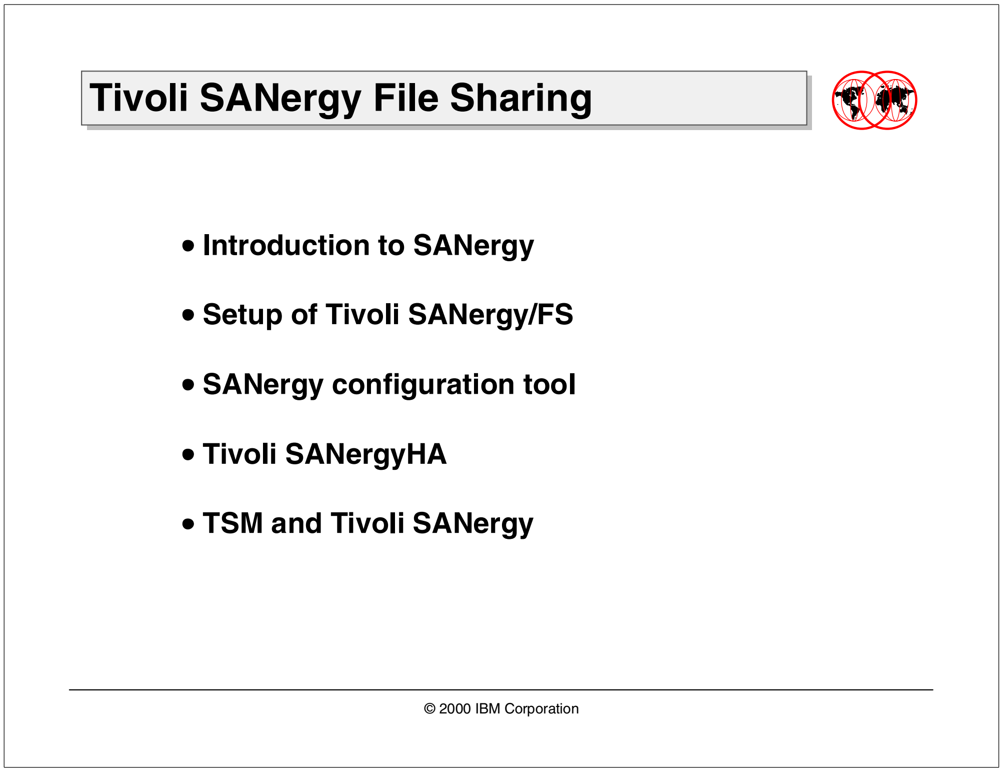

Tivoli SANergy File Sharing

- Introduction to SANergy
- Setup of Tivoli SANergy/FS
- SANergy configuration tool
- Tivoli SANergyHA
- TSM and Tivoli SANergy

© 2000 IBM Corporation

In April 2000, Tivoli announced the product Tivoli SANergy File Sharing, which introduced LAN file sharing technologies to Storage Area Networks (SAN). This chapter will describe the SANergy product; the architecture, data flow, and implementation of SANergy; and their interrelationship with Tivoli Storage Manager.

7.1 Introduction to SANergy File Sharing

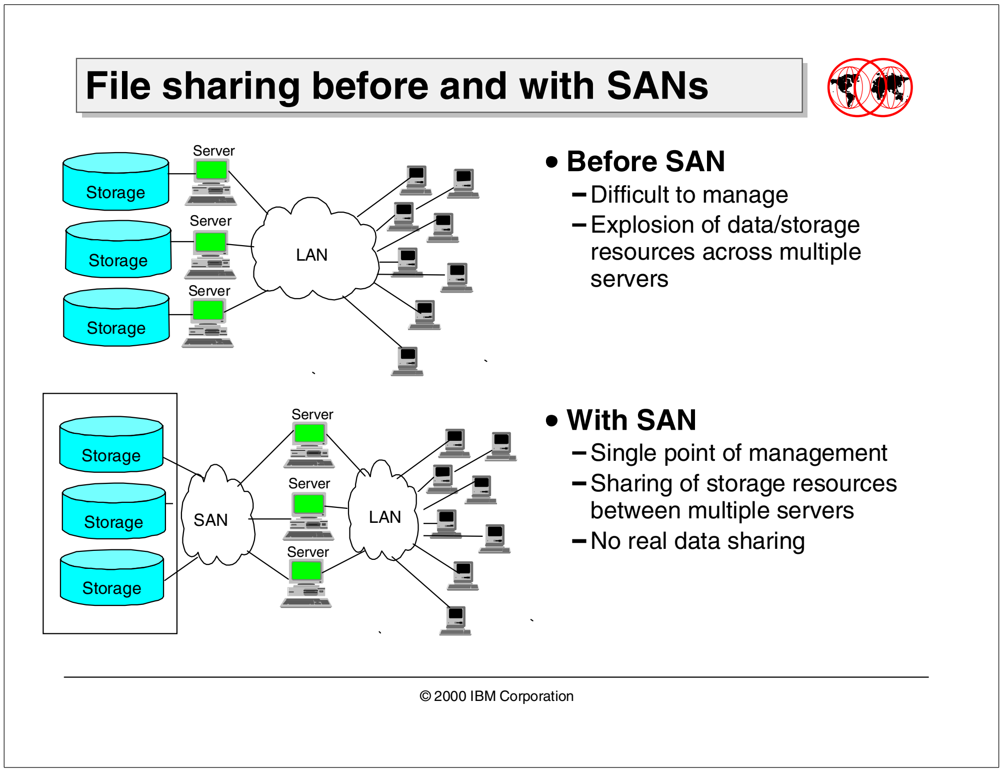

File sharing before and with SANs

- **Before SAN**
  - Difficult to manage
  - Explosion of data/storage resources across multiple servers
- **With SAN**
  - Single point of management
  - Sharing of storage resources between multiple servers
  - No real data sharing

7.1.1 File sharing before and with SAN’s

Before Storage Area Networks (SANs), each server had its own storage. The number of servers in most large organizations today has exploded, which in many cases has created serious technology management problems. While the costs of servers and storage devices has plummeted, the tools available to manage this huge numbers of servers in real-time has not kept up with the management requirements. IT customers around the world have found that providing backup, security, and recovery for large numbers of distributed servers, for example, is a much more complicated and difficult problem than it is for centralized mainframes.

The original LANs were installed largely to support clusters of PCs or workstations requiring straightforward file and printer sharing. But today’s major client-server applications such as ERP, e-business and business intelligence require a different level of support. These new client-server systems have made the enterprise’s distributed application and data servers every bit as critical as any of the enterprise’s mainframes.

Recently, organizations have begun to look at a more structured strategy of dealing with the many servers in their inventory. One of the most popular emerging strategies is server consolidation. This involves projects such as:

- Relocation of servers — The servers are relocated to fewer locations.
- Physical consolidation — Many small servers are consolidated into a few large servers along with associated data storage consolidation.
- Application consolidation — Different applications are consolidated to one large server.
- Data/application consolidation — This involves vertical consolidation to a different platform.

Tivoli SANergy allows file sharing between different heterogeneous platforms, including Windows NT, SUN Solaris, IBM AIX, and Apple Macintosh. Without SANergy, a piece of data held on disk to be used for Windows NT is often copied to another platform to perform some kind analysis. Files could not take full advantage of the capabilities of the SAN.

Tivoli SANergy File Sharing software helps computer users in networked environments speed transaction time by taking full advantage of the capabilities the SAN.

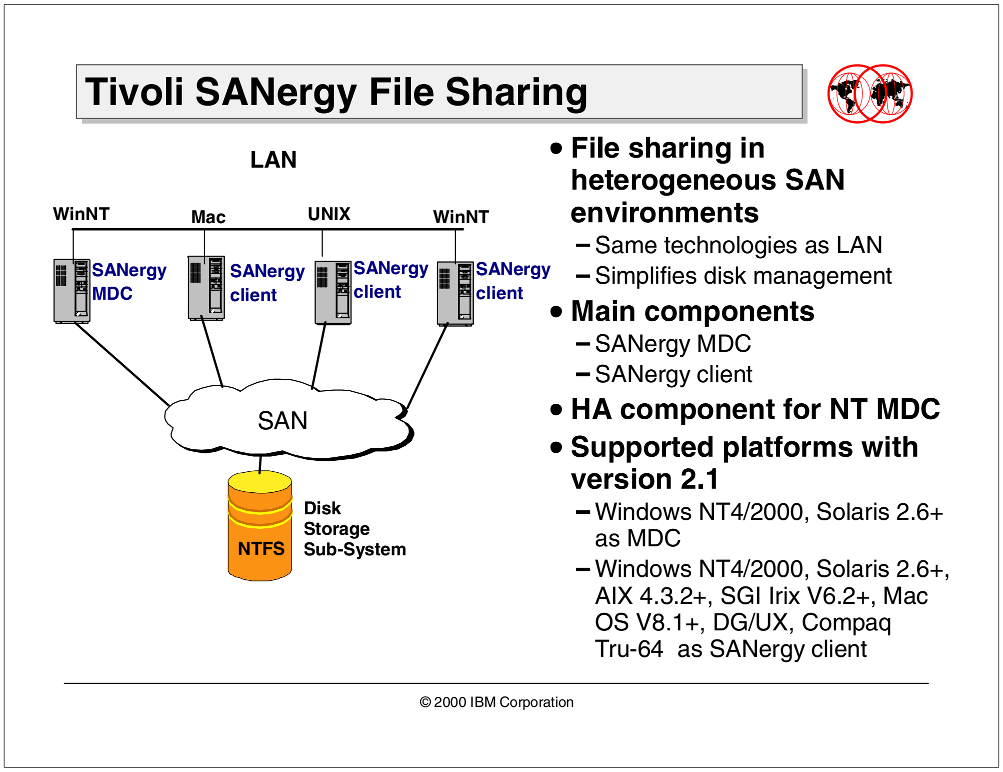

Tivoli SANergy File Sharing

- **File sharing in heterogeneous SAN environments**
  - Same technologies as LAN
  - Simplifies disk management
- **Main components**
  - SANergy MDC
  - SANergy client
- **HA component for NT MDC**
- **Supported platforms with version 2.1**
  - Windows NT4/2000, Solaris 2.6+ as MDC
  - Windows NT4/2000, Solaris 2.6+, AIX 4.3.2+, SGI Irix V6.2+, Mac OS V8.1+, DG/UX, Compaq Tru-64 as SANergy client

7.1.2 Tivoli SANergy File Sharing

SANs are fast becoming the preferred method for companies to manage the vast amounts of data created by their e-business systems — by creating networks to connect "islands of information" that can be shared quickly across the enterprise. SANs can link data centers and manage data across heterogeneous computing platforms for real-time backup and recovery, data migration, and resource sharing.

Tivoli SANergy File Sharing (SANergyFS) is the only SAN software that allows for the sharing of application files and data between a variety of heterogeneous servers and workstations connected to a SAN. In addition, Tivoli SANergyFS software uses only industry-standard file systems, enabling multiple computers simultaneous access to shared files through the SAN. This allows users to leverage existing technical resources instead of learning new tools or migrating data to a new file system infrastructure. This software allows SAN-connected computers to have the high-bandwidth disk connection of a SAN while keeping the security, maturity and inherent file sharing abilities of a LAN.

The SANergy software employs technology to combine the simplicity of LAN-based file sharing with the very high data transfer speeds afforded by modern Fibre Channel, SCSI, and SSA storage networks. This enables the use of high-speed, heterogeneous data sharing without the performance-limiting bottlenecks of file servers and traditional networking protocols.

The Tivoli SANergyFS product is unique in that it extends standard file systems and network services provided by the operating systems that it supports (Windows NT, MacOS, and UNIX). As an O/S extension built on standard systems interfaces, SANergyFS fully supports the user interface, management, access control, and security features native to the host platforms, providing all the file system management, access control, and security expected in a network. With SANergyFS, virtually any network-aware application can access any file at any time, and multiple systems can transparently share common data.

In addition to the SAN, Tivoli SANergyFS also uses a standard LAN for all the meta data associated with file transfers. Because Tivoli SANergyFS is NTFS-based, should the SAN fail, access to data via the LAN is still possible. Since each system has direct access to the Tivoli SAN-based storage, Tivoli SANergyFS can eliminate the file server as a single point of failure for mission-critical enterprise applications. Tivoli SANergyFS can also easily manage all data backup traffic over the storage network, while the users enjoy unimpeded LAN access to the existing file servers.

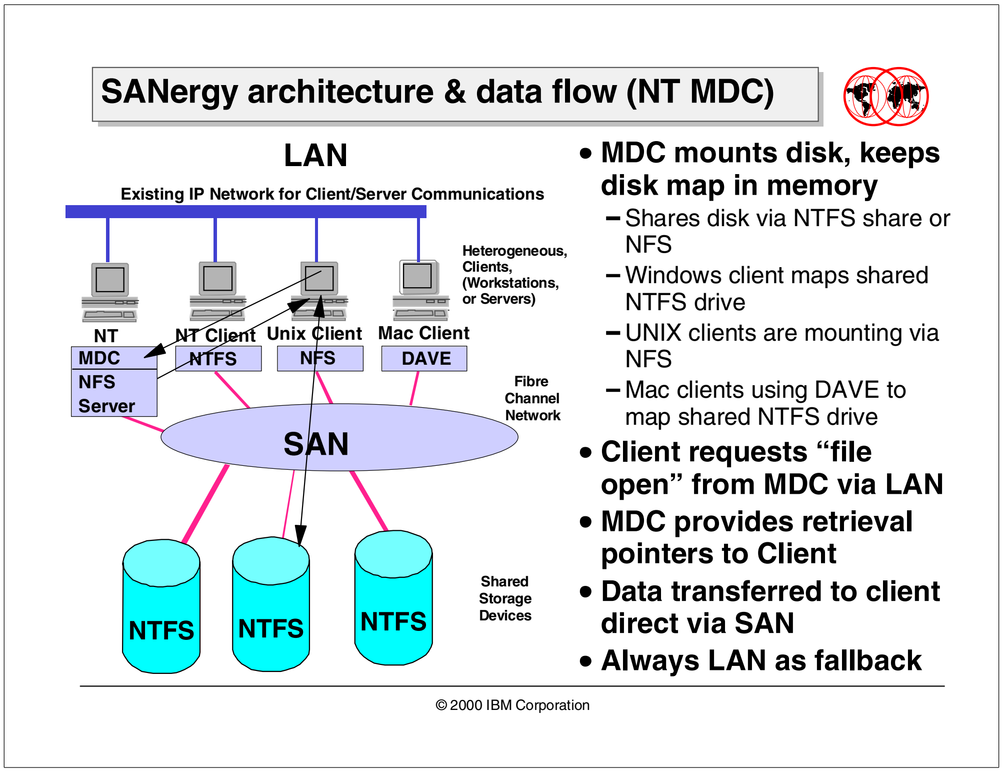

SANergy architecture & data flow (NT MDC)

- **MDC mounts disk, keeps disk map in memory**
  - Shares disk via NTFS share or NFS
  - Windows client maps shared NTFS drive
  - UNIX clients are mounting via NFS
  - Mac clients using DAVE to map shared NTFS drive
- **Client requests “file open” from MDC via LAN**
- **MDC provides retrieval pointers to Client**
- **Data transferred to client direct via SAN**
- **Always LAN as fallback**

7.1.3 Tivoli SANergyFS architecture and data flow

The basic problem in storage area networking at the file level is keeping the separate operating systems up to date with each other's independent and asynchronous use of the storage. Tivoli SANergyFS is a hybrid of conventional networking and direct attached storage. Conventional networking is rich with abilities for keeping many computers coherent. That is, if one computer has an open view of a directory, and another changes that directory (adds/deletes a file), the view on all computers will change.

Conventional networking allows administrators to establish centralized access control lists and other data management facilities. Data "about" data is referred to as "meta data". Examples include file names, file sizes, and access control lists. The Tivoli SANergyFS architecture lets meta data transactions take place over conventional LAN networking. The actual content of files moves on the high-speed direct SAN connection.

Tivoli SANergyFS works with Ethernet, ATM, or anything else that carries networking protocols. The network operating system can also be SMB (Windows NT), Appletalk, NFS (UNIX), or a combination. (Only certain protocols and combinations have been tested, however.)

Similarly, Tivoli SANergyFS supports any available disk-attached storage fabric. This includes Fibre Channel, SSA, SCSI, and any other disk-level connection. It is also possible for installations to use one set of physical wiring to carry both the LAN and storage traffic. (FC supports both IP and SCSI protocols.)

When you use Tivoli SANergyFS, one computer in the workgroup is tagged as the *Meta Data Controller* (MDC) for a particular volume. You can have a single computer as the MDC for all volumes, or it can be spread around. The other computers are *SANergy clients* and use conventional networking to "mount" that volume, and Tivoli SANergyFS on those clients separates the meta data from the raw data automatically.

7.2 Tivoli SANergy implementation

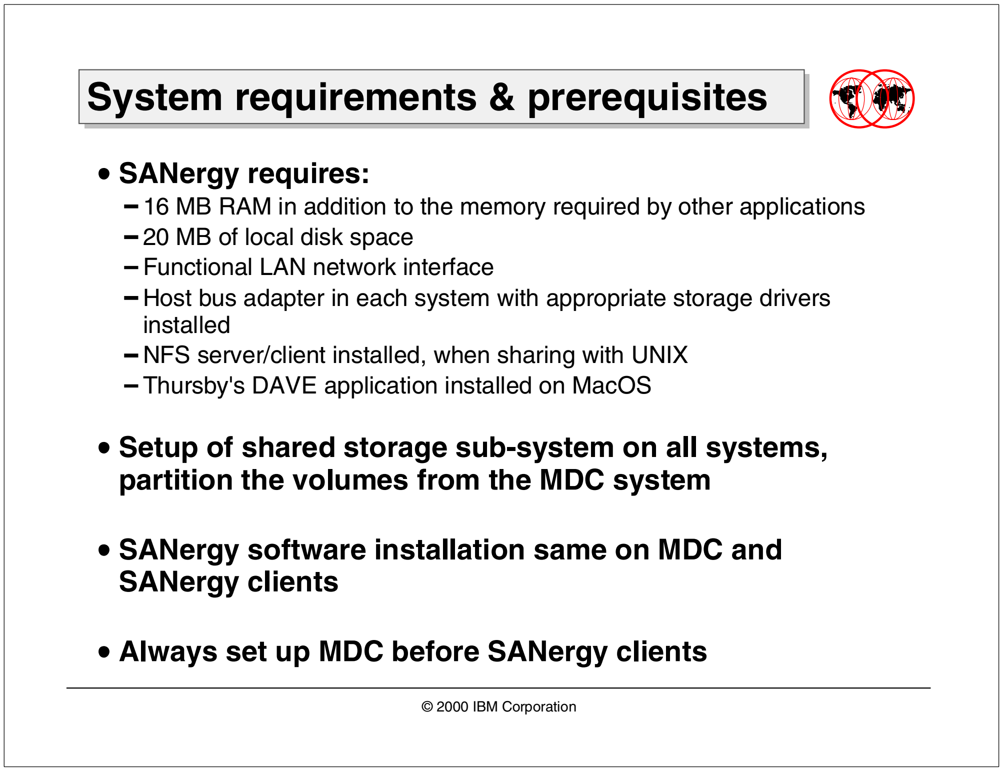

System requirements & prerequisites

- **SANergy requires:**
  - 16 MB RAM in addition to the memory required by other applications
  - 20 MB of local disk space
  - Functional LAN network interface
  - Host bus adapter in each system with appropriate storage drivers installed
  - NFS server/client installed, when sharing with UNIX
  - Thursby's DAVE application installed on MacOS
- **Setup of shared storage sub-system on all systems, partition the volumes from the MDC system**
- **SANergy software installation same on MDC and SANergy clients**
- **Always set up MDC before SANergy clients**

7.2.1 Overall system requirements

For each system that will be part of the Tivoli SANergy installation the following is required, both for a SANergy MDC or SANergy client:

- 16 MB RAM in addition to the memory required by other applications.
- 20 MB of local disk space.
- 1 PCI slot available for new host bus adapter.
- Functional LAN network interface.
- Floppy and CD drives for software installation.
- SAN host bus adapter in each system with appropriate SAN drivers installed.

At the time of writing this book, the requirements for the supported operating systems were as follows.

7.2.1.1 Macintosh system requirements

The requirements for SANergy on a Macintosh are:

- Model: PowerPC series systems with PCI bus.
- MacOS version 8.1 or higher. To determine which version of MacOS and how much memory is installed on your system, go to the Apple-menu and select "About This Computer...."
- All Macintosh systems on the SAN running the same version of MacOS.

7.2.1.2 Windows NT system requirements

The requirements for SANergy on an NT Pentium or an Alpha processor with PCIBus are:

- Windows NT 4.0 Workstation or Server and Service Pack 3 or higher. To determine the version of Windows NT and the amount of memory installed on your system, on the desktop double-click My Computer, then select Help About Windows NT.
- Integraph's DiskShare, NetManage's Chameleon UNIX Link 97, FTP Software's InterDrive, or Hummingbird Communications NFS Maestro, if there is a UNIX computer on the SAN.

7.2.1.3 UNIX system requirements

The system requirements for a UNIX system are as described below for SGI, Solaris, and AIX.

SGI

The requirements for SANergy on an SGI system are:

- Model: O2, Octane and others. Contact your sales representative for more details.
- Irix version 6.2 or later.
- ONC3/NFS version 2, from the SGI operating system diskettes.

Solaris

The requirements for SANergy on a Solaris system are:

- Model: SPARC system
- Solaris version 2.6 or later.
- Maximum shared memory (shmmax) set to 4 MB or greater.

AIX

The requirements for SANergy on an AIX system are:

- Model: RS6000 system.
- AIX version 4.3.2 or later.

7.2.1.4 Network system requirements

The SANergy network requirements are:

- Access to World Wide Web for software updates.
- **MDC** — You must have at least one MDC on the SAN. The MDC can be either:
  - **Windows NT** — The shared volumes must be formatted under Windows NT using NTFS. Refer to NT documentation on how to format disks.
  - **Solaris** — The shared volumes must be formatted as UFS and non-striped. Refer to UNIX documentation on how to format disks.
- **SAN** — All computers attached directly to the storage you want to share, not just by IP, but also by a SCSI-type connection, which includes SCSI, SSA, or FC.

- **Mapping** —You must have each client computer "map" the Meta Data Controller volume. This requires general networking hardware and software (for example, TCP/IP on Ethernet) to be installed and configured.
  - – With Windows NT, use conventional Windows NT networking and sharing.
  - – With MacOS, use the DAVE software from Thursby.
  - – With UNIX computers (SGI, Solaris SPARC, or AIX), use NFS. You will need an NFS-server for your Windows NT MDC computer (that is, InterDrive, NFS Maestro, DiskShare, or Chameleon UNIX Link 97).

7.2.2 Prerequisites for a Windows NT MDC

This section lists the actions which typically need to be done before installing Tivoli SANergy.

1. Power off all other computers.

Power off all additional computers, except for the first MDC (Meta Data Controller) Windows NT computer.

2. Configure the LAN Adapter.

Note

Note: Each computer using SANergy must have a unique name. SANergy uses the name to identify the workstation within the workgroup. Enter the name from the network setup program under Identification.

3. Install and Configure the Storage Adapter.
4. Install the Shared Storage.
5. Reboot the computer.
6. Set up Partition and Volume Configuration.

   The MDC must be able to see the actual disk system at this point.

   The importance of this step is **planning**. Plan your environment and decide the best way to configure the storage. How many volumes are there to be in the configuration? Do you want striping? One tip is to create many volumes to have greater flexibility.

Install and set up SANergy as MDC on NT

- **Install the SANergy code**
  - 'SANergy Key.txt' contains product Key
  - Install program files, SNMP and docs
- **Select managed busses from list**
- **Select or deselect volumes**
- **Define volumes as shared NT volumes using standard NT tools**

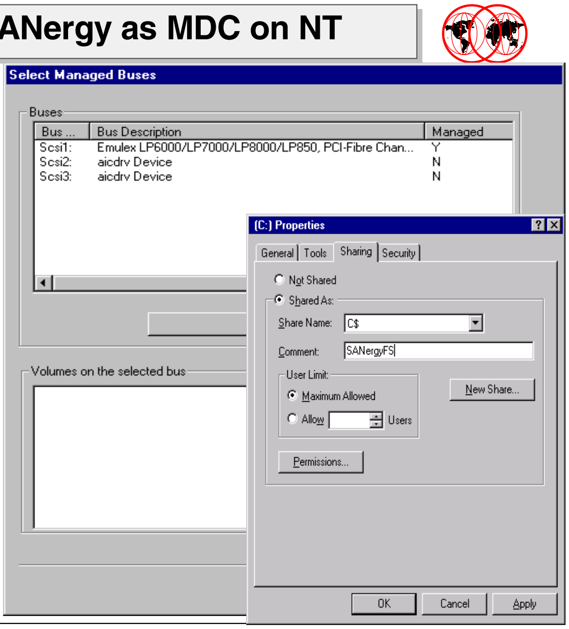
Select Managed Buses

Buses

Bus ...	Bus Description	Managed
Scsi1:	Emulex LP6000/LP7000/LP8000/LP850, PCI-Fibre Chan...	Y
Scsi2:	aicdrv Device	N
Scsi3:	aicdrv Device	N

Volumes on the selected bus

(C:) Properties

○ Not Shared

● Shared As:

Share Name: C$

Comment: SANergyFS

User Limit:

● Maximum Allowed

○ Allow Users

7.2.3 Installation of Tivoli SANergyFS product on NT

This section shows the steps to implement the Tivoli SANergyFS product on a Windows NT system. There may be a variety of computers to share the storage in the SAN, but it is recommended to configure a single computer at a time.

You need to start by configuring the MDC computer first and do one at a time. A single MDC can be set up for all volumes, or an MDC can be spread among several systems. From the MDC computer, use regular Windows NT tools to configure, stripe, and format the shared storage. Configure security and access control parameters at this computer also.

The steps to install the Tivoli SANergyFS MDC will be the same for the Tivoli SANergyFS client. The difference, when installing a client, will be the MDC that has already been installed (by following the instructions below). This will show in the *Volume Assignment* window.

7.2.3.1 SANergy software installation

To start installing the Tivoli SANergyFS MDC, follow these steps:

1. Start by inserting the CD into the server where the MDC will reside. A SANergy setup window will appear.
2. Click the **Next >** button, which will take you to the “Choose Destination Location Window”

   In this case, we chose the default `C:\Program Files\SANergy\SANergyFS`.

3. Click the **Next** > button, which will take you to the “Select Program Folder Window”.

   In this case, we chose the default SANergy

4. Click the **Next** > button, which will take you to the “User Information Window”

   Type in the relevant information for Name and Company. The license key is in a file named `SANergy Key.txt`.

   Cut and Paste this string to the Registration String window.

   Click the **Next** > button.

Note

The Registration String shown in the graphic is a sample key only.

5. The Select Component Window is the next screen as shown in Figure 51 and we selected all the Components: Program Files, SNMP, and Docs.

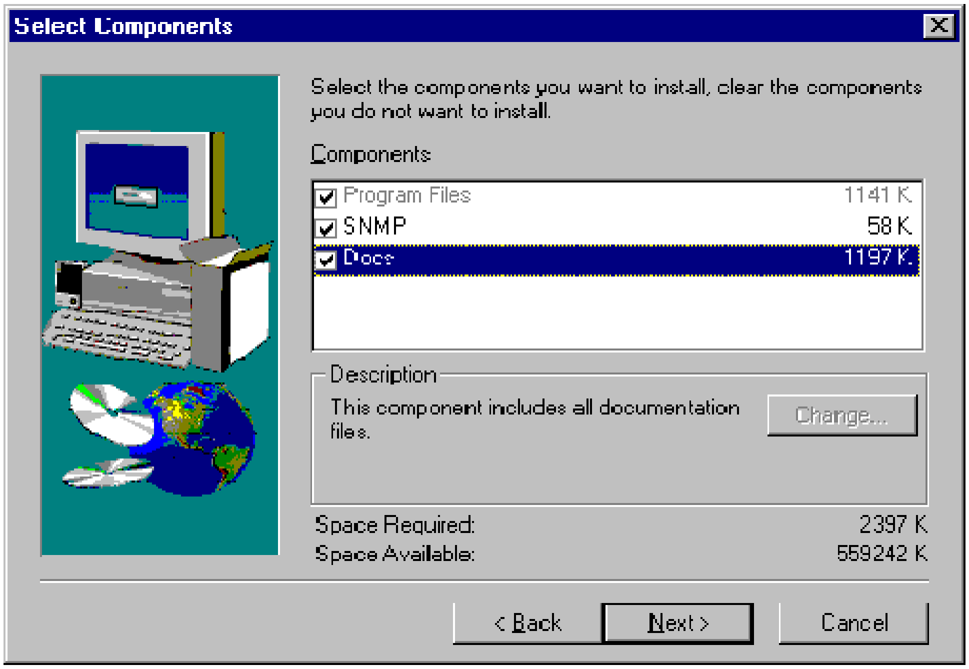

Select Components

Select the components you want to install, clear the components you do not want to install.

Components

☑ Program Files	1141 K
☑ SNMP	58 K
☑ Docs	1197 K

Description

This component includes all documentation files.

Space Required: 2397 K

Space Available: 559242 K

Figure 51. SANergy component selection window

Click the **Next** > button. The “Start Copying Files” window appears.

Click the **Next** > button again.

Select Managed Buses

- Mark SAN bus (here "Scsi1") and select "Change" button
- Click on "Next >" to get Volume Assignment Window
- Takes some time as SANergry is polling the volumes

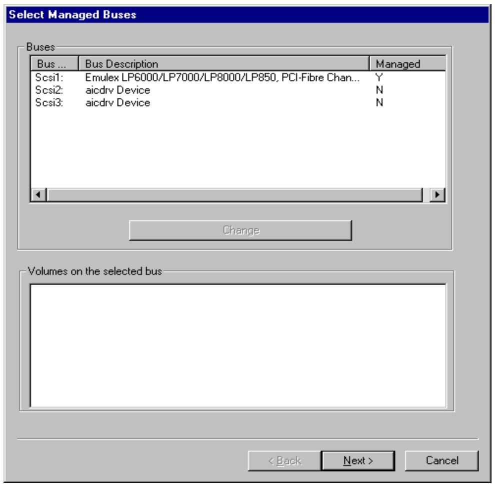

Bus ...	Bus Description	Managed
Scsi1:	Emulex LP6000/LP7000/LP8000/LP850, PCI-Fibre Chan...	Y
Scsi2:	aicdrv Device	N
Scsi3:	aicdrv Device	N

Click here for optional figure #

YRDDPPPPUUU

7.2.3.2 Select Managed Buses

Select the only the bus that is shared and click on the “Change” button. The Managed Field updates to Y to indicate that the bus is now being managed by Tivoli SANergyFS.

Once the “Next >” button is clicked, the “Volume Assignment” window should appear.

Note: At this point, do not worry if this takes some time. Tivoli SANergyFS gave no indication as to what was happening, but we assumed that the volumes were being polled.

Volume Assignment window

- Shows CARTEGENA as MDC for the volumes displayed
- Select "Finish" button to confirm
- Computer will automatically reboot

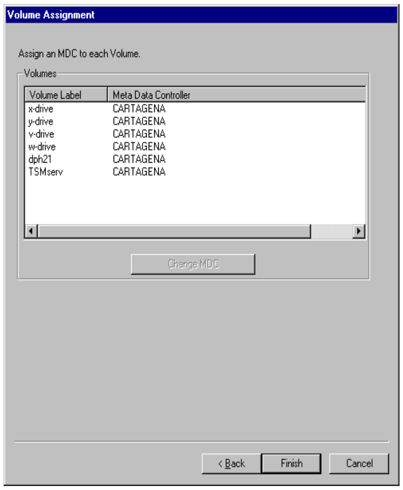

Assign an MDC to each Volume.

Volumes

Volume Label	Meta Data Controller
x-drive	CARTAGENA
y-drive	CARTAGENA
v-drive	CARTAGENA
w-drive	CARTAGENA
dph21	CARTAGENA
TSMserv	CARTAGENA

Click here for optional figure #

YRDDPPPPUUU

7.2.3.3 Volume Assignment window

By default, the first (and maybe the only) MDC is the present owner of all the volumes in your configuration. In this case, our MDC was CARTEGENA, a Windows NT 4 server on a Pentium Pro 200.

It is possible at this point to change the MDC owner if required. However, as this was the first system to install the Tivoli SANergyFS software, we did not need to do this.

Note

When the Tivoli SANergyFS software was installed as a SANergyFS client on our other Windows NT system, this Window automatically showed CARTEGENA as the MDC.

When the configuration is complete, click the “**Finish**” button. The computer will automatically reboot.

Changes to the Tivoli SANergyFS configuration can be made by using the SANergy Configuration Tools. This is detailed later in this chapter.

Initially, once the “**Finish**” button has been clicked, then the volumes must be set up for sharing.

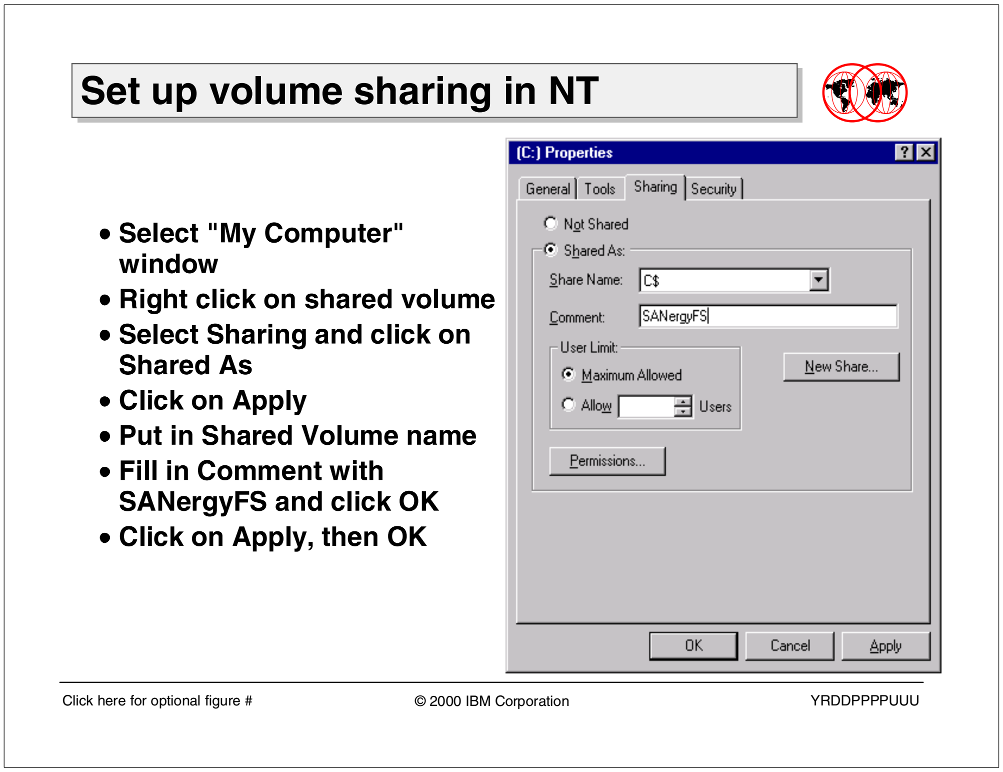

Set up volume sharing in NT

- Select "My Computer" window
- Right click on shared volume
- Select Sharing and click on Shared As
- Click on Apply
- Put in Shared Volume name
- Fill in Comment with SANergyFS and click OK
- Click on Apply, then OK

(C:) Properties

○ Not Shared
● Shared As:
Share Name: C$
Comment: SANergyFS
User Limit:
● Maximum Allowed
○ Allow Users

Click here for optional figure #

YRDDPPPPUUU

7.2.3.4 Set up volume sharing in NT

You need to set up volume sharing on the MDC for all shared volumes in the SANergyFS environment.

1. Select “My Computer” window.
2. Right-click on the shared volume.
3. Select Sharing and click on “Shared As”.
4. Click Apply.
5. Type in the shared volume name.
6. Fill in the comment with SANergyFS and click OK.
7. Click on Apply, then click OK.

Once this step has completed, the Tivoli SANergyFS **clients** need to MAP to the shared volumes and receive “fused access”. This is a SANergy term meaning that the volume is in the SAN. Fused access is controlled by the SANergy MDC.

Tivoli SANergy NT client setup

- **Follow same code installation steps as on MDC**
- **SANergy will automatically recognize MDC of the shared volumes**
- **Map the shared volumes using standard NT Explorer**
- **Be AWARE of mapping**

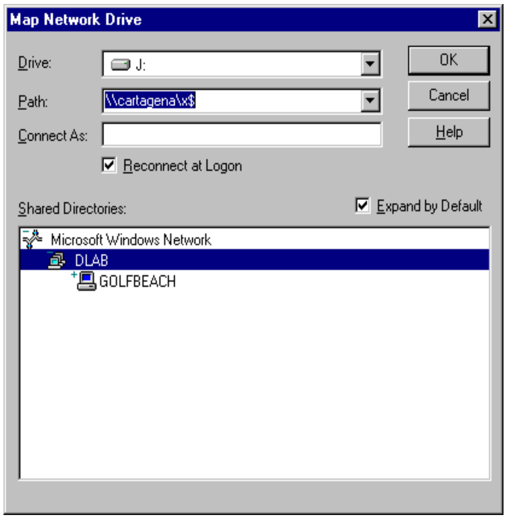

Map Network Drive

Drive: J:

Path: \\cartagena\x$

Connect As:

☑ Reconnect at Logon

Shared Directories:

☑ Expand by Default

Microsoft Windows Network
DLAB
GOLFBEACH

7.2.4 Tivoli SANergy Windows NT client setup

Meta Data Controller (MDC) computers are also SANergy clients. They can access their own MDC volumes as well as other volumes served by different MDC computers. A Windows NT Server can just be a SANergy client, an MDC, or both.

At Windows NT boot time, volumes that are part of the shared storage are mounted only on the server that is the MDC for that volume. Managed volumes on the MDC appear as Windows NT attached disks.

To set up a Windows NT SANergy client, follow these steps:

1. Configure the LAN adapter.
2. Install and configure the storage adapter.
3. Connect the cables to the Shared Storage. In our case, the storage was on an IBM Fibre Channel Raid Storage Server.
4. Reboot the computer.
5. Install the Tivoli SANergyFS software and reboot the computer. You need to follow the same steps in this chapter for setting up the MDC controller. Tivoli SANergy will recognize that the MDC is already defined.
6. Map the Networked drive to the SANergy client using Windows NT tool in NT Explorer.

Note

The drive mapping should be seriously thought out and planned.

Document the mapping in your environment. To make computers match their mapped volume letters, reboot the server. When mapping to a volume, select the letter matching the MDC's.

Keep a copy of the mapping by issuing either:

`winmsd` for Windows NT

or

`msinfo32` for Windows 2000

This is especially important for use with Tivoli Storage Manager. When using Tivoli Storage Manager in conjunction with Tivoli SANergyFS, it is important to keep the drive mapping consistent. Tivoli Storage Manager keeps track of the volume (filespace name) in its database and has no knowledge that this is a mapped network drive.

For example, we have a Windows NT server, which is a Tivoli Storage Manager Client backing up a mapped drive which is SANergyFS managed, such as “F:” This Windows NT server is also a Tivoli SANergy client.

The “F” drive could become disconnected. An administrator could MAP another drive to this Server using the same drive name “F”. But this is a different drive in the MDC configuration from what was mapped to “F” previously. Tivoli Storage Manager could overwrite data onto that mapped drive with the same name, in a restore operation.

7.3 Tivoli SANergyFS configuration tool

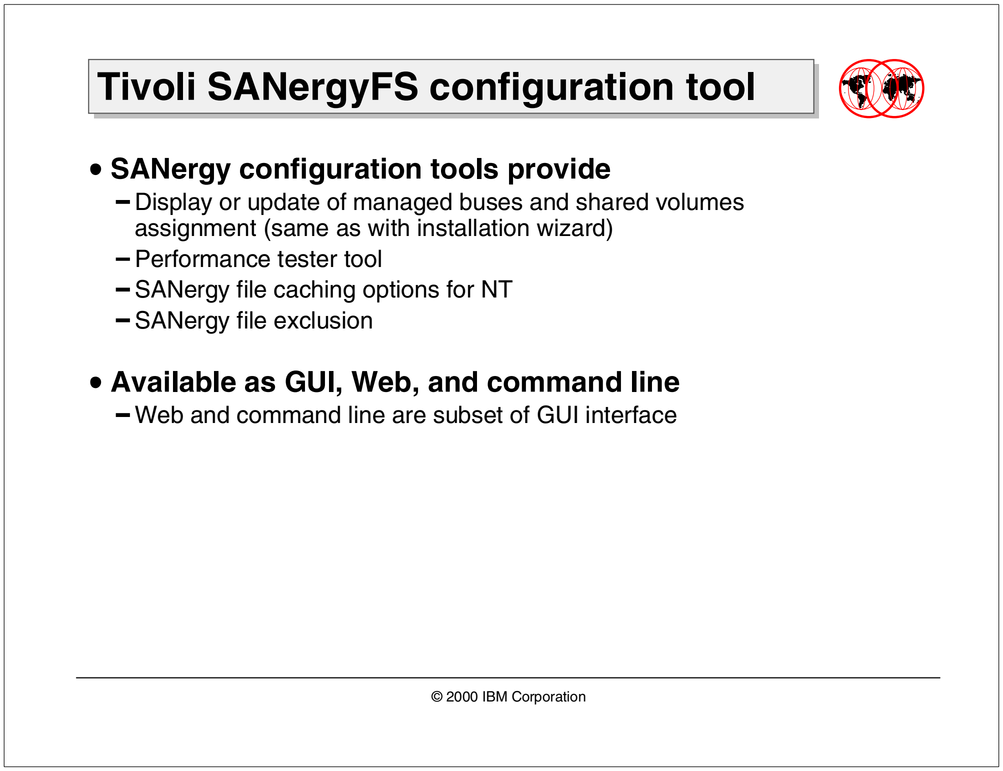

Tivoli SANergyFS configuration tool

- **SANergy configuration tools provide**
  - Display or update of managed buses and shared volumes assignment (same as with installation wizard)
  - Performance tester tool
  - SANergy file caching options for NT
  - SANergy file exclusion
- **Available as GUI, Web, and command line**
  - Web and command line are subset of GUI interface

Tivoli SANergyFS configuration tool has three interfaces to set up the shared storage for file-level access of all the shared volumes. There is the Tivoli SANergy Setup Tool (the GUI interface), the Web tool and the command line interface. In the following section we want to use the Setup Tool, which is detailed in the following sections.

7.3.1 Tivoli SANergy Setup Tool

This configuration tool lets you set up your shared storage for file-level access of all shared volumes. It provides functions for

- Display and update of all known managed buses and volume assignments
- Volume performance testing tool
- Managing SANergy behavior as caching or file exclusion

To access the tool:

```
Select ==> Start/Programs/SANergy/SANergyFS Setup Tool
```

The Tivoli SANergy Setup Tool appears.

7.3.2 Managed Buses folder

Use this folder to view or make changes to the shared bus configuration. To make any changes to the bus configuration follow these steps:

1. Power off all SANergy client servers.
2. In the Bus field, click on the shared bus to be managed.
3. Click on the Change button.

7.3.3 Volume Assignment Folder

The Volume Assignment tab displays a list of all shared volumes and their owners. This is probably the most useful screen. If the MDC field is blank, then the volume has no owner.

It took some time for SANergy to retrieve this information, and there is no progress bar. To change any assignments, follow the following instructions:

1. Warn the users that there is going to be a change in volume ownership. The users must close all applications that are using SANergy and delete all network mappings to that volume.
2. Start the SANergyFS Setup Tool and select the Volume Assignment tab.
3. Click on the volume that is changing ownership and select **Change**. The Change Meta Data Controller appears.
4. Delete the owner and select **OK** and confirm with the **Yes** button.
5. Go to the new MDC and start the SANergyFS Setup Tool there.
6. Click on the Volume Assignment tab.
7. Click on the volume where ownership must change and select **Change**. The Meta Data Controller window appears. Click **OK**.

Follow the previous steps to allow for volume sharing and re-establish any client mappings.

Performance Tester folder

- **Used for performance tests on the drives**
- **Fused Reads/Writes - shows the test is using SANergy**
- **Unfused Reads/Writes - shows the test is not using SANergy, but the LAN**

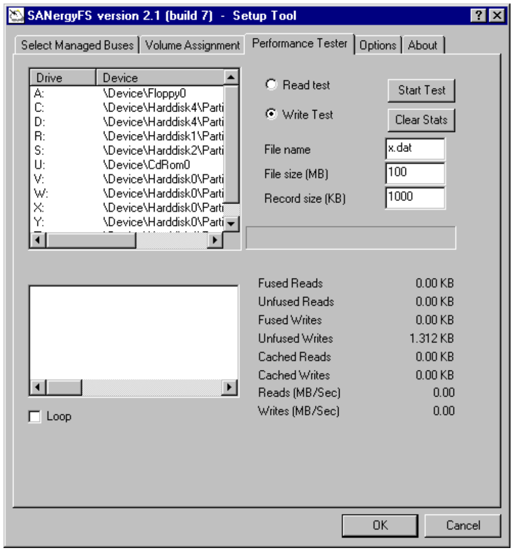

Drive	Device
A:	\Device\Floppy0
C:	\Device\Harddisk4\Parti
D:	\Device\Harddisk4\Parti
R:	\Device\Harddisk1\Parti
S:	\Device\Harddisk2\Parti
U:	\Device\CdRom0
V:	\Device\Harddisk0\Parti
W:	\Device\Harddisk0\Parti
X:	\Device\Harddisk0\Parti
Y:	\Device\Harddisk0\Parti

○ Read test
● Write Test

File name: x.dat
File size (MB): 100
Record size (KB): 1000

Fused Reads	0.00 KB
Unfused Reads	0.00 KB
Fused Writes	0.00 KB
Unfused Writes	1.312 KB
Cached Reads	0.00 KB
Cached Writes	0.00 KB
Reads (MB/Sec)	0.00
Writes (MB/Sec)	0.00

□ Loop

7.3.4 Performance Tester folder

This folder is used to perform read and write performance tests on the shared SANergy volumes.

On the drive field, select the shared volume for the test and choose the required test. If you need to run a read test then a write test has to be performed first. Select the **Start Test** button to begin the performance test.

The bottom of the window updates with the results until the test has completed and the **Clear Stats** button clears the statistics.

Note

Fused Reads and Fused Writes — these are SANergy read and writes.

Unfused Reads and Unfused Writes — this shows that the test is not using SANergy and using the LAN.

The **Loop** feature can be clicked on; this enables the test to run continuously.

Options folder

- Option settings can have a large effect on performance
- Fused Exclusion List - files not to use for SANergy file sharing
- Hyperextension Exclusion List - which files not to hyperextend
- Minimum Fused File Size - minimum file size for files over SAN
- Cache settings are for Windows NT SANergyFS clients only

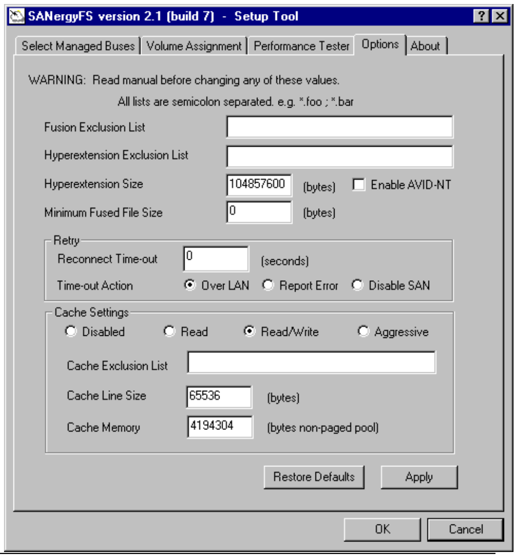

WARNING: Read manual before changing any of these values.

All lists are semicolon separated. e.g. *.foo ; *.bar

Fusion Exclusion List

Hyperextension Exclusion List

Hyperextension Size 104857600 (bytes) □ Enable AVID-NT

Minimum Fused File Size 0 (bytes)

Retry

Reconnect Time-out 0 (seconds)

Time-out Action ● Over LAN ○ Report Error ○ Disable SAN

Cache Settings

○ Disabled ○ Read ● Read/Write ○ Aggressive

Cache Exclusion List

Cache Line Size 65536 (bytes)

Cache Memory 4194304 (bytes non-paged pool)

7.3.5 Options folder

The Options folder lets you control which files are to be included/excluded from Tivoli SANergyFS. Also, the Options folder controls caching for NT clients. It has no effect on other clients or MDCs.

These changes can have a large effect on performance. Some of the tunable boxes are explained below.

- **Fused Exclusion List** tells Tivoli SANergyFS which files not to use for filesharing. This option is recommended only for applications that have an incompatibility problem with SANergyFS.
- **Hyperextension Exclusion List** tells Tivoli SANergyFS which files not to “hyperextend”. When Tivoli SANergy writes to a file and the file needs to grow then SANergy will hyperextend the file size, unless excluded. This is useful for most files and we did not make any exclusions.
- **Minimum Fused File Size** is the minimum size of a file for SANergy filesharing. In the case of very small files, the meta data can be larger than the data itself. These files do not benefit from Tivoli SANergyFS.
- **Retry** tells Tivoli SANergyFS what to do when connectivity is lost. The amount of time is specified in seconds. The **Time-out Action** mode **Over Lan** means the transfer is to use the LAN, this is the default. If **Report Error** is specified then the transfer stops and Tivoli SANergyFS reports an error. **Disable SAN** disables the SAN immediately and allows LAN transfers.

- **Cache Settings** refers to file caching to be performed by Windows NT SANergyFS clients. You can select the following:
  - **Disabled**: No caching. All transactions go directly to the storage medium.
  - **Read-Cache**: Only files opened for read-only mode, without sharing, can by cached.
  - **Read-Write**: Similar to “Read-Cache”, but also applies to files with write and read-write access, with no sharing.
  - **Aggressive**: Caches even those files tagged for sharing, with any access. This is not recommended. Use this mode only if you know your application is at no risk.

    Non-buffered transactions are never cached, so these settings do not apply.
  - **Cache Exclusion List**: Tells SANergyFS which files to not cache.
  - **Cache Line Size**: Tells SANergyFS how many lines of cache to give each file. The typical value is 3.
  - **Cache Memory**: Tells SANergyFS how much memory to use for file caches. This memory is non-paged-pool memory.

7.4 Tivoli SANergyHA

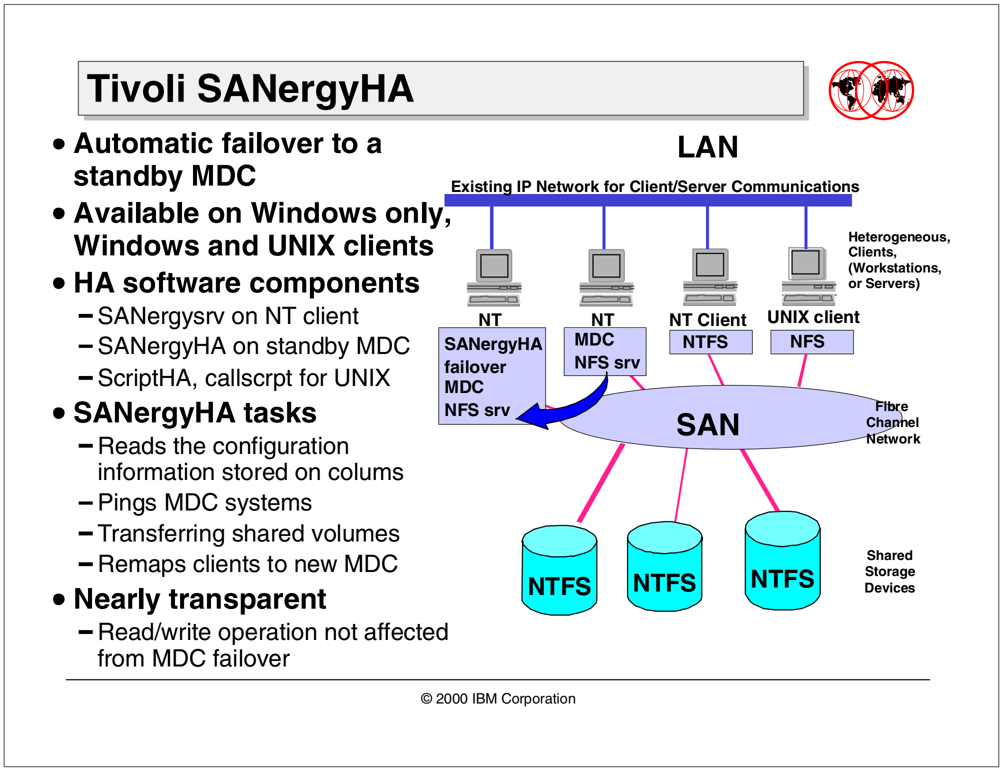

Tivoli SANergyHA

- **Automatic failover to a standby MDC**
- **Available on Windows only, Windows and UNIX clients**
- **HA software components**
  - SANergysrv on NT client
  - SANergyHA on standby MDC
  - ScriptHA, callscrpt for UNIX
- **SANergyHA tasks**
  - Reads the configuration information stored on colums
  - Pings MDC systems
  - Transferring shared volumes
  - Remaps clients to new MDC
- **Nearly transparent**
  - Read/write operation not affected from MDC failover

LAN
Existing IP Network for Client/Server Communications
Heterogeneous, Clients, (Workstations, or Servers)
NT
SANergyHA failover MDC NFS srv
NT
MDC NFS srv
NT Client
NTFS
UNIX client
NFS
SAN
Fibre Channel Network
NTFS
NTFS
NTFS
Shared Storage Devices

© 2000 IBM Corporation

Tivoli SANergyHA provides the ability for failover if your one MDC or multiple MDC's should fail, and allows for the continued use of the storage. If an MDC fails, then SANergyHA will automatically detect the failure and take over as the MDC. The MDC is a system that is the keeper of the meta data information, and therefore should be protected, because if it should fail, then access to the shared volumes would no longer be available. SANergyHA will monitor the health of designated MDCs. If the MDC should fail, the SANergyHA that is assigned to monitor that MDC will automatically detect the failure and take over as the MDC.

The SANergyHA system can be a SANergyFS client, but has the added responsibilty of monitoring any number of MDC's. The SANergyHA software must be installed on all standby systems. The SANergyHA software is currently only available for the Windows NT platform.

7.4.1 SANergy HA components

The SANergyFS utilizes two software components, or services, for failover support.

- The *SANergysvc* service keeps the meta data information about the files that are accessed using SANergyFS. This must have sufficient privileges to remove any active SANergy client mappings to the failed MDC and re-establish them to the new MDC.
- The *SANergyHA* service is part of the software install; it transfers the NT networking shares from the failed MDC to a local *ghost* volume. This service must have sufficient privileges to transfer these shares.

Note

Be sure that both of these services are enabled when logging-in as a user with a real account. This is necessary so that the service can perform failover operations that require administrator privileges

7.4.2 Tivoli SANergyHA operation

The base SANergyFS product stores the MDC name on the target volumes the MDC is controlling, allowing all systems to have access to the information. The SANergyHA software stores a table of standby MDC assignments for these volumes. SANergyHA monitors the MDC constantly by performing a variety of tasks:

- **Configuration Information:** Reads the configuration information stored on the SAN volumes. This includes the current MDC name of each volume and the list of standby systems.
- **Ping:** Performs a ping check of all the MDC's the SANergyHA system is monitoring.
- **Transferring Shares:** After a successful ping check, the standby system downloads the list of shares in the MDC.
- **Remapping Clients:** Downloads the list of connected clients to the MDC. SANergyHA then makes a ghost version of the volume (hard mount), sends a message to all SANergyFS clients, and remaps the clients existing mapped volumes to the Standby MDC.
- **Transparent:** Most clients should be unaware of this failover; it depends on what file system activity is in progress when the MDC fails. Files that are open when the MDC fails will continue to do reads and writes until the new meta data is needed (file close). Applications will experience typical Windows NT network disconnect errors at this time. File locks are not carried over.

If an MDC is rebooted (for installation of new software, for example), and this MDC is being monitored by SANergyHA, then automatic failover will result. If this is not necessary, it is recommended that the SANergyHA service be stopped on each standby machine.

7.5 TSM backups using Tivoli SANergy FS

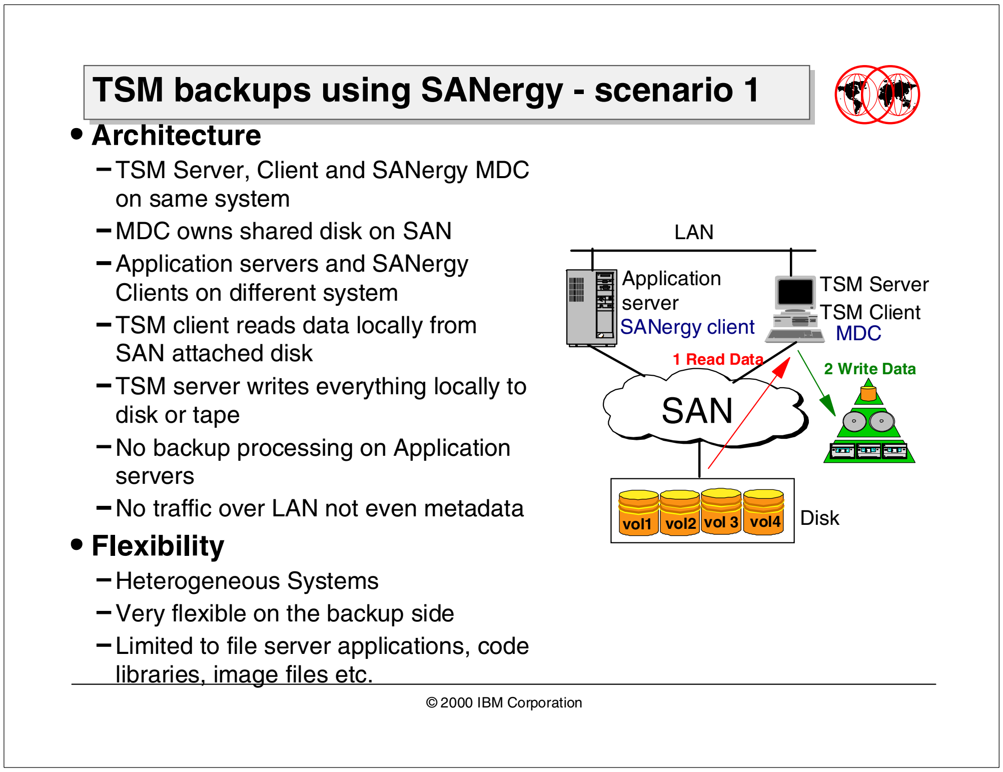

TSM backups using SANergy - scenario 1

- **Architecture**
  - TSM Server, Client and SANergy MDC on same system
  - MDC owns shared disk on SAN
  - Application servers and SANergy Clients on different system
  - TSM client reads data locally from SAN attached disk
  - TSM server writes everything locally to disk or tape
  - No backup processing on Application servers
  - No traffic over LAN not even metadata
- **Flexibility**
  - Heterogeneous Systems
  - Very flexible on the backup side
  - Limited to file server applications, code libraries, image files etc.

Tivoli SANergy, when used in conjunction with Tivoli Storage Manager, supports both LAN-free and serverless backup across a SAN. LAN-free and serverless backup allows the data traffic to be offloaded to the SAN, instead of moving the data through file servers over the LAN. This reduces LAN and server overhead, minimizing the impact of backups on the efficiency of the network. Similarly, when using Tivoli SANergy, data migration from legacy storage on servers to new storage on the SAN has no impact on the LAN performance.

In the following sections, we will discuss two possible configuration scenarios.

7.5.1 TSM backups using SANergy — Scenario 1

Applications that depend on file-served data can be backed up by placing a Tivoli Storage Manager client and server on the file server (MDC) as illustrated above. This configuration allows various heterogeneous SANergy client applications systems to share the SAN mounted volumes. This is a likely configuration for customers who currently have SANergy installed and are looking for a backup/archive solution.

In this example, the application hosts are running SANergy client code, are mounting the disks they need, and are sharing data among themselves. The TSM server and client are running on another processor that is also a SANergy MDC. This configuration is appropriate for heterogeneous environments, where we don't have an MDC on all of the platforms, and/or there is a real need to share data among the hosts.

Because the SANergy MDC "owns" the volumes, it has access to all the data from the other systems, which means the Backup-Archive client can back up all of the information it needs for each system's data. Restores are simplified, since they are local to the server.

The main issues revolve around backing up on-line databases. Oracle will not allow this to happen without the use of Oracle Parallel Server. All databases that will allow dynamic backups are susceptible to database corruption upon the restore. One possible workaround would be to have the connect agent on the machine that is the DB server. However, this does not get the "storage" traffic off the LAN. The SAN would simply be utilized for file sharing.

Another potential issue involves multiple backups of open files. If a client currently has a file open and it is backed up, you could potentially get three copies of the same file backed up from each client accessing the file. This occurs automatically with Windows, but the issue is magnified with SANergy File Sharing.

The system that will be set up as the MDC and TSM server should have adequate processing power. The application servers will be using this MDC to get access to their data. The TSM client and server will be using the system to back up the data in a consistent format. All these operations use system resources and need a system powerful enough to handle the workload. The factors mentioned above may vary for different environments, depending on the size and type of data, the number of application servers accessing the MDC, the attached backup devices, and so on. Therefore, it is advisable to consult your solution provider before sizing the system.

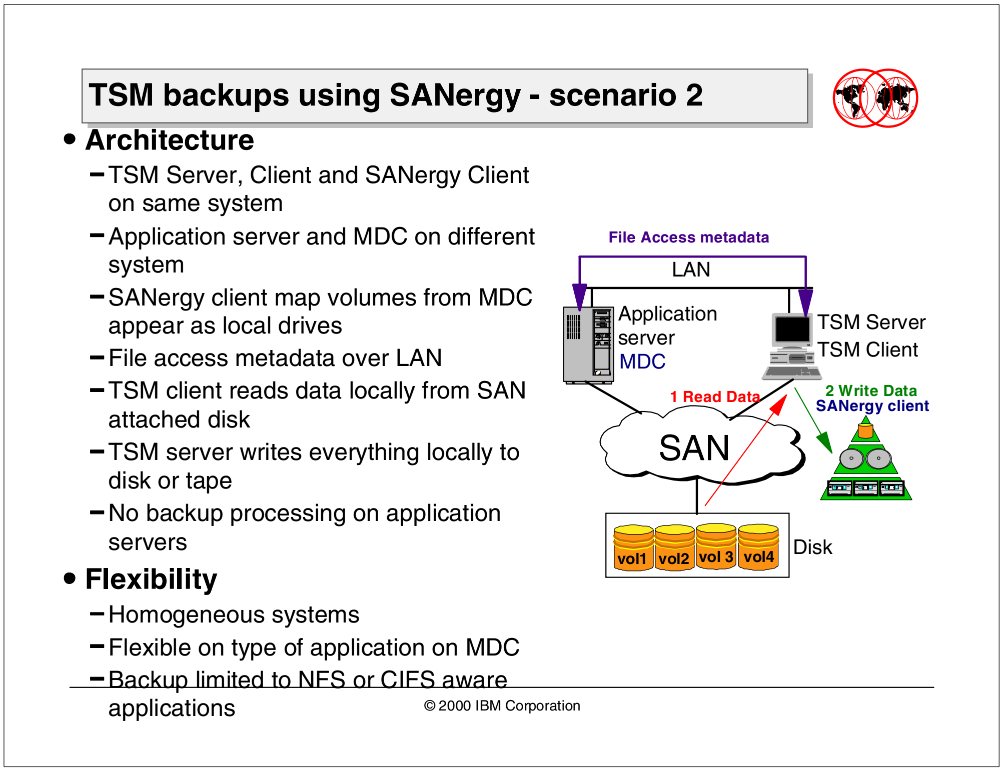

TSM backups using SANergy - scenario 2

- **Architecture**
  - TSM Server, Client and SANergy Client on same system
  - Application server and MDC on different system
  - SANergy client map volumes from MDC appear as local drives
  - File access metadata over LAN
  - TSM client reads data locally from SAN attached disk
  - TSM server writes everything locally to disk or tape
  - No backup processing on application servers
- **Flexibility**
  - Homogeneous systems
  - Flexible on type of application on MDC
  - Backup limited to NFS or CIFS aware applications

© 2000 IBM Corporation

7.5.2 TSM Backups using SANergy — Scenario 2

Tivoli Storage Manager and SANergy can be used to achieve a backup where the I/O is LAN-free, and also, no application server resources are used. For applications with NFS or CIFS mounted volumes, SANergy can be used to remove the backup traffic from the LAN.

Application servers run SANergy, and each server is configured as an MDC. These systems are sharing data between them, so we want the SANergy configuration that has the least overhead for them to access their data. These servers read their data as locally attached drives (no-NFS mounts are required).

The system running the TSM server and client runs SANergy code in client mode in order to have shared access to the other system's files. This system NFS mounts the volumes and then uses the Backup-Archive client to back up the data over the SAN using the SANergy technology. During the backup, the TSM client and SANergy client do LAN-free reads of application files from the file server with only the meta data flowing over the LAN.

Since the TSM client is mounting the files, this works best in a homogeneous environment. NT would probably be the best, although it should also work in a SUN environment, as long as the customer is happy with using the UFS file system.

The main issues revolve around database backups. A customer will NOT be able to back up a database dynamically. Oracle will not allow this backup to take place. The only way this scenario would work with Oracle is for the TSM connect agent to reside on the MDC that owns the Oracle database. Oracle Parallel Server and Oracle Fail Safe (Windows NT only) will allow the TSM Backup-Archive client to back up an off-line copy of the database while the on-line copy can continue to be updated. This information is true for any database package available.

Another potential issue involves multiple backups of open files. If a client currently has a file open and it is backed up, you could potentially get three copies of the same file backed up from each client accessing the file. This occurs automatically with Windows, but the issue is magnified with SANergy File Sharing.

There is no TSM client on the application server. The SANergy client maps the volumes from the MDC, and TSM sees them as locally attached drives/file systems.

Note

The SANergy client on the system with the TSM server maps the volumes from the MDC on its local system. It is important to keep a copy of those maps, because to TSM, these volumes appear as locally mounted volumes.

During restores, TSM has no information on the MDC from which the volume was mapped. Correct volume mapping from SANergy client (TSM server) to MDC has to exist beforehand in order to restore to the proper place.

Another point to note is that the SANergy MDC does utilize some CPU resources, and therefore will have some impact on the application server.

Chapter 8. Tivoli Storage Manager SAN exploitation

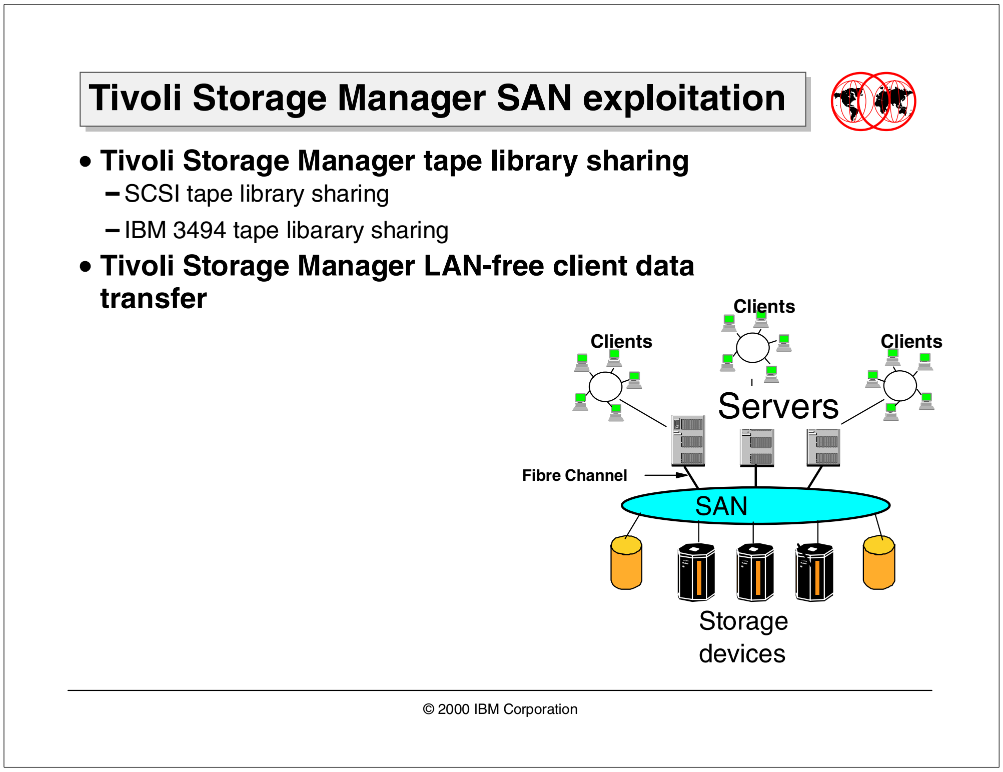

Tivoli Storage Manager SAN exploitation

- **Tivoli Storage Manager tape library sharing**
  - SCSI tape library sharing
  - IBM 3494 tape libarary sharing
- **Tivoli Storage Manager LAN-free client data transfer**

© 2000 IBM Corporation

Storage Area Network (SAN) is a new architecture that puts storage on a separate dedicated network to allow businesses of all sizes to provide access to important data, regardless of operating systems, as a significant step towards helping customers cope with the explosive growth of information in the e-business age.

A SAN is a dedicated network used for data movement or access purposes. This type of network is contrasted with the typical network, which in addition to data access (file serving), is used for communications functions like e-mail, terminal connection, and application program communication.

The intent of a storage area network is to isolate shared data access functions from other functions to gain performance advantages while allowing the storage sharing and distances available with typical network file sharing.

SAN topology enables "any-to-any" connections and consolidation among servers and storage systems in a networked environment. SANs de-couple ownership of storage resources from the servers.

SAN architectures facilitate storage consolidation, clustering, high availability, fault tolerance, remote management and flexible topologies, and can help eliminate scalability limitations and bandwidth concerns inherent in current environments.

The industry considers Fibre Channel as the architecture on which most SAN implementations will be built.

Fibre Channel is a technology standard that allows data to be transferred from one network node to another at very high speeds. Current implementations that transfer data at 100 MB/sec 200 and 400 MB/sec data rates have already been tested. This standard is backed by a consortium of industry vendors and has been accredited by the American National Standards Institute (ANSI). Many products are now on the market that take advantage of FC's high-speed and high-availability characteristics.

Fibre Channel architecture is often referred to as the Fibre version of SCSI. Fibre Channel is an architecture used to carry IPI traffic, IP traffic, FICON traffic, SCSI traffic, and potentially traffic using other protocols, all at the same level on the standard FC transport. FICON is expected to be the standard protocol for S/390, and FCP is expected to become the standard protocol for non-S/390 systems, both using Fibre Channel architecture to carry the traffic.

Tivoli Storage Manager provides a number of important SAN solutions continuing to exploit the technology. Many other solutions are in development and will be rolled out as SANs mature. All these solutions address the need for efficient and reliable data protection.

This chapter covers the following solutions:

- **Tivoli Storage Manager tape library sharing**

  The Tape Resource Sharing feature of Tivoli Storage Manager allows administrators to pool tape resources for use by many Storage Manager servers running on heterogeneous platforms. This can improve backup and recovery performance, and tape hardware asset utilization.

- **Tivoli Storage Manager LAN-free client data transfer**

  At the direction of the Tivoli Storage Manager server, tape storage pools are dynamically allocated to Tivoli Storage Manager clients. This allows backup information to be sent across the SAN directly from the client to the server storage pools. In this way the data path completely bypasses the LAN and the Tivoli Storage Manager server. A lot less IP communications bandwidth is used, and service levels for end users are improved. Tivoli Storage Manager's first implementation of LAN-free client data transfer protects Microsoft Exchange databases using Tivoli Data Protection for Exchange and SAP's R3 implementations using Tivoli Data Protection for R3.

8.1 SCSI tape library sharing

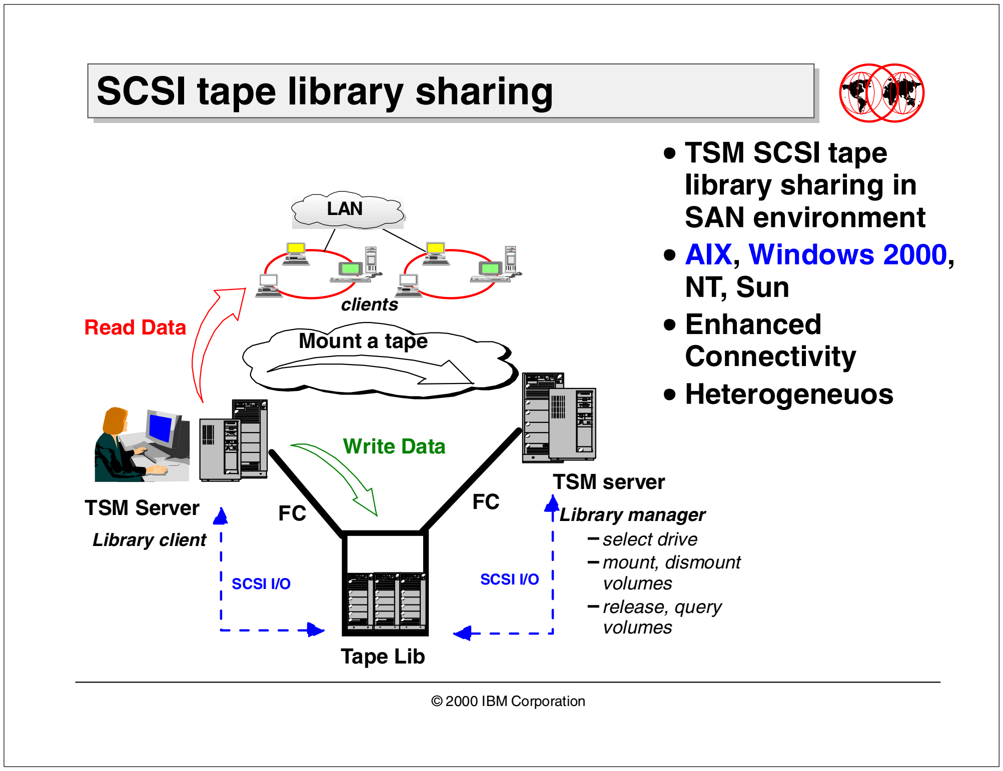

SCSI tape library sharing

- TSM SCSI tape library sharing in SAN environment
- AIX, Windows 2000, NT, Sun
- Enhanced Connectivity
- Heterogeneuos

© 2000 IBM Corporation

With Tivoli Storage Manager server 3.7 and above, multiple Tivoli Storage Manager servers can dynamically share library volume and tape drive resources of one physical SCSI connected tape library. The hosts can thus maintain high speed connections to the same devices through the SAN fabric. Applications which immediately benefit from this include backup and restore. The effect is pronounced for environments with large amounts of data to back up over shrinking windows of time and constrained LAN bandwidth.

Using the Fibre Channel technology of SANs, distances between the tape library and the Tivoli Storage Manager server(s) can be extended as well. Tape libraries can reside at an alternate location. Customers can use electronic vaulting for a quick and efficient way to protect against and recover from a disaster situation. Finally, since a reliable tape drive is quite an expensive device, tape sharing can be a very important economic factor. This section introduces the new platform and device support for SCSI tape library sharing.

For more details on how to configure Tivoli Storage Manager SCSI tape library sharing on a SAN, consult the *Tivoli Storage Manager V3.7: Technical Guide,* SG24-5477.

8.1.1 New device and platform support

AIX and Windows 2000 have been added as new platforms that are supported in the TSM SCSI tape library sharing solution.

Tape Library sharing for SCSI libraries is now supported across the following platforms:

- IBM AIX
- Windows NT
- Sun Solaris
- Windows 2000

Many elements are involved in a SAN solution. These include:

- Host operating system
- Cabling
- SAN topology
- Gateway, bridge, hub and switch
- FC adapter hardware, microcode, and device drivers
- Device adapter hardware, microcode, and device drivers
- Related hardware (tape library; disk array), microcode, and device drivers

All the elements must be compatible with each other in order to have working end-to-end communications.

As with other innovative technologies in their infancy, SAN's suffer from incompatibility issues. This has the unfortunate consequence that not just any SAN device will work with any other SAN device. IBM only supports sets of SAN products which have been verified to be working properly together. This is established after a series of comprehensive tests conducted in IBM interoperability laboratories.

The current list of supported combinations for Tivoli Storage Manager SCSI tape library sharing solution is available at:

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

You and your solution provider should use this information during the design phase to ensure a smooth implementation.

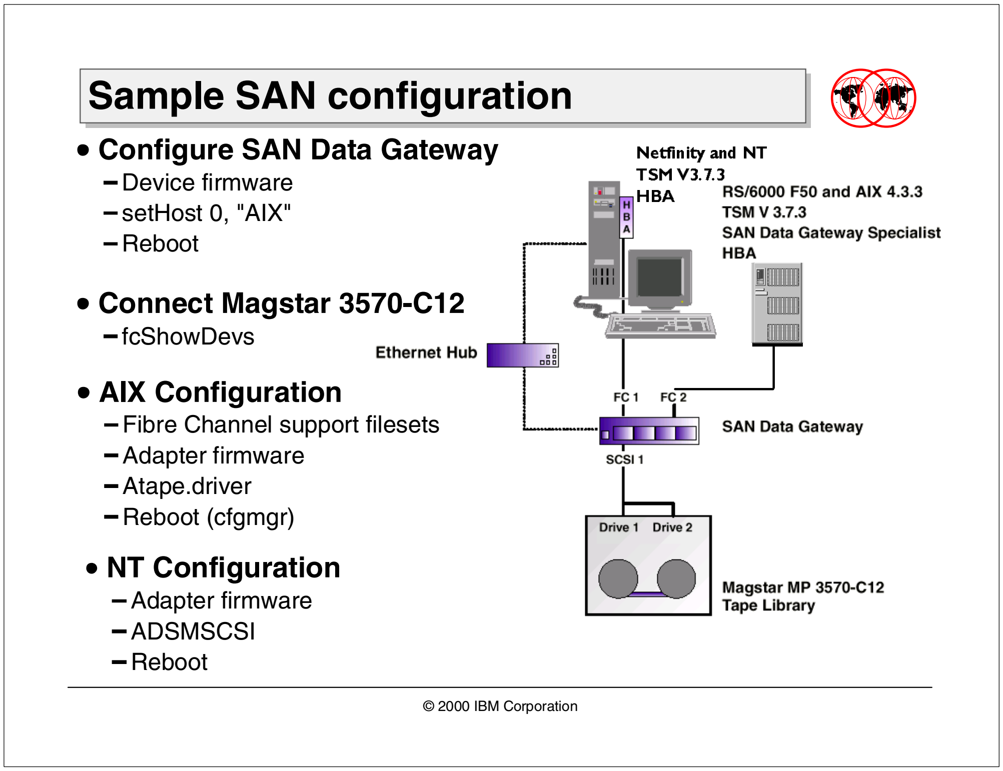

Sample SAN configuration

- **Configure SAN Data Gateway**
  - Device firmware
  - setHost 0, "AIX"
  - Reboot
- **Connect Magstar 3570-C12**
  - fcShowDevs
- **AIX Configuration**
  - Fibre Channel support filesets
  - Adapter firmware
  - Atape.driver
  - Reboot (cfgmgr)
- **NT Configuration**
  - Adapter firmware
  - ADSMSCSI
  - Reboot

© 2000 IBM Corporation

8.1.2 Sample SAN configuration

In this section we attempt to illustrate how to set up AIX and NT device drivers to work in a SAN environment with an IBM 3570 C12 connected using an IBM SAN Data Gateway.

8.1.2.1 SAN Data Gateway configuration

The first step is to enable the SAN Data Gateway so it can connect to AIX and/or Windows NT.

The SAN Data Gateway comes with the ethernet port disabled, so the only way to connect to it initially is through its service port. When you connect to the service port you get a command line console which lets you enter service port commands.

In our case, we use a crossed “null modem” cable to connect the service port to a 9-pin serial port (COM1) on a PC running NT 4.0. We configured a hyperterminal session on NT with the following settings: Bits per second - 19200, Data bits - 8, Parity - None, Stop bits - 1 and Flow Control - Xon/Xoff.

Open the terminal session, press Enter, and you get a command prompt. The network name of the SAN Data Gateway is set to `Router` by default. You can change it by using the `HOSTNAMESET` command:

```
Router > hostnameSet “Gateway”
```

The first step is to check the firmware level of the SAN Data Gateway. We issued the `VERSION` command as shown below:

```
Router > version

SAN Router Version 0340.10 Built Mar 31 2000, 10:17:21
VxWorks (for Pathlight (i960RD)) version 5.3.1.
Kernel: WIND version 2.5.
value = 26 = 0x1a
```

Your firmware level should be at 03.40.10 or above (The current level is 03.40.16). You can download the latest firmware for SAN Data Gateway including installation information, from:

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

If your firmware is not at the latest level, then you should upgrade it. The firmware can be upgraded using the service port.

Following are the steps you would need in order to upgrade the firmware of the SAN Data Gateway using the console session on the service port.

1. From the hyperterminal window, select `Transfer` and then `Send File` .
2. From the `Send File` dialog box, enter the path and filename where the firmware file is located. You should already have downloaded the file from the Web. In the `Protocol` field, select `Zmodem` (default option) and then click the **Send** button. It takes several minutes for the file to be transferred. The window will show the current status (see Figure 52).

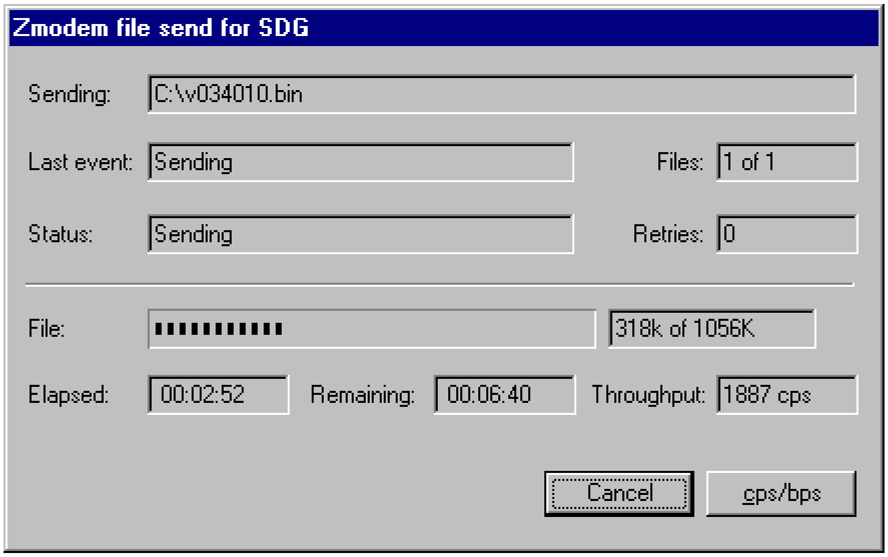

Sending: C:\v034010.bin

Last event: Sending

Files: 1 of 1

Status: Sending

Retries: 0

File: 318k of 1056K

Elapsed: 00:02:52

Remaining: 00:06:40

Throughput: 1887 cps

Figure 52. SDG firmware file transfer using hyperterminal

3. After the file is completely transferred, you will see a `Firmware Update in Progress` message on the service terminal console.
4. Wait for the `Firmware Update Complete` message, as shown in Figure 53, to appear, indicating successful completion.
5. Reboot the SAN Data Gateway for the update to take effect.

```
Router > reboot
```

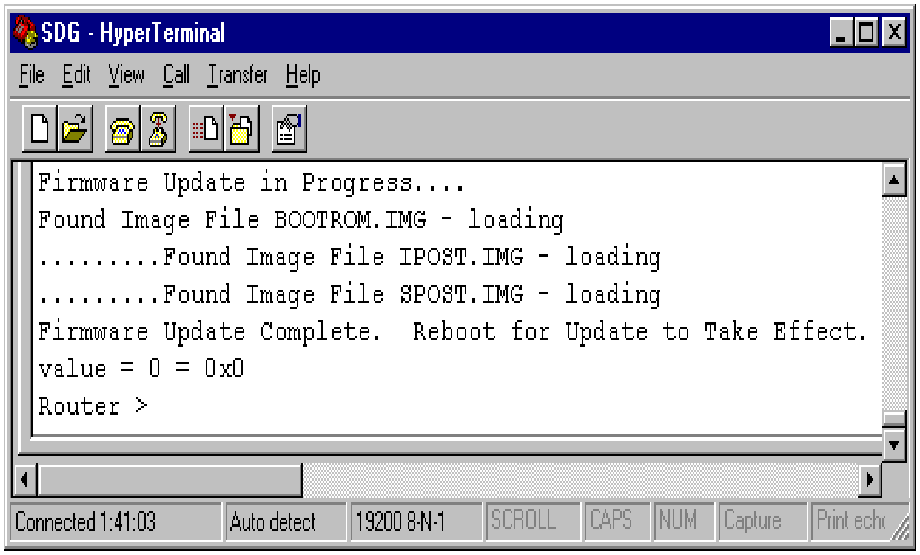

```
Firmware Update in Progress....
Found Image File BOOTROM.IMG - loading
........Found Image File IPOST.IMG - loading
........Found Image File SPOST.IMG - loading
Firmware Update Complete.  Reboot for Update to Take Effect.
value = 0 = 0x0
Router >
```

Figure 53. SAN Data Gateway firmware update

All the settings and installation/configuration instructions for the IBM SAN Data Gateway are documented in the *IBM Storage Area Network Data Gateway Installation and User's Guide*, SC26-7304, and in the *IBM Storage Area Network Data Gateway Router Service Guide*, SY27-7614.

After the SAN Data Gateway (SDG) has the correct level of firmware on it, there is one more command you need to issue in order to enable it to participate successfully in an AIX/NT environment. Below is the SETHOST command we issued.

```
Router > setHost 0, "AIX"
```

The syntax of the command is `setHost [port], "OS"`. The `setHost` command sets the operating system type for the specified SAN interface. This command provides some customizing in the way the SDG is presented to the particular OS. If the port is `0`, it applies to all SAN connections; otherwise, it is applied only to the SAN connection on the specified interface. The default setting is `NT`. The OS can be one of `NT`, `AIX`, or `Solaris`.

Note

We found out that setting the OS to **NT** (default setting) does not work well with AIX systems, while setting the OS to **AIX** works equally well with NT systems. If you have an all-NT environment, set the OS to NT; otherwise, if you have a mixed AIX/NT environment, set it to **AIX**.

8.1.2.2 SAN Data Gateway Magstar connectivity

After configuring the SDG, we powered it off and physically connected it to a Magstar 3570 tape library. Our Magstar tape library had two 3570 drives with an autochanger. The drives and autochanger were connected to a single SCSI port on the SAN Data Gateway.

Next we powered on the tape library. Once the library initialized itself, we powered on the SDG. It is important to power on the devices in the correct sequence otherwise the SDG may not recognize the tape devices properly.

Once the SDG came up, the 3570 devices were recognized. You can verify this by issuing the `fcShowDevs` command.

```
Router > fcShowDevs
FC 1:
LUN Chan  Id  Lun  Vendor   Product          Rev   SN
-----------------------------------------------------
  0   0    0    0  PATHLGHT SAN Router       328a  21081340156
  2   1    2    0  IBM      03570C12         5424  0000000B1982
  3   1    2    1  IBM      03570C12         5424  0000000B1982
  4   1    3    0  IBM      03570C12         5424  0000000B2070
value = 6 = 0x6
```

The `fcShowDevs` command displays the target devices with both their SCSI addresses and the Fibre Channel LUN assigned on the SDG.

8.1.2.3 Windows NT Fibre Channel configuration

There are a few steps required for Windows NT to work in the SDG environment with a Magstar 3570.

We installed a QLogic 2100F Fibre Channel host bus adapter (HBA) in our NT system. We used it to connect to the Fibre Channel port on the SDG. The first step is to get the latest firmware for the card. We downloaded the latest driver for NT from their Web site: `http://www.qlogic.com`

We exploded the executable 21nt0402.exe in the directory `C:\Qlogic`. As a result another directory `C:\Qlogic\NT` was created. The directory contents of NT included a driver installation script, QLogic QLA2100 PCI Fibre Channel miniport driver, and a readme file. The `readme.txt` file lists the procedure to install the driver for the first time.

1. Click on **Start**.
2. Under “Settings”, click on **Control Panel**.
3. Select “SCSI Adapters” from the “Control Panel”.
4. Select the **Drivers** tab and then click on **Add.**
5. Click on **Have Disk**.
6. Browse to find the correct path for the NT driver and click **OK**. In our case, the correct path was `C:\Qlogic\NT`
7. Choose “QLogic QLA2000/QLA2100 PCI Fibre Channel Adapter” and click **OK.**
8. Select **Yes** to restart system.

It is important to consult the latest readme files that are shipped with the vendor's driver before installation.

After our system came up, we were able to see the Fibre Channel adapter and all the devices attached to it across the SAN Data Gateway. The adapter was listed under "SCSI Adapters" under "Control Panel" as shown in Figure 54.

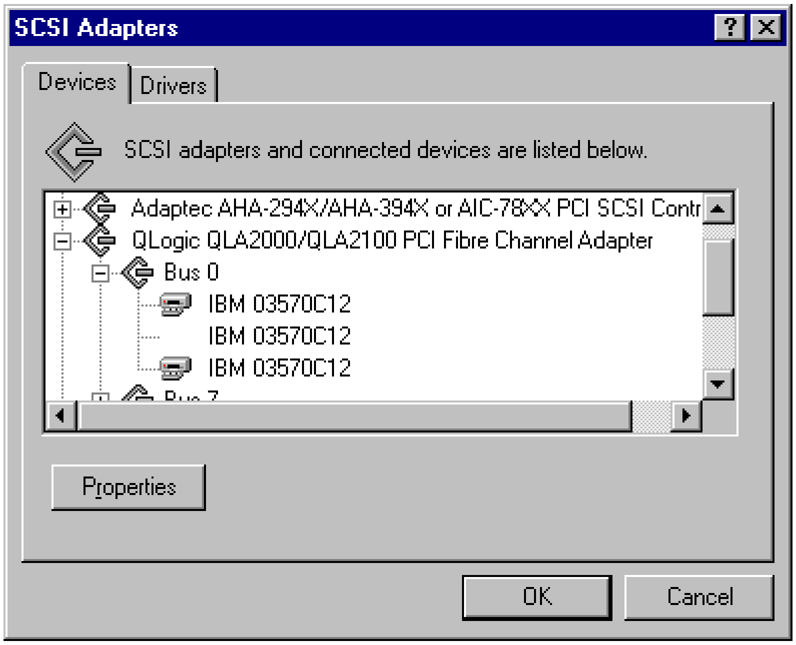
SCSI Adapters

SCSI adapters and connected devices are listed below.

```
⊞ Adaptec AHA-294X/AHA-394X or AIC-78XX PCI SCSI Contr
⊟ QLogic QLA2000/QLA2100 PCI Fibre Channel Adapter
   ⊟ Bus 0
        IBM 03570C12
        IBM 03570C12
        IBM 03570C12
   ⊞ Bus 7
```

Figure 54. SCSI Adapters — Devices

The Magstar tape drives were listed under "Tape Devices" under "Control Panel". However, the drivers for the tape drives were not loaded (refer to Figure 55). This is because Magstar 3570 drives under NT require ADSMSCSI driver for support. In our case although we had the system set up as a TSM server, we had not started the TSM device driver (ADSMSCSI).

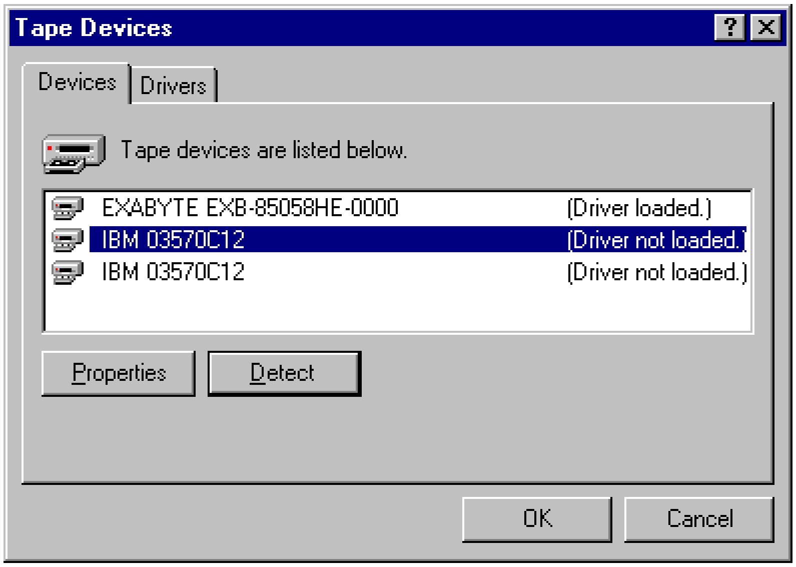
Tape Devices

Tape devices are listed below.

```
EXABYTE EXB-85058HE-0000        (Driver loaded.)
IBM 03570C12                    (Driver not loaded.)
IBM 03570C12                    (Driver not loaded.)
```

Figure 55. Tape Devices — Drivers not loaded

We went into “Control Panel”, then “Devices”, and then selected “ADSMSCSI”. We clicked on **Start** and started it. As soon as the driver started, the tape device drivers were loaded (see Figure 56).

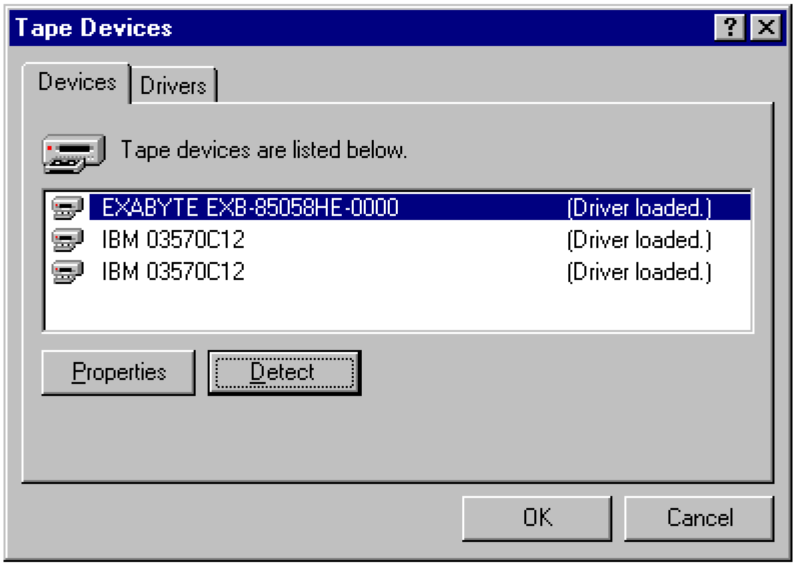

Tape Devices

Tape devices are listed below.

EXABYTE EXB-85058HE-0000	(Driver loaded.)
IBM 03570C12	(Driver loaded.)
IBM 03570C12	(Driver loaded.)

Figure 56. Tape Devices — TSM Device Driver started

The device descriptions (like those displayed in Figure 57) can be seen by clicking on **Properties** for each drive.

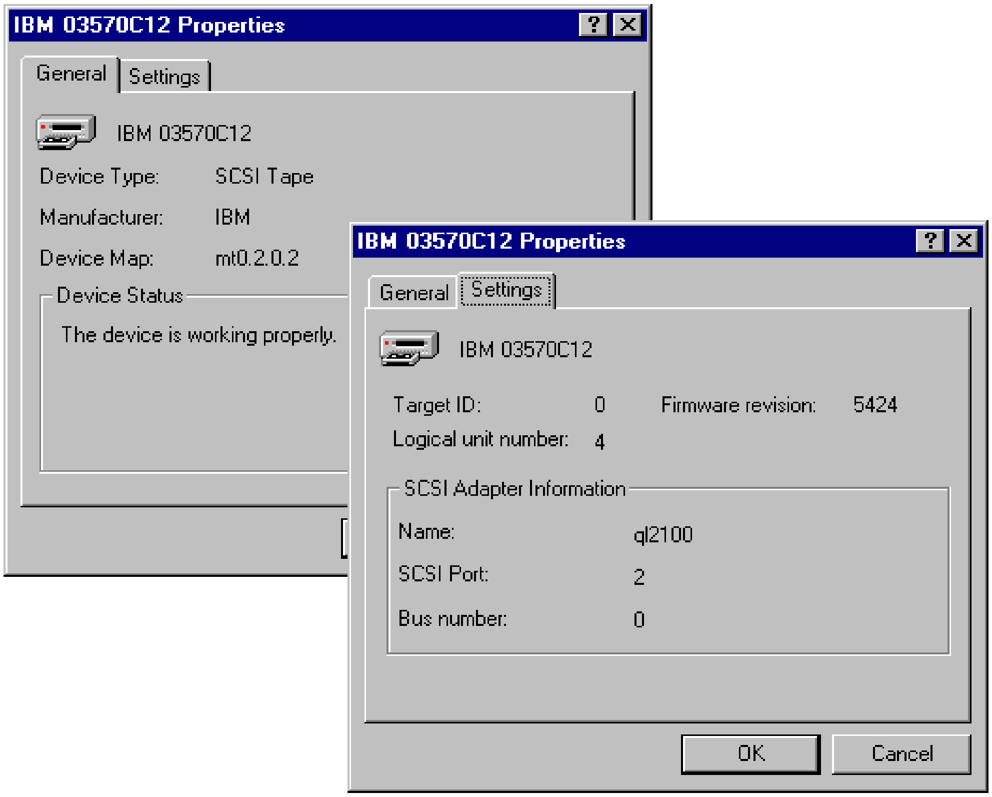

IBM 03570C12 Properties

IBM 03570C12

Device Type: SCSI Tape

Manufacturer: IBM

Device Map: mt0.2.0.2

Device Status

The device is working properly.

IBM 03570C12 Properties

IBM 03570C12

Target ID: 0

Firmware revision: 5424

Logical unit number: 4

SCSI Adapter Information

Name: ql2100

SCSI Port: 2

Bus number: 0

Figure 57. Tape Devices — Properties

We found the NT configuration to be pretty straightforward. The important thing is to make sure you have obtained and installed the latest level of firmware from the vendor for the Fibre Channel adapter, and also that the TSM device driver is running.

8.1.2.4 AIX Fibre Channel configuration

As before, the first step in setting up AIX is to install the Fibre Channel gigabit host bus adapter and related drivers and filesets.

We were running AIX 4.3.3 with the latest updates. The following list shows the installed software list on our system specific for Fibre Channel support in AIX.

```
devices.pci.df1000f7.com   4.3.3.0  C  Common PCI FC Adapter Device
devices.pci.df1000f7.diag  4.3.3.0  C  PCI FC Adapter Device
devices.pci.df1000f7.rte   4.3.3.0  C  PCI FC Adapter Device Software
devices.common.IBM.fc.rte  4.3.3.0  C  Common IBM FC Software
```

These software levels are necessary to support the Fibre Channel gigabit adapter.

After we physically installed the IBM Gigabit Fibre Channel adapter (FC #6227), we powered on the system. The system came up and recognized the adapter as `fcs0`. When we issued the command `lsdev -Cc adapter` we were able to see the following entry for the Fibre Channel adapter:

```
fcs0    Available 10-68    FC Adapter
```

The first thing we did, once the adapter was installed and configured, was to make sure we had the latest level of firmware installed.

We issued the command `lscfg -vl fcs0` to get information on the firmware level for fcs0.

```
fcs0           10-68            FC Adapter
        Part Number.................03N4167
        EC Level....................D
        Serial Number...............KT90601339
        Manufacturer................0010
        FRU Number..................03N4167
        Network Address.............10000000C920D21F
        ROS Level and ID............029022F0
        Device Specific.(Z0)........4002206D
        Device Specific.(Z1)........10020193
        Device Specific.(Z2)........3001506D
        Device Specific.(Z3)........01000909
        Device Specific.(Z4)........FF101270
        Device Specific.(Z5)........029022F0
        Device Specific.(Z6)........061122F0
        Device Specific.(Z7)........071122F0
        Device Specific.(Z8)........10000000C920D21F
        Device Specific.(Z9)........SS2.23
        Device Specific.(ZA)........S1F2.23
        Device Specific.(ZB)........S2F2.22X1
        Device Specific.(YL)........P2-I3
```

If one is using the correct AIX software, there will be `Z9`, `ZA`, and `ZB` items. The microcode level can be determined by looking at `ZB` and removing the second character. Thus, the above shows the microcode level of sf222. If the `lscfg` does not show a `ZB` field, then the software on this machine is down-level and should be upgraded first. The microcode level that is required for SAN Data Gateway support is SF222.x1.

In case your adapter firmware needs to be upgraded, download the microcode and instructions from `http://www.rs6000.ibm.com/support/micro/df1000f7.html`. The microcode file is `df1000f7.bin`. Download the file in the root directory. Unzip the file on AIX by entering the command:

```
/: > chmod +x df1000f7.bin  ./df1000f7.bin
```

You will be prompted for a password (which everyone can obtain from the download Web site mentioned above), after which the files are unpacked into `/etc/microcode`. Issue the command listed below to upgrade the microcode:

```
/: > diag -c -d fcs0 -T"download"
```

Run `lscfg -vl fcs0` to ensure that you have the proper level of firmware installed.

The file `readme.txt` on the Web site below contains a wealth of useful information:

`http://www.storage.ibm.com/hardsoft/products/sangateway/support/cdr/AIX/readme.txt`

This file contains detailed compatibility information and installation instructions and should be consulted before attempting the configuration.

Next we connected the RS/6000 to the SAN Data Gateway. When connecting the Fibre Channel optical cables, if the host and SAN Data Gateway are powered on and properly cabled (optical cables), there will be two green LEDs lit on the rear of the SAN Data Gateway for that port.

Next, we installed the software needed to support the Magstar 3570 device. The Magstar 3570 and 3575 tape libraries require the `Atape.driver` fileset for support in AIX. `Atape.driver` is the IBM AIX Enhanced Tape and Medium Changer Device Driver. For the SAN tape library solution, the driver must be at level 5.0 or better. Our driver level was 5.1.3.0. We obtained the latest version of the `Atape.driver` fileset from `ftp://index.storsys.ibm.com` and installed it.

To install the fileset, go to `smitty install_latest`. Input the name of the directory where you stored the downloaded file. Press F4 to select `Atape` and press `Enter` twice to install the fileset. You would need to reboot the system so the changes take effect. Once our system came up, we were able to see the devices. We issued a `lsdev -Cc tape` command. The output is shown below:

```
rmt0 Available 10-68-01 IBM Magstar MP Tape Subsystem (FCP)
rmt1 Available 10-68-01 IBM Magstar MP Tape Subsystem (FCP)
```

In order to ensure that proper host-to-tape device connection has been established, we ran the `tapeutil` program. We issued a sequence of commands. The numbers below represent the menu options:

- `1 Open drive /dev/rmt0`
- `30 Read/Write test`
- `2 Close drive`
  (necessary to enable further communications via other applications)
- `Q Quit`

8.2 IBM 3494 tape library sharing

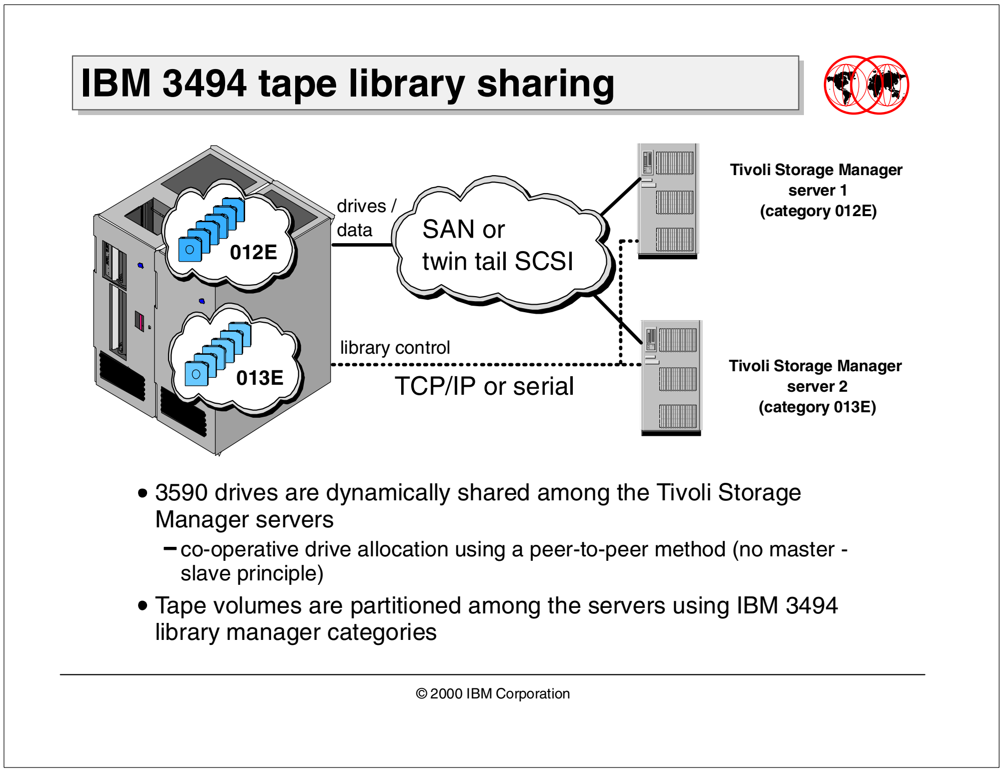

IBM 3494 tape library sharing

012E
013E
drives / data
SAN or twin tail SCSI
library control
TCP/IP or serial
Tivoli Storage Manager server 1 (category 012E)
Tivoli Storage Manager server 2 (category 013E)

- 3590 drives are dynamically shared among the Tivoli Storage Manager servers
  - co-operative drive allocation using a peer-to-peer method (no master - slave principle)
- Tape volumes are partitioned among the servers using IBM 3494 library manager categories

Tivoli Storage Manager SCSI tape library sharing support is implemented for libraries that use SCSI commands to control the library robotics and the tape management. This does not include the IBM 3494.

The IBM 3494 tape library can be shared today among multiple TSM servers on a SAN using features of the IBM 3494, Tivoli Storage Manager, and a supported SAN configuration. For the IBM 3494 product, tape sharing is part of its basic function. The library is shared in a physical way with each system thinking it really owns the entire library. The IBM 3494 library manager ensures that integrity is enforced.

For more information, consult the *Tivoli Storage Manager V3.7: Technical Guide*, SG24-5477, and the *Guide to Sharing and Partitioning IBM Tape Library Dataservers*, SG24-4409.

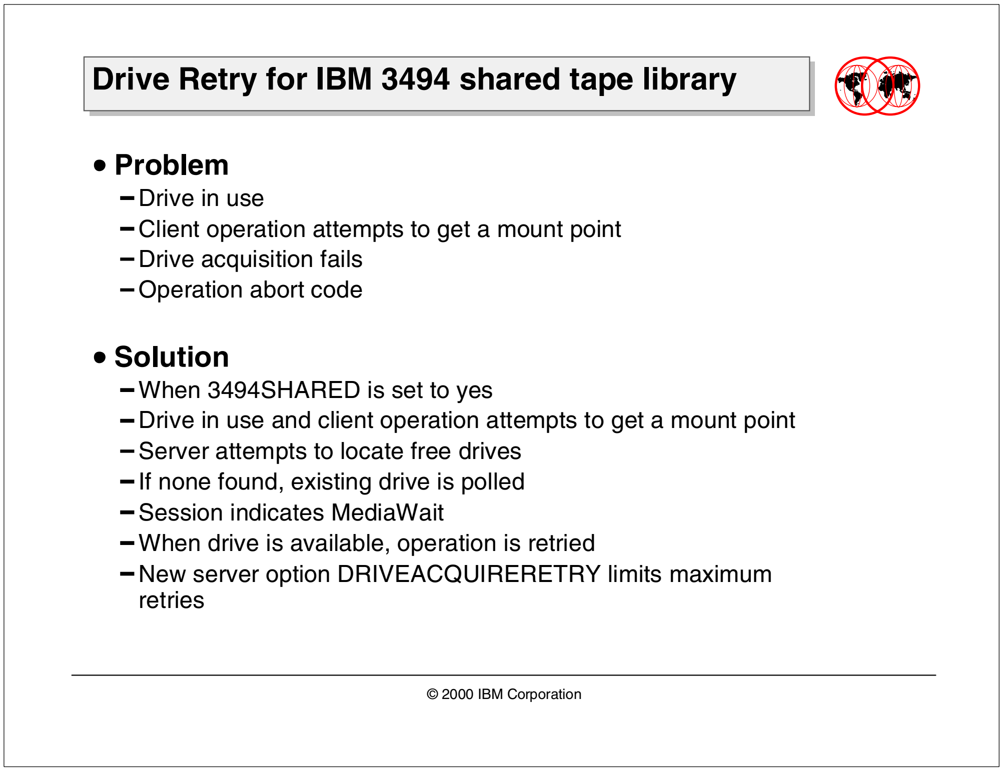

Drive Retry for IBM 3494 shared tape library

- **Problem**
  - Drive in use
  - Client operation attempts to get a mount point
  - Drive acquisition fails
  - Operation abort code
- **Solution**
  - When 3494SHARED is set to yes
  - Drive in use and client operation attempts to get a mount point
  - Server attempts to locate free drives
  - If none found, existing drive is polled
  - Session indicates MediaWait
  - When drive is available, operation is retried
  - New server option DRIVEACQUIRERETRY limits maximum retries

8.2.1 Drive Retry for IBM 3494 shared library

In an environment where a 3494 library is shared between two or more TSM servers, there are cases where a write or read tape operation from a given server fails, because the drive was in use by a different application. For example, if a client is backing up to the server and has a destination indicating that it should go directly to a sequential media storage pool, the operation can fail, and the client will terminate the action. This occurs because when the data is actually sent to the server, the server is unable to acquire a drive to satisfy the allocated mount point.

The sequence of events is as follows:

1. A tape operation is requested by the client. This can be either a read or a write operation.
2. The server attempts to allocate a "mount point" for this operation.
3. A mount point will be granted if the "mount limit" for the specified device class has not been exceeded on this server. So, if their 3590CLASS had a mount limit of 5 specified, if this server only had three other mount points in this device class allocated and assigned, it would be able to grant what it thought to be the fourth mount point.
4. The client begins the tape operation. In the case of a write, data is sent to the server to be stored — in this case, either a scratch volume or a private volume that is not full is sought to be mounted. And in the case of a read, a restore/retrieve request is made, and a specific volume is sought to be mounted.

5. As a side-effect of the volume open processing, the server actually acquires the drive and attempts the mount. In the case where this is a shared library and the drive is in fact already in use by a different host (typically, this is another TSM or ADSM server), the drive acquisition fails, and the underlying volume open processing on the server fails.
6. At this point, the client is returned an abort code indicating that the data is not available on the server.

8.2.2 Solution for TSM V3.7.3 and V4.1

Enhancements have been made to the TSM server so the server will change the way mount point allocation and volume open operations are working. The general approach is as follows:

To acquire an actual mount point for a read/write operation, TSM reviews each drive defined to the specified library that TSM believes is empty. TSM reviews each drive it believes is empty until it finds a drive for the operation or it has exhausted the MOUNTWAIT period with no drive becoming available. If there is a volume currently loaded in the drive, TSM will poll the drive. The polling thread provides an internal notification that there was a change in the state of the polled drive; that is, either it is or is not available. Any processes "listening" for that notification can then take the appropriate actions. Again, TSM does not wait for a particular drive to become available.

If TSM has a problem communicating with a drive in a 3494 library or the drive is unavailable to the 3494 library manager, TSM will poll the drive only for 10 minutes.

TSM will retry the actual mount point acquisition the number specified in the DRIVEACQUIRERETRY option. The reason is there is a race condition between the TSM servers that share the drives in a 3494 library. If, for example, TSM server 1 is waiting for a drive to become available, and the drive does become available, there is a chance that TSM server 2 could be using the same drive again before TSM server 1 retries to acquire the actual mount point.

In the case where the drive cannot be acquired even after the retry attempts, the client sees the same abort code, indicating that data is not available on the server.

During this entire mount point request/poll/retry/acquire operation, the session reflects a state of MediaWait.

8.2.3 Server enhancements

The server option DRIVEACQUIRERETRY has been added. This possible values for this option are:

`FOREVER` — The acquisition of the drive will be retried until a drive is successfully acquired. This is the default setting for this option and the default behavior that customers should expect to experience.

`NEVER` — Never attempt retry acquisition of a drive. If it fails, exhibit the current behavior that is done today and fail the operation.

`1 - 9999` — If a numeric value in the range of 1 to 9999 is entered, this is the MAXIMUM number of retries that will be attempted for re-acquisition of the drive. Supported server levels include TSM 3.7.3 and TSM 4.1.

8.3 LAN-free client data transfer for API client

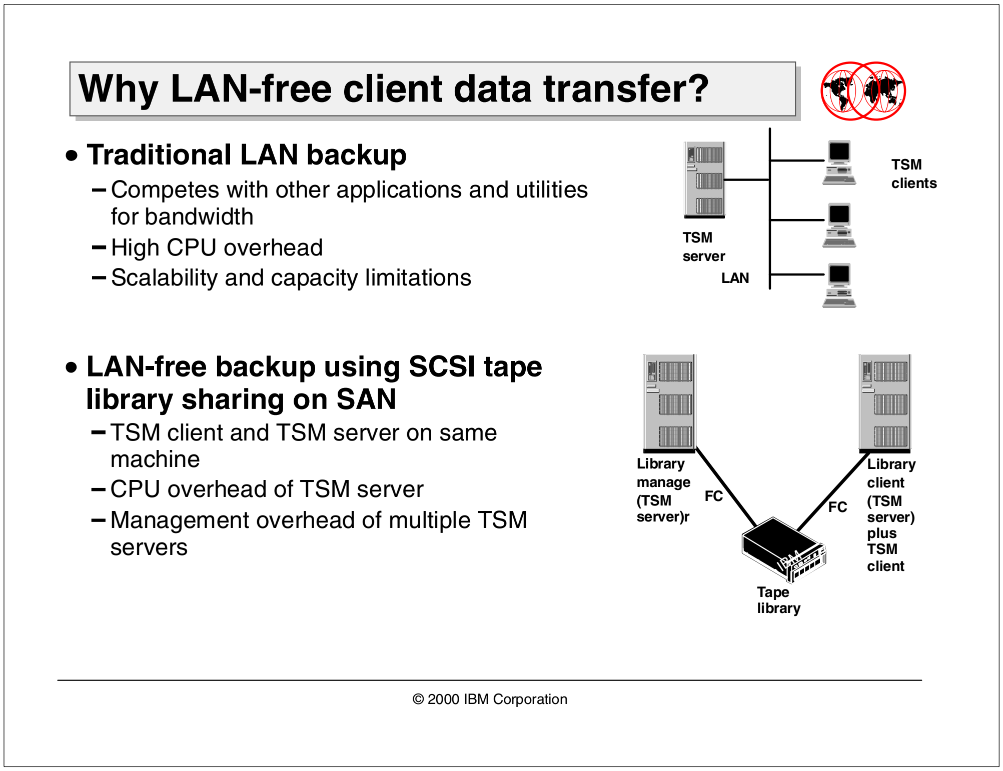

Why LAN-free client data transfer?

- **Traditional LAN backup**
  - Competes with other applications and utilities for bandwidth
  - High CPU overhead
  - Scalability and capacity limitations
- **LAN-free backup using SCSI tape library sharing on SAN**
  - TSM client and TSM server on same machine
  - CPU overhead of TSM server
  - Management overhead of multiple TSM servers

TSM clients
TSM server
LAN
Library manage (TSM server)r
FC
FC
Library client (TSM server) plus TSM client
Tape library

© 2000 IBM Corporation

The advent of SAN technology provides an alternative path for data movement between the TSM client and the server. Shared storage resources are accessible to both the TSM client and the server through the storage area network. It is possible, then, for the TSM client to write to storage managed by the server, or for the server to read directly from disks that are normally used by the client for filesystem/database applications.

LAN-free client data transfer is a feature of Tivoli Storage Manager that allows a Tivoli Data Protection application client to transfer data directly to a SAN attached storage device (tape only for now) that is known to both the client and server.

In this section, we will cover the architecture and implementation of this feature.

8.3.1 Why LAN-free client data transfer?

The purpose of LAN-free data movement is to off-load storage management data movement from the network so that more network bandwidth is available to other applications.

8.3.1.1 LAN backup limitations

The current Tivoli Storage Manager product utilizes network bandwidth for data movement between the client and the server. This approach has many disadvantages:

- LAN speeds are fast becoming insufficient to support today’s high growth and high volume backup environments.

- LAN backup suffers from high network overhead that is inherent in the prevalent LAN protocols.
- It also incurs high CPU overhead on the systems sending or receiving data.
- It decreases the available bandwidth available for database, mail, internet, file serving or other applications that also depend on the network.

SAN I/O transfer is much more efficient than that of a LAN. Furthermore SANs offer a high level of scalability and were designed to transfer large amounts of data with minimum delay.

8.3.1.2 LAN-free SCSI tape library sharing — backup limitations

Tivoli Storage Manager provides a library, drive, and media sharing solution for Tivoli Storage Manager servers.

Tivoli Storage Manager Version 3.7 exploits the Storage Area Network (SAN) technology for dynamic tape library sharing. With this function, multiple Tivoli Storage Manager servers can dynamically share library volume and tape drive resources of one physical tape library.

One Tivoli Storage Manager server, known as the library server, controls the tape library. It is responsible for the mounts and dismounts of tapes into the drives, for the serialization of the drive accesses, and for the library volume inventory. The server can also use all of the tape drives for its own operation.

All other Tivoli Storage Manager servers using this tape library are known as library clients. They have physical access to all the drives throughout the SAN. The library clients can directly read or write data to or from tapes, but they need the library manager server to get tapes mounted or unmounted.

While this is a useful solution for many environments, the hardware and management cost incurred due to having multiple Tivoli Storage Manager servers may make it prohibitive for small customers.

With the TSM LAN-free backup solution, the need to manage and maintain multiple servers has been removed. The library client concept can now be implemented on a TSM client.

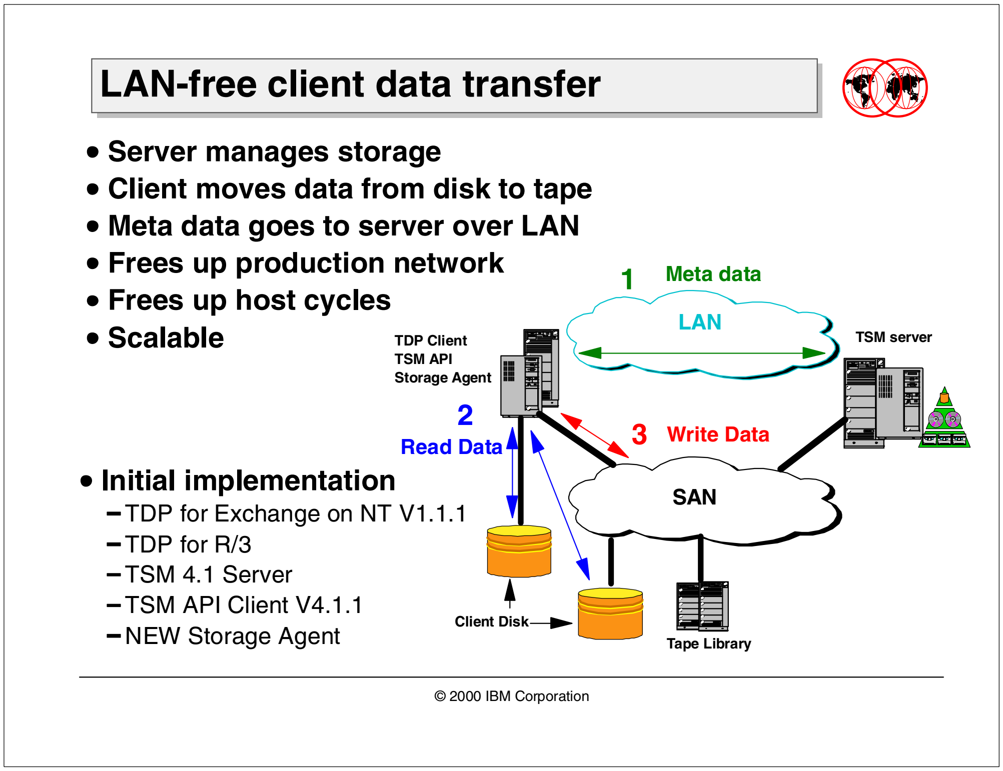

LAN-free client data transfer

- Server manages storage
- Client moves data from disk to tape
- Meta data goes to server over LAN
- Frees up production network
- Frees up host cycles
- Scalable

- Initial implementation
  - TDP for Exchange on NT V1.1.1
  - TDP for R/3
  - TSM 4.1 Server
  - TSM API Client V4.1.1
  - NEW Storage Agent

8.3.2 Tivoli Storage Manager LAN-free client data transfer

Tivoli Storage Manager offers a solution to off-load network data movement for TSM client-server operation. The architecture relies on shared devices between the TSM client and server, such as that available in a SAN environment.

The TSM API client in conjunction with the TSM V4.1 server and a new TSM storage agent has been enhanced with the ability to write directly to server-owned tape storage media in a format that is consistent with that written by the server today. TDP for Exchange and TDP for R/3 are the only supported applications in the initial release.

The existing client-server network connection is used to exchange control information such as policy information and meta-data about the objects being backed up, but the data movement utilizes the LAN-free path over the SAN to write directly to the tape storage media.

Data movement is off-loaded from the LAN and from the TSM server processor for greater scalability.

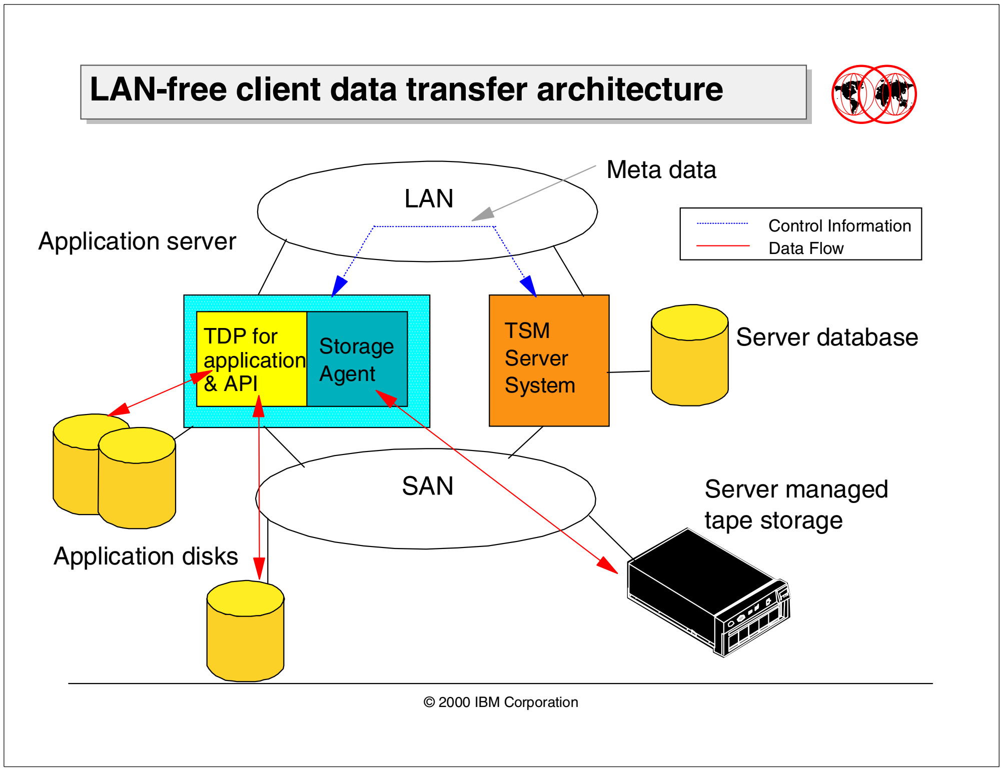

LAN-free client data transfer architecture

Meta data

LAN

Application server

Control Information

Data Flow

TDP for application & API

Storage Agent

TSM Server System

Server database

SAN

Server managed tape storage

Application disks

© 2000 IBM Corporation

8.3.3 LAN-free client data transfer architecture

Tivoli Storage Manager's LAN-free client data transfer solution provides direct movement between the TSM API client (TDP for application initially) and TSM server managed tape storage as controlled by the client system. A number of unique features and concepts were introduced both at the client and server end to architect this solution.

8.3.3.1 Storage agent

A special *storage agent* has been constructed and provided as a local service on the TSM client machine. The storage agent communicates with the actual TSM server over LAN using TCP/IP to obtain and store database information and coordinate device and volume access, much like a library client sends. But the storage agent has no database or storage hierarchy management features of its own. Meta data is passed between the storage agent and the TSM server and all events related to the storage agent are logged to the TSM server.

8.3.3.2 SCSI tape library sharing

Library sharing support is used to share SAN devices between the server and the storage agent. In order to set up the configuration, the server must assume the role of primary library manager. Therefore, the `DEFINE LIBRARY` command with the option `SHARED=YES` must be issued on the server. Extensions to the shared library mount/dismount functions allow the storage agent to request the server (library manager) for these functions for volumes in drives that are defined to it. The server tracks the drives that are being used by the storage agent and uses error processing to reacquire them if the storage agent cannot be contacted. Library volume ownership is tracked only by the TSM server (library manager).

8.3.3.3 Drive mapping

As a prerequisite to sharing libraries between the server and the storage agent, *drive mappings* (see details about `DRIVEMAPPING` commands in the next section) are set up on the server to indicate that a LAN-free path is available and to specify the device or use device name for the storage agent.

For this purpose, database entries on the server have been enhanced to track the `DEVICE` names for the server as well as the storage agent that is configured to share access to the device centrally on the server. The server stores this drive mapping information in a new database table. The new database table contains the drive name, its alias, its online status, and the storage agent identification.

The storage agent is able to determine which devices can be accessed in a LAN-free manner or which must be accessed in the traditional manner by contacting the server.

8.3.3.4 Backups

For backup operations, the client communicates with the storage agent if the `ENABLELANFREE` option is specified in the `dsm.opt` file for the TSM API client. The communication proceeds like this:

1. The client communicates with the server to inquire about the policy information for the data that is going to be backed up.
2. The server examines the device class defined for each storage pool destination that is referenced in the active policy set for the client. If the server finds that the drives in the device class have drive mappings defined for the client's storage agent, the server indicates to the client that a LAN-free path exists to the management class destination.
3. The client then communicates with the storage agent for data transfer.
4. The storage agent, by contacting the server, is aware of the devices that can be accessed in a LAN-free manner from this client and sends data directly over the SAN. The meta data and other control commands are still sent over the LAN.
5. In case no SAN devices are available to the management class for the backed-up data, the client sends the data directly to the server over the LAN as is done today.

8.3.3.5 Restores

For restore operations, the client always communicates with the storage agent to perform the restore if the `ENABLELANFREE` option is specified in the `dsm.opt` file for the TSM API Client. Since any TSM client never knows the location of the file in the server storage hierarchy, this feature makes the interaction much simpler between the TSM client and server. The data flows through the storage agent even if the file is not stored on the SAN-accessible device.

8.3.3.6 Failover

It is possible to set up the configuration in a way that, if the storage agent is unavailable at the time a backup or restore operation is being requested, then the client will failover the operation and send the data directly to the server using the LAN. If, however, the operation fails in the middle, the result is the same as today. The client software goes through retry logic that currently exists.

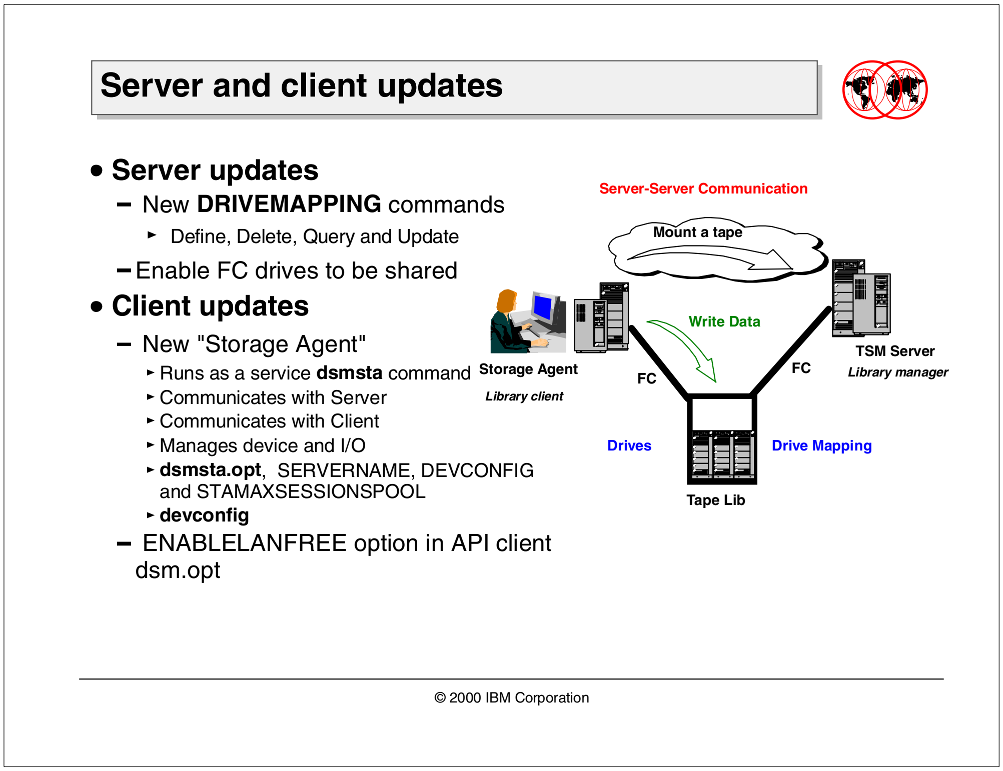

Server and client updates

- **Server updates**
  - New **DRIVEMAPPING** commands
    - Define, Delete, Query and Update
  - Enable FC drives to be shared
- **Client updates**
  - New "Storage Agent"
    - Runs as a service **dsmsta** command
    - Communicates with Server
    - Communicates with Client
    - Manages device and I/O
    - **dsmsta.opt**, SERVERNAME, DEVCONFIG and STAMAXSESSIONSPOOL
    - **devconfig**
  - ENABLELANFREE option in API client dsm.opt

8.3.4 Server updates

New DRIVEMAPPING commands have been implemented in the server to support this feature.

8.3.4.1 DEFINE DRIVEMAPPING command

The DEFINE DRIVEMAPPING command is used to define a device name alias for a drive. This command is used for devices that are shared between the TSM server and client. Devices accessible from different hosts in configurations like a storage area network may have different name specifications from each host, depending upon the adapter that is used to access the device.

This command is used to define the device name for a drive from the perspective of a specified client. The drive must already be defined to the TSM server (using the DEFINE DRIVE command) before an alias can be defined for it. Tivoli Storage Manager will use the alias information to automatically perform LAN-free data movement when accessing storage pools whose drives are aliased to the client.

To issue this command you must have system privilege or unrestricted storage privilege. The following parameters are applicable for the command.

storage_agent_name (required)

> Specifies the name of the storage agent as specified in the define server command.

lib_name (required)

Specifies the name of the library to which the drive is defined. The parameter is required for all drives, including stand-alone drives. The specified library must have been already defined using the `DEFINE LIBRARY` command.

drive_name (required)

Specifies the name that is assigned to the drive whose device name you wish to map. The drive must have already been defined using the `DEFINE DRIVE` command.

device_name (required)

Specifies the device name that is used to access the drive from the perspective of the storage agent machine.

online (optional)

Specifies if the drive is currently available for the storage agent:

`YES` — The drive can be used by the storage agent

`NO` — The drive can not be used.

8.3.4.2 DELETE DRIVEMAPPING command

The `DELETE DRIVEMAPPING` command is used to remove a device name mapping that was created with the `DEFINE DRIVEMAPPING` command. If you a remove a drive alias for a client node, the client will NOT be able to use LAN-free data movement to store or retrieve data to/from the TSM server using that drive. To issue this command, you must have system privilege or unrestricted storage privilege. The parameters for this command include:

storage_agent_name (required)

lib_name (required)

drive_name (required)

8.3.4.3 UPDATE DRIVEMAPPING command

The `UPDATE DRIVEMAPPING` command is used to update a device name alias that was created with the `DEFINE DRIVEMAPPING` command. Tivoli Storage Manager will use the alias information to automatically perform LAN-free data movement when accessing storage pools whose drives are aliased to the client. To issue this command you must have system privilege or unrestricted storage privilege. The requisite parameters include:

storage_agent_name (required)

lib_name (required)

drive_name (required)

device_name (optional)

online (optional)

8.3.4.4 QUERY DRIVEMAPPING command

The QUERY DRIVEMAPPING command is used to display a device name mapping that was created with the DEFINE DRIVEMAPPING command. To issue this command, you must have system privilege or unrestricted storage privilege. The requisite parameters include:

storage_agent_name (required if lib name or drive name are specified)

lib_name (required if drive name specified)

drive_name (optional)

Format (optional)

Standard is the default and Detailed adds more information.

8.3.5 Client updates

The major client updates include updates to the API client dsm.opt file and the addition of a new storage agent.

8.3.5.1 API client option file updates (dsm.opt)

The new ENEABLELANFREE option was introduced to the API client option file. It has the following syntax:

```
ENABLELANFREE Yes | No
```

This option is used to indicate that it is possible for a client to use SAN attached devices during backup and restore operations. The default setting is no.

If this option is specified in the dsm.opt file, the client will contact the storage agent for backups if the management class for the data indicates that a LAN-free path is available. Furthermore, the client will always contact the storage agent for restores.

If this option is disabled, the client will not know about the storage agent. In that case, the client behavior will be the same as today.

8.3.5.2 Storage agent

The setup files for the storage agent include dsmsta.opt and devconfig. A new command DSMSTA is used to start the service on the client.

dsmsta.opt setup file

The dsmsta.opt setup file can contain the following entries:

```
SERVERName <name>
```

> The SERVERNAME option defines the name of the TSM server. The server name is included in the dsmsta.opt file when executing the DSMSTA SETSTORAGESERVER command and matches with the entry in the devconfig file. The server name specified must accept the storage agent connection.

STAMaxsessionpool <number>

The STAMAXSESSIONPOOL option defines the maximum number of sessions the storage agent should keep open on an ongoing basis for accessing the TSM server. Default is 25, maximum is 150. Up to value of STAMAXSESSIONPOOL sessions will be held open indefinitely. These sessions do not count against any client maximum. If STAMAXPOOLSESSION is set to 0 then all connection are opened and closed after

DEVCONFIG <filename>

The DEVCONFIG option sets the name and location of the devconfig file. At startup, the storage agent gets information about the location of the devconfig file from dsmsta.opt.

dsmsta command

The DSMSTA command starts the storage agent on the client.

```
tsm > dsmsta
```

The DSMSTA command is also used for the storage agent to set up the necessary information in the devconfig file.

```
tsm > dsmsta SETSTORAGESERVER MYName=myservername MYPAssword=myserverpw
SERVERName=dbservername SERVERPAssword=dbserverpw
HLAddress=dbserver_tcpip_address LLAddress=dbserver_tcpip_port
```

The applicable parameters are shown below. All parameters are required.

MYName

Specifies the name of the storage agent. This value will be used as the server name in a SET STANAME command inserted into the devconfig file.

MYPASSword

Specifies the password of the storage agent. This value will be encrypted and used in a SET STAPASSWORD command inserted into the devconfig file.

SERVERName

Specifies the name of the TSM Server. This value will be used as the target server in a DEFINE SERVER command inserted into the devconfig file.

SERVERPassword

Specifies the password of the TSM Server. This value will be encrypted and used in the SERVERPASSWORD field of the DEFINE SERVER command inserted into the devconfig file.

HLAddress

Specifies the TCP/IP address of the TSM Server. This value will be used in the HLADDRESS field of the DEFINE SERVER command inserted into the devconfig file.

LLAddress

Specifies the TCP/IP port on which to access the TSM Server. This value will be used in the LLADDRESS field of the DEFINE SERVER command inserted into the devconfig file.

devconfig file

The devconfig file can be described as a bootstrap file which contains the information necessary for the storage agent to start up. At startup, the storage agent gets information about the location of the devconfig file from dsmsta.opt. The storage agent reads the devconfig file to get information about itself and about the server it has to access for device information.

The dsmsta command (mentioned above) logs the three commands in the devconfig file as follows:

```
SET STANAME myservername
SET STAPASSWORD myserverpw
DEFINE SERVER dbservername HLADDRESS=dbserver_tcpip_address
LLADDRESS=dbserver_tcpip_port SERVERPASSWORD=dbsrverpw
```

Figure 58. The devconfig file for the storage agent

Note

The SERVERPASSWORD is encrypted.

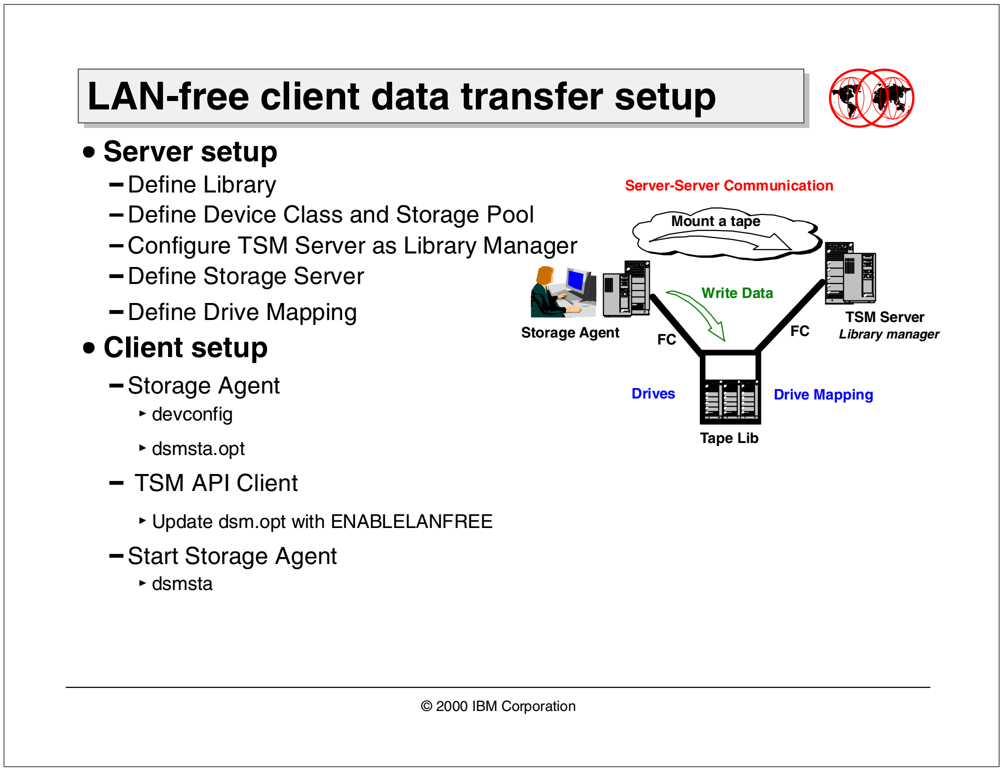

LAN-free client data transfer setup

- **Server setup**
  - Define Library
  - Define Device Class and Storage Pool
  - Configure TSM Server as Library Manager
  - Define Storage Server
  - Define Drive Mapping
- **Client setup**
  - Storage Agent
    - devconfig
    - dsmsta.opt
  - TSM API Client
    - Update dsm.opt with ENABLELANFREE
  - Start Storage Agent
    - dsmsta

© 2000 IBM Corporation

8.3.6 LAN-free client data transfer setup

An example of a LAN-free backup environment is given in the following sections. TSMSERVER and TSMNODE are two NT workstations. TSMSERVER is set up as a TSM server running version 4.1. TSMNODE is registered on the TSM server and set up as a TSM client. The storage agent, TDP for Exchange and updated TSM API client are installed on TSMNODE. STGSVC is the name of the storage agent on TSMNODE. Note that the storage agent is installed from its own CD-ROM.

The systems connect to an IBM MAGSTAR 3570 C12 Tape Library with two drives. The drives are connected to a SCSI port on the SAN Data Gateway. TSMSERVER and TSMNODE attach to two separate FC-AL ports on the same SAN Data Gateway to access the drives. The drives in the library are already configured to the operating systems on both systems. Please refer to Section 8.1.2, “Sample SAN configuration” on page 251 on how to set up the SAN.

To configure the LAN-free client data transfer you must do the following:

- TSM Server Configuration:
  1. Library and drive definitions
  2. Device class and storage pool definitions
  3. Library Manager setup
  4. Storage agent definitions
  5. Drive Mapping definitions

- TSM Client Configuration:
  1. TDP for Exchange options file
  2. Storage agent definitions

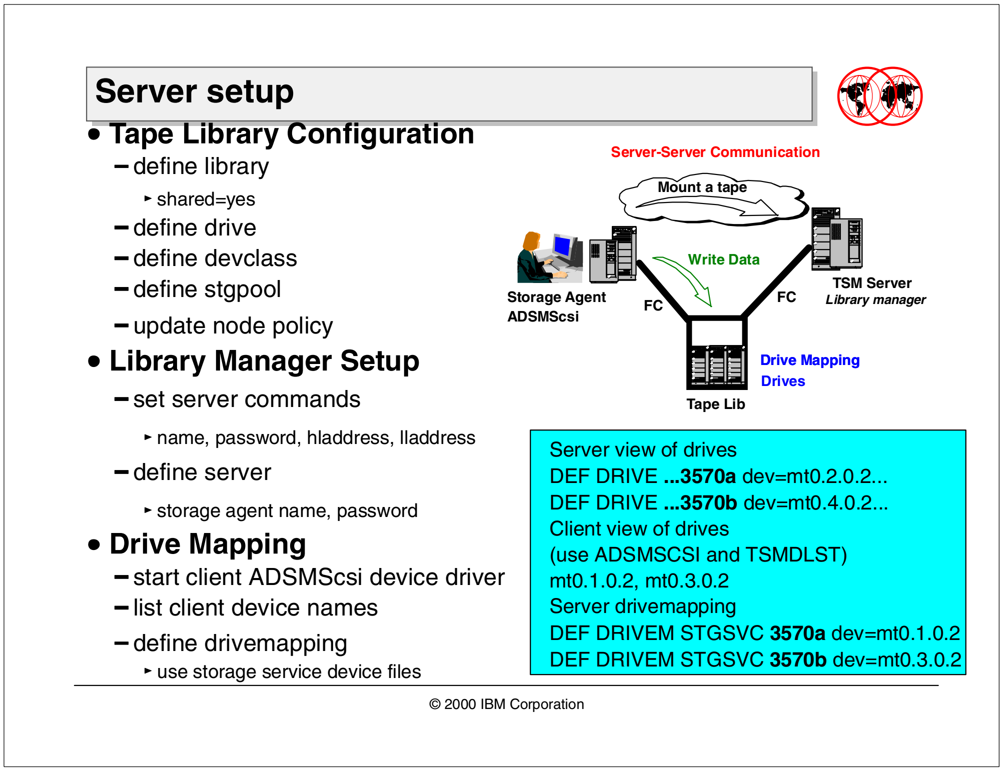

Server setup

- **Tape Library Configuration**
  - define library
    - shared=yes
  - define drive
  - define devclass
  - define stgpool
  - update node policy
- **Library Manager Setup**
  - set server commands
    - name, password, hladdress, lladdress
  - define server
    - storage agent name, password
- **Drive Mapping**
  - start client ADSMScsi device driver
  - list client device names
  - define drivemapping
    - use storage service device files

Server view of drives
DEF DRIVE **...3570a** dev=mt0.2.0.2...
DEF DRIVE **...3570b** dev=mt0.4.0.2...
Client view of drives
(use ADSMSCSI and TSMDLST)
mt0.1.0.2, mt0.3.0.2
Server drivemapping
DEF DRIVEM STGSVC **3570a** dev=mt0.1.0.2
DEF DRIVEM STGSVC **3570b** dev=mt0.3.0.2

8.3.6.1 TSM server configuration

As part of setting up a LAN-free environment, TSMSERVER is setup as a library manager with the correct device definitions. Secondly, the definitions for the storage agent are introduced on the server.

Step 1: Define tape library to TSM server

The first step is to define the tape library and the drives to the library manager (in our case, TSMSERVER). The following set of commands are issued:

```
tsm: TSMSERVER> def libr 3570lib libtype=scsi device=lb0.3.0.2 \
cont>shared=yes
tsm: TSMSERVER> def drive 3570lib 3570a device=mt0.2.0.2 element=16
tsm: TSMSERVER> def drive 3570lib 3570b device=mt0.4.0.2 element=17
```

The `shared=yes` parameter indicates that this library will be shared with other Tivoli Storage Manager Servers or, as in our case, storage agent Clients (STGSVC).

Note

The server must be licensed for tape library sharing, and the license must be registered using the `REGISTER LICENSE` command.

Step 2: Define device class and storage pool

A new device class and storage pool are created for the 3570 Tape Library. The existing disk storage pools are updated to point to the new storage pool. The following commands are issued:

```
tsm: TSMSERVER> def devc 3570devc devt=3570 libr=3570lib \
cont> mountlimit=2 mountreten=1
tsm: TSMSERVER> def stg tapepool 3570devc maxscratch=100
```

Step 3: Update node policy information

The policy for registered node TSMNODE is also updated so the data goes directly to `tapepool`. TSMNODE is registered with default policy settings in our environment. The following command is issued to update the default policy settings of the backup copy group in the default management class, so the data goes directly to `tapepool`.

```
tsm: TSMSERVER> upd stg archivepool next=tapepool
tsm: TSMSERVER> upd stg backuppool next=tapepool
tsm: TSMSERVER> update copygr standard standard standard standard \
cont> type=backup destination=tapepool
tsm: TSMSERVER> update copygr standard standard standard standard \
cont> type=archive destination=tapepool
tsm: TSMSERVER> validate policyset standard standard
tsm: TSMSERVER> activate policyset standard standard
```

Step 4: Configure server-to-server communications

TSMSERVER is setup as a library manager. Server-to-server communications must be established between TSMSERVER and the storage agent. The list of commands is given below. (For instructions on how to set up a library manager, including detailed information on each command, please refer to *Tivoli Storage Manager V3.7: Technical Guide*, SG24-5477.)

```
tsm: TSMSERVER> set servername TSMSERVER
tsm: TSMSERVER> set serverpassword TSMSERVER
tsm: TSMSERVER> set serverhla TSMSERVER.sanjose.ibm.com
tsm: TSMSERVER> set serverlla 1500
tsm: TSMSERVER> set serverurl http://TSMSERVER:1580
tsm: TSMSERVER> set crossdefine on
```

The definitions for the TSM server are also introduced in the devconfig file for the storage agent on the client at a later stage, and so it is important to note these for future reference.

Step 5: Define storage agent to TSM server

To use the tape sharing configuration between the TSM server and storage agent server-to-server communications must be established.

Define the storage agent (STGSVC) to the TSM server. This is done so that TSMSERVER has information on the storage agent, and so that the storage agent can communicate with the TSM server. In an environment with multiple storage agents, this command needs to be issued once for each system:

```
tsm: TSMSERVER> define server STGSVC desc="Storage agent" \
cont> serverpa=STGSVC hla=STGSVC.sanjose.ibm.com lla=1500 \
cont> crossdefine=no
```

A server definition for the TSM server (TSMSERVER) is also defined on the client at a later stage. It is important to ensure that this name and password matches at both TSM server and client.

Step 6: Define drive mapping for storage agent

This is a very important step in the configuration process. The `drivemapping` command is introduced so TSMSERVER can define how the devices are “seen” from the perspective of STGSVC. When the storage agent contacts the server about the device names of the drives in the SAN attached tape library, the server sends it the names of the drives with the device information that is consistent with how the storage agent sees it at its own end. Thus when the storage agent sends the actual data, the data goes to the correct tape drive.

```
tsm: TSMSERVER> define drivemapping STGSVC 3570lib 3570a \
cont> device=mt0.1.0.2
tsm: TSMSERVER> define drivemapping STGSVC 3570lib 3570b \
cont> device=mt0.3.0.2
```

As a prerequisite to this step, the address information is acquired from the client itself. Start the ADSMSCSI device driver on the client and use the tool `TSMDLST` command to get the actual device addresses. The tool comes together with the storage agent software distribution which is on a separate CD-ROM.

Note

For tape library sharing between multiple systems, particularly in SAN environments involving hubs and switches, sometimes the devices are mapped differently on each system. Each drive has a SCSI address. The drives are usually mapped in order of their SCSI IDs, either from lowest to highest or vice versa. The addressing is also a function of the Fibre Channel HBA and its discovery mechanisms, and the behavior may vary from vendor to vendor.

It is important to verify that the device name obtained from the client and used in the `drivemapping` command indeed maps to the same physical drive that is represented by the device name issued in the `define drive` command on the server for that drive.

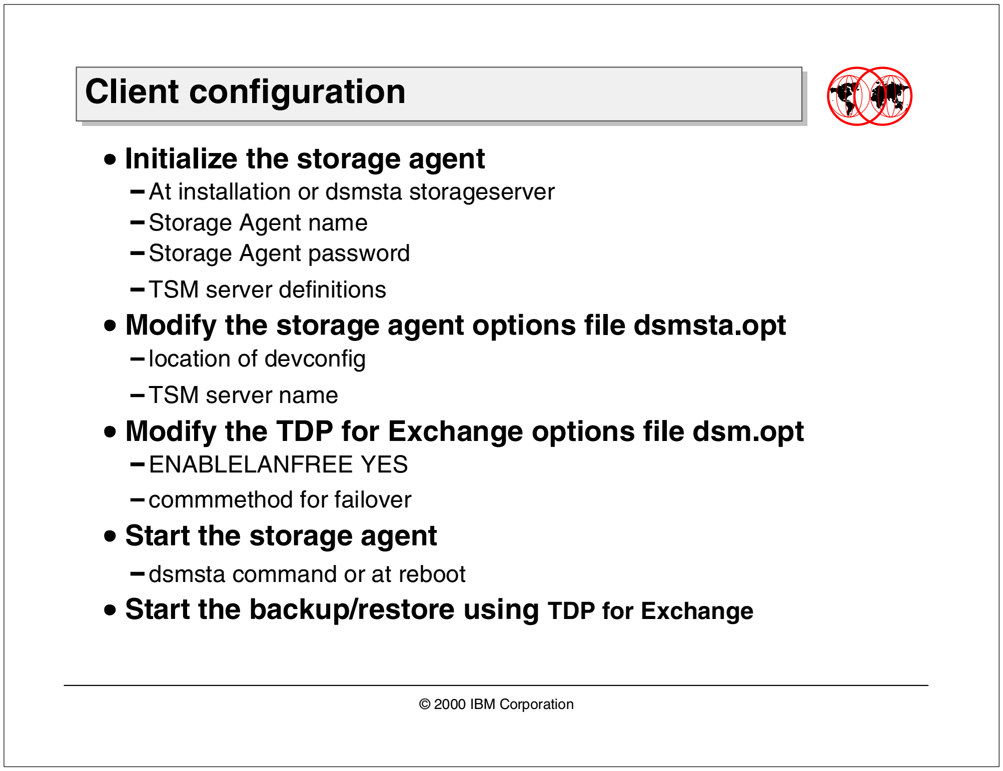

Client configuration

- **Initialize the storage agent**
  - At installation or dsmsta storageserver
  - Storage Agent name
  - Storage Agent password
  - TSM server definitions
- **Modify the storage agent options file dsmsta.opt**
  - location of devconfig
  - TSM server name
- **Modify the TDP for Exchange options file dsm.opt**
  - ENABLELANFREE YES
  - commmethod for failover
- **Start the storage agent**
  - dsmsta command or at reboot
- **Start the backup/restore using TDP for Exchange**

8.3.6.2 TSM client configuration

After the installation of the updated client (TSM Backup-Archive client, TSM API client, TDP for Exchange) code, the configuration basically includes setting up the definitions for the storage agent and modifying the TDP for Exchange options file to enable the client to be backed up LAN-free.

Step 1: Modify storage agent’s devconfig file

The devconfig file is a sort of bootstrap file for the storage agent. When the storage agent is started, it uses `dsmsta.opt` to get information on where the `devconfig` file is located. Next it locates and reads the `devconfig` file to get information about its name and password, as well as definitions for the TSM server it has to communicate with.

Usually the information for the defconfig file is gathered during the installation of the storage agent. However, you can also use the `DSMSTA SETSTORAGESERVER` command to change this the information in the `devconfig` file. The command is given below.

```
tsm > dsmsta SETSTORAGESERVER MYName=STGSVC MYPAssword=STGSVC
SERVERName=TSMSERVER SERVERPAssword=TSMSERVER
HLAddress=TSMSERVER.sanjose.ibm.com LLAddress=1500
```

The above values, `MYName` and `MYPAssword`, should match the name and password information used when the storage agent was defined on the server with the `DEFINE SERVER` command. These values are used in the `SET STANAME` and `SET STAPASSWORD` inserted in the devconfig file. The next four parameters define the name of the TSM server. These values should be the same as specified in the `SET SERVER` commands for the TSM server, and are used in the `DEFINE SERVER` command inserted in the devconfig file.

The `dsmsta` command stores three commands in the `devconfig` file, as shown in Figure 59:

```
SET STANAME STGSVC
SET STAPA 18e72f1903ce5a889b
DEFINE SERVER TSMSERVER HLADDRESS=TSMSERVER.SANJOSE.IBM.COM LLADDRESS=1500
SERVERPA=18f706dcb20fa58d345877
```

Figure 59. Sample storage agent devconfig file

The passwords are stored in `devconfig` in an encrypted form. The `dsmsta` command does not remove any previous entries in the `devconfig` file, but simply appends these three lines to the existing file. The `devconfig` file is located in `C:\Program Files\Tivoli\TSM\storage agent`.

Step 2: Modify storage agent's dsmsta.opt file

The dsmsta.opt file is the first file the storage agent looks for when it starts. The location of devconfig is also indicated in `dsmsta.opt`. The `dsmsta.opt` file is initially created at installation of the storage agent and can be modified by executing the `DSMSTA SETSTORAGESERVER` command (see "Step 1: Modify storage agent's devconfig file" on page 279) to update the `SERVERName` option to the name of the TSM server. The `dsmsta.opt` file in our example has two entries, as shown in Figure 60:

```
devconfig C:\Program Files\Tivoli\TSM\storage agent\devconfig
SERVERName TSMSERVER
```

Figure 60. Sample storage agent dsmsta.opt file

Step 3: Modify TDP for Exchange option file

The dsm.opt file for TDP for Exchange is modified to include `ENABLELANFREE` option. This option is set to YES indicating to the client that a LAN-free path exists, and that the storage agent should be utilized to send data directly over the LAN-free path to the SAN attached tape library. In case this option is not included in the `dsm.opt` file, the client does not know about the storage agent and sends data directly to the server over the LAN.

The storage agent is configured with the TCP/IP address (the high level address) and port number (the low level address) of the server. The client obtains this information from the storage agent and ignores its own `dsm.opt` communication settings. If, however, the storage agent is unavailable, the client will failover to the real server using communication settings specified in the `dsm.opt` file. Our `dsm.opt` file for TDP for Exchange is therefore modified to include the communication settings pointing to the real server.

Our dsm.opt file for the TDP for Exchange client is shown in Figure 61:

```
*======================================================================
* Tivoli Storage Manager
*
* Tivoli Data Protection for Microsoft Exchange Server

* Options File
*======================================================================

ENABLELANFREE YES

NODename          TSMNODE
CLUSTERnode       No
COMPRESSIon       Off
PASSWORDAccess    Generate
*======================================================================
* TCP/IP
*======================================================================
COMMMethod        TCPip
TCPPort           1500
TCPSERVERADDRESS    TSMSERVER.SANJOSE.IBM.COM
TCPBuffsize        31
TCPWindowsize      63

*======================================================================
* The Exchange Application Client requires tapeprompt to have a value
* of Yes in order to enable the *exc_mountwait parameter to work
* correctly.  If tapeprompt is set to No, then the application client
* will always wait for tape mounts regardless of the *exc_mountwait
* option.
*
* This restriction may be removed in a PTF.
*=======================================================================
TAPEPrompt         yes
*=======================================================================
* Exchange Application Client extended options -- Application Client
* Use Only
*
* Note that each of the keywords for these options begins with a *,
* which is a comment character for normal TSM options files.  This
* allows us to provide extended options in the options file that will be
* used by the Exchange Application Client and ignored by the TSM API.
*=======================================================================
*EXC_MOUNTWait    Yes
```

Figure 61. Sample TDP for Exchange option file

Step 4 — Start the storage agent

The last step is to start the storage agent on the client. Usually the storage agent runs as service, so it will started automatically at system startup; however, it can also be started using the command dsmsta. After that, you are able to initiate backup/restore functions over the SAN using your existing TDP for Exchange client.

Chapter 9. TSM and EMC Symmetrix hardware integration

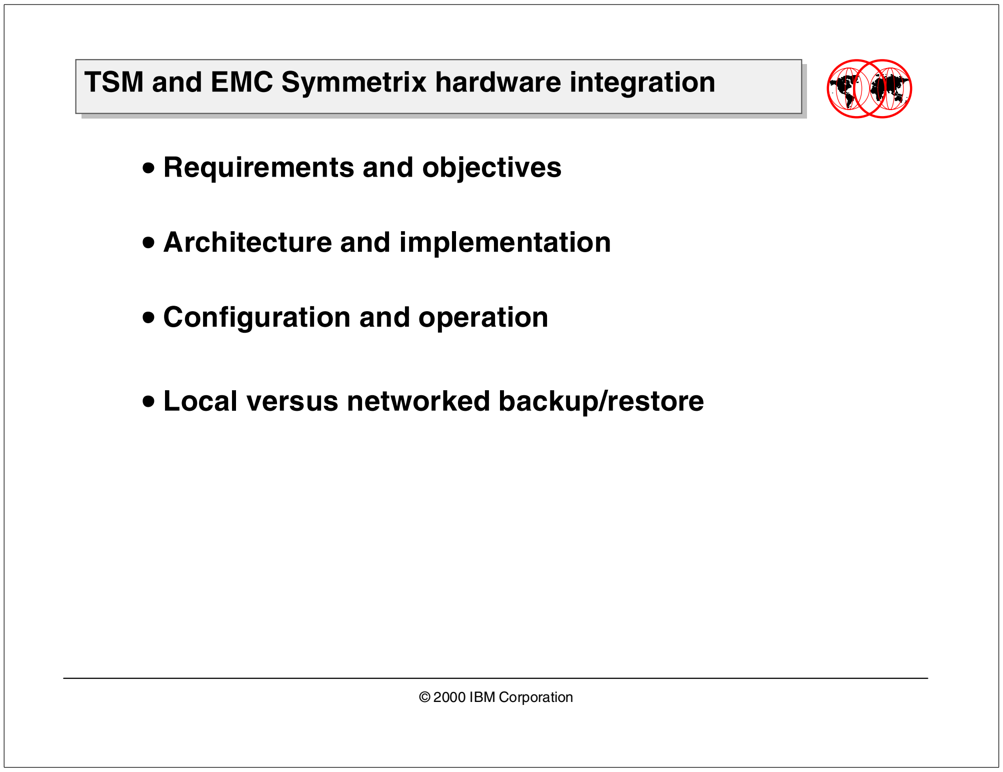

TSM and EMC Symmetrix hardware integration

- Requirements and objectives
- Architecture and implementation
- Configuration and operation
- Local versus networked backup/restore

© 2000 IBM Corporation

This chapter covers the customer requirements for new backup and restore methods. It illustrates the exploitation of the Tivoli Data Protection advanced copy function for EMC TimeFinder.

9.1 Requirements and objectives

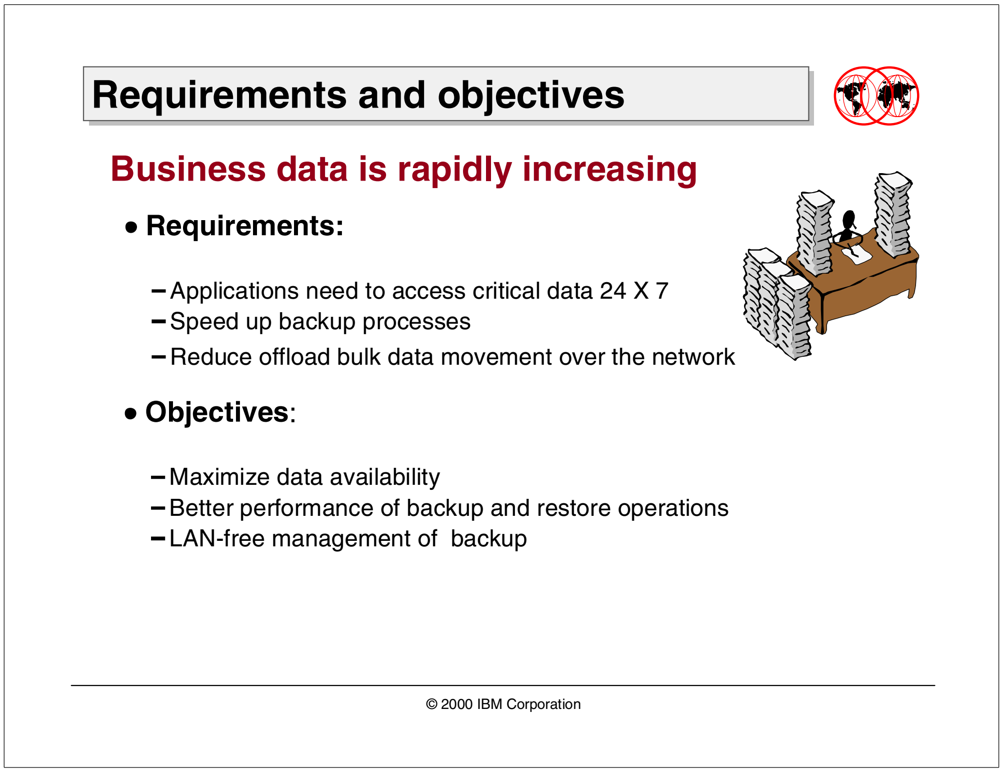

Requirements and objectives

Business data is rapidly increasing

- **Requirements:**
  - Applications need to access critical data 24 X 7
  - Speed up backup processes
  - Reduce offload bulk data movement over the network
- **Objectives**:
  - Maximize data availability
  - Better performance of backup and restore operations
  - LAN-free management of backup

9.1.1 Requirements

As databases grow in size, customers are looking for faster backup and restore solutions. Many today's applications need access to critical data and need to be available to users most of the time. The backup operation for databases with large amounts of data may take a long time to perform.

One important requirement is to speed up the backup in order to reduce the unavailability of data. Bottlenecks must be located and possibly removed, for example, by using faster disk drives or tape units.

If Local Area Network is used to back up data from the database server to the Tivoli Storage Manager server, this becomes the main bottleneck to remove.

9.1.2 Objectives

The objective is to minimize the impact of backing up data from an Oracle database on the database server. This is accomplished by exploiting the fact that the data resides on an Intelligent Disk Subsystem (IDS), such as the EMC Symmetrix, and by using Tivoli Storage Manager to perform scheduled backups of that database.

Since the transfer of backup data is streamed from the database server to the TSM server, usually over the network, the backup of large amounts of data may lead to a collapse of the network and time windows may become insufficient. One way to reduce network contention is to minimize the movement of backup data over the network by exploiting alternate methods provided by vendors of IDS.

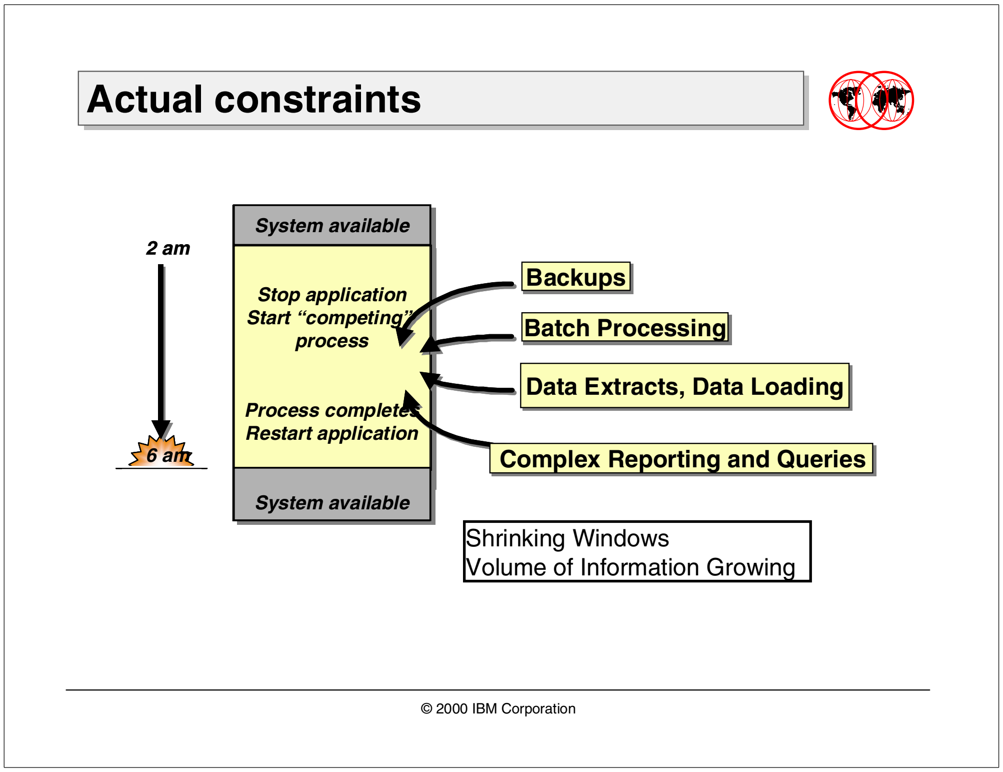

Actual constraints

2 am

6 am

System available

Stop application
Start “competing”
process

Process completes
Restart application

System available

Backups

Batch Processing

Data Extracts, Data Loading

Complex Reporting and Queries

Shrinking Windows
Volume of Information Growing

© 2000 IBM Corporation

9.1.3 Actual constraints

A full database backup operation is usually done in a time window during the night. The database is posted in a quiesced state, and the data is saved.

Batch processing (such as data extracts, data loading, and complex queries) is done when the database backup is completed.

At the completion of this operation sequence, the database is started and made available to the applications.

Since the volume of data is continuously growing, the time window reserved for the backup process can become insufficient; also, batch processes do not fit anymore in the shrinking windows.

The most common backup method moves data from the database server (TSM client) to a TSM server, over the network. Even if the network is quite fast, it can represent a bottleneck for data transfer.

If the database server and the TSM server are on the same machine, the backup operation is dramatically improved, but if using the TSM server to back up other machines over the network, it will contend with DBMS for CPU and system resources.

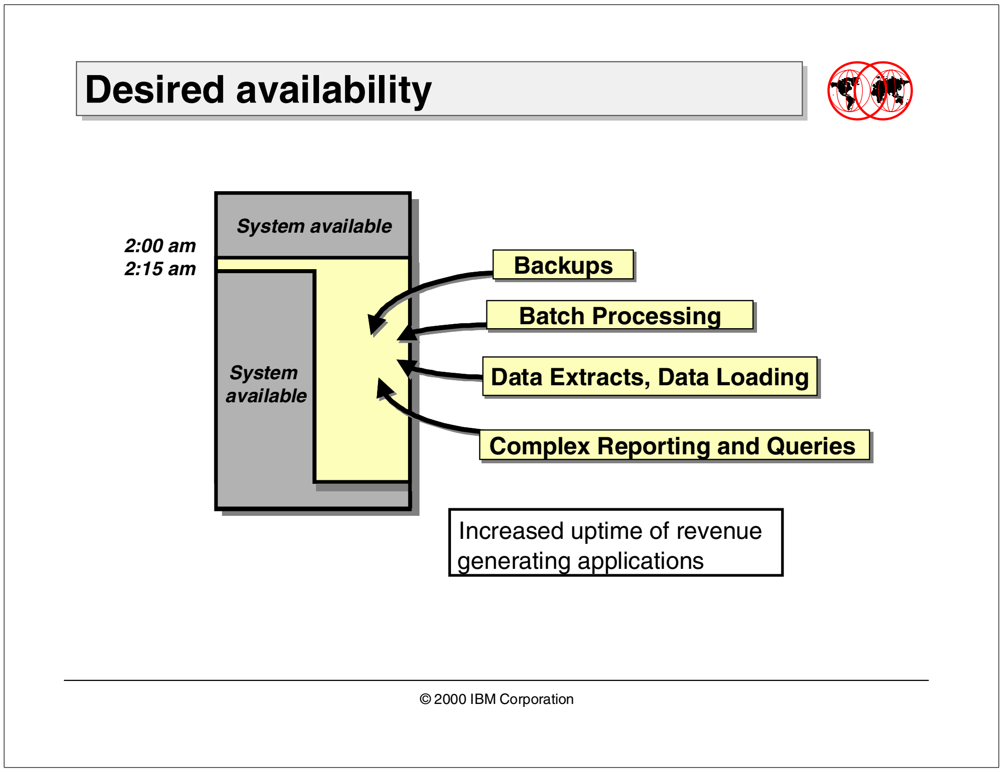

Desired availability

2:00 am

2:15 am

System available

System available

Backups

Batch Processing

Data Extracts, Data Loading

Complex Reporting and Queries

Increased uptime of revenue generating applications

9.1.4 Desired availability

Customers are desirous of innovative uses of fast backup methods to duplicate production databases to off-load decision support (DSS) and batch reporting.

Duplication of data from the production database to volumes that can be accessed by a different system (backup system) allows long and heavy activities, such as complex queries, data warehouse loading and reporting, to be done without affect the availability and performance of the production system.

The split process must be as fast as possible to minimize the time period for which the database is unavailable.

An efficient solution is the integration of the EMC TimeFinder copy function with Tivoli Storage Manager.

9.2 Architecture and implementation

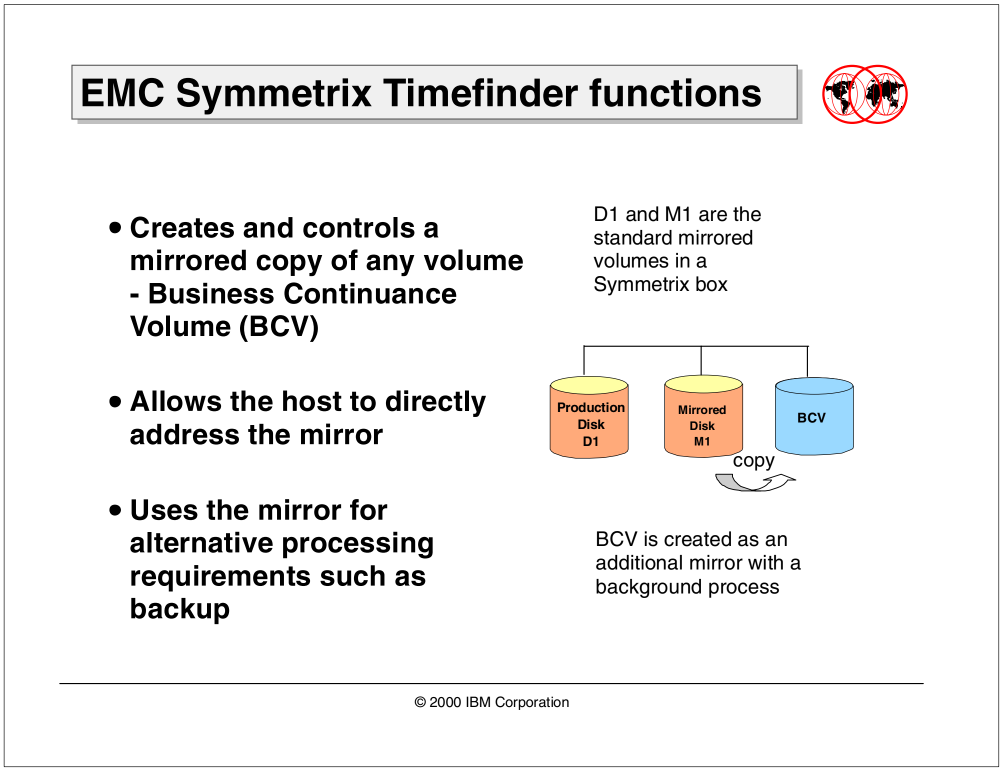

EMC Symmetrix Timefinder functions

- Creates and controls a mirrored copy of any volume - Business Continuance Volume (BCV)
- Allows the host to directly address the mirror
- Uses the mirror for alternative processing requirements such as backup

D1 and M1 are the standard mirrored volumes in a Symmetrix box

Production Disk D1
Mirrored Disk M1
BCV
copy

BCV is created as an additional mirror with a background process

© 2000 IBM Corporation

Creating a copy of the entire database on different physical volumes that are accessible by an alternate (backup) system can increase the availability of the database. The faster this copy operation is accomplished, the better the system availability.

This section covers the architectural design of the EMC Symmetrix TimeFinder copy function and how it is exploited with Tivoli Storage Manager. Two different backup solutions are illustrated.

9.2.1 EMC Symmetrix TimeFinder functions

EMC TimeFinder is a copy function allowing a dynamic mirror to be established, synchronized with its partner, de-established, and then made addressable to a host operating system. With EMC TimeFinder, a third copy of customer data at the logical volume level is possible within a single Symmetrix disk controller.

This new device type is called *Business Continuance Volume* (BCV). It is just an implementation of normal rules of Symmetrix mirroring that must be maintained for local protection. The BCV volume production is called BCV establishment. When the established operation is complete, the BCV volumes are synchronized with the production volumes. EMC TimeFinder allows the hosts to directly address the mirrored volume.

Duplication of data on volumes that can be accessed from a backup system enables the execution of a database backup, and applications that can perform complex queries, extract data, and create reports without affecting the availability and performance of the production system.

Oracle provides the ability to quiesce the database instance running on the primary device to assure that all host buffering and intermediate caching is flushed to the appropriate logical device on the EMC Symmetrix, thereby creating a consistent copy of the data.

This ability allows the Oracle application to continue to run while performing the BCV split accepting the execution of new transactions, and at the same time, writing all information on the redo log files (which may grow quickly if many update or insert transactions are executed frequently).

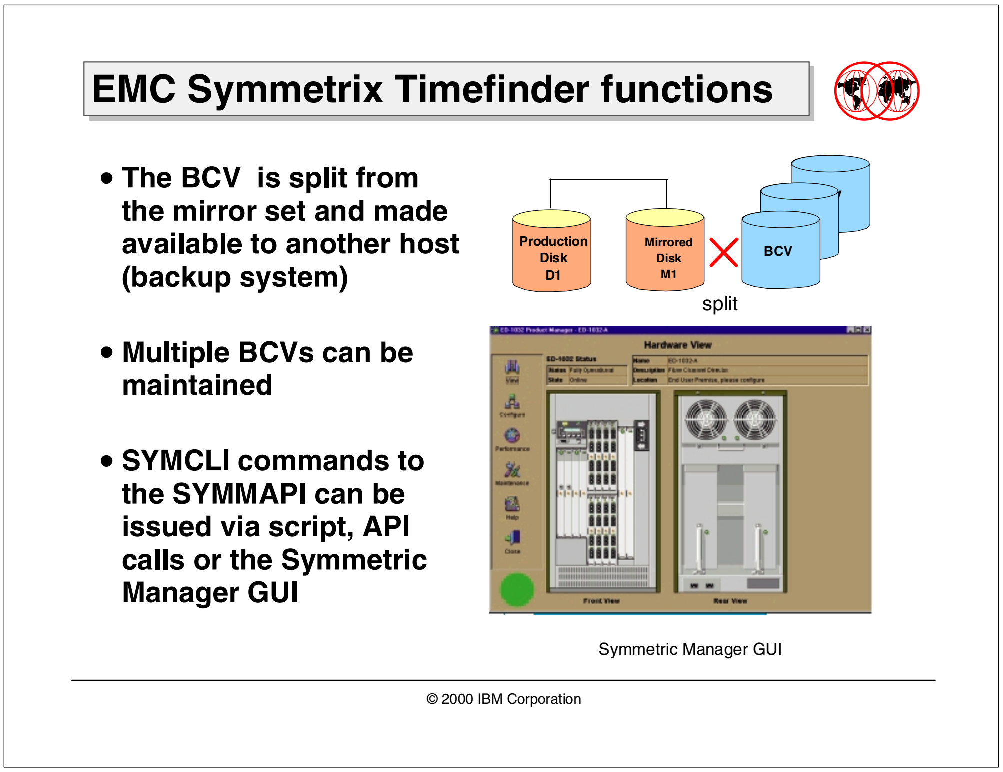

EMC Symmetrix Timefinder functions

- The BCV is split from the mirror set and made available to another host (backup system)
- Multiple BCVs can be maintained
- SYMCLI commands to the SYMMAPI can be issued via script, API calls or the Symmetric Manager GUI

Symmetric Manager GUI

If desired, an Oracle database can also be posted in a *freeze status.* This condition basically avoids the execution of new transactions while performing the split of BCV volumes.

- The mirror-breaking, or BCV splitting, makes the secondary copy available to another host. The backup system is SCSI-connected to the Symmetrix disks and can access data on the BCVs. It can also start a new database instance and execute a backup operation without impacting the performance or availability of the production database server and disks.
- Multiple BCV volumes can be maintained for a single production volume, representing multiple, different point-in-time copies of the production volume.
- SYMMAPI commands can be issued via script, command line, API calls, or the Symmetrix Manager GUI. The BCV split operation can be done only at completion of the BCV establishment process. In other words, the BCV must be synchronized and identical to the production volume before running the SPLIT command.

Mirrored copies on BCV volumes can provide a point-in-time database image on disk. Backups are performed from these copies, thus enabling the database to remain available during the backup operation.

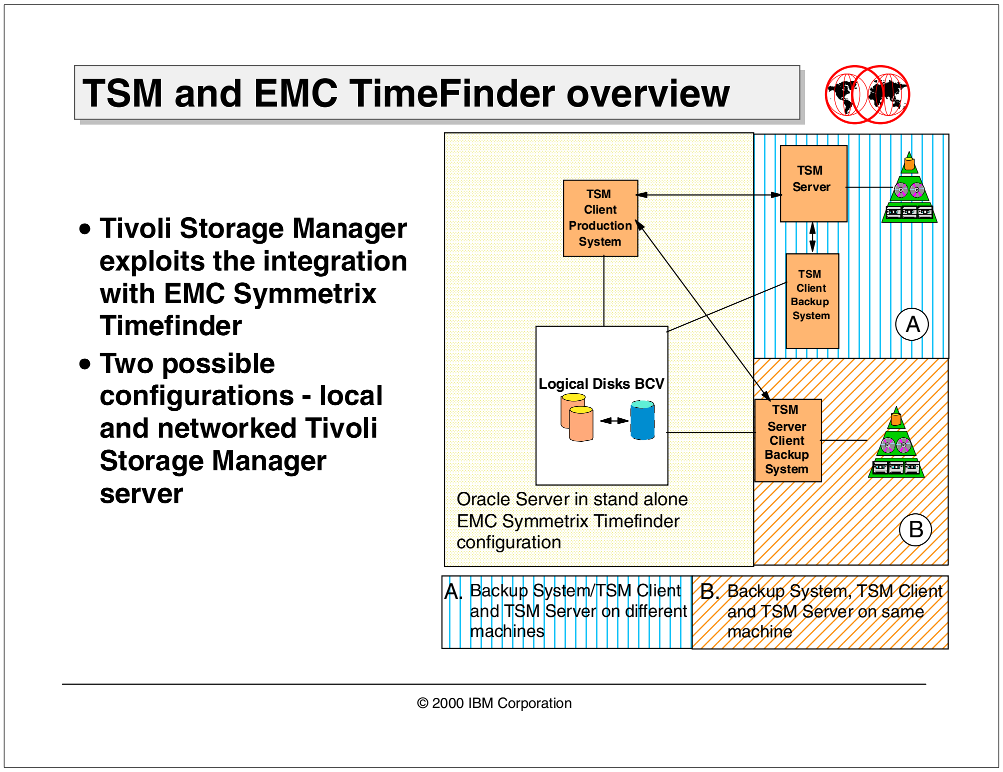

TSM and EMC TimeFinder overview

- Tivoli Storage Manager exploits the integration with EMC Symmetrix Timefinder
- Two possible configurations - local and networked Tivoli Storage Manager server

Oracle Server in stand alone EMC Symmetrix Timefinder configuration

A. Backup System/TSM Client and TSM Server on different machines

B. Backup System, TSM Client and TSM Server on same machine

© 2000 IBM Corporation

9.2.2 TSM and EMC TimeFinder overview

Tivoli Storage Manager is enhanced to provide automated exploitation of EMC TimeFinder copy function for Oracle Database Server.

The backup process of an Oracle Database on a Sun Solaris machine with an EMC Symmetrix storage subsystem can now be efficiently performed with the new function of TSM Server and the new TDP for EMC Symmetrix for Oracle.

The EMC Symmetrix TimeFinder copy function to be implemented requires a new backup server machine SCSI-connected to the storage subsystem, as is the production system. The backup system can be also a TSM client of a TSM Server reachable on the LAN. To minimize data movement of the backup data (from the database) on the network, a more powerful and integrated solution can be implemented by having the backup system and the TSM server on the same host machine.

Note:

If the backup system and the TSM server are collocated on the same machine, they must be implemented on the same operating system platform as the production server.

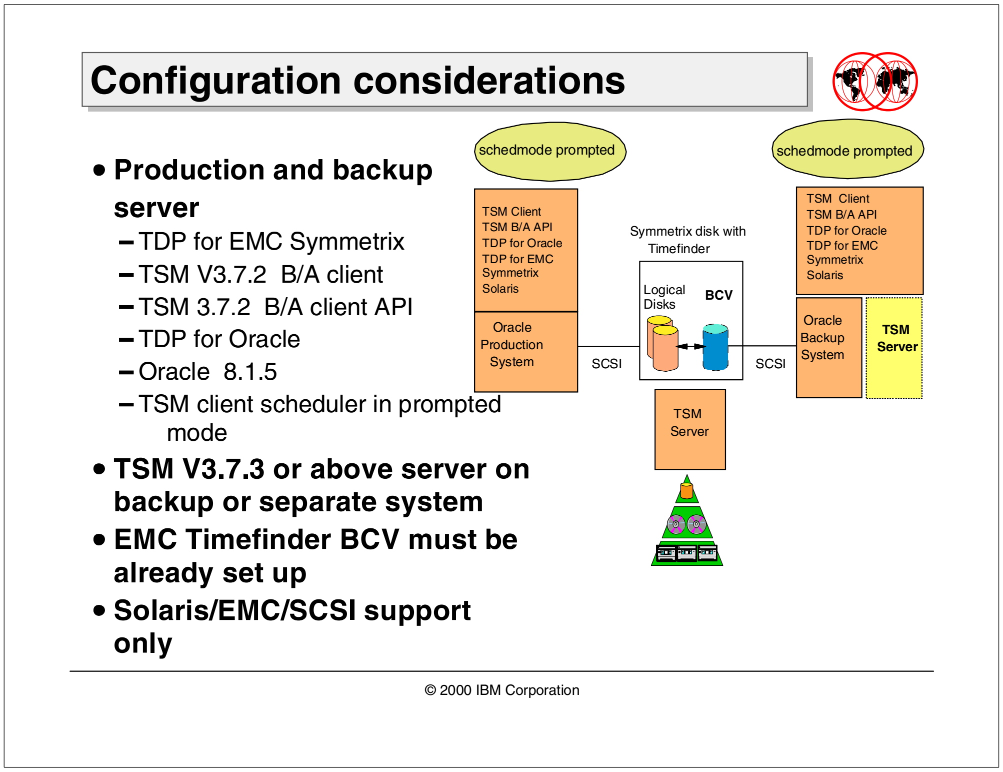

Configuration considerations

- **Production and backup server**
  - TDP for EMC Symmetrix
  - TSM V3.7.2 B/A client
  - TSM 3.7.2 B/A client API
  - TDP for Oracle
  - Oracle 8.1.5
  - TSM client scheduler in prompted mode
- **TSM V3.7.3 or above server on backup or separate system**
- **EMC Timefinder BCV must be already set up**
- **Solaris/EMC/SCSI support only**

© 2000 IBM Corporation

9.2.3 Configuration considerations

Currently, the new TDP for EMC Symmetrix only supports the Sun Solaris operating system environment. Tivoli Storage Manager Server must be at Version 3.7.3 or above, and Tivoli Storage Manager Backup-Archive client and API must be at Version 3.7.2. Both production system and backup system must have the same ability to manage Oracle instances and to back up database data, so they need same the software applications installed.

If an Oracle backup is scheduled on the TSM server, both the production and backup systems need the TSM client scheduler to be active in server-prompted mode. That is, the `schedmode` option in the client system option file (dsm.sys) must be set to prompted.

EMC Symmetrix TimeFinder is required to allow the ability to manage the BCV volumes. Oracle Database Server Version 8.1.5 is needed with support for EMC Symmetrix TimeFinder as well.

BCV volumes must be previously established or in synchronized state. It is the user responsibility to guarantee this condition. The user is also responsible to configure the storage controller so the proper LUNs are available on both machines.

The new TSM server code supports synchronous execution of immediate schedules. This means that the TSM server will not wait (with the `WAIT=YES` option) until the defined schedule session is executed. BCV Split, Oracle instances, management, and backup operations are controlled by the TSM server.

New TSM API application code has been developed and is available with a sample command script. TDP for EMC Symmetrix for Oracle or for SAP R/3 is used to do offline or online backups of the full database copied in the BCVs.

This release will support Veritas File System (VXFS), Veritas Volume Manager and Universal File System (UFS). For the detailed list of supported / unsupported features of these products, please review the readme file shipped with the product.

9.3 Backup process

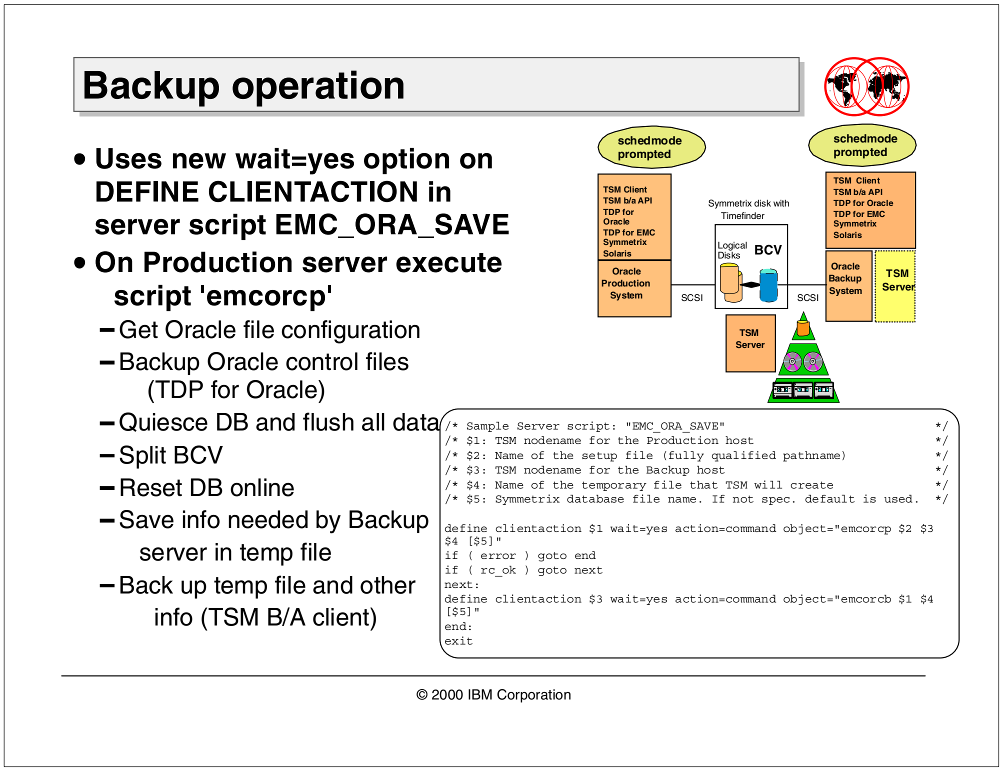

Backup operation

- Uses new wait=yes option on DEFINE CLIENTACTION in server script EMC_ORA_SAVE
- On Production server execute script 'emcorcp'
  - Get Oracle file configuration
  - Backup Oracle control files (TDP for Oracle)
  - Quiesce DB and flush all data
  - Split BCV
  - Reset DB online
  - Save info needed by Backup server in temp file
  - Back up temp file and other info (TSM B/A client)

```
/* Sample Server script: "EMC_ORA_SAVE"                                    */
/* $1: TSM nodename for the Production host                                */
/* $2: Name of the setup file (fully qualified pathname)                   */
/* $3: TSM nodename for the Backup host                                    */
/* $4: Name of the temporary file that TSM will create                     */
/* $5: Symmetrix database file name. If not spec. default is used.        */

define clientaction $1 wait=yes action=command object="emcorcp $2 $3
$4 [$5]"
if ( error ) goto end
if ( rc_ok ) goto next
next:
define clientaction $3 wait=yes action=command object="emcorcb $1 $4
[$5]"
end:
exit
```

A scheduled TSM server command script triggers execution from TSM. Using the TSM server scripting and scheduling enhancements as shown in Section 3.3, “Server scripting and scheduling enhancements” on page 102, synchronous execution of client actions will be available. Therefore, the server command script (shown in Figure 62) can wait for the response of each client action to complete before proceeding.

```
/* Sample Server script: "EMC_ORA_SAVE"                                    */
/* $1: TSM nodename for the Production host                                */
/* $2: Name of the setup file (fully qualified pathname)                   */
/* $3: TSM nodename for the Backup host                                    */
/* $4: Name of the temporary file that TSM will create                     */
/* $5: Symmetrix database file name. If not spec. default is used.        */

define clientaction $1 wait=yes action=command object="emcorcp $2 $3 $4
[$5]"
if ( error ) goto end
if ( rc_ok ) goto next
next:
define clientaction $3 wait=yes action=command object="emcorcb $1 $4 [$5]"
end:
exit
```

Figure 62. Server command script for TDP for EMC for Oracle

The server script can be run as an administrative schedule, or can be executed in an administrative CLI shell with the command:

```
tsm: TSM010> run EMC_ORA_SAVE prodsys /emc/bin/setup backsys EMC
```

1. The first action in the server command script is to execute an immediate client schedule on the production system to invoke the TSM API application `emcorcp`. The API application checks and registers the current mode of the production database.
2. The Oracle database is queried for the logical volumes associated with the requested backup.
3. The logical volumes or file systems are associated with the device serial numbers, and these are matched against the appropriate target devices / BCVs.
4. A backup of database control files is performed.
5. The Oracle database is quiesced, and all buffered data are flushed to the storage controller.
6. The last action executed by the API on the production system is to SPLIT the BCV.

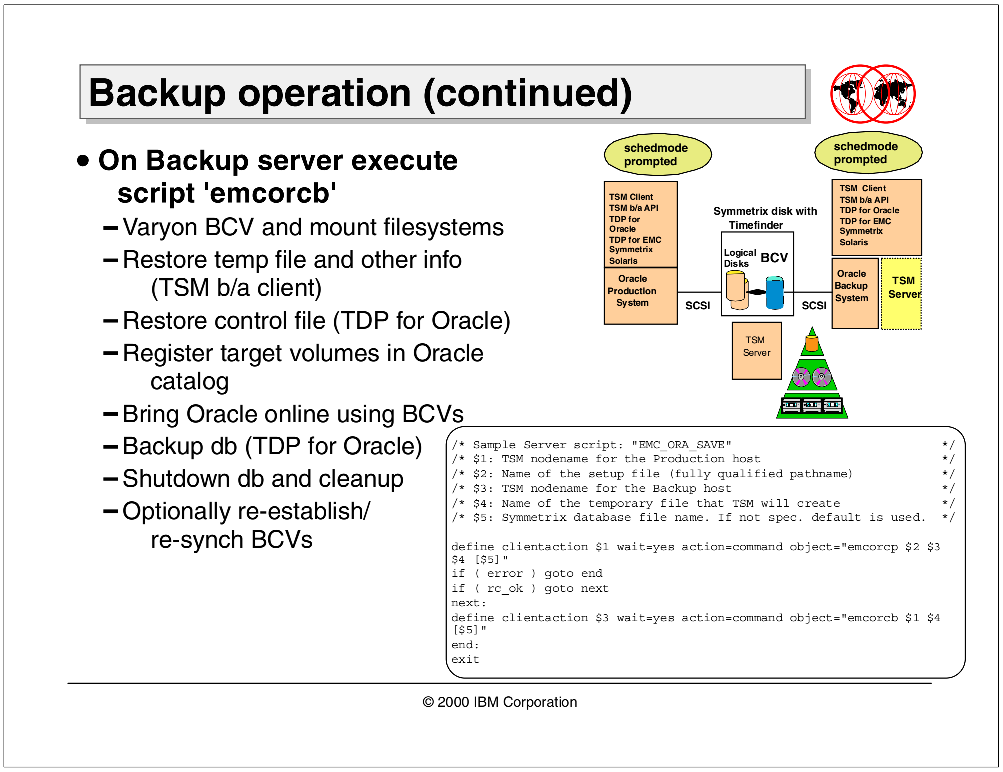

schedmode
prompted
TSM Client
TSM b/a API
TDP for Oracle
TDP for EMC Symmetrix
Solaris
Oracle Production System
SCSI
Symmetrix disk with Timefinder
Logical Disks
BCV
TSM Server
SCSI
schedmode
prompted
TSM Client
TSM b/a API
TDP for Oracle
TDP for EMC Symmetrix
Solaris
Oracle Backup System
TSM Server

Backup operation (continued)

- **On Backup server execute script 'emcorcb'**
  - Varyon BCV and mount filesystems
  - Restore temp file and other info (TSM b/a client)
  - Restore control file (TDP for Oracle)
  - Register target volumes in Oracle catalog
  - Bring Oracle online using BCVs
  - Backup db (TDP for Oracle)
  - Shutdown db and cleanup
  - Optionally re-establish/ re-synch BCVs

```
/* Sample Server script: "EMC_ORA_SAVE"                                */
/* $1: TSM nodename for the Production host                            */
/* $2: Name of the setup file (fully qualified pathname)               */
/* $3: TSM nodename for the Backup host                                */
/* $4: Name of the temporary file that TSM will create                 */
/* $5: Symmetrix database file name. If not spec. default is used.     */

define clientaction $1 wait=yes action=command object="emcorcp $2 $3
$4 [$5]"
if ( error ) goto end
if ( rc_ok ) goto next
next:
define clientaction $3 wait=yes action=command object="emcorcb $1 $4
[$5]"
end:
exit
```

After the mirror is broken (we assume to receive a positive return code from the production system script) the TSM server will initiate another immediate action to execute the API application `emcorcb`.

The application will do the following:

1. Mount the target volumes, start the backup database in mount mode.
2. Login to the Oracle catalog database to register the copy target volumes.
3. Use TDP for Oracle or TDP for SAP R/3 to restore the database control files.
4. Back up the database from the mirror copy and store data on the TSM Server.
5. At the completion of the database backup, the TSM backup system will shut down the backup database and return control of the process to the TSM Server.
6. As an optional final step, the TSM API application on the backup system can be instructed to re-sync the BCV.

9.4 Local versus network backup and restore

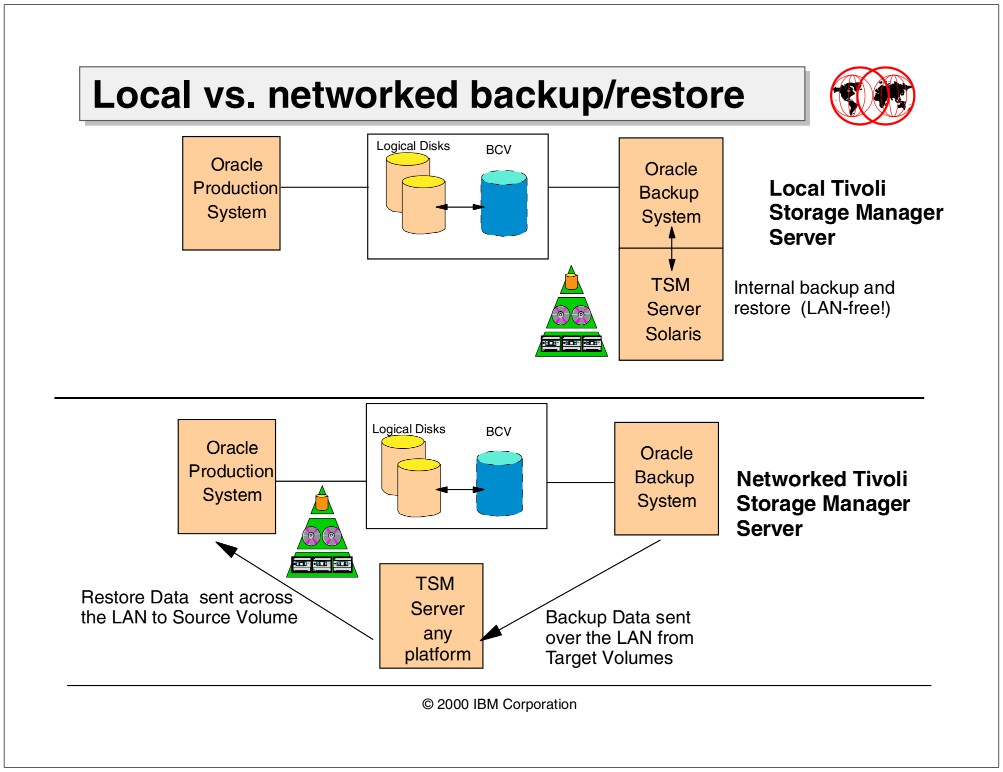

Local vs. networked backup/restore

Oracle Production System

Logical Disks

BCV

Oracle Backup System

TSM Server Solaris

Local Tivoli Storage Manager Server

Internal backup and restore (LAN-free!)

Oracle Production System

Logical Disks

BCV

Oracle Backup System

Networked Tivoli Storage Manager Server

Restore Data sent across the LAN to Source Volume

TSM Server any platform

Backup Data sent over the LAN from Target Volumes

© 2000 IBM Corporation

In the network scenario, the backup system and the TSM server system are two separate machines. The backup system is able to access data on BVC disks and is also a TSM client system. In this configuration Tivoli Data Protection for Oracle will perform the database backup to the TSM server over the network.

An efficient solution to speed up the backup operations is to have the TSM server installed on the backup system. In this configuration, the database BCV volumes are local to the TSM server, so no data is sent across the network. The database data files are backed up using the SCSI connection between the backup system and the BCV volume in the storage unit. The disadvantage of this solution is that the TSM server has to be a Solaris server.

If the TSM server is on a LAN attached host, the restore is typically done across the LAN to the source volumes, as that is significantly simpler and safer. The TDP specific agent will be used to restore the database at the desired point in time. This is a standard database restore in a network scenario, and the duration of the operation depends on the network speed.

Restore directly to the disk subsystem avoids data movement over the LAN. This can be done if the TSM server is installed on the backup system. Restore could be done to source volumes, or to the target volumes, but it is more complex and manual.

For restore to source volumes, the backup system needs to:

- Establish the access to those volumes

For restore to target volumes (BCV), the backup system needs to:

- Validate the relationship of source volumes to target volumes,
- Verify that meta data on target volumes has not changed.
- Verify that additional data has not been written to target volumes.

These manual steps will be addressed via documented recovery guidelines in the User's Manuals. Limited support will be offered.

Restore will be done under database administrator control via standard Oracle commands using RMAN and TDP for Oracle or SAPDBA and TDP for R/3, so the same degree of manual intervention as ordinarily required in this environment will be applicable.

Typically, a database restore and recovery will involve:

- A full database restore (from tape) or by re-syncing/copying from a previous point-in-time target volume copy.
- Restore and application (roll-forward recovery) of archived redo logs to roll forward from the point-in-time mirrored copy to the desired point-in-time.
- These steps will be accomplished as they ordinarily are with the RMAN and SAPDBA interfaces.

Another design point to consider is whether one should elect to retain multiple, different point-in-time broken mirror copies to facilitate restore back to different points in time without having to recover a full backup from tape (the so-called "one minute restore").

Chapter 10. TDS for Storage Management Analysis

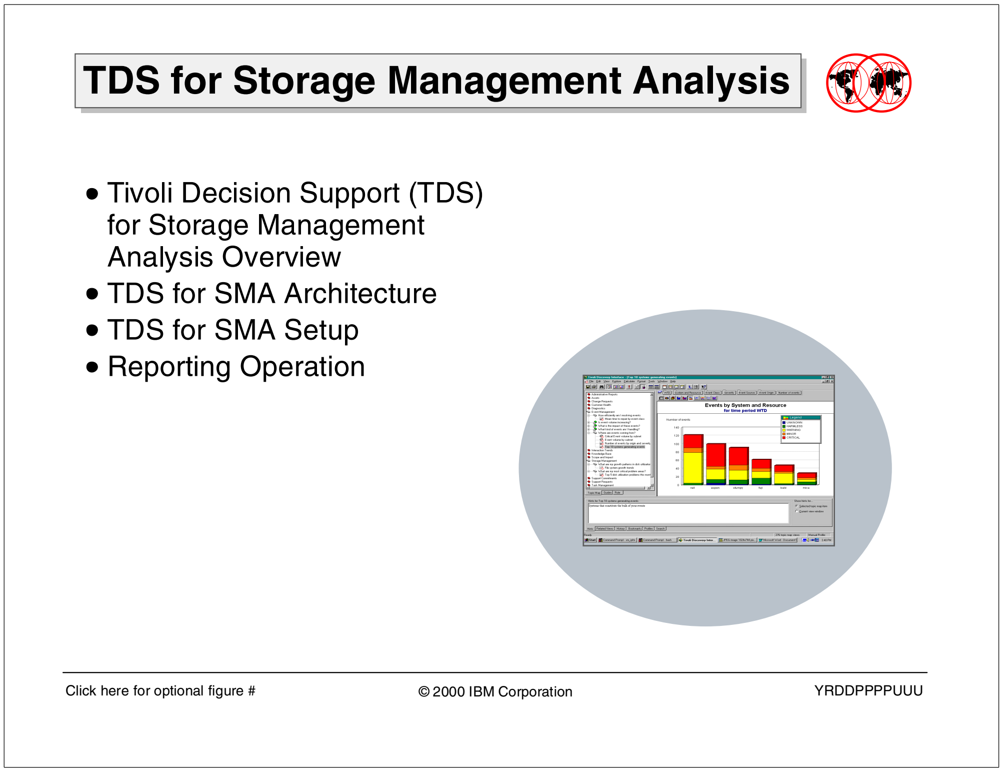

TDS for Storage Management Analysis

- Tivoli Decision Support (TDS) for Storage Management Analysis Overview
- TDS for SMA Architecture
- TDS for SMA Setup
- Reporting Operation

Click here for optional figure #

© 2000 IBM Corporation

YRDDPPPPUUU

Tivoli Decision Support for Storage Management Analysis is a new product of the Tivoli Storage Management product set. This new product provides tools and interfaces for the Tivoli Decision Support product to analyze and display data extracted from the Tivoli Storage Manager server. It is used to obtain snapshot and exception reports, and to verify the trends for historical data and forecasting of the storage management environment. Automation of the report creation is possible.

This chapter will cover:

- An introduction to reporting using Tivoli Decision Support
- The components, architecture, reporting, and data transfer model of Tivoli Decision Support for Storage Management Analysis
- The setup steps of Tivoli Decision Support for Storage Management Analysis
- Considerations when using the product for reporting

10.1 Introduction to reporting using TDS

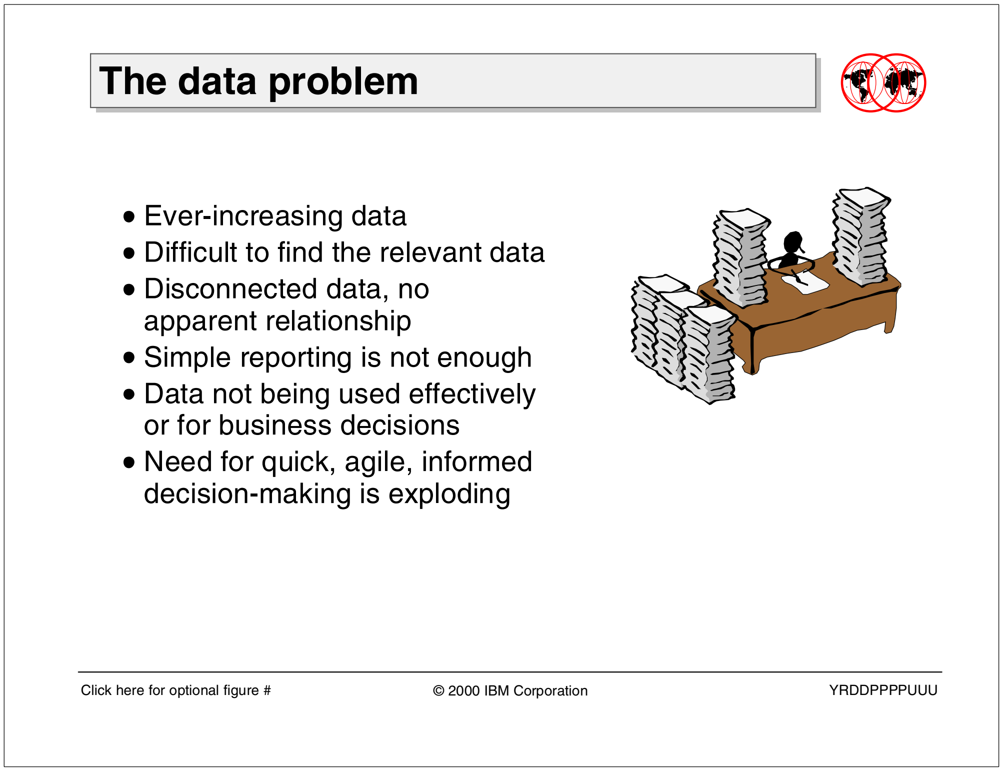

The data problem

- Ever-increasing data
- Difficult to find the relevant data
- Disconnected data, no apparent relationship
- Simple reporting is not enough
- Data not being used effectively or for business decisions
- Need for quick, agile, informed decision-making is exploding

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

10.1.1 The data problem

In the past few years, the ever-increasing growth of data has been enormous. The amount of information stored in network environments is doubling at least annually; the difficulty of finding the relevant data to maintain or control systems gets harder by the day; simple reporting is no longer sufficient; and the current data reports are not being used effectively to make intelligent business decisions.

The need for quick, agile, informed decision-making is exploding — and that is what TDS does best.

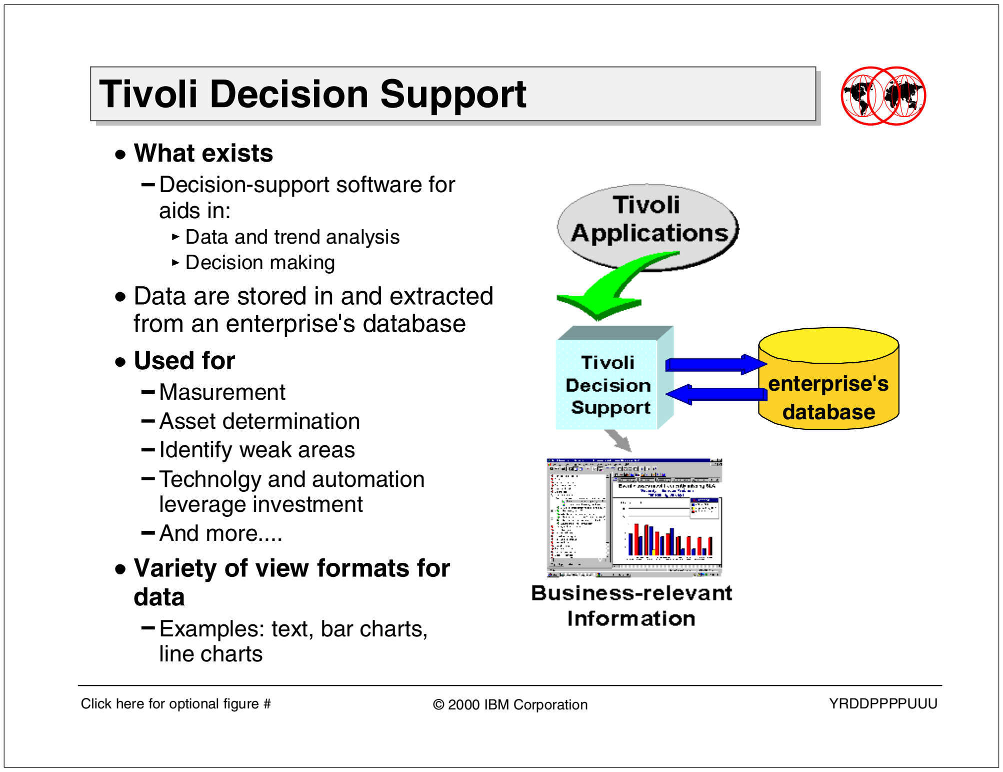

Tivoli Decision Support

- **What exists**
  - Decision-support software for aids in:
    - Data and trend analysis
    - Decision making
- Data are stored in and extracted from an enterprise's database
- **Used for**
  - Masurement
  - Asset determination
  - Identify weak areas
  - Technolgy and automation leverage investment
  - And more....
- **Variety of view formats for data**
  - Examples: text, bar charts, line charts

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

10.1.2 Tivoli Decision Support

Tivoli Decision Support (TDS) is a decision-support software that aids analysis and decision making in an organization.

Tivoli Decision Support enables users to find and use the data stored in an enterprise’s database. It enables business to extract the gathered data information, and to build content automatically, ready for the users.

Tivoli Decision Support is a *decision-support* framework, made by components that aid in finding and analyzing data.

By using Tivoli Decision Support to tap into your operational data, you can, for example:

- Measure the effectiveness of your operation.
- Determine how assets are distributed across the enterprise.
- Identify areas of weakness, to convert from reactive support activities to proactive support planning.
- Leverage your investment in technology and automation.

The data can be viewed in a variety of formats. It is possible to have summary viewing or detailed aspects, presented as text, bar charts, line charts, or other graphical formats.

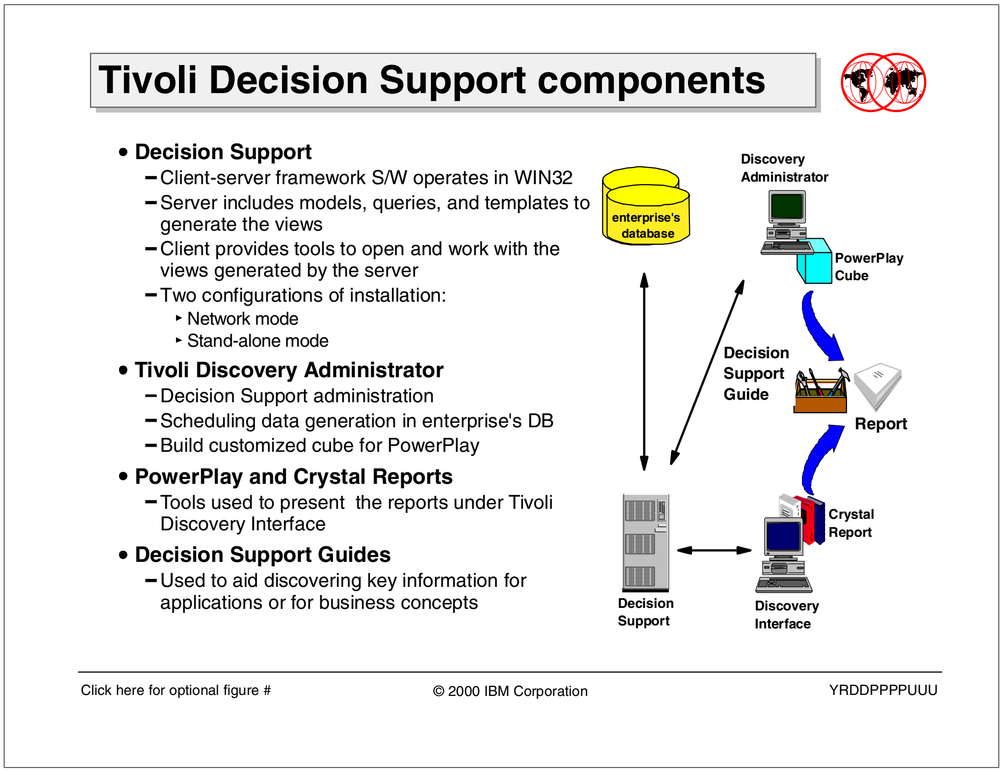

Tivoli Decision Support components

- **Decision Support**
  - Client-server framework S/W operates in WIN32
  - Server includes models, queries, and templates to generate the views
  - Client provides tools to open and work with the views generated by the server
  - Two configurations of installation:
    - Network mode
    - Stand-alone mode
- **Tivoli Discovery Administrator**
  - Decision Support administration
  - Scheduling data generation in enterprise's DB
  - Build customized cube for PowerPlay
- **PowerPlay and Crystal Reports**
  - Tools used to present the reports under Tivoli Discovery Interface
- **Decision Support Guides**
  - Used to aid discovering key information for applications or for business concepts

Click here for optional figure #

© 2000 IBM Corporation

YRDDPPPPUUU

10.1.3 Tivoli Decision Support components

The components that form Tivoli Decision Support are:

- **Tivoli Decision Support**: Client-server framework software operates in a Windows 32-bit environment, Windows 95, 98, or NT.

  The server component includes the models, queries, templates, and other information required to generate the views.

  The client component, also called **Tivoli Discovery Interface**, provides all the tools needed to open and work with the views of data from the enterprise's database.

  Depending on the configuration installation, the server component could be installed in a server machine, network mode, or in all the machines used by the clients component, stand-alone mode.

- **Tivoli Discovery Administrator**: This module is used to administer Decision Support, schedule data generation processes in the enterprise's database and build customized cubes, interfacing the
- **Cognos PowerPlay**: This component is used to generate multidimensional cubes, used by the Tivoli Discovery Interface.
- **Crystal Reports**: This component is another tool used to present the reports under the Tivoli Discovery Interface.
- **Decision Support Guides**: These are used by Decision Support to aid in discovering key information for applications or for a business concept, for example call center management. They provide a comprehensive set of views and methods for the extraction of the data from the enterprise's database. The Discovery Guides are used with Tivoli Discovery Interface.

10.2 TDS for Storage Management Analysis

.

TDS for Storage Management Analysis

- Best practice guides to analyze & display data from Tivoli Storage Manager
  - Event & Performance analysis
    - **Verify the operations and establish improvement and correction**
    - **Servers and clients**
  - Capacity planning in future releases
- RDBMS is used as a warehouse of Tivoli Storage Manager data, historical, trends & forecast
- Data displayed using TDS
  - Graphical and text formats
- Replaces Sterling's SAMS Vantage: ADSM edition

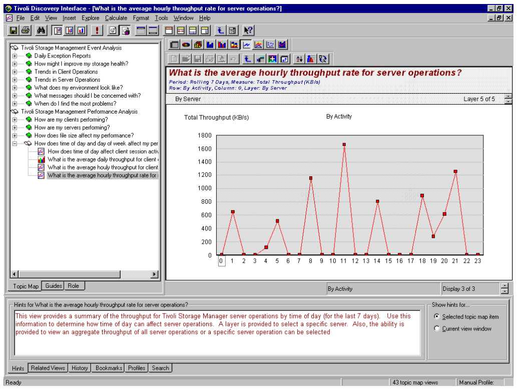
Tivoli Discovery Interface - [What is the average hourly throughput rate for server operations?]
File Edit View Insert Explore Calculate Format Tools Window Help
Tivoli Storage Management Event Analysis
Daily Exception Reports
How might I improve my storage health?
Trends in Client Operations
Trends in Server Operations
What does my environment look like?
What messages should I be concerned with?
When do I find the most problems?
Tivoli Storage Management Performance Analysis
How are my clients performing?
How are my servers performing?
How does file size affect my performance?
How does time of day and day of week affect my per
How does time of day affect client session activ
What is the average daily throughput for client
What is the average houly throughput for client
What is the average hourly throughput rate for
Topic Map Guides Role
What is the average hourly throughput rate for server operations?
Period: Rolling 7 Days, Measure: Total Throughput (KB/s)
Row: By Activity, Column: 0, Layer: By Server
By Server
Layer 5 of 5
Total Throughput (KB/s)
By Activity
1800
1600
1400
1200
1000
800
600
400
200
0
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23
By Activity
Display 3 of 3
Hints for What is the average hourly throughput rate for server operations?
Show hints for...
● Selected topic map item
○ Current view window
Hints Related Views History Bookmarks Profiles Search
Ready
43 topic map views
Manual Profile:

This view provides a summary of the throughput for Tivoli Storage Manager server operations by time of day (for the last 7 days). Use this information to determine how time of day can affect server operations. A layer is provided to select a specific server. Also, the ability is provided to view an aggregate throughput of all server operations or a specific server operation can be selected

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

The Tivoli Decision Support for Storage Management Analysis product provides "best-practices" guides to analyze data and display reports for the following Tivoli Storage Manager main areas:

- Event analysis
- Performance analysis

The purpose of the Event Analysis Guide is to gain insight into the health of a Tivoli Storage Manager environment, showing how well the environment is operating and what area needing improvement. The reports are generated for both Tivoli Storage Manager server and clients.

The purpose of the Performance Analysis Guide is to gain insight into the performance of the environment, showing how well the environment is performing, and what areas need improvement. Also, performance reports are generated for both server and clients.

Future Storage Management Analysis Guide releases will address other questions, particular in the area of Tivoli Storage Management capacity planning.

A relational database is used to warehouse the Tivoli Storage Manager data, historical and forecast. The information can then be displayed in a variety of formats, graphical and text, and it is possible to drill down from the summaries to get detailed aspects of an environment.

Tivoli Decision Support for Storage Management Analysis replaces Sterling's SAMS Vantage: ADSM edition product, used with ADSM V3.1 for reporting purposes.

The following section will discuss the architecture, the main components, and the data transfer model of Tivoli Decision Support for Storage Management Analysis.

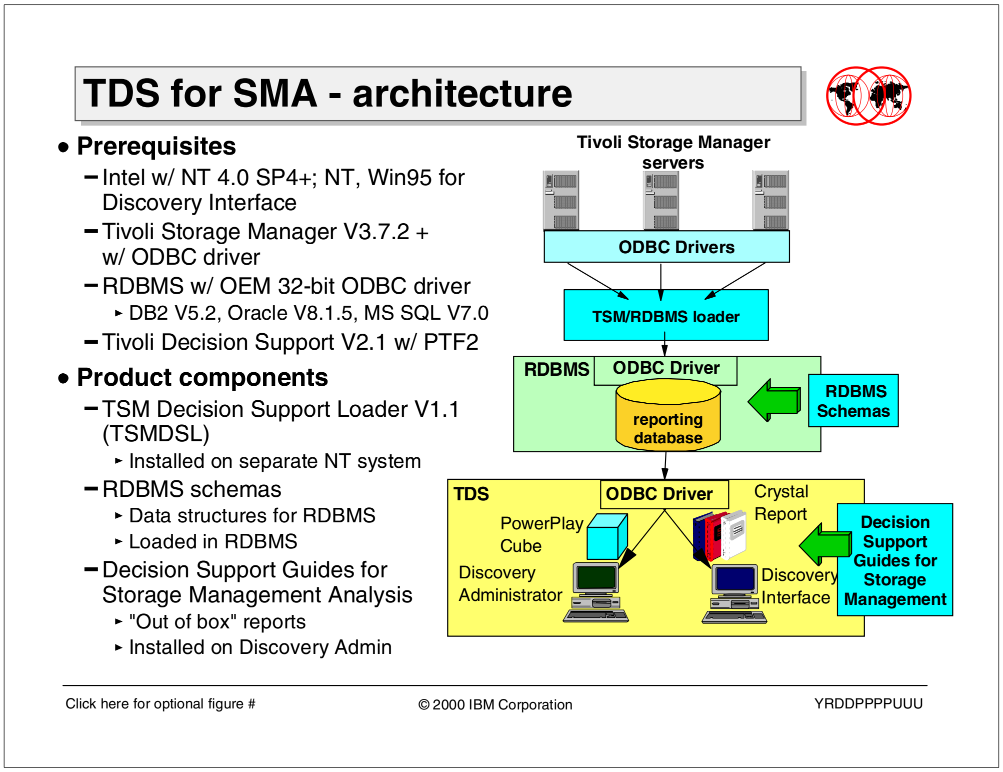

TDS for SMA - architecture

- **Prerequisites**
  - Intel w/ NT 4.0 SP4+; NT, Win95 for Discovery Interface
  - Tivoli Storage Manager V3.7.2 + w/ ODBC driver
  - RDBMS w/ OEM 32-bit ODBC driver
    - DB2 V5.2, Oracle V8.1.5, MS SQL V7.0
  - Tivoli Decision Support V2.1 w/ PTF2
- **Product components**
  - TSM Decision Support Loader V1.1 (TSMDSL)
    - Installed on separate NT system
  - RDBMS schemas
    - Data structures for RDBMS
    - Loaded in RDBMS
  - Decision Support Guides for Storage Management Analysis
    - "Out of box" reports
    - Installed on Discovery Admin

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

10.2.1 TDS for Storage Management Analysis — architecture

The Tivoli Decision Support for Storage Management Analysis product consists of the following components:

1. **Storage Manager Decision Support Loader Tool** to provide the function of the roll-up module. The module can be installed on any NT machine connected to the Tivoli Storage Manager servers and the reporting database machine with TCP/IP.
2. **RDBMS schemes** for the reporting database to create the data structure in the reporting database. The schemes must be loaded in the reporting database.
3. **Decision Support Guide for Storage Management Analysis** to provide the queries for Event and Performance reports. The guides must be imported into Discovery Administration

However, there are more components that make up a complete Tivoli Decision Support for Storage Management Analysis reporting solution. These base components are:

- Tivoli Storage Manager
- Tivoli Storage Manager / RDBMS interface
- A Relational Database System
- Tivoli Decision Support

10.2.1.1 Tivoli Storage Manager servers

The Tivoli Storage Manager database is the source of information used for event and performance information.

Version 3.7.2 or later of the Tivoli Storage Management server is required. This is because data is extracted from the Tivoli Storage Manager server from new tables included in this version. They are the SUMMARY and EVENTS tables. Therefore, this version is required to provide support for Storage Management Analysis information gathering.

You will have at least one Tivoli Storage Management database source, but you may have many of them. They can be on multiple processing platforms. These will be known to Tivoli Decision Support and Storage Management Analysis as source data servers. That will become important when setting up the parameters for the loader program, which we will discuss in the next section.

10.2.1.2 Tivoli Storage Manager / RDBMS interface

The major portion of this interface is the Decision Support Loader program. Its purpose is to move data from the Tivoli Storage Management source database to the Tivoli Decision Support reporting database. The recommendation is to install the Decision Support Loader on a dedicated NT machine. Running the decision Support Loader and the Tivoli Decision Support product on the same processor can adversely impact performance.

This program uses the Tivoli Storage Manager ODBC interface to read data from the Tivoli Storage Manager server and uses the RDBMS ODBC interface to save data into the Tivoli Storage Manager Reporting Database. The basic steps performed by the program during the data collection process are:

- Extract data from the Tivoli Storage Manager server using the Tivoli Storage Manager ODBC driver
- Transform data into the format required by the Tivoli Storage Manager reporting database. This includes processing such as providing a unique date/time key for the information extracted from the Tivoli Storage Manager server.
- Write data to the Tivoli Storage Manager Reporting database using the RDBMS ODBC driver.

The program also controls the expiration of old records from the RDBMS.

You can concurrently transfer data from multiple Tivoli Storage Management servers to a shared RDBMS database server by installing the Decision Support Loader on multiple machines. To use this configuration, you must set up each Tivoli Storage Management server as an ODBC data source, and configure each Decision Support Loader to access the appropriate Tivoli Storage Management server and shared RDBMS database server. See the Tivoli Storage Management Decision Support Loader Release notes Version 1.1.1 for more detail.

10.2.1.3 Relational database system

The source information is saved in a relational database called the Tivoli Storage Manager reporting database. One Tivoli Storage Manager reporting database can be used to contain data from multiple Tivoli Storage Manager servers. The Tivoli Storage Manager reporting databases that are initially supported are:

- Oracle (V8.1.5)
- DB2(V5.2)
- MS SQL(v7.0)

A schema is provided with Tivoli Decision Support for Storage Management Analysis that contains the definitions of the tables used in the RDBMS for these three systems.

10.2.1.4 Tivoli Decision Support

Tivoli Decision Support includes several components as already explained under Section 10.1.3, "Tivoli Decision Support components" on page 302. You must decide which components to install and on which machines based on how you plan to run the application. You will choose to run in either stand-alone mode or network mode.

You must also install Tivoli Decision Support Version 2.1 and PTF2 for Tivoli Decision Support to have Storage Management Analysis function properly.

The Storage Management Analysis Guide delivered with the Tivoli Decision Support for Storage Management Analysis product consists of:

- **"Out-of-box" reports** — This sets up the reporting links to the reporting data base, Cognos, and Crystal. This is done via Tivoli Decision Support and its GUI interface, which lets you access these reports with multiple viewing options.
- **Cube reporting** — This component sets up reporting from Cognos 3D graphics reporting.

Tivoli Decision Support and Storage Management Analysis must reside on a Windows NT 4.0 SP4 level platform, or better. Tivoli Decision Support operates on Windows 95, Windows 98, and Windows NT 4. Note that Tivoli Decision Support does **not** require the Tivoli Framework.You may have multiple Tivoli Decision Support servers in the same reporting system.

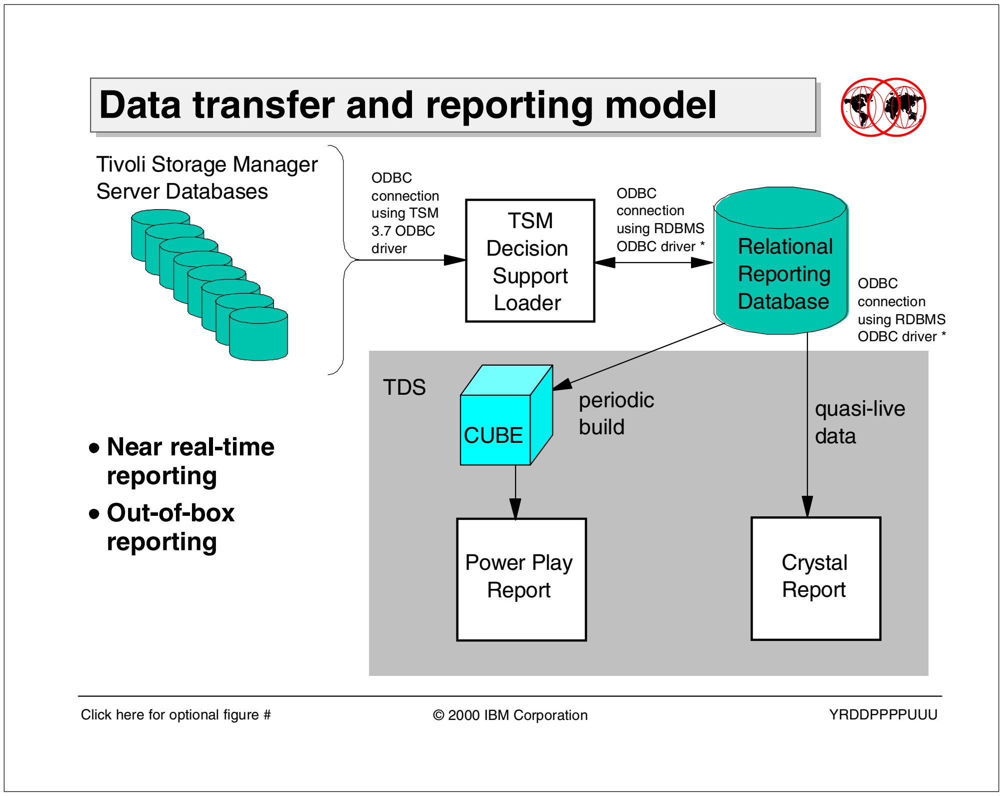

Data transfer and reporting model

Tivoli Storage Manager Server Databases

ODBC connection using TSM 3.7 ODBC driver

TSM Decision Support Loader

ODBC connection using RDBMS ODBC driver *

Relational Reporting Database

ODBC connection using RDBMS ODBC driver *

TDS

CUBE

periodic build

quasi-live data

Power Play Report

Crystal Report

- **Near real-time reporting**
- **Out-of-box reporting**

Click here for optional figure #

© 2000 IBM Corporation

YRDDPPPPUUU

10.2.2 Data transfer and reporting model

The information is extracted from the Tivoli Storage Manager server and transferred to a reporting relational database. An opposite roll-up module has been built to interface both the servers and the reporting database. The interfaces used by the roll-up module to interconnect with the servers and the relational database are:

- A specific ODBC driver installed and operating on the Tivoli Storage Manager servers
- The ODBC driver provided by the RDBMS product used to collect the data.

To extract the data from the Tivoli Storage Manager servers, the roll-up module must use a specific Tivoli Storage Manager administrator with the privilege level ANALYST.

To build the reports, the collect data are queried from the RDBMS using the specific guides Cognos PowerPlay and Crystal Reports. RDBMS schemes are provided to create the appropriate structure. To query the data from the relational database and obtain the reports, an ODBC connection is used in the Tivoli Decision Support.

Cognos PowerPlay is imbedded and installed as an integrated and core part of Tivoli Decision Support. The Discovery Interface uses Cognos' PowerPlay reports. The Cognos PowerPlay cubes provide numerous dimensions and measures that allow the user to view data in ways additional to the pre-defined views shipped with the Guide.

Seagate Crystal is Imbedded and installed as an integrated and core part of Tivoli Decision Support. The Discovery Interface uses Seagate's Crystal Reports to allow the user to view data in additional ways than the pre-defined views shipped with the Guide.

The reports that are generated are not in a real time or ‘live’ mode. The Storage Management Analysis mode used is known as ‘near real-time’. When you pull the data from the RDBMS, the reports will reflect data up to the point in time of the last data load.

Tivoli Decision Support and Storage Management Analysis supplies comprehensive pre-coded reports available upon install. In fact, Storage Management Analysis, as a Tivoli Decision Support Guide, is designed to ‘not’ be a custom solution. The complexity of a six-product solution (Tivoli Storage Manager, reporting database system, Tivoli Decision Support, Storage Management Analysis, Cognos PowerPlay, and Seagate Crystal Reporting) allows little room for custom reporting and introduces a great deal of risk if attempted.

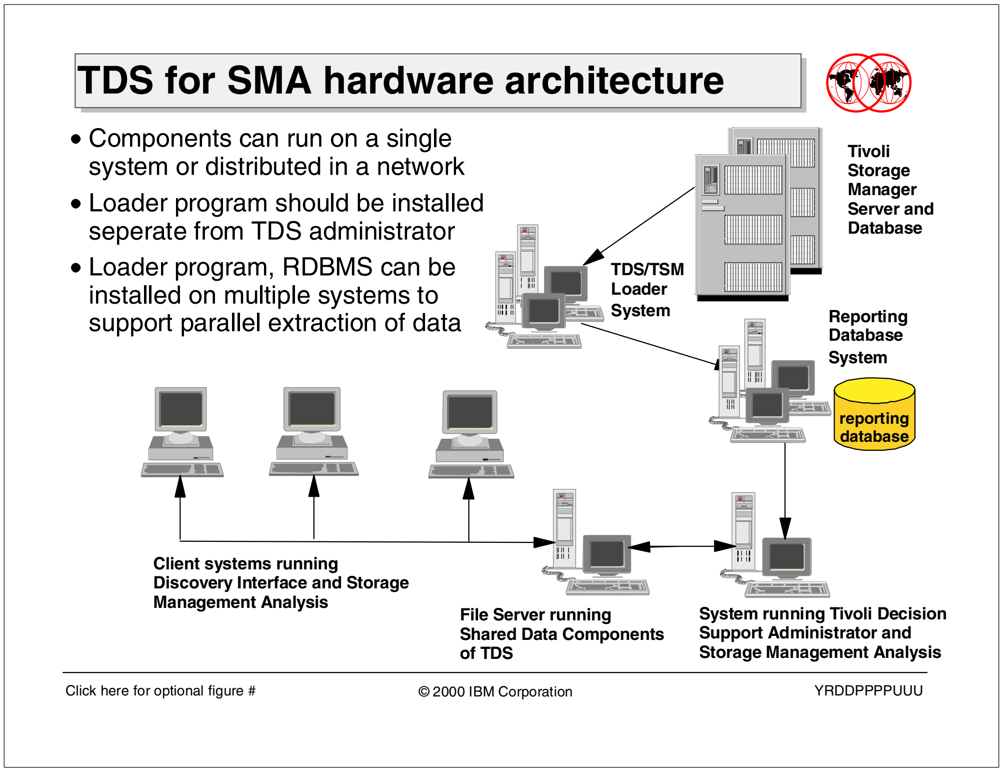

TDS for SMA hardware architecture

- Components can run on a single system or distributed in a network
- Loader program should be installed seperate from TDS administrator
- Loader program, RDBMS can be installed on multiple systems to support parallel extraction of data

Tivoli Storage Manager Server and Database

TDS/TSM Loader System

Reporting Database System

reporting database

Client systems running Discovery Interface and Storage Management Analysis

File Server running Shared Data Components of TDS

System running Tivoli Decision Support Administrator and Storage Management Analysis

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

10.2.3 TDS for SMA hardware architecture

There are multiple configurations related to TDS, the loader program, the RDBMS and the Tivoli Storage Manager servers. For example, the loader can run on the same machine as the RDBMS or on a different machine. Also, there can be multiple loader programs running on separate machines (extracting data from different Tivoli Storage Manager servers) and loading data into the same RDBMS. These different configurations can affect performance.

One of the first things to plan for is how many workstations need access to reports and how often. If there are multiple reporting workstations, you may want to consider a dedicated file server for reports, as described in the Tivoli Decision Support Installation Guide.

If you have a relatively small Tivoli Storage Management system or a system which is accessed by a single Tivoli Storage Management Administrator you may need only one machine configured for Tivoli Decision Support and Storage Management Analysis.

However, you may want only the Tivoli Storage Management Administrator to use Tivoli Decision Support, but then provide for multiple Storage Management Analysis machines as shown on the slide.

These factors will dictate what you decide on using stand-alone or network mode for Tivoli Decision Support. After that, you are ready to look at the remaining issues. Once the network or stand-alone decision made for Tivoli Decision Support, Tivoli Storage Management database size, network configuration, and reporting database size may be the major determinants in deciding how many machines should be configured.

We recommend that you install the Decision Support Loader on a dedicated workstation. Running the decision Support Loader and the Storage Management Analysis product on the same processor can adversely impact performance.

10.2.3.1 Access frequency and the network

Consider the approximate number of reports a given user of the system is likely to build on download from the file server or reporting database. Each report varies in the amount of data, depending on the content and the volume of elements included. Nevertheless, response will be poor if data traffic or bandwidth precludes rapid transfer. You may want to run some initial benchmarks to see exactly what impact your particular reporting needs have on your network.

Consider also, the impact of loader executions on the network. Again, we recommend that you install the Decision Support Loader on a dedicated workstation. How often you run this program will have a impact on the amount of data transferred and the time of day may have an impact on data traffic. Some Administrators may want to run on this on an hourly basis, but more often it will be run on a daily basis.

Avoid running this program when other activity such as backups and archives may be accessing the Tivoli Storage Management database. Also, avoid running building cubes during `tsmdsl` execution. Both of these activities can adversely affect performance.

10.2.3.2 Network configuration

If you plan to report from multiple Tivoli Storage Management machines to a common RDBMS machine you probably want to install `tsmdsl` on multiple machines, as well. In this way, you may do concurrent downloads of data from Tivoli Storage Management with a minimum of impact on a given loader processor. However, keep in mind the impact on the RDBMS server itself and the network paths to it.

10.2.3.3 System sizing

The RDBMS and loader machines themselves have minimum configuration specifications. However, depending on usage other processor components can be stressed. Consider I/O performance, CPU utilization, memory usage, and data storage availability when planning for Storage Management Analysis or trouble shooting problems.

If too many Storage Management Analysis reports are open simultaneously on the Discovery Interface system, this can effect memory and cause unpredictable results. Also, in the case of cube reporting, memory on the PC video card can be restrictive. This situation can be unpredictable, as well. Generally, these situations can be isolated by running only the failing set(s) of reports and noting results.

10.3 TDS for Storage Management Analysis setup

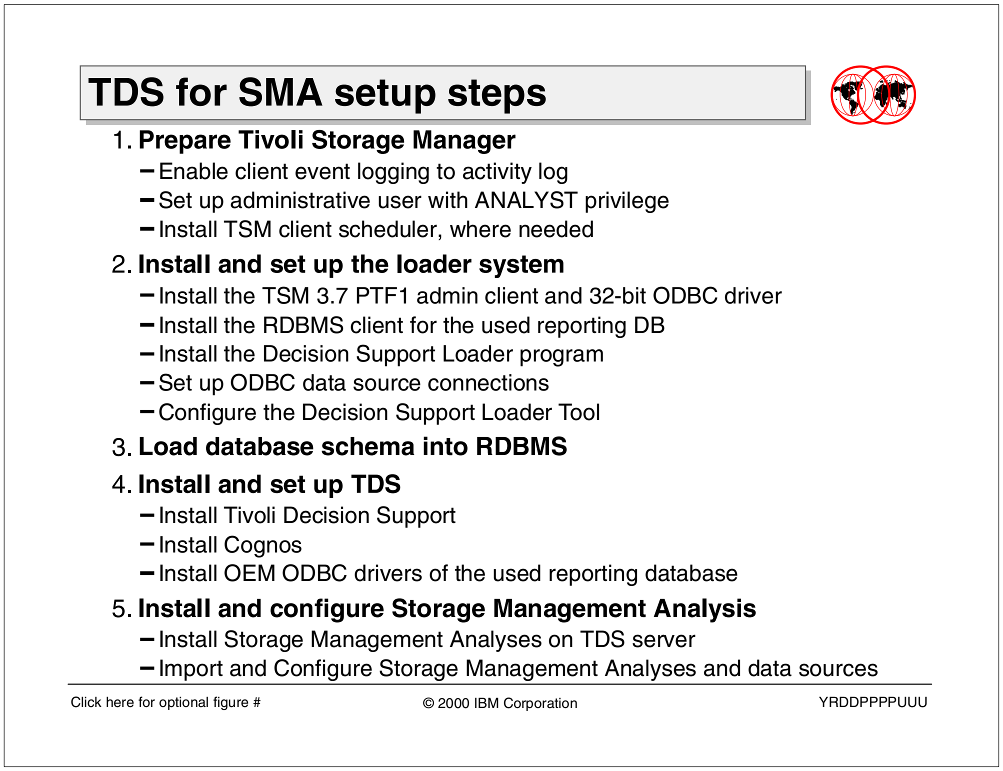

TDS for SMA setup steps

1. **Prepare Tivoli Storage Manager**
   - Enable client event logging to activity log
   - Set up administrative user with ANALYST privilege
   - Install TSM client scheduler, where needed
2. **Install and set up the loader system**
   - Install the TSM 3.7 PTF1 admin client and 32-bit ODBC driver
   - Install the RDBMS client for the used reporting DB
   - Install the Decision Support Loader program
   - Set up ODBC data source connections
   - Configure the Decision Support Loader Tool
3. **Load database schema into RDBMS**
4. **Install and set up TDS**
   - Install Tivoli Decision Support
   - Install Cognos
   - Install OEM ODBC drivers of the used reporting database
5. **Install and configure Storage Management Analysis**
   - Install Storage Management Analyses on TDS server
   - Import and Configure Storage Management Analyses and data sources

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

To set up a complete TDS for Storage Management Analysis solution, multiple setup steps must be performed:

1. Preparation of Tivoli Storage Manager for TDS reporting
2. Installation and setup of the Decision Support Loader program
3. Loading of the TDS for SMA schema into the selected reporting database
4. Installation and setup of the TDS product
5. Installation and setup of Storage Management Analysis guides into TDS using the administrator tools

Depending on the hardware architecture decision you have made, according to the considerations presented in Section 10.2.3, “TDS for SMA hardware architecture” on page 310, you might have to perform some tasks on different systems and/or multiple times.

For more information on the exact tasks that need to be performed, consult product documentation and check the IBM Redbook *Tivoli Storage Management Reporting*, SG24-6109.

In the following sections, we will outline some installation “gotchas” and discuss some automation aspects.

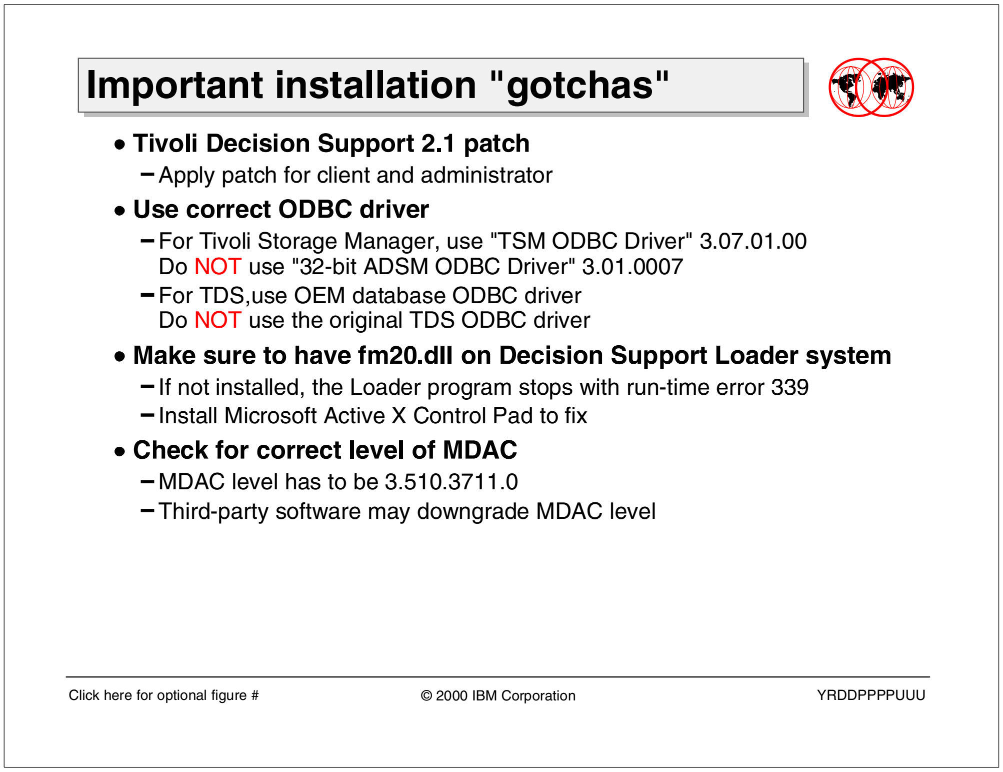

Important installation "gotchas"

- **Tivoli Decision Support 2.1 patch**
  - Apply patch for client and administrator
- **Use correct ODBC driver**
  - For Tivoli Storage Manager, use "TSM ODBC Driver" 3.07.01.00
    Do NOT use "32-bit ADSM ODBC Driver" 3.01.0007
  - For TDS,use OEM database ODBC driver
    Do NOT use the original TDS ODBC driver
- **Make sure to have fm20.dll on Decision Support Loader system**
  - If not installed, the Loader program stops with run-time error 339
  - Install Microsoft Active X Control Pad to fix
- **Check for correct level of MDAC**
  - MDAC level has to be 3.510.3711.0
  - Third-party software may downgrade MDAC level

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

10.3.1 Important installation “gotchas”

This section covers problems that may be encountered when setting up Tivoli Decision Support for Storage Management Analysis and what can be done to solve them.

10.3.1.1 Tivoli Decision Support 2.1 patch

The Tivoli Decision Support version 2.1 patch will update the current version of Tivoli Decision Support Process Scheduler, the Path Manager Utility and the relevant readme.txt files, up to and including Admin-gd.pdf document files. Should the patch not have been installed, the Scheduler may have some discrepancies. There are two parts to applying this patch, the client as well as the administrator, and they may be located in the following paths:

```
C:\Program Files\TDS 2.1\ default
```

The Tivoli Decision Support patch fixes the following problems:

- Error 78886: The user can only apply changes once per session. If the shared source directory exists in two separate location, the user cannot change between the two locations using the same instance of the Path Manager Utility.
- Error 80911: The user cannot set the path to a mapped network drive.
- Error 81616: The Admin users must be able to specify the user name, password, and the network share name that will be used by the Tivoli Decision support Scheduler service to create a network connection.
- Error 81727: The option on the Tools menu used to re-register the MS DA0. Library file ‘dao350.dll’ will not work if the path to the file contains a space.

- Error 81071: Error message “Unable to connect to the data source. The database does not exist or is read only.” Occurs when trying to connect to a shared data drive.

The Tivoli Decision Support Patch fixes the above errors. The patch is a self-extracting file and is run from the “`Start`” menu and select `Run`.

10.3.1.2 Use correct ODBC driver version

The Tivoli Storage Management server is the source of data used for event and performance information. This data is transferred to the relational database using the Tivoli Storage manager ODBC driver. One reporting database can be used to contain the data from multiple Tivoli Storage Manager servers.

The new Tivoli Storage Management ODBC data source is set up in the `Control Panel -> ODBC`. The latest version to be installed is the Tivoli Storage Management ODBC Driver 3.07.01.00. When adding a new Tivoli Storage Management data source, refer to Figure 63.

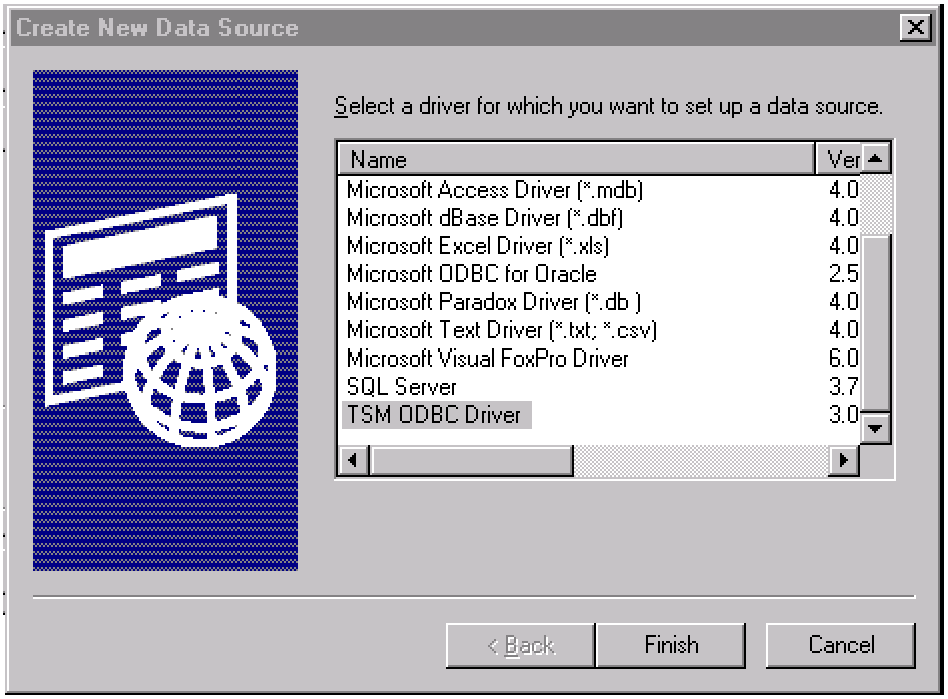

Select a driver for which you want to set up a data source.

Name	Ver
Microsoft Access Driver (*.mdb)	4.0
Microsoft dBase Driver (*.dbf)	4.0
Microsoft Excel Driver (*.xls)	4.0
Microsoft ODBC for Oracle	2.5
Microsoft Paradox Driver (*.db)	4.0
Microsoft Text Driver (*.txt; *.csv)	4.0
Microsoft Visual FoxPro Driver	6.0
SQL Server	3.7
TSM ODBC Driver	3.0

Figure 63. Setting up a new TSM data source

You may also find the older version of the ODBC driver that is listed as “32bit ADSM ODBC” driver version 3.01.0007. This ODBC driver should be upgraded to the version as mentioned above.

The ODBC drivers shipped on the Tivoli Decision Support CD-ROM will not work with Storage Management Analysis. Install the OEM ODBC drivers provided by the RDBMS vendors. To ensure that the latest drivers are installed, you may have to download these drivers from the vendor’s Web site.

10.3.1.3 fm20.dll module access

The run-time error 339 occurs (see Figure 64) when you install Tivoli Decision Support Loader and the fm20.dll module is not present on your NT system.

Some methods of obtaining the FM20.dll files are to install the suite of Microsoft Office or a portion thereof, like Microsoft Power Point. Another other way of fixing this is by going to the Microsoft Web site and downloading Microsoft Active X Control Pad from the following Web site `http://www.microsoft.com`.

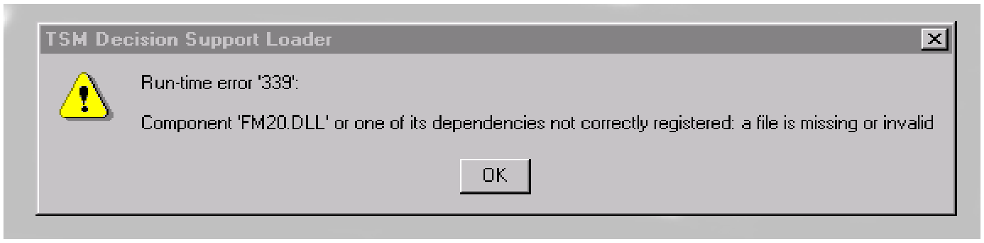

Run-time error '339':

Component 'FM20.DLL' or one of its dependencies not correctly registered: a file is missing or invalid

*Figure 64. The FM20.DLL error that appears if the required *.DLL files are missing*

10.3.1.4 MDAC levels

You should check the Microsoft Data Access Component (MDAC) levels regardless of what error appears on your monitor. Should you have errors and you suspect that it is a Windows NT error, check the operating system MDAC level. You can find this in `Control Panel ->ODBC` driver (as shown in Figure 65).

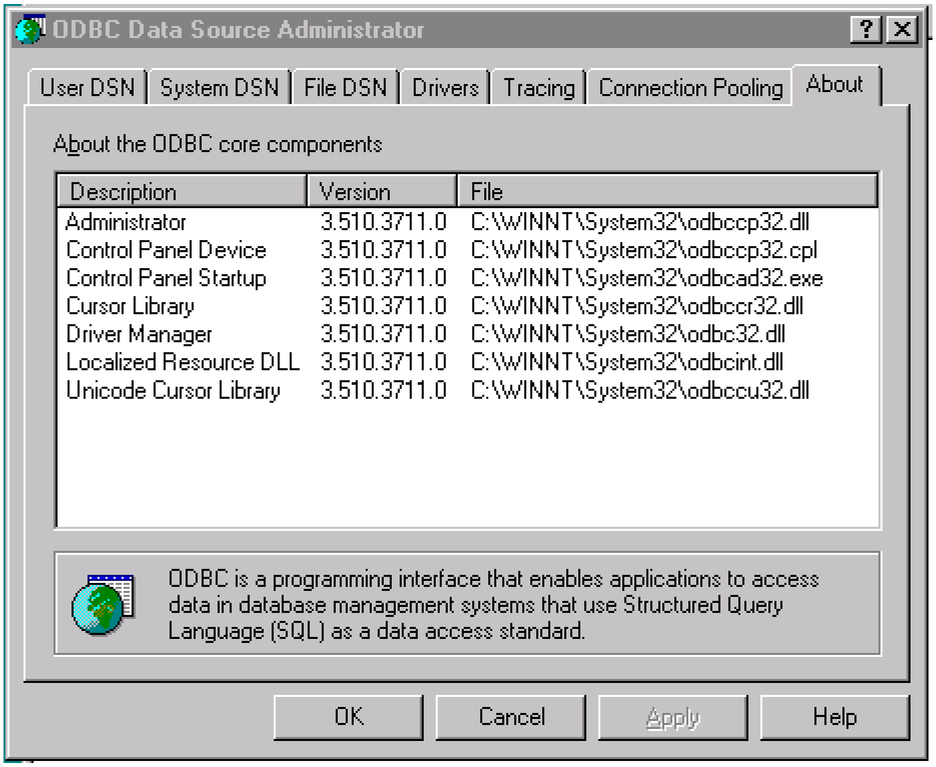

About the ODBC core components

Description	Version	File
Administrator	3.510.3711.0	C:\WINNT\System32\odbccp32.dll
Control Panel Device	3.510.3711.0	C:\WINNT\System32\odbccp32.cpl
Control Panel Startup	3.510.3711.0	C:\WINNT\System32\odbcad32.exe
Cursor Library	3.510.3711.0	C:\WINNT\System32\odbccr32.dll
Driver Manager	3.510.3711.0	C:\WINNT\System32\odbc32.dll
Localized Resource DLL	3.510.3711.0	C:\WINNT\System32\odbcint.dll
Unicode Cursor Library	3.510.3711.0	C:\WINNT\System32\odbccu32.dll

ODBC is a programming interface that enables applications to access data in database management systems that use Structured Query Language (SQL) as a data access standard.

Figure 65. Correct version of MDAC, with all items being at the same level

Verify the correct version of MDAC, which should be 3.510.3711.0. If you do not have the correct version, it is advisable to upgrade to the correct level to prevent operating system errors or other functional failures that could cause a fatal system error.

A more general point that should be noted is that any third-party software that is installed or downloaded from the Internet may downgrade your MDAC levels. For a further insight into MDAC, refer to the following Web site:
`http://www.microsoft.com.`

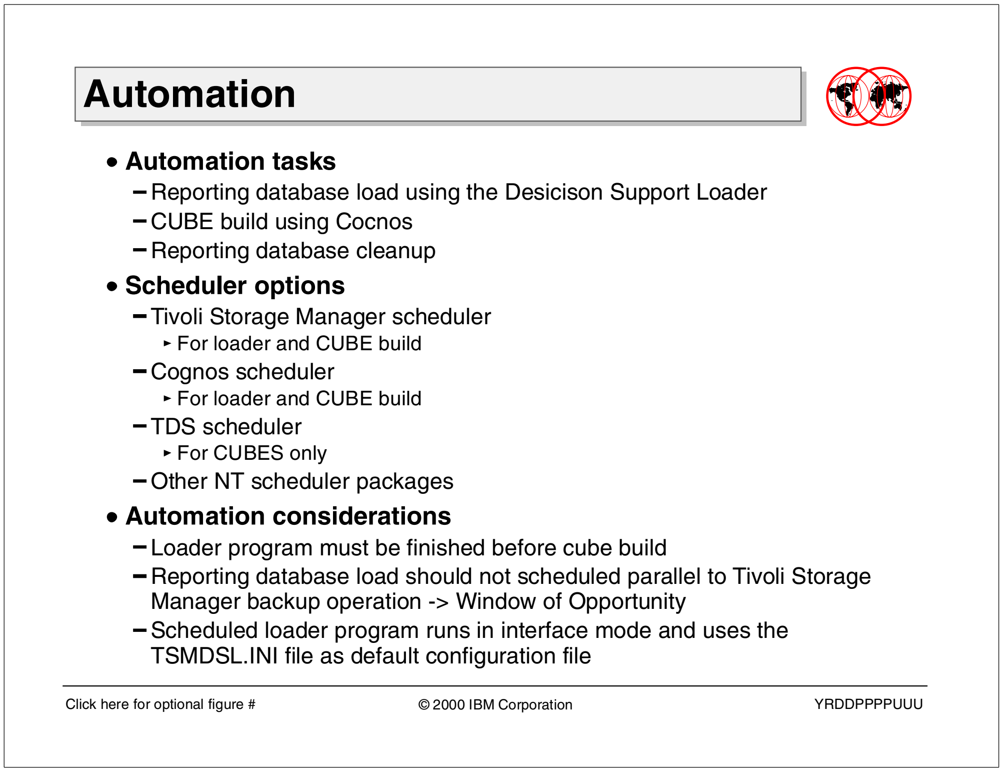

Automation

- **Automation tasks**
  - Reporting database load using the Desicison Support Loader
  - CUBE build using Cocnos
  - Reporting database cleanup
- **Scheduler options**
  - Tivoli Storage Manager scheduler
    - For loader and CUBE build
  - Cognos scheduler
    - For loader and CUBE build
  - TDS scheduler
    - For CUBES only
  - Other NT scheduler packages
- **Automation considerations**
  - Loader program must be finished before cube build
  - Reporting database load should not scheduled parallel to Tivoli Storage Manager backup operation -> Window of Opportunity
  - Scheduled loader program runs in interface mode and uses the TSMDSL.INI file as default configuration file

Click here for optional figure #

YRDDPPPPUUU

10.3.2 Automation

Using Tivoli Decision Support for Storage Management Analysis there are three periodical tasks you might want to automate:

- Load the reporting database with new event and performance data from the Tivoli Storage Manager databases
- Rebuild the Cognos CUBEs data base for the report creation task
- Clean up reporting database records using the Decision Support Loader program

The build process can be scheduled to execute automatically and at regular intervals. We recommend staggering the start times of the multiple cube builds so that they do not process simultaneously. This will improve performance for each build process.

Also, make sure that the cube builds are scheduled after the Tivoli Storage Management Decision Support Loader job runs because this job must finish in order to have new data available for the cubes. The loader job should also be executed via a scheduler. When you schedule the Decision Support Loader unattended, `TSMDSL.INI` is used as the default configuration file.

Note

Even when the Tivoli Storage Management Decision Support Loader is being executed unattended, the Tivoli Storage Manager Decision Support Loader graphical user interface appears on the screen. Do not use your mouse to click on the Decision Support Loader. This will stop the loader operation.

You have multiple scheduler options available to use in this component set.

10.3.2.1 Tivoli Storage Manager scheduler

The Tivoli Storage Manager scheduler is the recommended scheduler to choose for loader program execution.

When executing the scheduler from the install directory, you may have to add additional parameters to get the Tivoli Storage Management scheduler to function correctly. If you have the Tivoli Storage Management scheduler installed as a Windows NT service, make sure the ID has rights to the programs, including `tsmdsl`, and the data it must access. If you have installed the scheduler as a ‘System’ account, make sure you have checked the ‘allow desktop interface’ box.

To define a Tivoli Storage Manager schedule to automate the Decision Support Loader, you would issue the following command:

```
tsm: TSM010> define schedule SERVER TSMDSL type=client \
cont> descr=’execute loader program for TDS4SMA’ action=command \
cont> objects=’"c:\Program Files\Tivoli\tsm\decision\tsmdsl.exe"‘ \
cont> starttime=04:00 duration=10 durunits=Minutes period=1 \
cont> perunits=Days dayofweek=Weekday
```

The command will define a schedule which executes the loader program every weekday between 4:00 am and 4:10 am. The schedule needs to be connected with the client node name of the system, where the loader program should be executed.

To define a schedule for the CUBE build with the Tivoli Storage Manager, you would first need to create a CUBE build task using the Tivoli Decision Support Discovery Administrator. It is important that you record the task ID. For our example, we are assuming the task ID is 2.

To schedule the CUBE build using Tivoli Storage Manager, you would issue the following command:

```
tsm: TSM010> define schedule SERVER CUBE type=client \
cont> descr=’execute the CUBE build for TDS4SMA’ action=command \
cont> objects=’"c:\Program Files\TDS\EDAdmin.exe" /TaskID=2‘ \
cont> starttime=05:00 duration=10 durunits=Minutes period=1 \
cont> perunits=Days dayofweek=Weekday
```

The command will define a schedule which builds new Cognos CUBEs as described in task 2 every weekday between 5:00 am and 5:10 am. The schedule needs to be connected with the client node name of the system, where the TDS administrator program should be executed.

Using the Tivoli Storage Manager advanced scheduling functionality as described in Section 3.3, “Server scripting and scheduling enhancements” on page 102, it is possible to create a more complex schedule, which will first load the reporting database, and after successful completion, start the CUBE build.

10.3.2.2 Cognos scheduler

The Cognos scheduler can be set up to build the cubes. Be sure to schedule this at a time when the `tsmdsl` job has completed execution. Further information can be found in Chapter 4 of the IBM Redbook *Using Tivoli Decision Support Guides*, SG24-5506.

10.3.2.3 TDS scheduler

Tivoli Decision Support scheduler is configured for the Tivoli Discovery Administrator interface and can be used to add a schedule for cube building. Choose the name of the cube you want to build. Select daily, weekly, monthly, hourly or minutes, and/or how often in ‘days’ you want the execution to take place. Tivoli Decision Support will not execute until the service has started in `Control Panel -> Services -> TDS Scheduler`. Upon execution of the scheduler, the cubes will be re-built. Further information can be found in Chapter 4 of the IBM Redbook *Using Tivoli Decision Support Guides*, SG24-5506.

10.3.2.4 Cleanup of reporting database records

Cleanup of data records which you no longer want to report on in the RDBMS (Microsoft SQL, DB2, Oracle) is required in any database that is used. This could be performed by the Database Administrators.

Another method of doing data cleanup is called pruning. This is performed by the Tivoli Storage Management Administrators, by using a a Decision Support Loader parameter to specify the length of time data will be retained in the reporting database. When the Decision Support Loader is run after the number of days specified in this parameter, data exceeding the parameter is expired from the reporting database. The maximum length of time you can specify is 366 days.

The Decision Support Loader is set to prune approximately 1000 records at a time. You can change this default setting by modifying the number of records specified in the following line in the `TSMDSL.INI` configuration file:

```
ApproxRecordsToPrunePerBatch:=1000
```

Should you have a large number of records to prune, a small number of records will slow the Decision Support Loader processing time.

10.4 Reporting example using TDS for SMA

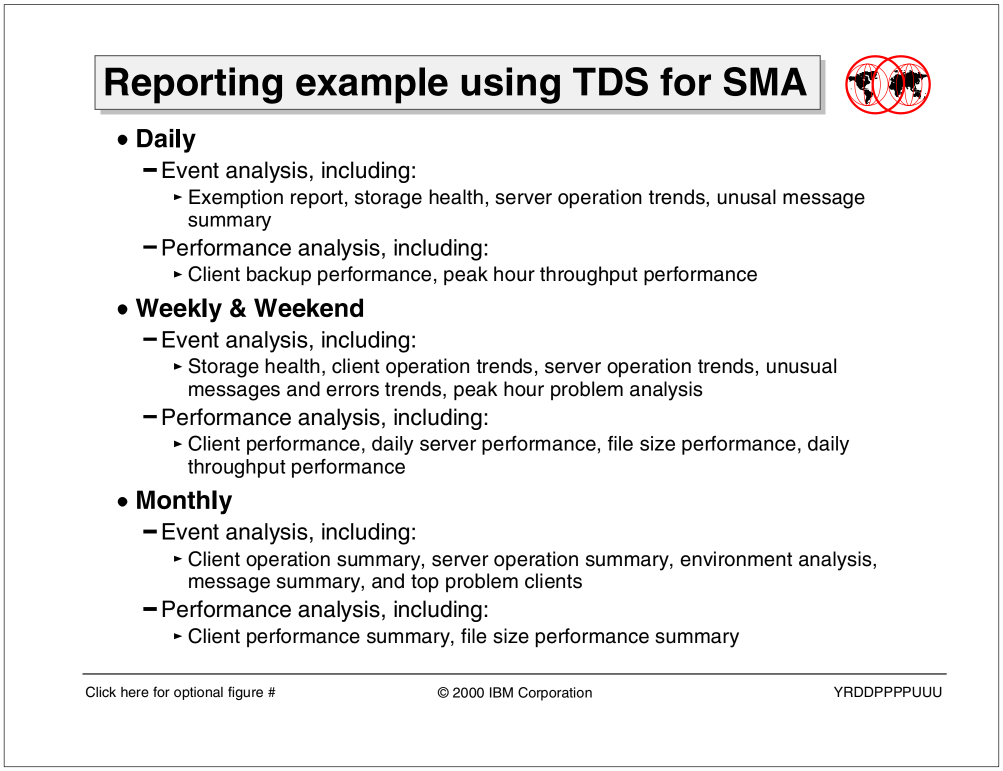

Reporting example using TDS for SMA

- **Daily**
  - Event analysis, including:
    - Exemption report, storage health, server operation trends, unusal message summary
  - Performance analysis, including:
    - Client backup performance, peak hour throughput performance
- **Weekly & Weekend**
  - Event analysis, including:
    - Storage health, client operation trends, server operation trends, unusual messages and errors trends, peak hour problem analysis
  - Performance analysis, including:
    - Client performance, daily server performance, file size performance, daily throughput performance
- **Monthly**
  - Event analysis, including:
    - Client operation summary, server operation summary, environment analysis, message summary, and top problem clients
  - Performance analysis, including:
    - Client performance summary, file size performance summary

Click here for optional figure # © 2000 IBM Corporation YRDDPPPPUUU

As Tivoli Storage Management Administrator, you should set up a regularly executed schedule for loading the reporting database and cube building. Keep in mind that you should try to avoid running the data loads at a time when the Tivoli Storage Manager database is being heavily used. Avoiding times when client backups and archives are running is preferred. Avoid running data loads during Tivoli Storage Manager administrative jobs. Expect degraded performance if this is not possible.

Reporting using System Management Analysis should be done on a regularly scheduled basis, as well. This will make trends and anomalies more apparent to the viewer and increase the value of the product to your organization. Specific reports to be drawn on a weekly, monthly, daily, or weekend basis are dependent on your preference and relevance to your operation.

The following is a recommended guide for a Storage Management Analysis report schedule:

10.4.1 Daily reporting

Having completed scheduling the Tivoli Storage Manager Decision Support Loader and the cube builds, the reports can be viewed in the Tivoli Discover Interface. These reports can show you possible server or client errors that should be checked on a daily basis. Use the following reports from the Tivoli Discovery Interface for daily reporting.

10.4.1.1 Tivoli Storage Management Event Analysis

1. Daily exception reports:
   - Classify outcome for client sessions.
   - Summarize the messages issued in my environment.
   - What are the results of all operations that occurred for a client?
   - What is the status of non-scheduled administrative operations?
   - What is the status of scheduled administrative operations?
   - What is the status of scheduled client sessions?
2. How might I improve my storage health?
   - How current are client backups?
   - Show me my server database backup record.
   - When did clients last contact the server?
3. Trends in server operations:
   - What is the average daily failed object count?
   - What is the average daily successful object count/
4. What messages should I be concerned with?
   - Summarize client messages by severity.
   - Summarize server error messages by severity.
   - What are the top 10 server messages?

10.4.1.2 Tivoli Storage Management Performance Analysis

1. How are my clients performing?
   - Which are my 10 fastest clients?
   - Which are my 10 slowest clients?
2. How does time of day and day of week affect my performance?
   - What is the average daily throughput for client operations?
   - What is the average hourly throughput for client operations?
   - What is the average hourly throughput take for server operations?

This set of reports should point to immediate problems which need attention. Look for unexplained ‘spikes’ in the performance and object failure.Also, error messages indicating severe and critical errors.

10.4.2 Weekly reporting

Use the following reports from the Tivoli Discovery Interface for weekly reporting:

10.4.2.1 Tivoli Storage Management Event Analysis

1. How might I improve my storage health?
   - How are clients schedules affected by the day of the week?
   - How current are client backups?
   - What administrators and clients have been locked out of the system?
2. Trends in client operations:
   - What is the average daily failed object count?
   - What time of day do peak client data transfers occur?
3. Trends in server operations:
   - How does time affect server session activity?
   - What is the daily failed object count for server operations?

4. What messages should I be concerned with?
   - Summarize client messages by severity.
   - What are the top 10 client messages with the most problems?
5. When do I find the most problems?
   - What day of week do peak client error rates occur?
   - What day of week do peak server error rates occur?
   - What time of day do peak clients error rates occur?
   - What time of day do peak server error rates occur?

10.4.2.2 Tivoli Storage Management Performance Analysis

1. How are my clients performing?
   - How does TCP/IP address affect throughput?

These reports should alert you to potential problems, unusual growth which could lead to problems, or beginning trends.

10.4.3 Weekend reporting

The following reports from the Tivoli Discovery Interface are recommended for weekend reporting:

10.4.3.1 Tivoli Storage Management Event Analysis

1. When do I find the most problems?
   - What day of week do peak client error rates occur
   - What day of week do peak server error rates occur

10.4.3.2 Tivoli Storage Management Performance Analysis

1. How are my servers performing?
   - How much time is spent daily waiting for media access
2. How does file size affect my performance?
   - How does file size affect client operation throughput
3. How does time of day and day of week affect my performance?
   - What is the average daily throughput for client operations?
   - What is the average hourly throughput for client operations?
   - What is the average hourly throughput take for server operations?

Weekend reporting is much like weekly in purpose. Some organizations prefer to look at data as it corresponds to the calendar week. You may find that this is appropriate. If not, run these along with the weekly reports.

10.4.4 Monthly reporting

The following from the Tivoli Discovery Interface are recommended for monthly reporting:

10.4.4.1 Tivoli Storage Management Event Analysis

1. Trends in client operations:
   - Summarize client messages issued by platform.
   - Summarize client messages issued by software level.
   - What time of day do peak client data transfers occur?
   - Which clients have the highest average failed objects count?
2. Trends in server operations:
   - How does time affect server session activity?
   - How much data is processed by daily server operations?
   - What is the daily failed object count for server operations?
3. What does my environment look like?
   - How many clients are registered with each server?
4. What messages should I be concerned with?
   - Summarize client messages by severity?
   - What are the top 10 clients with the most problems?

10.4.4.2 Tivoli Storage Management Performance Analysis

1. How are my clients performing?
   - How does platform affect client throughput rates?
   - How does TCP/IP address affect throughput?
2. How does file size affect my performance?
   - Classify sessions by file size and performance thresholds?
   - How does file size affect client operation throughput?

These reports are geared, with the proper date filters, to yield an analysis of performance and events within the month. They will give a picture of what has transpired during the month, as well as showing storage trends in the enterprise. These reports may be good candidates to be stored for further reference or year-end summary reporting.

Appendix A. Available TSM and TDP versions and platforms

This appendix provides an overview of currently available versions of Tivoli Storage Manager server and client, Tivoli Data Protection for Applications, and supported operating system levels. Also, it includes a table showing the compatibility for the new version 3.7.2 and 4.1 features when using different Tivoli Storage Manager client and server levels together.

A.1 Tivoli Storage Manager Server platforms

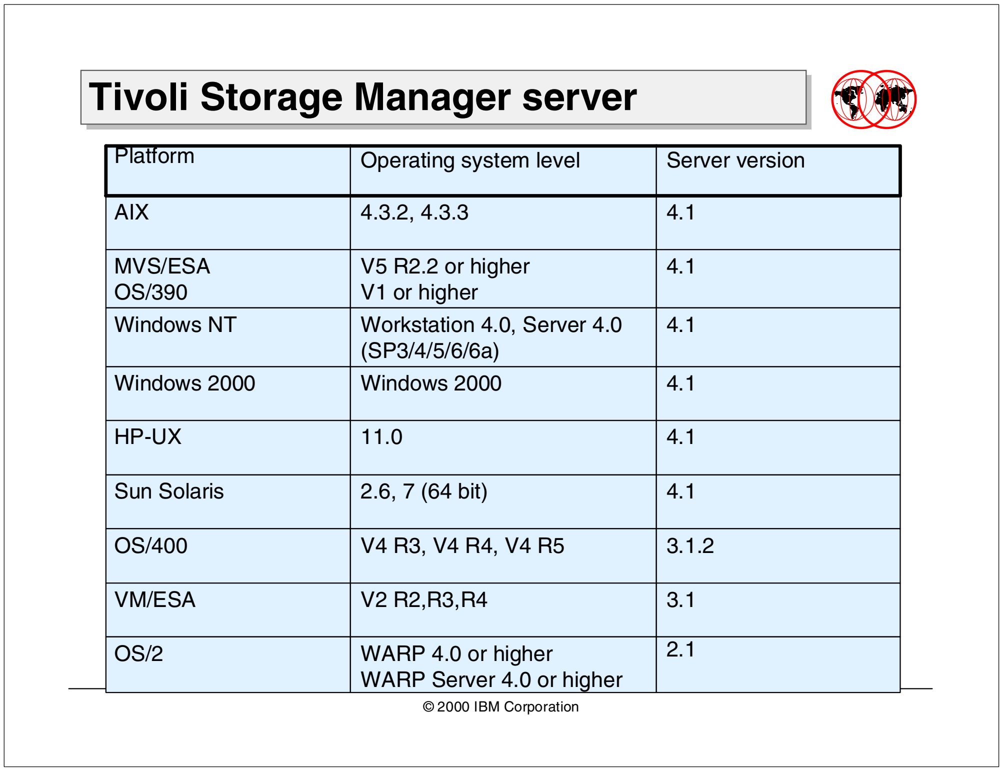

Tivoli Storage Manager server

Platform	Operating system level	Server version
AIX	4.3.2, 4.3.3	4.1
MVS/ESA OS/390	V5 R2.2 or higher V1 or higher	4.1
Windows NT	Workstation 4.0, Server 4.0 (SP3/4/5/6/6a)	4.1
Windows 2000	Windows 2000	4.1
HP-UX	11.0	4.1
Sun Solaris	2.6, 7 (64 bit)	4.1
OS/400	V4 R3, V4 R4, V4 R5	3.1.2
VM/ESA	V2 R2,R3,R4	3.1
OS/2	WARP 4.0 or higher WARP Server 4.0 or higher	2.1

This graphic shows all Tivoli Storage Manager supported operating system platforms and the latest available server versions.

A.2 Tivoli Storage Manager client platforms

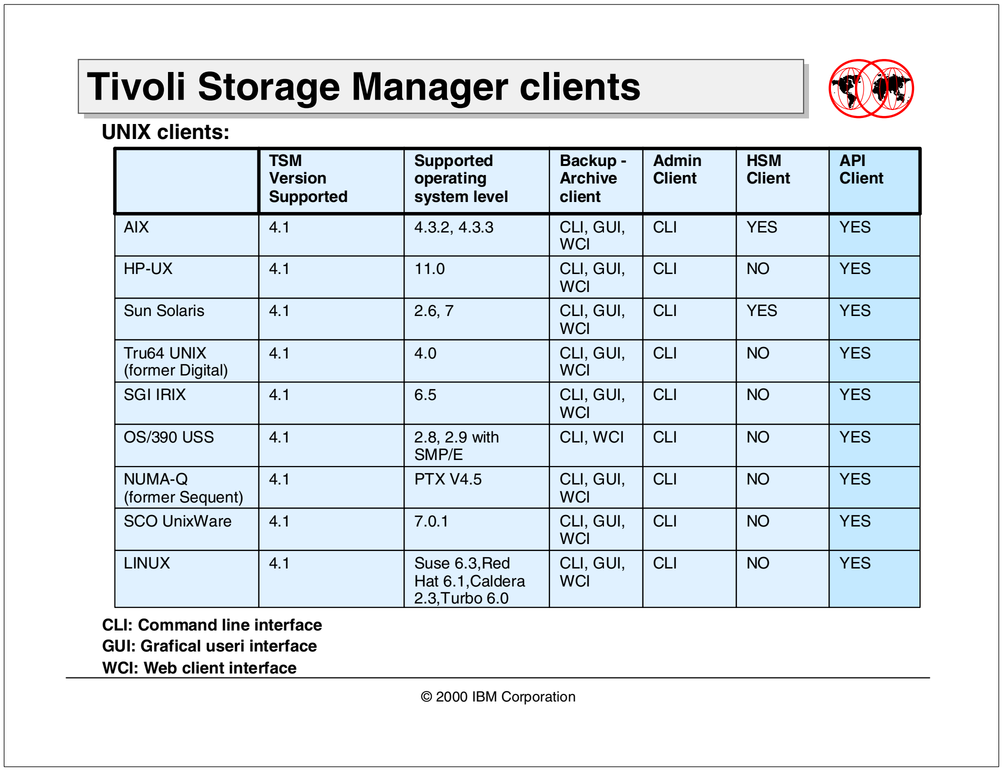

Tivoli Storage Manager clients

UNIX clients:

	TSM Version Supported	Supported operating system level	Backup - Archive client	Admin Client	HSM Client	API Client
AIX	4.1	4.3.2, 4.3.3	CLI, GUI, WCI	CLI	YES	YES
HP-UX	4.1	11.0	CLI, GUI, WCI	CLI	NO	YES
Sun Solaris	4.1	2.6, 7	CLI, GUI, WCI	CLI	YES	YES
Tru64 UNIX (former Digital)	4.1	4.0	CLI, GUI, WCI	CLI	NO	YES
SGI IRIX	4.1	6.5	CLI, GUI, WCI	CLI	NO	YES
OS/390 USS	4.1	2.8, 2.9 with SMP/E	CLI, WCI	CLI	NO	YES
NUMA-Q (former Sequent)	4.1	PTX V4.5	CLI, GUI, WCI	CLI	NO	YES
SCO UnixWare	4.1	7.0.1	CLI, GUI, WCI	CLI	NO	YES
LINUX	4.1	Suse 6.3,Red Hat 6.1,Caldera 2.3,Turbo 6.0	CLI, GUI, WCI	CLI	NO	YES

CLI: Command line interface
GUI: Grafical useri interface
WCI: Web client interface

A.2.1 UNIX client platforms

This graphic shows the highest currently available Tivoli Storage Manager version of UNIX clients. All version 3.1 clients can be used together with Tivoli Storage Manager V4.1 servers. In this case, Version 3.7 and 4.1 client function is not available.

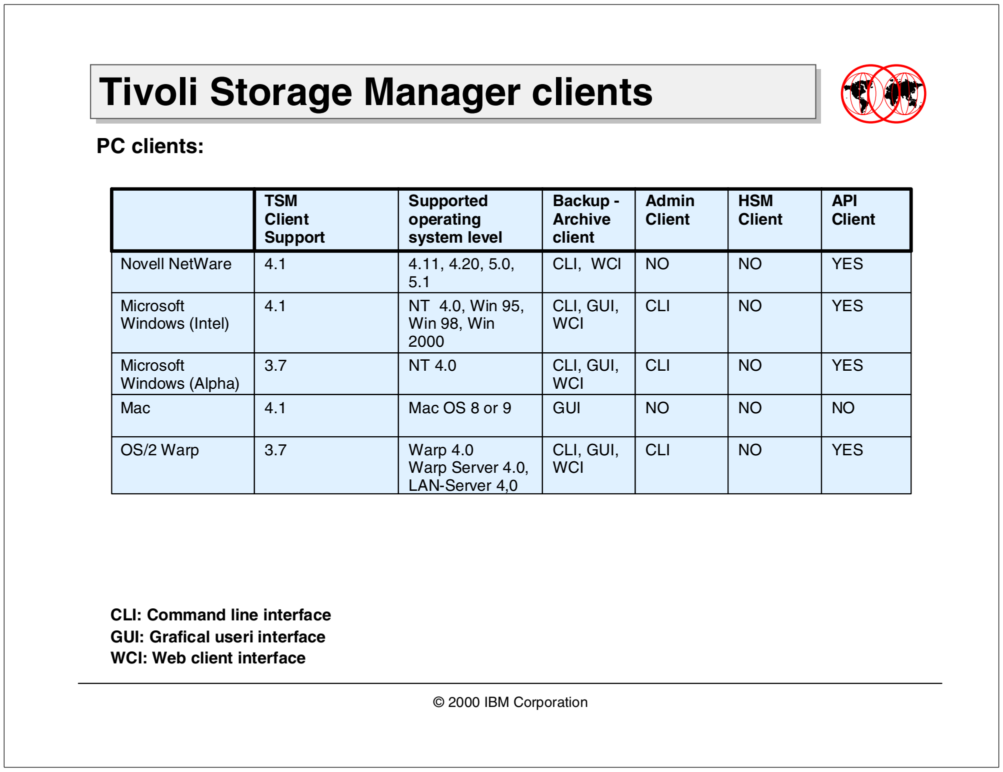

Tivoli Storage Manager clients

PC clients:

	TSM Client Support	Supported operating system level	Backup - Archive client	Admin Client	HSM Client	API Client
Novell NetWare	4.1	4.11, 4.20, 5.0, 5.1	CLI, WCI	NO	NO	YES
Microsoft Windows (Intel)	4.1	NT 4.0, Win 95, Win 98, Win 2000	CLI, GUI, WCI	CLI	NO	YES
Microsoft Windows (Alpha)	3.7	NT 4.0	CLI, GUI, WCI	CLI	NO	YES
Mac	4.1	Mac OS 8 or 9	GUI	NO	NO	NO
OS/2 Warp	3.7	Warp 4.0 Warp Server 4.0, LAN-Server 4,0	CLI, GUI, WCI	CLI	NO	YES

CLI: Command line interface
GUI: Grafical useri interface
WCI: Web client interface

A.2.2 PC client platforms

This graphic shows the highest currently available Tivoli Storage Manager version of PC clients. All version 3.1 clients can be used together with Tivoli Storage Manager V4.1 servers. In this case, Version 3.7 and 4.1 client function is not available.

A.3 Tivoli Data Protection products

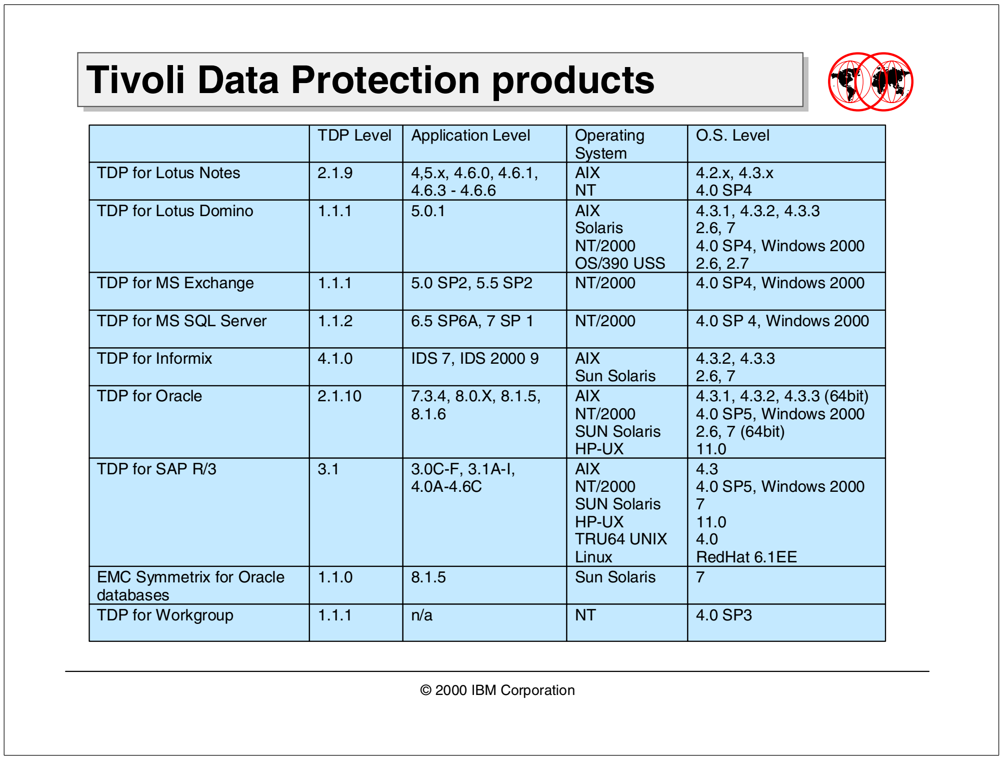

Tivoli Data Protection products

	TDP Level	Application Level	Operating System	O.S. Level
TDP for Lotus Notes	2.1.9	4,5.x, 4.6.0, 4.6.1, 4.6.3 - 4.6.6	AIX NT	4.2.x, 4.3.x 4.0 SP4
TDP for Lotus Domino	1.1.1	5.0.1	AIX Solaris NT/2000 OS/390 USS	4.3.1, 4.3.2, 4.3.3 2.6, 7 4.0 SP4, Windows 2000 2.6, 2.7
TDP for MS Exchange	1.1.1	5.0 SP2, 5.5 SP2	NT/2000	4.0 SP4, Windows 2000
TDP for MS SQL Server	1.1.2	6.5 SP6A, 7 SP 1	NT/2000	4.0 SP 4, Windows 2000
TDP for Informix	4.1.0	IDS 7, IDS 2000 9	AIX Sun Solaris	4.3.2, 4.3.3 2.6, 7
TDP for Oracle	2.1.10	7.3.4, 8.0.X, 8.1.5, 8.1.6	AIX NT/2000 SUN Solaris HP-UX	4.3.1, 4.3.2, 4.3.3 (64bit) 4.0 SP5, Windows 2000 2.6, 7 (64bit) 11.0
TDP for SAP R/3	3.1	3.0C-F, 3.1A-I, 4.0A-4.6C	AIX NT/2000 SUN Solaris HP-UX TRU64 UNIX Linux	4.3 4.0 SP5, Windows 2000 7 11.0 4.0 RedHat 6.1EE
EMC Symmetrix for Oracle databases	1.1.0	8.1.5	Sun Solaris	7
TDP for Workgroup	1.1.1	n/a	NT	4.0 SP3

This graphic shows currently available Tivoli Data Protection for Applications products their latest version, and supported platforms and application levels, if applicable.

Appendix B. Special notices

This publication is intended to help customers, consultants, IBM and Tivoli Business Partners and staff understand Tivoli Storage Manager Version 3.7.3 and 4.1. The information in this publication is not intended as the specification of any programming interfaces that are provided by Tivoli Storage Manager Version 3.7.3 or 4.1. See the PUBLICATIONS section of the IBM Programming Announcement for Tivoli Storage Manager Version 3.7 or 4.1. for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

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

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

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

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

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

IBM ®
OS/2
RS/6000
SP
XT
Lotus
Redbooks
Notes
OS/390
S/390
Wizard
400
Domino
Redbooks Logo

The following terms are trademarks of other companies:

Tivoli, Manage. Anything. Anywhere.,The Power To Manage., Anything. Anywhere.,TME, NetView, Cross-Site, Tivoli Ready, Tivoli Certified, Planet Tivoli, and Tivoli Enterprise are trademarks or registered trademarks of Tivoli Systems Inc., an IBM company, in the United States, other countries, or both. In Denmark, Tivoli is a trademark licensed from Kjøbenhavns Sommer - Tivoli A/S.

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

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

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

PC Direct is a trademark of Ziff Communications Company in the United States and/or other countries and is used by IBM Corporation under license.

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

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

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

Other company, product, and service names may be trademarks or service marks of others.

Appendix C. Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

C.1 IBM Redbooks

For information on ordering these publications see “How to get IBM Redbooks” on page 335.

- *Tivoli Storage Management Concepts,* SG24-4877
- *Getting Started with Tivoli Storage Manager: Implementation Guide* SG24-5416
- *Tivoli Storage Manager V3.7: Technical Guide,* SG24-5477
- *Tivoli Storage Management Reporting*, SG24-6109
- *Using Tivoli Decision Support Guides*, SG24-5506
- *Guide to Sharing and Partitioning IBM Tape Library Dataservers*, SG24-4409

C.2 IBM Redbooks collections

Redbooks are also available on the following CD-ROMs. Click the CD-ROMs button at ibm.com/redbooks for information about all the CD-ROMs offered, updates and formats.

CD-ROM Title	**Collection Kit Number**
IBM System/390 Redbooks Collection	SK2T-2177
IBM Networking Redbooks Collection	SK2T-6022
IBM Transaction Processing and Data Management Redbooks Collection	SK2T-8038
IBM Lotus Redbooks Collection	SK2T-8039
Tivoli Redbooks Collection	SK2T-8044
IBM AS/400 Redbooks Collection	SK2T-2849
IBM Netfinity Hardware and Software Redbooks Collection	SK2T-8046
IBM RS/6000 Redbooks Collection (PDF Format)	SK2T-8043
IBM Application Development Redbooks Collection	SK2T-8037
IBM Enterprise Storage and Systems Management Solutions	SK3T-3694

C.3 Other resources

These publications are also relevant as further information sources:

- *Tivoli Storage Manager for AIX, Quick Start,* GC35-0402
- *Tivoli Storage Manager for Sun Solaris, Quick Start,* GC35-0406
- *Tivoli Storage Manager for Windows, Quick Start*, GC35-0409
- *Tivoli Storage Manager for Windows Administrator’s Reference*, GC35-0381
- *Tivoli Storage Manager, Installing the clients*, SH26-4119
- *IBM Storage Area Network Data Gateway Installation and User’s Guide,* SC26-7304
- *IBM Storage Area Network Data Gateway Router Service Guide*, SY27-7614

C.4 Referenced Web sites

These Web sites are also relevant as further information sources:

- http://www.ibm.com/redbooks
  IBM ITSO home page
- http://www.tivoli.com/support/storage_mgr/clients.html#linux
  Tivoli Storage Manager Linux client downloads and latest maintenance releases (fixes)
- http://www.tivoli.com/support/storage_mgr/san/overview.html
  Tivoli Storage Manager fibre channel interoperability
- http://www.storage.ibm.com/hardsoft/products/sangateway/support/cdr/sdgcdr.htm
  IBM Storage Area Networks, SAN Data Gateway downloads
- http://www.qlogic.com/
  Qlogic Inc. home page
- http://www.rs6000.ibm.com/support/micro/df1000f7.html
  Microcode information for the IBM RISC System/6000 Gigabit FibreChannel PCI adapter
- http://www.storage.ibm.com/hardsoft/products/sangateway/support/cdr/AIX/readme
  IBM SAN Data Gateway on RS-6000 Version 1.07 release notes
- http://www.microsoft.com
  Microsoft.com home page

How to get IBM Redbooks

This section explains how both customers and IBM employees can find out about IBM Redbooks, redpieces, and CD-ROMs. A form for ordering books and CD-ROMs by fax or e-mail is also provided.

- **Redbooks Web Site** ibm.com/redbooks

  Search for, view, download, or order hardcopy/CD-ROM Redbooks from the Redbooks Web site. Also read redpieces and download additional materials (code samples or diskette/CD-ROM images) from this Redbooks site.

  Redpieces are Redbooks in progress; not all Redbooks become redpieces and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

- **E-mail Orders**

  Send orders by e-mail including information from the IBM Redbooks fax order form to:

	e-mail address
In United States or Canada	pubscan@us.ibm.com
Outside North America	Contact information is in the “How to Order” section at this site: http://www.elink.ibmlink.ibm.com/pbl/pbl

- **Telephone Orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	Country coordinator phone number is in the “How to Order” section at this site: http://www.elink.ibmlink.ibm.com/pbl/pbl

- **Fax Orders**

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	Fax phone number is in the “How to Order” section at this site: http://www.elink.ibmlink.ibm.com/pbl/pbl

This information was current at the time of publication, but is continually subject to change. The latest information may be found at the Redbooks Web site.

IBM Intranet for Employees

IBM employees may register for information on workshops, residencies, and Redbooks by accessing the IBM Intranet Web site at http://w3.itso.ibm.com/ and clicking the ITSO Mailing List button. Look in the Materials repository for workshops, presentations, papers, and Web pages developed and written by the ITSO technical professionals; click the Additional Materials button. Employees may access MyNews at http://w3.ibm.com/ for redbook, residency, and workshop announcements.

IBM Redbooks fax order form

Please send me the following:

Title	**Order Number**	**Quantity**

First name Last name

Company

Address

City Postal code Country

Telephone number Telefax number VAT number

□ Invoice to customer number

□ Credit card number

Credit card expiration date Card issued to Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

Index

IBM Redbooks review

Your feedback is valued by the Redbook authors. In particular we are interested in situations where a Redbook "made the difference" in a task or problem you encountered. Using one of the following methods, **please review the Redbook, addressing value, subject matter, structure, depth and quality as appropriate.**

- Use the online **Contact us** review redbook form found at **ibm.com**/redbooks
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to redbook@us.ibm.com

Document Number **Redbook Title**	SG24-6110-00 Tivoli Storage Manager Version 3.7.3 & 4.1: Technical Guide
Review	
What other subjects would you like to see IBM Redbooks address?	
Please rate your overall satisfaction:	O Very Good O Good O Average O Poor
Please identify yourself as belonging to one of the following groups:	O Customer O Business Partner O Solution Developer O IBM, Lotus or Tivoli Employee O None of the above
Your email address: The data you provide here may be used to provide you with information from IBM or our business partners about our products, services or activities.	 O Please do not use the information collected here for future marketing or promotional contacts or other communications beyond the scope of this transaction.
Questions about IBM's privacy policy?	The following link explains how we protect your personal information. **ibm.com**/privacy/yourprivacy/

IBM

Redbooks

Tivoli Storage Manager Version 3.7.3 & 4.1: Technical Guide

(0.5" spine)
0.475"<->0.873"
250 <-> 459 pages

IBM®

Tivoli Storage Manager Version 3.7.3 & 4.1: Technical Guide

Redbooks

Exploiting SAN environments for LAN-free backup of application data

Complete data protection solution for mobile systems

Windows 2000 exploitation and support

Tivoli Storage Manager, the backbone product of the Tivoli Storage Management product set, is an enterprise-wide network storage management solution. Tivoli Storage Manager Version 4.1 was introduced in July 2000 and is the successor of Tivoli ADSM Version 3.1 and Tivoli Storage Manager Version 3.7. It provides many new functions such as support for mobile systems backup, LAN-free backup and restore of data using a Storage Area Network (SAN) environment, and Windows 2000 exploitation and support.

This IBM Redbook presents an overview of Tivoli Storage Manager Version 3.7.3 and Version 4.1. The book provides updates on the Tivoli Storage Management product set, and it gives a detailed description of each of the new functions of Tivoli Storage Manager. The book also discusses in detail Tivoli Storage Manager Windows 2000 exploitation and support, and introduces two other members of the Tivoli Storage Management product set: Tivoli SANergy File Sharing, and Tivoli Decision Support for Storage Management Analysis.

This book is intended for customers, consultants, IBM Business Partners, IBM employees, and Tivoli staff who are familiar with ADSM Version 3.1 and Tivoli Storage Manager Version 3.7, and who need to understand what is new in Tivoli Storage Manager Version 3.7.3 and 4.1.

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

SG24-6110-00 ISBN 0738418153