IBM

IBM ProtecTIER Implementation and Best Practices Guide

- Implement scenarios for backup recovery applications with ProtecTIER
- Optimize data deduplication and storage efficiency
- Exploit the File System Interface using CIFS and NFS

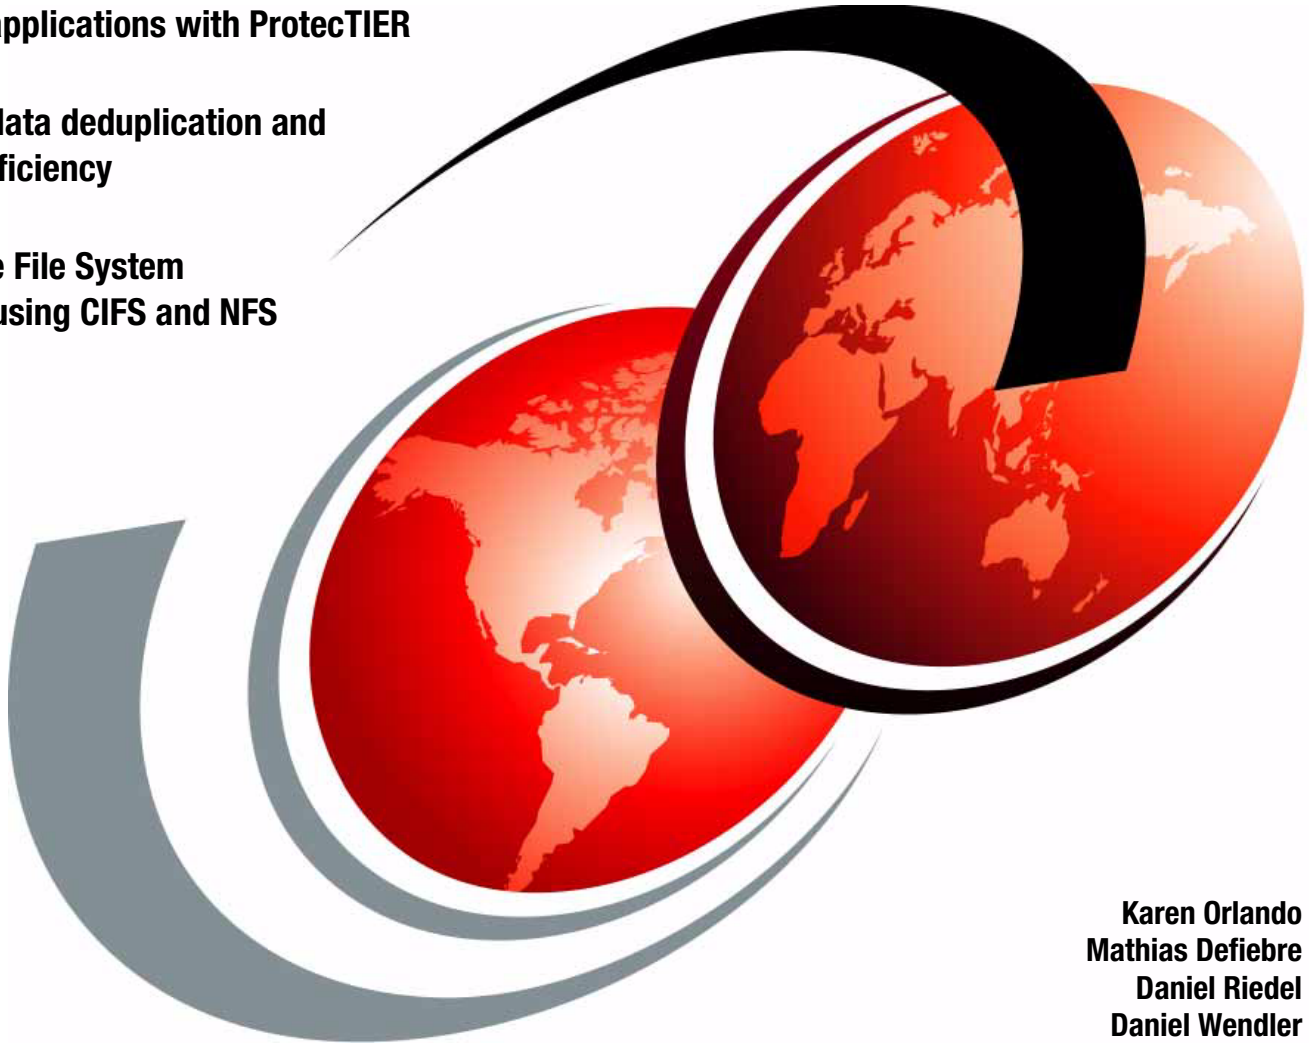

Karen Orlando
Mathias Defiebre
Daniel Riedel
Daniel Wendler

Redbooks

ibm.com/redbooks

International Technical Support Organization

IBM ProtecTIER Implementation and Best Practices Guide

May 2013

SG24-8025-01

Note: Before using this information and the product it supports, read the information in “Notices” on page xiii.

Second Edition (May 2013)

This edition applies to ProtecTIER Version 3.3

Contents

Notices

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

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

AIX®	HyperFactor®	Redbooks®
AS/400®	IBM®	Redbooks (logo) ®
DB2®	Informix®	RETAIN®
DB2 Connect™	iNotes®	Sametime®
DB2 Universal Database™	iSeries®	Storwize®
Domino®	Lotus®	System i®
DS4000®	Lotus Notes®	System p®
DS8000®	Notes®	System Storage®
Easy Tier®	POWER®	System x®
eServer™	Power Systems™	Tivoli®
FICON®	ProtecTIER®	XIV®
Global Technology Services®	Real-time Compression™	z/OS®

The following terms are trademarks of other companies:

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Ultrium, the LTO Logo and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

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

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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

Other company, product, or service names may be trademarks or service marks of others.

Preface

This IBM® Redbooks® publication provides best practice guidance for planning, installing, and configuring the IBM TS7600 ProtecTIER® family of products. This guide provides all the latest best practices for using ProtecTIER Software Version 3.3 and the revolutionary and patented IBM HyperFactor® deduplication engine, along with other data storage efficiency techniques, such as compression and defragmentation.

The IBM System Storage® TS7650G ProtecTIER Deduplication Gateway and the IBM System Storage TS7620 ProtecTIER Deduplication Appliance Express are disk-based data storage systems that are configured for three available interfaces:

- The Virtual Tape Library (VTL) interface is the foundation of ProtecTIER and emulates traditional automated tape libraries.
- The Symantec NetBackup OpenStorage (OST) API can be integrated with Symantec NetBackup to provide backup-to-disk without having to emulate traditional tape libraries.
- The newly available File System Interface (FSI) supports Common Internet File System (CIFS) and Network File System (NFS) as backup targets.

For your existing ProtecTIER solution, this guide provides best practices and suggestions to boost the performance and the effectiveness of the data deduplication with regards to your application platforms for your VTL, OST, and FSI systems.

When you build a ProtecTIER data deduplication environment, this guide helps your IT architects and solution designers plan for the best option and scenario for data deduplication for their environments. This guide helps you optimize your deduplication ratio, while reducing the hardware, power and cooling, and management costs.

This guide provides expertise that was gained from the IBM ProtecTIER Field Technical Sales Support (FTSS/CSS) Group, development, and Quality Assurance teams.

Note: The following ProtecTIER products are withdrawn and can no longer be directly ordered:

- TS7650 ProtecTIER Deduplication Appliance, 3958-AP1.
- TS7610 ProtecTIER Deduplication Appliance Express, 3959-SM1. This product was replaced by the TS7620 ProtecTIER Deduplication Appliance Express, 3959-SM2.

For more information, see http://www.ibm.com/common/ssi/index.wss.

The team who wrote this book

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, Poughkeepsie Center.

Karen Orlando is a Project Leader at the International Technical Support Organization, Tucson, Arizona Center. Karen has over 25 years in the I/T industry with extensive experience in open systems, and Information and Software development of IBM hardware and software storage. She holds a degree in Business Information Systems from the University of Phoenix and is Project Management Professional (PMP) certified since 2005.

Mathias Defiebre is a Lab Services IT Specialist from Germany. He joined IBM in 1998 and has worked extensively with storage products. He graduated from the University of Cooperative Education Mannheim with a German Diploma in Information Technology Management and a Bachelor of Science degree. His areas of expertise include developing and delivering services for IBM TS76xx ProtecTIER deduplication products, IBM Tivoli® Storage Manager integration, and healthcheck solutions. Mathias also conducts proof of concepts, workshops, and benchmarks with focus on data protection. He is an IBM Certified Specialist for TotalStorage Networking and Virtualization Architectures. Mathias previously co-authored multiple versions of the Tivoli Storage Productivity Center and ProtecTIER Implementation and Best Practices IBM Redbooks publication.

Daniel Riedel is an IBM Certified I/T Professional in the United States. He has 25 years of experience in the field and has worked at IBM for 20 years. His area of expertise includes ProtecTIER data deduplication and IBM Power Systems™. He has written extensively about the ProtecTIER product line and has published more than twenty technical articles on ProtecTIER implementation and replication.

Daniel Wendler is an IT Specialist within the IBM MTS Group in Germany. Daniel studied computer science and graduated at the University of Applied Science Wiesbaden in 2005. He wrote his diploma thesis about automated policy-based management of removable storage media with a focus on IBM storage. Daniel worked for IBM storage software development as a student and joined the IBM European Storage Competence Center as a permanent employee in 2005 after his graduation. He now works as a Product Field Engineer (PFE) for IBM removable media storage system products with a focus on data protection and retention. He provides post-sales support, implementation and education services for enterprise tape libraries, Open System virtualization engines, tape backup, archive, and encryption solutions.

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

Ahmed Almoustafas, Shira Ben-Dor, Aviv Caro, Joseph Dain, Doron Tal, James Thompson, Shmuel Vashdi
IBM Systems and Technology Group, Systems Development

Tom Chandler and Alexander Jung
IBM Sales & Distribution, STG

Jana Jamsek and Nancy Roper
IBM Sales & Distribution, TSS

Erick Kissel
IBM Software Group, Tivoli, Tivoli Storage Manager Server Development

Larry Fuss and Dietmar Schniering
IBM Systems and Technology Group, Storage Platform

Mervyn Venter
IBM Systems & Technology Group, Client Enablement & Systems Assurance

Bjoern Wesselbaum
IBM Global Technology Services®, High End Disk Solutions EMEA

Leigh Woods
IBM Global Technology Services, Tivoli Storage Manager

Authors of the first edition, *IBM ProtecTIER Implementation and Best Practices Guide* published on August 31, 2012 were:

Karen Orlando, Adriana Pellegrini Furnielis, Mathias Defiebre, Jane Lau, Libor Miklas, Angela Pholphiboun

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

  ibm.com/redbooks
- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

Part 1

General best practices

This part describes the general best practices that users of ProtecTIER can employ, such as recovery management and back-end storage concepts. This part also describes guidelines for Virtual Tape Library (VTL), File System Interface (FSI) for the Common Internet System (FSI-CIFS) and for the Network File System (FSI-NFS), and for OpenStorage (OST). This part describes configuration options with regards to host attachment.

This part describes the following topics:

- ProtecTIER basics
- Deduplication considerations
- Virtual Tape Library guidelines
- ProtecTIER File System Interface: General introduction and NFS
- OpenStorage guidelines
- Host attachment considerations for VTL

1

ProtecTIER basics

This chapter describes the general concepts of ProtecTIER as related to its basic operation, including a list of terms and definitions that are used in this book and in the ProtecTIER environment. A brief overview is also provided about the existing models, and the benefits of the ProtecTIER command-line interface (ptcli).

In addition, this chapter provides readers with a basic introduction to the ProtecTIER engine for deduplication, which is HyperFactor. This chapter also describes the Remote Support and Call Home capabilities.

This chapter describes the following topics:

- ProtecTIER basic concepts
- ProtecTIER models for open systems
- Terminology
- Remote support and Call Home
- ProtecTIER command-line interface overview

1.1 Terminology

The following terms and definitions are used in this publication and are common for all products from the ProtecTIER family:

ProtecTIER When used by itself, this expression points to the IBM patented deduplication solution based on HyperFactor. Depending on the context, it can mean the family of products, a specific device, or just the deduplication engine.

TS7600 When used alone, this term signifies the IBM family of virtualization solutions that operate on the ProtecTIER platform, including the ProtectTIER appliances and gateways.

Factoring ratio The factoring ratio is the ratio of nominal capacity to physical capacity in the ProtecTIER repository. For example, if you have 100 TB of user data (nominal capacity) and it is stored on 10 TB of physical capacity, your factoring ratio is 10:1.

HyperFactor The patented IBM algorithm that eliminates data duplicates from the incoming backup data streams. The factoring ratios are the result of HyperFactor processing.

VTL Virtual Tape LIbrary. The ProtecTIER VTL service emulates traditional tape libraries. By emulating tape libraries, you can use ProtecTIER VTL to migrate to disk backup without having to replace your entire backup environment. Your existing backup application can access virtual robots to move virtual cartridges between virtual slots and drives. The backup application perceives that the data is being stored on cartridges, but ProtecTIER actually stores data on a deduplicated disk repository.

Shelf A container of VTL cartridges within a ProtecTIER repository. This container is analogous to a shelf or a rack where physical tapes are kept outside of the automated tape library in cartridge slots. This container is only applicable with the VTL mode and is not available in OST or FSI modes.

OST Open Storage Technology. This configuration option allows ProtecTIER to be integrated with Symantec NetBackup to provide a backup-to-disk capability without using a VTL emulation. Using a plug-in that is installed on an OST enabled media server, you can use ProtecTIER to implement a communication protocol that supports data transfer and control between the backup server and the ProtecTIER server. Thus, you use ProtecTIER to implement a storage server emulation.

FSI File System Interface. This configuration option allows ProtecTIER to present disk repository storage as a virtualized hierarchy of file systems.

FSI-CIFS	FSI Common Internet File System. ProtecTIER emulates Windows file system behavior and presents a virtualized hierarchy of file systems, directories, and files to Windows CIFS clients. When configured for FSI-CIFS, ProtecTIER emulates a network-attached storage (NAS) backup target capable of using both HyperFactor and ProtecTIER Native Replication bandwidth reduction techniques for storing and replicating deduplicated data.
FSI-NFS	FSI Network File System. ProtecTIER emulates UNIX file system behavior and presents a virtualized hierarchy of file systems, directories, and files to UNIX based clients using the NFS protocol. When configured for FSI-NFS, ProtecTIER emulates a network-attached storage (NAS) backup target that can use both HyperFactor and ProtecTIER Native Replication bandwidth reduction techniques for storing and replicating deduplicated data.
System console	The system console that is used with the TS7650 Appliance products and TS7650G Gateway is the IBM TS3000 System Console (TSSC). The ProtecTIER TS7620 SMB appliance has integrated system console functions and therefore does not require a TS3000 system console.
Metadata	Metadata is the information that is used to track the user data that is sent from the backup servers, including where it is stored on the disk.
User data	User data consists of the backup files and data sets that are stored in the ProtecTIER system. It is the data that the backup applications are storing on disk.
Repository	The repository is the physical disk that holds the ProtecTIER factored data. There are two types of file systems that make up the ProtecTIER Repository: metadata and user data.
Front end	The connection from ProtecTIER to the backup server.
Back end	The connection from ProtecTIER to the attached disk storage subsystem that acts as a repository.
Node and server	A single ProtecTIER system. It can be a TS7650G Gateway, TS7650 Appliance, TS7610 Appliance Express, or TS7620 Appliance Express, and is represented as a node from the ProtecTIER Manager software. Stand-alone node or dual-node cluster configurations are available. This book uses the terms node and server interchangeably.
IBM Tivoli Assist On-site (AOS)	A web-based tool that enables a remote support representative in IBM to view or control the management node desktop. For more information, go to the Tivoli AOS website at http://www.ibm.com/support/assistonsite.

1.1.1 ProtecTIER Appliance terminology

The ProtecTIER Appliance is the IBM self-contained virtualization solution that includes an embedded pre-configured disk storage repository. The following terms are specific to the TS7650 Appliance:

Disk controller The disk controller for the TS7650 Appliance is the 4.8 TB Fibre Channel Disk Controller (Feature Code 3708). Use this feature code for ordering or service purposes.

Disk expansion unit The disk expansion unit for the TS7650 Appliance is the 4.8 TB Fibre Channel Disk Expansion Unit (Feature Code 3707). Use this feature code for service purposes.

Server The 3958 AP1 server, which is based on the IBM System x3850 X5 Type 7145-AC1, was withdrawn as of 30 November 2012. During this interim process, some replacements are available (TS7650G Server 3958-DD5 and TS7620 Appliance 3959-SM2) with ProtecTIER Version 3.3. The link for the withdrawal announcement for AP1 can be found at http://www-01.ibm.com/common/ssi/index.wss?request_locale=en. Look for announcement number 912-173.

3959 SM1 TS7610 ProtecTIER Deduplication Appliance Express. A self-contained virtualization solution that includes an embedded pre-configured disk storage repository. It is no longer available for purchase.

3959 SM2 TS7620 ProtecTIER Deduplication Appliance Express. A self-contained virtualization solution that includes an embedded pre-configured disk storage repository. It has a base unit, which has two capacity versions (6 TB and 12 TB), and it can have up to two expansion units providing more capacity (23 TB and 35 TB).

3959 EXP The 3959 EXP expansion drawer enhances the capacity and improves the performance of the TS7620 Appliance Express. The base unit with one field expansion drawer offers 23 TB repository capacity. The base unit with two field expansion drawers offers 35 TB of repository capacity.

1.1.2 ProtecTIER gateway terminology

The TS7650G ProtecTIER Deduplication Gateway (TS7650G) is the IBM virtualization solution that does not ship with a disk storage repository. The customer can choose a solution from various storage options in order to build the back-end disk repository. IBM supports two clustered pairs of TS7650G servers in a single frame.

There are four types of servers that can be used in the TS7650G:

3958 DD5 This server is the newest, high performance server (available since May 2012), shipped with ProtecTIER Version 3.2 or higher. This server is based on the IBM System x7145 model. When used as a server in the TS7650G, its machine type and model are 3958 DD5. Use this machine type and model for service purposes.

3958 DD4 This server is a newer, higher performance server (available since December 2010). This server is based on the IBM System x3850 X5 Type 7145-AC1. When used as a server in the TS7650G, its machine type and model are 3958 DD4. Use this machine type and model for service purposes.

Note: Effective August 31, 2012, IBM withdrew the 3958 DD4 from marketing. For more information, see the IBM United States Withdrawal Announcement 912-096, June 4, 2012

3958 DD3 This server (available since March 2009) is based on the IBM System x3850 M2 Type 7233. When used as a server in the TS7650G, its machine type and model are 3958 DD3. It is no longer available for purchase.

3958 DD1 This server is the original server (available since August 2008). This server is based on the IBM System x3850 M2 Type 7141. When used as a server in the TS7650G, its machine type and model are 3958 DD1. IBM withdrew this product from market and terminated support.

Disk controller The customer must choose a disk controller for use with the TS7650G gateway. A list of compatible controllers is at the IBM Tape Systems Resource Library website found at http://www.ibm.com/systems/storage/tape/library.html.

Disk expansion unit The customer must choose a disk expansion unit for use with the TS7650G gateway. A list of compatible expansion units is at the IBM Tape Systems Resource Library website fond at http://www.ibm.com/systems/storage/tape/library.html.

1.1.3 ProtecTIER replication terminology

Replication enables you to set rules (depending on your required replication needs) for replicating data objects across ProtecTIER repositories. The ProtecTIER repositories can be different in size and physical layout. Because ProtecTIER deduplicates data before storing it, only the changes are transferred to the remote site. These rules for replicating data objects are defined in replication policies on each ProtecTIER system.

The replication function allows TS7600 deployment to be distributed across sites. Each site has a stand-alone or clustered ProtecTIER environment. Each ProtecTIER environment has at least one ProtecTIER server. ProtecTIER servers come with two dedicated replication ports. Replication ports are connected to the customer's wide area network (WAN) and are configured on two subnets by default.

The following terms define replication in a ProtecTIER context:

ProtecTIER Replication Manager A software component that is installed on a ProtecTIER server or a dedicated host. The ProtecTIER Replication Manager remotely manages the configuration of the grid (for example, grid creation/deletion and repository membership in the grid) In most cases, the ProtecTIER Replication Manager is on the ProtecTIER server. An agent on each ProtecTIER server interacts with the ProtecTIER Manager and maintains a table of its grid members.

Replication grid	A set of repositories that shares a common replication ID and can potentially transmit and receive logical objects through replication. A replication grid includes up to 24 ProtecTIER repositories and the connections between them. The replication grid is configured using the ProtecTIER Replication Manager.
Replication group	Also known as a Topology group. A replication group defines the relationships between nodes in a grid and determines which nodes are allowed to replicate to other nodes. Currently, there are four types of Topology groups: VTL many-to-one, VTL Bidirectional (many-to-many), OST Hub mesh, and FSI Bidirectional.
Many-to-One Topology Group	For VTL, this configuration is also known as *Spoke and Hub*. Up to 12 ProtecTIER system (spokes) can replicate to a single ProtecTIER node. (hub). The Hub node can go into DR mode for one or more spoke nodes simultaneously while still receiving replication data from the remaining nodes. If the hub is a TS7620, the maximum number of connected spokes is four.
Many-to-Many Topology Group	Bidirectional replication among multiple ProtecTIER nodes. Each node can define multiple replication targets, with up to 12 nodes per group in an OST topology, up to eight nodes per group in an FSI topology, and up to four nodes per group in a VTL topology.
Hub	For VTL, this item is a replication *and* a backup target. It receives replicated data from up to 12 spokes in a many-to-one replication group.
Spoke	For VTL, this item is a backup source that can replicate only to a single hub in many-to-one replication groups. Spokes are not applicable in many-to-many replication groups, as all nodes are considered to be both hubs and spokes.
Replication grid ID	A number 0 - 63 that identifies a replication grid within an organization.
Replication grid member	A repository that is a member in a replication grid.
Replication pairs	Two repositories within a replication grid that replicate from one to another.
Replication policy	A replication policy is defined on a ProtecTIER server and is made up of rules that define a set of objects (for example, VTL cartridges) from a source repository to be replicated to one or more target repositories.
Repository unique ID (RID)	A number that uniquely identifies the repository. The RID is created from the replication grid ID and the repository internal ID in the grid.
Replication time frame	A scheduled period for replication to take place for all policies.

Visibility	This term represents whether an application backup server can see or has visibility to a VTL cartridge. This construct is unique to VTL technology; ProtecTIER ensures that a tape is accessible only in one place at a time.
Visibility switching	The automated process that transfers the visibility of a VTL cartridge from its master to its replica and vice versa. Visibility switching is defined in the replication policy. The process is triggered by moving a cartridge to the source library Import/Export (I/E) slot. The cartridge then disappears from the I/E slot and appears at the destination library's I/E slot. To move the cartridge back to the source library, the cartridge must be ejected to the shelf from the destination library. The cartridge then disappears from the destination library and reappears at the source I/E slot.
Principality/ownership	An attribute that is set at the repository where an individual cartridge can be updated or written to by a backup application. A cartridge at its principal repository can be in read/write (RW) or read-only (RO) mode. At other sites, it is always RO. Each cartridge has enabled principality/ownership for one site only.
Dirty bit	The dirty bit attribute (in-sync) helps identify consistency points during disaster recovery (DR). When the dirty bit is off for a cartridge at the hub, this cartridge is fully synchronized with the spoke. During DR, that cartridge is fully synchronized, not only at the consistency point, but also after the last replication occurred. If a cartridge is out of sync during DR, determine the particular consistency point to which the cartridge adheres. Generally, the cartridge adheres to the consistency point that was established the last time that the cartridge was fully synchronized.
Disaster recovery	Disaster recovery (DR) is the process of recovering production site data at a remote location. It includes a way to indicate to a remote repository that the production site went down and notifies an administrator to initiate data recovery process.
Failover	Failover is a process of enabling the production at a remote site when there is a critical event or disaster at the primary site. It can be initiated intentionally if the primary site is under threat of a catastrophe and it is beneficial to perform takeover at the remote site with full control.
Failback	A process that is initiated from the remote site when the source site is again available for operation. The process ensures that the paired repositories are resynchronized using the least amount of bandwidth and maintaining the most recent copies of backups.

Remote Destination Directory	Applies only to FSI replication. It is a dedicated directory at the remote replication destination. Used to replicate a file system's directories and all objects that are contained in those directories recursively.
Cloning	Cloning creates a space-efficient, writable point-in-time copy of a Replication Destination Directory (RDD). Cloning an RDD is used for disaster recovery (DR) testing without disrupting ongoing replication to the RDD.

1.2 ProtecTIER basic concepts

ProtecTIER is a Data Protection and Retention (DP&R) product that appears to the backup servers as one of three standard interfaces:

Virtual Tape Library (VTL)	A standard tape library with a robot, cartridges, and tape drives with support for *LAN-free* data streams from the hosts.
Open Storage Technology (OST)	With OST, ProtecTIER can be integrated with Symantec NetBackup to provide backup-to-disk without having to emulate traditional tape libraries.
File System Interface (FSI)	The ProtecTIER FSI emulates file system behavior and presents a virtualized hierarchy of file systems, directories, and files to clients.

As data is written to the ProtecTIER device, it is examined for identical blocks of information that already were added to the repository. This identical data is not stored again in the repository; it is referenced as duplicate data and reduces the amount of disk space that is required. This process is known as *deduplication*. The engine for ProtecTIER deduplication is called *HyperFactor*.

For the deduplication process to work effectively, the data that is sent to the ProtecTIER system must not be manipulated, that is, modified as it passes from the disk drive on the client to the ProtecTIER system. Any change reduces or eliminates the ability of the HyperFactor engine to recognize a subsequent version of it.

Figure 1-1 shows the relationship between HyperFactor (HF) processing and the respective storage savings.

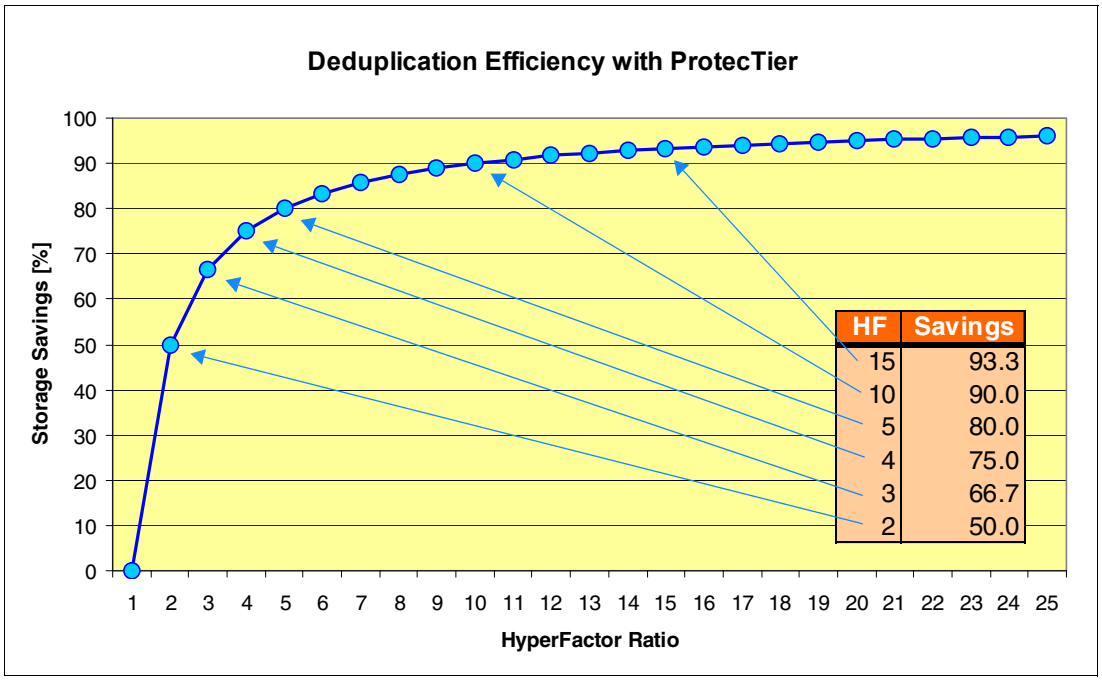

Figure 1-1 Storage savings with ProtecTIER HyperFactor

The effect and result of HyperFactor processing is a *factoring ratio*. In simple words, the factoring ratio is the ratio of nominal data (as a sum of all user data backup streams) to the occupied physical storage in ProtecTIER repository.

There are two key elements that impact your factoring ratio and effectiveness of your ProtecTIER environments:

Data retention period The period, which is measured in days, for which the backup application keeps the data available in its disk or tape repository. Typical retention periods of user data and applications are not longer than 90 days. Typical retention periods for database backups are not longer than 30 days, but specific weekly or monthly snapshots can be stored for months or even years.

Tip: A longer retention period increases the factoring ratio of your ProtecTIER product.

Data change rate The average percentage of data that is received from the backup application that changed from the previous backup iteration. This parameter is important when you size the backup repository and planning for optimal backup windows of various applications or systems. Examples of typical applications to be sized are progressive forever incremental backups and online database backups (IBM Lotus®, IBM Domino®, Oracle, SAP, IBM DB2®, Microsoft SQL, MaxDB, and so on). The data change rate might vary from platform to platform (from 1% to greater than 25%).

Tip: Less changing of data improves the deduplication process and boosts the factoring ratio.

For more information about the concepts of IBM ProtectTIER deduplication products, see *IBM System Storage TS7600 with ProtecTIER Version 3.1*, SG24-7968 or go to the following website:

http://www.ibm.com/systems/storage/tape/enterprise/virtual.html

1.2.1 ProtecTIER with deduplication disabled

The TS7650G can exclude the purchase of the capacity license and use the ProtecTIER gateway as a VTL only. This enablement reduces your costs when you implement a small to medium sized backup solution and improves restore performance.

For certain scenarios, especially in small to medium business environments, clients might want to use existing SAN networks and disk storage subsystems for backup purposes by using a VTL, but without deduplication techniques. This requirement is valid for backup solutions with short-term retention periods, where the high frequency and amount of data that is changed on the hosts daily makes restore performance from physical tape not optimal. This situation includes, for example, the backup of large file servers with millions of small files.

In this case, clients implement the ProtecTIER gateway that is connected to the existing supported disk storage subsystem as a repository, but with the deduplication function set to *disabled*. Your environment can then benefit from the VTL by using existing storage products, eliminating the need to buy more automated tape libraries (ATL), while still offering the capability to use replication features.

Note: It is not possible to exclude the capacity license in ProtecTIER appliances (models AP1, SM1, and SM2) where minimal available capacity must be always chosen.

1.3 ProtecTIER models for open systems

This section provides an overview of the existing models of ProtecTIER servers on the market (Figure 1-2). This section does not go in to all the technical details. For more technical details and comprehensive information about each model, see *IIBM System Storage TS7600 with ProtecTIER Version 3.1*, SG24-7968 or refer to the following website:

http://www.ibm.com/systems/storage/tape/midrange/virtual.html

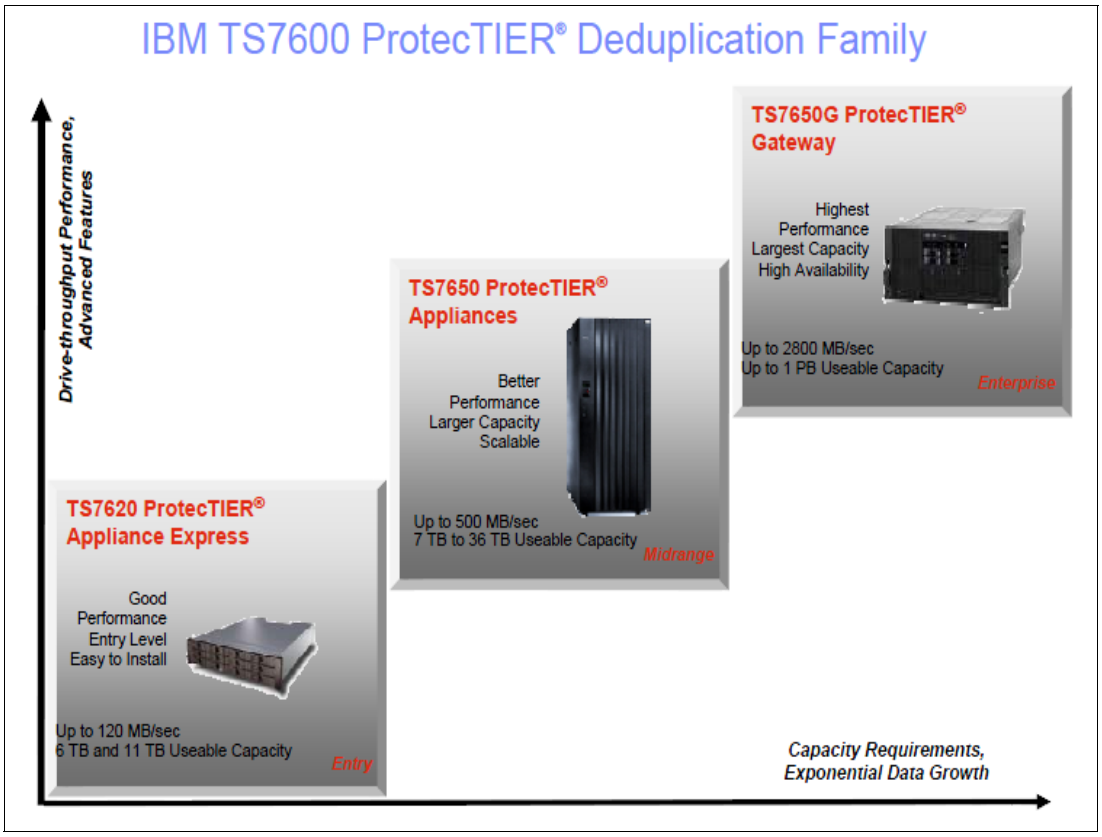

Figure 1-2 ProtecTIER product line

1.3.1 ProtecTIER appliances

The IBM System Storage TS7600 ProtecTIER Appliances are designed and built with the repository pre-configured on internal disk drives. Customers select one of two possible capacity options when ordering.

IBM System Storage TS7620 ProtecTIER Deduplication Appliance Express

Available in two configuration options, the IBM System Storage TS7620 ProtecTIER Deduplication Appliance Express (TS7620) model 3959-SM2, is an integrated server and storage hardware platform that ships with IBM ProtecTIER deduplication software preinstalled. With an available repository capacity 5.5 or 11 TB, the appliance is targeted at small or medium business backup environments (SMB). Clients can choose between VTL configuration, OST, or the newly available File System Interface (FSI) as a backup target.

The ProtecTIER Deduplication Appliance Express is shown in Figure 1-3.

Figure 1-3 BM System Storage TS7620 ProtecTIER Deduplication Appliance Express

The product features and hardware summary include the following items:

- VTL, OST, or FSI support
- A single 3U integrated appliance in 5.5 TB and 11 TB physical capacities
- For VTL/FSI Performance, up to:
  - 145 MBps backup
  - 190 MBps restore
- OST Performance up to:
  - 130 MBps backup
  - 170 MBps restore
- Plug and Play installation
- One LSI MegaRAID card and battery backup and six 8 GB memory DIMMs (a total of 48 GB of memory)
- One-to-One, Many-to-One, and Many-to-Many replication features

IBM System Storage TS7650 ProtecTIER Deduplication Appliance

The IBM System Storage TS7650 ProtecTIER Deduplication Appliance (TS7650) is a preconfigured solution of IBM storage products, IBM servers, and the IBM ProtecTIER data deduplication software that is preinstalled in a standard 19-in. rack. It improves backup and recovery operations, including remote replication with disaster recovery functions.

The solution is available in four configurations (7 TB, 18 TB, 36 TB, or 36 TB dual node cluster) to meet the disk-based data protection needs of various organizations, from mid-sized IT environments to enterprise data centers.

Figure 1-4 shows a fully equipped TS7650 Appliance.

Figure 1-4 IBM System Storage TS7650 ProtecTIER Deduplication Appliance

Here are the hardware product features:

- Emulation of up to 12 virtual libraries, 256 virtual drives, and 128,000 virtual cartridges when in a cluster configuration.
- Up to 500 MBps or more inline data deduplication performance.
- VTL, OST, or FSI support.
- One-to-One, Many-to-One, and Many-to-Many replication to support disaster recovery.
- Preconfigured with an IBM System Storage DS5020 storage subsystem as a repository.

1.3.2 ProtecTIER deduplication gateways

IBM ProtecTIER Deduplication Gateway products offer clients the capability to use their existing IBM or third-party disk storage subsystems and SAN infrastructure as a repository. For a list of supported disk storage subsystems, see the TS7650/TS7650G ISV and Interoperability Matrix at:

http://www.ibm.com/systems/storage/tape/resources.html

This section provides a brief overview that describes the available gateways.

IBM System Storage TS7650G ProtecTIER Deduplication Gateway

The IBM System Storage TS7650G ProtecTIER Deduplication Gateway (TS7650G) shown in Figure 1-5 is available in stand-alone and clustered configurations. For a stand-alone configuration, one IBM machine type and model 3958-DD5 server is required (or any previously supported model). For a clustered configuration, two 3958-DD5, 3958-DD4, or 3958-DD3 servers are required (or any combination), along with a Cluster Connection Kit, which includes two required Ethernet switches and one network attached power switch. The existing, withdrawn, and stand-alone 3958-DD1 servers may be upgraded to a clustered configuration by clustering a 3958-DD1 server and a 3958-DD3 or 3958-DD5 server. The 3958-DD1 and 3958-DD4 gateway servers cannot be clustered. In all cases, the clustered servers must be installed in the same physical frame.

Tip: To gain the maximum performance from your clustered ProtecTIER product, use two of the same model as cluster nodes.

The disk storage array attaches to the TS7650G through Fibre Channel connections and holds the repository of deduplicated backup data. The amount of cache that is available depends on your disk subsystem and configuration.

Figure 1-5 IBM System Storage TS7650G ProtecTIER Deduplication Gateway

The TS7650G offers the following features and functions:

- Emulation of up to 16 virtual tape libraries and 256 tape drives per cluster node.
- Four LSI MegaRAID cards and battery backup and sixteen 4 GB memory DIMMs (for a total of 64 GB of memory).
- Scales to 1 PB of physical back-end storage.
- Cloning of virtual cartridges to physical tapes at remote sites.
- VTL, FSI, or OST configuration, where only one of the modes can be enabled on the same TS7650G.

Note: ProtecTIER offers true highly available active-active dual-node clustering for VTL and OST models. The initial configuration of the FSI model is available as a single-node configuration.

1.4 Remote support and Call Home

Remote support is available for the TS7650G and the TS7650 Appliance through the *Call Home* capability on the IBM TS3000 System Console (TSSC). The Call Home feature reports failures that are detected by the ProtecTIER servers. Whenever a failure is detected, Call Home sends detailed error information to the IBM Service Center (home). An IBM System Service Representative (SSR) can then prepare an action plan to handle the problem before he travels to the affected installation. The appliance or gateway might also periodically send support information (such as configuration, code versions, and error logs) to IBM. Doing so speeds up the problem determination and fault resolution process. When enabled on the appliance and gateway, Call Home uses a connection on your Ethernet network to transmit hardware and software problem reports to IBM. Call Home is enabled and tested by SSRs during the initial system installation.

When the Reliability, Availability, and Serviceability (RAS) software on the ProtecTIER server detects an error condition, Call Home sends detailed error information through the TSSC to IBM. If the error indicates a problem with a field replaceable unit (FRU), an SSR can pre-order the affected unit for optional replacement at the site.

The TS7650G and the TS7650 Appliance provide four Call Home capabilities:

Test Call Home The SSR sends a Test Call Home signal after you enable the Call Home feature during the initial installation. You can also send a Test Call Home to ensure that the setup is correct and that the appliance or gateway can successfully open a Problem Management Record (PMR) in the IBM Remote Technical Assistance Information Network (IBM RETAIN®).

Problem Call Home When the RAS software detects a problem, it initiates a Call Home operation to create a PMR in RETAIN. The PMR is a single page of text data that enables the Support Center or an SSR to access an action plan and a list of applicable FRU components.

Heartbeat Call Home To ensure proper ongoing Call Home functionality, the system can be configured to send a Heartbeat Call Home on a regularly scheduled basis. The heartbeat interval is user-defined.

Heartbeat interval: The optional Heartbeat Call Home can be an interval of 1 - 14 days. Failed attempts are logged in the system console TSSC and can be monitored by an administrator.

User-Initiated Call Home You can manually initiate Call Home from the TSSC GUI to collect a product engineering (PE) package.

The RAS software sends data files that might be helpful to IBM Support Center personnel for all four types of Call Home. These data files include error logs and configuration information, such as the Machine Reported Product Data (MRPD) log.

Important: You must configure RAS software on the ProtecTIER server before you set up Call Home.

The TSSC is a service tool that must be present in an IBM supported TS7650G or TS7650 Appliance. You can either order a TSSC with your appliance or gateway, or use a TSSC already installed at your site.

Important: Although it is possible to operate a TS7650G or TS7650 Appliance without a connected TSSC, IBM does not support this configuration.

All TS3000/TSSC related tasks are IBM support personnel responsibility.

1.4.1 How to set up Call Home

This section provides an overview and the necessary steps to configure and test the Call Home feature and how to find these events in the system console.

Important: Be sure to complete the Call Home test message process. Failure to verify that Call Home is working correctly on the new system can result in Call Home messages that fail to send.

You must configure RAS software on the ProtecTIER server before you set up Call Home.

To set up Call Home, complete the following steps:

1. Log on to the ProtecTIER system and access the ProtecTIER service menu by enter the **menu** command. The menu that is shown in Example 1-1 opens.

Example 1-1 Initial screen of ProtecTIER service menu

```
--------------------------------------------------------------------------
ProtecTIER Service Menu running on tuc1dd5
--------------------------------------------------------------------------
1) ProtecTIER Configuration (...)
2) Manage ProtecTIER services (...)
3) Health Monitoring (...)
4) Problem Alerting (...)
5) Version Information (...)
6) Generate a service report
7) Generate a system view
8) Update ProtecTIER code
E) Exit
--------------------------------------------------------------------------
>>> Your choice?
```

2. Select the Problem Alerting option and press Enter. The screen that is shown in Example 1-2 opens. Select Enable/Disable Call Home.

Example 1-2 Call Home options in ProtecTIER service menu

```
--------------------------------------------------------------------------
ProtecTIER Service Menu running on rasddx
Problem Alerting (...)
--------------------------------------------------------------------------
1) Enable/Disable Call Home
2) Send a Test Call Home
3) Configure Call Home Heartbeat frequency
4) Send a Heartbeat Call Home
```

```
5) Enable/Disable Notification by email
6) Test Email Notification
7) Activate Call Home Polling Function
B) Back
E) Exit
--------------------------------------------------------------------
>>> Your choice?
```

If you receive a message that states that Call Home is already enabled, ignore the message and continue.

3. Enter the number that corresponds to Send a Test Call Home and press Enter. After a few seconds, you should receive the following message:

   ```
   Test Call Home sent successfully
   ```

   The TS3000 system console's Call Home Queue should have the test message listed as a pending transmission.

4. Right-click an empty area of the TSSC desktop and click **System Console Actions** → **Console Configuration Utility** → **Call Home Queue**. If prompted for a user name and password, enter service for both.

   If the Call Home Queue is empty, the call already might have been sent to the IBM Customer Configuration Profile File (CCPF) system.

   Important: Each TS7650 server must be under valid warranty or Maintenance Agreement (MA) coverage or it is rejected and no record or Problem Management Report (PMR) is generated for the Call Home event.

5. Click **Call Home Event Log** in the Call Home Queue and navigate to **Browser Functions** → **Call Home Log**. Check for the records that contain the serial number of the applicable server.
   - If the record appears in the log, the test completed successfully.
   - If the record does not appear in the log, and service code 20 is registered for this system, and you are certain that the server has a valid warranty and MA coverage, contact the next level of support. If service code 20 is not registered, disregard this check.
6. Close any open browser windows and initiate the configuration of Call Home Heartbeat frequency.

1.5 ProtecTIER command-line interface

This section provides a brief overview of the ProtecTIER command-line interface (ptcli). The ptcli is run from the workstation or on the ProtecTIER node. The ptcli is a low-level entry point into the ProtecTIER system. It provides information that can be formatted and queried to provide wanted data, and direct the output to a file for persistent storage. This information provides the administrator with data that provides insight into the long-term operation of the system's performance, capacity, configuration, and effect of operational changes on the system. The information can be used as input to the administrator's other management applications.

The ptcli is loaded during the installation of ProtecTIER software and ProtecTIER Manager software. The ptcli can be used to complete any of the following tasks:

- Configure ProtecTIER (including configuration of a ProtecTIER repository and configuration of ProtecTIER virtual tape libraries).
- Monitor ProtecTIER (including statistics of virtual tape libraries and statistics about the repository).
- Take snapshots of and filter ProtecTIER virtual tape cartridges (mostly used for DR scenarios).

For more details about how to install and use the ptcli, refer to Appendix A, “Command Line Interface” of *IBM System Storage TS7600 with ProtecTIER Version 3.3*, SG24-7968.

2

Deduplication considerations

This chapter describes the ProtecTIER data deduplication concepts, methods, and system components. This chapter also elaborates on the benefits of data deduplication with HyperFactor, and on general ProtecTIER deduplication considerations. This chapter also describes data type candidates for high factoring ratios, and describes data types that can have a negative impact on factoring ratios.

This chapter describes the following topics:

- HyperFactor data deduplication
- ProtecTIER HyperFactor deduplication processing
- Components of a ProtecTIER system
- Benefits of ProtecTIER HyperFactor
- General ProtecTIER deduplication considerations
- Data types

2.1 HyperFactor data deduplication

Data deduplication is used to reduce the space that is needed to store data on disk. This situation is achieved by storing only a single instance of data that is backed up repetitively. Data deduplication is not a storage device; it is a function of a system, for example, a VTL, an OST API interface, or a File System Interface (FSI). Data deduplication is not an I/O protocol, but it requires an I/O protocol for data transfer, for example, Fibre Channel Protocol (FCP), Common Internet File System (CIFS), Network File System (NFS), or an application programming interface (API).

With deduplication, repeated instances of identical data are identified and stored in a single instance, as shown in Figure 2-1. This process saves storage capacity and bandwidth. Data deduplication can provide greater data reduction than previous technologies, such as Lempel-Ziv (LZ) compression and differencing, which is used for differential backups. It does not always make sense to use data deduplication because not all types of data can be deduplicated with identical efficiency. Data deduplication might interfere with other technologies, such as compression, encryption, or data security requirements. Data deduplication is not apparent to users and to applications.

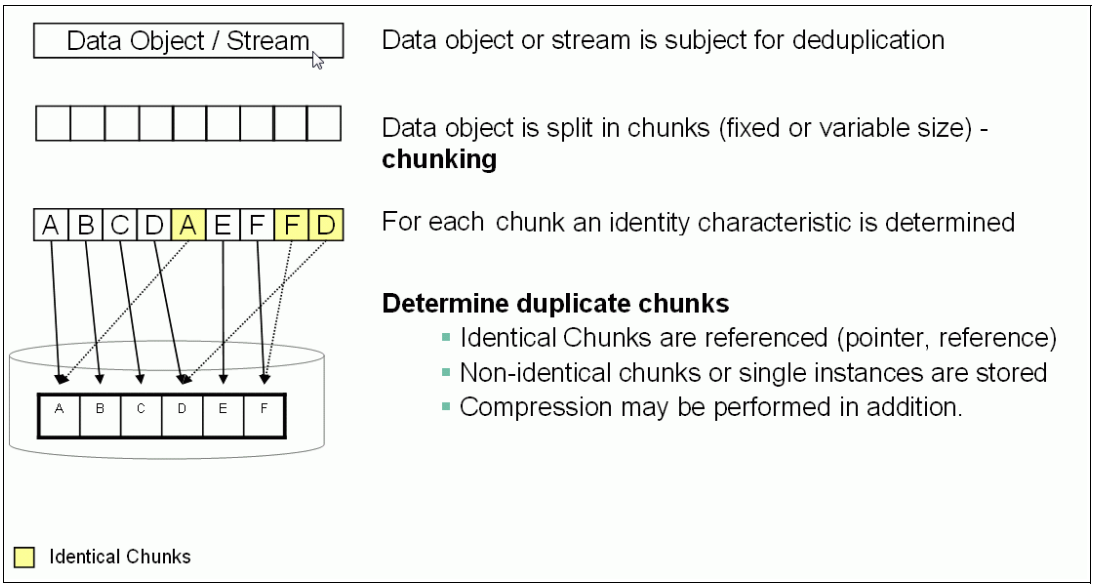

Figure 2-1 Simplified data deduplication process

With data deduplication, the incoming data stream is read and analyzed by the ProtecTIER HyperFactor algorithm while it looks for duplicate data. Using inline processing, ProtecTIER ensures high performance, scalability, and 100% data integrity while it compares data elements of variable size to identify duplicate data. After the duplicate data is identified, one instance of each element is stored, pointers are created for the duplicate items, and the duplicate items are not stored but only referenced.

The effectiveness of data deduplication depends on many variables. The data change rate, the amount of new data that is backed up, the number of backups, the amount of repetitive or similar data in your backups, and the data retention period have a major impact on the effectiveness of data deduplication. For example, if you back up the exact same uncompressible data once a week for six months, you save the first copy and do not save the next 24, which would provide a 25:1 data deduplication ratio. If you back up an uncompressible file on week one, back up the exact same file again on week two and never back it up again, you have a 2:1 deduplication ratio.

A more likely scenario is that some portion of your data changes from backup to backup so that your data deduplication ratio changes over time. For example, assume that you take weekly full and daily differential incremental backups. Assume that your data change rate for the full backups is 15% and for the daily incrementals is 30%. After 30 days, your deduplication ratio might be around 6:1, but if you kept your backups up to 180 days, your deduplication ratio might have increased to 10:1.

These examples, and the remainder of this book, describes the deduplication ratio as being the nominal data (total backup data that is received) divided by the physical data (amount of disk space that is used to store it).

Data deduplication can provide storage savings, but the benefit that you derive is determined by your data and your backup policies. Workloads with a high database content generally have the highest deduplication ratios. However, product functions such as IBM Tivoli Storage Manager Incremental Forever, Oracle RMAN, or LiteSpeed for SQL Server, can affect the deduplication ratio. Compressed, encrypted, or otherwise scrambled workloads typically do not benefit from deduplication because the potential deduplication candidates are no longer similar. For more information, see 2.6, “Data types” on page 34.

2.1.1 HyperFactor, deduplication, and bandwidth savings

The cornerstone of ProtecTIER is HyperFactor, the IBM technology that deduplicates data inline as it is received from the backup application. ProtecTIER bandwidth-efficient replication, inline performance, and scalability directly stem from the technological breakthroughs inherent to HyperFactor. HyperFactor is based on a series of algorithms that identify and filter out the elements of a data stream that was stored by ProtecTIER. Over time, HyperFactor can increase the usable capacity of an amount of physical storage by 25 times or more.

With replication, the data reduction value of HyperFactor is extended to bandwidth savings and storage savings for the Disaster Recovery (DR) operation. These performance and scalability attributes are critical for the DR operation, in addition to the primary site data protection operation.

When new data is received by the ProtecTIER native replication technology, HyperFactor finds any similar data elements that are already stored. This search is quick and uses a small and efficient memory-resident index. After similar data elements are found, HyperFactor can compare the new data to the similar data to identify and store only the byte-level changes.

With this approach, HyperFactor can surpass the reduction ratios that are attainable by any other data reduction method. HyperFactor can reduce any duplicate data, regardless of its location or how recently it was stored. When new data is received, HyperFactor checks to see whether similar data is already stored. If similar data is already stored, then only the difference between the new data and previously stored data must be retained. This technique is an effective and high performance one for identifying duplicate data.

Data deduplication using the HyperFactor technology identifies data similarities and checks those similarities against the fixed size Memory Resident Index every time new data is received. When similar matches are found, a binary differential comparison is performed on similar elements. Unique data with corresponding pointers is stored in the repository and the Memory Resident Index is updated with the new similarities. Existing data is not stored again.

HyperFactor data deduplication uses a fixed size 4 GB Memory Resident Index to track similarities for up to 1 PB of physical disk in a single repository. Depending on the data deduplication ratio for your data, you could store much more than 1 PB of data on your disk array. For example, with a ratio of 12:1, you could store 12 PB of data on 1 PB of a disk array. With the Memory Resident Index, HyperFactor can identify potentially duplicate data quickly for large amounts of data and it does this action on data ingest, or inline, reducing the amount of processing required for your data.

The read-back rate of the ProtecTIER deduplication technology is generally higher than the write rate to the system because there is no risk of fragmentation. No access to the index or heavy computation is required during a restore activity. It requires only that you open metadata files and fetch the data according to the pointers that they contain.

2.2 ProtecTIER HyperFactor deduplication processing

Data deduplication is performed while the data is being backed up to the ProtecTIER (inline) server, in contrast to after the data is written to it (post processing). The advantage of inline data deduplication is that the data is processed only once and there is no additional processing that is performed after the backup window. Inline data deduplication requires less disk storage because the native data is not stored before data deduplication.

2.3 Components of a ProtecTIER system

The ProtecTIER data deduplication system consists of three main components. Two of these components, the ProtecTIER server and the HyperFactor deduplication software, are always bundled together for your convenience. Depending on the model of ProtecTIER you look at, you might have a bundled disk storage subsystem. You can provide an individual storage subsystem, or even to share an existing one. For an overview of ProtecTIER models, see 1.3, “ProtecTIER models for open systems” on page 13.

The components that are shown in Figure 2-2 are covered in the following three sections.

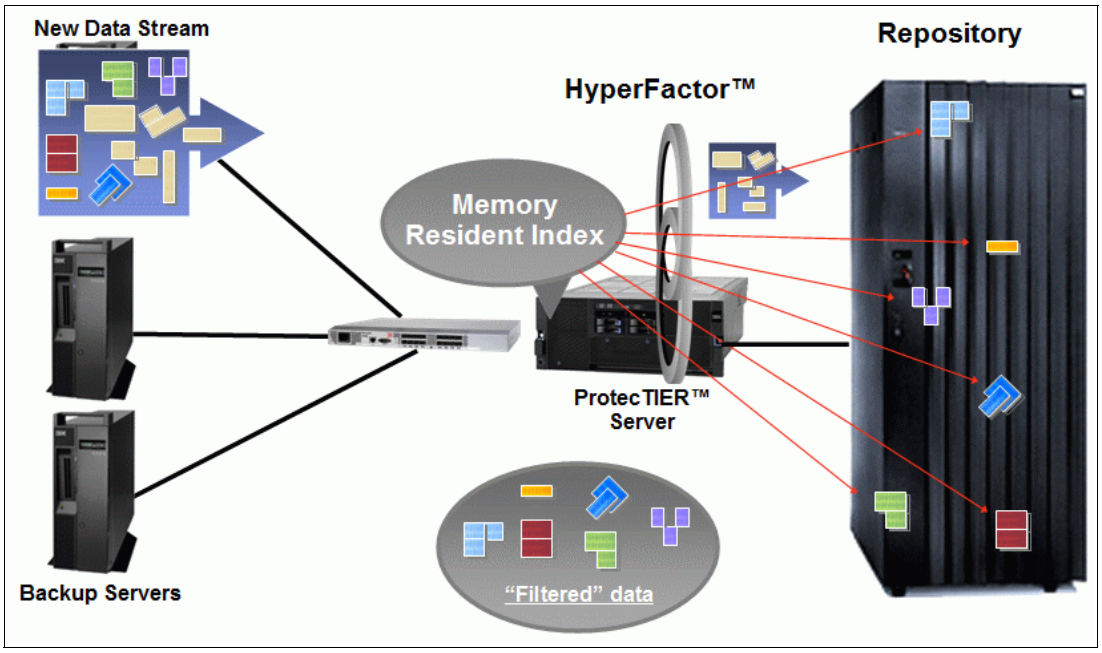

Figure 2-2 ProtecTIER components

2.3.1 ProtecTIER server

Every ProtecTIER deduplication system uses a server with an operating system on which the HyperFactor software runs. The TS7620 Appliance Express comes as a bundle that allows up to 145 MBps performance. The TS7650 Appliance comes with servers configured to allow up to 500 MBps performance. The TS7650G server is a high-performance configuration with four 10-core processors (DD5) and 64 GB of memory. A ProtecTIER high-availability, active-active cluster configuration is available with the TS7650 Appliance and the TS7650G solution. You can have up to two servers in a ProtecTIER cluster solution.

Upgrades: With ProtecTIER Version 3.2, the TS7620 Appliance Express (SM2) can replace the TS7610 Appliance Express (SM1). For a comparison of the TS7610 Appliance Express and the TS7620 Appliance Express, see the *IBM System Storage TS7610 and TS7620 ProtecTIER Deduplication Appliance Express ProtecTIER User's Guide for VTL Systems*, GA32-0916, the *IBM System Storage TS7610 and TS7620 - ProtecTIER Deduplication Appliance Express - ProtecTIER User's Guide for OpenStorage*, GA32-2230, or the *IBM System Storage TS7610 and TS7620 - ProtecTIER Deduplication Appliance Express - ProtecTIER User's Guide for FSI and NFS*, GA32-2231.

2.3.2 HyperFactor deduplication algorithm

The HyperFactor data deduplication solution can be ordered in three different styles. There is a virtual tape library (VTL) interface, the open storage (OST) API interface, and the file system interface (FSI) with common internet file system (CIFS) and network file system (NFS) support. As the different interface methods cannot be intermixed, you must choose one or deploy multiple ProtecTIER models simultaneously.

Note: With the ProtecTIER FSI model, you can have CIFS and NFS connectivity with the same machine.

2.3.3 Disk storage subsystem

Data that is processed by the ProtecTIER HyperFactor data deduplication software is stored on disk. The ProtecTIER appliance systems, for example, TS7610 Appliance Express, TS7620 Appliance Express, and TS7650 Appliance, come prebundled with disk storage included in a ready to run configuration. The ProtecTIER TS7650G Gateway server attaches to a wide variety of disk storage subsystems that separately must be made available. For a list of supported disk systems, see TS7650/TS7650G ISV and Interoperability Matrix at the following website:

http://www.ibm.com/systems/storage/tape/resources.html

Compression: ProtecTIER does not require more compression to be effective. ProtecTIER performs compression after the deduplication process completes, unless you decide to turn off compression, which is not recommended. Allocation of disk arrays for ProtecTIER that perform their own compression is not recommended, as there is no additional benefit.

If you want to use encryption, attach ProtecTIER to a back-end storage system that supports encryption instead of backing up encrypted data to ProtecTIER. This action has a beneficial effect on your data deduplication ratio.

The ProtecTIER back-end storage subsystem must be a random access disk storage subsystem from the supported back-end storage list. Consult the TS7650/TS7650G ISV and Interoperability Matrix.

Attaching a physical tape drive: Attaching a physical tape drive directly to a ProtecTIER system is not supported. However, operating your backup application with a combination of ProtecTIER and physical tape drives is feasible.

2.4 Benefits of ProtecTIER HyperFactor

When appropriately deployed, data deduplication can provide benefits over traditional backups to disk or virtual tape libraries. Data deduplication enables remote vaulting of backup data using less bandwidth because only changed data is shipped to the remote site. Long-term data retention for local or offsite storage might still be achieved most economically with physical tape.

2.4.1 Flexibility

ProtecTIER is independent from the backup application. You can combine multiple backup applications from different vendors to work with one single ProtecTIER solution. All attached backup solutions directly benefit from the whole ProtecTIER deduplication potential. There is no drawback in sharing the repository with multiple backup applications.

2.4.2 High availability

ProtecTIER offers true highly available active-active dual-node clustering for VTL and OST models. Mimicking the behavior of a physical tape library, you can use the ProtecTIER solution to access your data even if a node is unavailable. The initial configuration of the FSI model is available as single node.

2.4.3 High performance, reduced storage requirements, and environment friendly

Data deduplication can reduce the amount of disk storage that is required to store data and keep it online. Performing restores from disk can be faster than restoring from tape, and having the data online for longer periods reduces the possibility that the required data might be shipped offsite. Inline deduplication has no need for more post-processing space and therefore further reduces space requirements. If data deduplication reduces your disk storage requirements, the environmental costs for running and cooling the disk storage are also reduced.

2.5 General ProtecTIER deduplication considerations

The following general considerations of ProtecTIER deduplication can help you better understand the “do's and don'ts” of deduplication.

2.5.1 Rethinking your overall backup strategy

The best practices for ProtecTIER can be achieved by adopting your environment by using the examples that are provided in this chapter. Be sure to revisit your backup strategy from a greater perspective. One of the biggest benefits of ProtecTIER is fast restore performance. Most clients are more interested in quickly restoring their data if the need should arise, as opposed to quickly backing up their data. Restoring your data quickly and efficiently is crucial to business continuity.

Rethink the method that you use to do backups to allow the fastest restore possible. For example, backing up data to a ProtecTIER server with only a few streams and by using only a few mount points is no longer necessary. Think big; it is all virtual and virtual tape drives are available at no additional cost.

Keeping the number of used cartridges low to save money no longer applies in the virtual world of ProtecTIER. Using as many cartridges in parallel as possible, to some extent, is a good idea. The maximum number of cartridges in a VTL is greater than 65,000. You do not need to use all of them, but you should plan on using more virtual resources than you would use physical resources. This guideline is true for virtual tape drives and virtual tape cartridges. This general approach is also true for OST and FSI deployments.

If you use methodologies such as client compression to reduce the load on your network, you might want to rethink compression as well. Most pipes are “fat”, meaning that your infrastructure has plenty of bandwidth to allow many uncompressed backups. This situation ensures faster backups and faster restores. This situation is true for network and Fibre Channel infrastructures. LAN-free backups within your data center can be possible if you do not have infrastructure bandwidth congestion.

If you perform incremental backups, especially for your databases, you might also want to rethink this process for critical applications. Multiple full backups, especially on a high frequency schedule, might appear to be a waste of space, but this situation is where you can benefit the most from ProtecTIER deduplication. A ProtecTIER server has the best deduplication, the highest backup speed, and the highest restore speed if you write multiple full backups of the same objects to it. Your infrastructure should be up to the challenge because resources tend to sit idle during non-backup hours. So why not increase the usage of your already available resources?

As an additional benefit, the restore performance is further increased by the reduced number of restore steps. With ProtecTIER technology, you do not need to restore your database by first restoring the latest full backup, then multiple incremental backups, and finally applying the database logs. Simply restoring the latest full backup and applying the logs is sufficient to be at your recovery point objective (RPO).

Evaluate these suggestions with your data protection peers, staff, and infrastructure strategists to transform your data protection strategy and achieve the most out of your solution. This task is not always easy. If you cannot change the current setup now, at least make sure that you have an impact on the planning for future deployments.

2.5.2 The number 32: The ProtecTIER product is not physical tape

When you use the ProtecTIER product, you encounter the number 32 many times. The potential number of arrays in the back end is 32. The number of parallel streams to use with your backup application is 32. The recommended number of tape drives to emulate with your ProtecTIER system is also 32.

Exception: Every rule has its exceptions. The estimated number of client sessions that are needed to saturate a maximum performance configuration of ProtecTIER is 32 for VTL, 128 for OST, and 16 for FSI.

What all of these statements have in common is that they highlight that the general mode of operation of the ProtecTIER system is different from a physical tape drive. Be it virtual tape libraries (VTL), an OST API interface, or the FSI, all of the modes of interfacing with the ProtecTIER system behave differently from a physical tape drive. The modes also behave differently than classic random access disks. A ProtecTIER system prefers to work with multiple parallel streams, with each of the streams using a sequential pattern of its own. Because of this situation, a massively parallel way of interfacing with a ProtecTIER system is preferred for a single stream operation. This situation is true for the front-end communication between your backup application and ProtecTIER and for the back-end communication of the ProtecTIER Gateway with its back-end storage. If you, for example, are limited to four streams per backup client, back up multiple clients to the ProtecTIER system at the same time. The throughput of one single client might not range up to your expectation, but the combined throughput of multiple clients with multiple streams achieves your wanted performance.

The guideline of the number 32 is a starting point, and gives you an idea of the dimensions we are talking about. The average number of physical tape drives used might be one, two, four, or eight in the most standard cases. A ProtecTIER system should be used with 32 streams, maybe even 64 or 128, or even more, especially for high performance gateway installations. This general guideline is true for the VTL, OST, and FSI versions of a ProtecTIER system. As stated above, the estimated number of client sessions that are needed to saturate a maximum performance configuration of a ProtecTIER system is 32 for VTL, 128 for OST, and 16 for FSI.

2.5.3 Data reduction technologies should not be combined

ProtecTIER data deduplication is a data reduction technology. Compression is another data reduction technology. IBM Tivoli Storage Manager is an example of an application that provides its own brand of compression and deduplication. Tivoli Storage Manager also offers incremental forever backup, which can be thought of as a data reduction technology. With incremental forever backup, only changed data is backed up. There are many other potential data reduction technologies.

Attention: Do not combine multiple data reduction technologies, as there is no benefit in compressing or deduplicating data multiple times. If your goal is to achieve a high deduplication ratio, disable all other data reduction technologies.

If you prefer to combine another data reduction technology with a ProtecTIER solution, there is a solution without deduplication that is also available. Ask your IBM marketing representative for a ProtecTIER solution without a capacity license.

Some of the scenarios that allow the combination of data reduction technologies are described in this section. For example, you can combine the IBM Real-time Compression™ Appliance (RTCA) with a ProtecTIER solution to even further increase your storage savings. For more information about a ProtecTIER solution with RTCA, see 13.2.1, “Combining RTCA and IBM ProtecTIER deduplication”, in *Introduction to IBM Real-time Compression Appliances*, SG24-7953.

Tivoli Storage Manager can combine both compression and deduplication within itself. The details are explained in Chapter 4, “Introduction to IBM Tivoli Storage Manager deduplication”, in *Implementing IBM Storage Data Deduplication Solutions*, SG24-7888.

IBM DB2 database software can handle data in a way so that it can be compressed within DB2 but still achieves high deduplication ratios. For more information about using DB2 compression with a ProtecTIER repository, see 20.4.1, “Combining DB2 compression and ProtecTIER deduplication” on page 328.

2.5.4 Data streams must be in order

Many technologies that are available for improving performance and throughput for physical tape drives do not work well with deduplication. Multiplexing, for example, shuffles the data, so you cannot identify potential candidates for deduplication in the data stream. If you aim for a narrow backup window, increase the number of streams, increase parallelism, and disable multiplexing. Disabling multiplexing improves the HyperFactor process and increases performance.

Encryption also results in shuffled data. A small change in an encrypted file results in a file that appears different to a deduplication solution. Potential deduplication candidates cannot be identified, as the patterns do not match anymore. If you plan to use encryption with your ProtecTIER deduplication solution, implement storage subsystem-based encryption for your ProtecTIER repository. This situation does not affect your HyperFactor duplication factoring ratio.

Analyze your environment for other potential data shuffling causes and aim at eliminating them.

2.5.5 Data organization within your ProtecTIER repository

The ProtecTIER repository is the place where your deduplicated data is stored. You can define one or many virtual tape libraries (VTLs) with multiple slots and cartridges. You can define one or many storage units for the OST API, or you can have multiple file shares for the FSI. No matter what type of ProtecTIER repository you use, logically segment your data and group similar backup types together. This setup allows detailed deduplication analysis that is based on cartridge granularity that done by your SSR. If you can supply a list of virtual cartridges, or a virtual cartridge range that contains one special type of backed up data for detailed analysis, this setup provides valuable data that you use to improve your data protection environment.

Organization: Apply a meaningful organization scheme to your backup data. For VTL, multiple slots and cartridges should align to different barcode ranges. For FSI, dedicated directories with meaningful names should align to dedicated backup servers.

2.5.6 The dynamics of the ProtecTIER repository

In addition to the data that you write into your ProtecTIER repository, there are two other major effects of the ProtecTIER repository that must be understood. First, the repository dynamically reacts to the quality of your data. If the data you back up to the ProtecTIER repository suddenly changes and allows a higher deduplication ratio, the repository adapts and can store more data. If the quality of your data changes and allows only a reduced deduplication ratio, the ProtecTIER repository also reacts to this change, and less data can be stored.

Repository size: The ProtecTIER repository size is calculated by the formula “Physical Repository Size” x “HyperFactor Ratio” = “Available Free Space for you to write to” (Nominal Free Space).

If your HyperFactor ratio changes, the available space for you to write to adapts.

A ProtecTIER repository is not directly aware of your data retention requirements. A ProtecTIER repository stores all data unless told otherwise. It is especially important for VTL emulations to specify whether you still need the data. As an example, Tivoli Storage Manage uses the relabel scratch option of its library definition to communicate to the ProtecTIER repository that the space of the virtual cartridge can be freed up. Other backup applications might rely on housekeeping scripts to initiate a label sequence / write sequence from the beginning of tape, which has the same effect. Make sure to regularly free up unused space.

After you release the unused space, it becomes marked as “Pending” within the ProtecTIER repository. The ProtecTIER repository then automatically uses internal processes to optimize the available space for future backups. Internal “Deleter” processes reduce the “Pending” space and, within the process, create “Fragmented” space. Internal “Defragger” processes then reduce the “Fragmented” space. In the right side of the pie chart in the ProtecTIER GUI (Figure 2-3), you can see the Nominal data, which displays 16.9 TB of Pending space.

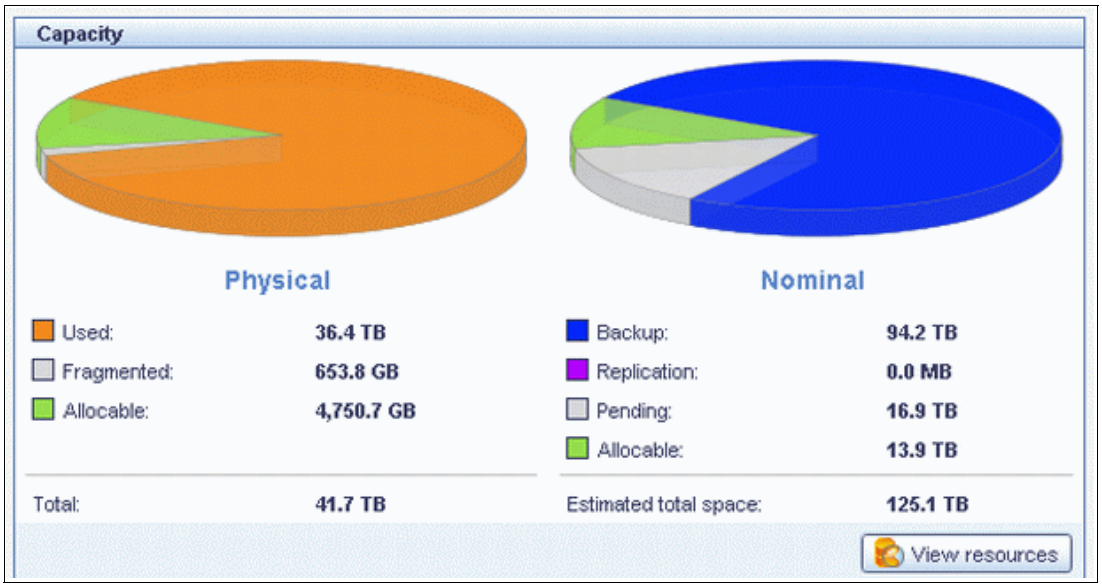

Figure 2-3 ProtecTIER repository with pending nominal space

As a result of this operation, some fragmented space can occur, as shown in the pie chart on the left in Figure 2-3. Further ProtecTIER internal housekeeping eliminates that space. The newly reworked repository is perfectly aligned to your next incoming backup.

2.5.7 ProtecTIER repository usage

From a technical standpoint, there is no problem with having a ProtecTIER repository that is 100% used. After you reach a steady state where daily housekeeping frees up enough space to allow daily backups, a high usage is possible. But in reality, data tends to grow, so sooner or later you might face changed requirements for your ProtecTIER repository size. You should configure an automated message that informs you when the repository usage crosses a specified threshold value. Depending on the time you need to prepare a capacity upgrade of the ProtecTIER back-end storage, values greater than 80% or 90% can be selected to allow ample time for preparation.

To reach the configuration window, click **System** → **Configuration** within the ProtecTIER Manager GUI, as shown in Figure 2-4.

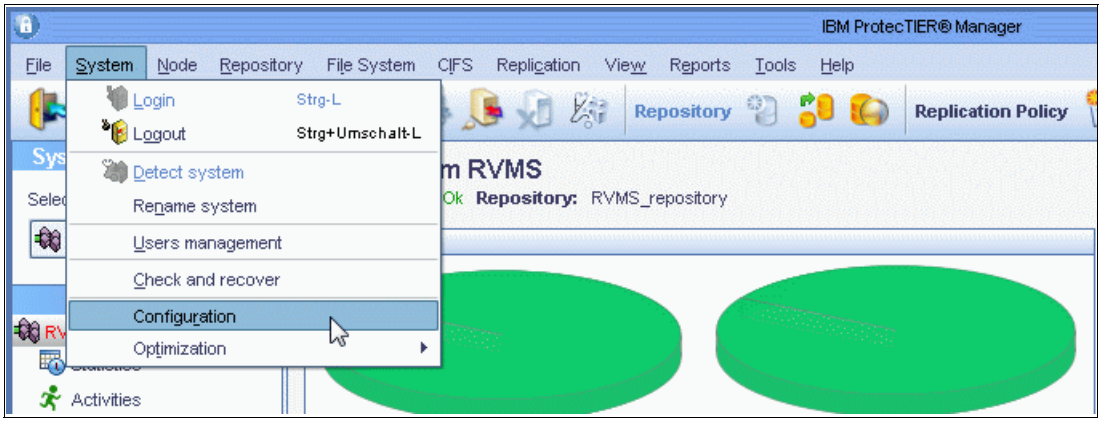

Figure 2-4 ProtecTIER Manager GUI

Click **Physical Space Threshold Alerts** and select the values for information and warning messages, as shown in Figure 2-5.

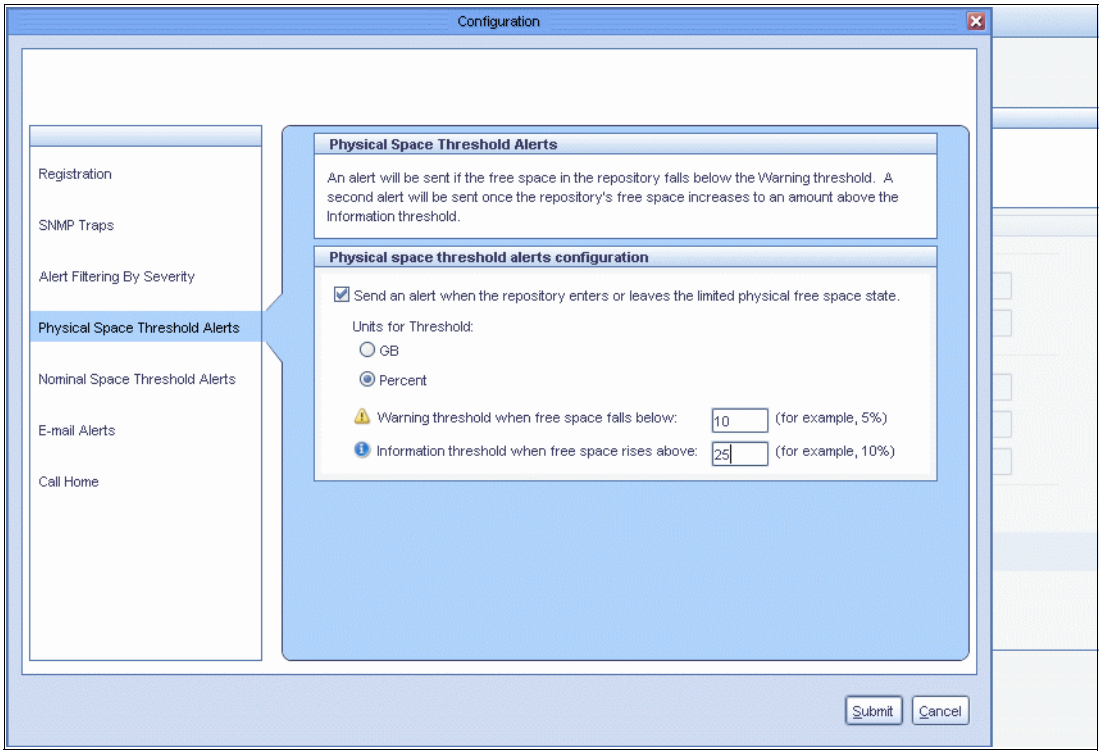

Figure 2-5 Physical Space Threshold Alerts

Important: If you face an out-of-space condition, adding more virtual cartridges to a VTL does not allow you to store more data to your ProtecTIER repository. You must expand your repository by adding more physical disks to the back end to store more data.

2.5.8 Compression

Compression has a negative effect on your deduplication ratio. It effectively shuffles the data sent to the ProtecTIER repository, making pattern matching difficult. As expected, this action affects data matching rates. The ProtecTIER repository compresses the data before it is written to the back-end physical disk. To avoid this negative effect, disable any compression features for the ProtecTIER that are defined in the backup server. Client compression should be disabled as well.

Compression: Compression can hide in unexpected places. Table and Row compression features of databases, IBM Lotus Notes® compaction technology, compressed files, and `*.mpeg` files are all examples of compressed data. Compressed data files are not necessarily easily identified, but still lower your HyperFactor deduplication ratio.

2.5.9 Encryption

Encryption has a negative effect on your deduplication ratio. It makes each piece of data that is sent to the ProtecTIER repository unique, including duplicate data. This situation affects the data matching rates and the factoring performance. Even if the same data is sent each time, it appears differently to the deduplication engine, as shown in Figure 2-6. To avoid this negative effect, disable any encryption features working with data that is sent to ProtecTIER.

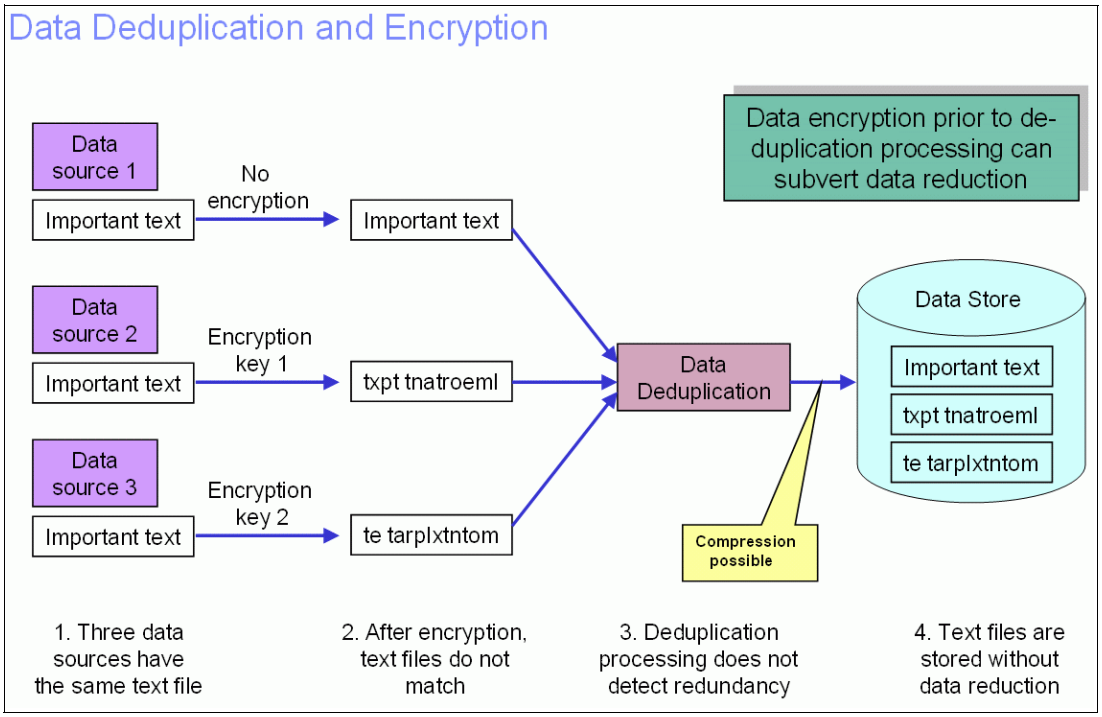

Figure 2-6 Challenges combining encryption with data deduplication

Encryption: If you prefer to run your environment with encryption, consider enabling disk storage-based encryption, for example, IBM System Storage DS8000® encryption. If you prefer to have client side encryption enabled, consider using a ProtecTIER solution without deduplication, as described in 2.5.3, “Data reduction technologies should not be combined” on page 29.

2.5.10 Database logs and other data types with high data change rates

If you have specific data with high change rates, you might decide to point the backup of this data to a target other then ProtecTIER repository to maximize your deduplication ratio within ProtecTIER. For example, database logs are known to have a high change rate, namely 100%. As database logs track all changes within the database, they are never identical. Consider multiple ProtecTIER deployments, some with deduplication enabled and some with deduplication disabled if you prefer to store data on virtual tape libraries (VTL).

Backing up database logs: There is no issue with backing up database logs to a ProtecTIER repository, but be aware that it has an impact on your deduplication ratio.

2.5.11 Multiplexing

Multiplexing has a negative effect on your deduplication ratio. It mixes up the bits of data from many different clients. This situation makes it harder to detect segments of data that exist, so the HyperFactor and compression are greatly reduced. If you want to avoid this situation, disable any multiplexing features in your backup environment. To meet your backup window needs, increase the number of streams and the parallelism of the backup operation.

2.5.12 Tape block size

A large tape block size positively affects your deduplication ratio. To optimize the backup server, set the block size for data that is sent to the (virtual) tape drives to be at least 256 KB. This situation positively affects your HyperFactor deduplication ratio.

2.5.13 File size

Many small files (less than 32 KB in size) have a negative effect on your deduplication ratio. They do not factor well, although the built-in compression might reduce their stored size. If you have a special application that generates many of these files, they are probably not good deduplication candidates.

2.6 Data types

Deduplication is primarily influenced by the type of data you have. Depending on whether the data is structured to a high degree or unstructured, and possibly already compressed, deduplication yields a higher or lower ratio. For more information about data types, see Chapter 20, “Application considerations and data types” on page 315.

2.6.1 Candidates for a high factoring ratio

Potential candidates for a high deduplication ratio are all kinds of structured data. For example, databases are perfect candidates for deduplication, as is email. Most applications that deal with structured data, such as databases and email, offer some compression to reduce the amount of storage the application needs. Because these types of data are good candidates for data reduction in general, many application vendors already have implemented some compression, compaction, or defragmentation. Turning off these application internal data reduction technologies or ensuring that they do not affect the backup data stream allows for high deduplication ratios.

For an example of effectively using DB2 compression with a ProtecTIER repository, see 20.4.1, “Combining DB2 compression and ProtecTIER deduplication” on page 328.

2.6.2 Candidates for a low factoring ratio

Data types that are unstructured have a negative impact on the achievable data deduplication ratio. Image data is an example of this type of data. Some image formats are `*.jpg`, `*.exif`, `*tiff`, or `*gif`. All of them come with compression that shuffles the data and reduces the achievable deduplication ratio. This situation is also true for video formats such as `*.mpg`, `*.mp4`, `*.3gp`, `*.flv`, or `*.asf`. All of these data types are also compressed, which affects your deduplication ratio in a negative way.

The same situation generally applies to voice or audio data. Formats such as `*.mp3`, `*.aac`, `*.ogg`, `*.wma`, or `*.m4a` are also compressed. Backing up image files, video files, or audio files to a ProtecTIER repository results in a combination of data reduction technologies. This situation produces low deduplication ratios, as already reduced data cannot be reduced again (for more information, see 2.5.3, “Data reduction technologies should not be combined” on page 29).

All the mentioned file types include compression. This compression does not work well with data deduplication. For the same reason, archives are also not good deduplication candidates because most archives are already compressed. File types such as `*.zip` (Phil Katz zip, such as **`pkzip`** and **`pkunzip`**), `*.gz` (GNU zip, such as **`gzip`** and **`gzip -d`**), `*.rar`, or `*.tgz` all use a compression algorithm.

Note: If you back up a compressed or encrypted file multiple times without changing it between the backup cycles, you have a high deduplication ratio. Multiple full backups of identical data yield high deduplication ratios. Changing only one single file within a huge compressed archive affects the whole data structure of that archive, which does not result in good deduplication.

3

Virtual Tape Library guidelines

This chapter describes general best practices for optimizing the ProtecTIER Virtual Tape Library (VTL). It describes VTL concepts, methods, and system components.

This chapter also describes the procedure and best practices for creating and configuring virtual libraries for optimal performance.

This chapter describes the following topics:

- ProtecTIER Virtual Tape Library introduction
- General best practices for the Virtual Tape Library
- Setting up the virtual library and cartridges

3.1 ProtecTIER Virtual Tape Library introduction

The ProtecTIER VTL service emulates traditional tape libraries. By emulating tape libraries, you can use ProtecTIER VTL to switch to disk backup without having to replace your entire backup environment. Your existing backup application can access virtual robots to move virtual cartridges between virtual slots and drives. The backup application perceives that the data is being stored on cartridges while the ProtecTIER product stores data on a deduplicated disk repository.

3.2 General best practices for the Virtual Tape Library

A ProtecTIER Virtual Tape Library (VTL) can be optimized by the following these simple rules:

- Create more slots than are needed for future growth in the number or cartridges. Adding cartridges while there are free slots available is an online procedure, while adding more slots to the library is an offline procedure.
- Create cartridges with a fixed cartridge size. You can accomplish this task by selecting **Max Cartridge Growth** in the tape cartridges creation menu.
- Create small cartridge sizes to allow as much parallelism as possible. The VTL emulates a physical tape library. So virtual cartridges behave in the same manner as physical cartridges with sequential access. If you have many drives and large tape sizes, you might encounter a situation where the backup or restore is waiting for a large sized tape that is being used by another backup session. If you have small tapes with a capacity of 100 GB, for example, you decrease the probability of wait times for backup and restore operations.
- If you have a physical tape library that is connected to the same backup application that is using ProtecTIER, make sure that the libraries use different barcode ranges. This action avoids tape access conflicts and facilitates the identification of which tape is physical and which tape is virtual.
- Create only the number of cartridges that your repository can handle, maybe even fewer to control the repository allocation of different VTLs. You can estimate the size of a repository by multiplying the real size of the repository by the HyperFactor ratio. Then, divide it by the tape size and determine the optimized number of tapes.

Important: Be careful not to overestimate the repository size. Wait until the backup application sends some data to provide a better view of the real deduplication ratio.

3.3 Setting up the virtual library and cartridges

You can use the ProtecTIER Manager to create virtual tape libraries where the backup application stores your data. These libraries and their components are part of the virtual tape service.

3.3.1 Creating libraries

Tip: Use the **Scan** button of the Fibre Channel port attributes pane to verify that the ports of the ProtecTIER system to which the virtual devices of the library are assigned are connected to the correct host. For more information, see Chapter 11, “Monitoring and reporting”, in *IBM System Storage TS7600 with ProtecTIER Version 3.1*, SG24-7968.

To create a library on a ProtecTIER system, complete the following steps:

1. Log on to the system on which you want to add a library.
2. From the menu bar, click **VT** → **VT Library** → **Create new library**. The Create new library wizard Welcome window opens (Figure 3-1).

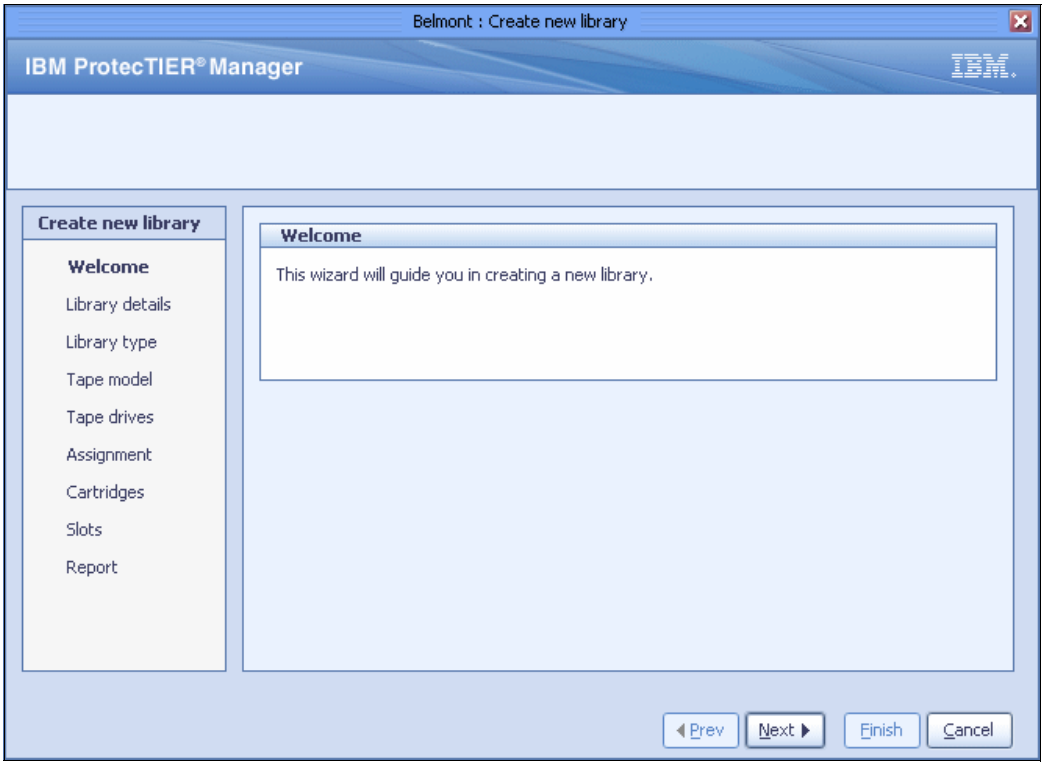

Figure 3-1 Create new library - Welcome window

3. Click **Next**. The Library details window opens (Figure 3-2).

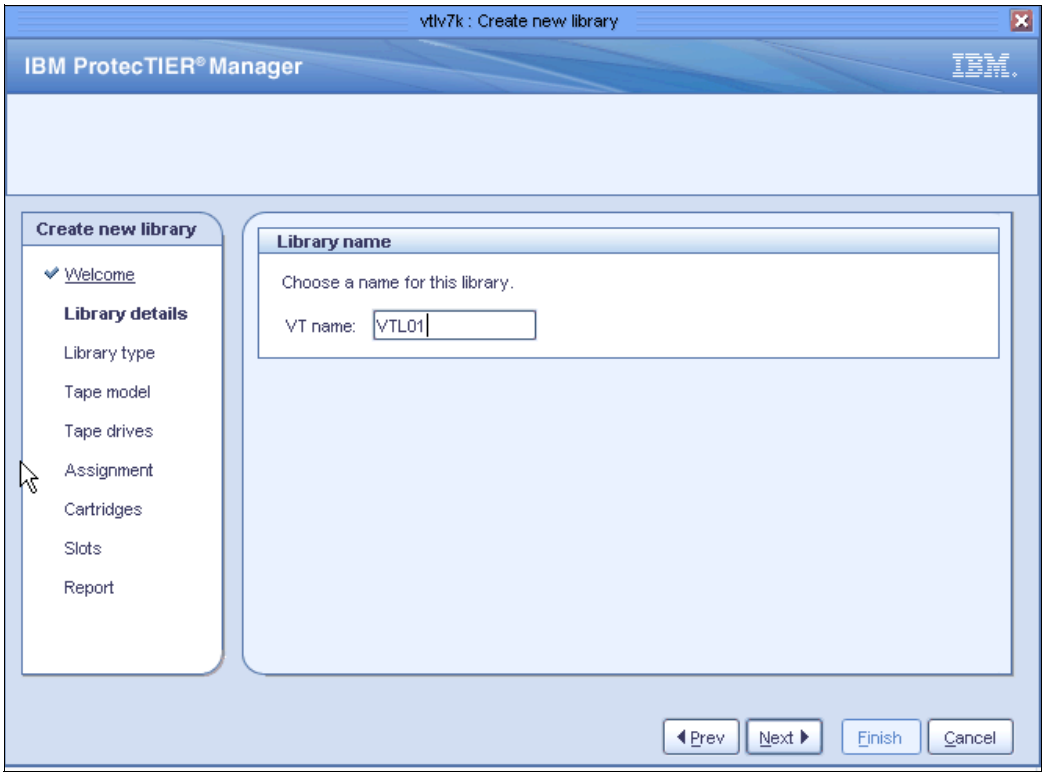

Figure 3-2 Create new library - Library details window

4. In the ProtecTIER VT name field, enter a name for the library.

5. Click **Next**. The Library Type window opens (Figure 3-3).

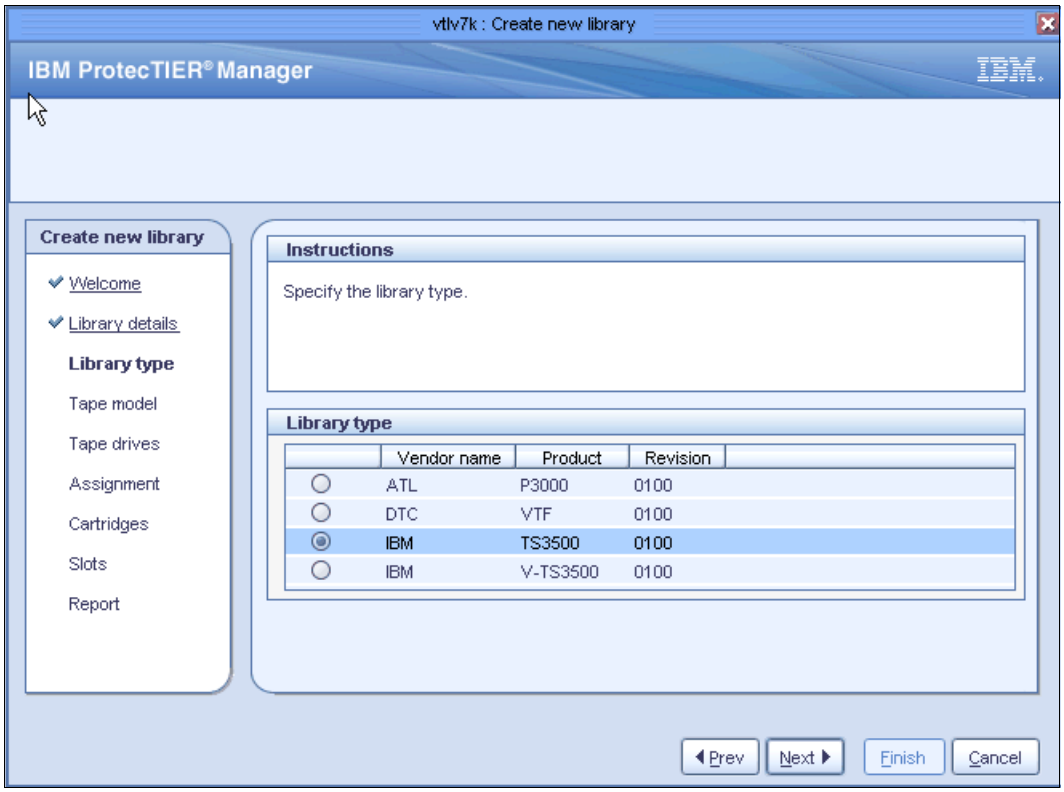

Figure 3-3 Create new library - Library type window

By default, IBM TS3500 is selected. The functionality of the TS3500 and the IBM V-TS3500 is the same.

Attention: Verify that the backup application that you are using supports the type of library model that you select:

- If you are using Symantec NetBackup software, you should select **V-TS3500**.
- If you are using Tivoli Storage Manager software, or any other backup application, you should select **TS3500**.

6. Click **Next**. The Tape model window opens (Figure 3-4).

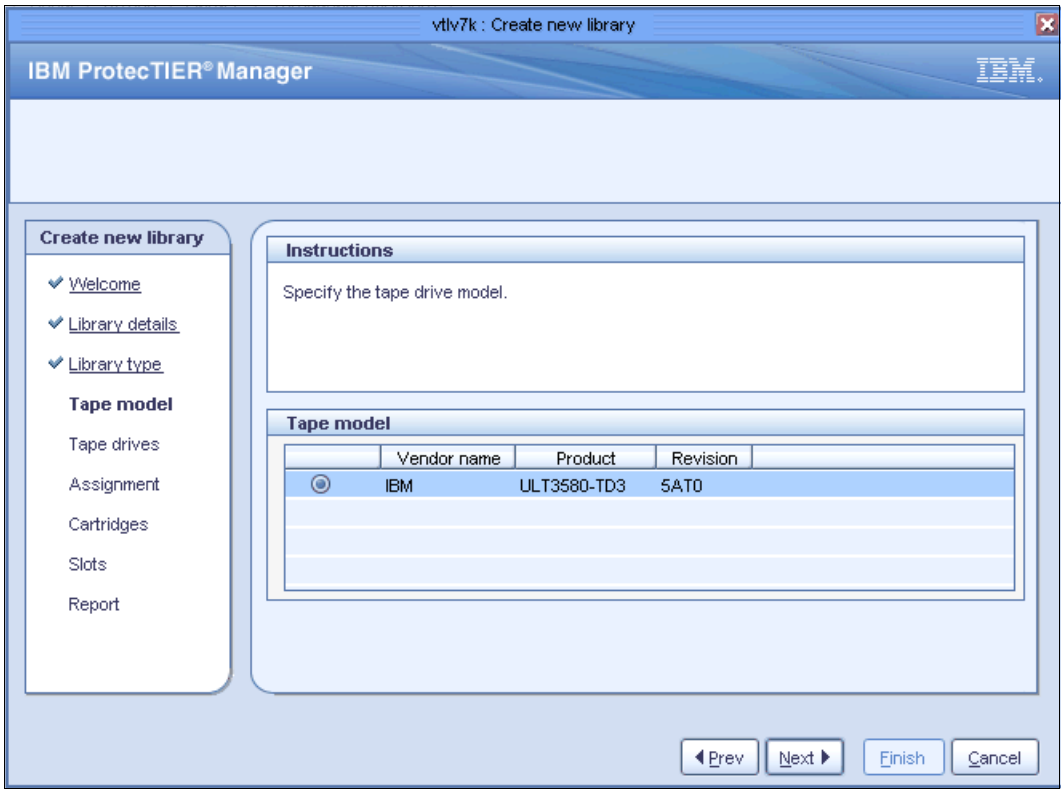

Figure 3-4 Create new library - Tape model window

Select the tape drive model that you want to use for your virtual library.

The type of tape drive model that is displayed depends upon the type of library that was selected in the previous step.

7. Click **Next**. The Tape drives window opens (Figure 3-5).

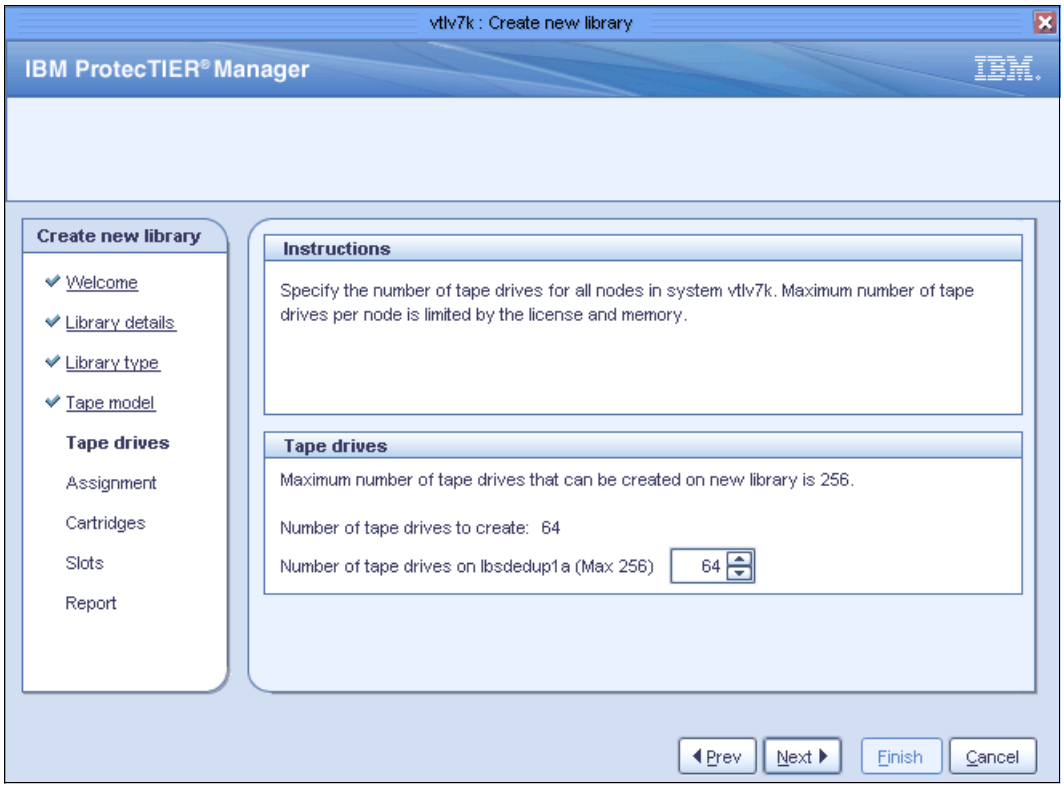

Figure 3-5 Create new library - Tape drives window

In the Number of tape drives field, enter the number of tape drives to assign to the library.

The maximum number of virtual tape drives depends on the model of ProtecTIER. A ProtecTIER gateway (TS7650G) supports a maximum of 512 virtual tape drives on a dual node cluster.

In this example, we create 64 drives on the libsdedup1 node.

Number of drives: Check with your backup application administrator to verify the number of drives and cartridges that is supported by your application. The value of the maximum number of tape drives per TS7650 ProtecTIER node is 256.

TS7620 Appliance Express supports up to 64 virtual drives.

8. Click **Next**. The Port assignment window opens (Figure 3-6).

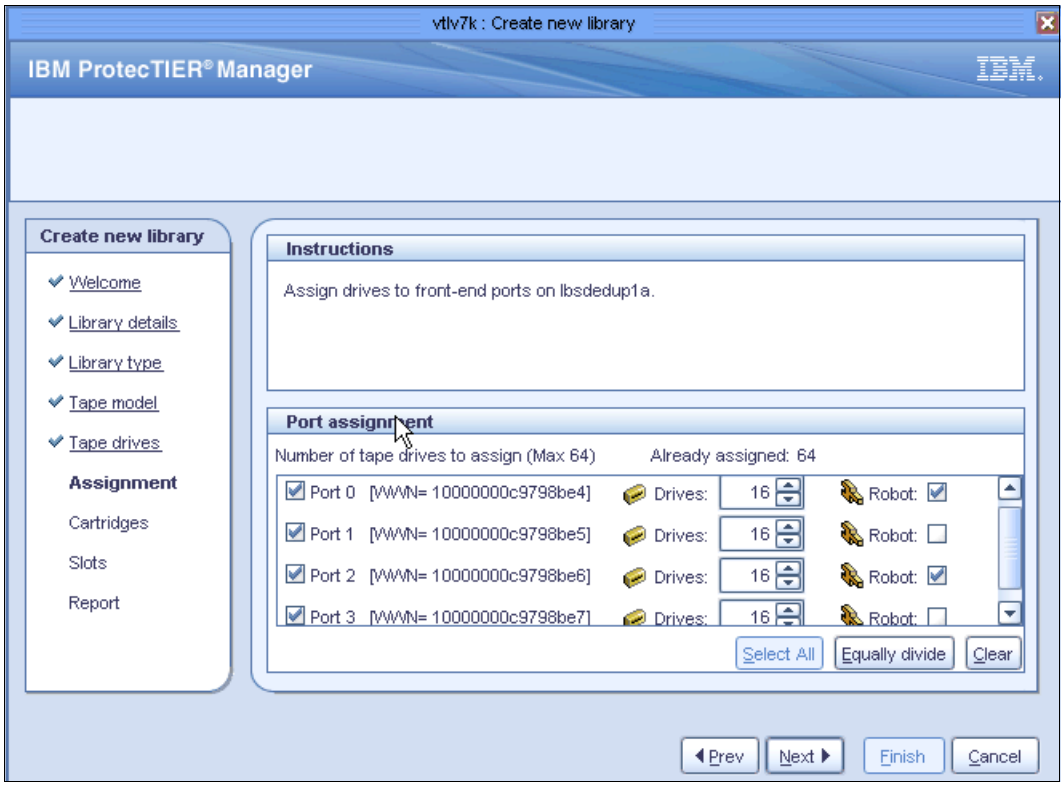

Figure 3-6 Create new library - Port assignment window

Select or clear the check boxes next to each port to define which of the node's ports are assigned virtual devices. By default, the IBM TS3500 supports Control Path Failover (CPF), so all the robots are selected and enabled. If you chose a library model other than IBM, the robots are not checked and only one robot must be chosen.

In the Drives fields corresponding to each selected port, select the number of virtual tape drives that are assigned to each port.

Optionally, click **Select All** to automatically select both ports. Click **Equally divide** to evenly divide the number of drives between the ports. Check the **Robot** check box if you want the library virtual robot to be accessible through this port.

9. Click **Next**. If a second node exists in your cluster, the Assignment (2) window opens. Complete the same steps as you did in Figure 3-6 for your first cluster.

High availability: For high availability purposes, the IBM System Storage TS7600 with ProtecTIER supports the assignment of the virtual robot to multiple ports.

The backup application can access only the virtual robot through the specific node and port to which the robot is assigned. Verify that the port is connected to the appropriate host in your backup environment by clicking **Scan** in the Port attributes pane.

10. Click **Next**. The Cartridges window opens (Figure 3-7). In the **No. of cartridges** field, enter the number of cartridges that you want to have in the library.

    By default, ProtecTIER calculates a virtual cartridge size by dividing all of the free Nominal space in the repository by the number of requested cartridges.

    It is a best practice to create fixed size cartridges. Do this task by selecting the **Max Cartridge Growth** check box and entering a cartridge size in gigabytes.

    Although there is no magic number for cartridge size, the consensus opinion is smaller is better for efficient repository usage. Creating small cartridges can result in many cartridges for the backup application to manage.

    Tip: The recommended guideline for cartridge maximum size is 100, 200, or even 400 GB, depending on the workload. Be aware of what could happen if a customer uses 400 GB cartridges and sets up reclamation to occur when 50% of data on this cartridge is no longer valid. Because of the large size of the cartridge, it could take a long time for this 400 GB cartridge to be reclaimed. While such cartridges with expired data are not reclaimed, the backup application cannot scratch the cartridge to free up the repository space, which means wasted space. 100 GB is a good balance between smart management and wasted space.

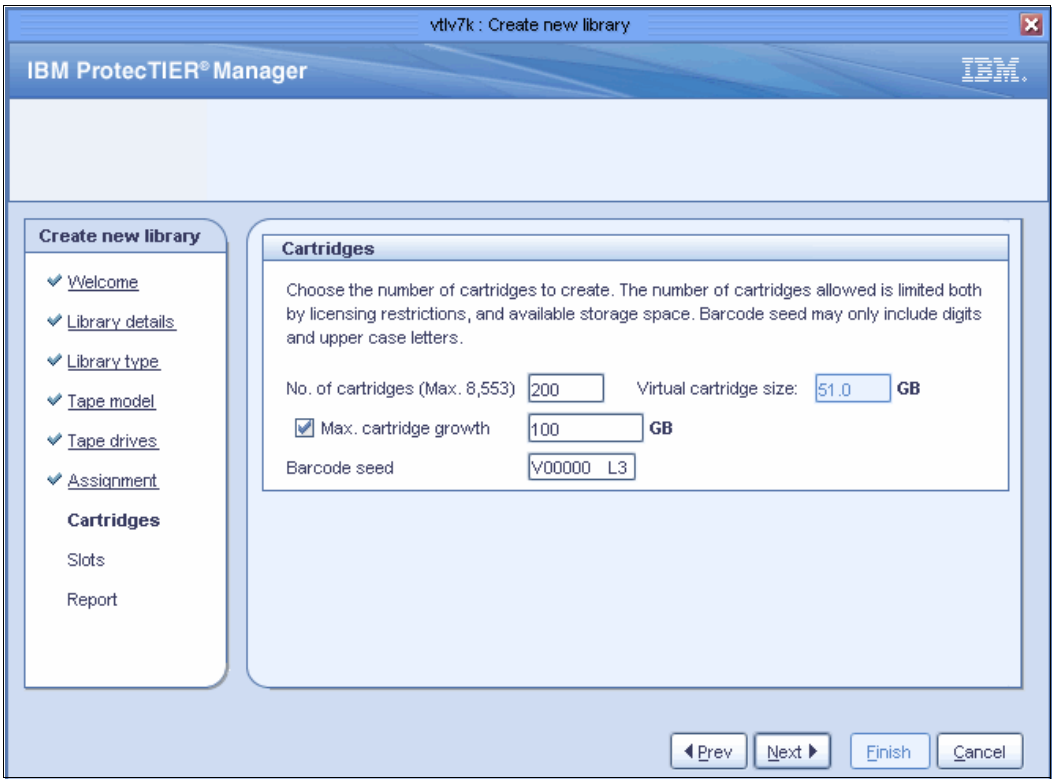

Figure 3-7 Create new library - Cartridges window

11. In the Barcode seed field (Figure 3-7), enter a value for the barcode seed. The barcode seed is the barcode that is assigned to the first cartridge created. Every cartridge that is added after the first cartridge is incrementally assigned a barcode that follows the previous one.

Tips:

- The barcode seed must contain only numbers and capital letters and be only six characters in length (for example, DS0006).
- Do not define the same barcode range that is in use by an existing library. Following this suggestion avoids conflicts and makes administration easier.
- ProtecTIER data replication in a VTL is defined at the cartridge level. Using barcode ranges to group data logically makes replication planning easier.

12. Click **Next**. The Slots window opens (Figure 3-8).

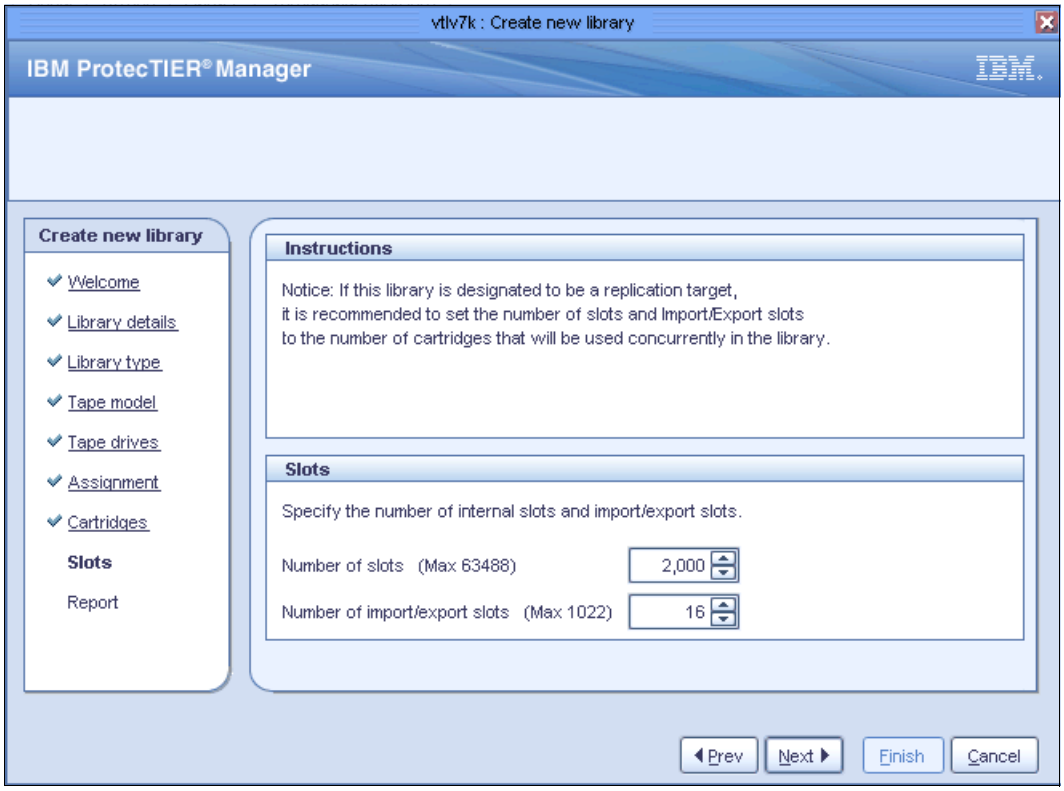

Figure 3-8 Creating new library - Slots window

In the Number of slots field, enter the number of cartridge slots that you want to have in the library.

Tips:

- The number of cartridge slots must be equal to or more than the number of cartridges that you are creating. Create more slots now, if you expect the number of cartridges to increase later.
- Libraries that are attached to IBM i can have a maximum of 4096 positions where media can be stored. The total number of drives, number of convenience IO slots plus the number of media slots plus 1 (for the picker), must not exceed 4096.

In the Number of import/export slots field, enter the number of import/export slots that you want to have in the library. The maximum number of import/export slots that can be defined is 1022 per virtual library.

13. Click **Next**. The Create new library wizard closes and a summary report opens. Click **Finish**. For an example of the summary report output, see Figure 3-9 and Figure 3-10 on page 48.

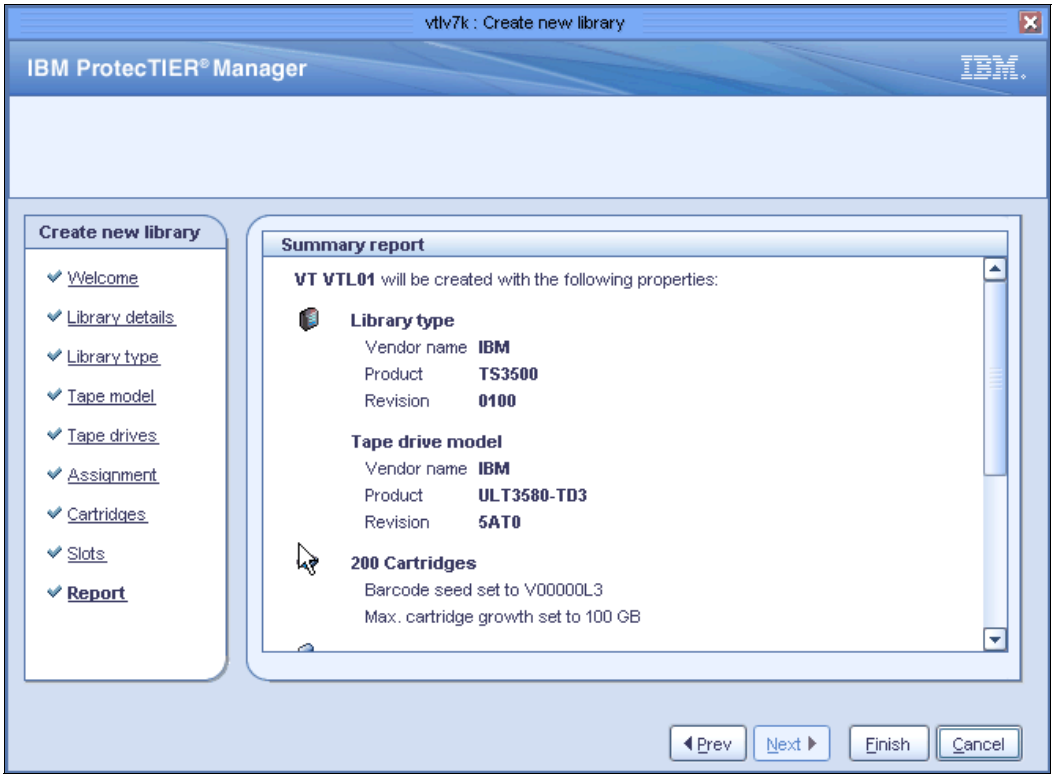

Figure 3-9 Summary report part 1

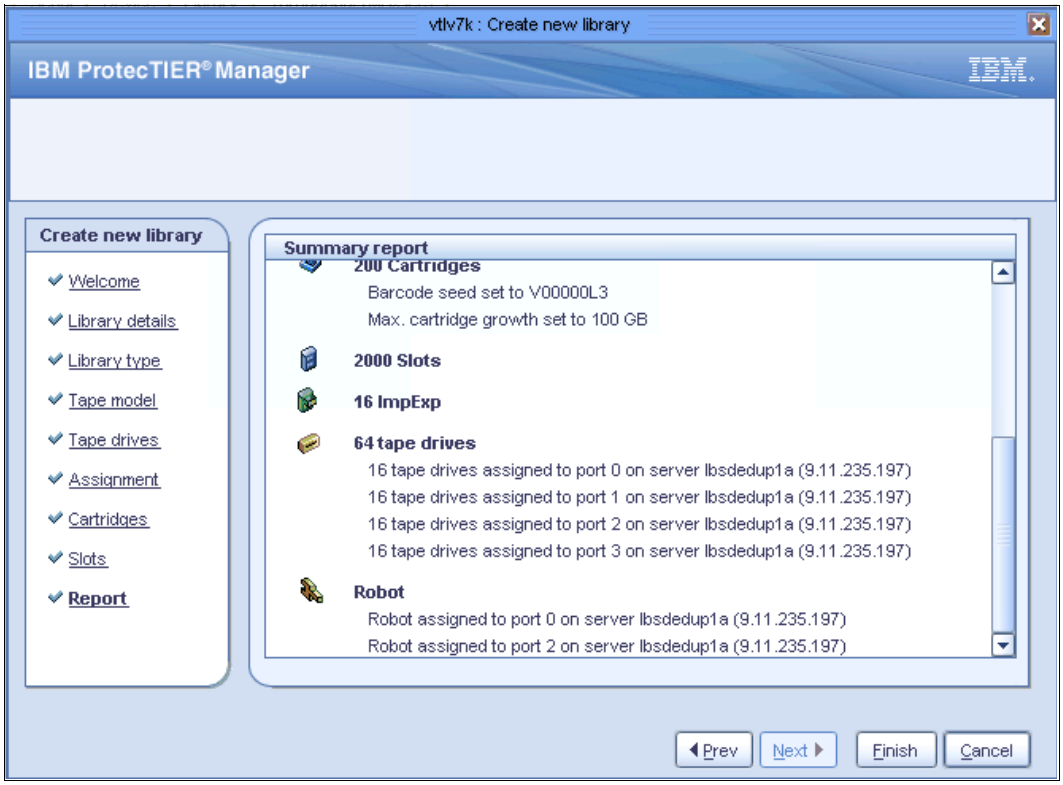

Figure 3-10 Summary report part 2

14. The Confirm operation window opens (Figure 3-11). Click **Yes**. The ProtecTIER system temporarily goes offline to create the library.

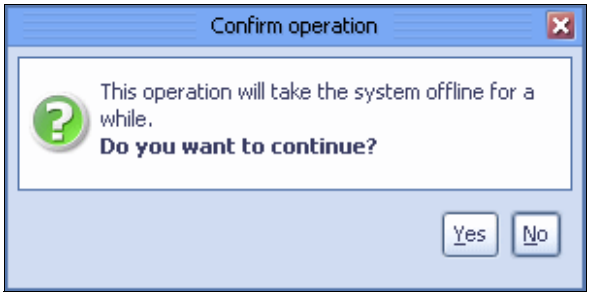

Figure 3-11 Create new library - Confirm operation window

15. The newly created library is displayed in the left Navigation pane of the ProtecTIER Manager GUI. The right Context pane shows the details of the virtual library that you just created (Figure 3-12).

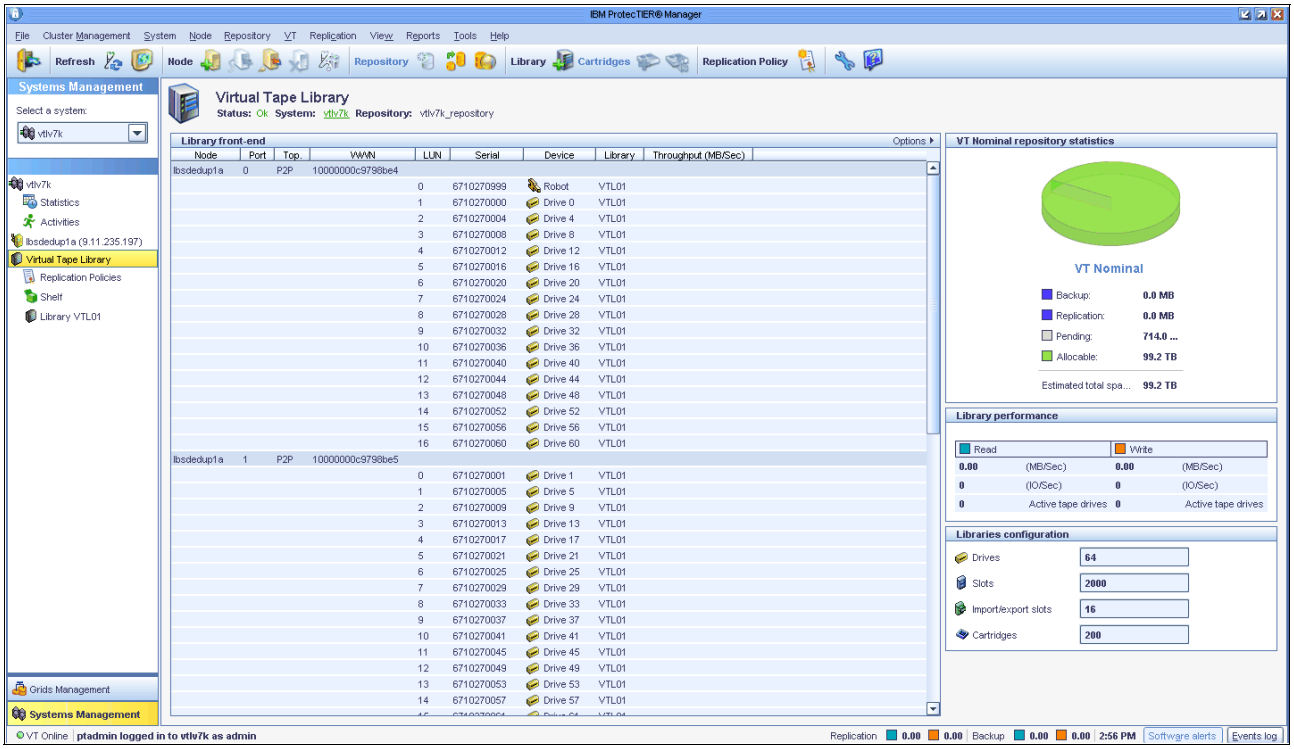

Figure 3-12 Newly created Virtual Tape Library

Selecting the new library (**Library VTL01**) in the Navigation shows the library's details. The General tab of the logical library is shown in Figure 3-13.

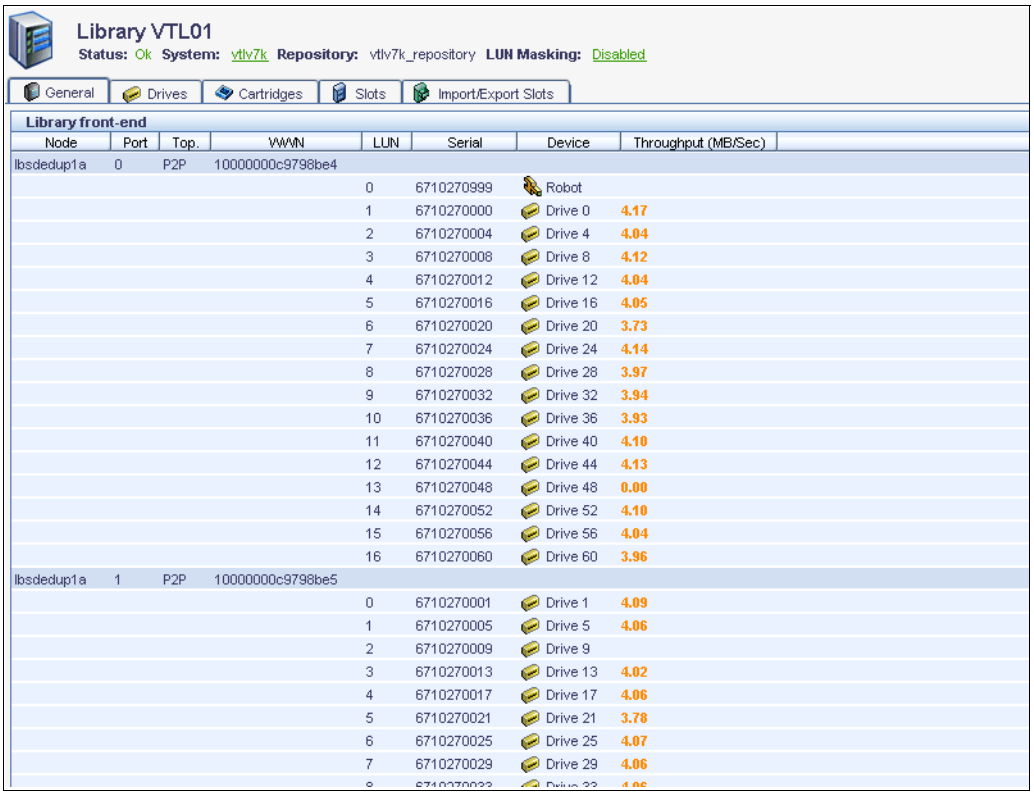

Figure 3-13 Logical library details

In this window, you can click the tabs for details about each component, and also take the corresponding actions for the component, as though it is a physical library. Here are the components:

- Drives: Shows mounted tapes and the throughput of each logical tape drive. From this tab, you can, for example, reset a drive.
- Cartridges*:* Shows the cartridge inventory, including the tape size, its usage, and whether it is full or not. From the Cartridge tab, you can also see the cartridge replication synchronization status. For more information, see 22.8.3, "Gauging the replication completion status" on page 424.
- Slots: Represents the logical slots where the virtual tapes are stored. For example, if a virtual tape is mounted, the slot where the tape was located is empty while the tape is mounted, as though it is a real library slot.

  Slots tab: From the Slots tab, you can move tapes manually to a drive, to another slot, to the shelf, or to Import/Export slots. Although it is possible to do a manual activity, you should move tapes by using the backup application to avoid mismatched library inventory within the backup application and the virtual tape library.

- Import/Export Slots: The logical representation of the I/O station or bulk, where tapes are inserted or removed from the library. The Import/Export slots are used during the replication process to send the virtual tape to the secondary site.

To the right of the General tab, you have a representation of the drives, online throughput statistics, and the dimensions of the VTL (Figure 3-14).

Backup traffic: When a backup application is using the virtual tape drives, it shows a small blue or orange square for each virtual tape drive. A blue square means read activity from the drive, and an orange square means write activity. Furthermore, it shows the overall performance of the system for both backup/restore and replication activity.

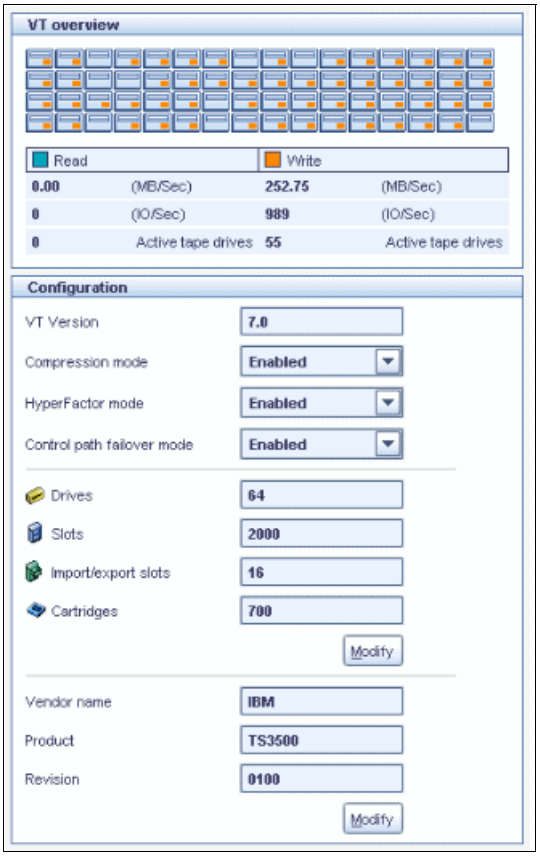

Figure 3-14 Logical tape library details and dimensions

4

ProtecTIER File System Interface: General introduction

This chapter provides network configuration considerations for setting up your network with ProtecTIER FSI. In addition, this chapter describes the steps and guidelines to configure ProtecTIER components, such as application IP interfaces and your ProtecTIER system. This chapter describes how to create FSI NFS and FSI CIFS shares on ProtecTIER and how to connect your backup host to the FSI shares.

This chapter describes the following topics:

- ProtecTIER NFS authentication and security management
- File System Interface guidelines for NFS
- File System Interface guidelines for CIFS

The ProtecTIER File System Interface (FSI) presents ProtecTIER as a network-attached storage backup and recovery target that can use the HyperFactor algorithm and ProtecTIER native replication bandwidth reduction techniques for storing and replicating deduplicated data. The ProtecTIER FSI interface is intended to be used for backup and restore of data sets by using a backup application.

As of ProtecTIER Version 3.2, support is provided for Windows based servers through the CIFS protocol, and with ProtecTIER Version 3.3, support is provided for UNIX clients through the NFS protocol. ProtecTIER emulates a UNIX or Windows file system behavior and presents a virtualized hierarchy of file systems, directories, and files to UNIX NFS or Windows CIFS clients. These clients can perform file system operations on the emulated file system content.

You can use the ProtecTIER FSI to create multiple user file systems in a single ProtecTIER repository. When you create a user file system, the maximum size of the user file system is dynamically calculated by determining the total free nominal space in the repository and comparing it to the overall maximum user file system size of 256 TB. The size of all file systems shrinks proportionally if the deduplication ratio goes lower than expected. If the deduplication ratio goes beyond the expected size, it is possible to extend the file system size up to the 256 TB limit by using the ProtecTIER Manager.

The FSI interface of ProtecTIER for UNIX and Windows clients is supported on a single node. No dual node cluster support is available currently. However, a single node can serve multiple CIFS and NFS exports in the same repository. Exporting a single FSI share through CIFS or NFS protocol is mutually exclusive. To change the export type from NFS to CIFS, you must delete the NFS share definition before you export it through CIFS and vice versa; disabling the share definition alone is not sufficient.

Important: ProtecTIER FSI support is intended for storing backup images that are produced by backup applications and not for primary storage deduplication. ProtecTIER performs best when sequential streams are delivered to ProtecTIER instead of random I/O.

4.1 ProtecTIER FSI network overview

This section provides an overview for the ProtecTIER FSI network configuration.

4.1.1 ProtecTIER network

ProtecTIER servers have several physical network ports. The number of ports varies based on the ProtecTIER model. Ports are used for management, replication, or file system related operations from the hosts. Each port is assigned to one of these uses. This configuration is achieved by assigning the physical ports to a virtual interface on the ProtecTIER server. The set of virtual interfaces on the ProtecTIER product includes External, Replication1, Replication2, FSI1, FSI2, on to FSI_N. Each one of the virtual interfaces can be assigned to one or more physical network ports.

The default setup of the ProtecTIER product assigns all of the FSI physical ports to a single virtual interface by using round robin load balancing. This setup can be changed as needed. If there is more than one physical port that is assigned to a virtual interface, it is important to configure the bonding methodology in this interface to align with the network environment and fulfill the wanted behavior in terms of performance and redundancy. For more information about the bonding methods available with the ProtecTIER product, see Chapter 5, “Networking essentials” on page 79.

4.1.2 Network configuration considerations

This section describes network configuration considerations and best practices for FSI. The following guidelines are valid for CIFS and NFS configurations. As ProtecTIER IP replication with FSI is realized on a file share level, you should create a dedicated file share for each backup server that you use with ProtecTIER FSI.

- Make sure that the backup application runs in the context of a user that has read and write permissions on the FSI share.
- You must have *at least* two different network subnets to separate the ProtecTIER management IP interface from the ProtecTIER file system (application) interfaces.
- For FSI workloads, you must have sufficient TCP/IP infrastructure for the incoming and outgoing traffic of backup servers. Ensure that you do not suffer from network bandwidth congestion.

- If bonding of network adapters is implemented, it must be implemented on all involved network devices, that is, the ProtecTIER server, the backup server, and the network switches. Simply enabling bonding on the ProtecTIER server might not be enough to achieve the best results.

4.1.3 Connecting a ProtecTIER server to the network

To better understand the requirements that arise from using the FSI model of ProtecTIER in your network environment, look at the potential connections that you deal with during the initial deployment. As shown in Figure 4-1, we use in our example the connection (labeled “1”) to attach the ProtecTIER server to the customer network. Through this connection, you use the ProtecTIER Manager GUI and connect to the ProtecTIER server for management and configuration purposes.

For the ProtecTIER IP replication feature, there are two Ethernet connections (labeled “21” and “22”). As you can see, the interfaces that are used for ProtecTIER IP replication are distributed across different physical hardware adapters for high availability reasons. By default, the replication workload is balanced across both ports.

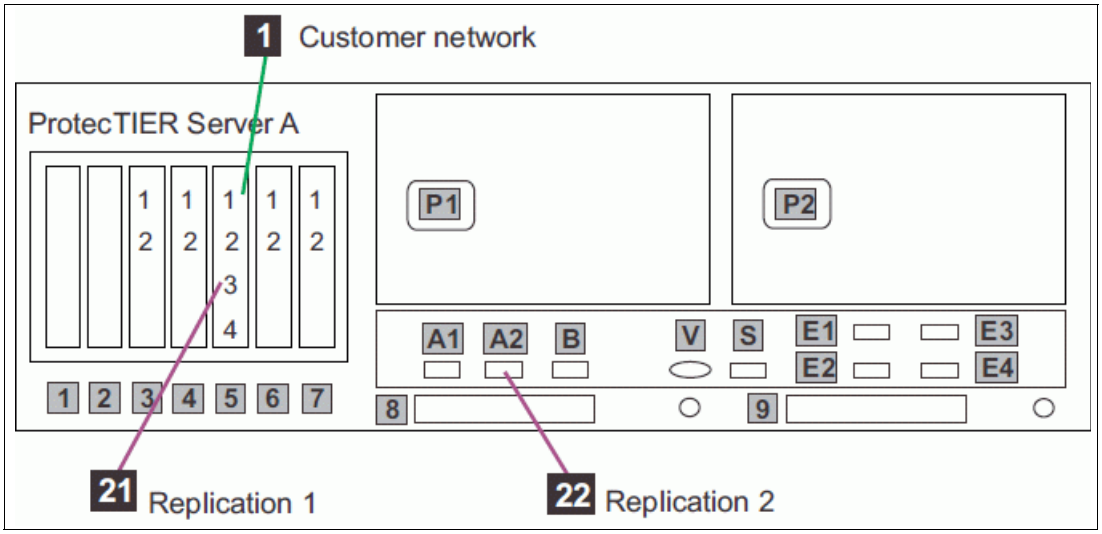

Figure 4-1 ProtecTIER network interfaces for customer and replication network

To use the FSI feature, you must prepare at least one (and as many as eight) dedicated subnets for the backup server traffic to the ProtecTIER server. The data that is transferred from the backup server to the FSI interface must not use the customer network IP interface or the replication interfaces.

The interfaces that are labeled “13”, “14”, “15”, and “16” in Figure 4-2 are available to use the FSI traffic between the backup servers and the ProtecTIER server. Figure 4-2 shows the 10 GB interface model of the ProtecTIER server.

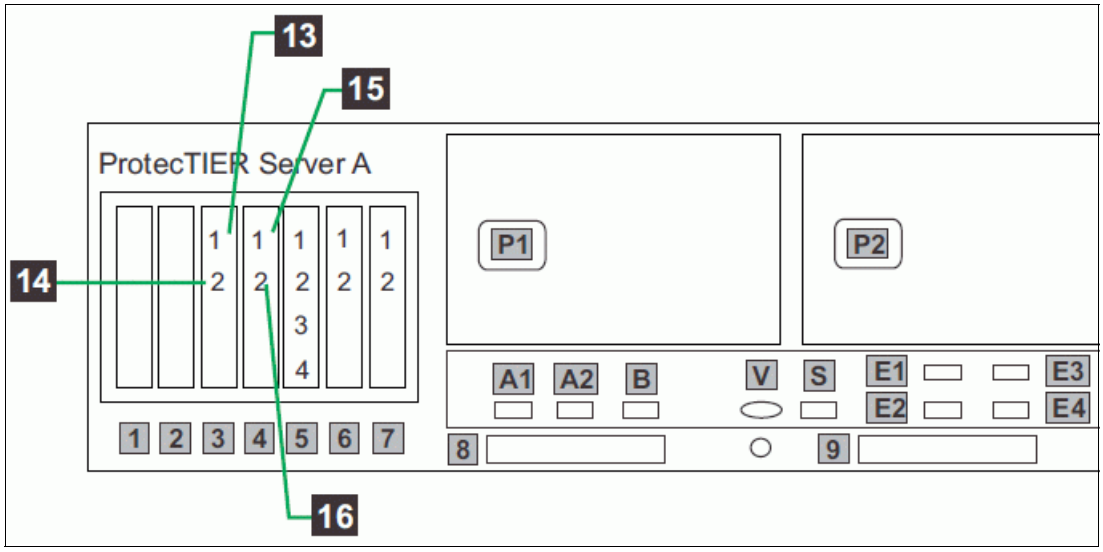

Figure 4-2 ProtecTIER network interfaces for FSI traffic from backup servers

Now that you see all of the important interfaces for potential network traffic, you can review the configuration through the ProtecTIER Manager GUI. To configure all networking-related aspects of the ProtecTIER server, open the ProtecTIER Manager GUI and click **Node** → **Network configuration** (Figure 4-3).

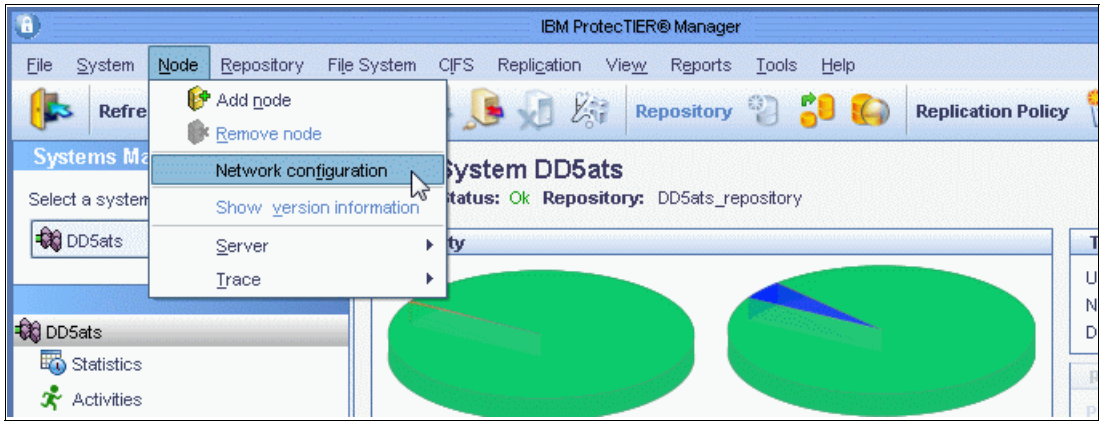

Figure 4-3 ProtecTIER Manager GUIs network configuration window

The Network Configuration menu provides the options to change the networking parameters (Figure 4-4).

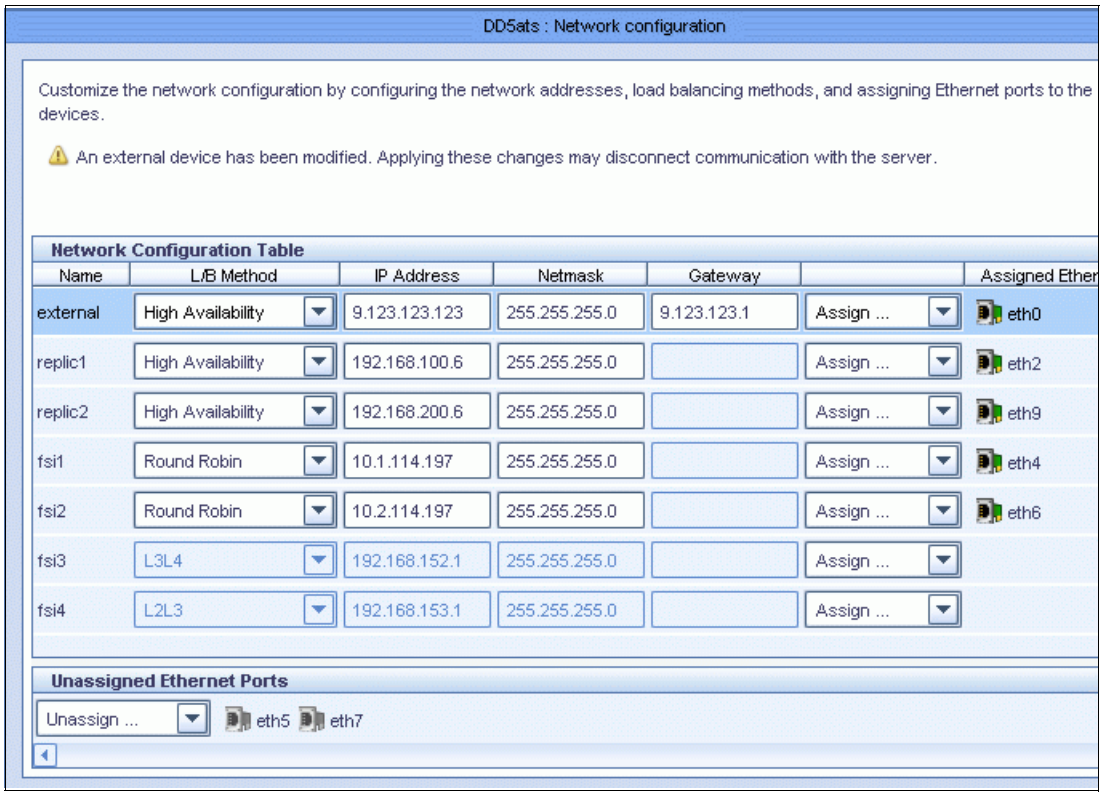

Figure 4-4 ProtecTIER network configuration window

When you perform ProtecTIER network configuration, you can assign a physical device to a ProtecTIER virtual device for all interfaces, even if the virtual interface contains only one physical device.

There are multiple ways to set up your networking to ensure that you have a high availability configuration and that you distribute the load across all available resources. The default setup is a single virtual interface *fsi1*, which consists of all four physical 10 Gb ports (Table 4-1).

Table 4-1 DD5 10 Gb Ethernet default port assignments

Network types	**Virtual interfaces**			**Assigned physical ports**			
	Network IP	**LB**	**Subnet**	**Name**	**Speed**	**Slot**	**Port**
External	External IP	RR	1	Eth0	1 Gb	5	Top
Application	fsi1	RR	2	Eth4	10 Gb	3	Top
				Eth5	10 Gb	3	Bottom
				Eth6	10 Gb	4	Top
				Eth7	10 Gb	4	Bottom
Replication	replic1	N/A	3	Eth2	1 Gb	5	Bottom-Middle
	replic2	N/A	4	Eth9	1 Gb	E2	Onboard

Separation of networks: Again, you must separate your external customer management network from your backup OST or FSI network. It is important to configure the ProtecTIER network so that each virtual interface (IP) is on a different network and preferably a different VLAN in a multitier network infrastructure.

If you use the configuration that is shown in Table 4-1 on page 57, all of your backup servers connect to the IP of the ProtecTIER application virtual interface fsi1. The default load-balancing (LB) method of round robin (RR) mode 1 works without special network infrastructure hardware requirements. This LB mode allows, depending on your network infrastructure, a unidirectional bandwidth increase. This configuration means that from the perspective of a single data stream that flows outbound from a ProtecTIER server, you potentially could benefit from up to 40 Gb of bandwidth, which is essentially the combined throughout of all four aggregated interfaces. It also means that restoring data from your ProtecTIER server to your backup server could be fast. Backing up your data creates a data stream that is directed towards the ProtecTIER server. Single data streams directed towards a ProtecTIER server do not benefit from the potential bandwidth increase when you use the round robin LB method in this example.

To fully use the ProtecTIER server resources in this configuration, you must use multiple backup servers that back up to their respective file systems on the ProtecTIER server.

To further optimize the potential throughput of single backup server environments, you must understand the different link aggregation methods that can be used for load balancing and increasing throughput, as described in Table 5-1 on page 84.

4.2 Configuring components for ProtecTIER FSI

This section describes the steps that are required to configure your ProtecTIER application IP interfaces. It also describes the steps to configure your ProtecTIER server, to create CIFS and NFS shares on your ProtecTIER server, and to connect to these shares. The details of best practices and configuring applications (Tivoli Storage Manager, EMC NetWorker, Symantec NetBackup, CommVault, and Symantec BackupExec) for ProtecTIER FSI are described in Part 3, “Backup management, VTL, OST, and FSI best practices” on page 209.

4.2.1 Configuring a network

This section describes the steps to configure application IP interfaces and static routes on a ProtecTIER server for a host.

Configuring the application IP interfaces

To configure the replication and FSI virtual interfaces, use the ProtecTIER Service menu or the Network configuration wizard of the ProtecTIER Manager GUI (Figure 4-4 on page 57).

Configuring the static routes for FSI on ProtecTIER and host

When you configure multiple subnets for ProtecTIER FSI addresses, it might be necessary to configure static routes on both the ProtecTIER servers and the backup hosts. Complete the following steps on the ProtecTIER server. On the backup hosts, configure static routes to direct IP traffic to the FSI addresses of the ProtecTIER.

1. Log on to the ProtecTIER server command line by using Secure Shell (SSH) with user name “ptconfig” and password “ptconfig”. To start the ProtecTIER Service Menu, run the following command:

   ```
   [root@BUPKIS ~]# menu
   ```

2. The ProtecTIER Service Menu screen opens (Figure 4-5). Select option 1) ProtecTIER Configuration.

```
+-----------------------------------------
| ProtecTIER Service Menu running on DD5_H
+-----------------------------------------
|   1) ProtecTIER Configuration (...)
|   2) Manage ProtecTIER services (...)
|   3) Health Monitoring (...)
|   4) Problem Alerting (...)
|   5) Version Information (...)
|   6) Generate a service report
|   7) Generate a system view
|   8) Update ProtecTIER code
|
|   E) Exit
+-----------------------------------------
Your choice? 1
```

Figure 4-5 ProtecTIER Service Menu

3. The ProtecTIER Configuration screen opens (Figure 4-6). Select option 8) IP Network configuration.

```
+-------------------------------------------------------------------------------+
| ProtecTIER Service Menu running on thekla                                     |
| ProtecTIER Configuration (...)                                                |
+-------------------------------------------------------------------------------+
|   1) Configure ProtecTIER node                                                |
|   2) Recover Configuration for a replaced server                              |
|   3) Configure machine serial number for a replaced server                    |
|   4) Configure RAS                                                            |
|   5) Update Time, Date, Timezone and Timeserver(s)                            |
|   6) Scan storage interconnections                                            |
|   7) File Systems Management (...)                                            |
|   8) Configure replication (...)                                              |
|   9) IP Network configuration (...)                                           |
|  10) Update Firmware                                                          |
|  11) Update the System's name                                                 |
|  12) Validate configuration                                                   |
|  13) Single Node - code upgrade                                               |
|                                                                               |
|   B) Back                                                                     |
|   E) Exit                                                                     |
+-------------------------------------------------------------------------------+
Your choice? 9
```

Figure 4-6 ProtecTIER Service Menu - ProtecTIER configuration IP network configuration

4. The ProtecTIER IP Network configuration screen opens (Figure 4-7). Select option 3) Configure Static Routes.

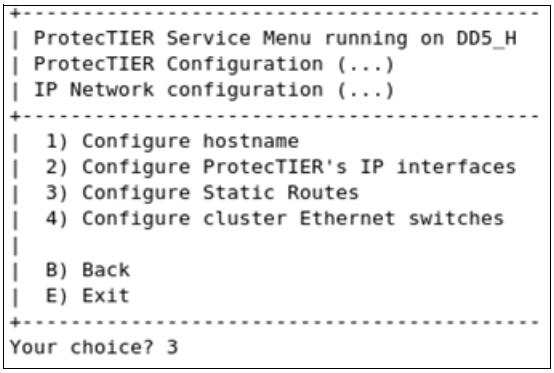

Figure 4-7 ProtecTIER Service Menu - IP Network configuration

5. The ProtecTIER Available Options screen opens, where you can configure static routes (Figure 4-8). Select option (a)dd a new record.

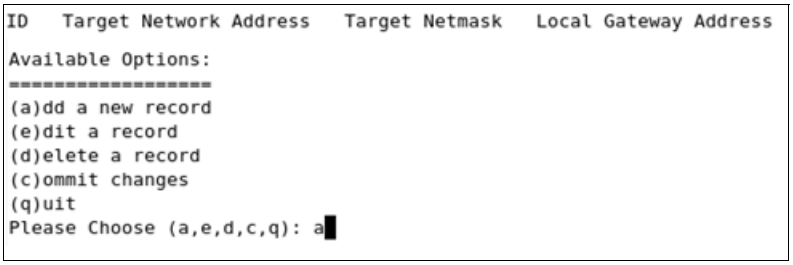

Figure 4-8 ProtecTIER Service Menu - (a)dd a new record

On a Windows host, you configure static routes by running the **route add** and using the **-p** flag to make the change permanent:

```
route add -p "target network address" mask "target network netmask" "local gateway address"
```

Example 4-1 shows an example of this procedure.

Example 4-1 Configure the static route for FSI on a windows host

```
C:\Users\Administrator>route add -p 10.200.104.0 mask 255.255.255.192 10.200.31.62
```

4.2.2 Replication

You can use ProtecTIER to define replication policies to replicate a file system's directories and all the objects that are contained in these directories recursively to remote ProtecTIER repositories without any disruption to the operation of the file system as a target for backup. It is possible to define up to 64 source directories per one replication policy, and to define up to three remote ProtecTIER destinations. The replicated data in the remote destination can be easily used to restore data in the case of a disaster recovery, or in the case of a disaster recovery test (without any interruption to the backup and replication procedures).

It is important to allow the ProtecTIER system to supervise all the changes that are made to a directory, or to a set of directories, that is constantly defined within a replication policy. Therefore, you should not disable a replication policy unless this policy is no longer considered relevant. If there is a scheduled maintenance of the network that is used for replication, it is possible (though not mandatory) to suspend the replication to a specific destination. Suspending replication allows the ProtecTIER system to continue supervising all of the changes, but it does not attempt to send the replication data through the network for the time that is defined by the suspend operation. The suspend operation is limited in time, with a maximum suspend time of 72 hours.

If a policy is disabled for some reason, then a new Replication Destination Directory (RDD) must be defined to re-enable the policy. The ProtecTIER system does not need to replicate all of the data from scratch if the old RDD is not deleted; it needs to create only the structure and metadata within the new RDD. Therefore, you should not delete the old RDD until at least a new cycle of replication to the new RDD is complete.

4.2.3 Disaster recovery: Test

Use the ProtecTIER cloning function for disaster recover (DR) testing in an FSI environment. Cloning creates a space-efficient, writable, and point-in-time copy of the data without disruption to the ongoing replications and recovery point objective (RPO). The DR test can be performed on the cloned data while the source repository continues replicating data without modifying any data on the cloned copy.

4.2.4 Disaster recovery: Event

If there is a real disaster recovery event where the primary repository that owns the backup data is temporarily or permanently down, the data can be restored from the replicated copy. If you want to take new backups at the DR ProtecTIER system during the DR event, then you must take ownership of the replication destination directory (RDD) in order to have write privileges.

Taking ownership of an RDD means that the replication directory can be accessed through shares with read/write permissions. After an RDD is modified to be read/write accessible, the source repository can no longer replicate data to the modified RDD. The modified RDD now becomes a “regular” directory and can be used as a source for replication and can have shares that are defined to it with writing permissions.

For more information about this procedure, see Chapter 19, “Native replication management (for FSI)”, in *IBM System Storage TS7600 with ProtecTIER V3.2 User's Guide*, GS32-0922.

4.2.5 General FSI recommendations

You should disable any encryption features in the backup server when you use ProtecTIER as the backup target, as shown in Table 4-2.

Table 4-2 Recommended settings

Parameter	Value in backup application
Compression	Disable
Deduplication	Disable

Parameter	Value in backup application
Encryption	Disable
Multiplexing	Disable

4.3 File System Interface guidelines for NFS

This section provides an introduction and best practices for configuring the ProtecTIER File System Interface (FSI) for Network File System (NFS) protocol. The ProtecTIER FSI for NFS emulates a Network File System that is accessed by UNIX Operating Systems. The FSI NFS file system presents a virtualized hierarchy of file systems, directories, and files to UNIX NFS clients. The ProtecTIER FSI interface is intended to be used for backup and restore of data sets by using a backup application.

4.3.1 ProtecTIER NFS authentication and security management

ProtecTIER Version 3.3 implements FSI NFS exports with NFS protocol Version 3. Access to the export is granted either for a single host or a host group. Before we guide you through the process of creating and mounting an FSI NFS share, we describe the most important options that you must specify when you use the create NFS share wizard.

- Port Security
- Root squash / no root squash

Port Security

In the Properties tab of the export NFS share wizard, you find the Port Security option under the Details section. Here you can select if you want to allow NFS clients to connect to ports higher than 1023. The port numbers 0 - 1023 are the well-known ports, also known as system ports. These TCP/IP port numbers allow only root users and services to run servers on these ports. Port 1024 and higher are also known as user ports.

Keep the default setting and leave the check box selected.

Root squash

Enabling this option prevents root users from the NFS client systems from having root privileges on the NFS export shares that are provided by ProtecTIER. If you do not enable root squash, any root user of a remote system could delete any user data on the mounted NFS shares because root can delete any data of foreign users. To prevent this action, **`root squash`** maps the root user ID 0 (UID) and group ID 0 (GUI) to a customized UID. By doing this task, the remote root user cannot delete or modify any other data than the one that is created with the customized UID. Typically, the root squash function by default maps the UID to nfsnobody, but in the ProtecTIER implementation, the UID of that user is higher than the value that you are allowed to enter in the wizard's field. Alternatively, the **`no_root_squash`** option turns off root squashing.

Mapping the root user to an UID that does not exist on the ProtecTIER system is possible but not recommended. Instead, map it to an existing user such as *nobody.* The nobody user has limited permissions and is not allowed to log in to the system. Alternatively, you can create a user and a group with limited permissions and map the root users of the client host systems to these IDs.

Example 4-2 shows how to determine the UID and the GUI of the user nobody. This user exists in the ProtecTIER system. You must log on to the ProtecTIER command line using Secure Shell (SSH) to query the user account information.

Example 4-2 Determine the user ID and group ID of user nobody on the ProtecTIER server

```
root@BUPKIS]# grep nobody /etc/passwd
nobody:x:99:99:Nobody:/:/sbin/nologin
nfsnobody:x:4294967294:4294967294:Anonymous NFS User:/var/lib/nfs:/sbin/nologin
[root@BUPKIS]# grep nobody /etc/group
nobody:x:99:
```

The output of the commands in Example 4-2 shows that the numeric value for the user and group are both 99. You can use this number to configure the root user ID mapping or create a customized user account and a dedicated group to map one or more root accounts of the remote NFS clients. If you decide not to use the existing nobody account, you can create your own customized group and several users, as shown in Example 4-3.

Example 4-3 Create a customized group and user

```
[root@BUPKIS]# groupadd -g 65536 nfsanonymous
[root@BUPKIS]# useradd -u 65536 -g nfsanonymous -M -s /sbin/nologin
                       -c "Anonymous PT NFS client user" nfsanonymous
```

4.3.2 Configuration of a ProtecTIER FSI-NFS share

This section describes the steps that are required to configure your ProtecTIER server, to create an NFS share on your ProtecTIER server, and how to connect your NFS backup client to the share. The details of best practices and configuring applications (Tivoli Storage Manager, EMC NetWorker, Symantec NetBackup, CommVault, and Symantec BackupExec) for ProtecTIER FSI are described in Chapter 15, “Symantec NetBackup and BackupExec” on page 239 and higher.

Before you start with the FSI NFS configuration, make sure that you review the following chapter and sections:

- Chapter 5, “Networking essentials” on page 79
- Section 4.1.3, “Connecting a ProtecTIER server to the network” on page 55
- Section 4.2, “Configuring components for ProtecTIER FSI” on page 58

Creating a file system

You can create file systems through the Systems Management view of the ProtecTIER Manager GUI. NFS shares and CIFS shares can be on the same file system and share a common file space. The file system settings refer to the entire file system content.

Creating an NFS share

To create an NFS export on an existing FSI file share through the ProtecTIER Manager, choose the menu **Create NFS export** from the top menu. In this section, we create an NFS export with the characteristics shown in Table 4-3.

Table 4-3 Values to create an NFS export

Option	Value	Description
Name	thekla_tsm6	Name of the NFS export.
Path	/thekla_tsm6	Defines the NFS mount point.
Port Security	checked	Allows the NFS client to connect only on the well-known ports 0 - 1023. **Note:** when using the AIX® default NFS mount, this option should be unchecked. If checked, AIX NFS needs to set the following option on OS level: nfso -p -o nfs_use_reserved_ports=1
Host IP	192.170.150.6	IP address of the NFS client where the backup application is running. Wildcard characters such as “*” and “?” are allowed.
Permission	write enabled	Grants the NFS client read and write access to the NFS share.
User ID mapping	root squash The default is for root squash to be disabled. We advise that root squash be enabled as illustrated in Figure 4-9 on page 65. See 4.3.3, “Understanding root squash” on page 68.	Choose either root squash or no_root_squash. For more information about this option, see “Root squash” on page 62. The ProtecTIER wizard does not check whether this user ID and the group ID exist on the ProtecTIER system. However, it checks that the entered UID and GID are -65536 to 65536.
Anonymous User Mapping	User ID: 65536 Group ID: 65536	Enter the user ID and the group ID that you want to map the root user to.

After you create an NFS export, the NFS share overview window in the ProtecTIER Manager looks similar to Figure 4-9. The picture shows two different host access definitions for the same NFS share. The first one defines a host with read/write access and root squash (anonymous user mapping) options, while the second definition grants only read access without root squash.

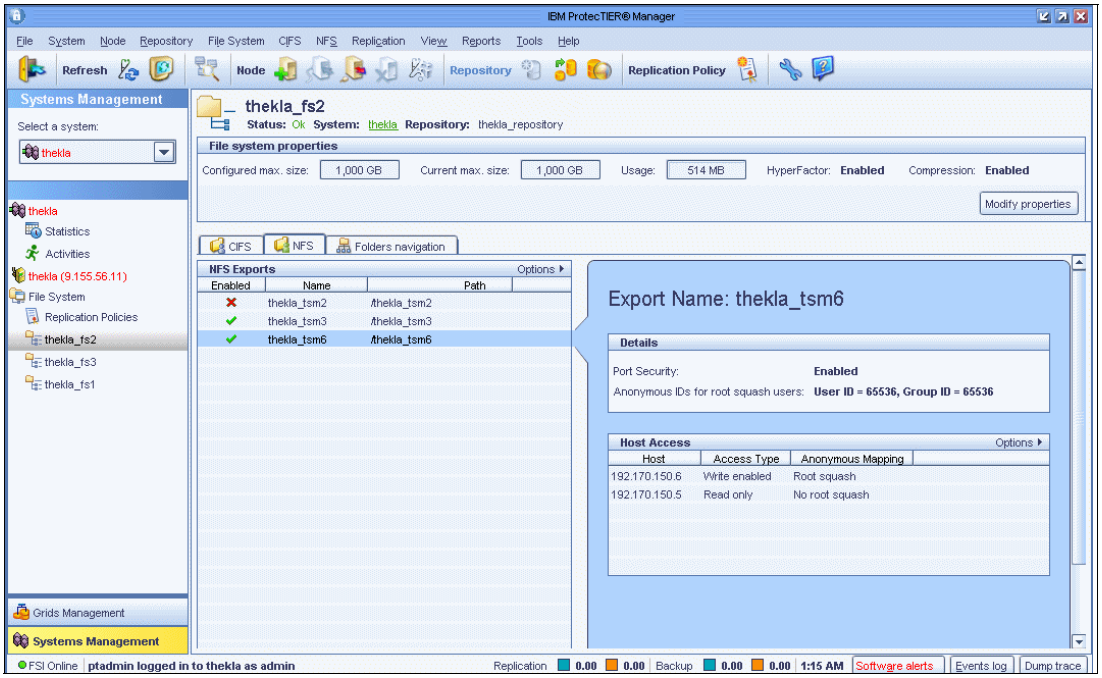

Figure 4-9 NFS export overview pane

If you want to configure an NFS export through the command-line interface to automate the configuration or create multiple shares without using the ProtecTIER Manager, you can use the command-line tool ptcli. Example 4-4 shows a command that creates an NFS export with the same characteristics shown in Figure 4-9.

Example 4-4 Create an NFS export through the command-line tool ptcli

```
[ptadmin@BUPKIS] ./ptcli CreateExportPath --ip localhost --loginFile
        /opt/dtc/ptcli/ptadmin --fsName thekla_fs2 --exportName thekla_tsm6
        --path thekla_tsm6 --secured yes --uid 65536 --gid 65536
           ip=192.170.150.6,permission=rw,mapping=root_squash\),
        \(ip=192.170.150.5,permission=r,mapping=no_root_squash\)
```

For more information about the usage of the ptcli tool and its parameters, see Appendix A, “Command-Line Interface”, in *IBM System Storage TS7600 with ProtecTIER Version 3.1*, SG24-7968.

Configuring UNIX to mount the FSI NFS export

This section describes how to connect an NFS client to the FSI NFS export share. We demonstrate the required steps and parameters for a Linux System Red Hat Enterprise Server.

Mounting the NFS share on the NFS client system

Before you mount the ProtecTIER FSI NFS export, create a host alias of the ProtecTIER FSI application IP addresses on your NFS client system by completing the following steps. You can skip this process if the IP address of the FSI interface is registered in a domain name server (DNS).

1. Add a DNS entry for each FSI interface IP in your domain name server or add an alias to the local hosts definition file of your backup server. Example 4-5 shows how to add aliases for two different FSI interfaces. Edit the `/etc/hosts` file on your backup application server, add the IP addresses of the FSI interface you configured on the ProtecTIER server, and define a host alias. Afterward, verify the connectivity by pinging the host alias.

Example 4-5 Create a host alias and verify connectivity to the FSI interface

```
[root@BUPKIS ~]# vi /etc/hosts
192.170.150.40 thekla_fsi1
192.170.151.40 thekla_fsi2

[root@bupkis ~]# ping thekla_fsi1
PING thekla_fsi1 (192.170.150.40) 56(84) bytes of data.
64 bytes from thekla_fsi1 (192.170.150.40): icmp_seq=1 ttl=64 time=1.02 ms
```

2. Create a mount point and mount the NFS export,as shown in Example 4-6. The **`mount`** command uses the host alias instead of an IP address.

Example 4-6 Mount the NFS share on a Linux host

```
[root@BUPKIS ~]# mkdir /mnt/thekla_tsm6
[root@BUPKIS ~]# mount -o rw,soft,intr,nolock,timeo=3000,nfsvers=3,proto=tcp
thekla_fsi1:/thekla_tsm6 /mnt/thekla_tsm6/
```

Table 4-4 explains the recommended **`mount`** options. The additional and optional parameters are specified after the **`-o`** option. The **`-o`** option is followed by a comma-separated string of options, as shown in Example 4-6.

Table 4-4 Recommended mount options

Recommended value	Explanation
`rw`	The NFS share is mounted for read and write operations.
`soft`	Valid values are **`hard`** or **`soft`**. The value specifies whether the program using a file through an NFS connection should stop and wait (**`hard`**) for the server to come back online, if the host serving the exported file system is unavailable, or if it should report an error (**`soft`**) after the specified timeout. If hard is specified, the user cannot terminate the process that is waiting for the NFS communication to resume unless the **`intr`** option is also specified. This configuration might lead to hang conditions in Tivoli Storage Manager when the NFS share on the ProtecTIER system is unavailable. Thus, the value **`hard`** is not recommended. When you use **`soft`**, the user or program can set an additional **`timeo=<`*`value`*`>`** option, where <value> specifies the number of seconds to pass before the error is reported. If neither option is specified, NFS requests are retried indefinitely.

Recommended value	Explanation
intr	Valid options are **intr** or **nointr**. If the recommended NFS mount option **intr** is specified, it is possible to interrupt NFS requests in case the NFS server is unreachable. Using the **intr** option is preferred with the **soft** option because it reduces the risk of inconsistent data during a write failure.
nolock	Valid options are **nolock** or **lock**. Specify the **nolock** parameter instead of the **lock** parameter because when you use **nolock**, then applications can lock files, but such locks provide exclusion only against other applications that are running on the same client.
timeo=3000	The time (in tenths of a second) that the NFS client waits for a response from the ProtecTIER server before it retries an NFS request. If this option is not specified, requests are retried every 60 seconds. The NFS client does not perform any kind of timeout backoff for NFS over TCP. The recommended value **timeo=3000** specifies a timeout of 5 minutes.
nfsvers=3	The only valid option is NFS Version 3.
proto=tcp	Specifies for the NFS mount to use the TCP protocol and not UDP.

The generic NFS mount syntax with the recommended parameters is shown in Example 4-7.

Example 4-7 Generic mount syntax for Linux

```
mount -o rw,soft,intr,nolock,timeo=3000,nfsvers=3,proto=tcp <server>:<path>
/<mountpoint>
```

Example 4-8 shows the mount syntax for IBM AIX systems. The AIX **mount** command requires AIX V6.1 TL8 or AIX V7.1 TL2.

Example 4-8 Generic mount syntax for AIX

```
mount -o
rw,soft,intr,llock,timeo=3000,nfsvers=3,proto=tcp,rsize=262144,wsize=262144
<server>:<path> /<mountpoint>
```

Example 4-9 Generic mount syntax for Solaris

```
mount -o rw,soft,intr,llock,timeo=3000,nfsvers=3,proto=tcp <server>:<path>
/<mountpoint>
```

Note: In contrast to an FSI CIFS environment, it is not an absolute requirement to use host name aliases or DNS names instead of static IP addresses for mounting the share. However, we recommend the usage of host aliases over static IP addresses for FSI NFS, as well when defining and mounting the NFS share on the backup clients. If you correct the IP aliases on your backup host after you change the ProtecTIER FSI configuration or after you switch to an RDD in a disaster recovery scenario, this approach might save time during reconfiguration. Instead of using a host alias, you could also adjust the **mount** commands or the fstab definition to reflect the changed IP address.

4.3.3 Understanding root squash

The basics of NFS root squash and no root squash are explained in “Root squash” on page 62. The following section demonstrates the effects of turning root squash on or off.

Example 4-10 shows a directory listing of a ProtecTIER NFS share. The `file1` file was created by a root user. Usually, the user ID of root is 0, but because we turned on root squash when we defined the NFS export, the root user ID is mapped to a defined UID (in our example, they are user ID 65536 and group ID 65536). The `file2` file was created by the tsminst1, which belongs to the tsmsrvrs group.

Example 4-10 Directory listing on an NFS share

```
[tsminst1@Amsterdam thekla_tsm6]$ ls -ltrh
total 1.0K
-rw-r--r--. 1    65536    65536 12 Nov  7 02:07 file1
-rw-r--r--. 1 tsminst1 tsmsrvrs 12 Nov  7 02:08 file2
```

When root squash is enabled, the root user looses the authority to delete files that belong to any other user ID than the root squash user ID. In this example, the root user is not allowed to delete files of tsminst1 anymore. Turning on root squash is an important security feature. It prevents the possibility that any root user of any host can mount the share and delete data that belongs to other systems and users.

Example 4-11 demonstrates that the root user ID is not allowed to delete `file2`, which belongs to tsminst1. The **`delete`** command fails with an error message `Permission denied`.

Example 4-11 Deleting files with root squash enabled in the NFS export definition

```
[root@Amsterdam thekla_tsm6]# rm file2
rm: remove regular file `file2'? y
rm: cannot remove `file2': Permission denied
```

To demonstrate the power of the root user without the root squash function enabled, we modified the NFS export share definition and disabled root squash. In comparison to Example 4-11 the root user can delete `file2` even if the file is owned by tsminst1. The result of the delete operation is shown in Example 4-12. The `file2 was` deleted without any error.

Example 4-12 Deleting files with root squash disabled in the NFS export definition

```
[root@Amsterdam thekla_tsm6]# rm file2
rm: remove regular file `file2'? y
[root@Amsterdam thekla_tsm6]# ls -ltr
total 1
-rw-r--r--. 1 65536 65536 12 Nov  7 02:07 file1
```

4.4 File System Interface guidelines for CIFS

This section provides a general introduction and best practices for configuring the ProtecTIER File System Interface (FSI) for CIFS. The ProtecTIER FSI emulates Windows file system behavior and presents a virtualized hierarchy of file systems, directories, and files to Windows CIFS clients. Clients can perform all Windows file system operations on the emulated file system content. The ProtecTIER FSI interface is intended to be used for backup and restore of data sets using a backup application.

This section describes how to create a CIFS share on ProtecTIER, connecting to a CIFS share, and shows best practices.

4.4.1 ProtecTIER authentication and user management

The ProtecTIER product supports two modes of authentication and user management in a CIFS environment:

- Active directory
- Workgroup

In the Active directory mode, the ProtecTIER system joins an existing domain that is defined by the user. The domain users can work with the file systems if they are authenticated by the Active Directory server. In Workgroup mode, the ProtecTIER system manages the users that can access the file systems. In Workgroup mode, you define the users through the ProtecTIER Manager GUI.

Active Directory and user IDs

The ProtecTIER system assigns user IDs to Windows users that access the system through CIFS. The only control that you have is to set the range of *user IDs* that are generated. You should set a range that is not overlapping with UIDs used for existing UNIX users in the organization.

Active Directory realm

One of the parameters that must be provided to the ProtecTIER system when you define authentication mode to Active Directory is the *realm*. In most cases, the name of the realm is the DNS domain name of the Active Directory server. The realm should always be in uppercase and should not be a single word (for example, add .COM or .LOCAL to the domain name).

Some helpful commands can be used to define the realm:

- From the Active Directory server, run the following command:

```
C:\>ksetup
default realm = RNDLAB02.COM ---------------> The realm
```

- From the ProtecTIER server, run the following command:

```
net ads lookup -S IP_Address_of_ADServer
```

Example 4-13 shows output for the **net ads lookup** command.

Example 4-13 Output of the net ads lookup command

```
net ads lookup -S 9.148.222.90
output:
Information for Domain Controller: 9.148.222.90
Response Type: SAMLOGON
GUID: 9a1ce6f2-17e3-4ad2-8e41-70f82306a18e
Flags:….
Forest:                   rndlab02.com
Domain: R                    NDLAB02.COM ---------------> The realm
Domain Controller:  RNDAD02.rndlab02.com
Pre-Win2k Domain:       RNDLAB02
Pre-Win2k Hostname:     RNDAD02
Server Site Name :              Default-First-Site-Name
Client Site Name :              Default-First-Site-Name
```

4.4.2 Configuring your ProtecTIER system for FSI-CIFS

The following sections provide guidance for configuring your ProtecTIER system for FSI-CIFS. These sections describe the steps for creating file systems and for setting up CIFS authentication for either workgroups or active directories. We also describe how to create a CIFS share on a ProtecTIER system and how to connect to a CIFS share from a Windows client.

Creating a file system

You can create file systems through the Systems Management view of the ProtecTIER Manager GUI. To create the file system for system Bupkis, from the Systems Management window, click **File System** → **Create file system**, as shown in Figure 4-10.

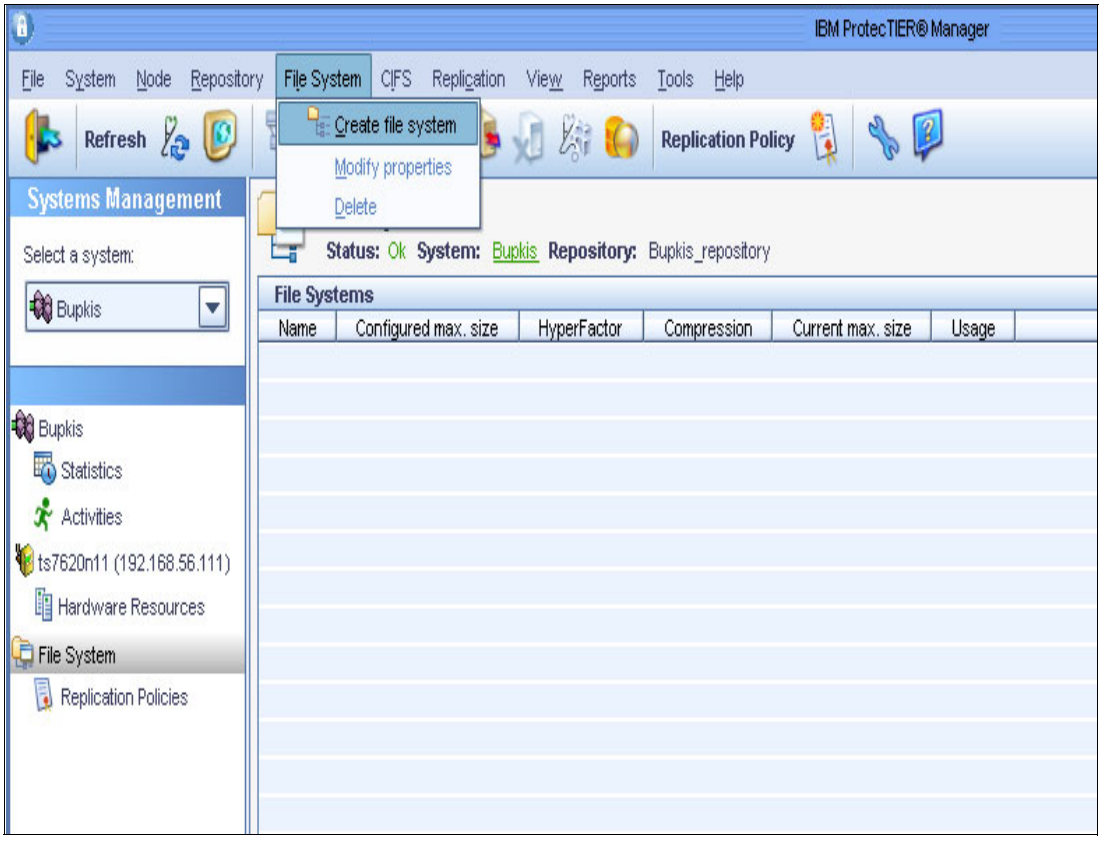

Figure 4-10 Create CIFS file system

FSI-CIFS file system scalability

A FSI-CIFS file system can scale to the following values:

- Maximum virtual file systems per repository: 128, 4 on SM2
- Maximum nominal virtual file system size: 256 TB
- Maximum files per virtual file system: 1 million
- Maximum files per repository: 16 million, 4 million on SM2
- Maximum “open files” per replication (streams): 192, 64 on SM2

Setting CIFS authentication to either a workgroup or Active Directory

Before you set up the ProtecTIER FSI CIFS configuration, consider what is the best authentication mode for your environment. If a decision is made to use Workgroup authentication, verify that Active Directory is not used in your environment. If Active Directory is used in your environment, check which UIDs are used by it, and ensure that you use a unique range of UIDs in the ProtecTIER workgroup that is not being used in the Active Directory. This situation is also true if you late install Active Directory in your environment.

After an authentication method is implemented and used by ProtecTIER FSI CIFS, do not change it (from workgroup to an Active Directory and vice versa) because this action can lead to various authorization issues.

Attention: If an authentication method change must be done, and cannot be avoided, do the change with the involvement and guidance of your IBM System Service Representative (SSR).

When you decide to use Active Directory authentication with ProtecTIER FSI CIFS, use SFU mode to allow seamless compatibility between CIFS and NFS. SFU mode enables you to set the UID/GID of the user to be compatible with the one that is defined within the NIS servers. This configuration enables the same user to access his data from both Windows and UNIX/Linux.

Use the authentication setup wizard to set the authentication that the ProtecTIER system uses for working with CIFS clients. To accomplish this task, complete the following steps:

1. From the Systems Management menu, for system Bupkis, click **CIFS** → **Authentication** → **Set the authentication mode**, as shown in Figure 4-11.

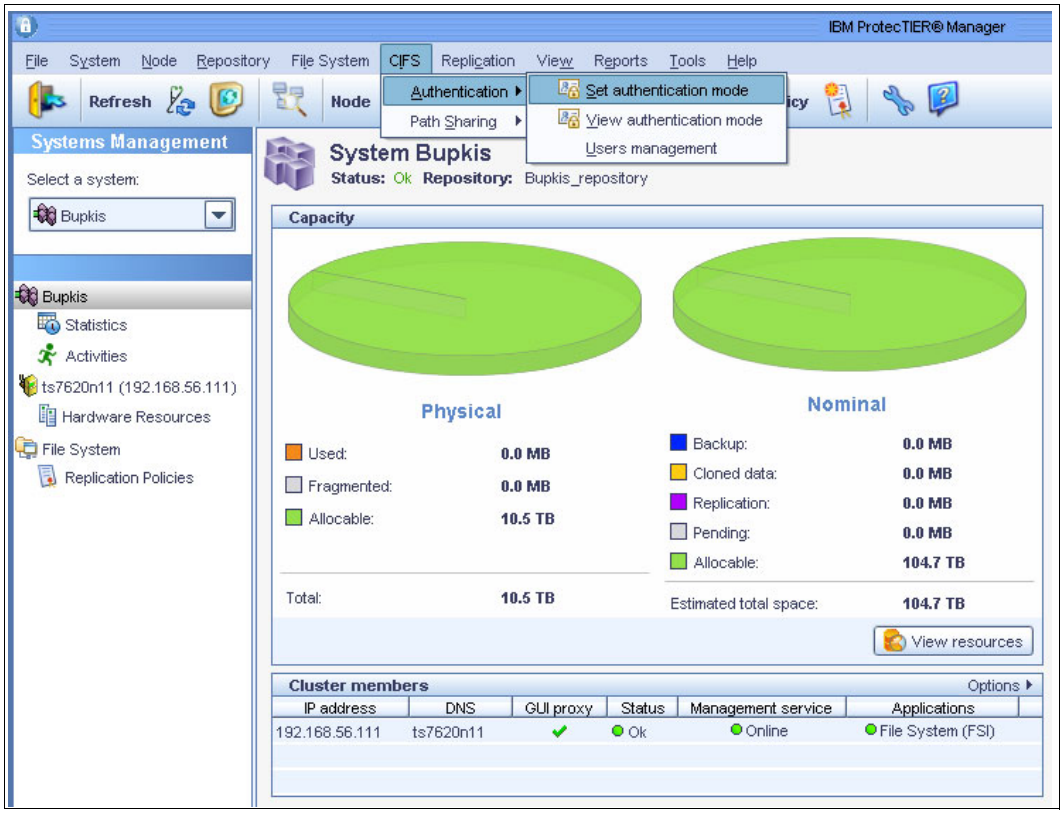

Figure 4-11 Authentication mode

2. To manage data with CIFS, you must set the authentication to either Workgroup or Active Directory. In Figure 4-12, ProtecTIER is configured to join an Active Directory domain. Click **Next**.

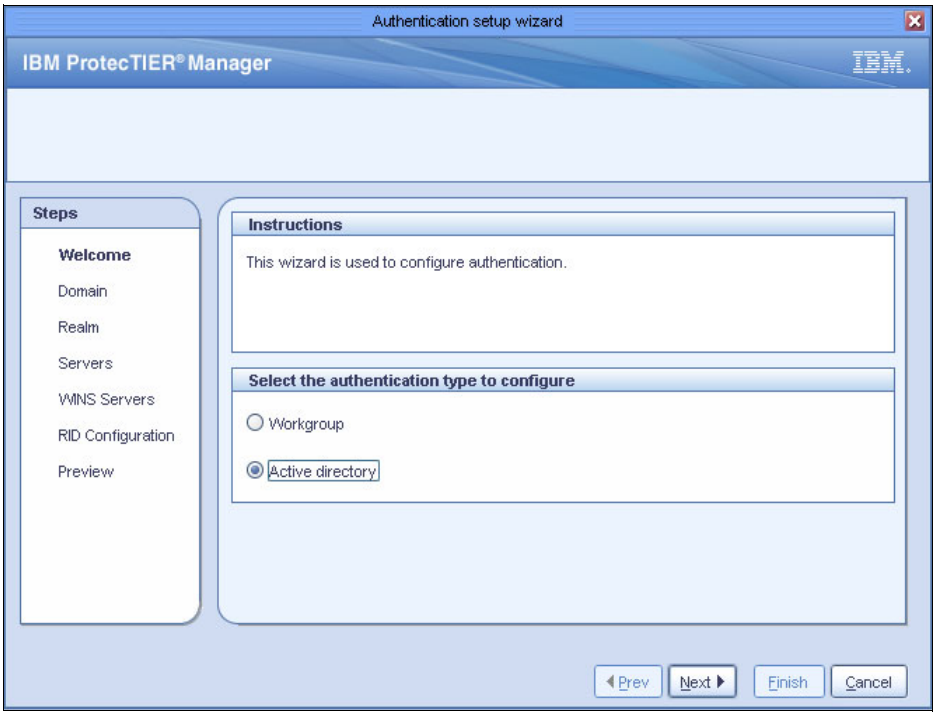

Figure 4-12 Authentication mode

3. To set up the Active Directory, enter the Active Directory server credentials and password. Enter the Active Directory server address and click **Next**, as shown in Figure 4-13.

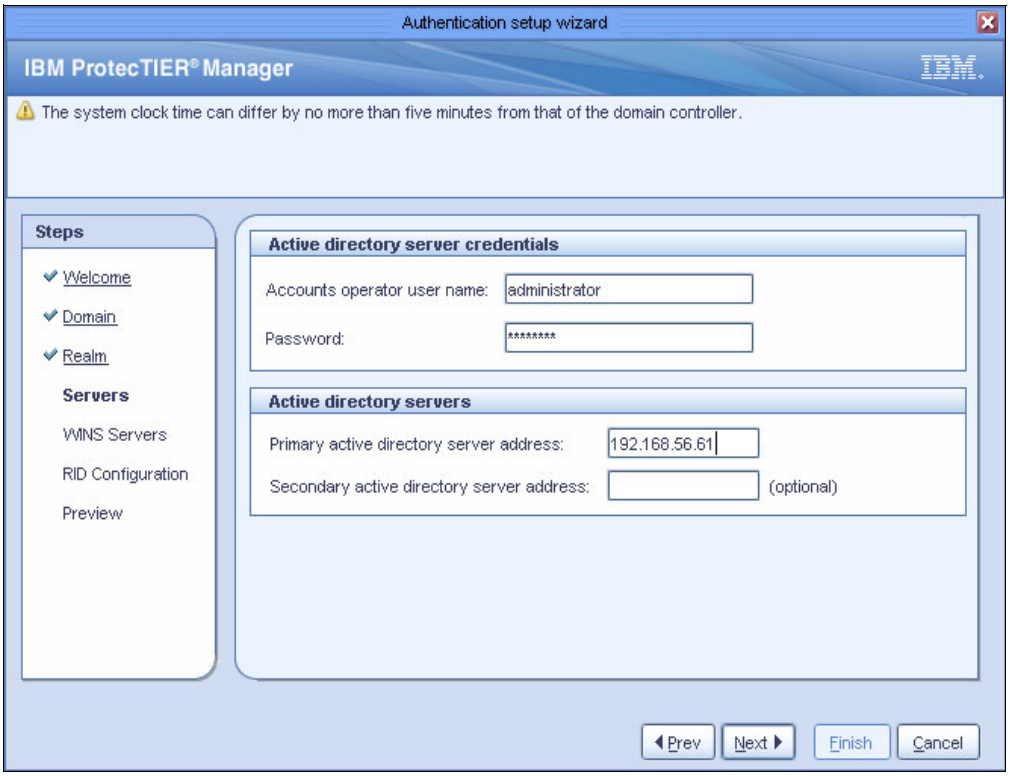

Figure 4-13 Active Directory setup

4. A report is displayed that summarizes your authentication setup, as shown in Figure 4-14.

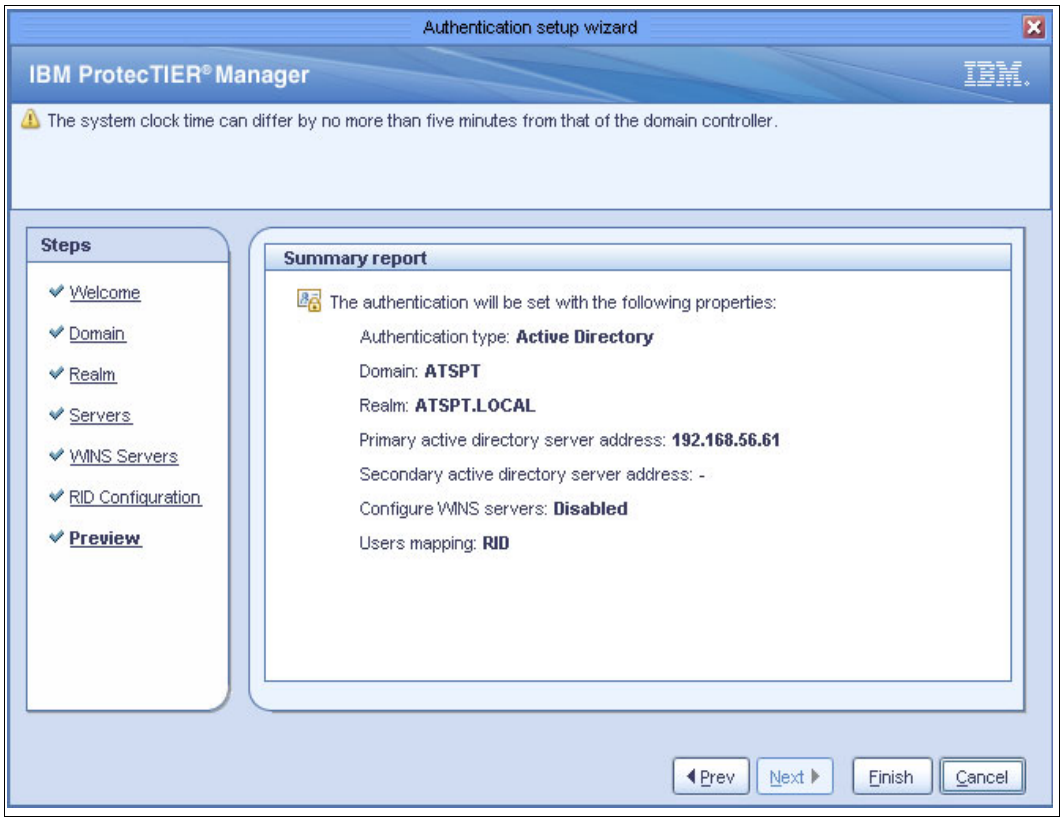

Figure 4-14 Authentication setup summary

Adding a CIFS user for workgroups

Use this task to add users to a workgroup for access to a CIFS shared path. This task is done from the Systems Management view of ProtecTIER Manager, as shown in Figure 4-15.

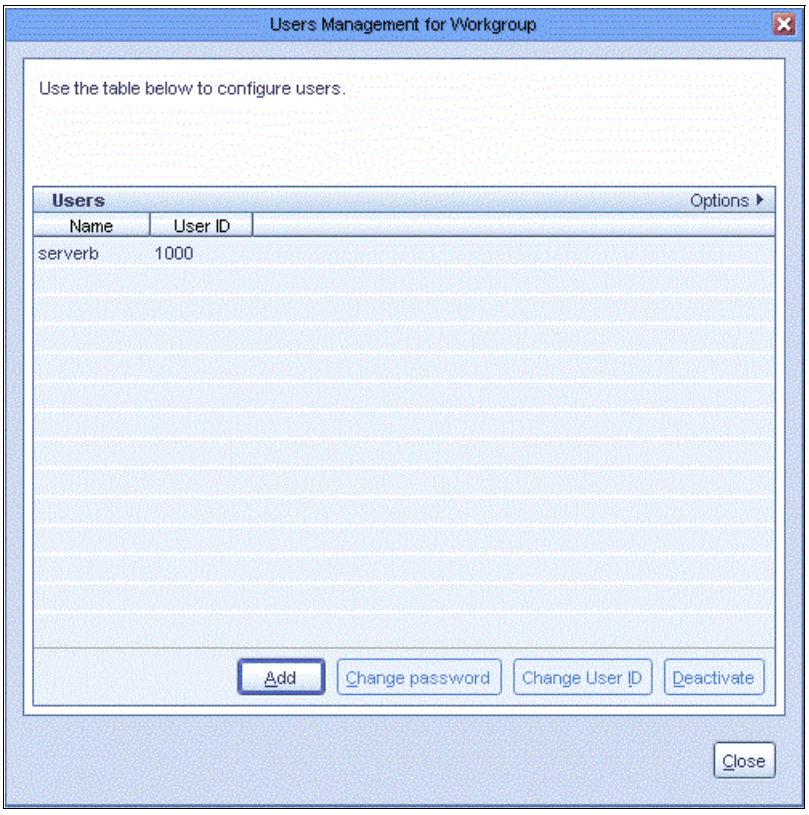

Figure 4-15 Workgroup authentication users

User account: Log on to the backup application with the same user account that is used for the CIFS user account.

Creating a CIFS share on a ProtecTIER server

To enable CIFS access and share data, you must grant a host and a user or group access to shared directories by creating and configuring a CIFS shared path. This task is done from the Systems Management view of ProtecTIER Manager GUI, as shown in Figure 4-16.

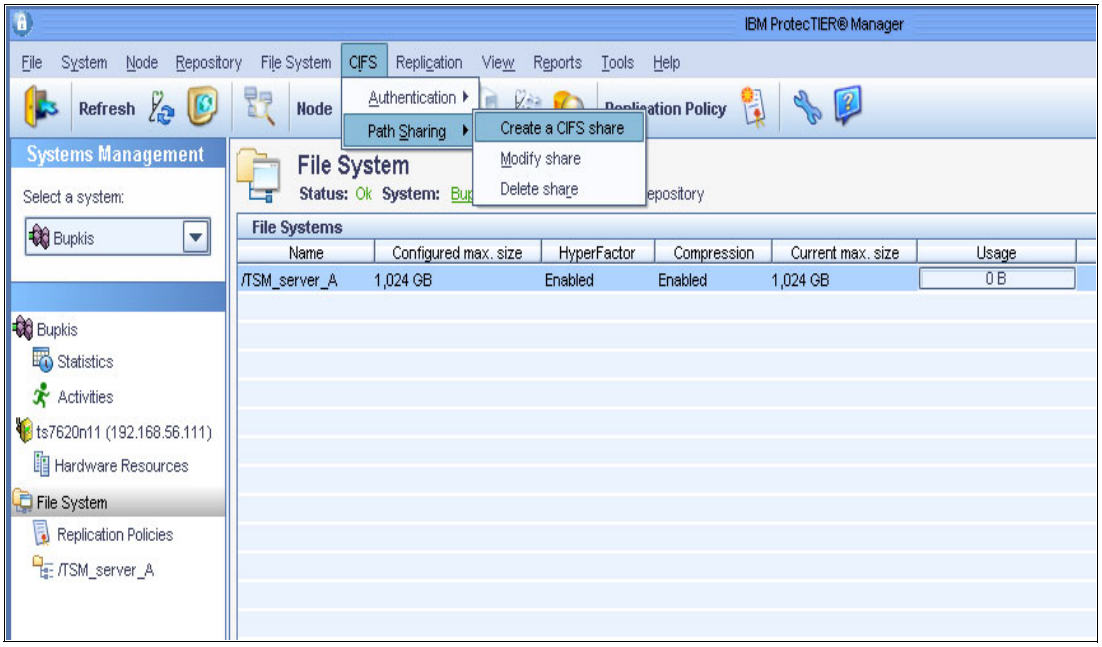

Figure 4-16 Create a CIFS share

Assign the user permissions by selecting **Write-enabled** from the **Permissions** menu (Figure 4-17).

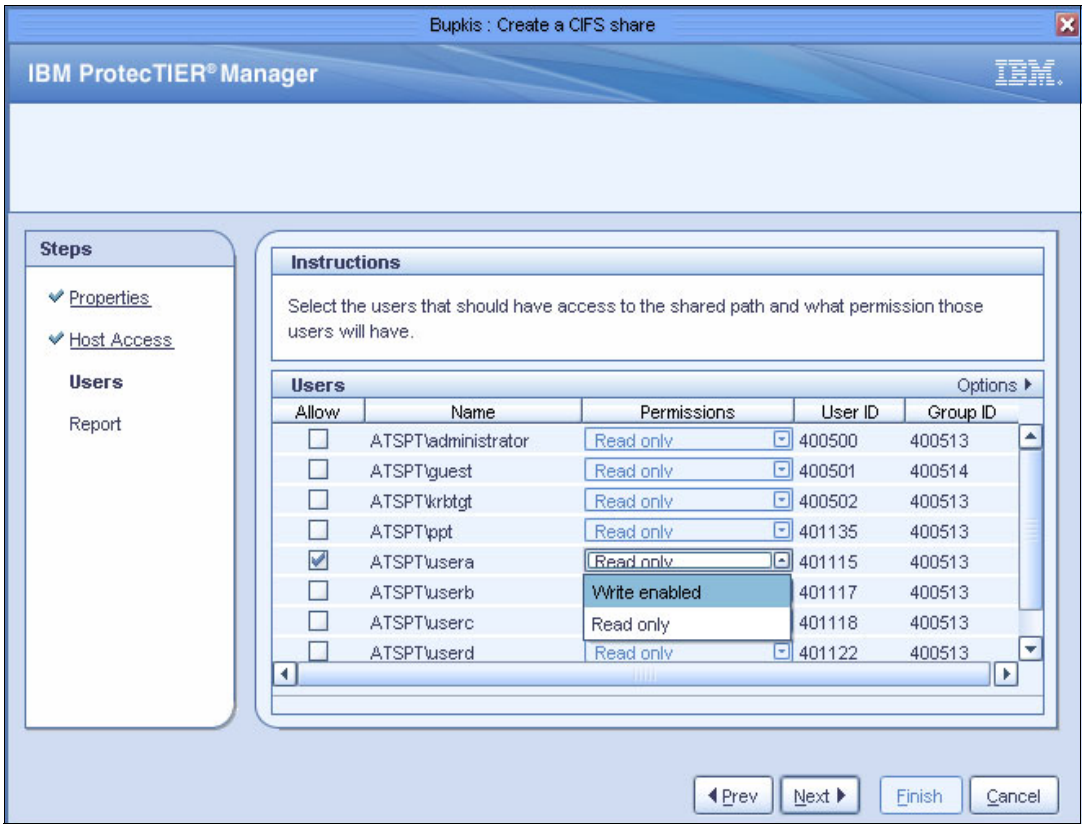

Figure 4-17 Configure Share Permissions

Connecting to a CIFS share

To connect to a share from a Windows client, complete the following steps:

1. Click **Start** → **Run**.
2. The Run window opens (Figure 4-18). In the **Open** field, enter \\<FSI_IP> (or the ProtecTIER server name) and click **OK**.

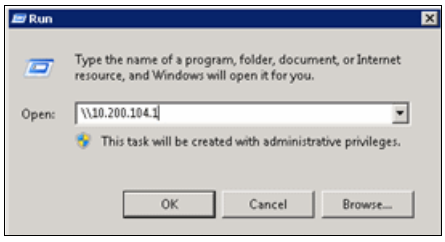

Figure 4-18 FSI IP

The list of shares window opens (Figure 4-19).

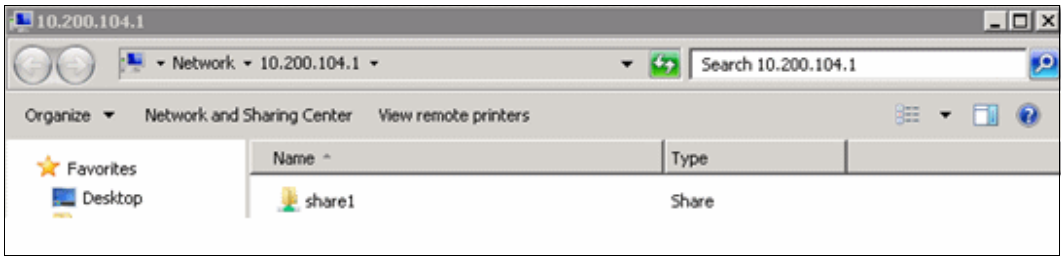

Figure 4-19 List of shares

5

Networking essentials

This chapter describes the ProtecTIER basic networking considerations, terminology, and concepts that are common to all ProtecTIER configurations using non-Fibre Channel network interfaces for front-end connectivity. These interfaces include OpenStorage (OST), and File System Interface (FSI) with the options of Network File System (NFS), and Common Internet File System (CIFS).

In a non-VTL environment, the ProtecTIER server connects to the network with multiple Ethernet adapters, and supports a total throughput of hundreds of megabytes per second over the network per node. This section describes network configuration best practices to support different IP configurations. It provides the scope and objectives, along with the basic acronyms that are used in the subsequent sections. The goal is to describe the networking technologies that are used to set up the ProtecTIER front-end interfaces to the network. In addition, this section describes teaming or bonding in the network servers, 802.3ad link aggregation, and similar technologies on the LAN switches, along with network topologies for stand-alone and redundant architectures.

Some of these networking concepts also apply for ProtecTIER network replication environments.

This chapter describes the following topics:

- FSI and OST overview and concepts
- Networking overview and description of the following items:
  - Bonding and teaming
  - Redundant mode
  - 802.3ad topologies
  - Intel Network Interface Cards (NIC)
  - Technologies for LAN switches
- Setting up your network in a ProtecTIER environment

No support as primary storage: Using the ProtecTIER FSI as an NAS in terms of storing primary data directly to it is not supported currently. ProtecTIER FSI must be used with an application that is supported by ProtecTIER.

Terminology: In this chapter, we use the term *bonding* when we talk about aggregating multiple network links to a single logical interface. You might have heard terms such as network aggregation, link aggregation, port trunking, trunking, link bundling, Ethernet bonding, network bonding, NIC bonding, or NIC teaming, which also describe the same concept.

Cabling: Multiple members of a single bond must be connected to the same network switch. If you use multiple physically independent networks, similar to a dual-fabric approach in Fibre Channel terms, you are required to have at least two bonds. Each bond must be cabled to only one network switch.

5.1 Network terminology

This topic defines terms and abbreviations that are relevant to network configurations in general:

Bonding	A method for grouping several physical adapters into a single virtual adapter for load sharing, throughput enhancement, and redundancy enhancement. “Bonding” is the term that is typically used in Linux/UNIX operating systems.
Teaming	An alternative term for bonding, typically used in Microsoft operating systems.
Trunking	An alternative term for bonding or teaming.
Bundling	An alternative term for bonding or teaming.
Link aggregation	A method for grouping several interfaces into a single virtual interface for load sharing between the interfaces. Also known as network aggregation.
IEEE 802.3ad	IEEE standard for Link Aggregation for LAN connectivity.
Gigabit Ethernet	Ethernet that runs in a gigabit per second bandwidth.
VLAN	Virtual LAN is a software-defined LAN that groups network elements in the same broadcast domain.
Host	An entity that is connected to the network. For example, the NetBackup media servers are referred to as hosts.
Bonding/Teaming	Teaming or bonding in servers, and link aggregation (802.3ad) or Cisco Etherchannel in LAN switches. The purpose of these mechanisms is to achieve higher bandwidth on the connection, as close as possible to the multiplication of the ports' bandwidth, along with redundancy between the ports.

5.2 General configuration considerations

This section describes network configuration considerations and best practices for FSI-CIFS. Here are some general best practices:

- As ProtecTIER IP replication with FSI is realized on a file share level, you should create a dedicated file share for each backup server that you use with ProtecTIER FSI. This file share allows for more sophisticated load balancing scenarios.
- Make sure that the backup application runs in the context of the user that mounted the file system with the write permission. If you suffer from access rights issues, this configuration most probably is the reason for them.
- You must have at least two different network subnets to separate the ProtecTIER management IP interface from the ProtecTIER file system IP interfaces. Otherwise, you cannot ensure that your backup traffic is using the FSI or OST interfaces and not the management interface.
- For FSI workloads, you might not have dedicated infrastructures for the incoming and outgoing traffic of backup servers. Ensure that you do not suffer from infrastructure congestion.
- Bonding must be implemented on all involved network devices, whether the devices are the ProtecTIER server, the backup server, or even the network switches. Simply enabling bonding on the ProtecTIER server might not be enough to achieve the best results.

5.3 Bonding and teaming

To achieve high availability, load balancing, and increased throughput, you can use a network technology that is known by many names. You might have heard of terms such as network aggregation, link aggregation, port trunking link bundling, Ethernet bonding, network bonding, NIC bonding, 802.3ad, or NIC teaming. All these terms describe solutions that you can use to achieve high availability, load balancing, or increased throughput by combining multiple network interfaces and using them as one logical link. Table 5-1 on page 84 shows the available ProtecTIER link aggregation modes, bonding options, and features for each link aggregation mode.

Connectivity in a ProtecTIER environment is based on bonding (usually on Linux or UNIX platforms) or teaming (usually in Microsoft platforms) in servers, and link aggregation (802.3ad) or Cisco Etherchannel in LAN switches. These mechanisms achieve higher bandwidth on the connection and provide redundancy between the ports.

Teaming and bonding group interfaces in Layer 2 (on the Ethernet layer), and the whole team has a single IP address. The following modes are common:

- Transmit and Receive load balancing modes: In these modes, a receive or transmit load is distributed independently of the LAN switch to which they are connected to.
- 802.3ad mode: In this mode, the servers and the switches they are connected to must support the 802.3ad standard and load distribution is performed according to this standard.

5.3.1 The three different bonding modes of ProtecTIER

The ProtecTIER product supports three basic modes of bonding. For summaries of the bonding methodologies that are supported by the ProtecTIER product, see Table 5-1 on page 84.

Note: For the following three modes, the first modes are topologies that are switch-less because the switch does not have to support any specific standard. In the last mode, the switches in the topology must support the 802.3ad standard, or in some cases, the Cisco Etherchannel implementation.

Method one: High availability

The high availability load balancing method uses an active-backup policy. Only one interface in this bond is active. If one of the interfaces fails, the other interface becomes active, and takes over communication. This bond's MAC address is only visible on one port so that the switch is not confused. With this method, you can achieve fault tolerance. It is also called redundant mode or active-backup mode.

Important: This mode does not attempt to perform load balancing.

Method two: Round robin

The round robin load balancing method uses a round robin policy. Outgoing network traffic is distributed across all members of the bond in sequential order. Incoming network traffic is still limited to one single network port (the primary). If one of the network interfaces fails, the other bond members take over. Outgoing traffic is distributed across the remaining bond members. If the primary adapter for incoming traffic fails, the bond automatically selects a new primary adapter, and incoming traffic is handled from that one single network adapter. With this mode, you can achieve fault tolerance and load balancing. A potential unidirectional bandwidth increase for outgoing traffic is possible if you have multiple backup servers in your environment.

Round robin mode: With the round robin mode, you do not see any bandwidth increase if you have only two communication partners, for example, the ProtecTIER server and one backup server. Even enabling this mode on both of these machines does not allow you to use more than the bandwidth of one single interface, as the incoming traffic of both machines is dealt with only by one single network interface. If you have enough communication partners, for example, multiple backup servers and one ProtecTIER server, the increased bandwidth is used only during parallel restore of multiple backup servers at the same time.

Method three: L2, L2L3, L3L4, 802.3ad, mode=4, and dynamic link aggregation

The third load balancing method is the most advanced method. If you set it up correctly, you can use the combined throughput of all the involved network interfaces for incoming and outgoing traffic, balance the load across all available interfaces, and have a fault tolerant system at the same time.

To use this method, you must have a network infrastructure that fully supports it end-to-end. This bonding method relies on the IEEE 802.3ad dynamic link aggregation standard. It is therefore also known as the 802.3ad mode. The servers and the switches they are connected to must support the 802.3ad standard, and load distribution is performed according to this standard. You can it use to aggregate a group of interfaces of the same speed and duplex setting. Table 5-1 on page 84 has a summary of the bonding methods and descriptions of the associated network layers. The 802.3ad standard does not mandate any particular distribution algorithms. However, any distribution algorithm ensures that the following actions do not occur:

- Misordering frames that are part of any conversation
- Duplicating frames

The standard suggests, but does not mandate, that the algorithm may assign one or more conversations to the same port; however, it must not allocate some of the frames of a conversation to one port and the remainder to different ports. The information that is used to assign conversations to ports could include the following items:

- Source MAC address
- Destination MAC address
- Source IP address
- Destination IP address
- The reception port
- The type of destination address (individual or group MAC address)
- Ethernet Length/Type value (protocol identification)
- Higher layer protocol information (for example, addressing and protocol identification information from the LLC sub layer or above)
- Combinations of these items

The hash policy decides, according to parameters or a combination of parameters, the frames that are distributed. For example, when we have a server that exchanges information with several hosts on the same subnet, configuring a source/destination MAC hash usually produces a reasonable load distribution. If we want to use load balancing over a router, then a Layer 3 hash does not help because the server sees only one IP address (of the router), and therefore all traffic is sent over the same interface. In this case, a Layer 4 hash must be used.

Link aggregation: The 802.3ad dynamic link aggregation method is suitable to increase your throughput when you use 1 Gb network interfaces in combination with a single backup server.

It is not a best practice to combine all 1 Gb network interfaces into one single link aggregation group. Use multiple groups of two interfaces or four interfaces instead.

For options, and features for each link aggregation mode, see Table 5-1.

Table 5-1 Bonding methods and available ProtecTIER link aggregations modes

ProtecTIER GUI CLI	CLI	BONDING_OPTS	Bonding mode	Features
High Availability	HA	miimon=100 mode=1	Mode 1	Fault tolerance
Round Robin	RR	miimon=100 mode=0	Mode 0	► Load balancing ► Fault tolerance ► Unidirectional bandwidth increase
L2	L2	miimon=100 mode=4 xmit_hash_policy=layer2 802.3ad switch support needed	Mode 4 Layer 2 only (based on MAC address)	► Load balancing ► Fault tolerance ► Bidirectional bandwidth increase[a]
L2L3	L2L3	miimon=100 mode=4 xmit_hash_policy=layer2+3 802.3ad switch support needed	Mode 4 Layer 2 and Layer 3 (based on IPs)	► Load balancing ► Fault tolerance ► Bidirectional bandwidth increase[b]
L3L4	L3L4	miimon=100 mode=4 xmit_hash_policy=layer3+4 802.3ad switch support needed	Mode 4 Layer 3 and Layer 4 (based on TCP Ports)	► Load balancing ► Fault tolerance ► Bidirectional bandwidth increase possible[c]

a. Outgoing traffic is spread by using a default transmit hash policy of Layer 2. Network administrators might understand the formula (source MAC ⊕ destination MAC)%N (number of subordinates), which should be read as (source MAC “XOR” destination MAC) “MODULO” N (number of subordinates).

b. Outgoing traffic is spread by using a transmit hash policy of MAC addresses and IP addresses of the source and the destination.

c. Outgoing traffic is spread by using a transmit hash policy of IP addresses and ports of the source and the destination.

Bonding modes: L2, L2L3, and L3L4 require a network infrastructure (backup server, network switches, and ProtecTIER) that supports IEEE 802.3ad dynamic link aggregation. All of these modes use the 802.3ad standard.

For more information about bonding and teaming, see Appendix C, “Networking” on page 477.

5.4 Recommended ProtecTIER bonding configuration

This section describes the recommended ProtecTIER bonding configuration to use with FSI and OST.

When your performance requirement for ProtecTIER FSI is below 500 MBps, you can use the default ProtecTIER bonding configuration of having all four 10 Gbit interfaces within one single aggregate. If you need more than 500 MBps, optimize the way that you use the ProtecTIER FSI and configure at least two shares that are exported across two individual IPs. This configuration allows you to distribute the load across all the available resources in the ProtecTIER. This configuration is only viable if your environment allows that setup to be efficiently integrated. If you, for example, have two or more backup servers, each of them should use its own FSI file system and file share. With this setup, you have no problems using multiple file shares / IPs on ProtecTIER at the same time.

If your environment consists of only one backup server, the usage of multiple IP addresses can be a challenge. Tivoli Storage Manager, for example, allows you to create one device class that uses multiple IPs at the same time. With this setup, you can use the best practices ProtecTIER IP setup.

The backup application creates a backup or restore stream for each backup set that is copied to or from the ProtecTIER server. In an OST environment, the ProtecTIER plug-in on the media server opens, by default, up to eight TCP connections for each stream. The best practice for the ProtecTIER server is to use the default single team/bond configuration. The Layer 3 and Layer 2+3 configurations are recommended only if single bond is not applicable.

Whatever network methodology you decide to use, connect only one IP to each subnet. This requirement is an IP requirement, which is the protocol that we use on the Open Systems Interconnection (OSI) model Layer 3. Bonding all interfaces and assigning one IP to them to connect them to the network is the easiest way of attaching ProtecTIER to your network.

Table 5-2 clarifies the minimum number of interfaces, subnets, bonds, and IP addresses for different environments and performance goals. Using more of these items might be viable.

Table 5-2 Minimum numbers interfaces, subnets, bonds, and IP addresses for FSI and OST

FSI / OST interface[a]	ProtecTIER performance goal[b]	Minimum # of subnets	Minimum # of bonds and IP addresses[c]
6x 1G	< 110 MBps	1	1
6x 1G	> 110 MBps < 500 MBps[d]	1	1
6x 1G	> 500 MBps[c]	2	2
4x 10G	< 500 MBps	1	1
4x 10G	> 500 MBps	2	2

a. Interface numbers cover single node configurations. For OST cluster configurations, you can work with double the number of interfaces.
b. The maximum performance that you can reach with your ProtecTIER setup is determined by the ProtecTIER sizing.
c. This column stresses the importance of assigning only one IP address per node (whether it is ProtecTIER or your backup server) to a subnet.
d. Assuming the maximum speed of a single 1 Gbit network link is 110 MBps, you need a working 802.3ad setup or multiple subnets to reach these numbers.

IP configuration: Only one IP address per subnet is allowed. This situation is true for all backup servers and ProtecTIER.

5.4.1 Single team/bond (Layer 2) configuration

In this configuration, all front-end interfaces on a ProtecTIER server are grouped into one team. This configuration maximizes the high availability and ease of use, as a single IP and subnet is used for the entire team. This configuration is recommended if you use two or eight interfaces on the ProtecTIER side. By default, ProtecTIER DD5 servers are configured with eight ports that are bonded together, and SM2 servers are configured with two ports that are bonded together. If a different interface number is used (for example, all six interfaces in the case of 1 Gb NICs), this method might not maximize the performance, so dual/three teams (Layer 2+3) configuration is recommended instead. A single bond configuration is shown in Figure 5-1.

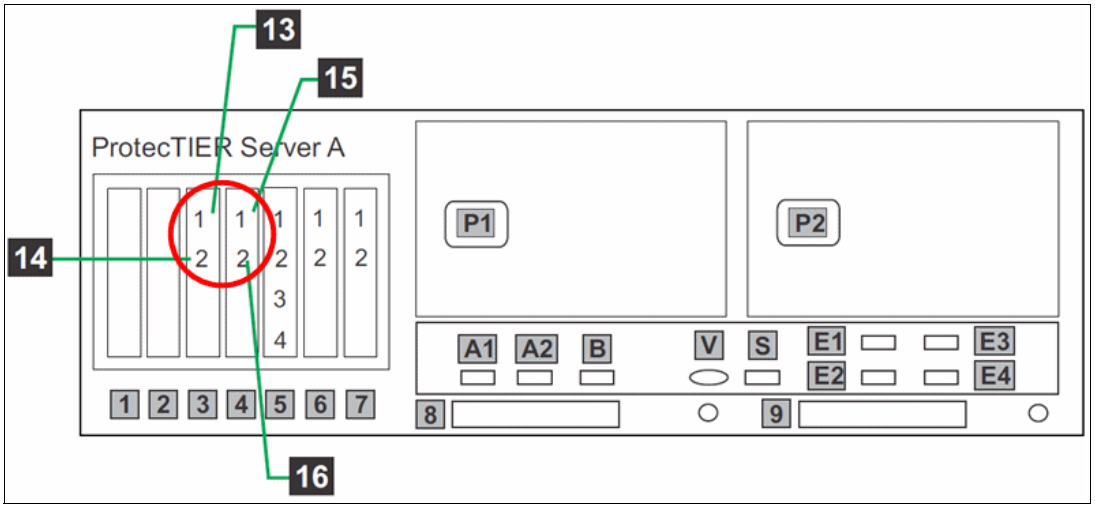

Figure 5-1 Illustration of a ProtecTIER single bond configuration

5.4.2 Individual IPs (Layer 3) configuration

In this configuration, each interface gets its own IP address. This configuration maximizes the performance but does not offer high availability of the interfaces. It also requires multiple IPs and subnets. No teaming or bonding is used in this configuration. This configuration is recommended if no switch support for the other options is available.

Figure 5-2 illustrates the individual IPs (Layer 3) configuration.

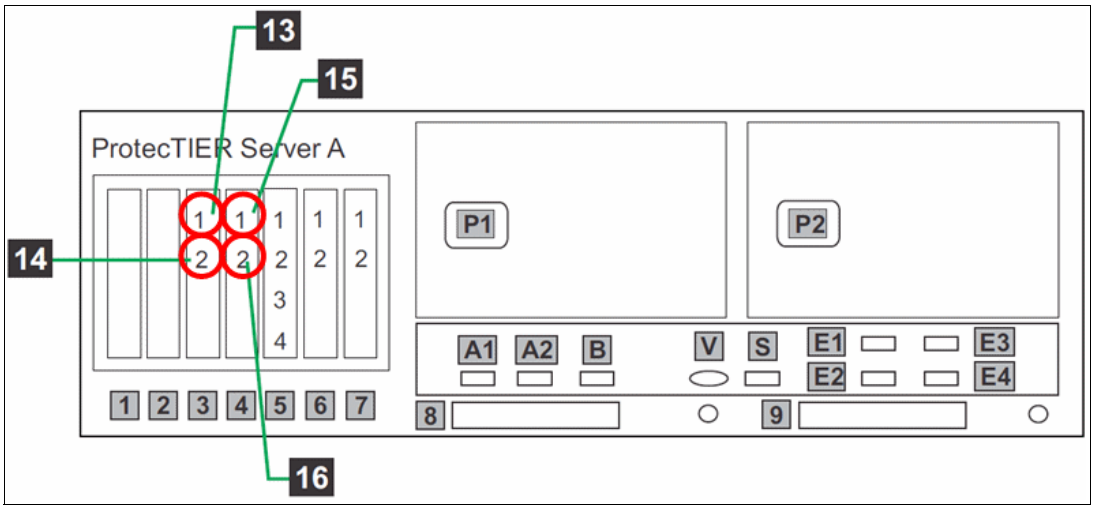

Figure 5-2 Illustration of ProtecTIER quad bond configuration

5.4.3 Dual/Three teams (Layer 2+3) configuration

In this configuration, two or three teams are configured. This configuration balances high availability, ease of use, and performance. This configuration is recommended if the number of interfaces that are used on the ProtecTIER side is different from two or eight. In testing with a Cisco 6500 switch that is configured with Etherchannel, three teams that contain two interfaces each maximized the performance. If you are configuring teams, for each one of the teams, an IP address should be configured on the team, and the team should be configured with load balancing. Figure 5-3 shows a schematic example.

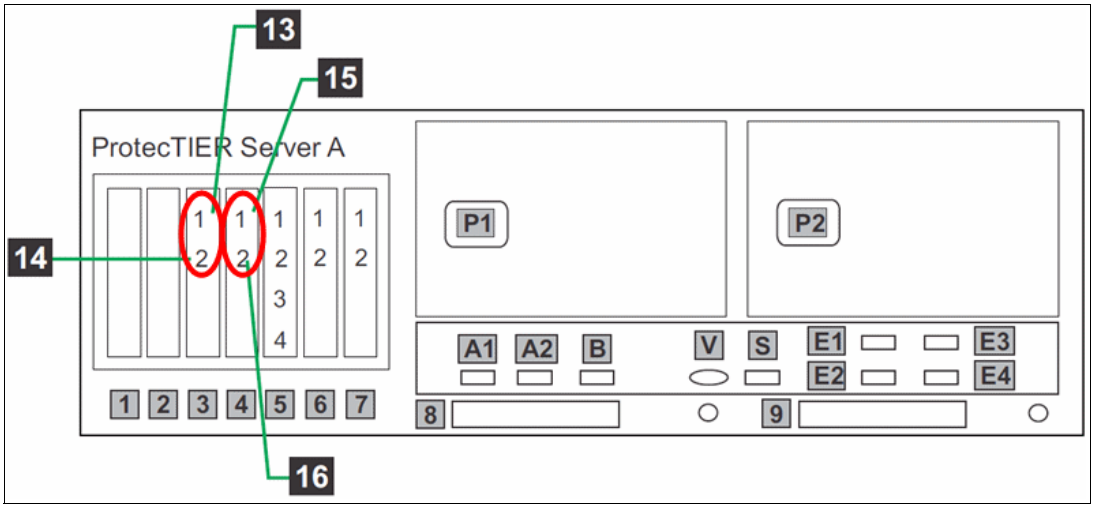

Figure 5-3 Illustration of ProtecTIER dual bond configuration

IP requirements: Whatever network methodology you decide to use, connect only one IP to each subnet. This requirement is an IP requirement, which is the protocol that we use on the Open Systems Interconnection (OSI) model Layer 3. Bonding four interfaces and assigning one IP to them to connect them to the network is suggested.

Subnetwork and VLAN separation

It is important to configure the ProtecTIER network so that each virtual interface (IP) is on a different subnetwork and preferably a different VLAN in a multitier network infrastructure. This configuration is important to segregate the backup traffic and other types of traffic for security and administrative reasons.

Selecting the bonding type and mode for file system interfaces (application interfaces)

The default setting for the application interfaces is one application virtual interface that is assigned several physical network ports, depending on the model. This interface is configured in a bond, mode 0 (round robin). You should change the mode to L3/4 (IEEE 802.3ad) when a supporting switch is available.

Switch configuration: The switch must be configured for the L3/4 (IEEE 802.3ad) mode as well. If there is no switch support for IEEE 802.3ad, or the hosts are directly connected to the ProtecTIER server, the default mode should not be changed.

Using several application interfaces for backup/restore versus using a single interface

The default setting for the application interface is one application virtual interface that has several physical network ports that are assigned. The advantage of this configuration is that only a single IP address is assigned to the backup/restore activity, and all hosts and all shares are mounted to the same IP. This configuration is the simplest one to set up.

The main problem with this setup is related to performance if the ProtecTIER server is using the 1 Gb ports (versus using the 10 Gb configurations). Although the ports are configured to share the load on the ProtecTIER server side, the hosts (even if they are part of a bond or team) do not always know to load balance the activity to get the full throughput from the ports. Load balancing mainly depends on the network cards that are installed on the hosts, and their implementation of teaming. Thus, you should perform the following activities:

- In a 1 x 1 setup (one host to one ProtecTIER server), if the ProtecTIER server is using 1 Gb ports and the performance target is more than 125 MBps, consider changing the default setup and define several application interfaces. Divide the physical ports between the interfaces, and define a different IP address and subnetwork for each IP. In this case, the host must choose to mount the shares on different IPs to benefit from them. For redundant ports, include at least a pair of ports for each application interface.
- In a Mx1 setup (many hosts to one ProtecTIER server), if the aggregate performance is important (versus the performance of a specific single host), leave the default setup as it is, except the bonding type and mode, as explained in this section.
- If the ProtecTIER server is configured with 10 Gb ports, the throughput can be satisfied by a single interface. However, if you need more than 500 MBps performance, define at least two FSI IPs on the ProtecTIER server by dividing the physical ports between the IPs. This configuration provides better throughput because the CIFS traffic flows in two different paths from the host to the ProtecTIER server.

Network setup example

As shown in Figure 5-4, the network setup requires thoughtful planning. In this example with five subnets, we took all the necessary steps to separate the management network from all other networks. We use two dedicated subnets for backup traffic to the ProtecTIER server, and we have two more dedicated subnets for ProtecTIER IP replication traffic. We also use bonding for the backup traffic subnet to aggregate two interfaces into one logical link that uses only one IP.

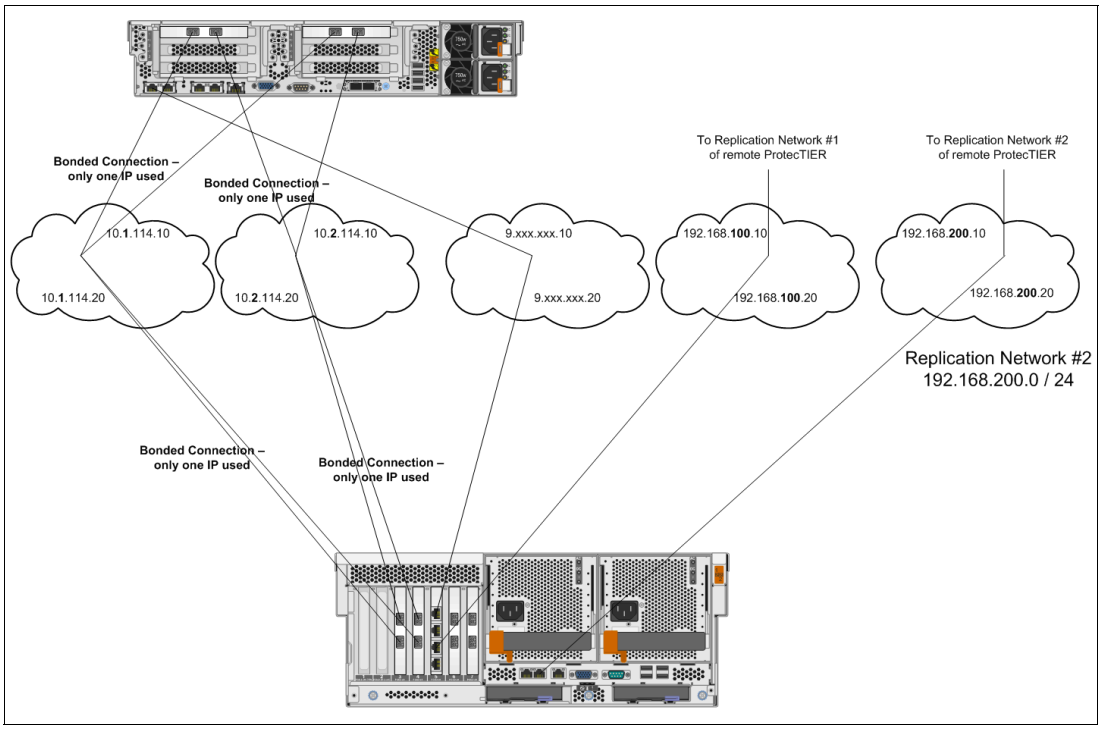

Figure 5-4 Potential network layout with five subnets

Tip: Aggregate ports from different physical adapters to avoid physical adapter failure. Aggregating the two ports from a dual port adapter does not avoid failure of the whole adapter.

If your physical network layout does not allow a similar setup, virtual private networks (VPNs) are an excellent method of segmenting traffic and achieving individual subnets that are separated from each other.

Static routes can be configured through the ProtecTIER service command-line menu to allow static routes for network segmentation. This option might be especially useful when you use an existing WAN connection together with other devices.

5.4.4 VLANs

When you connect the ProtecTIER server on a single site with the hosts, you can connect it on the same VLAN as the hosts or on separate VLANs.

As shown in Figure 5-5, a single switch topology, the ProtecTIER servers, and the hosts are connected to the same VLAN, with the same IP subnet, on the same physical switch.

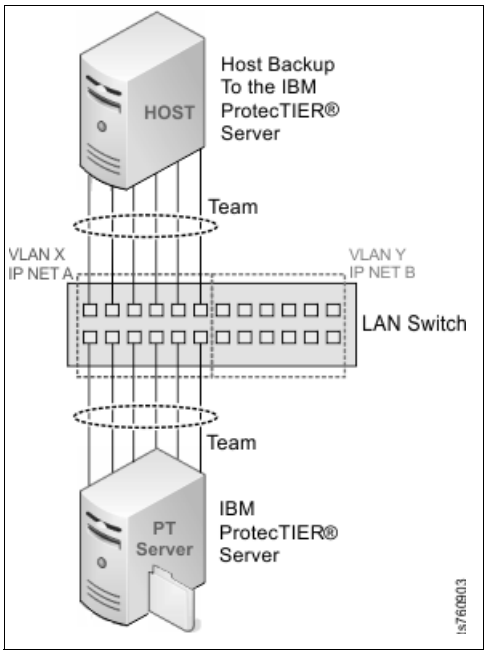

Figure 5-5 Single switch configuration

When you connect the hosts and the ProtecTIER servers on multiple LAN switches, the connectivity between the switches must be able to transfer the data rate that is required for the backup. For best results, use 10 Gb Ethernet connectivity between the switches. Another option is to define another link aggregation between the switches so that they can transfer the required bandwidth (Figure 5-6).

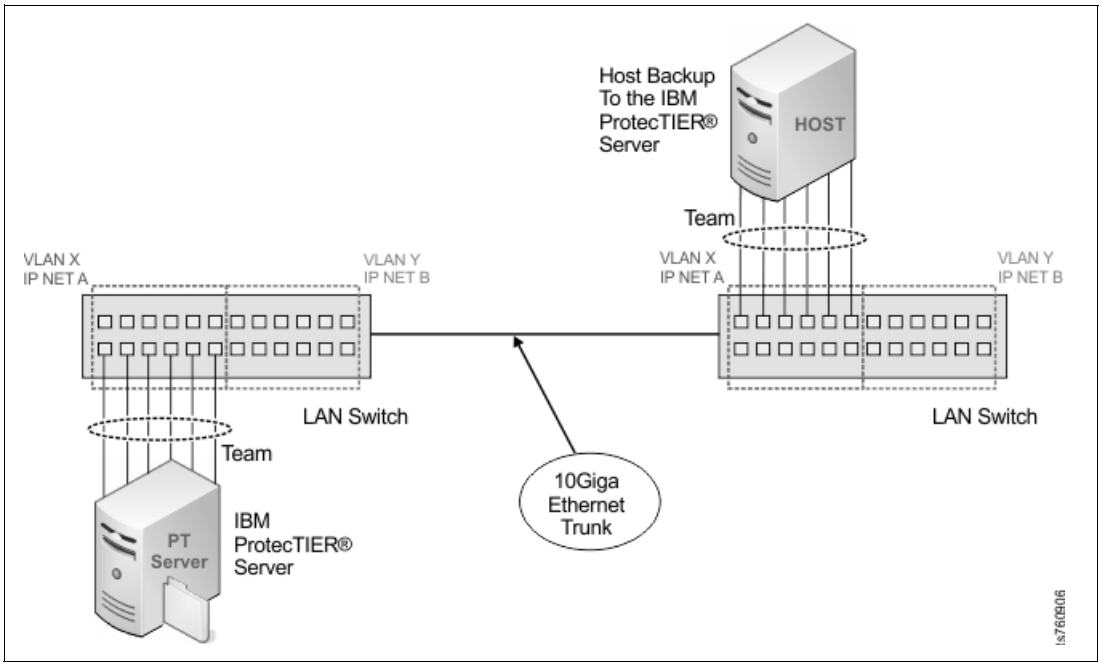

Figure 5-6 Multiple LAN switches configuration

The dual switch configuration can be used for high availability with switch redundancy (Figure 5-7).

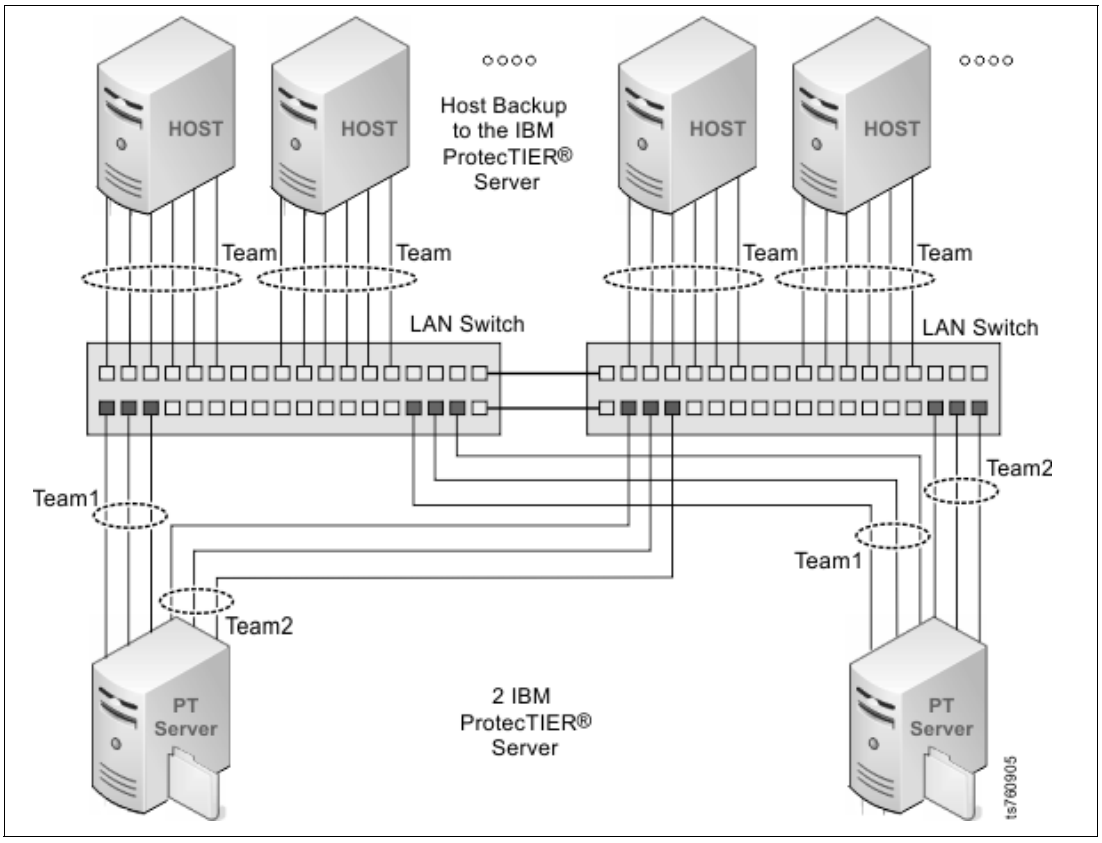

Figure 5-7 Dual switch configuration

When using different VLANs with different IP subnets, the host and the ProtecTIER server are connected on separate VLANs and subnets. The switch has Layer 3 support. Routing is performed between VLANs (Figure 5-8).

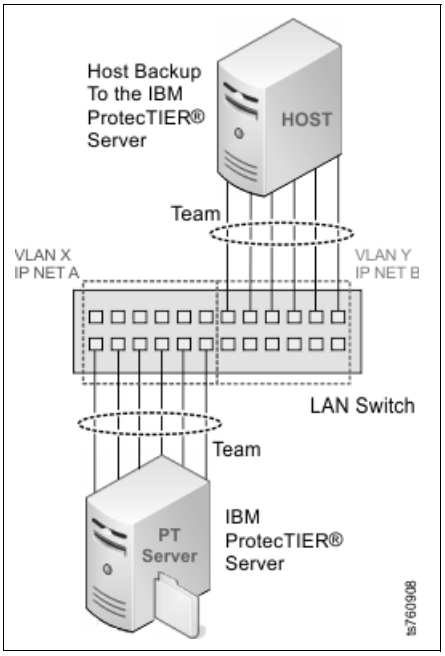

Figure 5-8 Separate VLANs and subnets configuration

5.4.5 IP addresses

You must configure unique IP addresses on the hosts and on the ProtecTIER servers if bonds are configured. If you are configuring bonds, each bond (or team) must be assigned a single IP address. Otherwise, each physical interface must be assigned a unique IP address. On each system, host or ProtecTIER, each IP address that is configured must be on a different subnet. Additional hosts and ProtecTIER servers can share the subnet. For example, on the first ProtecTIER server, you can configure the following IP addresses:

- `192.168.151.1/24`
- `192.168.152.1/24`
- `192.168.153.1/24`
- And so on

In this case, the second ProtecTIER node can use the following addresses:

- `192.168.151.2/24`
- `192.168.152.2/24`
- `192.168.153.2/24`
- And so on

In this example, the first network is `192.168.151.0`, and you can define 255 subnet addresses. Therefore, the first ProtecTIER server is using an address in this subnet (`192.168.151.1`), and the second ProtecTIER server can use a different address on the same subnet (`192.68.151.2`).

5.4.6 Routing the IP traffic

Static routes are a simple and effective way of instructing the host IP stack how to route IP traffic that is destined for specific subnets. This configuration is necessary whenever traffic to any specific subnet must be sent through a different gateway and possibly a different network-interface than the default-gateway definition would otherwise dictate. If required, configure your static routes so that each port on the host can reach one virtual port on each ProtecTIER server to which it is connected. If possible, configure all IP addresses on the media servers on the same subnets that you defined on the ProtecTIER servers.

6

OpenStorage guidelines

This chapter describes the ProtecTIER OpenStorage (OST) concepts, methods, and system components. This chapter also provides a networking overview that illustrates OST network configuration best practices to support different IP configurations.

This chapter describes the following topics:

- OST overview and main components
- Networking overview and description
- NetBackup Storage Lifecycle Policies
- OST functionality

6.1 OpenStorage overview

OpenStorage Servers are intelligent storage devices that interface with NetBackup Media Servers through the Symantec OpenStorage (OST) application programming interface (API). The vendor that supplies the Storage Server appliance provides a software plug-in, which is installed on each NetBackup Media Server that is attached to the OpenStorage Server.

With OST, ProtecTIER can be integrated with Symantec NetBackup to provide backup-to-disk without having to emulate traditional tape libraries. By using a plug-in that is installed on an OST-enabled NetBackup media server, the ProtecTIER product can implement a communication protocol that supports data transfer and control between the backup server and the ProtecTIER server.

6.1.1 Main components

The OST environment is composed of the following main components (Figure 6-1).

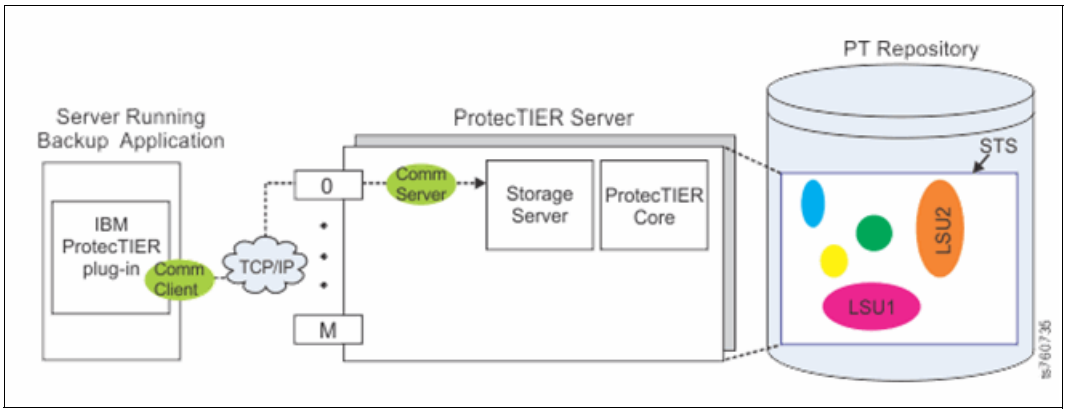

Figure 6-1 ProtecTIER components

- Storage Server (STS): The STS is an entity that runs on the ProtecTIER servers and uses the major internal ProtecTIER platform functions, such as HyperFactor, clustering, and replication.
- Logical Storage Unit (LSU): LSUs are configured on the STS to divide the system into one or more logical units of space. The LSU is the container of storage and images. The LSU storage capacity is defined as a nominal percentage of the overall nominal capacity of the repository. The LSU nominal percentage considers the configured percentages of the other LSUs that all share the physical storage of the repository.
- Plug-in: The plug-in is a shared library (a stateless software component) that is on the NetBackup host system. The plug-in is dynamically linked to the NetBackup application for data transfer to the ProtecTIER storage server emulation.

  The current version of the ProtecTIER OST Plug-in is Version 3.3, which supports the NetBackup Optimized Deduplication function.

6.2 Networking overview

In an OST environment, the ProtecTIER server connects to the network with multiple Ethernet adapters, and supports a total throughput of hundreds of megabytes per second over the network per node. This section describes basic OST network terminology.

For a more thorough description of configuration best practices to support different IP configurations, see Chapter 5, “Networking essentials” on page 79. It provides the scope and objectives, along with the networking technologies that are used to set up the ProtecTIER front-end interfaces to the network.

For a description of typical considerations and best practices for replication in an OST configuration, see 22.8, “Replication best practices for OST” on page 422.

6.2.1 Definitions and acronyms

For definitions, terms, and abbreviations that are relevant to OST network configuration, see 5.1, “Network terminology” on page 80.

6.2.2 Load distribution methodology

The ProtecTIER plug-in for OST runs on the media server and distributes the load on Layer 4. The network load balancing should also be performed on this layer to achieve the best performance over the bond group.

6.2.3 Bonding configuration

The ProtecTIER server default installation bonds all the Ethernet ports to be used as OST front-end ports together in a single bond or team. The default bonding configuration is round robin. When you use switchless topology, use the Individual IPs (Layer 3) configuration. When you use Etherchannel or 802.3ad, change the default round robin bonding configuration to Layer 3/4 hash mechanism.

6.2.4 Broadcom NICs with Microsoft platforms

If you use Ethernet teaming, the Broadcom NICs should be configured with the Broadcom network utility. From the utility, select **Link Aggregation 802.3ad** for the team type.

6.2.5 IBM AIX platforms

If you use bonding, use the AIX configuration tools to create a link aggregation of type 802.3ad. Set the hash mode to `src_dst_port`.

6.2.6 Solaris platforms

If you use bonding, use the **`dladm`** command to create a link aggregation. The LACP mode should be active and the policy should be L4.

6.2.7 Configuring the network interfaces on the host media server

On the hosts, similar to the ProtecTIER servers, each physical interface should be assigned its own IP address or grouped within a bond with a single IP address. Each interface or bond on the host side must have its own IP address on a separate subnet. As in the ProtecTIER server, different hosts can share the subnet.

6.2.8 Configuring a ProtecTIER server to work with the OpenStorage environment

For instructions about configuring a ProtecTIER server to work with the OST environment, see the *IBM System Storage TS7600 - ProtecTIER User's Guide for OpenStorage Systems, v.3.3*, GA32-2234.

6.3 Performance optimization

ProtecTIER configuration settings can be tuned in order to achieve optimal performance when you use 10 Gb Ethernet adapters. To tune the settings, complete the following steps:

1. Modify the OST `user.xml` settings:
   a. Stop the vtfd services on all nodes by running the following command:

      ```
      service vtfd shutdown
      ```

   b. Locate and edit the `user.xml` file:

      ```
      find /mnt |grep OST |grep pri |grep user.xml
      ```

      The following lines show the `user.xml` file before you edit it:

      ```
      <nAsyncThreadPoolThreads default="64" max-value="256" min-value="16"
      value="use-default">
      <nostNetThreads default="64" max-value="256" min-value="16"
      value="use-default">
      ```

      The following lines show the `user.xml` file after you edit it:

      ```
      <nAsyncThreadPoolThreads default="64" max-value="256" min-value="16"
      value="256">
      <nostNetThreads default="64" max-value="256" min-value="16" value="256">
      ```

   c. Restart vtfd on all nodes by running the following command:

      ```
      service vtfd init
      ```

2. Modify the ProtecTIER plug-in on the NetBackup Host by running the following command:

   ```
   /opt/IBM/ost_plugin_tools/ostp_cli cfg_change
   min-physical-per-logical-connections-complex 2
   ```

6.4 NetBackup Storage Lifecycle Policies

A Storage Lifecycle Policy (SLP) is a plan or map of where and how long backup data is stored. Storage Lifecycle Policies are used to provide an intuitive interface that controls the entire data lifecycle. The Storage Lifecycle Policy determines to where the backup is initially written and where it is then duplicated. It also automates the duplication process and determines how long the backup data is in each location to which it is duplicated.

SLPs offer the opportunity for users to assign a classification to the data at the policy level. A data classification represents a set of backup requirements, which makes it easier to configure backups for data with different requirements, for example, email data and financial data.

SLPs can be set up to provide staging behavior. They simplify data management by applying a prescribed behavior to all the backup images that are included in the SLP. This process allows the NetBackup administrator to use the advantages of disk-based backups in the near term. It also preserves the advantages of tape-based backups for long-term storage.

An SLP thus replaces both the duplication process and the staging process by introducing a series of storage locations or destinations that use different types of storage with different retention periods and by ensuring that data always exists at the appropriate locations at the appropriate phases of the lifecycle.

Each location can be regarded as a distinct service-level tier for Recovery Time Objective (RTO); thus, the RTO is initially a “Platinum” service level, but it degrades over time to a “Bronze” service level instead of remaining at Platinum indefinitely. As the service level degrades, the cost of storing the data decreases. This is an acceptable trade-off, as the value of backup data decreases with time. Backup data is at its most valuable immediately after the backup is made, and it is then that the RTO must be kept to a minimum. Once a more recent backup exists, the previous backup has less value because it does not offer the best RPO. As more time passes, the likelihood that a restore from the backup is required decreases, and even should a restore be required, it is unlikely to be an urgent requirement. It is therefore reasonable to allow the RTO to increase (see Figure 6-2).

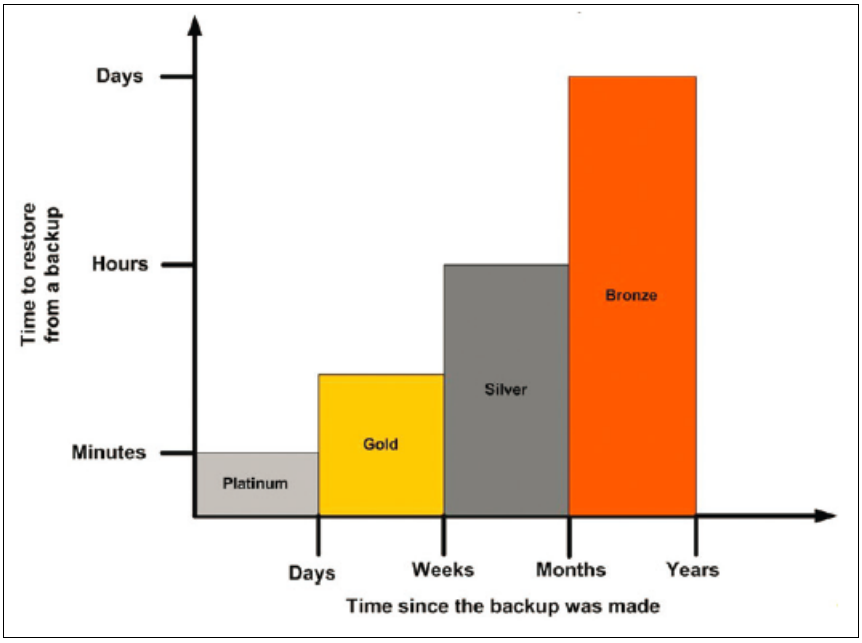

Figure 6-2 Recovery time versus the age of the backup

Figure 6-3 shows how the initial cost of storing a backup image rapidly decreases as the backup ages. The total cost of storage over the life of the backup is lower than it would be if the backup was held on Platinum storage for its entire life. Less higher-cost storage is required overall because the available storage can be reused more often with the tiered model.

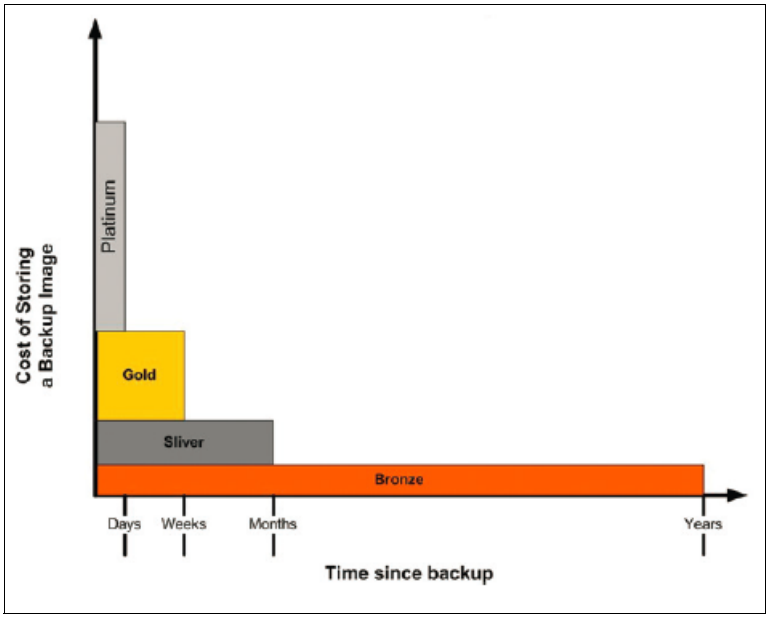

Figure 6-3 Cost of storing a backup image over time

6.5 OST functionality

The OST API allows intelligent disk devices to integrate with NetBackup. The NetBackup OST API defines several modalities for duplicating backup images. ProtecTIER supports only NetBackup Optimized Duplication and Granular Recovery Technology (GRT).

Note: Other NetBackup OST Functions are included here for completeness and are not a statement of support:

- Auto Image Replication (A.I.R.)
- Accelerator
- Optimized Synthetic Backup

For the latest OST plug-in support for OS support, OpenStorage functionality for the ProtecTIER product family, go to the following website:

`http://www.ibm.com/systems/support/storage/config/ssic/index.jsp`

6.5.1 Optimized duplication

The Symantec NetBackup OpenStorage Optimized Duplication feature takes advantage of the replication capabilities that are built into many storage servers by allowing images that are replicated by the storage servers to be registered as duplicated copies in the NetBackup catalog. This feature requires at least one NetBackup Media Server with connectivity to both the source storage server and the destination storage server. Optimized duplication cannot be performed across NetBackup domains.

ProtecTIER native replication integrates with the NetBackup OST API to provide a many-to-many replication strategy. Up to 12 ProtecTIER systems can be defined in a bidirectional OST mesh replication group. See Figure 6-4.

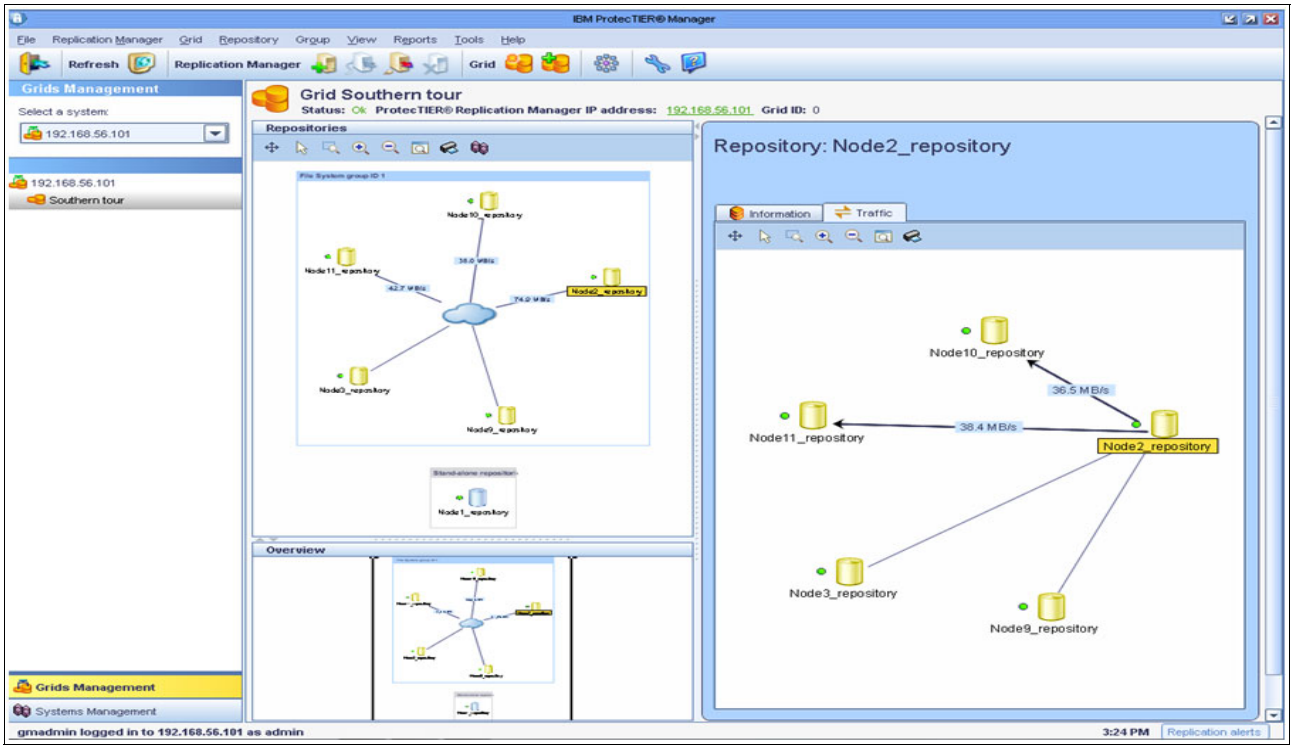

Figure 6-4 ProtecTIER native replication with OST

6.5.2 Granular Recovery Technology

Note: GRT requires ProtecTIER plug-in Version 3.3.

Symantec's patent-pending Granular Recovery Technology (GRT) eliminates the pain from the recovery process. Exclusive to Symantec Backup Exec 12, GRT enables businesses to recover critical data in mere seconds for key Microsoft applications, such as Exchange, SharePoint, and Active Directory from a single interface.

Now, any piece of Exchange data that is needed (including individual emails, folders, or mailboxes) can be quickly recovered without doing time-consuming mailbox backups. GRT simply performs a single-pass database backup, and supports local, remote, or removable backup to disk backups. It is easy to customize the technology according to your needs. For example, when it is time to make room for new backups, old backups can be archived or offloaded to tape. Also, the Exchange administrator can browse a backup set just like browsing Exchange and select the items to restore, even down to the individual email.

6.5.3 Auto Image Replication

A NetBackup domain is a collection of NetBackup clients and media servers that are controlled by a single NetBackup master server. Before NetBackup 7.1, SLPs work exclusively within the confines of a single NetBackup domain. Starting with NetBackup 7.1, SLPs support the duplication of images from one NetBackup domain to another. This feature is known as Auto Image Replication (A.I.R.).

A.I.R. requires compatible OpenStorage devices (including NetBackup deduplicating devices) to transfer the data between the source and target NetBackup domains because it uses the unique capabilities of the OpenStorage API to make the duplication process efficient. OpenStorage devices need a plug-in that supports the A.I.R. feature to use it.

6.5.4 Accelerator

NetBackup Accelerator facilitates intelligent and streamlined backups to disk by increasing the speed of full backups. The increase in speed is made possible by change detection techniques on the client. The client uses the change detection techniques and the client's current file system to identify changes that occurred since the last backup. The client sends the changed data to the media server in a more efficient backup stream. The media server combines the changed data with the rest of the client's data that is stored in previous backups to create a new full backup image without needing to transfer all the client data.

6.5.5 Optimized Synthetic Backup

NetBackup Optimized Synthetic Backup is a feature where a full backup can be synthesized on storage server by using previous full and subsequent incremental backups without reading those component images and writing a new image. This technology has been in NetBackup since Version 6.5.4. It is available on all NetBackup appliances, Media Server Deduplication Pool, and PureDisk Deduplication Option Pool. Recently, some of the OpenStorage partners also have announced support for this feature.

7

Host attachment considerations for VTL

This chapter describes the best practices for connecting hosts to the ProtecTIER Virtual Tape Library (VTL), including device driver specifications for various operating system platforms, such as AIX, UNIX/Linux, and Solaris. This chapter also describes the recommended settings for LUN masking, persistent device name binding, and considerations about control path failover (CPF) and data path failover (DPF).

Note: This chapter applies only when you use the VTL emulation feature of ProtecTIER. It does not apply to OST or FSI.

This chapter describes the following topics:

- General recommendations for connecting any backup application host to a ProtecTIER VTL system
- Device driver specifications
- Control path failover and data path failover guidance
- Persistent device name binding
- LUN masking for VTL systems

7.1 General recommendations

When you use the VTL emulation of the ProtecTIER system, there are several general recommendations that you can follow to connect any backup application host to the ProtecTIER system:

- Ensure that the operating system (OS) platform and version of your host, and your backup server version, are listed as supported in the IBM Interoperability Matrix and in the Backup Application ISV Support Matrix. For more information, see 7.2, “Device driver specifications” on page 104.
- Install the recommended device driver on the host, as specified in the IBM Interoperability Matrix. For more information, see 7.2, “Device driver specifications” on page 104.
- When possible, configure Control Path Failover (CPF) to enable redundancy to access virtual tapes robot devices. For more information, see 7.2, “Device driver specifications” on page 104.
- Set up persistent device naming to avoid changes on devices that are recognized by the operating system after a system reboot. For example, persistent naming can be configured under Linux by using the udev device manager. For more information, see 7.2, “Device driver specifications” on page 104. When you set up persistent naming, do not use SAN discovery in Tivoli Storage Manager. The Tivoli Storage Manager SAN discovery function discovers IBM devices that are based on the original operating system device name and not based on customized devices names as they are created, for example, with udev.
- When you share a virtual tape library across several backup hosts, enable the LUN masking feature and configure LUN masking groups, as described in 7.3, “LUN masking for VTL systems” on page 114.

7.2 Device driver specifications

Select the appropriate device driver, depending on the backup application, OS platform, and version of the host that you attach to the ProtecTIER server.

Ensure that your server hardware is also listed as supported in the System Storage Interoperation Center (SSIC). This interoperability matrix also provides information about the required or minimum firmware versions. The SSIC can be found at the following website:

`http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss`

Each application has a set of specific instructions to be followed. It is important to understand the requirements for each backup application and operating system version. Those requirements are listed in the IBM System Storage TS7650 / TS7650G / TS7620 ProtecTIER Deduplication Appliance / Gateway / Appliance Express - Backup Application ISV Support Matrix, which can be found, along with information about other tape devices, under the Compatibility information section at the following website:

`http://www.ibm.com/systems/storage/tape/resources.html#compatibility`

This matrix can also be found at the following link:

`http://public.dhe.ibm.com/common/ssi/ecm/en/ivl12348usen/IVL12348USEN.PDF`

Review the Notes section of the IBM ProtecTIER Backup Application ISV Support Matrix. This section specifies which device driver (either IBM Tape Device Driver or native OS driver) must be installed on the host to work with the ProtecTIER VTL. All ProtecTIER releases are listed in the same document. You must scroll down to find the section that is related to previous releases of the ProtecTIER product.

Table 7-1 summarizes which device driver should be chosen for each application. To confirm detailed information about version and specific configurations, see the latest release of the ISV Support Matrix.

Table 7-1 Summary of device drivers by each backup application

Backup application	IBM Tape Device Drivers	Native OS driver or ISV driver
IBM Tivoli Storage Manager	All platforms	N/A
Symantec Veritas NetBackup (NetBackup)	AIX with NBU 6.5.2 and later. When an IBM tape driver is used, its multipath function must be disabled.	AIX with NBU older than Version 6.5.2: Requires the Symantec ovpass driver. Solaris: Requires the Symantec sg driver for drives and solaris st driver for the robot changer devices. All other platforms.
EMC NetWorker	Windows AIX Solaris (or native)	Linux HP-UX Solaris
Commvault	Windows AIX	All other platforms: Native OS drivers
HP Data Protector	Windows	All other platforms: Native OS drivers

The following sections list the specifications grouped by OS, but always refer to the latest release of the ISV Support Matrix.

7.2.1 AIX specifications to work with VTL

The following items describe backup and recovery applications and AIX specifications that work with VTL:

- Tivoli Storage Manager backup application on all AIX OS versions requires IBM Tape Device Drivers for the TS3500 Library medium changer and for LTO3 drives.
- The EMC NetWorker (Legato) backup application on all AIX OS versions requires IBM Tape Device Drivers for the LTO3 tape drives.
- The HP Data Protector backup application requires the native OS driver for changer and drive devices.
- Symantec NetBackup (NetBackup) in Version 6.5.2 and higher uses the IBM tape driver with TS3500 Library medium changer and LTO3 drives. Earlier releases require the Symantec ovpass driver and the V-TS3500 library.
- For all other backup applications on AIX platforms, use the native SCSI pass-through driver on all existing VTL emulations.

7.2.2 Solaris specifications to work with VTL

The following items describe backup recovery applications and Solaris specifications that work with VTL:

- The Tivoli Storage Manager backup application on all Solaris platforms requires IBM Tape Device Drivers.
- The EMC NetWorker (Legato) backup application supports either the IBM Tape Device Driver or the native **st** driver.
- The HP Data Protector backup application requires a Solaris **sst** driver for the TS3500 medium-changer and the native driver for the drives.
- All other backup applications on Solaris use the native driver for all existing VTL emulations.

7.2.3 Linux specifications to work with VTL

The following items describe backup recovery applications and Linux specifications that work with VTL:

- The Tivoli Storage Manager backup application on all Linux platforms requires IBM Tape device drivers.
- The EMC NetWorker (Legato) backup application requires the native **st** driver only, and it can support up to 128 tape drives per host.
- For all other backup applications on Linux platforms, use the native SCSI pass-through driver for all existing VTL emulations.
- Implementation of control path failover (CPF) and data path failover (DPF) features are possible only with the Tivoli Storage Manager backup application on all Linux platforms.

7.2.4 Windows specifications to work with VTL

The following items describe backup recovery applications and Windows specifications that work with VTL:

- Tivoli Storage Manager, EMC NetWorker, and CommVault require IBM Tape Device Drivers.
- Symantec NetBackup (NetBackup) and all other backup applications that are not previously listed use the native Windows driver for the VTL emulations.

7.2.5 IBM Tape Device Driver

For the IBM Tape Device Driver, there is an installation and user's guide that contains detailed steps to install, upgrade, or uninstall the device driver for all supported OS platforms. The *IBM Tape Device Drivers Installation and User's Guide* can be downloaded from the following website:

`http://www.ibm.com/support/docview.wss?rs=577&uid=ssg1S7002972`

The IBM Tape Device Drivers can be downloaded from the Fix Central website. Fix Central also provides fixes and updates for your systems software, hardware, and operating system.

To download the IBM Tape Device Driver for your platform, go to the IBM Fix Central website at:

http://www.ibm.com/support/fixcentral

After you access the website, complete the following steps, as shown in Figure 7-1:

1. Click the **Product Group** drop-down menu and select **Storage Systems**.
2. Click the **Product Family** drop-down menu and select **Tape Systems**.
3. Click the **Product Type** drop-down menu and select **Tape device drivers and software**.
4. Click the **Product** drop-down menu and select **Tape device drivers**.
5. Click the **Platform** drop-down menu and select your operating system. You can select the generic form of the platform (Linux) and *all* device drivers for that platform appear.
6. Click **Continue**. In the window that opens, select the download that you need.

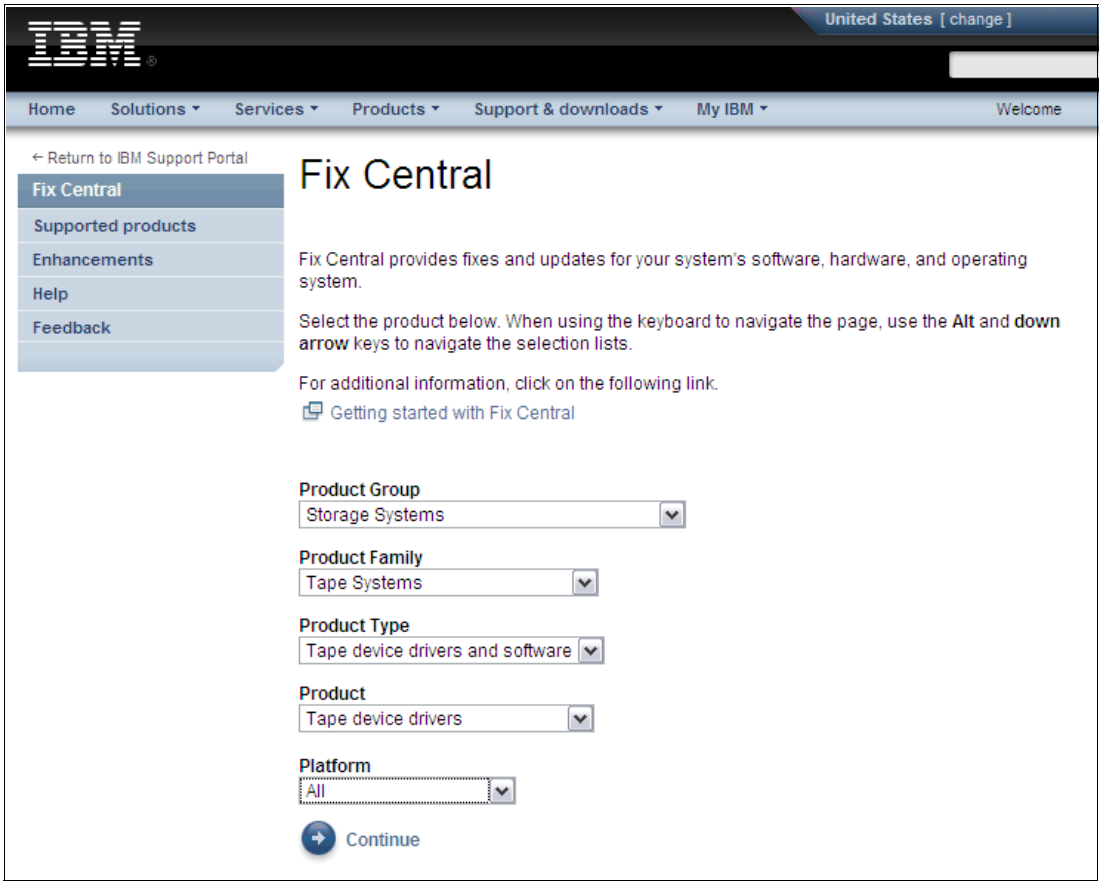

Figure 7-1 Fix Central website with the tape device driver selection

7.2.6 Control path failover and data path failover

Device driver path failover support configures multiple physical paths to the same device within the device driver. It also provides automatic failover to an alternative physical path when a permanent error occurs on one path.

The automatic failover support provides error recovery on an alternative path when a permanent error occurs on the primary path. This action is transparent to the backup application.

There are two types of path failover:

- Data path failover (DPF)
- Control path failover (CPF)

DPF is automatic failover support for the transfer of data, which provides error recovery for systems that are connected to tape drives. CPF is automatic failover support for the transfer of commands to move tape cartridges, that is, commands to the robot (media changer). For an example of control path failover, see Figure 7-2.

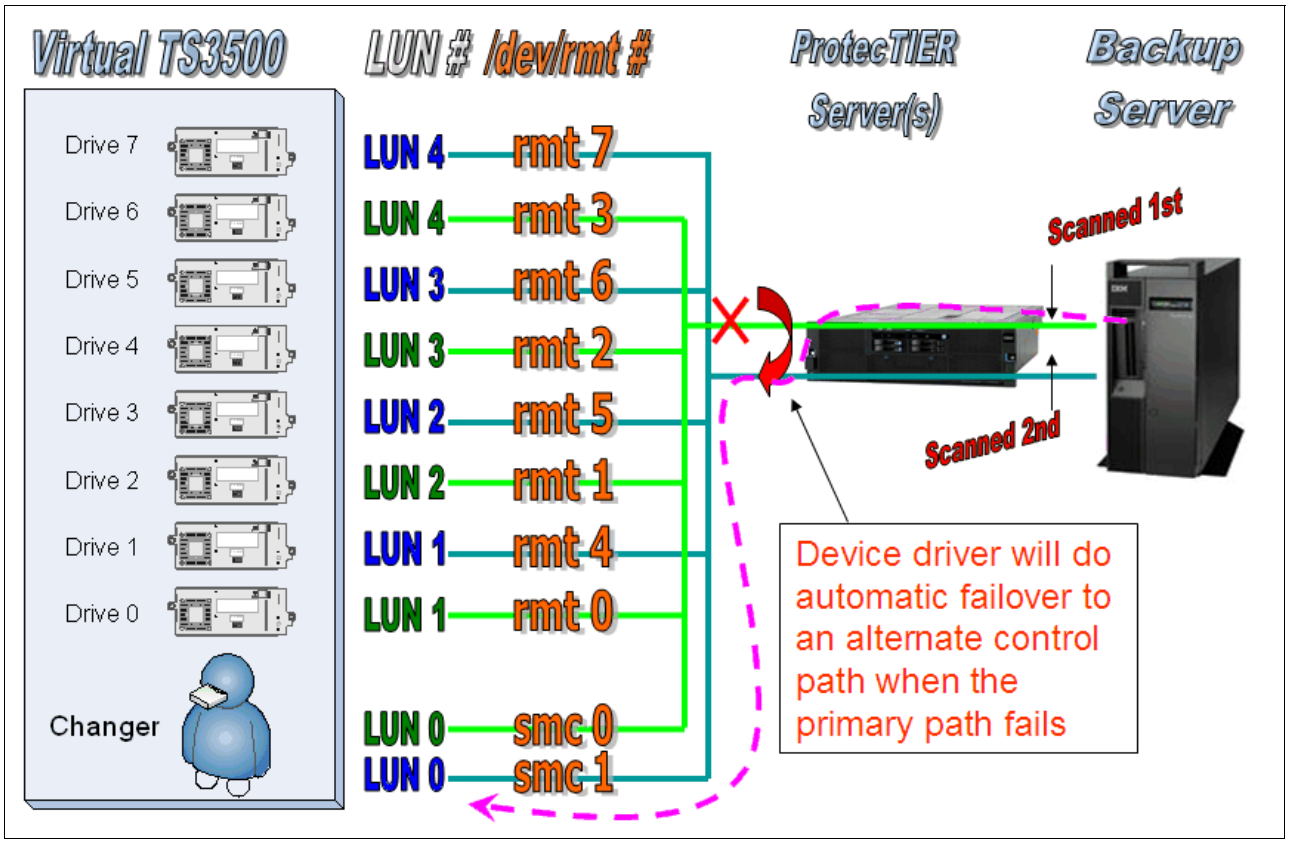

Figure 7-2 Control path failover

When path failover support is enabled on all logical devices, the device driver configures them internally as a single device with multiple paths. The application can still open and use only one logical device at a time. If an application opens the primary device and a permanent path error occurs, the device driver initiates failover error recovery automatically on an alternative path. If successful, the current operation continues on an alternative path without interrupting the application.

Data path failover (DPF) for tape drives is the same as a redundancy IBM driver mechanism for the tape drives. DPF is supported on AIX, Linux, Windows, HP UX, and Solaris. It has the following attributes:

- It functions as though it is a physical LTO-3 tape drive. There are redundant paths from the host, through the fabric, to a single FC interface at the drive.
- It requires the IBM Tape Device Driver for the OS that is hosting the backup application.
- A single virtual tape drive is still instantiated on a single front-end port at a specific LUN.
- The data path failover error recovery first restores the previous device state, SCSI Reservation, and tape position, and then tries the failing operation again.

Control path failover is a redundancy IBM driver mechanism that causes failover to occur to an alternative path if there is a path failure to the tape library (robot) control path. CPF is supported on Microsoft Windows, Solaris, Linux, HP UX, and AIX. It has the following attributes:

- The TS3500 media changer emulation enables CPF.
- It requires the IBM Tape Device Driver for the OS hosting the backup application.
- The media changer can be instantiated multiple times across multiple front-end interfaces across both nodes.
- Log entries for media changer commands contain originating path information.
- Symantec qualified NetBackup with the “V-TS3500” media changer emulation setting. Because NetBackup does not support the IBM device driver for medium changers, NetBackup relies on a built-in failover mechanism.

IBM Path Failover (CPF and DPF) features require IBM Tape Device Drivers. Table 7-2 summarizes the backup application and operating systems that support CPF and DPF. Always refer to latest release of the ISV Support Matrix for the latest information.

Table 7-2 CPF and DPF supported applications

Backup application	AIX	Solaris	Windows	Linux	HP UX
IBM Tivoli Storage Manager	CPF and DPF [a]	CPF and DPF	CPF and DPF	CPF and DPF	CPF and DPF[b]
Symantec Veritas NetBackup (NetBackup)					
EMC NetWorker	DPF	DPF	DPF		
CommVault			CPF and DPF		
HP Data Protector					

a. DPF is not supported for Tivoli Storage Manager V6.3. For Tivoli Storage Manager V6.2 and earlier, the following fscsi device attributes changes are required: `dyntrk=yes` and `fc_err_recov=fast_fail`.

b. When you use CPF and DPF, specify the following IBM Tape Device Driver parameters: **`-kctune atdd_disable_reserve=1`** and **`-Kctune atdd_reserve_type=3`**. For more information, see *IBM Tape Device Drivers Installation and User’s Guide*, GC27-2130.

In the ProtecTIER Manager, you can verify that the control path failover is enabled, which is the default, by checking the properties of a defined library in the Configuration window (Figure 7-3).

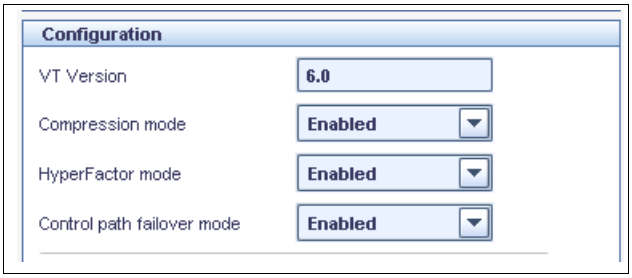

Figure 7-3 Control path failover mode enabled at ProtecTIER Manager

Enabling control path failover in Tivoli Storage Manager

To enable CPF/DPF in an AIX system with Tivoli Storage Manager, complete the following steps:

1. Ensure that the Tivoli Storage Manager option **SANDISCOVERY** is set to OFF by running the following command:

   ```
   q opt sandiscovery
   ```

   If it is not disabled, you can disable it by running the following command:

   ```
   setopt sandiscovery off
   ```

   SANDISCOVERY setting: Ensure that the Tivoli Storage Manager server setting **SANDISCOVERY** is disabled. It is not needed for CPF functionality. The ProtecTIER product exports multiple tape drives with the same WWPN, so the **SANDISCOVERY** feature does not work as expected and must be turned off.

   The **SANDISCOVERY** setting must be turned ON temporarily for Tivoli Storage Manager V6.3 to run the **libaction** command. It can then be turned off when you use CPF/DPF in Tivoli Storage Manager. For more information about the **libaction** command and the **SANDISCOVERY** setting, see 14.2.12, “Tivoli Storage Manager version with VTL” on page 227.

2. Enable path failover support on each SCSI medium changer by running the **chdev** command in AIX:
   - **chdev –l smc0 –aalt_pathing=yes**
   - **chdev –l smc1 –aalt_pathing=yes**

Primary and alternative paths

When the device driver configures a logical device with path failover support enabled, the first device that is configured always becomes the primary path.

On AIX systems, on SCSI attached devices, **-P** is appended to the location field. On Fibre attached devices, **-PRI** is appended to the location field of the device (Example 7-1 on page 111). When a second logical device is configured with path failover support enabled for the same physical device, it configures as an alternative path.

On SCSI attached devices, **-A** is appended to the location field. On Fibre attached devices, **-ALT** is appended to the location field of the device (Example 7-1 on page 111). A third logical device is also configured as an alternative path with either **-A** or **-ALT** appended, and so on. The device driver supports up to 16 physical paths for a single device.

If `smc0` is configured first, and then `smc1` is configured, the **`lsdev -Cc tape`** command output is similar to Example 7-1.

Example 7-1 Primary and alternative path example for Fibre attached devices

```
aixserver> lsdev -Cc tape | grep smc
smc0 Available 06-09-02-PRI IBM 3584 Library Medium Changer (FCP)
smc1 Available 0B-09-02-ALT IBM 3584 Library Medium Changer (FCP)
```

Configuring CPF: Detailed procedures about how to configure control path failover for AIX and other platforms can be found in 6.2, “Installing and configuring OS device drivers”, in *IBM System Storage TS7600 with ProtecTIER Version 3.3*, SG24-7968, and in the *IBM Tape Device Drivers Installation and User's Guide,* found at:

`http://www.ibm.com/support/docview.wss?rs=577&uid=ssg1S7002972`

Redundant robots with Symantec NetBackup V6.5.2

Symantec NetBackup (NetBackup) V6.0 became the first release to support multiple paths to tape drives. In Symantec NetBackup V6.5.2, the method for handling multiple robots is enhanced.

This version of NetBackup can handle multiple robot instances without the IBM tape driver because the path failover mechanism is implemented in the NetBackup software.

The V-TS3500 library type presents redundant robots to NetBackup V6.5.2, which eliminates the single robot limitation.

After you configure your storage devices (use the Configure Storage Devices wizard), only the first path that is detected by the robot is stored in the Enterprise Media Manager database.

If other paths to the Tape Library Device (TLD) robot exist, you can configure them as alternative paths by enabling multiple path support in NetBackup. Use the NetBackup **`robtest`** utility to enable and manage multiple path support for TLD robots.

If all paths fail and the robot is marked as down, then, in multiple path automatic mode, NetBackup regularly scans for the robot until it becomes available again. Automatic mode is the default. If you use multiple path manual mode, NetBackup regularly attempts to access the robot through all the paths that are configured in the multipath configuration.

To enable multiple paths for TLD robots, complete the following steps:

1. Start the **`robtest`** utility:
   - For UNIX, run **`/usr/openv/volmgr/bin/robtest`**.
   - For Windows, run **`install_path\Volmgr\bin\robtest.exe`**.
2. Select the TLD robot for which you want to enable multiple paths.
3. At the `Enter tld commands` prompt, enter the following command:

   ```
   multipath enable
   ```

When the multipath feature is enabled, it defaults to running in automatic mode. The automatic mode automatically scans for all paths for each TLD robot at each **`tldcd`** daemon start, requiring no additional setup.

CPF setup: The complete procedure to set up the control path failover with Symantec NetBackup can be found at:

http://www.symantec.com/business/support/index?page=content&id=TECH60395

7.2.7 Persistent device naming

Persistent device naming from a hardware perspective is a way of permanently assigning SCSI targets identifiers (IDs) to the same Fibre Channel logical unit numbers (LUNs). With persistent naming, these devices are discovered across system reboots, even if the device's ID on the fabric changes. Some host bus adapter (HBA) drivers have this capability built in, and some do not; therefore, you must rely on additional software for persistent binding.

From a software perspective, the device files that are associated with the Fibre Channel LUNs can be symbolically linked to the same secondary device file based on the LUN information. This setup ensures persistence upon discovery even if the device's ID on the fabric changes.

Operating systems and upper-level applications (such as backup software) typically require a static or predictable SCSI target ID for storage reliability, and persistent device naming.

An example where persistent naming is useful is a specific host that always assigns the same device name to the first tape library and drives it finds (Figure 7-4).[1]

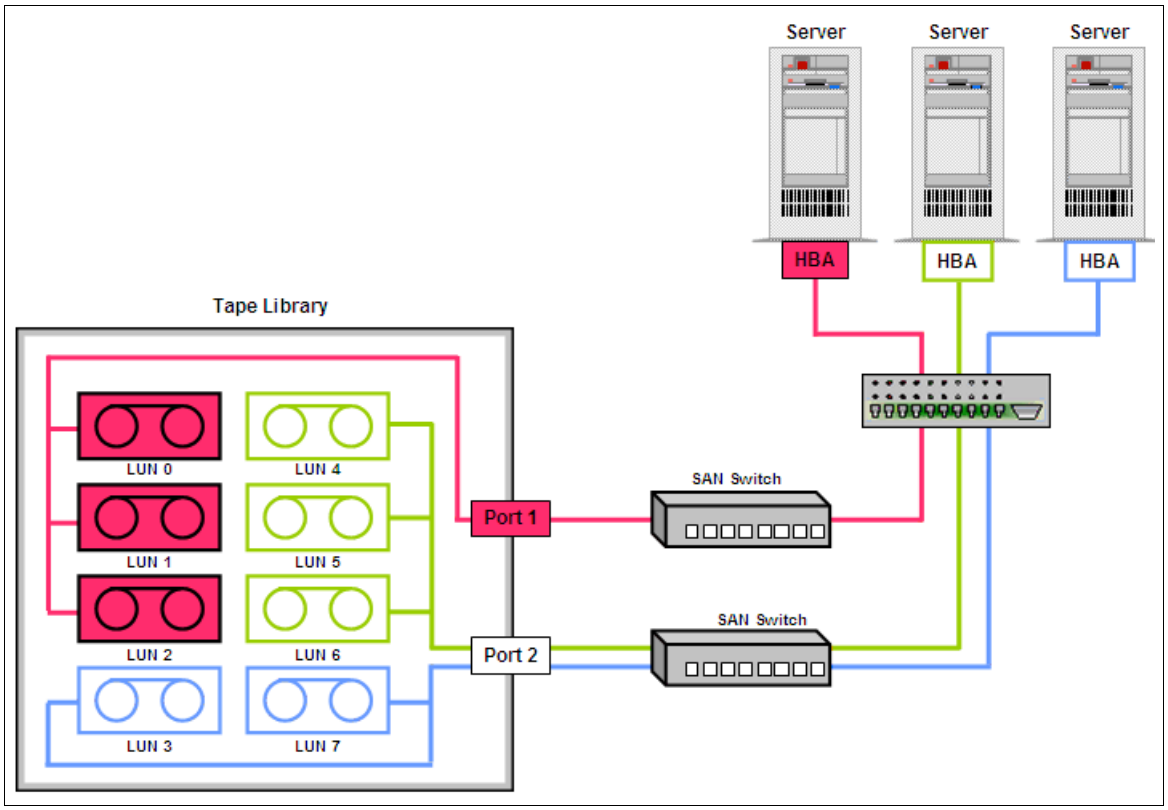

Figure 7-4 Persistent device name binding

[1] Reference: http://www.storagesearch.com/datalink-art1.html

Why persistent device naming matters

Persistent device naming support ensures that attached devices are always configured with the same logical name that is based on the SCSI ID, LUN ID, and HBA. You want to be certain that the same logical names are assigned to your device, even when the system is rebooted.

For example, when the AIX operating system is booted, the HBA performs a device discovery and assigns a default logical name to each device that is found, in sequential order.

Assume that an AIX system is connected to a tape library with two tape drives, with a LUN ID of 0 and target addresses of 0, 1, and 2. The HBA initially configures them as Available with the following logical names:

- `rmt0 target 0, lun 0 Available`
- `rmt1 target 1, lun 0 Available`
- `rmt2 target 2, lun 0 Available`

Suppose that the tape devices are deleted from the system (by running **`rmdev -dl rmt1`** and **`rmdev -dl rmt2`**) before you reboot the machine. On the next reboot, if the existing `rmt1` target 1 device is powered off or not connected, the HBA initially configures two devices as Available with the following logical names:

- `rmt0 target 0, lun 0 Available`
- `rmt1 target 2, lun 0 Available`

If the previous rmt1 target 1 device is powered on after reboot and the **`cfgmgr`** command is run, the HBA configures the device as rmt2 instead of rmt1:

```
rmt2 target 1, lun 0 Available
```

This example is a simple one. Imagine if you have a system with 200 tape drives, and with every system reboot, each device is assigned a different name. This situation could cause extra work for a system administrator to correctly reconfigure all the devices after each reboot or device reconfiguration, such as changing the characteristics of a virtual tape library.

For applications that need a consistent naming convention for all attached devices, use persistent device naming support by defining a unique logical name (other than the AIX default names) that is associated with the specific SCSI ID, LUN ID, and HBA that the device is connected to.

In AIX, you can change the logical name of a device by running the **`chdev`** command. For example, to change the logical name of the device rmt1 to rmt-1, run the following command:

```
chdev –l rmt1 –a new_name=rmt-1
```

This command allows the system to understand that rmt-1 is not detected by the HBA but is predefined at the SCSI ID and LUN ID. The rmt-1 device remains in the defined state and is not configured for use, but the next rmt-2 tape drive is configured with the same name at the same location after reboot.

Path failover: When path failover is enabled, if you change the logical name for either a primary or alternative device, only the individual device name changes.

Detailed procedures about how to configure persistent device naming for AIX and other platforms can be found in the *IBM Tape Device Drivers Installation and User's Guide*, found at:

`http://www.ibm.com/support/docview.wss?rs=577&uid=ssg1S7002972`

7.3 LUN masking for VTL systems

Administrators can manage the visibility of specific devices to specific hosts within the IBM ProtecTIER environment. This ability is called *LUN masking*.

LUN masking allows specific devices (such as tape drives or robots) to be seen by only a select group of host initiators. You can use this feature to assign specific drives to a specific host that runs backup application modules. It enables multiple initiators to share the target Fibre Channel (FC) port without having conflicts on the devices that are being emulated.

The LUN masking setup can be monitored and modified at any time during system operation. Every modification to LUN masking in a ProtecTIER server that might affect the host configuration requires rescanning by the host systems. By default, LUN masking is disabled. Without LUN masking, all of the devices in the environment are visible to all of the Fibre Channel attached hosts within the fabric if SAN zoning is set up accordingly. When you enable LUN masking, no LUNs are assigned to any backup host and the user must create LUN masking groups and associate them with the backup hosts.

The example that is shown in Figure 7-5 shows the management of a ProtecTIER environment. The ProtecTIER system includes several devices, such as tape drives and robots. Each device is assigned a LUN ID. The administrator manages two hosts and each host has two HBA ports, where each HBA port has a unique worldwide name (WWN).

A host initiator is equivalent to a host port. The host initiator uses the port's WWN for identification. By default, all the devices in the environment are visible to all the hosts. For security purposes, you must hide some of the devices from one of the ports. To accomplish this task, you must create a LUN masking group, and assign a host initiator and specific devices to that group. Performing this process ensures that the selected devices are only visible to the selected hosts.

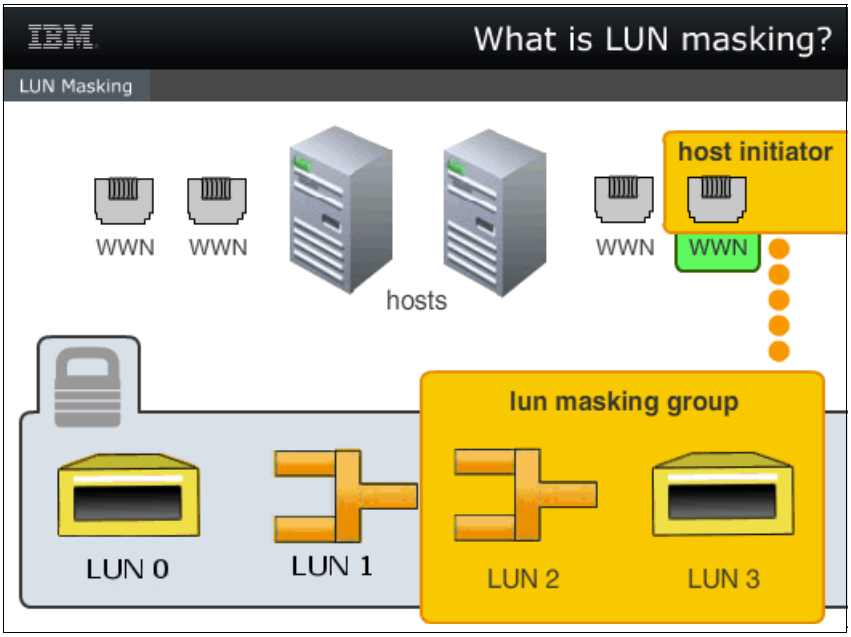

Figure 7-5 LUN masking scenario

7.3.1 LUN masking methods and best practices

Use LUN masking to manage device visibility. LUN masking conceals specific devices (tape drives or robots) from the view of host initiators while allowing a selected host initiator group to view them. Here are some best practices for LUN masking:

- Define host aliases to identify the host ports. When you define backup host aliases, use a practical naming scheme. For example:
  - `hostname-FE0` (for front-end port 0)
  - `hostname-P0` (for port 0)
- With more than two backup hosts, use LUN masking to load balance VTL performance across multiple front-end ports.
- Regardless of LUN masking, virtual drives are physically assigned to one front-end port, so backup hosts must be attached to that single port. For load balancing purposes, distribute drives across multiple front-end ports. If possible, distribute drives across all four front-end ports.
- Use LUN masking to establish two or more front-end paths to a backup server for redundancy. For example:
  - In environments with up to four backup servers, you could dedicate a single front-end port to each backup server rather than using LUN masking, but with the disadvantage of missing load balancing across multiple front-end ports and missing redundancy.
  - In environments where front-end ports are shared, and you want to prevent backup hosts from sharing, use LUN masking to isolate each backup host.

7.3.2 LUN masking configuration steps

Figure 7-6 shows the steps for a LUN masking configuration.

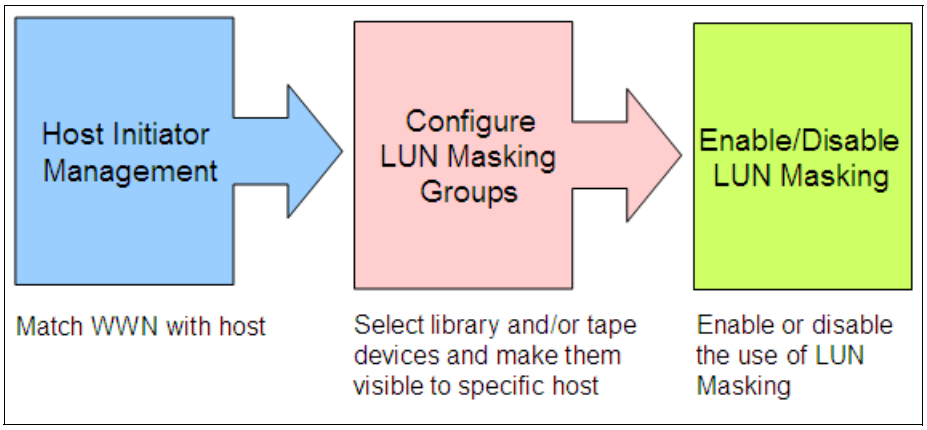

Figure 7-6 LUN masking configuration steps

Host initiator management

You must first perform host initiator management by completing the following steps:

1. From ProtecTIER Manager, click **VT** → **Host Initiator Management**. A list of available host initiators is displayed.

2. Select one or more host initiators from the list, or manually add the host initiator by entering the appropriate worldwide name (WWN), as shown in Figure 7-7.

Maximum host initiators: You can define a maximum of 1024 host initiators on a ProtecTIER system.

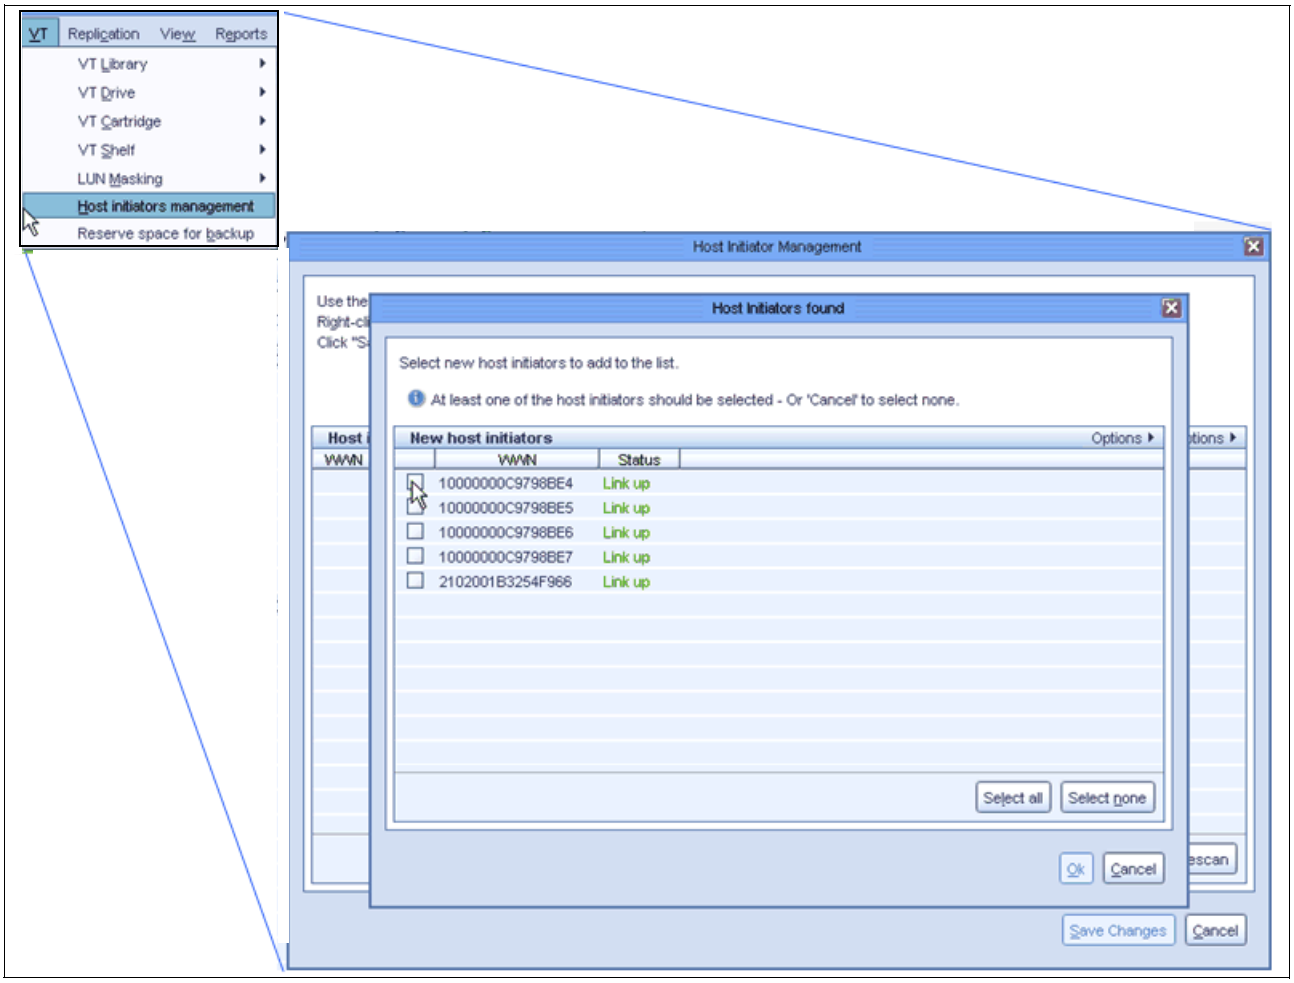

Figure 7-7 Host initiator management

3. You can also assign an alias to the WWN by clicking **Modify**. Aliases make it easier for you to identify which host is related to the WWN, as shown in Figure 7-8. The ProtecTIER worldwide port names (WWPNs) are found here.

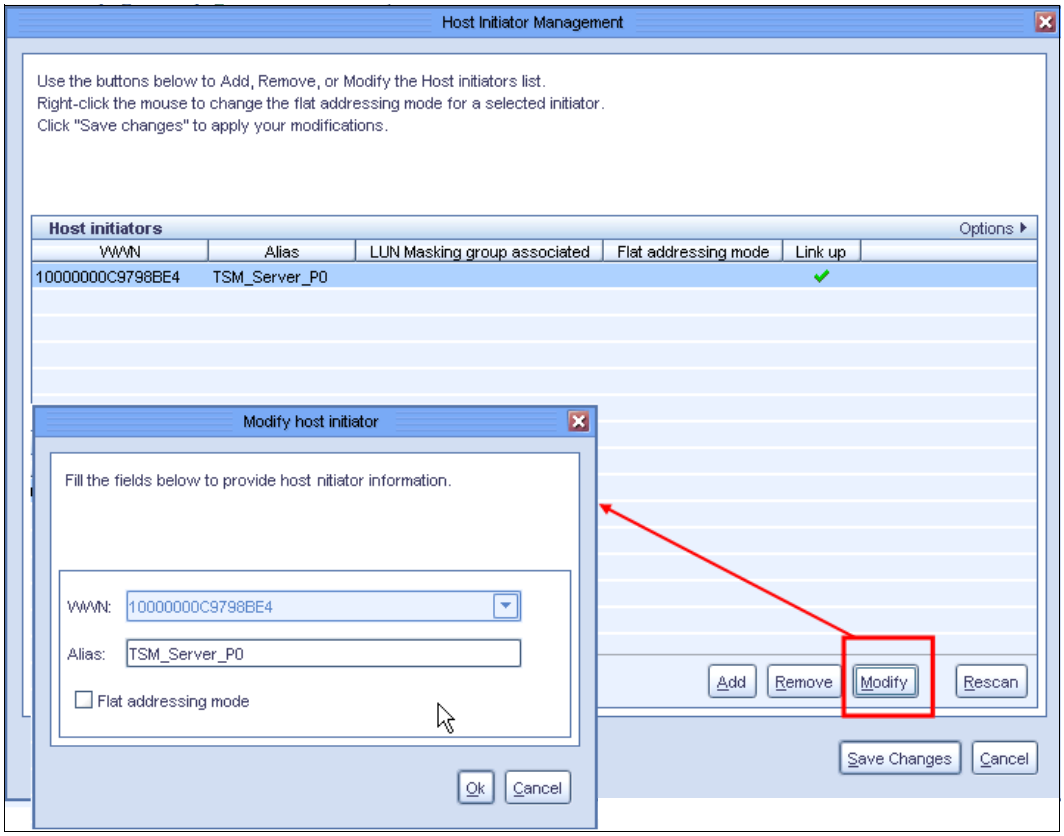

Figure 7-8 Modifying a host initiator alias

LUN masking groups

Now that you defined a host initiator, you can create a LUN masking group. Figure 7-9 shows the LUN Masking Group window.

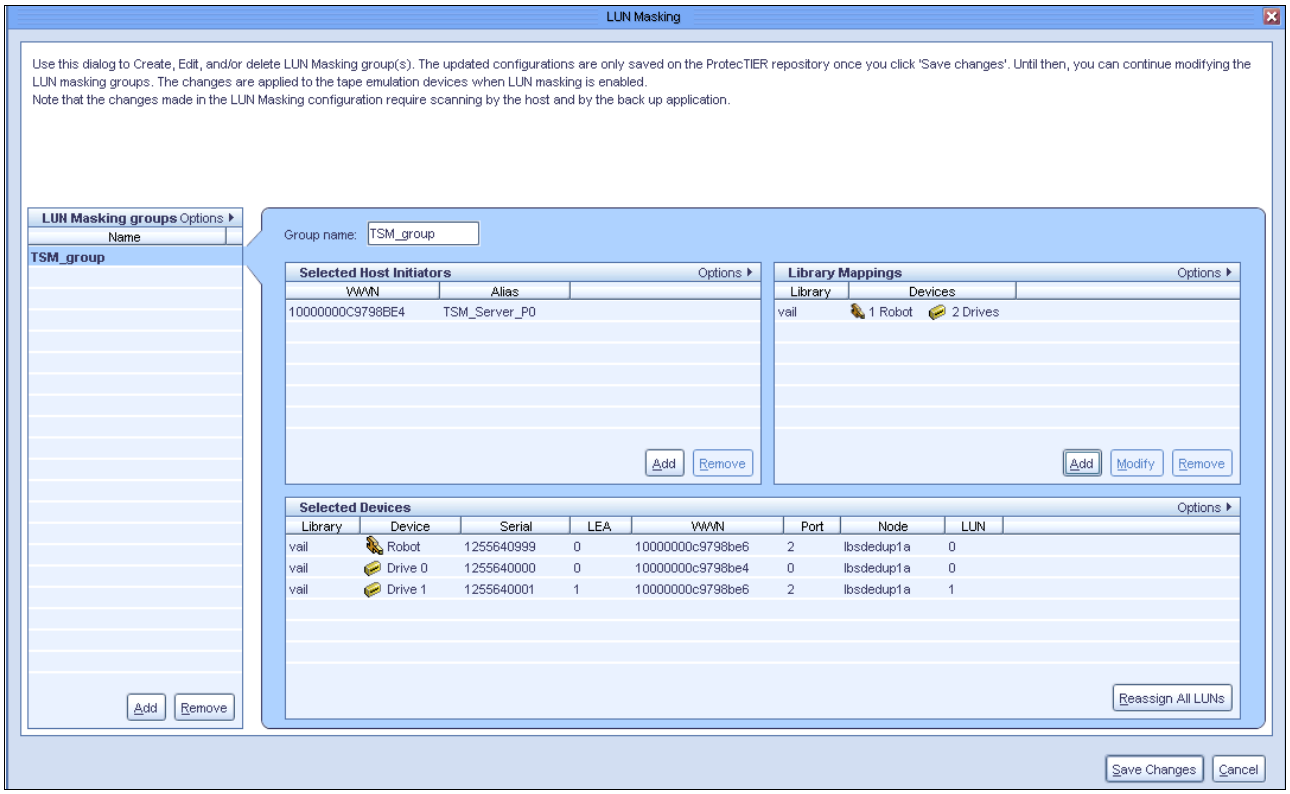

Figure 7-9 LUN Masking Group window

To create a LUN masking group, complete the following steps:

1. From ProtecTIER Manager, click **VT** → **LUN Masking** → **Configure LUN Masking Groups**.
2. At the LUN Masking Group pane, click **Add** and enter a name for the group.
3. Click **Add** in the “Selected node initiators” pane to add one or more host initiators to the group. The list of Host Initiators is displayed, and you can check the boxes of the necessary hosts.
4. Click **Add** in the “Library mappings” pane to select the library that contains the devices that you want to make visible to the hosts. Then, select the devices to assign to that group.
5. After you select all the necessary options, click **Save Changes** to create your LUN masking group.
6. If LUN masking is not enabled, the Protectier Manager asks, “LUN masking is disabled. Would you like to enable it?”. You can click **Yes** if you are ready to enable it, or **No** if you do not want to enable it yet. Even if you click **No**, the LUN masking group that you created is saved.

You can create more LUN masking groups, or you can modify an existing group for adding or removing devices, libraries, or host initiators.

Important:

- A maximum of 512 LUN masking groups can be configured per system.
- A maximum of 512 drives can be configured per LUN masking group.
- Each group must contain at least one host initiator and one device (tape drive or robot). Robots can be added as required.
- A specific host initiator can belong to one LUN masking group, but you can have multiple host initiators in a group, and multiple groups.
- A device can belong to multiple LUN masking groups, but a host initiator can belong to only one LUN masking group.

Reassigning LUNs

After you modify a LUN masking group, unwanted gaps could occur within the LUN numbering sequence.

For example, removing a device from an existing group causes gaps in the LUN numbering scheme if this device does not have the highest LUN number. As a result, the backup application might have trouble scanning the devices. If your backup application has trouble scanning the devices, you should renumber the LUN.

To reassign a LUN, complete the following steps:

1. From ProtecTIER Manager, click **VT** → **LUN Masking** → **Configure LUN Masking Groups**. The LUN Masking Group window opens.
2. Select one of the existing groups, and click **Reassign LUNs** at the bottom of the Select Devices pane.

3. The system displays the Reassign LUNs window, which has the message, “You are about to renumber all the LUN values of the available devices in the group and all host connected must be rescanned”, as shown in Figure 7-10.

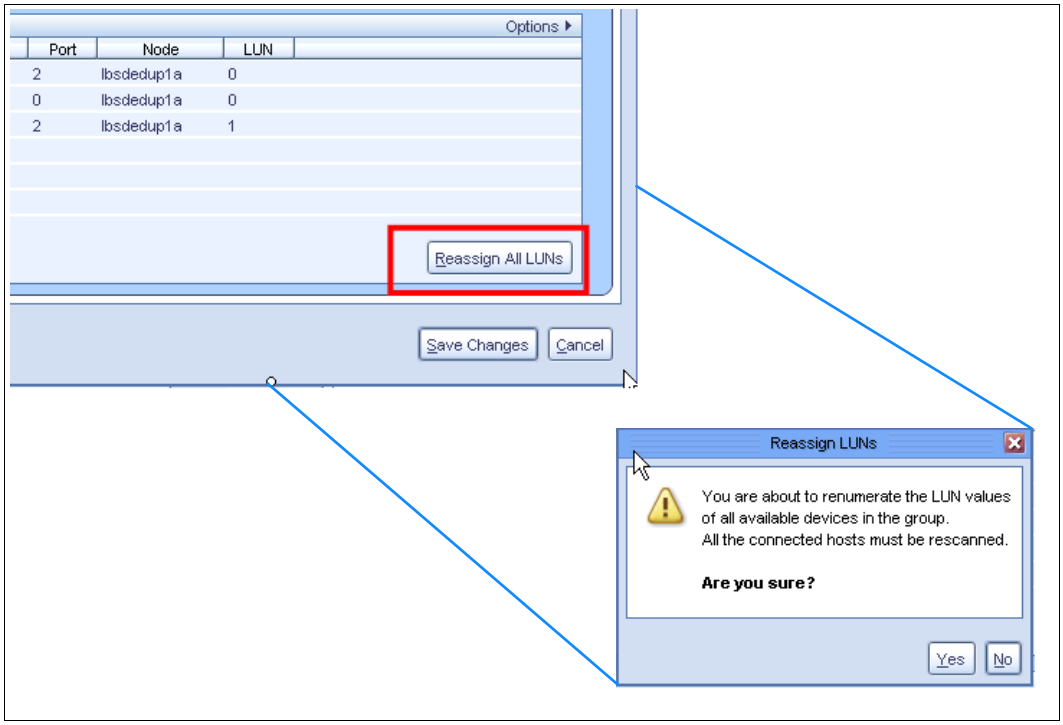

Figure 7-10 Reassigning LUNs

4. Click **Yes** to renumber. The LUN values are sequentially renumbered and all the devices in the Selected Devices pane are assigned new LUN numbers, sequentially, starting with zero.

Enabling or disabling LUN masking

LUN masking is disabled by default. When LUN masking is disabled, devices are accessible by all hosts that are zoned to the respective front-end ports. When LUN masking is enabled for the first time, all devices are masked/hidden from all hosts. You can then create LUN groups to associate host initiators with specific VTL devices, and open paths between hosts and devices. You can also enable or disable LUN masking at anytime.

To enable or disable LUN masking, complete the following steps:

1. From the ProtecTIER Manager, click **VT** → **LUN Masking** → **Enable/Disable LUN Masking**.
2. If no LUN masking groups are created, ProtecTIER Manager notifies you that if you enable the LUN masking feature without configuring LUN masking groups, the devices are hidden from the hosts. ProtecTIER Manager prompts you to confirm whether you want to proceed with this process.

3. When the Enable/Disable LUN masking window opens, select **Enable LUN masking**, and click **OK**, as shown in Figure 7-11.

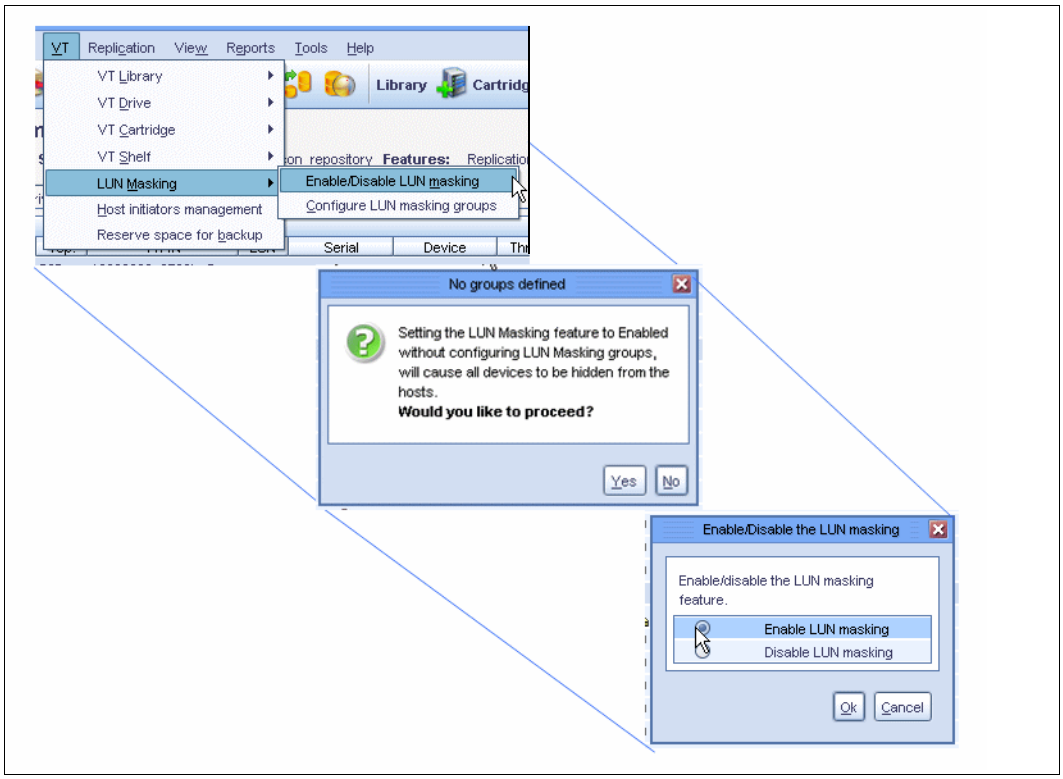

Figure 7-11 Enabling / Disabling LUN masking

You can use this same procedure to disable the LUN masking. After you enable or disable the LUN masking option, rescan the devices from the host systems. Rescanning sends the updated information for the list of visible devices and their associated LUN numbers.

Important: Every modification to LUN masking in a ProtecTIER server might affect the host configuration and might require rescanning by the hosts.

For more information about how to configure LUN masking, see 6.1.5, “LUN Masking for VTL systems”, in *IBM System Storage TS7600 with ProtecTIER Version 3.3*, SG24-7968.

Part 2

Back-end storage subsystems

This part describes ProtecTIER best practices and configuration guidelines for specific back-end storage subsystems. This section also provides a list of basic rules to follow for storage resources that are used in a ProtecTIER environment.

This part describes the following concepts and topics:

- Back-end storage overview
- IBM Storwize V3700
- IBM SAN Volume Controller, IBM Storwize V7000, and IBM Storwize V7000 Unified Storage
- IBM XIV Storage System
- IBM System Storage DS8000

Supported storage systems: For a list of IBM and third-party disk storage subsystems that are supported by the IBM System Storage TS7600 ProtecTIER deduplication servers, see the TS7650/TS7650G ISV and Interoperability Matrix, found at:

`http://www.ibm.com/systems/storage/tape/resources.html`

You can also find the most recent information about supported devices at the IBM System Storage Interoperability Center (SSIC), found at:

`http://www.ibm.com/systems/support/storage/config/ssic`

8

Back-end storage overview

Back-end storage is one of the key components of the IBM System Storage TS7650G ProtecTIER Deduplication Gateway (TS7650G) implementation. It consists of the storage subsystems that attach to the ProtecTIER gateway nodes. In contrast to the ProtecTIER appliance models, which come preconfigured with disks, the ProtecTIER gateways attaches to back-end storage subsystems that then contain the repository. This chapter lists the key factors and common configuration requirements for back-end storage that is supported by the TS7650G.

This chapter describes the following topics:

- Overview of file systems that make up the ProtecTIER repository
- ProtecTIER Performance Planner
- General rules for configuring an IBM System Storage TS7650G ProtecTIER Deduplication Gateway (TS7650G) with back-end storage
- Storage arrays configuration
- Integrating a ProtecTIER server into a SAN fabric

8.1 Overview

ProtecTIER back-end storage is a critical hardware component that holds the ProtecTIER repository. There are three types of file systems that make up the ProtecTIER repository:

- Cluster database: Single 1 GB LUN that holds the cluster configuration information. It is in a metadata RAID group and is seen as a metadata file system. This item is also called a quorum.
- Metadata: Stores all the aspects about the data that is backed up and indexed, but not the actual data. The references for restore are kept here. These file systems must be on RAID 10 arrays.
- User Data: Stores the actual data that is backed up, and referenced by new generations of backups. These file systems must be on RAID 6 or RAID 5 arrays.

Tip: The cluster database or quorum has a storage requirement that can be fulfilled by allocating 1 GB of data. You do not need to do a conversion to GiB or allocate more than 1 GB.

The file systems are created on the LUNs of storage arrays, so the construction and layout of the LUNs influence the overall ProtecTIER system performance.

Important: The file system layout is defined during the pre-sales cycle with the ProtecTIER Planner tool based on customer capacity and performance requirements. Figure 8-1 shows the Performance Planner, and Figure 8-2 shows the Metadata Planner. The outcome of a ProtecTIER planning session is used to size and configure storage arrays.

Review your storage array configurations with a trained ProtecTIER specialist before the ProtecTIER installation commences onsite. Only ProtecTIER trained specialists (IBM Pre-sales and IBM Business Partners) should do sizing and planning.

All aspects of the disk repository, including configuration, monitoring, code updates, and repair actions, are the responsibility of the client in conjunction with a disk storage vendor.

Figure 8-1 shows an example of the ProtecTIER Performance Planner. Based on the planner, to achieve 500 MBps and accommodate a 120 TB repository with a 10:1 HyperFactor ratio, the storage array should be configured with 32 x 300 GB 15,000 rpm of Fibre Channel (FC) drives in 4 x 4+4 RAID groups for metadata, and 176 x 1 TB 7,200 rpm of SATA drives for User Data in 22 x 6+2 RAID groups.

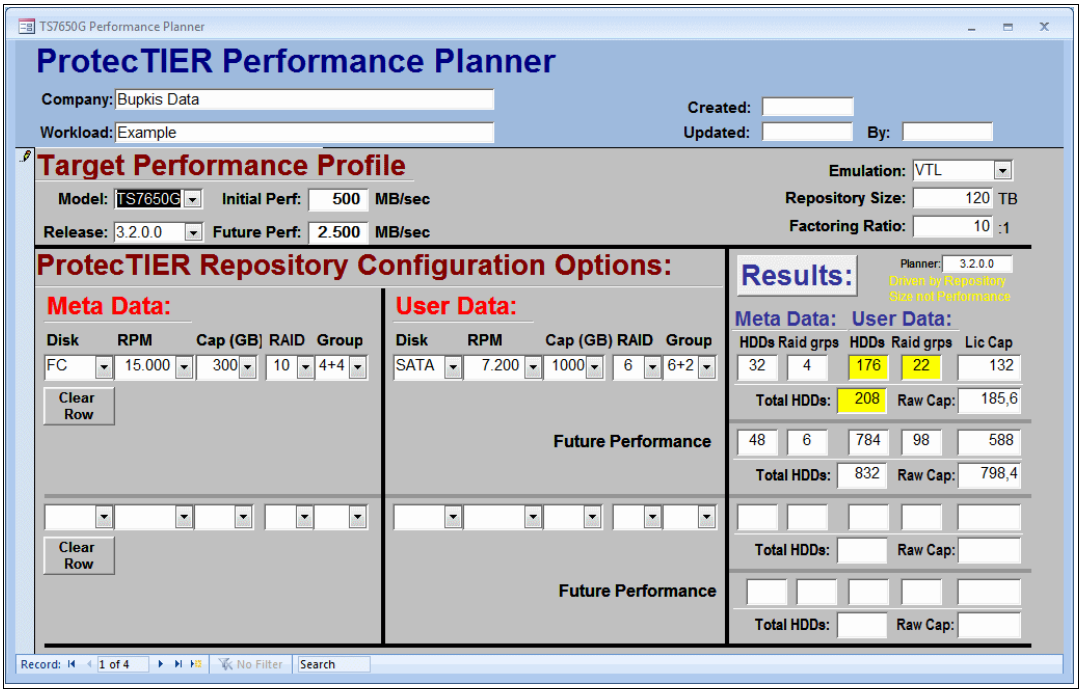

Figure 8-1 Example of ProtecTIER Performance Planner

Figure 8-2 shows an example of the ProtecTIER Metadata Planner. Based on the information that is shown in Figure 8-2, there should be five metadata file systems, including the 1 GB LUN for cluster database that is in four 4+4 RAID groups. The minimum sizes of metadata file systems are 1 GB, 886 GB, 1038 GB, 944 GB, and 365 GB.

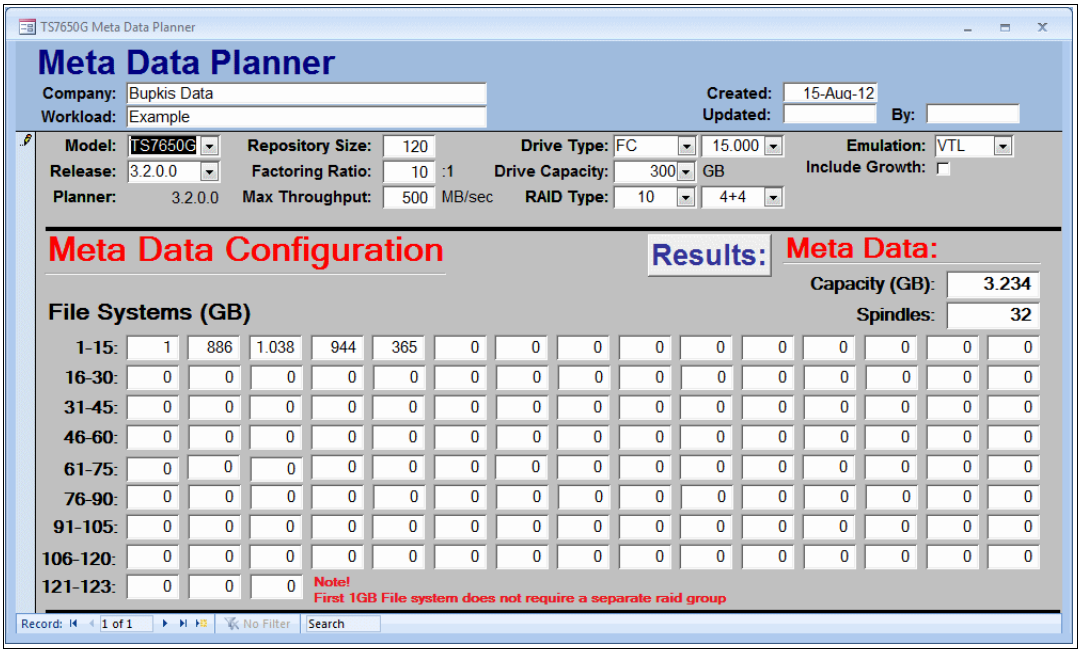

Figure 8-2 Example of ProtecTIER Metadata Planner

Important: Always assign the full array capacity to the LUNs that are used for ProtecTIER metadata and user data file systems. Follow the rule: One LUN per array and one file system per LUN. The only exception to this rule is the cluster database file system of 1 GB, which is the only file system that is allowed to be co-located on the same array with any other metadata file system.

8.2 Dependencies from a back-end storage subsystem view

Independent from the back-end storage subsystem, there are some general guidelines about how to lay out your storage for optimal usage of ProtecTIER.

Depending on whether you plan to deploy a small or large ProtecTIER gateway solution, there are certain limits and effects to be aware of.

All storage subsystems come with multiple controllers. So, you must use at least use two LUNs in order to distribute them across both controllers of a storage subsystem.

All IBM storage subsystems are based on standard hardware, for example, the DS8000 is based on the IBM System p® architecture, and the SAN Volume Controller, V7000, and V3700 are based on the IBM System x® architecture. With these architectures, the controller nodes of the storage subsystems tend to have multiple processors driving the I/O, and multiple cores per processor. To use all of these resources, you must have multiple LUNs per controller at the same time. For example, with SAN Volume Controller and V7000, you must have at least four arrays per controller to ensure the optimal usage of all computing cores.

Based on these calculations, assume that a minimum number of eight LUNs allows an optimal usage of the available resources. These are some more general assumptions; your individual system might be fine with less than eight LUNs if sizing and planning were done for your individual requirements.

8.3 Dependencies from a ProtecTIER view

The method that ProtecTIER uses to write data to disk is special. ProtecTIER writes data to disk by creating data in or appending data to the repository structure. These are the two basic modes of operation. Data that can be expired from a ProtecTIER perspective is not overwritten; it is deleted by the processes and is reshaped by the defragger practices before it is used again by write or append operations. This concept of data management directly benefits clients by having multiple file systems that it can work with at the same time. If your ProtecTIER repository allows backup speeds of 2500 MBps or more, you need at least 32 file systems in your back end for ProtecTIER to work with. If you aim for a medium performance solution, you can use fewer than 32 file systems. Again, these assumptions are more general ones, and your individual system will be fine if these assumptions are implemented according to your individual sizing and planning.

8.4 Smart storage subsystems

All the latest generation storage subsystems offer some sort of technology that adds an additional layer of “virtualization” to the distribution of data across the storage.

On V3700, V7000, and V7000 Unified, this technology is called *extent striping*. If you create a VDisk with the **`-vtype striped`** parameter (which is the default), all managed disks in the managed disk group are used to create the virtual disk. The striping is at an extent level; one extent from each managed disk in the group is used. For example, a managed disk group with 10 managed disks uses one extent from each managed disk, and then it uses the 11th extent from the first managed disk.

On DS8000, this technology is called *storage pool striping* (SPS) or *rotate extents*. The storage pool striping function stripes new volumes across all ranks of an extent pool. A similar approach is also used with V3700, V7000, and V7000 Unified.

So what is the benefit of performing striping? What benefits do extent striping, rotate extents, and so on, provide?

These features provide the following primary three enhancements:

- The striped volume layout reduces workload skew in the system without requiring manual tuning by a storage administrator (hot spots).
- This approach can increase performance with minimal operator effort (high performance).
- You can extend your existing arrays relatively easily by adding more extents to the existing pools (easy expansion).

These striping features allow you to avoid “hot spots” within your storage, to allocate more turning spindles of the storage to drive your workload, and to allow easy expansion of your environment even for existing arrays and LUNs.

Because ProtecTIER itself behaves like a storage subsystem by emulating a virtual tape library, by providing an open storage API interface, or by exporting CIFS and NFS shares like a NAS box, it has specific requirements to the back-end storage attached to it.

ProtecTIER *cannot* create hotspots. Because of the method that ProtecTIER uses to write data to disk, it is impossible for ProtecTIER to create a higher load on a single array only while others stay idle. ProtecTIER does not benefit from automatic hot spot avoidance.

ProtecTIER needs multiple LUNs to work with. Some of these LUNs are metadata, which must be RAID 10 in order to drive the heavily random write I/O; some of these LUNs are user data, which can be RAID 6 or RAID 5 in order to drive the random read I/O. ProtecTIER evenly distributes the load across the storage system and across all the available LUNs. For this process to work, you specifically must assign arrays of a certain RAID type to the specific workloads (RAID 10 for metadata, and RAID 6 or RAID 5 for user data). You must assign specific arrays by following the ProtecTIER planning process that is outlined in 21.14, “ProtecTIER Planner tool” on page 405.

ProtecTIER also uses a special way to extend the back-end storage during the ProtecTIER repository expansion process. The repository consists of metadata and user data. Depending on the amount of storage that you want to add, might extent metadata, user data, or both. Depending on how the initial deployment of ProtecTIER was done, extending the metadata might be done, along with adding storage to existing metadata LUNs, through the ptconfig menu. You could also add some new metadata file systems and leave the existing ones untouched. For the user data, because of the padding process that allocates all of the LUNs storage immediately during implementation, the repository growth includes adding new file systems. Expanding the ProtecTIER user data does not involve a logical volume manager (LVM). It does not involve growing individual file systems. To grow the ProtecTIER user data, more user data file systems must be added.

If you compare these three major differences to a classical disk workload, it is obvious that the special way ProtecTIER uses its back-end storage does not benefit from storage pool striping, rotate extents, or similar technologies.

It is theoretically possible that the usage of storage pool striping or rotate extents could collocate multiple and popular deduplication candidates within the repository on to a reduced number of spindles.

With these new features (growing disk sizes and placing one LUN on one dedicated array), you can implement storage arrays by using the various disk sizes, array layouts, and LUN sizes shown in Table 8-1. Create one dedicated LUN per array. LUNs can be up to 6 TB in size.

Table 8-1 Storage array implementation options with one LUN in a dedicated array

Disk size	Array layout	LUN size	LUNs per array	Viable?
300 GB	4+4	1200 GB	1	Yes
600 GB	4+4	2400 GB	1	Yes
900 GB	4+4	3600 GB	1	Yes
900 GB	8+8	7200 GB	1	Yes[a]
1000 GB	6+2	6000 GB	1	Yes
2000 GB	6+2	12000 GB	1	Yes[b]

Disk size	Array layout	LUN size	LUNs per array	Viable?
3000 GB	6+2	18000 GB	1	Yes[b]
2000 GB	8+2	16000 GB	1	Yes[b]

a. The LUN size can be larger than 6 TB (best practice) but must be below 8 TB.
b. LUNs that are bigger than 6 TB or 8 TB are managed by ProtecTIER starting Version 3.2.

General rule: Create only one LUN per array, even if the LUNs get bigger then 6 TB. Starting with Version 3.2, ProtecTIER manages LUNs greater than 6 TB or 8 TB in an optimal way. The only exception to this rule is the 1 GB quorum LUN, which can be collocated with any metadata LUN.

8.4.1 Rotate extents: Striping

The following section describes the **rotateexts** (rotate extents) feature and when to use in your ProtecTIER environment. The **rotateexts** feature is also referred to as storage pool striping (SPS.) In addition to the rotate volumes extent allocation method, which remains the default, the rotate extents algorithm is an additional option of the **mkfbvol** command. The rotate extents algorithm evenly distributes the extents of a single volume across all the ranks within a multirank extent pool. This algorithm provides the maximum granularity that is available on the DS8000 (that is, on the extent level, the granularity is equal to 1 GB for fixed block (FB) volumes), spreading each single volume across multiple ranks, and evenly balancing the workload within an extent pool.

Depending on the type and size of disks that you use within your DS8000 server and your planned array size for your ProtecTIER repository, you can consider using **rotateexts**. Because the ProtecTIER product equally distributes the load to the back-end disks, there are some potential scenarios where you should not use **rotateexts**.

Attention: For ProtecTIER performance, the most critical item is the number of spinning disks in the back end. The spindle count has a direct impact on the ProtecTIER performance. Sharing disk arrays between ProtecTIER and some other workload is *not* supported. This situation directly impacts your ability to reach your wanted performance.

Because you do not share disks between ProtecTIER and other workloads, assigning the full array capacity to the ProtecTIER server is recommended.

ProtecTIER prefers a high number of LUNs as back-end storage. For considerations about the potential number of arrays in the back end, see 2.5.2, “The number 32: The ProtecTIER product is not physical tape” on page 28. The recommended LUN size should not exceed 6 TB.

With these considerations, you can easily decide when to use **rotateexts** and when not to use it. Within the DS8000, the following array types should be used with ProtecTIER, taking the host spare (HS) requirements into account:

- 4+4 RAID 10
- 7+1 RAID 5 or 6+1+HS RAID 5
- 6+2 RAID 6 or 5+2+HS RAID 6

Tip: The DS8000 server creates four spares per device adapter (DA) pair. If you have a spare requirement while you create your RAID 10 arrays, you must create 3+3+2HS RAID 10 arrays. You should redesign your layout to allow all metadata arrays to be 4+4 RAID 10 arrays only. Do not create 3+3+2HS RAID 10 arrays for DS8000 repositories.

If you use a 2 TB SATA disk to create your arrays, you could have the following array dimensions:

- Creating a 7+1 RAID 5 with a 2 TB disk results in a potential LUN size of 14 TB.
- Creating a 6+1+HS RAID 5 with a 2 TB disk results in a potential LUN size of 12 TB.
- Creating a 6+2 RAID 6 with 2 TB disk results in a potential LUN size of 12 TB.
- Creating a 5+1+HS RAID 5 with 2 TB disk results in a potential LUN size of 10 TB.

All of the above LUN sizes exceed the recommended LUN size of 6 TB. In this case, you should use **`rotateexts`** to equally distribute the ProtecTIER load to the DS8000 equally across all available resources.

Important: You should not use SATA disks for metadata RAID 10 arrays. ProtecTIER metadata that is on the RAID 10 arrays has a heavily random write I/O characteristic. ProtecTIER user data that is on RAID 5 arrays has a heavily random read I/O characteristic. You should use high-performance and high-reliability enterprise-class disk for your metadata RAID 10 arrays.

If you work with a larger ProtecTIER repository, for example, bigger then physical 300 TB in the back end, the high spindle count that is combined with **`rotateexts`** means that you can use an all SATA repository.

8.5 Basic rules for a ProtecTIER server

A ProtecTIER server uses storage resources heavily. Therefore, the storage resources must be dedicated to the ProtecTIER server's activities only. The following list provides a list of basic rules for configuring a ProtecTIER server with your back-end storage:

- Disk-based replication is supported only by Request for Product Quotation (RPQ). Use the ProtecTIER native IP replication feature that is available in Version 2.5 and later. For more information about replication, see Part 5, "Replication and disaster recovery" on page 369.
- If you use SAN Point-to-Point topology to connect the TS7650G to the disk array, create dedicated zoning for ProtecTIER back-end ports. Do not mix the back-end ports (QLogic) with the front-end ports (Emulex) or any other SAN devices in the same zone.
- Dedicate the whole storage array to the TS7650G. If this configuration is not possible, make sure that the I/O requirements of ProtecTIER can be ensured and are not affected by other applications. Make sure that there is no congestion or oversubscription of resources because of other applications that might affect ProtecTIER arrays. Use zoning and LUN masking to isolate the TS7650G traffic from other applications. The TS7650G may never share RAID groups/arrays or LUNs with other applications.
- ProtecTIER creates a heavily random-read I/O pattern. About 80 - 90% of all I/O requests in a typical TS7650G environment are random reads. Because of the binary differential comparison, ProtecTIER creates this pattern even during backup traffic to ProtecTIER. The I/O pattern resembles the one for a database application. Therefore, implement suitable performance optimizations and tuning as recommended by the disk vendor for database I/O profiles.

- Because of the increased flexibility and the robustness against cabling errors, use WWPN-based zoning (soft zoning). Direct SAN attachment of backup servers is supported. Direct attachment of back-end storage subsystems is supported for V3700 and the DS8000 product family.
- RAID 10 must be used for metadata. Use RAID 6 for user data.

8.6 Storage arrays configuration

The following section describes general requirements for configuring storage arrays in a ProtecTIER environment. This section also describes some guidelines for setting up RAID groups and LUNs to get optimal performance with metadata and user data. This section then describes guidelines for placing user data on SATA disks, and expanding your repository.

8.6.1 General requirements

This section describes some general requirements for configuring storage arrays in a ProtecTIER environment, including firmware levels, host mapping, and controller support.

Storage array firmware

The storage array firmware level should be equal to or greater than the firmware version listed in the ProtecTIER interoperability matrix, found at:

http://www.ibm.com/common/ssi/cgi-bin/ssialias?subtype=WH&infotype=SA&appname=STGE_IV_IV_USEN&htmlfid=IVL12348USEN&attachment=IVL12348USEN.PDF

Storage cache

You should have 8 GB for a storage cache for every 500 MBps of planned ProtecTIER performance.

Host mapping

The ProtecTIER server host type (host connectivity settings) must be tuned for a Red Hat Linux device-mapper-multipath client. For example, you should map metadata LUNs and user data LUNs as LNXCLUSTER (which has AVT disabled) for host mapping. For more information, see the TS7650/TS7650G ISV and Interoperability Matrix at:

http://www.ibm.com/systems/storage/tape/resources.html

Storage arrays with active-active controller support

In storage arrays with active-active controller support, where a LUN can be accessed from both disk controllers simultaneously, LUNs must be mapped to both controllers for the optimum load balancing and redundancy.

Storage arrays with only active-passive support

In storage arrays with only active-passive support, where LUNs can be accessed only by one disk controller at a time, LUN mapping must be interleaved between controllers to establish effective load balancing. Always assign all LUNs in an array to have the same preferred path controller. Build the array and assign the LUNs such that:

- LUN 0 is from an array that is built using even and odd drives from the odd-numbered enclosures.
- LUN 1 is from an array that is built using even and odd drives from the even-numbered enclosures.

8.6.2 RAID considerations

It is critical to implement RAID for data protection and performance.

The RAID group size must be consistent because a smaller RAID group holds back the performance of a larger RAID group. For example, do not mix 4+1 User Data LUN with 7+1 User Data LUN. The same rules apply to metadata LUNs.

Important: Do not share RAID groups with other applications. Create only one LUN per arrays. The only exception to this rule is the cluster database / quorum of 1 GB. It can be co-located with any metadata LUN together on the same array.

Fine-tuning: RAID 5 and RAID 6 tend to cause the least performance penalty if created with even data spindles that are paired with additional parity spindles. For example, a 4+1 RAID 5 or an 8+1 RAID 5 is considered optimal. Arrays with odd data spindles tend to cause a more severe performance impact. For example, a 5+1 RAID5 or a 5+2 RAID6 is considered suboptimal.

Metadata

Here are the RAID considerations regarding metadata:

- The number of metadata RAID groups is defined with the ProtecTIER Planner Tool during the pre-sales cycle. This number ranges from 2 to more depending on repository size, the factoring ratio, and performance requirements.
- Use at least eight disk members in the group (4+4).
- Use RAID 10 with a layout that meets your planning requirements.
- Use FC drives or SAS drives for metadata, even though SATA or NL-SAS drives are used for user data.
- If needed, metadata file systems can be grown during ProtecTIER repository expansion.

Important: As the average SATA disk provides only a limited amount of IOPS in comparison to a SAS or FC spindle, the usage of SATA drives for metadata has a negative impact on performance.

User data

Here are the RAID considerations regarding user data:

- With FC drives or SAS drives, use RAID 5 with at least five disk members (4+1) per group.
- With SATA or NL-SAS drives, use RAID 6 with eight disk members (6+2) per group.
- User data file systems are padded before initial usage and therefore cannot be grown. Adding more capacity during ProtecTIER repository expansion is realized by adding user data file systems.

8.6.3 LUNs

Create only one LUN per RAID group. The only exception is for the 1 GB metadata LUN for a cluster database. It can be created on any metadata RAID group.

Important: Do not share LUNs with other applications. If possible, dedicate the storage array to the TS7650G.

Metadata

Here are guidelines for LUN creation regarding metadata:

- Metadata LUNs must be created on a different RAID group than user data LUNs.
- Create a 1 GB LUN for a cluster database on any RAID with a metadata LUN.
- The number and size of metadata LUNs are determined during the pre-installation discussion with a trained ProtecTIER specialist with the Metadata Planner (see Figure 8-2 on page 128).

User data

Here are guidelines for LUN creation regarding user data:

- As with metadata, the number of user data LUNs is determined during the pre-installation discussion with a trained ProtecTIER specialist. For optimal performance, create at least 24 LUNs.
- Keep the LUN size no greater than 6 TB for optimal performance.
- Having a repository with a number of 24 LUNs or more is considered to be optimal. With the best practice of a 6 TB LUN size, the maximum ProtecTIER configuration of 1 PB repository would require around 175 LUNs.

  Tip: Because a high number of LUNs attached to ProtecTIER increases the boot time, keep the number of LUNs at about 24 - 175 while you aim for a 6 TB LUN size.

- The size of user data LUNs must be consistent.

8.6.4 User data on SATA disks

The following statements apply specifically to the user data that is placed on SATA disks in the ProtecTIER repository:

- RAID 5 is supported, but for new systems, RAID 6 (enabling dual-parity to sustain dual disk failure) with 6+2 disk members is recommended for increased availability and faster recovery from disk failure.
- For optimal performance, create at least 24 user data RAID groups.

- The size of user data RAID groups/LUNs must be consistent. For example, do not mix 7+1 SATA user data LUNs with 3+1 SATA LUNs. Smaller disk groups hold back the performance of the larger groups and degrade the overall system throughput. For example, using storage from 2+2 or 4+1 RAID groups for the expansion might result in performance degradation because of input/output operations per second (IOPS) bottlenecks.
- When you use SATA disks, create LUNs with a maximum size of 6 TB. A bigger LUN increases the amount of metadata and impacts the performance of your ProtecTIER solution.

Important: Starting with ProtecTIER Version 3.2, the management of LUNs greater than 8 TB is improved. When ProtecTIER V3.2 works with LUNs greater than 8 TB, it splits them in to logical volumes of smaller size, which means that you can work with LUNs greater than 8 TB. There is no benefit in performance in doing this action.

You should always use RAID 6 for SATA or NL-SAS drives for the user data LUNs.

8.6.5 Expanding the repository

When you expand the repository, use the same spindle type and quality of RAID groups for metadata and user data. For example, if existing metadata LUNs were built on 4+4 RAID groups, then new metadata RAID groups must be at least 4+4. In this example, if 2+2 or 4+1 RAID groups are used, it degrades overall system performance because of an IOPS bottleneck.

8.7 Storage area network fabric

It is possible to directly connect ProtecTIER nodes to hosts (backup servers) and storage arrays. You also can connect the components into a storage area network (SAN) fabric. The connection between ProtecTIER nodes and hosts (backup servers) is referred as a front-end connection, and the connection between ProtecTIER nodes and storage arrays is referred as a back-end connection. For the updated list of supported SAN switches of ProtecTIER, see the IBM System Storage Interoperation Center (SSIC), found at:

http://www-03.ibm.com/systems/support/storage/ssic

Figure 8-3 illustrates an example of SAN fabric and zoning.

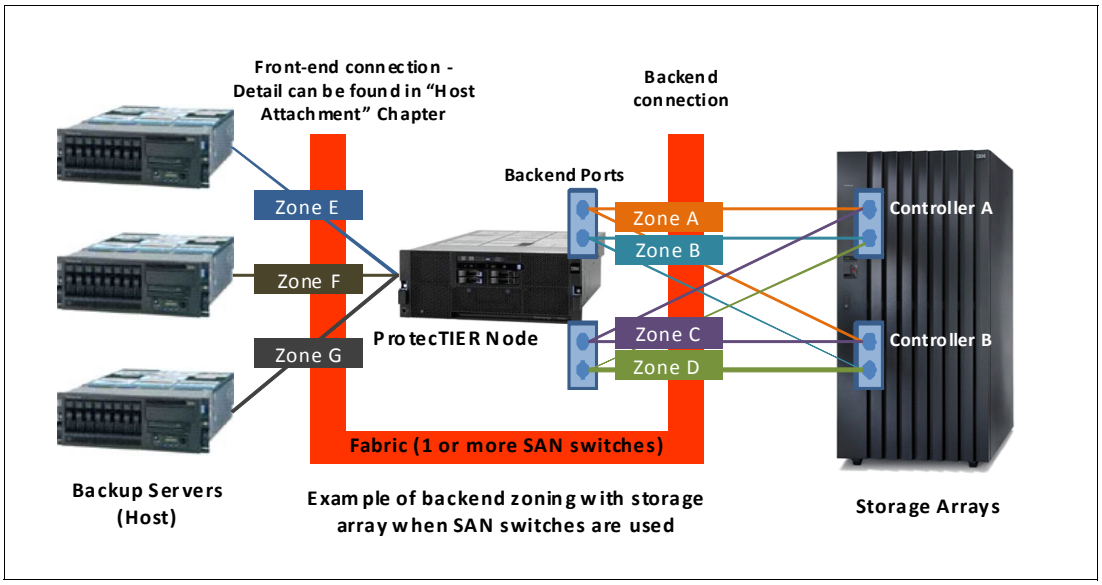

Figure 8-3 Example of SAN fabric and zoning

8.7.1 Two Fibre Channel paths to each storage controller

Each ProtecTIER node has four FC ports on the front-end FC Host Bus Adapter (HBA) for host (backup application) connectivity and four FC ports on a back-end HBA for back-end storage connectivity.

You should use two paths per storage controller.

HBAs: When you use only two back-end FC links, use separate HBAs on the TS7650G to avoid HBA hardware failure.

8.7.2 Inter-Switch Link

You should not use Inter-Switch Link (ISL) between SAN switches that connect ProtecTIER nodes and storage arrays. If ISL is used, ensure that it is not oversubscribed. Do not use ISL if IBM System Storage DS3000, IBM System Storage DS4000®, or IBM System Storage DS5000 systems are used as back-end storage.

Important: As of ProtecTIER V3.3.0, the entry-level and midrange family of disk storage subsystems for DS3000, DS4000, and DS5000 are not supported.

Current DS3000, DS4000, and DS5000 storage systems that are attached to ProtecTIER systems that run earlier releases are still supported. For best practices and guidelines for those systems, see Appendix B, "Entry-level and midrange disks" on page 457.

The list of supported entry-level and midrange disk storage subsystems can be found in the TS7650/TS7650G ISV and Interoperability Matrix, found at:

`http://www.ibm.com/systems/storage/tape/resources.html`

8.7.3 Dedicated zones

To connect a ProtecTIER node to a storage array, create dedicated zones (one zone per initiator) for the ProtecTIER back-end port. This configuration is also known as single-initiator zoning. For example, a zone that connects ProtecTIER with a DS8000 should contain one single ProtecTIER back-end FC port and all DS8000 ports that are available within that fabric.

Important: Do *not* mix the back-end ports (QLogic) with front-end ports (Emulex) or any other SAN devices in the same zone.

8.7.4 Front-end zones

Create a dedicated zone for ProtecTIER front-end ports to the host (backup application). Do *not* mix this zone with other SAN devices. It should have only one initiator per zone (single-initiator zoning). For example, one port of your Tivoli Storage Manager server is zoned to all the available ProtecTIER front-end ports. For more information about front-end connectivity to a host, see “Chapter 7, “Host attachment considerations for VTL” on page 103.

8.7.5 Back-end zones

Create one zone per initiator. For example, a dual-node cluster has eight initiators, so there are eight zones, and each zone includes the following ports:

- A single ProtecTIER back-end FC port.
- Multiple storage arrays host ports, at least one from each storage controller of the subsystem.

The zoning topology that is shown in Figure 8-3 on page 137 is one example of front-end and back-end connectivity of a ProtecTIER server. For more information about front-end connectivity requirements, see Chapter 7, “Host attachment considerations for VTL” on page 103. For more information about back-end zoning of different storage systems, see Part 3, “Backup management, VTL, OST, and FSI best practices” on page 209.

8.7.6 SAN paths

Keep the number of SAN paths to a reasonable amount. It is not helpful to have 64 paths to a storage subsystem. As a rule of thumb, do not use more than eight paths per storage subsystem, with approximately four paths per storage subsystem controller.

9

IBM Storwize V3700

This chapter describes the IBM Storwize® V3700. The Storwize V3700 is an affordable, easy to use, and self-optimizing storage solution with advanced functionality and reliability that is found only in more expensive enterprise systems. The Storwize V3700 is an entry level disk system that is based on the same software code as IBM Storwize V7000 and V7000 Unified virtualized disk systems. The Storwize V3700 is designed for the SMB market, and has the advanced capabilities that are required by small and midsize businesses, including internal virtualization, thin provisioning, data migration, and an easy to use GUI. Designed to deliver advanced functionality at a breakthrough price, this entry-level system provides an exceptional solution for workgroup storage applications.

This chapter describes the newly available IBM System Storage V3700 and how this storage virtualization product can be connected to the ProtecTIER system as back-end storage. This chapter also describes the recommended topology, volume (LUNs) configuration, and settings.

This chapter describes the following topics:

- Storage virtualization introduction and terminology
- General notes and best practices for the DS3700
- Guidelines for configuring these storage virtualization products and the ProtecTIER product in a SAN
- General recommendations for configuring metadata and user pools
- Steps to configure the ProtecTIER repository

9.1 V3700 overview

The newest member in the entry storage family, the IBM Storwize V3700 brings some of the features of the midrange storage products into the entry range at an affordable price. The V3700 can perform in iSCSI and Fibre Channel environments. It can be extended by small form factor (SFF) enclosures with 24 x 2.5-in. drives, and large form factor (LFF) enclosures with 12 x 3.5-in. drives.

IBM Storwize V3700 uses the proven technology of IBM Storwize V7000 to provide great performance and advanced features, such as virtualization, thin provisioning, copy services, and nondisruptive migration.

For more information about the storage concepts to use and the general recommendations for ProtecTIER in combination with V3700, see 10.1, “Storage virtualization introduction” on page 160 to 10.5, “User data and metadata pool: General recommendations” on page 168.

9.2 General V3700 considerations

To tailor your V3700 storage to ProtecTIER, you must use the V3700 command line to configure parts of the system. For more information about how to use the GUI to support your configuration, see 10.1, “Storage virtualization introduction” on page 160 to 10.5, “User data and metadata pool: General recommendations” on page 168.

9.3 Configuration steps

Figure 9-1 shows the steps that configure the ProtecTIER repository on a V3700.

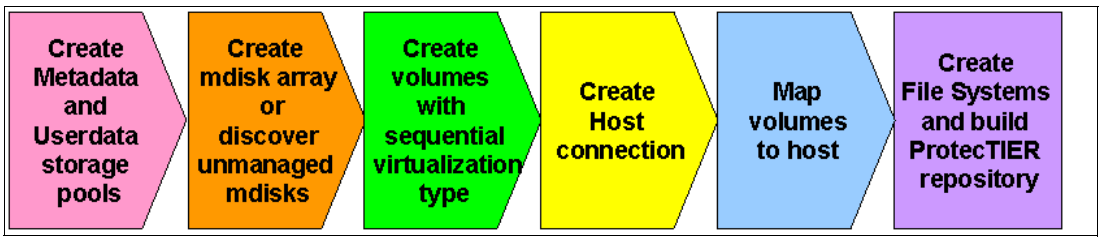

Figure 9-1 Configuration steps

Here are the steps:

1. Create empty user data and metadata storage pools.
2. Create the MDisk arrays.
3. Create volumes (VDisks) with a sequential virtualization type on the command-line interface (CLI).
4. Create a host connection for the ProtecTIER nodes.
5. Map volumes to the host.
6. Create file systems by using the File Systems Management under the **`ptconfig`** menu and build the ProtecTIER repository by using the ProtecTIER Manager.

9.3.1 Creating empty user data and metadata storage pools

You should create *only* two storage pools (also known as managed disk groups), where one is used for metadata and the other is used for user data. When you create more pools, the storage system cache is segmented among all pools. By having only two pools, you optimize the storage subsystem cache allocation.

To create these items, complete the following steps:

1. Hover your cursor over the Pools icon, and the system displays the Pools menu (Figure 9-2). Click **MDisks by Pools**.

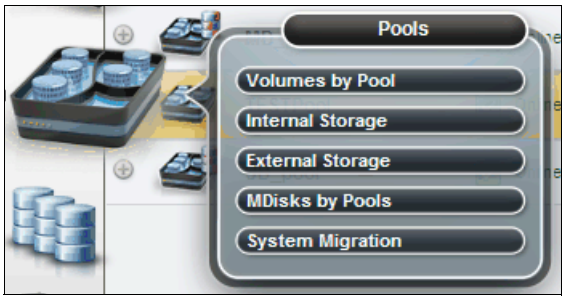

Figure 9-2 Pool menu

2. For each pool, click **New Pool** to create an empty pool. Insert a name for the pool, such as MD_pool or UD_pool. (Figure 9-3). Click **Next** to continue with the pool creation process.

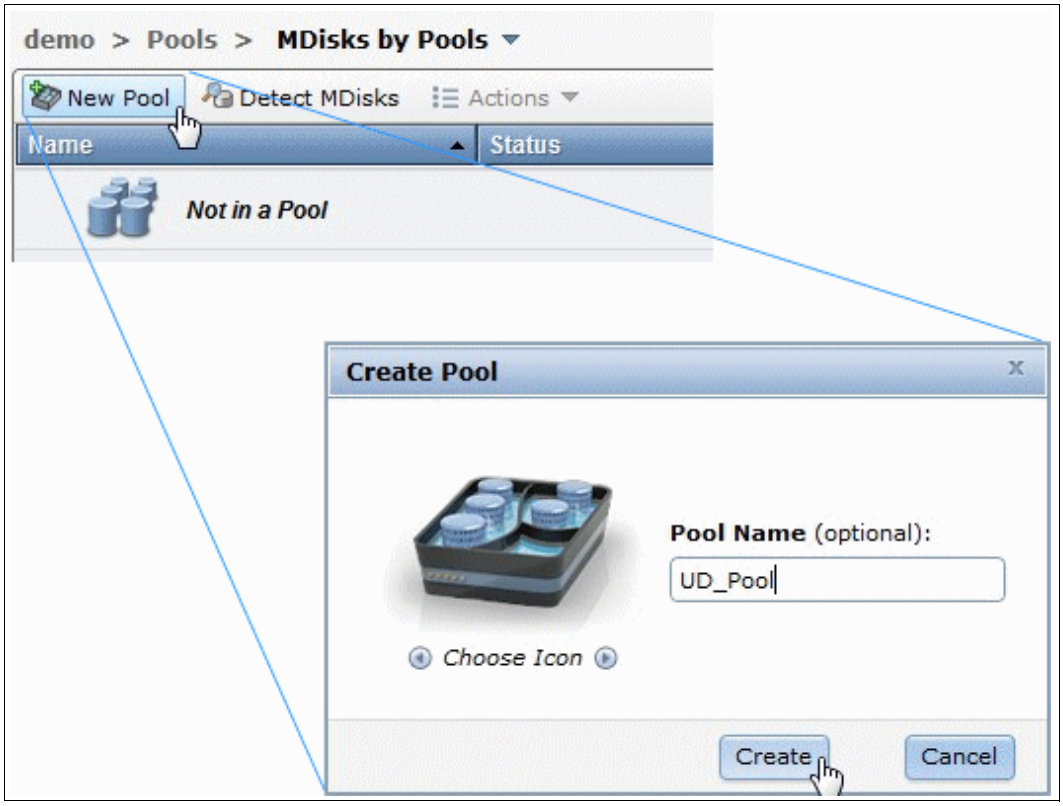

Figure 9-3 Creating pools

Tip: Set the extent size settings to 256 MB. This is the default. It is needed to allow large repository dimensions of 1 PB or more of virtualization within the storage. The default warning threshold of 80% should be changed to 100% to not receive alerts about pool usage. The pool is fully allocated by the ProtecTIER LUNs. To do this task, you must use the CLI to create the pools, as shown in Example 9-1 on page 142

3. Click **Next**. If you created your MDisk arrays or use MDisks from external storage, select the MDisks to be included in the pool. Otherwise, click **Finish** to create an empty pool. (Figure 9-4).

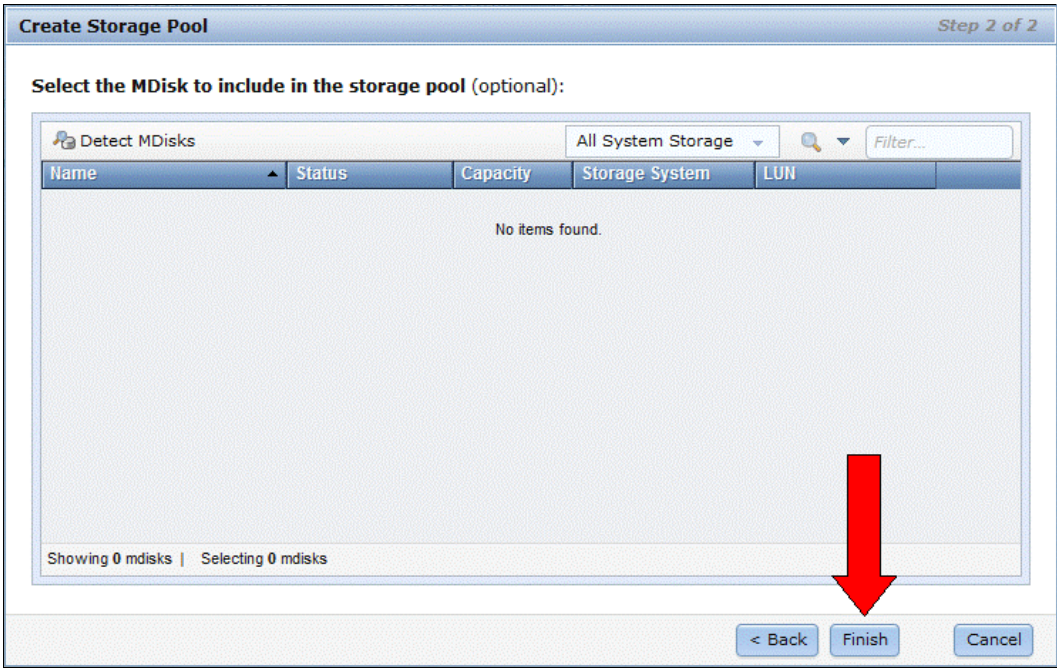

Figure 9-4 Select the MDisks to be included in the pool

4. Repeat this procedure for the next pool. You need two pools, one for metadata, one for user data.

The steps can also be done by using the CLI, as shown in Example 9-1.

Example 9-1 Creating a pool using the CLI

```
# svctask mkmdiskgrp -ext 256 -warning 100% -name UD_pool
```

9.3.2 Creating the MDisk arrays or discovering unmanaged MDisks

When you use the Storwize V3700 server, you can use internal or external disks to create MDisks. The SAN Volume Controller allows only MDisks coming from external virtualized storage.

To accomplish this task, complete the following steps:

1. Hover your cursor over the Pools icon and the system displays the Pools menu (Figure 9-5).

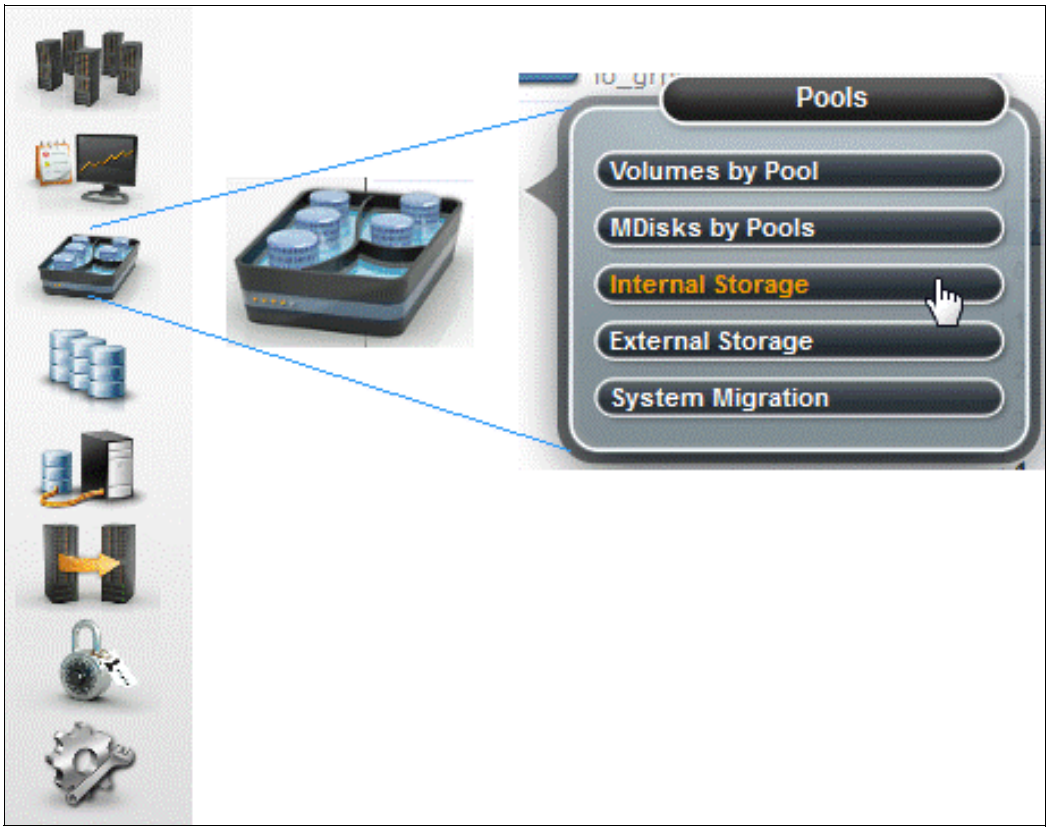

Figure 9-5 Pools menu

2. To create an MDisk with internal disks for IBM Storwize V3700, click **Internal Storage** to display the list of candidate internal disks. Select the appropriate Drive Class that you want to configure, for example, SAS disks for metadata and NL_SAS (or SATA) for user data. Then, click **Configure Storage** (Figure 9-6).

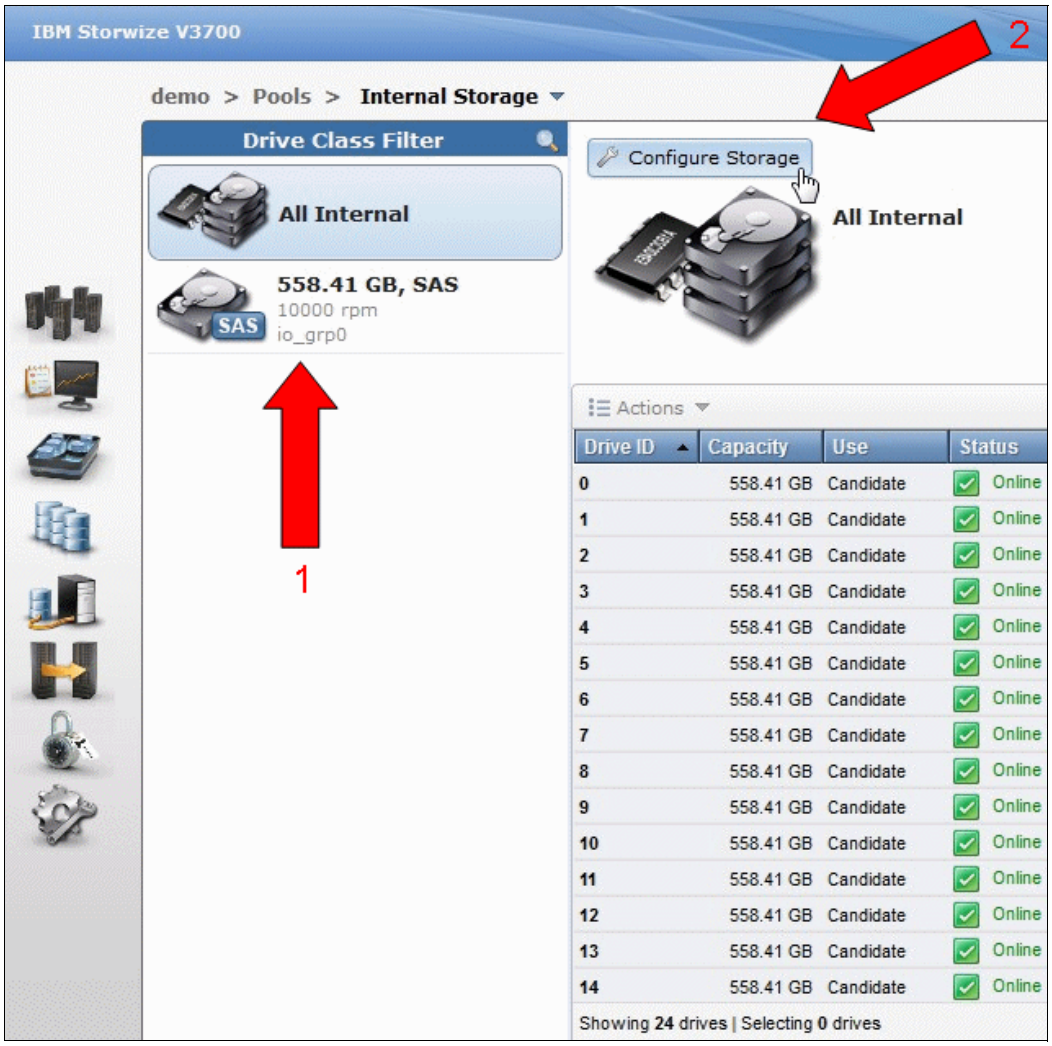

Figure 9-6 Configure Storage

3. If a SAN Volume Controller is being used, or if the MDisk is created from external storage in a Storwize V3700, click **External Storage** to display the list of unmanaged MDisks. If no MDisks are displayed, click **Detect MDisks**. When the list of available MDisks opens, select the appropriate disks, right-click, and click **Add to pool**.
4. The Configure Internal Storage window is displayed. The default settings are shown in the window. Proceed as shown in Figure 9-7 on page 145.
   a. To use different settings, click **Select a different configuration**.
   b. Select the RAID type, for example, **Balanced RAID 10** for a metadata array.
   c. Click **Optimized for Performance** to force the GUI into creating RAID arrays of equal dimensions.
   d. In the Number of drives to provision field, enter the number of disks to configure in your RAID arrays.

e. Ensure that the automatically calculated RAID layout resembles the RAID arrays that you want to create. For example, having eight disks in a Balanced RAID 10 creates a 4+4 array. Having eight disks in a RAID 6 creates a 6+2 array. Having five disks in a RAID 5 creates a 4+1 array.

> **Important:** Closely monitor the output of the GUI in step e. Especially while you try to create multiple arrays concurrently, the GUI can get a little too creative when it comes to deciding the array dimensions. Change the number of disks that you want to configure in a single step to avoid unwanted array sizes.

f. Click **Next**.

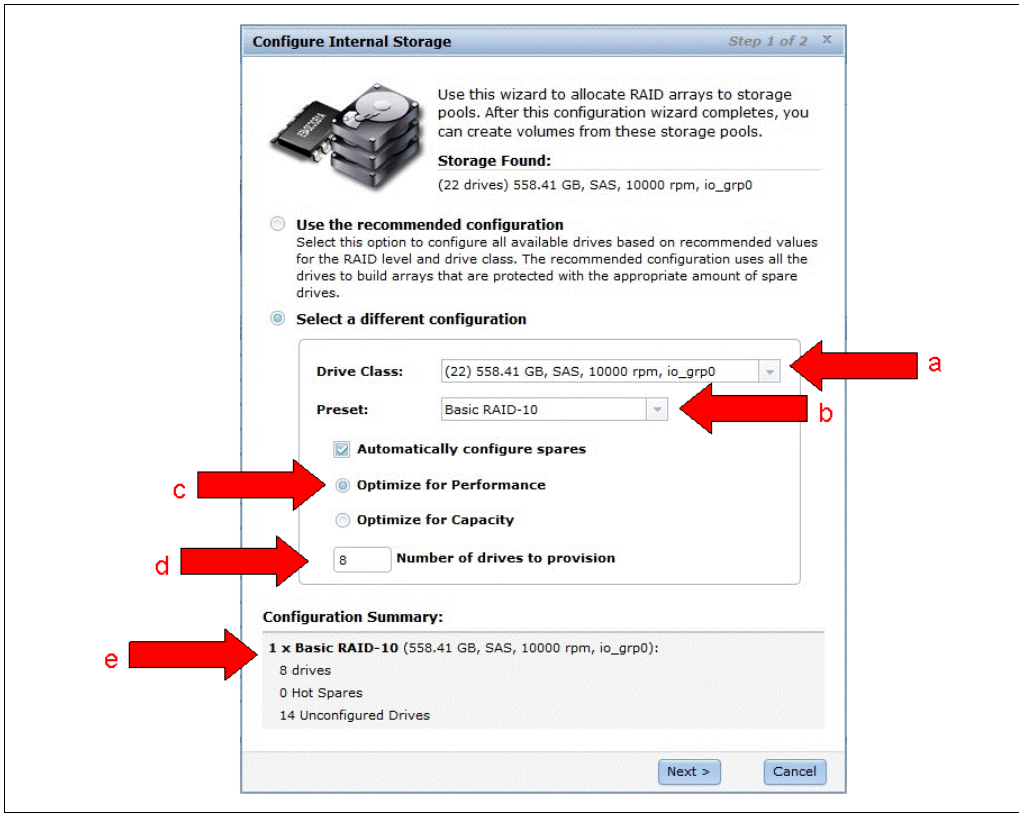

Figure 9-7 Configure Internal Storage window

These steps can also be done by using the CLI, as shown in Example 9-2.

Example 9-2 MDisk array creation command

```
# svctask mkarray -drive 28:27:26:25 -level raid5 -sparegoal 0 -strip 128 UD_pool
# svctask mkarray -drive 8:6:5:10 -level raid10 -sparegoal 0 -strip 128 MD_pool
```

The numbers after the **-drive** flag represents the physical disk's ID, which is part of the array.

Important:

The parameter **-sparegoal 0** allows array creation whether or not there are hot spare drives available. With this setting, you need to manually ensure that there are enough hot spare resources that are available within your storage subsystem to satisfy your needs.

To manually create a balanced RAID 10 array, you need to select the array members in a specific order. The **-drive** parameter uses the drives in the following order: Primary:Mirror:Primary:Mirror:Primary:Mirror, and so on. Make sure that each Primary disk is attached through one SAS chain and each Mirror disk is attached through the other SAS chain.

9.3.3 Creating volumes with a sequential virtualization type

The ProtecTIER architecture stripes the repository among all file systems and uses all the disks simultaneously. With this architecture, create the volume to accommodate the full size of each MDisk. Additionally, because of the nature of the ProtecTIER file system architecture and workload, *sequential* virtualization has better performance than striped virtualization.

To create volumes (VDisks) with the sequential virtualization type, you must use the CLI with the root user ID. The GUI allows only the default configuration, which is striped. To create sequential volumes, complete the following steps:

1. Identify your MDisks by listing them. Run the **svcinfo lsmdisk** command (Example 9-3) to list them.

Example 9-3 svcinfo lsmdisk output

```
IBM_2072:demo:webguest>lsmdisk
id name   status mode  mdisk_grp_id mdisk_grp_name capacity ctrl_LUN_#
controller_name UID tier
0  mdisk1 online array 2          MD_pool      2.2TB
generic_hdd
1  mdisk0 online array 2          MD_pool      2.2TB
generic_hdd
2  mdisk3 online array 1          UD_pool      2.2TB
generic_hdd
IBM_2072:demo:webguest>lsmdisk -unit mb -filtervalue
"mdisk_grp_name=MD_pool"
id name   status mode  mdisk_grp_id mdisk_grp_name capacity ctrl_LUN_#
controller_name UID tier
0  mdisk1 online array 2          MD_pool      2.2TB
generic_hdd
1  mdisk0 online array 2          MD_pool      2.2TB
generic_hdd
IBM_2072:demo:webguest>lsmdisk -unit mb -filtervalue
"mdisk_grp_name=UD_pool"
id name   status mode  mdisk_grp_id mdisk_grp_name capacity ctrl_LUN_#
controller_name UID tier
2  mdisk3 online array 1          UD_pool      2.2TB
generic_hdd
```

2. As explained in 10.1, “Storage virtualization introduction” on page 160, MDisks are compose of extents. The default extent size is 256 MB. The number of extents from one MDisk to another MDisk can vary according to the SAN Volume Controller / Storwize quorum information. To use all the extents in the MDisk, you must verify the number of the free extents by running the **lsfreeextents** command (Example 9-4).

Example 9-4 lsfreeextents output

```
IBM_2072:demo:webguest>lsfreeextents mdisk1
id 0
number_of_extents 8929
IBM_2072:demo:webguest>lsfreeextents mdisk0
id 1
number_of_extents 8934
IBM_2072:demo:webguest>lsfreeextents mdisk3
id 2
number_of_extents 8837
```

3. Take the number of free extents that are available in each MDisk and multiply them by the size of the extent, which is 256 MB. For example:
   - For mdisk1, the volume size in MB is 8929 (number_of_extents) x 256 MB (extent size) = 2285824 MB.
   - For mdisk0, the volume size in MB is 8934 (number_of_extents) x 256 MB (extent size) = 2287104 MB.
   - For mdisk3, the volume size in MB is 8837 (number_of_extents) x 256 MB (extent size) = 2262272 MB.

> **Information:** You need to do these calculations because the V3700 CLI requires you to specify the **-size** parameter if you use **-vtype seq**.

4. Create the volume (VDisk) by using the **-vtype=seq** flag, which means sequential type, using the value that was discovered in step 3 (Example 9-5).

Example 9-5 Sequential volume creation

```
IBM_2072:demo:webguest>mkvdisk -name MD_Quorum -mdiskgrp MD_pool -iogrp io_grp0
-mdisk mdisk0 -size 1024 -unit=mb -vtype seq
Virtual Disk, id [0], successfully created
IBM_2072:demo:webguest>mkvdisk -name MD_vol01 -mdiskgrp MD_pool -iogrp io_grp0
-mdisk mdisk0 -size 2284800 -unit=mb -vtype seq
Virtual Disk, id [1], successfully created
```

> **Hint:** The **-size** value of 2284800 for MD_vol01 is derived from having the quorum on the same array. The number of free extents must be recalculated after quorum creation. mdisk0 has 8934 extents of 256 MB in size. After you use four of those extents to create the quorum (4 x 256 MB = 1024 MB), the number of free extents is 8930. Multiplying 8930 (number of free extents) with 256 (extents size) leads to 2286080 MB as the value of **-size** for the metadata LUN that collocates with the quorum.

5. Run the **svctask mkvdisk** command to create all user data and metadata volumes.

> **Tip**: The 1 GB metadata quorum volume can be created on any metadata MDisk.

6. Run the **svcinfo lsvdisk** command to list all the created volumes. In Example 9-6, you can see the VDisk UID, which identifies the LUN in the OS. You can also see the volume type as seq (sequential) and the volume size.

Example 9-6 svcinfo lsvdisk output

```
IBM_2072:demo:webguest>lsvdisk
id name      IO_group_id IO_group_name status mdisk_grp_id mdisk_grp_name capacity type
FC_id FC_name RC_id RC_name vdisk_UID                         fc_map_count copy_count
fast_write_state se_copy_count RC_change compressed_copy_count
0 MD_Quorum 0       io_grp0     online 1        MD_pool     1.00GB  seq
                                                              1          empty            0
no        0
1 MD_vol01  0       io_grp0     online 1        MD_pool     2.18TB  seq
                                                              1          empty            0
no        0
IBM_2072:demo:webguest>lsmdisk
```

7. You can also see the volumes that are created by using the GUI and going to the Volumes menu (Figure 9-8).

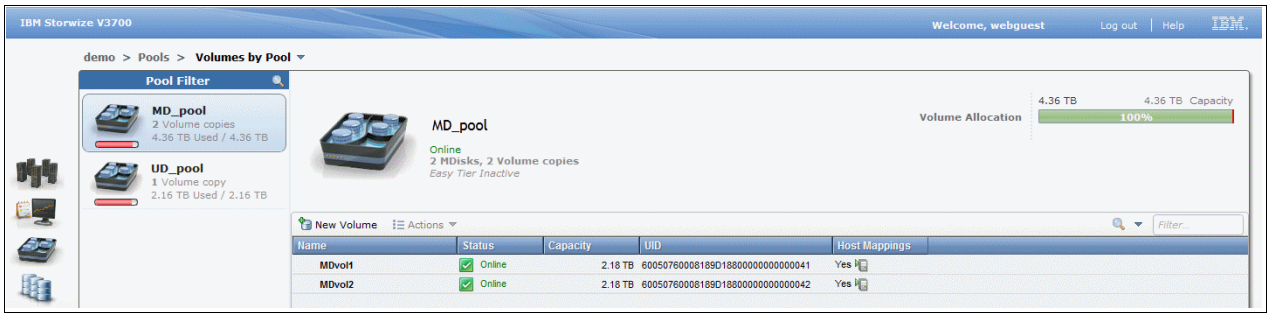

Figure 9-8 Listing volumes

Tip: You can customize columns by right-clicking above the column bar and selecting the wanted columns as shown in Figure 9-9.

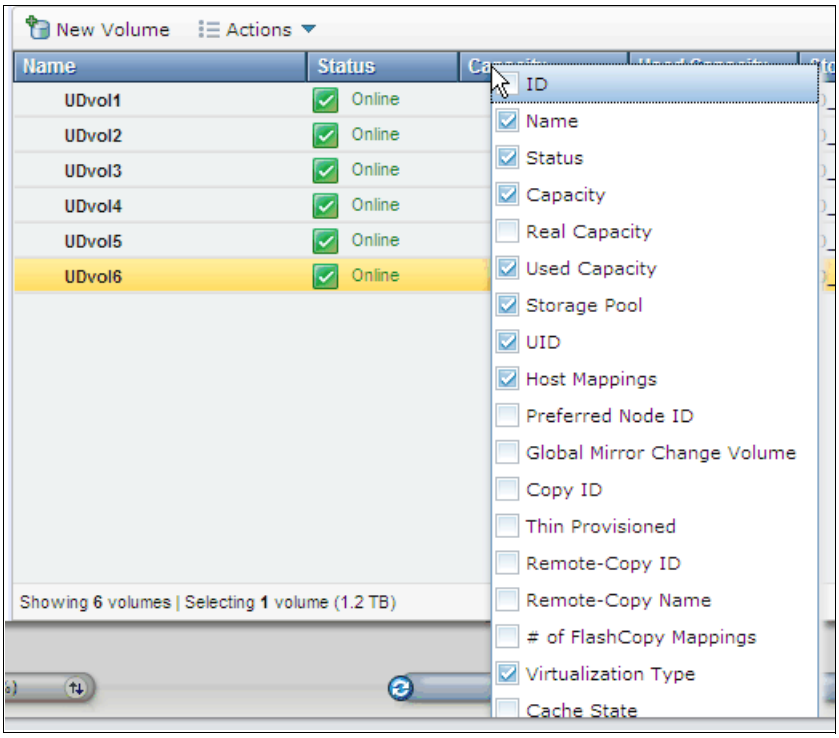

Figure 9-9 Customizing columns

9.3.4 Creating a host connection for the ProtecTIER nodes by using the Storwize V3700 GUI

To set up the host in the SAN Volume Controller or Storwize V3700 GUI, you must first know the worldwide port names (WWPNs) of the ProtecTIER nodes. To accomplish this task, complete the following steps:

1. You can see the label that is attached to the FC adapter or download the `system_view.html` file by using option “7) Generate a system view” from the ptconfig menu of your system (Example 9-7).

Example 9-7 Generating a system view report

```
-------------------------------------------------------------------------------+
| ProtecTIER Service Menu running on lbsdedup1a                                 |
+-------------------------------------------------------------------------------+
|  1) ProtecTIER Configuration (...)                                            |
|  2) Manage ProtecTIER services (...)                                          |
|  3) Health Monitoring (...)                                                   |
|  4) Problem Alerting (...)                                                    |
|  5) Version Information (...)                                                 |
|  6) Generate a service report                                                 |
|  7) Generate a system view                                                    |
|  8) Update ProtecTIER code                                                    |
|                                                                               |
|  E) Exit                                                                      |
+-------------------------------------------------------------------------------+
```

```
Your choice? 7
Begin Processing Procedure

SystemView version 1.2
Generating Service Report                                    [ Done ]
Systemview saved to /pt_work/systemview.html

End Processing Procedure Successfully
```

2. Open the `systemview.html` file with a browser and you see a list of items that you can select. Click **QLogic HBAs**, as shown in Figure 9-10.

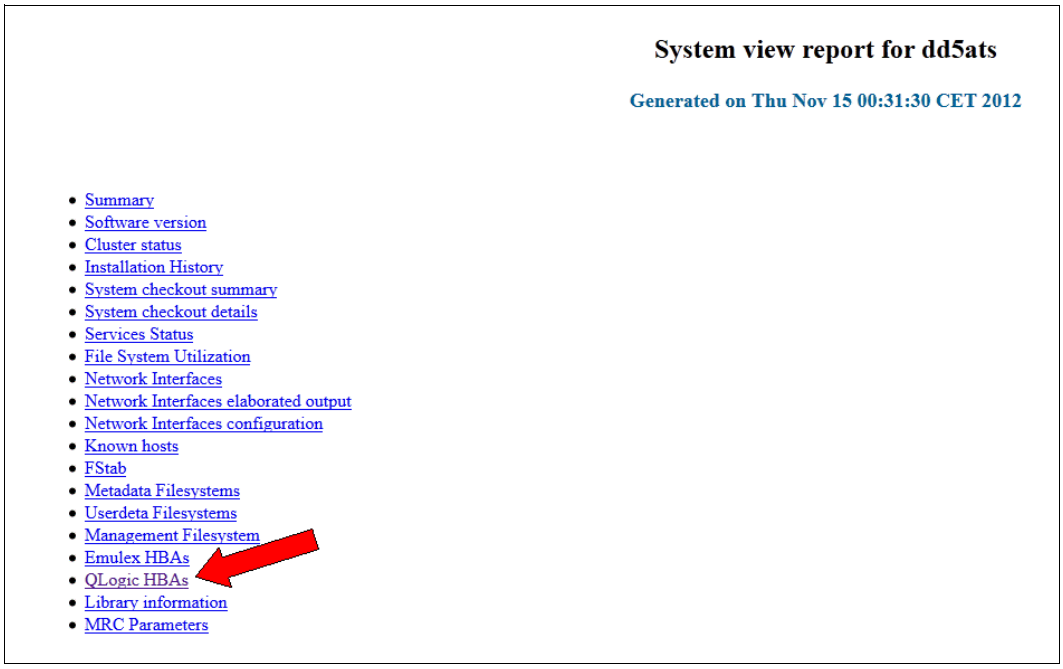

Figure 9-10 Browser showing the systemview.html file

3. The WWPN information of the QLogic back-end HBA is displayed, as shown in Figure 9-11. Be sure to use the port names as marked in the figure. Do not use the node names.

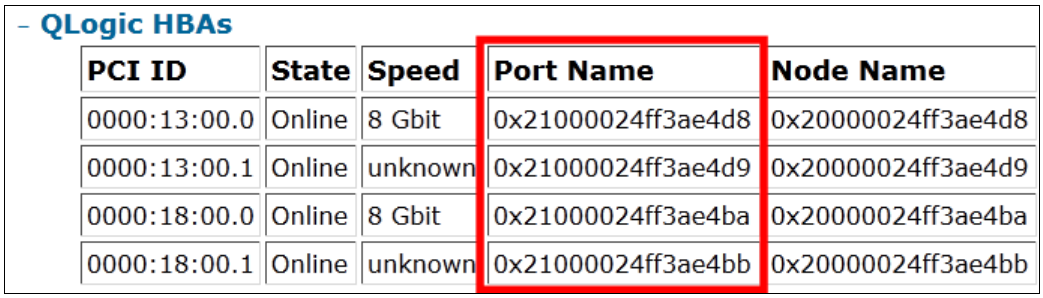

– QLogic HBAs

PCI ID	State	Speed	Port Name	Node Name
0000:13:00.0	Online	8 Gbit	0x21000024ff3ae4d8	0x20000024ff3ae4d8
0000:13:00.1	Online	unknown	0x21000024ff3ae4d9	0x20000024ff3ae4d9
0000:18:00.0	Online	8 Gbit	0x21000024ff3ae4ba	0x20000024ff3ae4ba
0000:18:00.1	Online	unknown	0x21000024ff3ae4bb	0x20000024ff3ae4bb

Figure 9-11 QLogic HBA section of the systemview.html file

To create the host connection on your Storwize V3700 server, complete the following steps:

1. Hover your cursor over the host type icon. The system displays the Hosts menu (Figure 9-12).

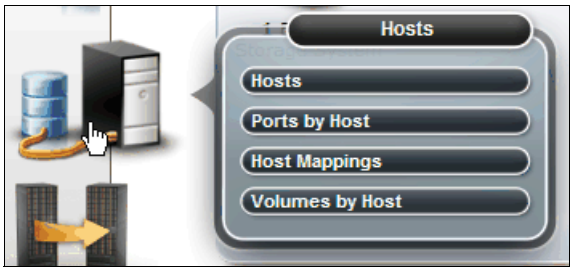

Figure 9-12 Hosts menu

2. Click **Hosts** to open the Hosts window. Click **New Host** to open the Create Host window (Figure 9-13).

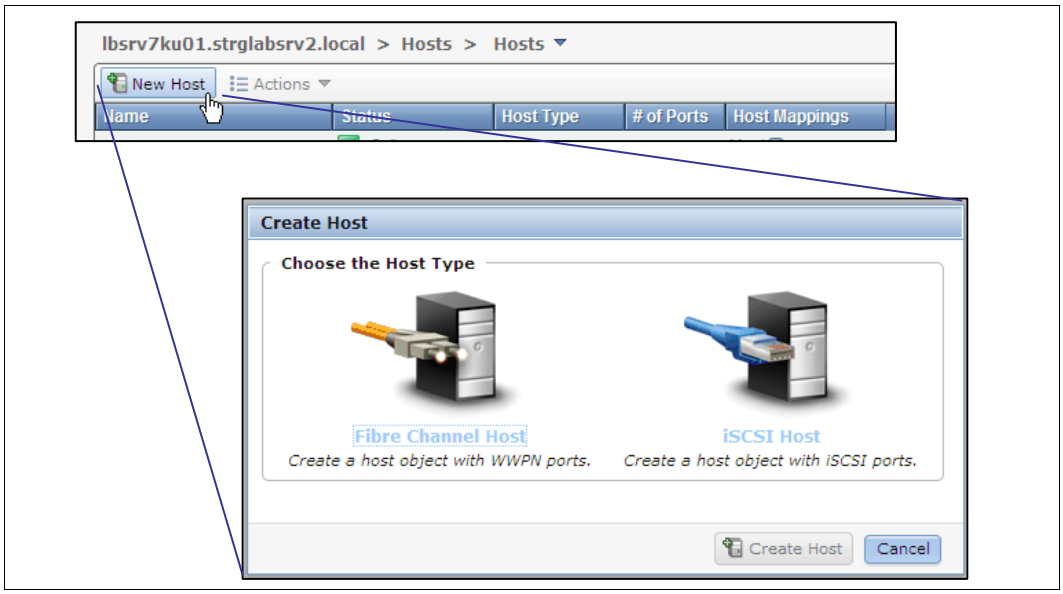

Figure 9-13 Create Host window

3. Click **Fibre Channel Host**, and then click **Create Host** to open the Create Host window (Figure 9-14).

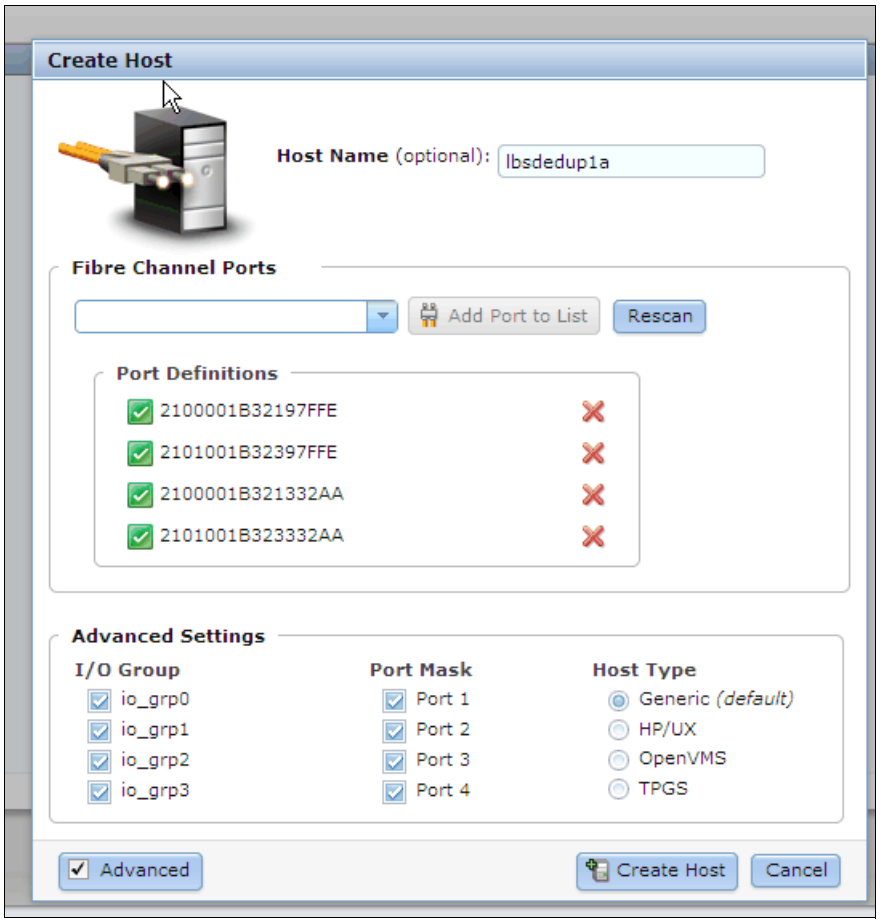

Figure 9-14 Expanded Create Host window

4. Provide the name of your host. The host name is a name that is already specified.
5. Add your worldwide port names (WWPNs) to the Fibre Channel Ports selection box.
6. Leave the **Generic (default)** radio button selected as the Host Type.
7. Click **Create Host** to return to the main menu.

Host groups and clusters: The Storwize V3700 storage subsystems do not use the concept of host group or cluster, which is used to map a volume to more than one host at the same time.

If you have a ProtecTIER dual-node server, you can choose between two options:

- Create one host for each ProtecTIER node and, when you perform the volume mapping, use the same LUN ID for both nodes.
- Create a single host, but add the ports from both ProtecTIER nodes. This method is simpler, but if there are issues with one FC port, it is more difficult to identify which node the FC port belongs to.

9.3.5 Mapping volumes to a host

To make the volumes available to the ProtecTIER node, you must map the volumes to the ProtecTIER host by completing the following steps:

1. Hover your cursor over the Pools icon. The system displays the Pools menu (Figure 9-15).

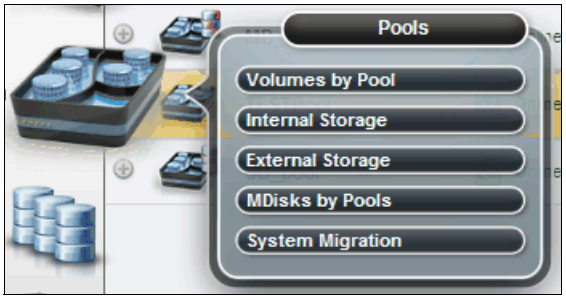

Figure 9-15 Pools menu

2. Click **Volumes by Pool** to open the list of volumes separated by pool (Figure 9-16). Mark the volumes, right-click, and click **Map to Host**.

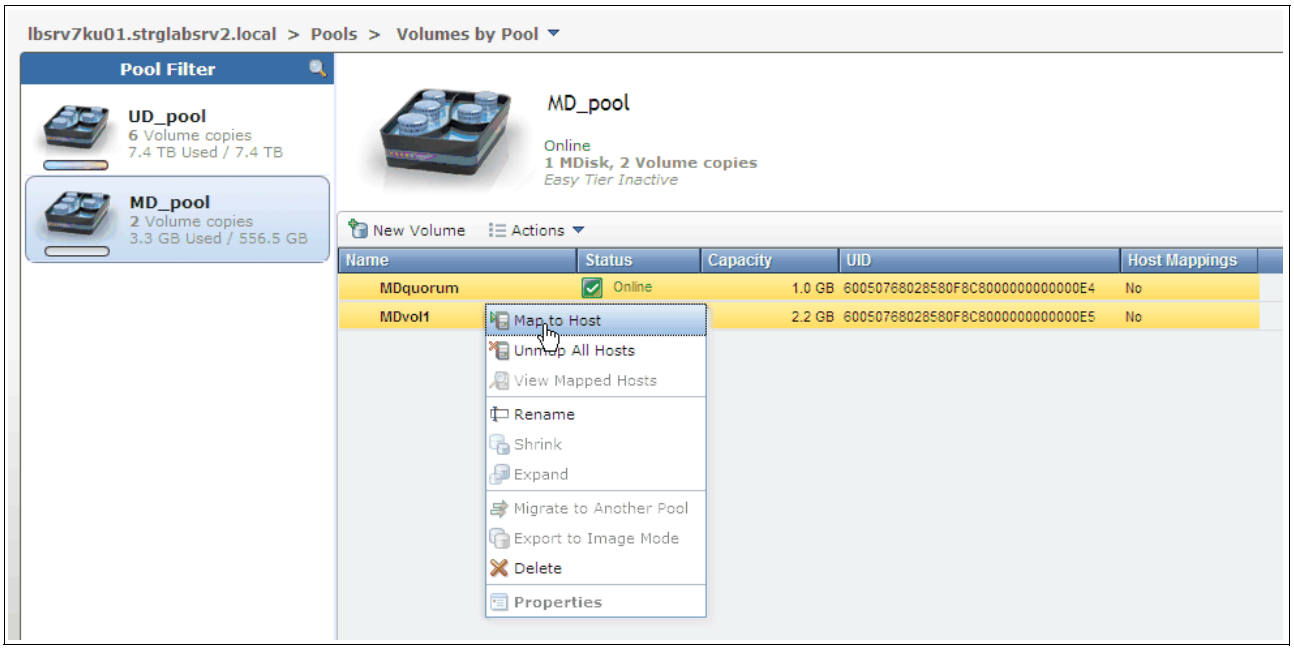

Figure 9-16 Volumes by pool

3. The Modify Host Mapping window opens. Click the **Host** drop-down menu to select the name of the host that you want to map volumes to. If the volumes are not highlighted in yellow, select volumes from the Unmapped Volumes pane and click the arrow (>) to move the volumes to the Volumes Mapped to the Host pane. Click **Map Volumes** or **Apply** (Figure 9-17).

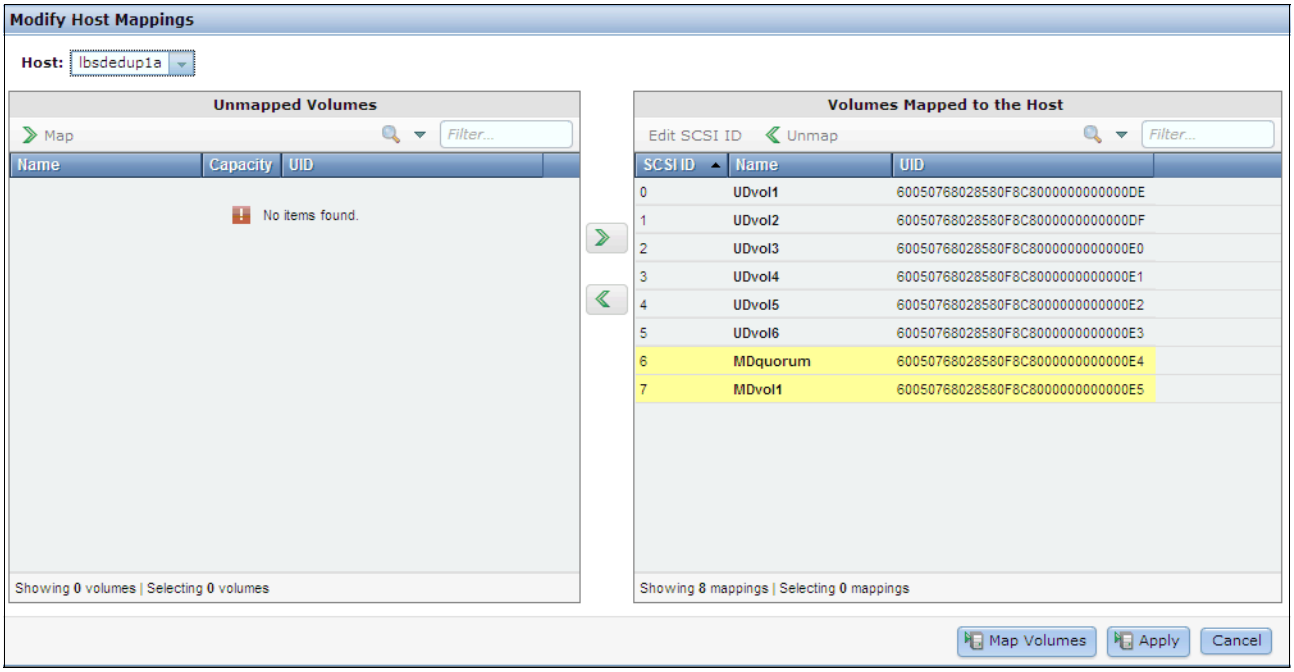

Figure 9-17 Modify mapping

Dual-node server: If you have a ProtecTIER dual-node server, you must select the other node and proceed with the volume mapping again for the second node. You might receive a warning that states that the specific volume is already mapped. Verify that the mapping is correct and continue (Figure 9-18).

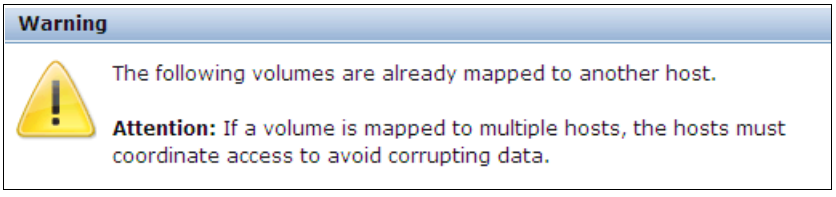

Figure 9-18 Volume mapping warning

9.3.6 Creating file systems and building the ProtecTIER repository

Beginning with ProtecTIER Version 3.2, the **`fsCreate`** tool is removed. All of the file system management is now integrated into the ProtecTIER Configuration Menu.

Tip: If you have a version of ProtecTIER earlier than Version 3.2, the **fsCreate** tool can be used to automatically create files systems for repositories by completing the following steps:

1. Log in to a ProtecTIER server.
2. At the command prompt, enter **cd /opt/dtc/app/sbin** and press Enter.
3. To open the list of file systems that are already a part of the repository, enter **./fsCreate -r** and press Enter.
4. To open the list of available new multipath devices, enter **./fsCreate -u** and press Enter.
5. Create file systems on available multipath devices, register them to /etc/fstab, and mount them. At the command prompt of Server A, enter **./fsCreate -n** and press Enter.
6. For verification, open the GFS file systems that are not part of the repository by entering ./**fsCreate -g** and pressing Enter.

To create the file systems using ProtecTIER V3.2, complete the following steps:

1. Verify that the ProtecTIER node recognizes the volumes that are presented by the SAN Volume Controller and Storwize V3700 server. The **multipath -ll** command shows the LUNs and the paths that are connected to the storage subsystem. If the LUNs do not appear to be mapped to the storage, you must run a rescan on the FC adapters or reboot the node.

   Note: You can scan the FC adapters by running the following command:

   ```
   echo "- - -" > /sys/class/scsi_host/<host??>/scan
   ```

   In this command, **<host??>** should be replaced by each FC adapter port.

   You can also filter to see only part of the data, as shown in Example 9-8. The first command filters only the volumes of type 2145, which is the SAN Volume Controller or Storwize V3700 type. The second command shows the details of one of the devices.

   Example 9-8 multipath command output

```
[root@lbsdedup1a ~]# multipath -ll | grep 2145
mpath9 (360050768028580f8c8000000000000e1) dm-8 IBM,2145
mpath8 (360050768028580f8c8000000000000e0) dm-7 IBM,2145
mpath14 (360050768028580f8c8000000000000e7) dm-27 IBM,2145
mpath7 (360050768028580f8c8000000000000df) dm-6 IBM,2145
mpath6 (360050768028580f8c8000000000000de) dm-5 IBM,2145
mpath12 (360050768028580f8c8000000000000e4) dm-11 IBM,2145
mpath11 (360050768028580f8c8000000000000e3) dm-10 IBM,2145
mpath10 (360050768028580f8c8000000000000e2) dm-9 IBM,2145

[root@lbsdedup1a ~]# multipath -ll | grep -A11 mpath9
mpath9 (360050768028580f8c8000000000000e1) dm-4 IBM,2145
[size=1.2T][features=1 queue_if_no_path][hwhandler=0][rw]
\_ round-robin 0 [prio=200][active]
 \_ 2:0:1:3 sdak 66:64  [active][ready]
 \_ 3:0:2:3 sdal 66:80  [active][ready]
 \_ 1:0:1:3 sdam 66:96  [active][ready]
 \_ 4:0:1:3 sdbs 68:96  [active][ready]
```

```
\_ round-robin 0 [prio=40][enabled]
 \_ 4:0:0:3 sdbk 67:224 [active][ready]
 \_ 1:0:0:3 sde  8:64   [active][ready]
 \_ 3:0:1:3 sdk  8:160  [active][ready]
 \_ 2:0:0:3 sdq  65:0   [active][ready]
```

2. In the ProtecTIER Service Menu, select the options “1) ProtecTIER Configuration (...)”and then “6) File Systems Management (...)” (Example 9-9).

Example 9-9 ProtecTIER configuration menu

```
[root@lbsdedup1a ~]# su - ptconfig
+------------------------------------------------------------------------------+
| ProtecTIER Service Menu running on lbsdedup1a                                |
+------------------------------------------------------------------------------+
|  1) ProtecTIER Configuration (...)                                           |
|  2) Manage ProtecTIER services (...)                                         |
|  3) Health Monitoring (...)                                                  |
|  4) Problem Alerting (...)                                                   |
|  5) Version Information (...)                                                |
|  6) Generate a service report                                                |
|  7) Generate a system view                                                   |
|  8) Update ProtecTIER code                                                   |
|                                                                              |
|  E) Exit                                                                     |
+------------------------------------------------------------------------------+
Your choice? 1

+------------------------------------------------------------------------------+
| ProtecTIER Service Menu running on lbsdedup1a                                |
| ProtecTIER Configuration (...)                                               |
+------------------------------------------------------------------------------+
|  1) Configure ProtecTIER node                                                |
|  2) Recover Configuration for a replaced server                              |
|  3) Configure machine serial number for a replaced server                    |
|  4) Configure RAS                                                            |
|  5) Update Time, Date, Timezone and Timeserver(s)                            |
|  6) File Systems Management (...)                                            |
|  7) Configure replication (...)                                              |
|  8) IP Network configuration (...)                                           |
|  9) Update Firmware                                                          |
| 10) Update the System's name                                                 |
| 11) Validate configuration                                                   |
|                                                                              |
|  B) Back                                                                     |
|  E) Exit                                                                     |
+------------------------------------------------------------------------------+
Your choice? 6
```

3. The File Systems Management menu opens. You have options to display information about the devices or file systems configuration. You have choices to configure the file systems on all available devices or for a single device. Example 9-10 shows the creation of a file system on a single unused device.

Example 9-10 File System Management menu

```
-------------------------------------------------------------------------------+
| ProtecTIER Service Menu running on lbsdedup1a                                 |
| ProtecTIER Configuration (...)                                                |
| File Systems Management (...)                                                 |
+-------------------------------------------------------------------------------+
|  1) Configure file systems on all available devices                           |
|  2) Create file system(s) on a single unused device                           |
|  3) Extend a file system with a new unused device                             |
|  4) Update /etc/fstab                                                         |
|  5) Display configured devices                                                |
|  6) Display unused devices                                                    |
|  7) Display GFS repository file systems                                       |
|  8) Display unused GFS file systems                                           |
|  9) Increase capacity completion (applicable for a second cluster node)       |
|                                                                               |
|  B) Back                                                                      |
|  E) Exit                                                                      |
+-------------------------------------------------------------------------------+
Your choice? 2
Begin Processing Procedure
Display of the available devices
Device:         Size:           Status
1. mpath0       1024.00M        Unused
2. mpath10      1286144.00M     Unused
3. mpath14      568832.00M      Unused
4. mpath11      1286144.00M     Unused
5. mpath6      1284864.00M     Unused
6. mpath7      1286144.00M     Unused
7. mpath8      1286144.00M     Unused
8. mpath9      1286144.00M     Unused

Please select device: 3
Please enter number of file systems[1-4]:1
Last file system - using the remaining free space
The file system to be created will have 568832.00 MiB size
Please confirm:? (yes|no) yes
Creating physical volume                                        [ Done ]
Creating volume group                                           [ Done ]
Creating logical volume                                         [ Done ]
Creating file system                                            [ Done ]

End Processing Procedure Successfully
```

4. After you create the file system for all the devices that are needed, you can go to the ProtecTIER Manager, add the node to it, and then select the Repository menu to create the repository. For more information, see *IBM System Storage TS7650 ProtecTIER Deduplication Appliance Installation Roadmap Guide*; GA32-0921.

9.3.7 Expanding the repository

When you expand the repository, use the same spindle type and quantity of RAID groups for metadata or user data. For example, if the original two metadata LUNs were built on RAID 4+4 groups, then the added metadata RAID groups must be at least 4+4 to maintain the same level of performance. Using storage from 2+2 or 4+1 RAID groups, for example, for the expansion might degrade performance because of an input/output operations per second (IOPS) bottleneck.

Important: The total number of volumes for both the ProtecTIER repository for metadata (MD) and user data (UD) should not exceed 170. Each volume size should not exceed 6 TB to comply with the 1 PB ProtecTIER repository size limit.

With ProtecTIER V3.2, the 8 TB restriction is removed. When the individual disk size is large (for example, 2 or 3 TB), use RAID 6 with 6+2 disk members.

10

IBM SAN Volume Controller, IBM Storwize V7000, and IBM Storwize V7000 Unified Storage

Storage virtualization technology complements a virtual server environment. The IBM SAN Volume Controller, IBM Storwize V7000 (Storwize V7000), and IBM Storwize V7000 Unified Storage (Storwize V7000 Unified) are the storage virtualization products from IBM.

This chapter describes how these three storage virtualization products can be connected to the ProtecTIER system as back-end storage. This chapter also describes the recommended topology, volume (LUNs) configuration, and settings.

This chapter describes the following topics:

- Storage virtualization introduction and terminology
- General notes and best practices for SAN Volume Controller, Storwize V7000, and Storwize Unified V7000
- Guidelines for configuring these storage virtualization products and the ProtecTIER product in a SAN
- General recommendations for configuring metadata and user pools
- Steps to configure the ProtecTIER repository in a storage virtualization product

Information: The primary interface methods that are used to work with Storwize V7000 Unified are NAS protocols that allow file level I/O. ProtecTIER requires block level I/O for its potential back-end storage subsystems. Storwize V7000 Unified offers both file level I/O and block level I/O capabilities.

10.1 Storage virtualization introduction

Storage virtualization brings intelligence to the storage environment by implementing a virtualization layer between storage back-end disks and hosts. System administrators can view and access a common pool of storage on the storage area network (SAN). This functionality helps administrators use storage resources more efficiently and provides a common base for advanced functions, such as copy services, thin-provisioned volumes, and easy tiering.

Here are some key benefits of using storage virtualization:

- Improves and optimizes capacity usage.
- Enables data migration of virtualized data with minimum application disruption.
- Facilitates a common platform for data replication, even if the back-end storage does not have this functionality.
- Provides a centralized point of control for storage provisioning across heterogeneous SAN environments.
- Reduces license costs for advanced functions in the virtualized storage systems.
- Increases storage administrator productivity.

Here are the IBM products that provide this technology:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified
- Storwize V3700 (See Chapter 9, “IBM Storwize V3700” on page 139 for details.)

SAN Volume Controller can virtualize multivendor storage disks to perform LUN provisioning in a high-end environment. In addition to performing external storage virtualization, Storwize V7000 has internal disks that constitute its storage capacity for a mid-range environment. Storwize V7000 Unified consolidates block (SAN) virtualization and file workloads (NAS) in to a single storage system.

Supported vendor storage: The list of supported vendor storage disks in SAN Volume Controller Version 6.3 can be found at the following website:

`http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003907`

Ensure that you refer to the System Storage Interoperation Center (SSIC) website for the latest information:

`http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss`

10.1.1 Terminology

This section describes the terminology that we use in this chapter that relates to storage virtualization. Here are the terms:

Storage pool Also known as managed disk group (MDG). An MDG is a collection of storage capacity (MDisks) that provides the capacity requirements (extents) for volumes that are based on MDisks characteristics, such as disk speed and type.

Managed disk (MDisk) An MDisk is composed of a Small Computer System Interface (SCSI) logical unit (LU) that a redundant array of independent disks (RAID) controller provides and that a clustered system manages. The MDisk is not visible to host systems on the storage area network (SAN).

Volume /(VDisk) Also known as virtual disk (VDisk). A volume is the representation of a SCSI disk that is presented to a host system.

Extents A unit of data that manages the mapping of data between managed disks (MDisks) and volumes. By default, its size is 256 MB.

Figure 10-1 shows a summary of storage virtualization concepts. The disk RAID array from an external storage system (or from internal disks, when you use Storwize V7000) is presented to an SAN Volume Controller or Storwize V7000 as an MDisk. A set of MDisks forms an array, from which extents are taken to create the volumes (LUNs). The volumes, now in virtualized mode, are presented to the hosts. In this sense, the hosts no longer see the back-end disks directly, and the SAN Volume Controller or Storwize V7000 behaves like a controller provisioning LUNs to the hosts.

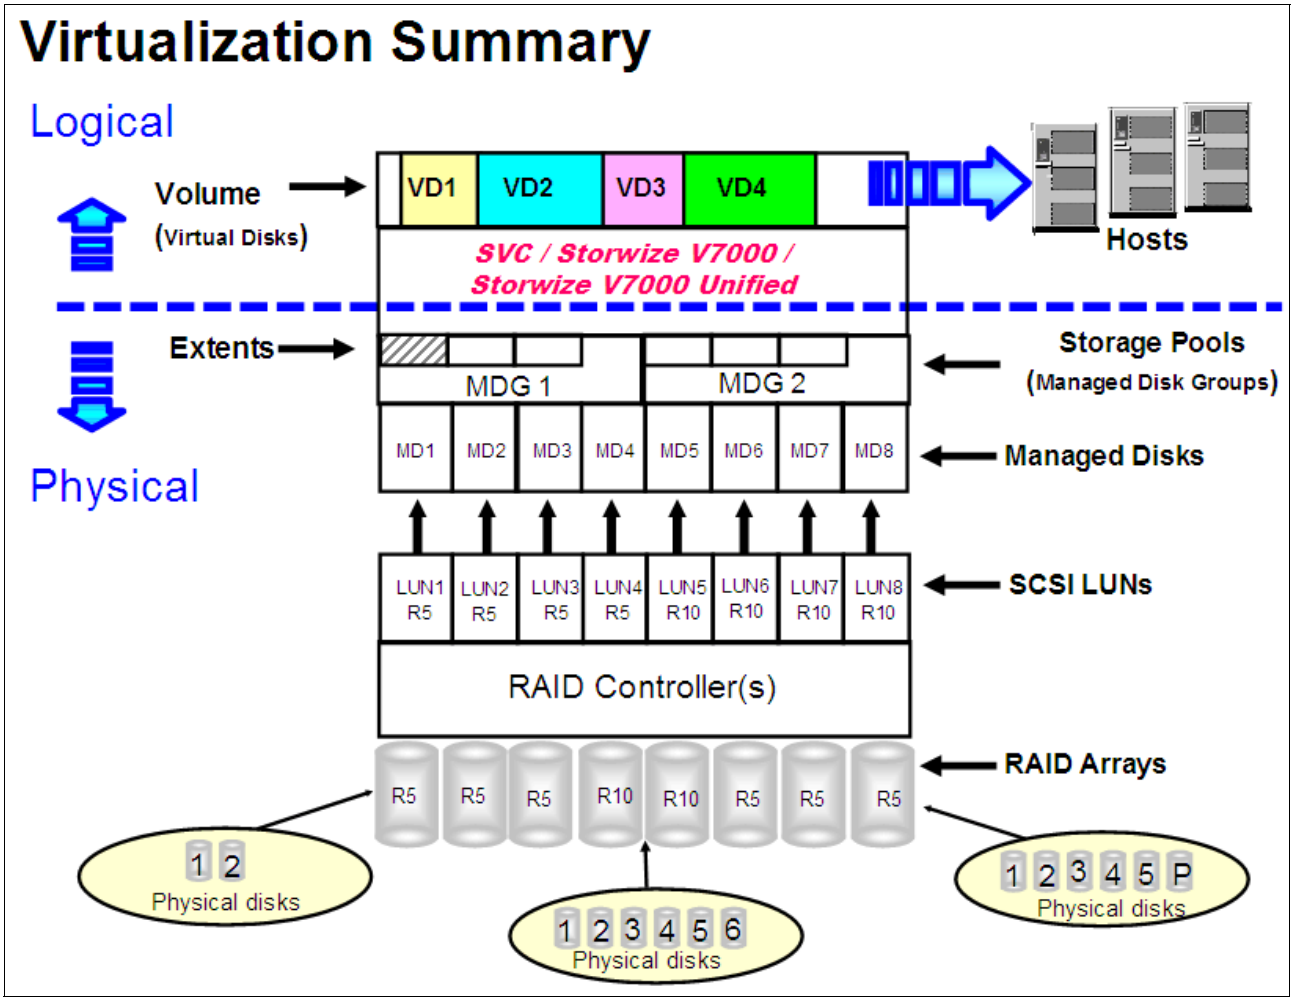

Figure 10-1 Storage virtualization concepts summary

There are three types of virtualization for volumes: striped, sequential, and image, as shown in Figure 10-2.

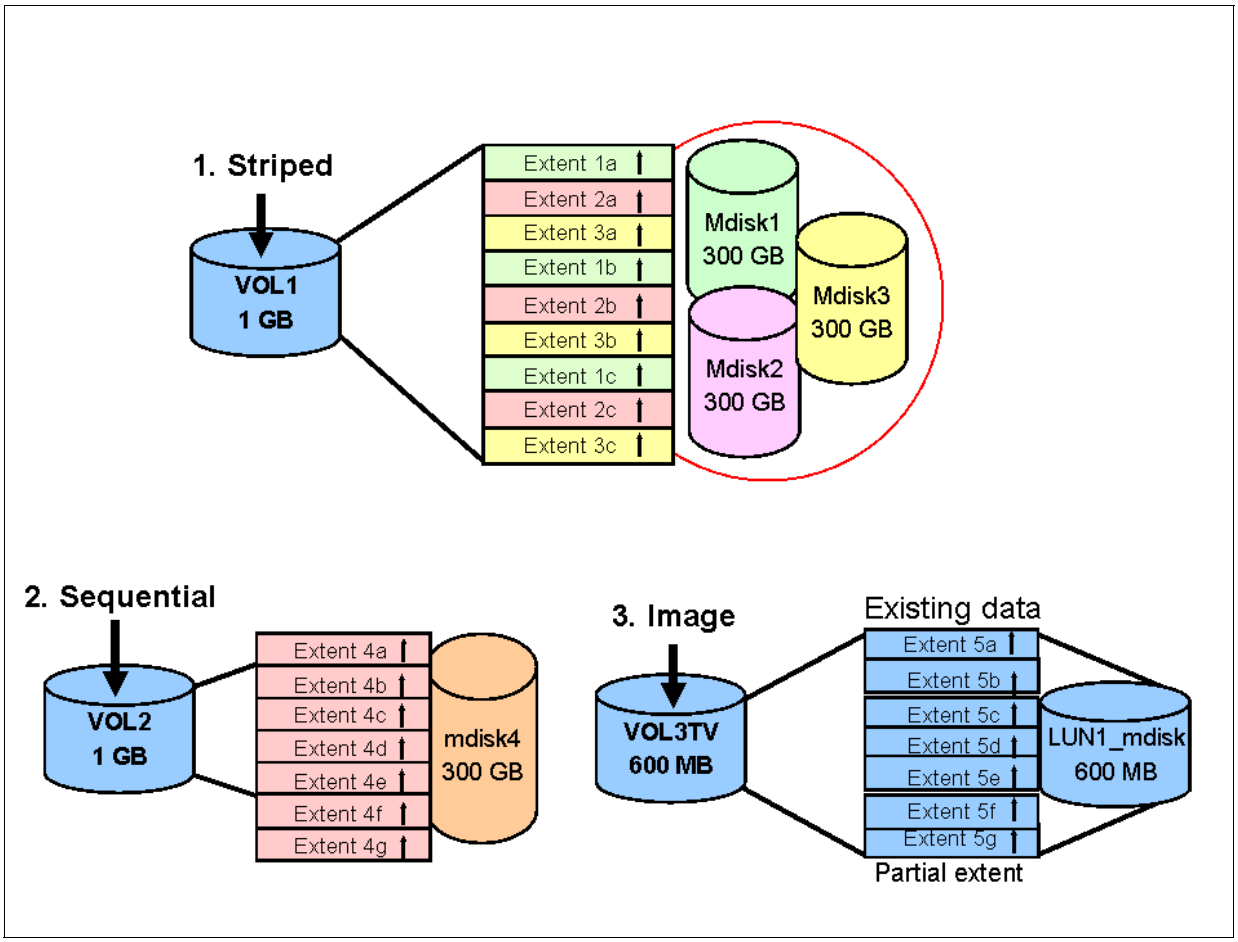

Figure 10-2 Three virtualization types for volumes

- Striped (default): Takes an extent in turn from each managed disk that has free extents (or a subset of managed disks, known as a striped set) in the pool (managed disk group).
- Sequential: Allocates extents from one managed disk to create the volume (VDisk) if enough consecutive free extents are available on the chosen managed disk.
- Image: Creates a one-to-one direct mapping between a virtual disk and the managed disk that contains existing data, for example, an existing LUN from an external virtualized storage.

Note: More information about storage virtualization products can be found in the IBM information Center for these products, which can be found at the following websites:

- SAN Volume Controller:

  http://publib.boulder.ibm.com/infocenter/svc/ic/index.jsp

  Support information, including configuration limits and restrictions for SAN Volume Controller, can be found at:

  http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003903

- Storwize V7000:

  http://publib.boulder.ibm.com/infocenter/storwize/ic/index.jsp

  Support information, including configuration limits and restrictions for Storwize V7000, can be found at:

  http://www.ibm.com/support/docview.wss?uid=ssg1S1003741

- Storwize V7000 Unified:

  http://publib.boulder.ibm.com/infocenter/storwize/unified_ic/index.jsp

  Support information, including configuration limits and restrictions for Storwize V7000 Unified, can be found at:

  http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003906

10.2 General notes

The following general notes apply for SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified storage subsystems:

- Because of the nature of the ProtecTIER product, storage resources are heavily used, so the SAN Volume Controller or Storwize V7000 storage subsystem should be dedicated solely to ProtecTIER activity.
- If it is not possible to dedicate the array to ProtecTIER, use zoning and LUN masking to isolate the TS7650G from other applications.

  Important: The TS7650G must not share pools, MDisks, or volumes with other applications.

- Disk-based replication is supported only by RPQ. Use the ProtecTIER native IP replication feature that is available in Version 2.5 and later. For more information about replication, see Part 5, “Replication and disaster recovery” on page 369.
- When you use SAN peer-to-peer (P2P) topology to connect the TS7650G to the disk array, create a dedicated zone for the ProtecTIER back-end ports. Do not mix the back-end ports (QLogic) with the front-end ProtecTIER ports (Emulex) or any other SAN devices in the same zone.
- Ensure that you have redundant connections to both SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified nodes or node canisters.
- Create *only* two storage pools (also known as managed disk groups), where one is used for metadata and the other is used for user data. Creating more pools causes the storage system cache to be segmented among all pools. By having only two pools, you optimize the storage subsystem cache allocation.

- Use the *sequential* virtualization type for each MDisk instead of striped virtualization. Because of the nature of ProtecTIER file system architecture and workload, sequential virtualization shows better performance than striped. When you expand the ProtecTIER repository, if you use sequential virtualization, the new sequential volumes maintain the same performance characteristics as the existing volumes. If you use the striped virtualization, you must manually check the extents distribution among the MDisks to ensure that they are balanced.
- Do not use thin-provisioning volumes. All of the storage that is presented to the ProtecTIER system is run through a “padding” process as part of the repository creation. This padding process immediately fills the thinly provisioned LUNS to 99% capacity.
- The number of paths to each volume *must not exceed eight paths* between the host and the SAN Volume Controller, Storwize V7000, or Storwize V7000 Unified environments.
- Each V7000 / SAN Volume Controller volume has a preferred V7000 controller node to be bound to. Balance the volumes between the SAN Volume Controller, Storwize V7000, or Storwize V7000 Unified controller nodes. The ProtecTIER server has connections and SAN zoning for both nodes. By default, the allocation should equally divide all volumes between the controller nodes; do not interfere with this behavior.
- As mentioned before in the general recommendations for back-end storage in 8.6.4, “User data on SATA disks” on page 135, the volume size should not exceed 6 TB.

Important: Always use RAID 6 in combination with SATA or NL-SAS drives.

LUN management: Starting with ProtecTIER Version 3.2, the management of LUNs greater than 8 TB is improved. When ProtecTIER V3.2 and later works with LUNs greater than 8 TB, it splits them in to logical volumes of smaller size, which means that you can work with LUNs greater than 8 TB. There is no benefit in performance in doing this action.

- ProtecTIER is a *random-read* application. 80 - 90% of I/O on a typical TS7650G environment is random read at 64 KB block size. Implement suitable performance optimizations and tuning as recommended by the disk vendor for this I/O profile.
- There is no benefit in using IBM Easy Tier® “On” on the metadata or user data pool. The workload pattern of ProtecTIER does not produce enough hot extents to justify the usage of solid-state disks (SSDs).

Number of back-end LUNs: The total number of volumes for both the ProtecTIER repository for metadata (MD) and user data (UD) should not exceed 170. Each volume size should not exceed 6 TB to comply to the 1 PB ProtecTIER repository size limit.

10.3 Firmware level

The array firmware level must be equal to or greater than the firmware version that is listed in the ProtecTIER Interoperability Matrix. You can access the ProtecTIER Interoperability Matrix, along with information about other tape devices, by going to the following website and navigating to the Compatibility information section:

http://www.ibm.com/systems/storage/tape/resources.html#compatibility

You can also access the ProtecTIER Interoperability Matrix directly at the following website:

http://public.dhe.ibm.com/common/ssi/ecm/en/ivl12348usen/IVL12348USEN.PDF

10.4 Fibre Channel connection topology

SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified all support redundant connections to clustered ProtecTIER nodes. To ensure full protection against the loss of any Fibre Channel paths from the ProtecTIER nodes to the Storwize V7000, use redundant host connections.

To meet the business requirements for high availability, use SAN design practices to build a dual fabric network, two independent fabrics, or SANs. To increase the total number of switch ports in a fabric, switches can be connected together with one or more interswitch links (ISLs). Multiple ISLs between two switches do not increase the total number of paths to the volumes.

Direct attachment: SAN Volume Controller does not support direct attachment connection to a ProtecTIER server. A Fibre Channel (FC) switch is required for attaching a Storwize V7000 to a ProtecTIER server. For V3700, Storwize V7000, and Storwize V7000 Unified, direct attachment might be available by RPQ.

Connect each host to the appropriate single-ported host channels on the SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified controllers, also known as nodes or node canister. These configurations have host and drive path failover protection. To ensure redundancy and stability, every ProtecTIER node must be configured to obtain two paths to each controller.

One port of each ProtecTIER server host adapter should be connected to one SAN fabric. For a volume (VDisk or LUN), which is owned by an I/O group, the number of paths from the SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified nodes to a host must not exceed eight. Use SAN zoning to limit the number of paths among ports across the SAN. This setting reduces the number of instances that the same volume (VDisk or LUN) is reported to a host operating system.

No more than four HBA ports of the ProtecTIER servers in total should be connected to the SAN. Each port should be zoned to a different set of the SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified ports to maximize performance and redundancy.

Figure 10-3 is an example of a redundant SAN fabric Fibre Channel configuration with a ProtecTIER single-node configuration. Figure 10-4 on page 167 is an example of a redundant SAN fabric Fibre Channel configuration with a ProtecTIER dual-node configuration.

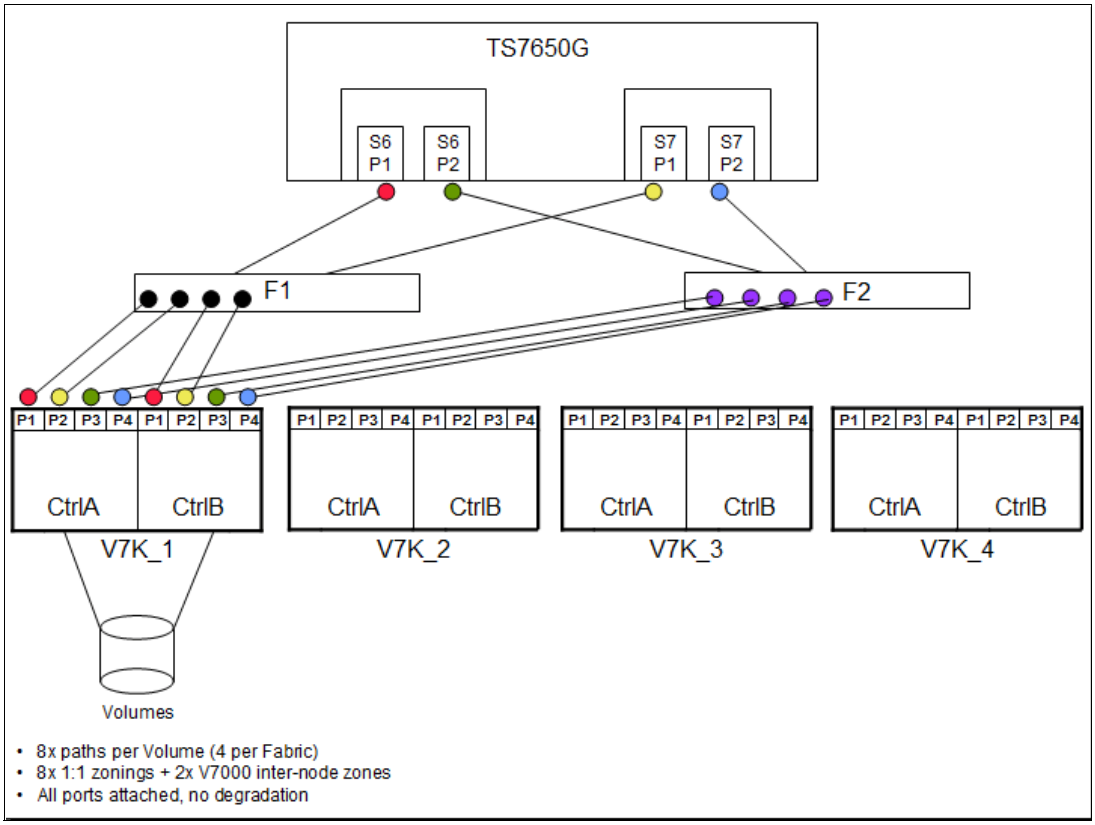

Figure 10-3 Example of a redundant SAN fabric Fibre Channel configuration with a ProtecTIER single-node configuration

Figure 10-4 illustrates that each host HBA port is zoned with two ports from each SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified node, providing a total of eight paths.

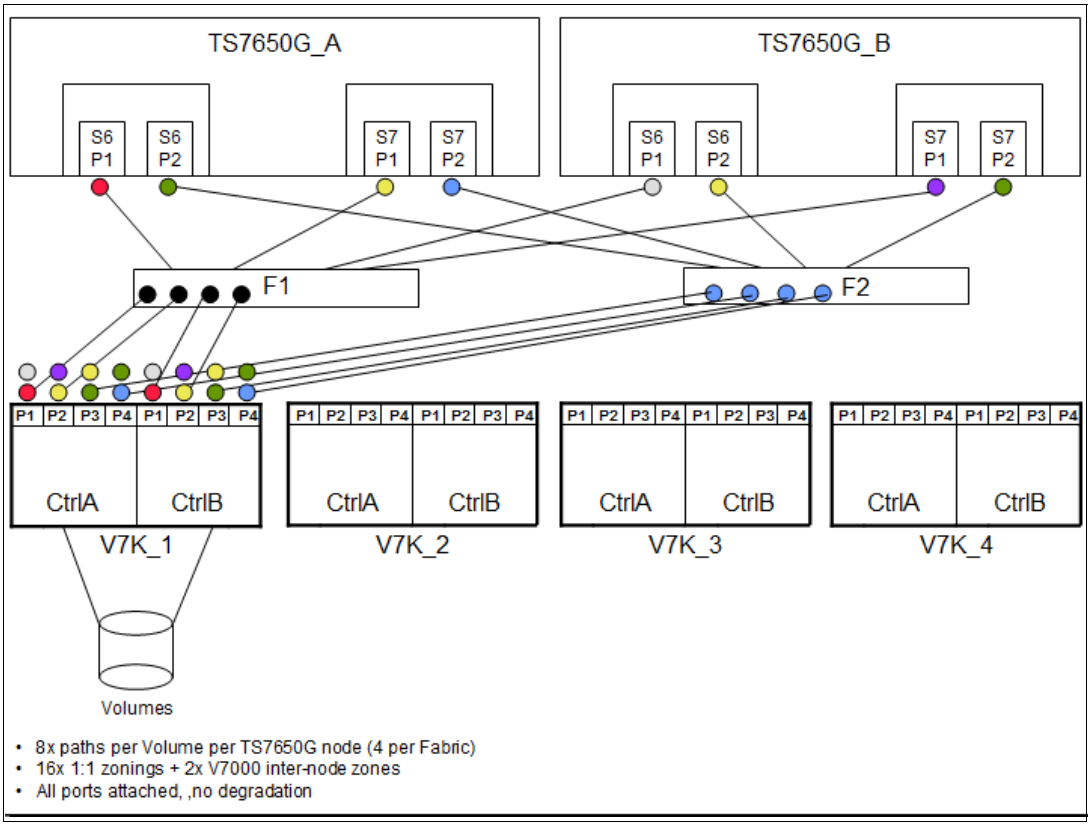

Figure 10-4 Example of redundant SAN fabric Fibre Channel configuration with a ProtecTIER dual-node configuration

Multipath settings: There is no need to manually configure the multipath specific settings (by running the **`multipath.conf`** command). Configuration is done automatically by an installation script of the ProtecTIER product called **`autorun`**.

Example 10-1 depicts how the output looks when you run the **`multipath -ll`**, depending on the physical cabling and the SAN zoning config:

Example 10-1 Example output from the multipath -ll command

```
mpath2 (360050768028109ed8800000000000009) dm-2 IBM,2145
[size=3.8T][features=1 queue_if_no_path][hwhandler=0][rw]
\_ round-robin 0 [prio=200][active]
\_ 6:0:1:8  sddb 70:144  [active][ready]
\_ 5:0:1:8  sddc 70:160  [active][ready]
\_ 4:0:1:8  sddd 70:176  [active][ready]
\_ 3:0:1:8  sdde 70:192  [active][ready]
\_ round-robin 0 [prio=40][enabled]
\_ 6:0:0:8  sdah 66:16   [active][ready]
\_ 3:0:0:8  sdai 66:32   [active][ready]
\_ 5:0:0:8  sdaj 66:48   [active][ready]
\_ 4:0:0:8  sdak 66:64   [active][ready]
```

Here are the corresponding details for Example 10-1 on page 167 for SCSI multipath information for each path:

```
HOST (=ProtecTIER ports) : CHANNEL : ID (=different disk subsystem ports per Host port) : LUN (=Volume)
```

- The mpath2 device alias represents the LUN/volume with ID 8 (unique ID in brackets shows disk subsystem WWN+LUNID).
- Eight paths are available to reach the mpath2 device.
- All four ProtecTIER QLogic ports (3,4,5,6) can reach this LUN ID 8.
- Each ProtecTIER port can reach two different disk subsystem ports (for example, PT port 6 sees disk subsystem ports 0 and 1).
- This is an active/active capable disk subsystem (there is no ghost status as the DS4000 active/passive system would show).
- Only one PT to disk subsystem port path group is actively used (active) at one time; the other one is available for failover scenarios (enabled).
- I/O is load balanced among the [active] path group (round robin).

10.5 User data and metadata pool: General recommendations

This section provides general recommendations for setting up metadata and user data pools for optimum performance in your ProtecTIER environment. We also provide general guidelines for expanding the ProtecTIER repository.

10.5.1 Metadata pool

Here are some items to consider regarding the metadata pool:

- Use balanced RAID 10 groups for metadata MDisks (use a layout according to your planning requirements) with at least 4+4 members. The recommended number of metadata RAID groups (MDisks) is determined by the Capacity Planning Tool during the pre-sales process. This number can be 2 - 10, based on repository size, factoring ratio, and performance needs.

  Terminology:

  - Balanced RAID 10 is defined as RAID group creation in Storwize V7000 and Storwize V7000 Unified, which uses drives to form the array from different enclosures. This setup protects against drive and enclosure failures, considering the SAS chain and the disk spare protection.
  - A SAS chain is defined as a cabling scheme for a string of expansion enclosures that provides redundant access to the drives inside the enclosures. This setup is achieved by having both node canisters in the control enclosure in the Storwize V7000.

- Do not share the metadata pool with the user data pool.
- Create a single pool for metadata that contains all the metadata MDisks.
- Specify an extent size default of 256 MB for each metadata volume created.
- For each metadata MDisk, create a single metadata (MD) volume to accommodate the full size of each metadata MDisk, by using the *sequential* virtualization type.

Quorum volume: The 1 GB MD quorum volume can be created on any metadata MDisk.

10.5.2 User data pool

Here are some items to consider regarding user data pools:

- Use RAID for data protection and performance. The RAID type depends on the disk category, but is usually RAID 5 or RAID 6.
- Do not share the user data pool with the metadata pool.
- Create a single pool for user data that contains all the user data MDisks.
- Specify an extent size default of 256 MB for each user data volume that is created.
- For each user data MDisk, create a single user data (UD) VDisk / volume to accommodate the full size of each user data MDisk by using the *sequential* virtualization type.
- The size of user data volumes must be consistent. All user data volumes should be the same size.
- Comply with the suggestions that are related to the number of ProtecTIER file systems, as described in 8.3, “Dependencies from a ProtecTIER view” on page 129.

10.6 Configuration steps

Figure 10-5 shows the steps to configure the ProtecTIER repository in a storage virtualization product.

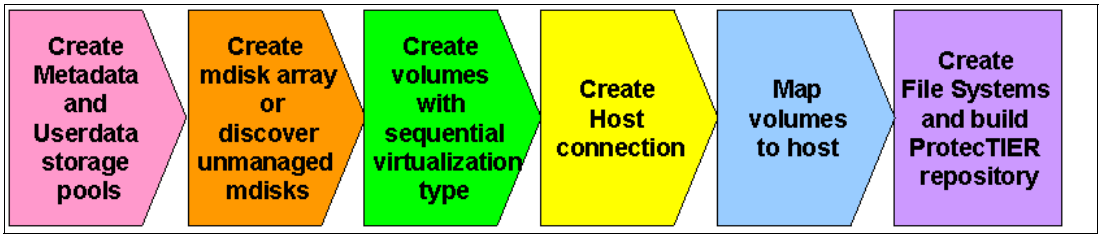

Figure 10-5 Configuration steps

Here are the steps:

1. Create empty user data and metadata storage pools.
2. Create the MDisk arrays or discover unmanaged MDisks.
3. Create volumes (VDisks) with a sequential virtualization type by using the command-line interface (CLI).
4. Create a host connection for the ProtecTIER nodes.
5. Map volumes to the host.
6. Create file systems by using the File Systems Management under the **`ptconfig`** menu and build the ProtecTIER repository by using the ProtecTIER Manager.

10.6.1 Creating empty user data and metadata storage pools

You should create *only* two storage pools (also known as managed disk groups), where one is used for metadata and the other is used for user data. When you create more pools, the storage system cache is segmented among all pools. By having only two pools, you optimize the storage subsystem cache allocation.

To create these items, complete the following steps:

1. Hover your cursor over the Pools icon. The system displays the Pools menu (Figure 10-6). Click **MDisks by Pools**.

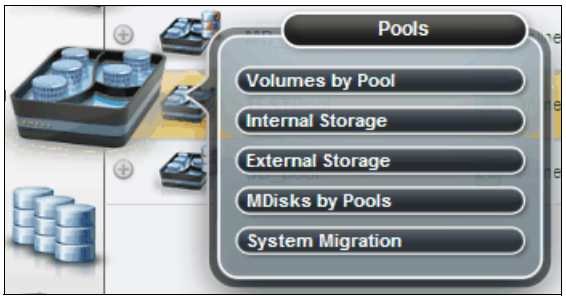

Figure 10-6 Pool menu

2. For each pool, click **New Pool** to create an empty pool. Insert a name for the pool, such as MD_pool or UD_pool. (Figure 10-7). Click **Next** to continue with the pool creation process.

Tip: Set the extent size settings of 256 MB. This is the default. It is needed to allow large repository dimensions of 1 PB or more of virtualization within the storage. You should also change the warning threshold to 100% to not receive alerts about pool usage. The default is 80%. The pool is fully allocated by the ProtecTIER LUNs. To do this task, you need to use the CLI to create the pools, as shown in Example 10-2 on page 171.

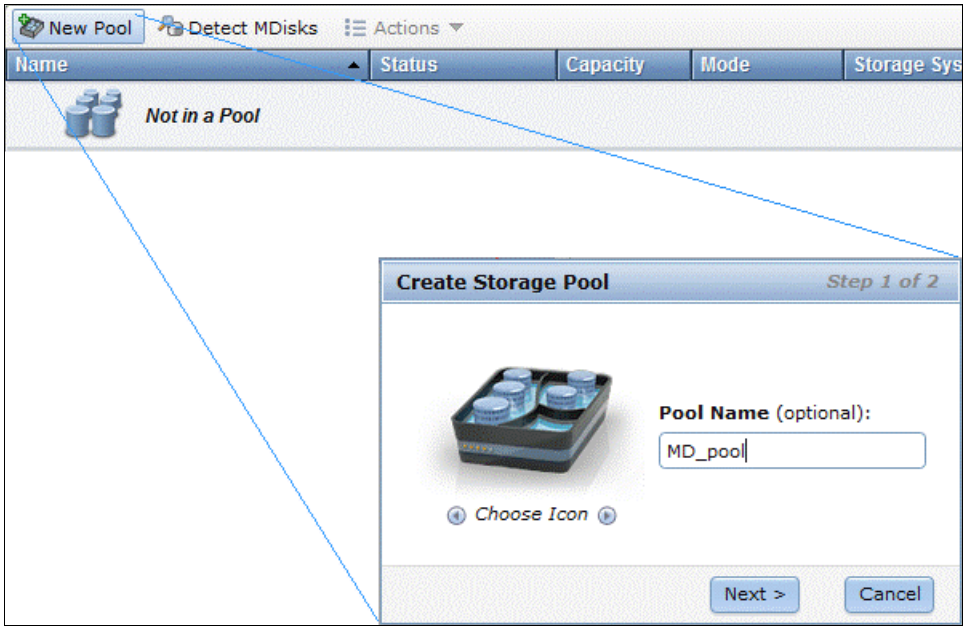

Figure 10-7 Create pools

3. Click **Next**. If you created your MDisk arrays or use MDisks from external storage, select the MDisks to be included in the pool. Otherwise, click **Finish** to create an empty pool. (Figure 10-8).

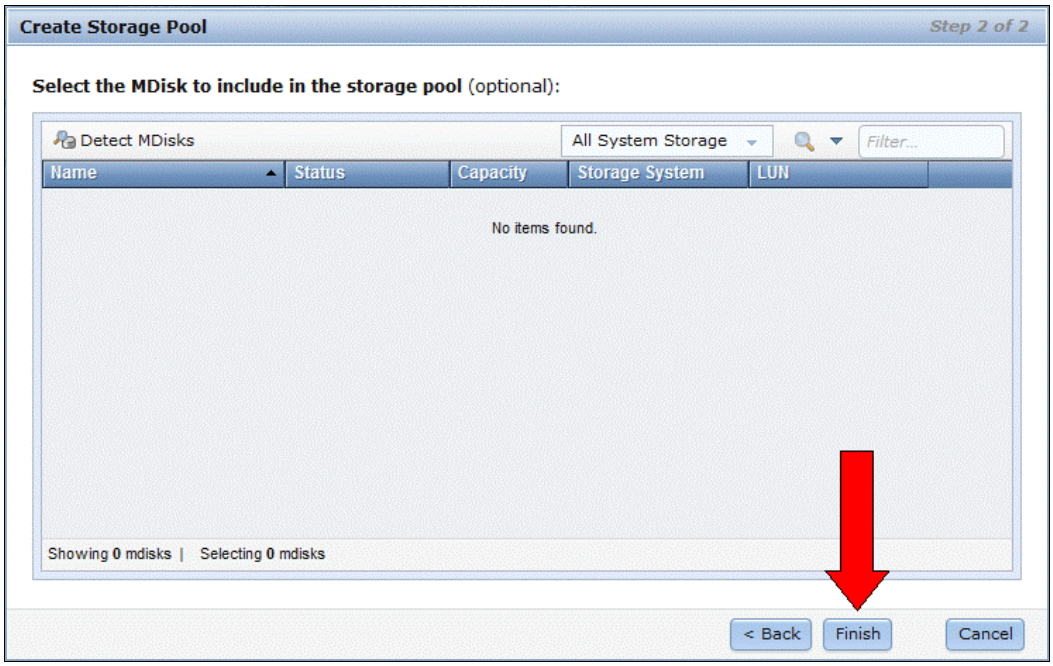

Figure 10-8 Create a storage pool

4. Repeat this procedure for the next pool. You need two pools only: one for metadata, and one for user data.

The steps can also be done by using the CLI, as shown in Example 10-2.

Example 10-2 Creating a pool by using the CLI

```
# svctask mkmdiskgrp -ext 256 -warning 100% -name UD_pool
```

10.6.2 Creating the MDisk arrays or discovering unmanaged MDisks

When you use the Storwize V7000 and Storwize V7000 Unified servers, you can use internal or external disks to create MDisks. The SAN Volume Controller allows only MDisks coming from an external virtualized storage.

To accomplish these tasks, complete the following steps:

1. Hover your cursor over the Pools icon. The system displays the Pools menu (Figure 10-9).

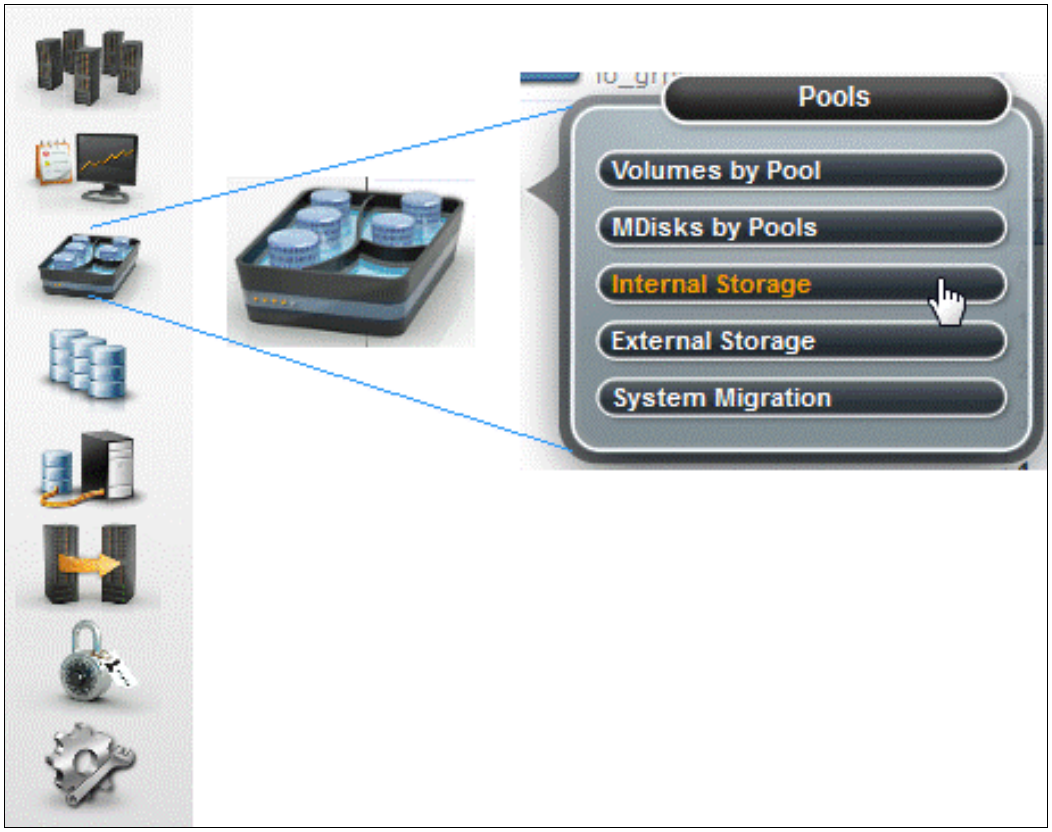

Figure 10-9 Pools menu

2. To create an MDisk with internal disks with IBM Storwize V7000 and IBM Storwize V7000 Unified, click **Internal Storage** to display the list of candidate internal disks. Select the appropriate Drive Class that you want to configure, for example, SAS disks for metadata and NL_SAS (or SATA) for user data. Then, click **Configure Storage** (Figure 10-10).

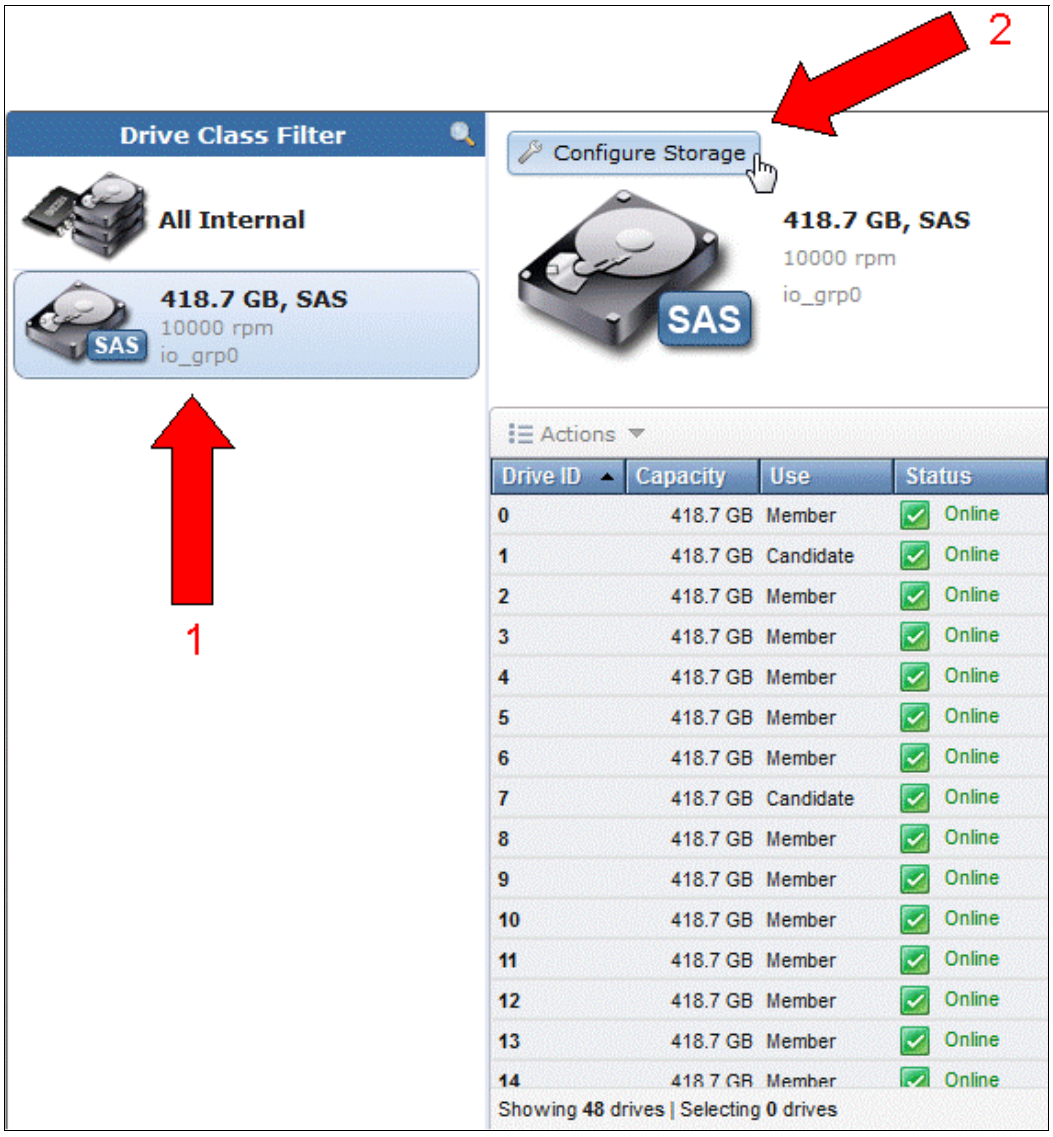

Figure 10-10 Configure Storage

3. If a SAN Volume Controller is being used, or if the MDisk is created from external storage in a Storwize V7000 or Storwize V7000 Unified server, click **External Storage** to display the list of unmanaged MDisks. If no MDisks are displayed, click **Detect MDisks**. When the list of available MDisks opens, select the appropriate disks, right-click, and click **Add to pool**.

4. The Configure Internal Storage window opens. The default settings are shown in the window. Proceed as shown in Figure 10-11 on page 175 and complete the following steps:
   a. To use different settings, click **Select a different configuration**.
   b. Select the RAID type, for example, **Balanced RAID 10** for a metadata array.
   c. Click **Optimized for Performance** to force the GUI into creating RAID arrays of equal dimensions.
   d. In the Number of drives to provision field, enter the number of disks to configure in your RAID arrays.
   e. Ensure that the automatically calculated RAID layout resembles the RAID arrays that you want to create. For example, having eight disks in a Balanced RAID 10 creates a 4+4 array. Having eight disks in a RAID 6 creates a 6+2 array. Having five disks in a RAID 5 creates a 4+1 array.

      Important: Closely monitor the output of the GUI in step e. Especially while you try to create multiple arrays concurrently, the GUI can get a little too creative when it comes to deciding the array dimensions. Change the number of disks that you want to configure in a single step to avoid unwanted array sizes.

   f. Click **Next**.

      Important: Use the ProtecTIER Planner Tool to calculate the correct number of metadata file systems and follow this guideline: For Storwize V7000 metadata planning, use the planner with 15 K Fibre Channel drives, even though the Storwize V7000 disk has 10 K drives.

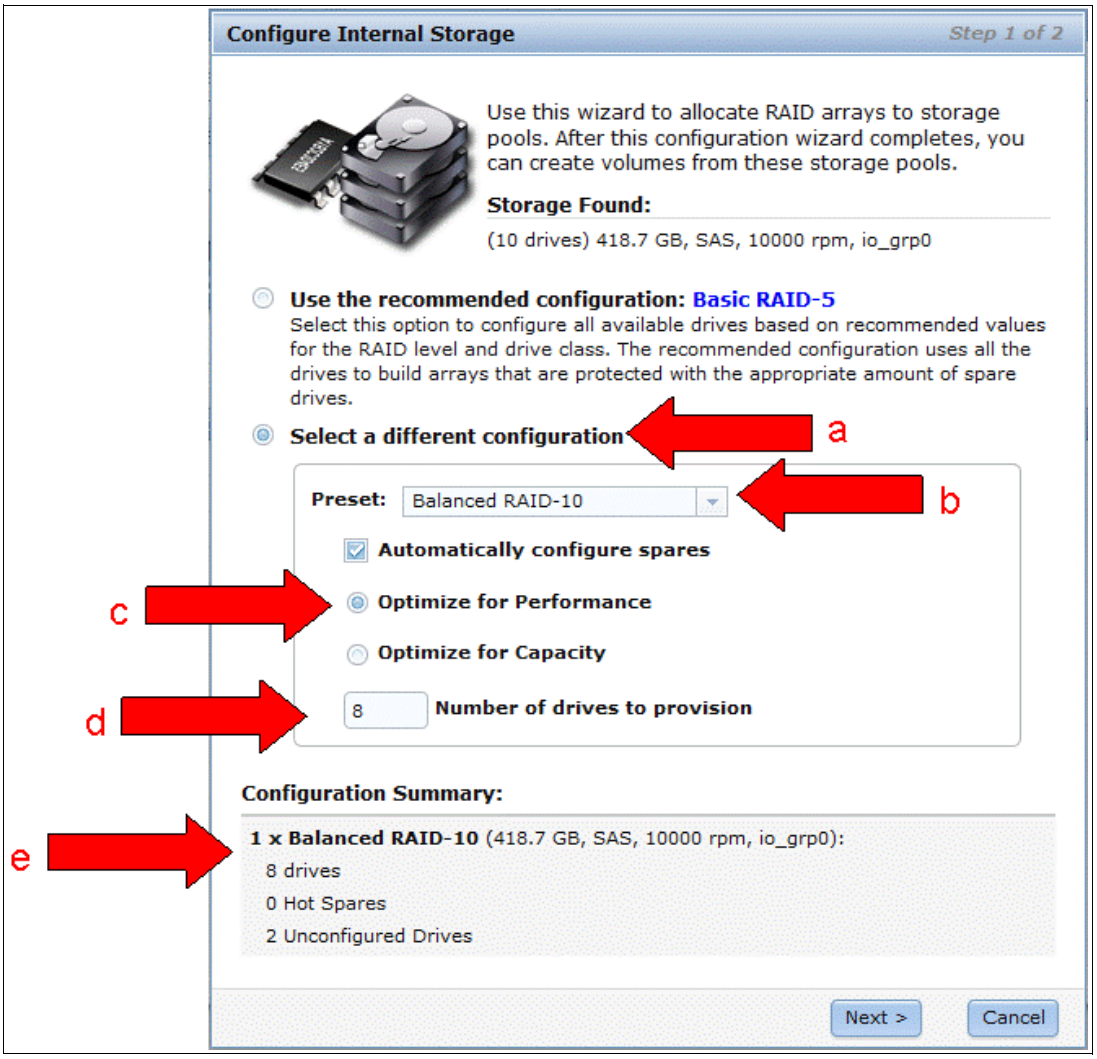

Figure 10-11 Configure Internal Storage window

5. Repeat the steps to create all the MDisks arrays that are needed to create the repository.

These steps can also be done by using the CLI, as shown in Example 10-3.

Example 10-3 MDisk array creation command

```
# svctask mkarray -drive 28:27:26:25 -level raid5 -sparegoal 0 -strip 128 UD_pool
# svctask mkarray -drive 8:6:5:10 -level raid10 -sparegoal 0 -strip 128 MD_pool
```

The numbers after the **-drive** flag represents the physical disk's ID, which is part of the array.

Important:

The **-sparegoal 0** parameter allows array creation whether there are hot spare drives available or not. With this setting, you need to manually ensure that there are enough hot spare resources available within your storage subsystem to satisfy your needs.

To manually create a balanced RAID 10 array, you need to select the array members in a specific order. The **-drive** parameter uses the drives in the following order: Primary:Mirror:Primary:Mirror:Primary:Mirror, and so on. Make sure that each Primary disk is attached through the one SAS chain and each Mirror disk is attached through the other SAS chain.

10.6.3 Creating volumes with a sequential virtualization type

The ProtecTIER architecture stripes the repository among all file systems and uses all the disks simultaneously. Because of this architecture, you should create the volume to accommodate the full size of each MDisk. Additionally, because of the nature of the ProtecTIER file system architecture and workload, *sequential* virtualization has better performance than striped virtualization.

To create volumes (VDisks) with the sequential virtualization type, you must use the CLI with the root user ID. The GUI allows only the default configuration, which is striped. To create sequential volumes, complete the following steps:

1. Identify your MDisks by listing them. Run the **svcinfo lsmdisk** command (Example 10-4) to list them.

Example 10-4 svcinfo lsmdisk output

```
# svcinfo lsmdisk -unit mb -filtervalue "mdisk_grp_name=UD_pool"
id name    status mode  mdisk_grp_id mdisk_grp_name capacity ctrl_LUN_#
controller_name UID tier
6  mdisk7  online array 1            UD_pool        1.2TB generic_hdd
7  mdisk8  online array 1            UD_pool        1.2TB generic_hdd
8  mdisk6  online array 1            UD_pool        1.2TB generic_hdd
9  mdisk10 online array 1            UD_pool        1.2TB generic_hdd
10 mdisk9  online array 1            UD_pool        1.2TB generic_hdd
11 mdisk11 online array 1            UD_pool        1.2TB generic_hdd

# svcinfo lsmdisk -unit mb -filtervalue "mdisk_grp_name=MD_pool"
id name    status mode  mdisk_grp_id mdisk_grp_name capacity ctrl_LUN_#
controller_name UID tier
5  mdisk5  online array     0        MD_pool        557.8GB generic_hdd
```

2. As explained in 10.1, “Storage virtualization introduction” on page 160, MDisks are compose of extents. The default extent size is 256 MB. The number of extents from one MDisk to another MDisk can vary according to the SAN Volume Controller / Storwize V7000 quorum information. To use all the extents in the MDisk, you must verify the number of the free extents by running the **lsfreeextents** command (Example 10-5).

Example 10-5 lsfreeextents output

```
# lsfreeextents mdisk7
id 6
number_of_extents 5019
```

```
# lsfreeextents mdisk8
id 7
number_of_extents 5024
```

3. Take the number of free extents available in each MDisk and multiply them by the size of the extent, which is 256 MB. For example:
   - For mdisk7, the volume size in MB is 5019 (number_of_extents) x 256 MB (extent size) = 1284864 MB.
   - For mdisk8, the volume size in MB is 5024 (number_of_extents) x 256 MB (extent size) = 1286144 MB.
4. Create the volume (VDisk) by using the flag **-vtype=seq**, which means sequential type, by using the value that was discovered in step 3 (Example 10-6).

Example 10-6 Sequential volume creation

```
# svctask mkvdisk -name UDvol1 -mdiskgrp UD_pool -iogrp io_grp0 -mdisk mdisk7
-size 1284864 -unit mb -vtype seq
Virtual Disk, id [0], successfully created

# svctask mkvdisk -name UDvol3 -mdiskgrp UD_pool -iogrp io_grp0 -mdisk mdisk8
-size 1286144 -unit mb -vtype seq
Virtual Disk, id [2], successfully created
```

5. Run the **svctask mkvdisk** command again to create all user data and metadata volumes.

Tip: The 1 GB metadata quorum volume can be created on any metadata MDisk.

6. Run the **svcinfo lsvdisk** command to list all the created volumes. In Example 10-7, you can see the VDisk UID, which identifies the LUN in the OS. You can also see the volume type as seq (sequential) and the volume size.

Example 10-7 svcinfo lsvdisk output

```
# svcinfo lsvdisk
id name   IO_group_id IO_group_name status mdisk_grp_id mdisk_grp_name capacity type
FC_id FC_name RC_id RC_name vdisk_UID fc_map_count copy_count fast_write_state
se_copy_count RC_change
0  UDvol1 0           io_grp0       online 1            UD_pool        1.23TB   seq
60050768028580F8C8000000000000DE 0            1          empty            0 no
1  UDvol2 0           io_grp0       online 1            UD_pool        1.23TB   seq
60050768028580F8C8000000000000DF 0            1          empty            0 no
2  UDvol3 0           io_grp0       online 1            UD_pool        1.23TB   seq
60050768028580F8C8000000000000E0 0            1          empty            0 no
3  UDvol4 0           io_grp0       online 1            UD_pool        1.23TB   seq
60050768028580F8C8000000000000E1 0            1          empty            0 no
4  UDvol5 0           io_grp0       online 1            UD_pool        1.23TB   seq
60050768028580F8C8000000000000E2 0            1          empty            0 no
5  UDvol6 0           io_grp0       online 1            UD_pool        1.23TB   seq
60050768028580F8C8000000000000E3 0            1          empty            0 no
6  MDquorum 0           io_grp0       online 0            MD_pool        1.00GB seq
60050768028580F8C8000000000000E4 0            1          empty            0 no
7  MDvol1 0           io_grp0       online 0            MD_pool        555.50GB seq
60050768028580F8C8000000000000E7 0            1          not_empty        0 no
```

7. You can also see the volumes that are created by using the GUI and going to the Volumes menu (Figure 10-12).

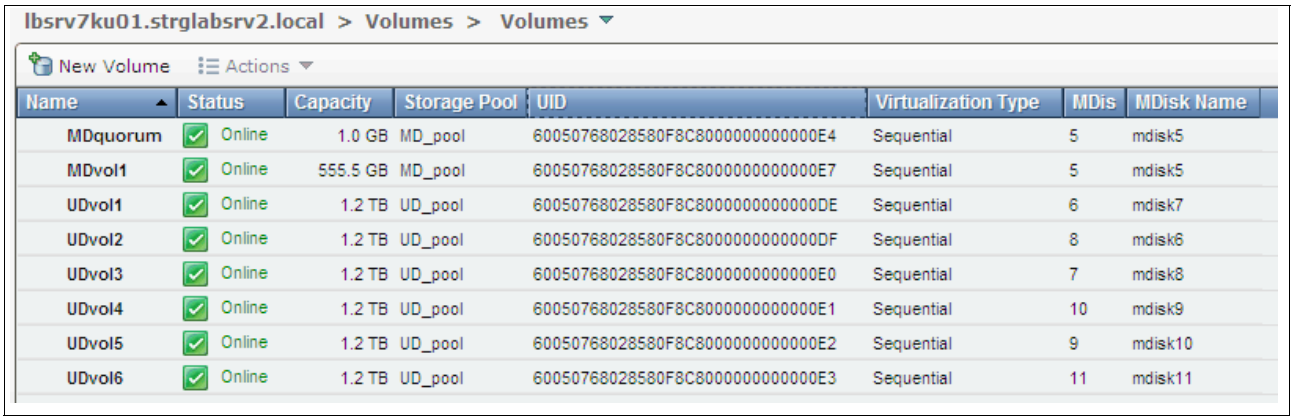

Name	Status	Capacity	Storage Pool	UID	Virtualization Type	MDis	MDisk Name
MDquorum	Online	1.0 GB	MD_pool	60050768028580F8C8000000000000E4	Sequential	5	mdisk5
MDvol1	Online	555.5 GB	MD_pool	60050768028580F8C8000000000000E7	Sequential	5	mdisk5
UDvol1	Online	1.2 TB	UD_pool	60050768028580F8C8000000000000DE	Sequential	6	mdisk7
UDvol2	Online	1.2 TB	UD_pool	60050768028580F8C8000000000000DF	Sequential	8	mdisk6
UDvol3	Online	1.2 TB	UD_pool	60050768028580F8C8000000000000E0	Sequential	7	mdisk8
UDvol4	Online	1.2 TB	UD_pool	60050768028580F8C8000000000000E1	Sequential	10	mdisk9
UDvol5	Online	1.2 TB	UD_pool	60050768028580F8C8000000000000E2	Sequential	9	mdisk10
UDvol6	Online	1.2 TB	UD_pool	60050768028580F8C8000000000000E3	Sequential	11	mdisk11

Figure 10-12 Listing volumes

Tip: You can customize columns by right-clicking above the column bar and selecting the wanted columns, as shown in Figure 10-13.

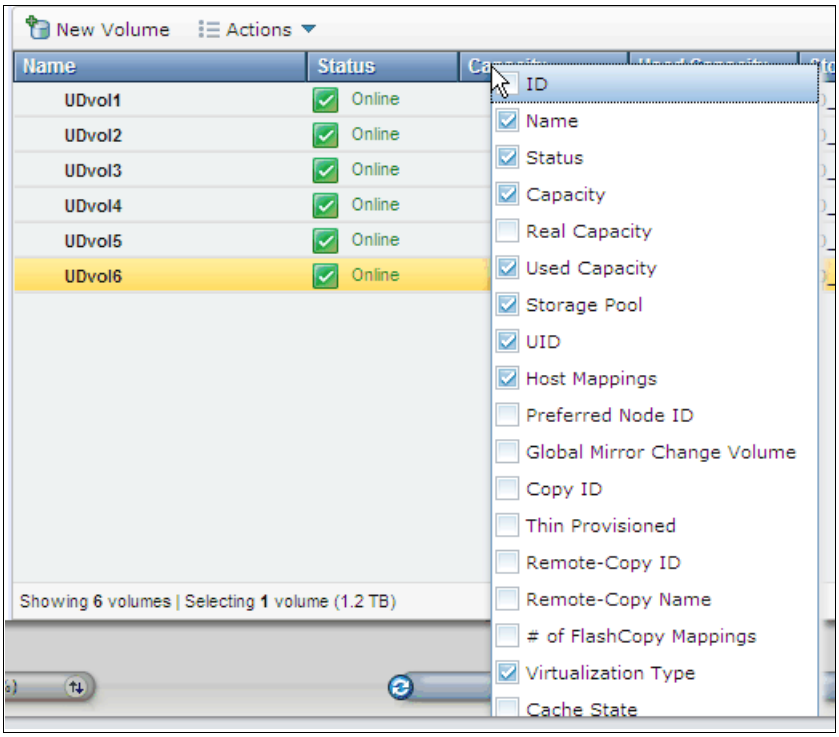

Figure 10-13 Customizing columns

10.6.4 Creating a host connection for ProtecTIER nodes in Storwize V7000 GUI

To set up the host in the SAN Volume Controller or Storwize V7000 GUI, you must first know the worldwide port names (WWPNs) of the ProtecTIER nodes.

Complete the following steps:

1. You can see the label that is attached to the FC adapter or download the `system_view.html` file by using option “7) Generate a system view” from the ptconfig menu of your system (Example 10-8).

Example 10-8 Generating system view report

```
--------------------------------------------------------------------------------+
| ProtecTIER Service Menu running on lbsdedup1a                                  |
+--------------------------------------------------------------------------------+
|  1) ProtecTIER Configuration (...)                                             |
|  2) Manage ProtecTIER services (...)                                           |
|  3) Health Monitoring (...)                                                    |
|  4) Problem Alerting (...)                                                     |
|  5) Version Information (...)                                                  |
|  6) Generate a service report                                                  |
|  7) Generate a system view                                                     |
|  8) Update ProtecTIER code                                                     |
|                                                                                |
|  E) Exit                                                                       |
+--------------------------------------------------------------------------------+
Your choice? 7
Begin Processing Procedure

SystemView version 1.2
Generating Service Report                                        [ Done ]
Systemview saved to /pt_work/systemview.html

End Processing Procedure Successfully
```

2. Open the `systemview.html` file with a browser and you see a list of items that you can select. Click **QLogic HBAs**, as shown in Figure 10-14.

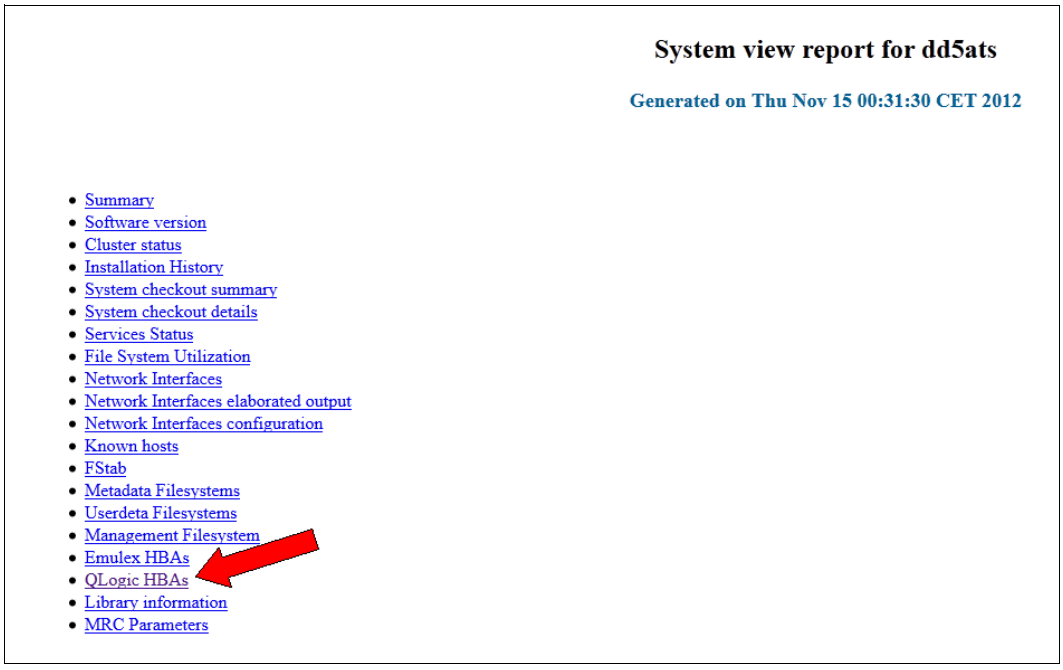

Figure 10-14 Browser showing the systemview.html file

3. The WWPN information of the QLogic back-end HBA is displayed, as shown in Figure 10-15. Be sure to use the port names as marked in the figure. Do not use the node names.

– QLogic HBAs

PCI ID	State	Speed	Port Name	Node Name
0000:13:00.0	Online	8 Gbit	0x21000024ff3ae4d8	0x20000024ff3ae4d8
0000:13:00.1	Online	unknown	0x21000024ff3ae4d9	0x20000024ff3ae4d9
0000:18:00.0	Online	8 Gbit	0x21000024ff3ae4ba	0x20000024ff3ae4ba
0000:18:00.1	Online	unknown	0x21000024ff3ae4bb	0x20000024ff3ae4bb

Figure 10-15 QLogic HBA section of the systemview.html file

To create the host connection on your SAN Volume Controller, Storwize V7000, or Storwize V7000 Unified server, complete the following steps:

1. Hover your cursor over the host type icon. The system displays the Hosts menu (Figure 10-16).

Figure 10-16 Hosts menu

2. Click **Hosts** to open the Hosts window. Click **New Host** to open the Create Host window (Figure 10-17).

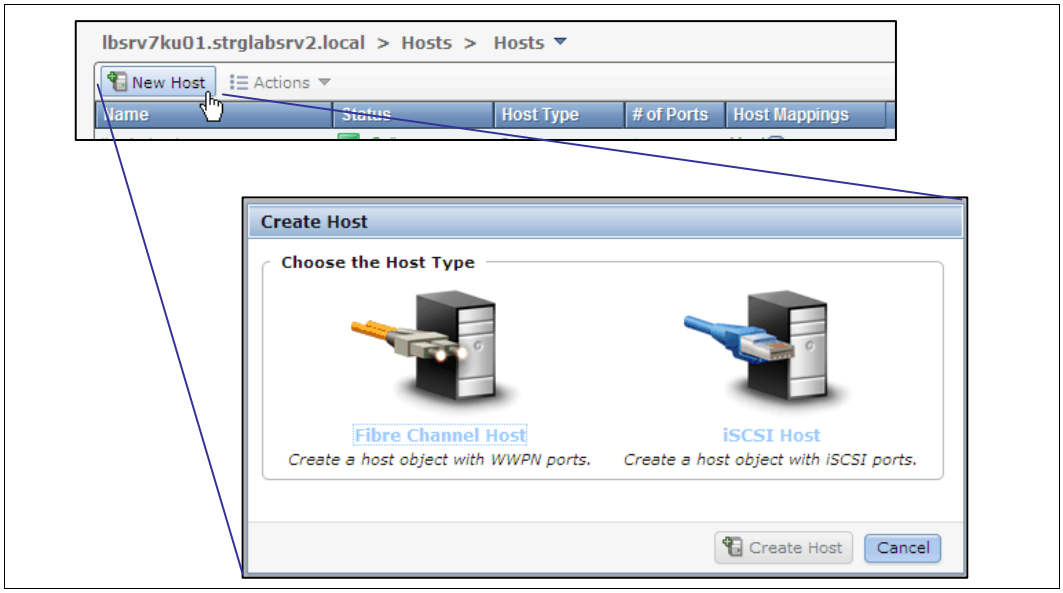

Figure 10-17 Create Host window

3. Click **Fibre Channel Host**, and then click **Create Host** to open the Create Host window (Figure 10-18).

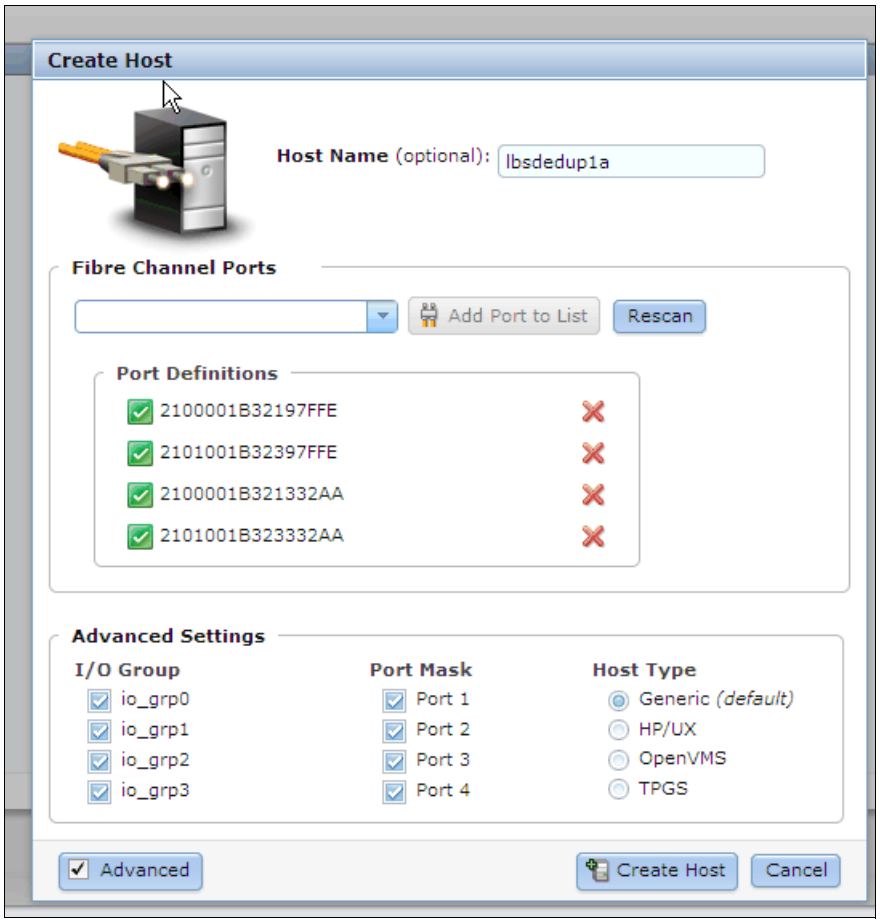

Figure 10-18 Expanded Create Host window

4. Provide the name of your host. The host name is a name that is already specified.
5. Add your worldwide port names (WWPNs) in the Fibre Channel Ports selection box.
6. Leave the **Generic (default)** radio button selected as the Host Type.
7. Click **Create Host** to return to the main menu.

Host groups and clusters: The Storwize V7000 and Storwize V7000 Unified storage subsystems do not use the concept of host group or cluster, which is used to map a volume to more than one host at the same time.

If you have a ProtecTIER dual-node server, you can choose between two options:

- Create one host for each ProtecTIER node and, when you perform the volume mapping, use the same LUN ID for both nodes.
- Create a single host, but add the ports from both ProtecTIER nodes. This method is simpler, but if there are issues with one FC port, it is more difficult to identify which node the FC port belongs to.

10.6.5 Mapping volumes to a host

To make the volumes available to the ProtecTIER node, you must map the volumes to the ProtecTIER host by completing the following steps:

1. Hover your cursor over the Pools icon. The system displays the Pools menu (Figure 10-19).

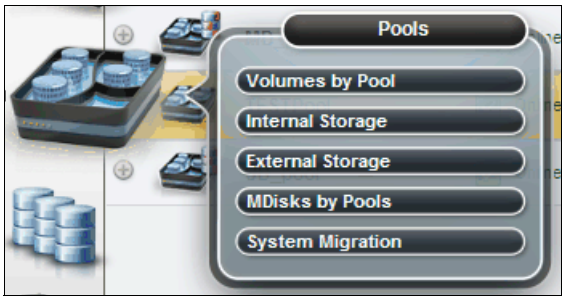

Figure 10-19 Pools menu

2. Click **Volumes by Pool** to open the list of volumes separated by pool (Figure 10-20). Mark the volumes, right-click, and click **Map to Host**.

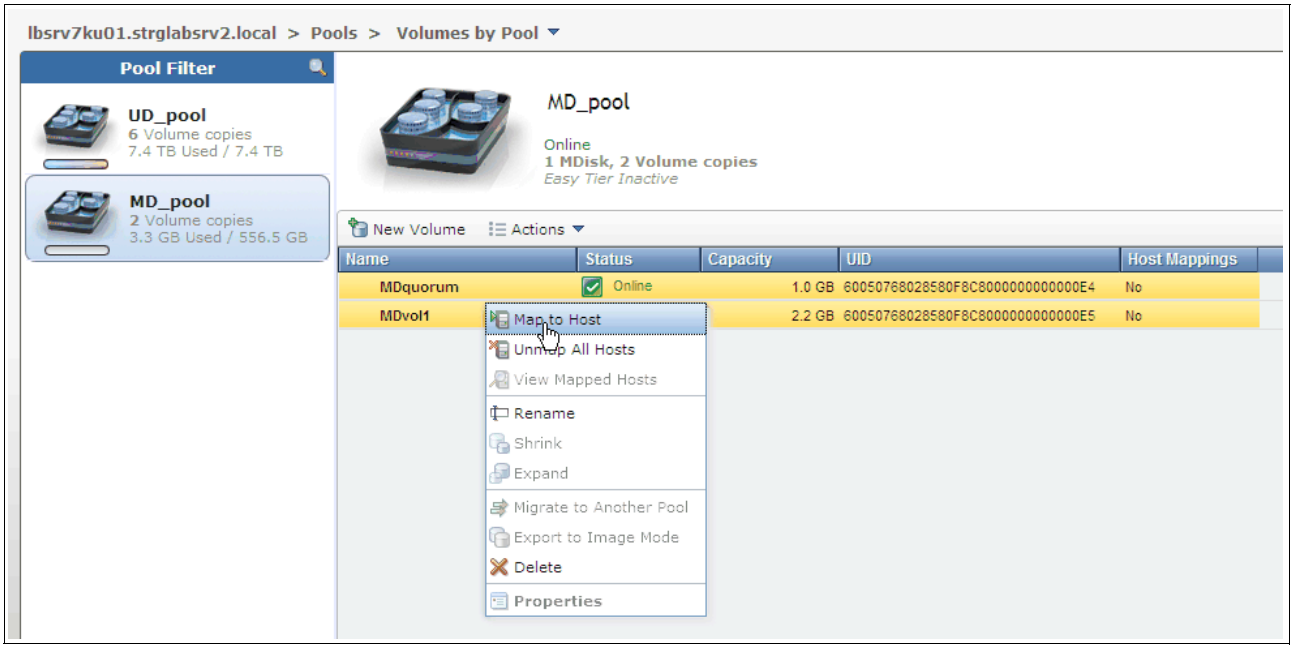

Figure 10-20 Volumes by pool

3. The Modify Host Mapping window opens. Click the **Host** drop-down menu to select the name of the host that you want to map volumes to. If the volumes are not highlighted in yellow, select volumes from the Unmapped Volumes pane and click the arrow (>) to move the volumes to the Volumes Mapped to the Host pane. Click **Map Volumes** or **Apply** (Figure 10-21).

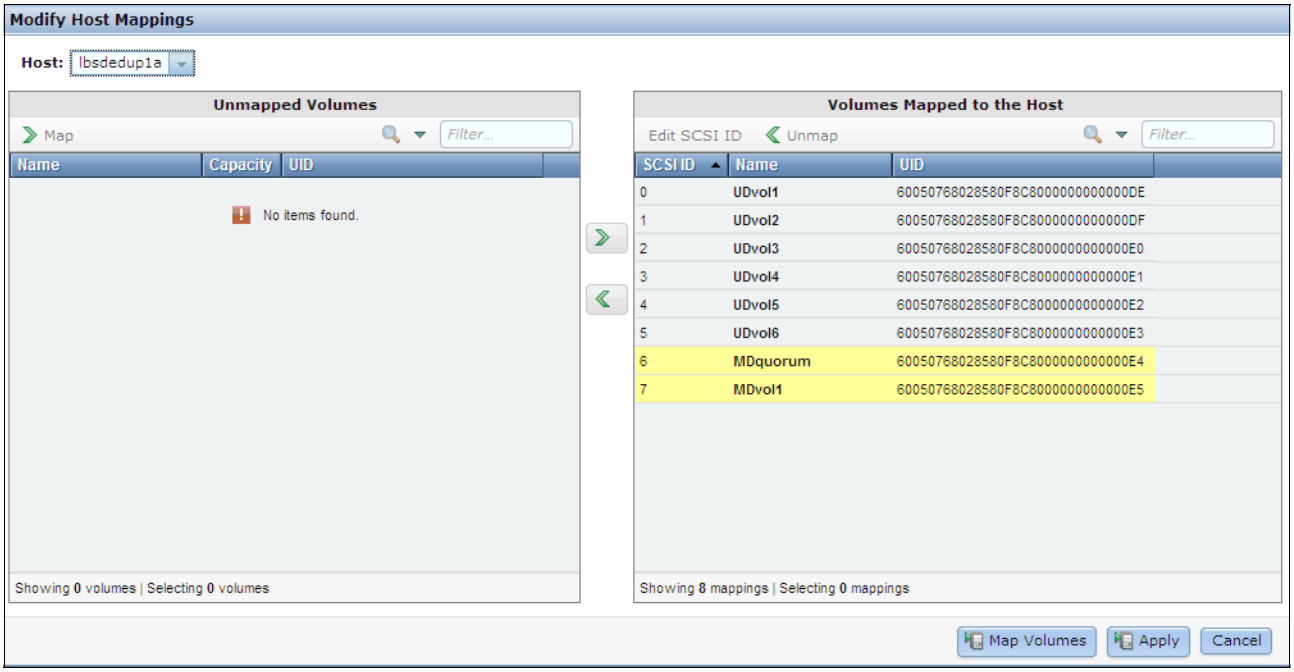

Figure 10-21 Modify mapping

Dual-node server: If you have a ProtecTIER dual-node server, you must select the other node and proceed with the volume mapping again for the second node. You might receive a warning that states that the specific volume is already mapped. Verify that the mapping is correct and continue (Figure 10-22).

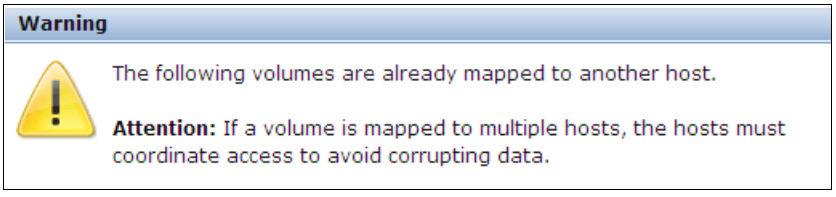

Figure 10-22 Volume mapping warning

10.6.6 Creating file systems and building the ProtecTIER repository

Beginning with ProtecTIER Version 3.3, the **`fsCreate`** tool is removed. All of the file system management is now integrated in to the ProtecTIER Configuration Menu.

Tip: If you have a version of ProtecTIER earlier than Version 3.2, the **fsCreate** tool can be used to create files systems automatically for repositories by completing the following steps:

1. Log in to a ProtecTIER server.
2. At the command prompt, enter **cd /opt/dtc/app/sbin** and press Enter.
3. To open the list of file systems that are already a part of the repository, enter **./fsCreate -r** and press Enter.
4. To open the list of available new multipath devices, enter **./fsCreate -u** and press Enter.
5. Create file systems on available multipath devices, register them to /etc/fstab, and mount them. At the command prompt of Server A, enter **./fsCreate -n** and press Enter.
6. For verification, open the GFS file systems that are not part of the repository by entering ./**fsCreate -g** and pressing Enter.

To create the file systems with ProtecTIER V3.2, complete the following steps:

1. Verify that the ProtecTIER node recognizes the volumes that are presented by the SAN Volume Controller, Storwize V7000, or Storwize V7000 Unified server. The **multipath -ll** command shows the LUNs and the paths that are connected to the storage subsystem. If the LUNs do not appear to be mapped to the storage, you must run a rescan on the FC adapters or reboot the node.

   Note: You can scan the FC adapters by running the following command:

   ```
   echo "- - -" > /sys/class/scsi_host/<host??>/scan
   ```

   In this command, **<host??>** should be replaced by each FC adapter port.

   You can also filter to see only part of the data, as shown in Example 10-9. The first command filters only the volumes of type 2145, which is the SAN Volume Controller or Storwize V7000 type. The second command shows the details of one of the devices.

   Example 10-9 multipath command output

```
[root@lbsdedup1a ~]# multipath -ll | grep 2145
mpath9 (360050768028580f8c8000000000000e1) dm-8 IBM,2145
mpath8 (360050768028580f8c8000000000000e0) dm-7 IBM,2145
mpath14 (360050768028580f8c8000000000000e7) dm-27 IBM,2145
mpath7 (360050768028580f8c8000000000000df) dm-6 IBM,2145
mpath6 (360050768028580f8c8000000000000de) dm-5 IBM,2145
mpath12 (360050768028580f8c8000000000000e4) dm-11 IBM,2145
mpath11 (360050768028580f8c8000000000000e3) dm-10 IBM,2145
mpath10 (360050768028580f8c8000000000000e2) dm-9 IBM,2145

[root@lbsdedup1a ~]# multipath -ll | grep -A11 mpath9
mpath9 (360050768028580f8c8000000000000e1) dm-4 IBM,2145
[size=1.2T][features=1 queue_if_no_path][hwhandler=0][rw]
\_ round-robin 0 [prio=200][active]
 \_ 2:0:1:3 sdak 66:64  [active][ready]
 \_ 3:0:2:3 sdal 66:80  [active][ready]
 \_ 1:0:1:3 sdam 66:96  [active][ready]
 \_ 4:0:1:3 sdbs 68:96  [active][ready]
```

```
\_ round-robin 0 [prio=40][enabled]
 \_ 4:0:0:3 sdbk 67:224 [active][ready]
 \_ 1:0:0:3 sde  8:64   [active][ready]
 \_ 3:0:1:3 sdk  8:160  [active][ready]
 \_ 2:0:0:3 sdq  65:0   [active][ready]
```

2. In the ProtecTIER Service Menu, select the options “1) ProtecTIER Configuration (...)”and then “6) File Systems Management (...)” (Example 10-10).

Example 10-10 ProtecTIER configuration menu

```
[root@lbsdedup1a ~]# su - ptconfig
+------------------------------------------------------------------------------+
| ProtecTIER Service Menu running on lbsdedup1a                                |
+------------------------------------------------------------------------------+
|  1) ProtecTIER Configuration (...)                                           |
|  2) Manage ProtecTIER services (...)                                         |
|  3) Health Monitoring (...)                                                  |
|  4) Problem Alerting (...)                                                   |
|  5) Version Information (...)                                                |
|  6) Generate a service report                                                |
|  7) Generate a system view                                                   |
|  8) Update ProtecTIER code                                                   |
|                                                                              |
|  E) Exit                                                                     |
+------------------------------------------------------------------------------+
Your choice? 1

+------------------------------------------------------------------------------+
| ProtecTIER Service Menu running on lbsdedup1a                                |
| ProtecTIER Configuration (...)                                               |
+------------------------------------------------------------------------------+
|  1) Configure ProtecTIER node                                                |
|  2) Recover Configuration for a replaced server                              |
|  3) Configure machine serial number for a replaced server                    |
|  4) Configure RAS                                                            |
|  5) Update Time, Date, Timezone and Timeserver(s)                            |
|  6) File Systems Management (...)                                            |
|  7) Configure replication (...)                                              |
|  8) IP Network configuration (...)                                           |
|  9) Update Firmware                                                          |
| 10) Update the System's name                                                 |
| 11) Validate configuration                                                   |
|                                                                              |
|  B) Back                                                                     |
|  E) Exit                                                                     |
+------------------------------------------------------------------------------+
Your choice? 6
```

3. The File Systems Management menu opens. You have options to display information about the devices or file systems configuration. You have choices to configure the file systems on all available devices or for a single device. Example 10-11 shows the creation of a file system on a single unused device.

Example 10-11 File System Management menu

```
--------------------------------------------------------------------------------+
| ProtecTIER Service Menu running on lbsdedup1a                                  |
| ProtecTIER Configuration (...)                                                 |
| File Systems Management (...)                                                  |
+--------------------------------------------------------------------------------+
|  1) Configure file systems on all available devices                            |
|  2) Create file system(s) on a single unused device                            |
|  3) Extend a file system with a new unused device                              |
|  4) Update /etc/fstab                                                          |
|  5) Display configured devices                                                 |
|  6) Display unused devices                                                     |
|  7) Display GFS repository file systems                                        |
|  8) Display unused GFS file systems                                            |
|  9) Increase capacity completion (applicable for a second cluster node)        |
|                                                                                |
|  B) Back                                                                       |
|  E) Exit                                                                       |
+--------------------------------------------------------------------------------+
Your choice? 2
Begin Processing Procedure
Display of the available devices
Device:         Size:           Status
1. mpath0       1024.00M        Unused
2. mpath10      1286144.00M     Unused
3. mpath14      568832.00M      Unused
4. mpath11      1286144.00M     Unused
5. mpath6      1284864.00M     Unused
6. mpath7      1286144.00M     Unused
7. mpath8      1286144.00M     Unused
8. mpath9      1286144.00M     Unused

Please select device: 3
Please enter number of file systems[1-4]:1
Last file system - using the remaining free space
The file system to be created will have 568832.00 MiB size
Please confirm:? (yes|no) yes
Creating physical volume                                        [ Done ]
Creating volume group                                           [ Done ]
Creating logical volume                                         [ Done ]
Creating file system                                            [ Done ]

End Processing Procedure Successfully
```

4. After you create the file system for all the devices that are needed, you can go to the ProtecTIER Manager, add the node to it, and then select the **Repository** menu to create the repository. For more information, see *IBM System Storage TS7650 ProtecTIER Deduplication Appliance Installation Roadmap Guide*; GA32-0921.

10.6.7 Expanding the repository

When you expand the repository, use the same spindle type and quantity of RAID groups for metadata or user data. For example, if the original two metadata LUNs were built on RAID 4+4 groups, then the added metadata RAID groups must be at least 4+4 to maintain the same level of performance. Using storage from 2+2 or 4+1 RAID groups, for example, for the expansion might degrade performance because of an input/output operation per second (IOPS) bottleneck.

Important: The total number of volumes for both the ProtecTIER repository for metadata (MD) and user data (UD) should not exceed 170. Each volume size should not exceed 6 TB to comply with the 1 PB ProtecTIER repository size limit.

As of ProtecTIER V3.2, the 8 TB restriction is removed. When the individual disk size is large (for example, 2 or 3 TB), use RAID 6 with 6+2 disk members.

11

IBM XIV Storage System

This chapter addresses specific considerations for using the IBM XIV® Storage System as storage for ProtecTIER servers.

This chapter describes the following topics:

- An overview of the XIV Storage System hardware
- Fibre Channel cabling and zoning configuration for maximum performance with an XIV Storage System
- Configuring an XIV Storage System for a ProtecTIER server

11.1 XIV Storage System hardware

Figure 11-1 shows an example of the IBM XIV Storage System hardware and supported modules.

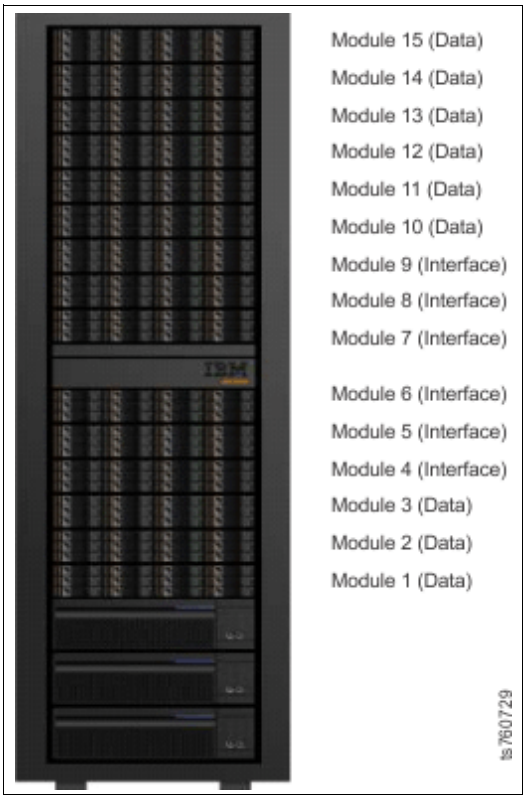

Figure 11-1 XIV Storage System hardware

XIV Storage System supports configurations of 6, 9, 10, 11, 12 13, 14, or 15 modules (Table 11-1):

- Modules 1 - 3 and 10 - 15 are disks only and are called data modules.
- Modules 4 - 9 have disks and host interfaces and are called interface modules.

Table 11-1 Configurations of modules

Number of modules	6	9	10	11	12	13 - 15
Interface Module 9 state	Empty	Disabled	Disabled	Enabled	Enabled	Enabled
Interface Module 8 state	Empty	Enabled	Enabled	Enabled	Enabled	Enabled
Interface Module 7 state	Empty	Enabled	Enabled	Enabled	Enabled	Enabled
Interface Module 6 state	Disabled	Disabled	Disabled	Disabled	Enabled	Enabled
Interface Module 5 state	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled

Number of modules	6	9	10	11	12	13 - 15
Interface Module 4 state	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
FC ports	8	16	16	20	20	24
Net capacity (decimal) - 1T B	27 TB	43 TB	50 TB	54 TB	61 TB	66 TB
Net capacity (decimal) - 2 TB	55 TB	87 TB	102 TB	111 TB	125 TB	134 TB

11.2 Fibre Channel switch cabling

For maximum performance with an XIV Storage System, connect all available XIV Storage System Interface Modules and use all of the back-end ProtecTier ports. For redundancy, connect Fibre Channel cables from the ProtecTIER server to two Fibre Channel (FC) switched fabrics.

If a single XIV Storage System is being connected, each Fibre Channel switched fabric must have six available ports for Fibre Channel cable attachment to the XIV Storage System. Generally, there are two connections that are used for each interface module in XIV Storage System. Typically, XIV Storage System Interface Module port 1 is used for Fibre Channel switch 1, and port 3 for switch 2 (Figure 11-2).

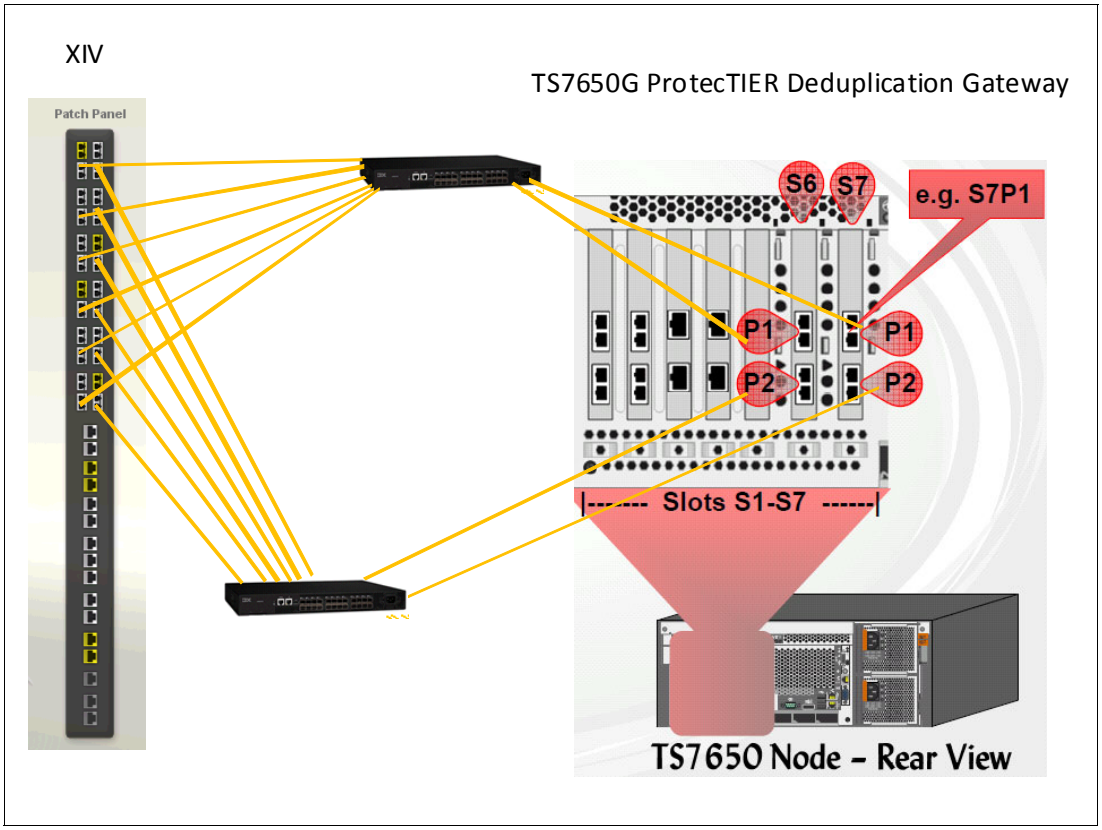

Figure 11-2 Cable diagram for connecting a TS7650G to XIV Storage System

When you use a partially configured XIV Storage System rack, see Figure 11-2 on page 191 to locate the available FC ports.

11.2.1 Zoning configuration

For each ProtecTIER disk attachment port, multiple XIV Storage System host ports are configured into separate isolated zone pairing in a 1:1 manner:

- All XIV Storage System Interface Modules on port 1 are zoned to the ProtecTIER host bus adapter (HBA) in slot 6, port 1 and HBA in slot 7, port 1.
- All XIV Storage System Interface Modules in port 3 are zoned to the ProtecTIER HBA in slot 6, port 2 and HBA in slot 7, port 2.

Information: We suggest that you connect port 1 and port 3 of the XIV Storage System Interface Modules, as they are predefined for host I/O. Other ports might be predefined for XIV Storage System replication. Also, ports 1 and 3 are distributed across the XIV internal dual-port adapters. So, using ports 1 and 3 minimizes administrative impact and ensures protection against an adapter failure of XIV I/O modules.

Each Interface Module in the XIV Storage System has a connection with both ProtecTIER HBAs. A typical ProtecTIER configuration uses 1:1 zoning (one initiator and one target in each zone) to create zones. These zones connect a single ProtecTIER server with a 15 module XIV Storage System with all six Interface Modules (Example 11-1).

Example 11-1 Zoning example for an XIV Storage System attachment

```
Switch 1:
Zone 01: PT_S6P1, XIV_Module4Port1
Zone 02: PT_S6P1, XIV_Module6Port1
Zone 03: PT_S6P1, XIV_Module8Port1
Zone 04: PT_S7P1, XIV_Module5Port1
Zone 05: PT_S7P1, XIV_Module7Port1
Zone 06: PT_S7P1, XIV_Module9Port1

Switch 02:
Zone 01: PT_S6P2, XIV_Module4Port3
Zone 02: PT_S6P2, XIV_Module6Port3
Zone 03: PT_S6P2, XIV_Module8Port3
Zone 04: PT_S7P2, XIV_Module5Port3
Zone 05: PT_S7P2, XIV_Module7Port3
Zone 06: PT_S7P2, XIV_Module9Port3
```

This example has the following characteristics:

- Each ProtecTIER back-end HBA port sees three XIV Storage System Interface Modules.
- Each XIV Storage System Interface Module is connected redundantly to two separate ProtecTIER back-end HBA ports.
- There are 12 paths (4 x 3) to one volume from a single ProtecTIER server.

11.2.2 Configuring the XIV Storage System for a ProtecTIER server

An IBM System Service Representative (SSR) uses the ProtecTIER Capacity Planning Tool to size the ProtecTIER repository metadata and user data. Capacity planning varies because it depends heavily on your type of data and expected deduplication ratio. The planning tool output includes the detailed information about all volume sizes and capacities for your specific ProtecTIER installation. If you do not have this information, contact your IBM SSR to get it.

Tip: The Factoring Ratio number is directly related to the size of the metadata volumes, and can be estimated with the IBM ProtecTIER Performance Calculator.

Take the maximum throughput and repository size values into account during the calculations for both the initial installation and future growth.

You must configure the XIV Storage System before the ProtecTIER system is installed by an SSR. To configure the system, complete the following steps:

1. Configure an XIV Storage System for the ProtecTIER system. Set the snapshot space to zero because creating snapshots on XIV Storage System is not supported by a ProtecTIER server.
2. Set up volumes in the XIV Storage System with the ProtecTIER Capacity Planning Tool and the Create repository planning wizard output. Starting with ProtecTIER V3.2.0, you can select XIV MD Device 8 +8 for the MD RAID Configuration from the ProtecTIER Manager when you are using the Create repository planning wizard. The ProtecTIER Capacity Planning Tool output gives you the metadata volume size and the size of the 32 user data volumes. Configure a Quorum volume with a minimum of 1 GB (17 GB for XIV Storage System because that is the smallest volume size that can be created) as well, in case the solution needs more ProtecTIER servers in the future.

   Important: Use the XIV Volume Sizing Spreadsheet Tool to calculate the size of the XIV volumes accurately to ensure that you get the expected volume sizes. This tool can be found at the following website:

   `http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TD105544`

3. Map the volumes to the ProtecTIER server or ProtecTIER cluster.

Example of configuring an XIV Storage System

If you want to set up a ProtecTIER environment with a 79 TB XIV Storage System and a deduplication factoring ratio of 12, use the following volumes sizes:

- 2 x 1571 GB volumes for metadata: Make these volumes equal to each other, and nearest to the XIV Storage System allocation size, in this case, 1583 (see Figure 11-4 on page 195).
- 1 x 17 GB volume for Quorum (see Figure 11-5 on page 195). It must be 17 GB because that is the XIV Storage System minimum size.
- 32 x <*Remaining Pool Space available*>, which is 75440. Dividing 75440 by 32 means that user data LUNs on the XIV Storage System should be 2357 GB each (see Figure 11-6 on page 196).

Memory totals: When you have an XIV Storage System Gen 3 full rack, you can have up to 243 TB of available space. With that configuration, and also if you have more than one XIV Storage System connected to the ProtecTIER server, you might need more than 32 LUNS. For the best performance, do not exceed the LUN size of 6 TB. For example, two full racks equal 486 TB. Dividing this number by 6 (as in the 6 TB recommended LUN size), you need roughly 81 LUNs. Create the necessary metadata LUNs of the recommended size and the 6 TB LUNs for user data.

Always upgrade the firmware of your XIV Storage System to the latest supported level. For more information, see the System Storage Interoperation Center (SSIC) website at:

`http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss`

For XIV Storage System series, capacity, and connectivity details, go to the following website:

`http://www-03.ibm.com/systems/storage/disk/xiv/specifications.html`

As illustrated in the following examples, the XIV Storage System V3.0 client GUI makes this calculation easy for you. Enter the number of volumes to create, then drag the slider to the right to fill the entire pool. The GUI automatically calculates the appropriate equivalent amount.

Create a pool size for the capacity that you want to use for the ProtecTIER Deduplication Gateway with the XIV Storage System GUI, as shown in Figure 11-3. Normally, this would be 100% of the available space on the XIV. Create one single pool only; there is no benefit to having multiple pools.

Information: On older firmware levels of XIV, you might be required to have multiple pools to allocate all the space that you want to use with ProtecTIER. Besides a slightly increased administrative impact, there is no drawback to doing this action.

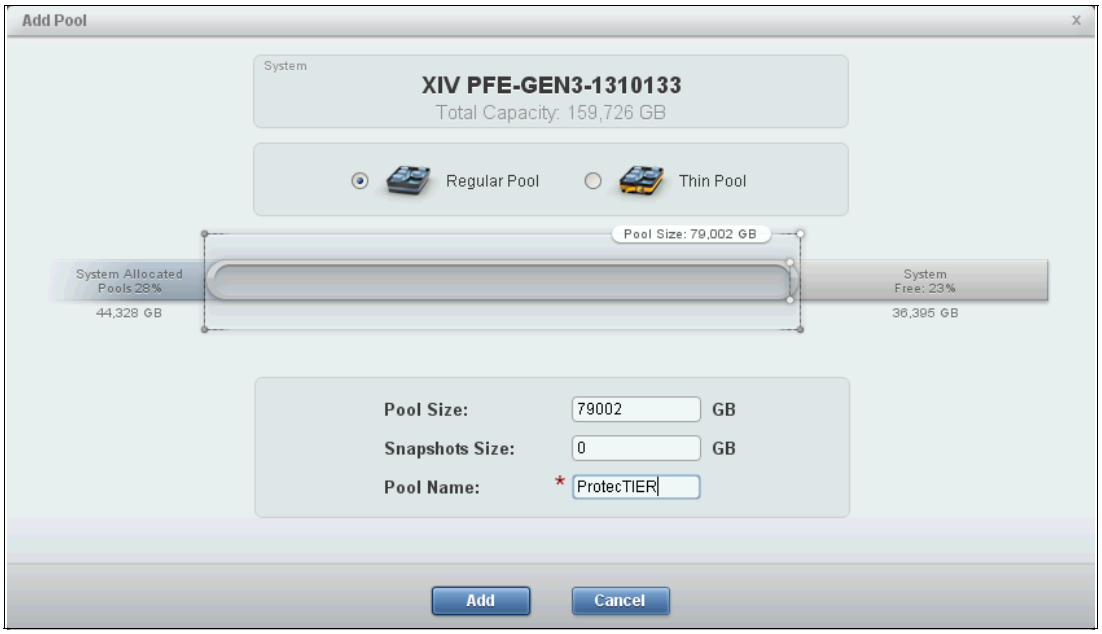

Figure 11-3 Creating a pool

Tip: Use a regular pool and zero the snapshot reserve space. Snapshots and thin provisioning are not supported when XIV Storage System is used with a ProtecTIER server.

Figure 11-4 shows the creation of the metadata volumes.

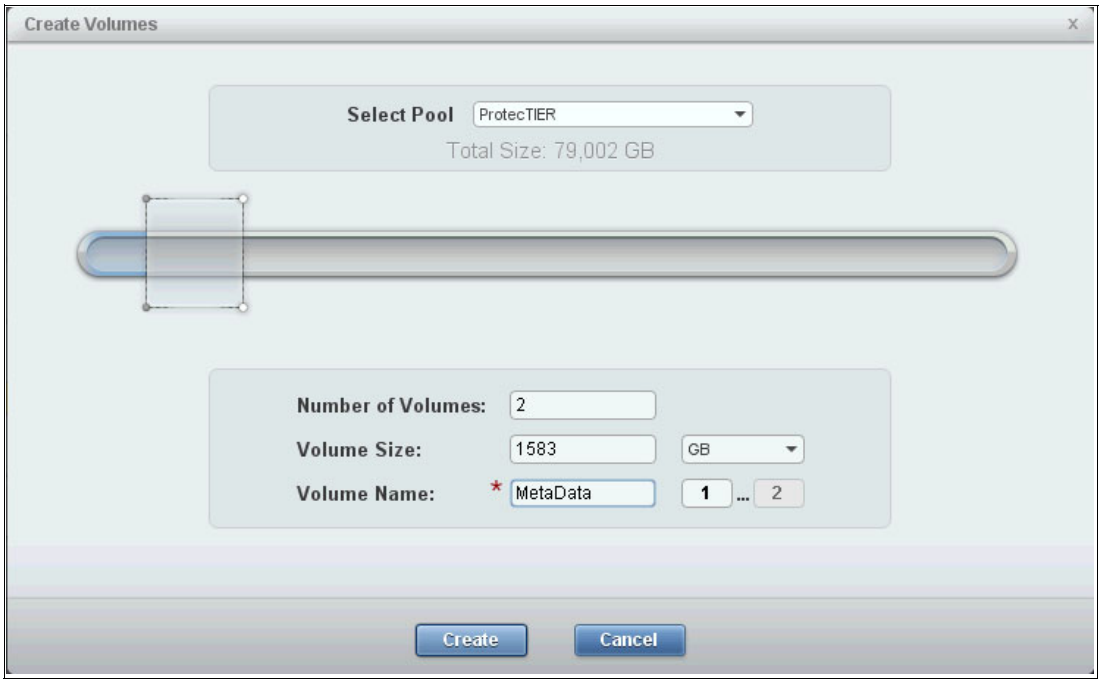

Figure 11-4 Creating metadata volumes

Figure 11-5 shows the creation of the quorum volume.

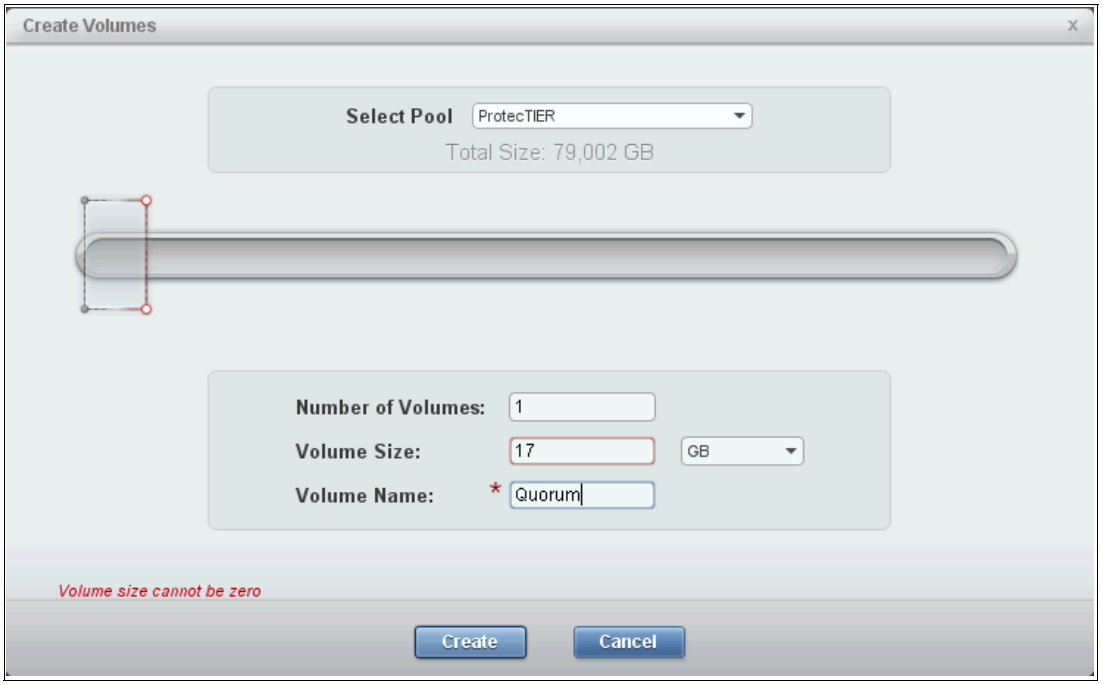

Figure 11-5 Creating a quorum volume

Figure 11-6 shows the creation of volumes for user data. The arrows show dragging the slider to use all of the pool. This action automatically calculates the appropriate size for all volumes.

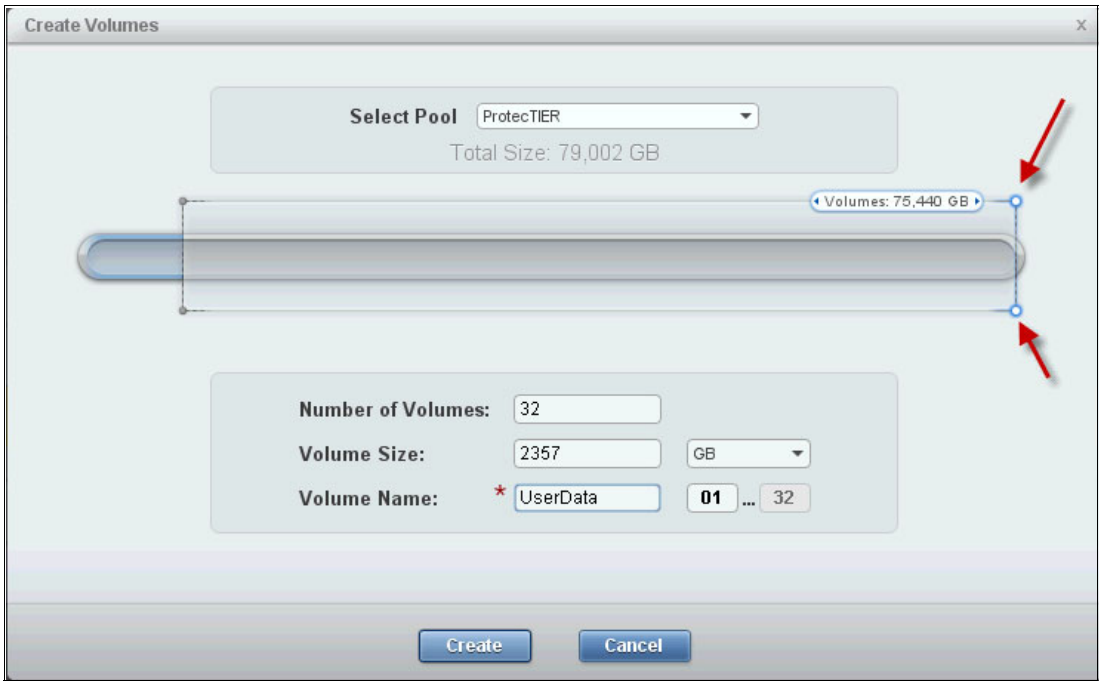

Figure 11-6 Creating user data volumes

If you have a ProtecTIER cluster (two ProtecTIER servers in a high availability solution), complete the following steps:

1. Create a cluster group (Figure 11-7).

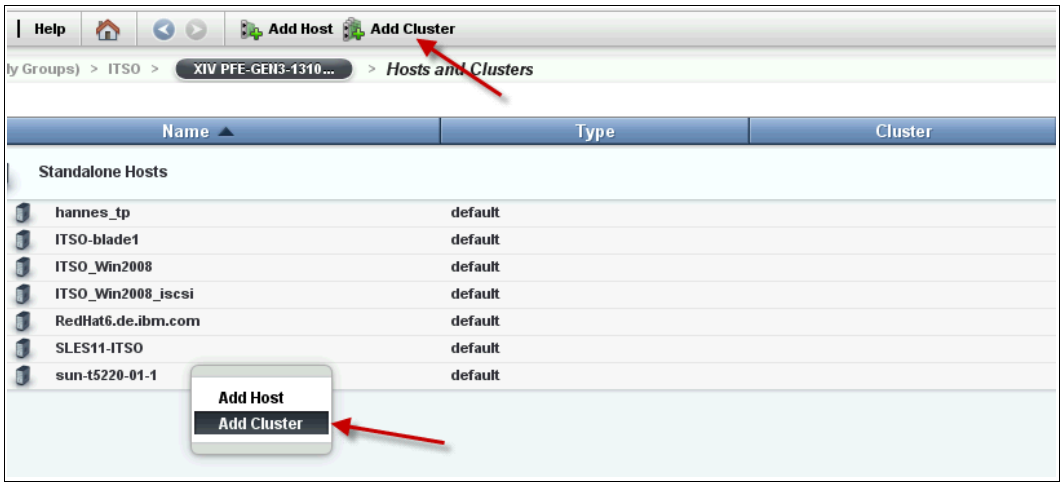

Figure 11-7 Creating a cluster group

2. Add a host that is defined for each node to that cluster group.

3. Create a cluster definition for the highly available ProtecTIER cluster (Figure 11-8).

Figure 11-8 Adding a cluster definition

4. Right-click the cluster and select **Add Host** (Figure 11-9).

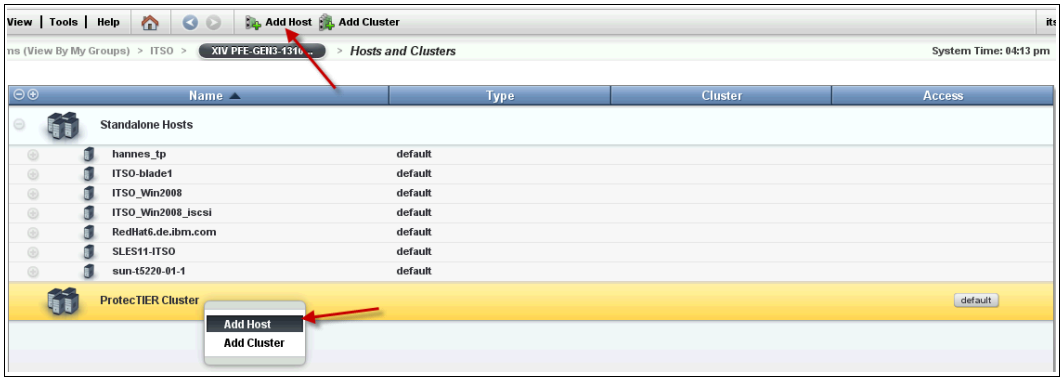

Figure 11-9 Adding a host to the cluster

5. Enter the information for the new ProtecTIER host and click **Add** (Figure 11-10).

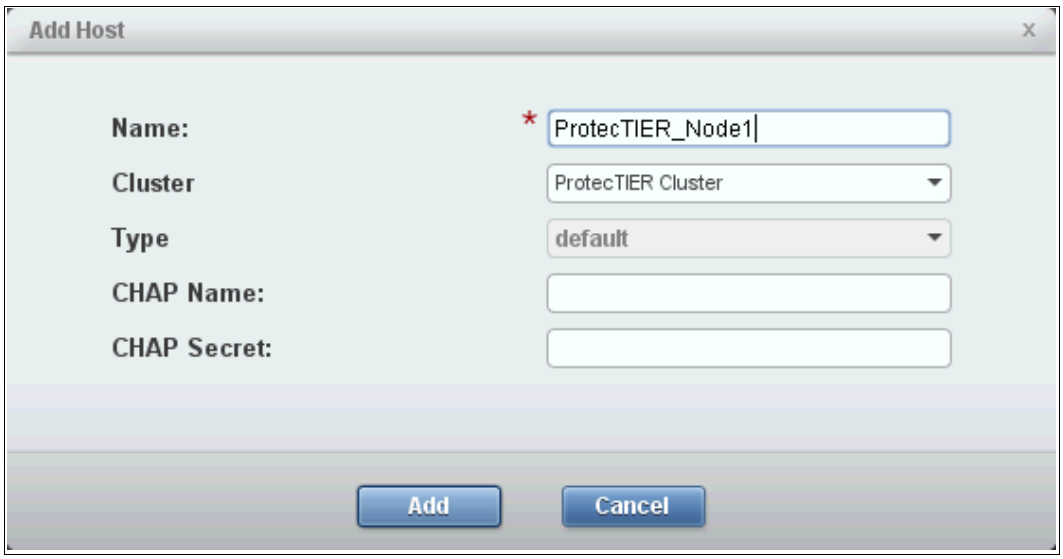

Figure 11-10 Adding a ProtecTIER host to a cluster

6. Find the worldwide port names (WWPNs) of the ProtecTIER servers. WWPNs can be found in the name server of the Fibre Channel switch. If zoning is in place, they are selectable from the menu. Alternatively, they can also be found in the BIOS of the HBA cards and then entered manually in to the XIV Storage System GUI.
7. Add the WWPNs to the ProtecTIER servers (Figure 11-11).

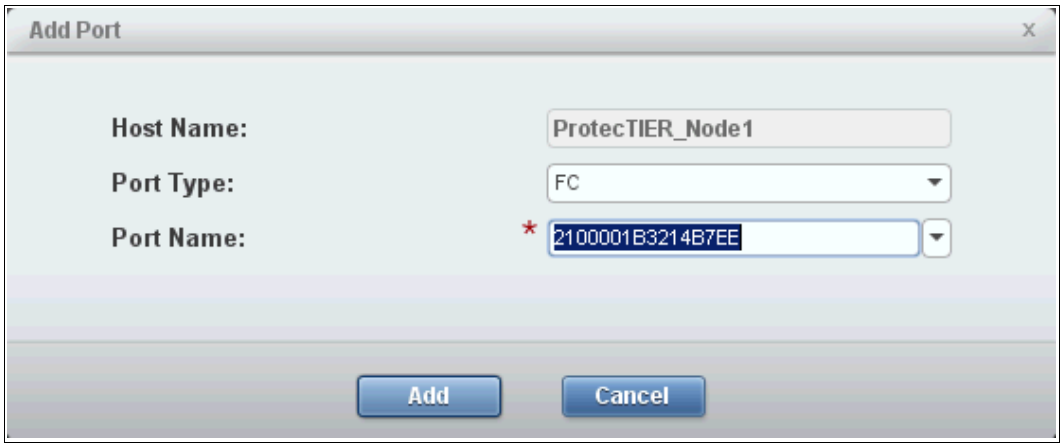

Figure 11-11 Adding the WWPN of ProtecTIER server 1 to the cluster

Figure 11-12 shows the WWPNs that are added to the hosts.

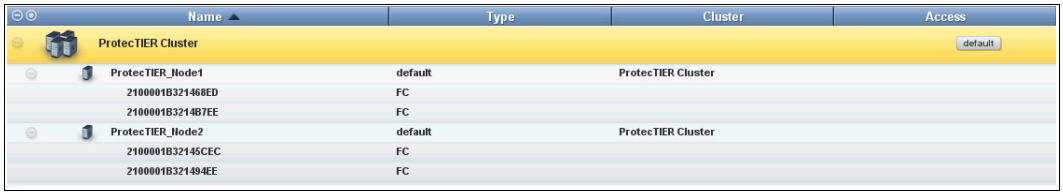

Figure 11-12 ProtecTIER WWPNs added to host and cluster definitions

8. Map the volumes to the ProtecTIER cluster. In the XIV Storage System GUI, right-click the cluster name or on the host if you have only one ProtecTIER server, and select **Modify LUN Mapping**. Figure 11-13 shows you what the mapping view looks like.

Tip: If you have only one ProtecTIER server, map the volumes directly to the ProtecTIER server.

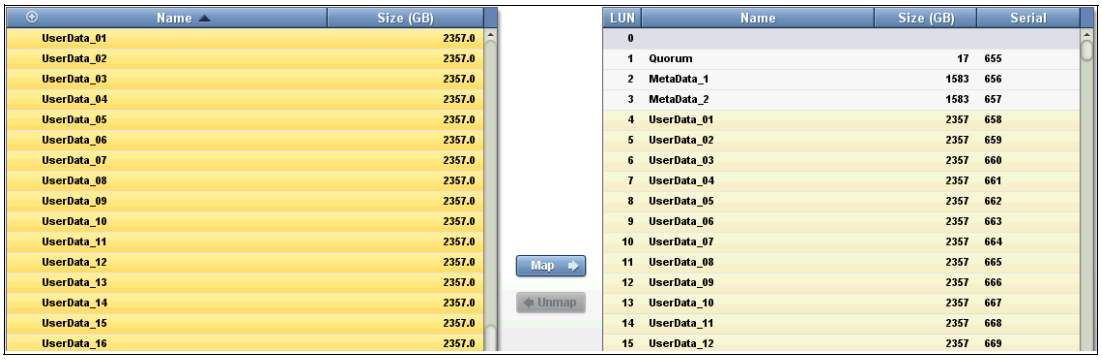

Figure 11-13 Mapping LUNs to the ProtecTIER cluster

12

IBM System Storage DS8000

This chapter addresses specific considerations for using the IBM System Storage DS8000 as a storage system for ProtecTIER servers.

This chapter describes the following topics:

- An overview of the DS8000 family and recommended RAID levels
- General considerations for planning tools, metadata, user data, firmware levels, and replication
- When to use rotate extents (striping)

12.1 DS8000 series overview

The DS8000 family is a high performance, high capacity, and resilient series of disk storage systems. It offers high availability, multiplatform support, and simplified management tools to provide a cost-effective path to an on-demand configuration.

12.1.1 Disk drives

This section briefly describes the DS8000 family products and available drives for the DS8000 products.

For more information, see the IBM system storage website:

http://www-03.ibm.com/systems/storage/disk/ds8000/overview.html

IBM System Storage DS8870

The DS8870 supports an intermix of 146 and 300 GB (15 K RPM) Enterprise disk drives and 600 or 900 GB (10 K RPM) disk drives. The DS8870 also supports Full Disk Encryption (FDE) disk drives. 3 TB (7200 K RPM) Near line (SATA) disk drives are also supported. The DS8700 also supports 400 GB solid-state drives (SSDs). 73 GB and 146 GB SSDs are also supported, but they were withdrawn from marketing as of the latest microcode release.

Solid-state drives

Solid-state drives (SSDs) are the best choice for I/O-intensive workloads. They provide up to 100 times the throughput and 10 times lower response time than 15 K RPM spinning disks. They also use less power than traditional spinning disks.

SAS and Fibre Channel disk drives

Because model 8800 Enterprise SAS drives are supported, the 8700 model supports FC drives high performance, reliability, availability, and serviceability. Enterprise drives rotate at 15 K or 10 K RPM. If an application requires high performance data throughput and continuous, intensive I/O operations, enterprise drives are the best price performance option.

SATA

The 3 TB near-line drives are both the largest and slowest of the drives available for the DS8870. Near-line drives are a cost-efficient storage option for lower intensity storage workloads and are available since the DS8870. Because of the lower usage and the potential for drive protection throttling, these drives are a not the optimal choice for high performance or I/O-intensive applications.

12.1.2 Host adapters

Each DS8870 Fibre Channel adapter offers 8 Gbps Fibre Channel Adapters; the DS8800 offers up to four 4 Gbps or 8 Gbps Fibre Channel ports. Each 4 Gbps port independently auto-negotiates to either 1, 2, or 4 Gbps link speed. Each 8 Gbps port independently auto-negotiates to either 2, 4, or 8 Gbps link speed. Each of the four ports on a DS8800 adapter can also independently be either Fibre Channel protocol (FCP) or IBM FICON®.

Each DS8870 Fibre Channel adapter offers four or eight 8 Gbps Fibre Channel ports. Each 8 Gbps port independently auto-negotiates to either 2, 4, or 8 Gbps link speed. Each of the ports on a DS8870 host adapter can also independently be either Fibre Channel protocol (FCP) or FICON.

12.1.3 RAID levels

The DS8000 series offers RAID 5, RAID 6, and RAID 10 levels. There are some limitations:

- RAID 10 for SSD is not standard and is only supported through a special RPQ.
- SSD disks cannot be configured in RAID 6.
- Nearline disks cannot be configured in RAID 5 and RAID 10.

RAID 5

Normally, RAID 5 is used because it provides good performance for random and sequential workloads and it does not need much more storage for redundancy (one parity drive). The DS8000 series can detect sequential workload. When a complete stripe is in cache for destaging, the DS8000 series switches to a RAID 3 like algorithm. Because a complete stripe must be destaged, the old data and parity do not need to be read. Instead, the new parity is calculated across the stripe, and the data and parity are destaged to disk. This action provides good sequential performance. A random write causes a cache hit, but the I/O is not complete until a copy of the write data is put in non-volatile storage (NVS). When data is destaged to disk, a write in RAID 5 causes four disk operations, the so-called write penalty:

- Old data and the old parity information must be read.
- New parity is calculated in the device adapter.
- Data and parity are written to disk.

Most of this activity is hidden to the server or host because the I/O is complete when data enters the cache and NVS.

RAID 6

RAID 6 is an option that increases data fault tolerance. It allows additional failure, compared to RAID 5, by using a second independent distributed parity scheme (dual parity). RAID 6 provides a read performance that is similar to RAID 5, but has more write penalty than RAID 5 because it must write a second parity stripe. RAID 6 should be considered in situations where you would consider RAID 5, but there is a demand for increased reliability. RAID 6 is designed for protection during longer rebuild times on larger capacity drives to cope with the risk of having a second drive failure within a rank while the failed drive is being rebuilt. It has the following characteristics:

- Sequential read of about 99% x RAID 5 rate.
- Sequential write of about 65% x RAID 5 rate.
- Random 4 K 70% R/30% W IOPS of about 55% x RAID 5 rate.
- The performance is degraded with two failing disks.

RAID 10

A workload that is dominated by random writes benefits from RAID 10. Here, data is striped across several disks and concurrently mirrored to another set of disks. A write causes only two disk operations compared to the four operations of RAID 5. However, you need nearly twice as many disk drives for the same capacity compared to RAID 5. Thus, for twice the number of drives (and probably cost), you can do four times more random writes, so it is worth considering using RAID 10 for high performance random write workloads.

12.2 General considerations

The following section describes general considerations for the ProtecTIER Capacity Planning Tool and some guidelines for the usage and setup of metadata and user data in your DS8000.

12.2.1 Planning tools

An IBM Field Technical Sales Support expert uses the ProtecTIER Capacity Planning Tool to size the ProtecTIER repository metadata and user data. Capacity planning is always different because it depends heavily on your type of data and the expected deduplication ratio. The planning tool output includes detailed information about all volume sizes and capacities for your specific ProtecTIER installation. If you do not have this information, contact your IBM sales representative to get it.

Tip: The Factoring Ratio number is directly related to the size of the metadata volumes, and can be estimated with the IBM ProtecTIER Performance Calculator.

Be sure to take the Max Throughput and Repository Size values into account during the calculations for both the initial install and future growth.

12.2.2 Metadata

Consider the following items about metadata:

- Use the ProtecTIER Capacity Planning Tool and the Create repository planning wizard output to determine the metadata requirements for your environment.
- You must use RAID 10 for metadata. Use high-performance and high-reliability enterprise class disks for metadata RAID 10 arrays.
- When possible, do not use SATA disks, as RAID 10 is not supported by SATA disks and because ProtecTIER metadata has a heavily random read I/O characteristic. If you require a large physical repository and have only SATA drives available in the storage subsystem, you must use the **`rotateextents`** feature for all LUNs to ensure the equal distribution of ProtecTIER workload across all available resources. For more information about the usage of the **`rotateextents`** feature, see 12.3, “Rotate extents: Striping” on page 203.

12.2.3 User data

Consider the following items about user data:

- ProtecTIER is a random-read application. 80 - 90% of I/O in a typical ProtecTIER environment is random read. Implement suitable performance optimizations and tuning as recommended for this I/O profile.
- For SATA drives or large capacity disk drives, use RAID 6 with 6 + 2 disk members for increased availability and faster recovery from disk failure.
- Do not intermix arrays with different disk types within the metadata and the user data because smaller disk types hold back the performance of larger disk types and degrade the overall system throughput.
- For smaller capacity FC or SAS drives, use RAID 5 with at least five disk members per group.
- Create an even number of LUNs in each pool.

- Use LUNS that are 2 - 6 TB in size to reduce the number of file systems and produce optimal performance.

> **Important:** With ProtecTIER and DS8000, create LUNs that are all the same size to avoid performance degradation.
>
> Starting with ProtecTIER V3.2, the management of LUNs greater than 8 TB is improved. When ProtecTIER uses LUNs greater than 8 TB, it splits them in to logical volumes of smaller size, which means that you can work with LUNs greater than 8 TB. There is no benefit in performance in doing this action.
>
> You should always use RAID 6 for SATA or NL-SAS drives for the user data LUNs. With SAS drives, only RAID 6 is supported.

- Do not use thin provisioning.

12.2.4 Firmware levels

Ensure that you are using supported firmware levels. When possible, use the latest supported level. For compatibility information, see the following website:

`http://www.ibm.com/systems/support/storage/config/ssic/displayesssearchwithoutjs.wss?start_over=yes`

12.2.5 Replication

Do not use disk-based replication because disk-based replication features are not supported by the ProtecTIER product. Instead of using the replication feature of the DS8000, use the ProtecTIER native replication, which is available with Version 2.5 and higher. For more information about replication, see Part 5, “Replication and disaster recovery” on page 369.

12.3 Rotate extents: Striping

The following section describes the **`rotateexts`** feature and when to use or not use it in your ProtecTIER environment. The rotate extents **`rotateexts`** feature is also referred to as Storage Pool Striping (SPS.) In addition to the rotate volumes extent allocation method, which remains the default, the rotate extents algorithm is an additional option of the **`mkfbvol`** command. The rotate extents algorithm evenly distributes the extents of a single volume across all the ranks within a multirank extent pool. This algorithm provides the maximum granularity that is available on the DS8000 (that is, on the extent level that is equal to 1 GB for FB volumes), spreading each single volume across multiple ranks, and evenly balancing the workload within an extent pool.

Depending on the type and size of disks you use within your DS8000 server and your planned array size to create your ProtecTIER repository, you can consider using **`rotateexts`**. Because the ProtecTIER product already does a good job at equally distributing the load to the back-end disks, there are some potential scenarios where you should not use **`rotateexts`**.

Attention: For ProtecTIER performance, the most critical item is the number of spinning disks in the back end. The spindle count has a direct impact on the ProtecTIER performance. Sharing disk arrays between ProtecTIER and some other workload is *not* recommended. This situation directly impacts your ability to reach your wanted performance.

Because you do not share disks between ProtecTIER and other workloads, assigning the full array capacity to the ProtecTIER server is recommended.

ProtecTIER prefers a high number of LUNs as back-end storage. For considerations about the potential number of arrays in the back end, see 2.5.2, “The number 32: The ProtecTIER product is not physical tape” on page 28. The recommended LUN size should not exceed 6 TB.

With these considerations, you can easily decide when to use **rotateexts** and when not to use it. Within the DS8000, the following array types should be used with ProtecTIER, taking the host spare (S) requirements into account:

- 4+4 RAID 10
- 7+1 RAID 5 or 6+1+S RAID 5
- 6+2 RAID 6 or 5+2+S RAID 6

Tip: The DS8000 server creates four spares per device adapter (DA) pair. If you have a spare requirement when you create your RAID 10 arrays, you must create 3+3+2S RAID 10 arrays. You should redesign your layout to allow all metadata arrays to be 4+4 RAID 10 arrays only. Do not create 3+3+2S RAID 10 arrays for DS8000 repositories.

If you use 3 TB SATA disk to create your arrays, you could have the following array dimensions:

- Creating a 6+2 RAID 6 with a 3 TB disk results in a potential LUN size of 18 TB.
- Creating a 5+1+H RAID 5 with a 3 TB disk results in a potential LUN size of 15 TB.

These LUN sizes exceed the recommended LUN size for ProtecTIER of 6 TB. In this case, you should use **rotateexts** to equally distribute the ProtecTIER load to the DS8000 equally across all available resources. The rotate extend feature helps you create smaller LUNs when ProtecTIER code older than Version 3.2 is installed. Starting with ProtecTIER V3.2, it is possible to create multiple partitions on a LUN in order to cope with the 6 TB file system size limitation in the ProtecTIER system.

Important: ProtecTIER metadata that is on the RAID 10 arrays has a heavily random write I/O characteristic. ProtecTIER user data that is on RAID 5 or RAID 6 arrays has a heavily random read I/O characteristic. You should use high-performance and high-reliability enterprise-class disk for your metadata RAID 10 arrays.

12.3.1 When not to use rotate extents

Rotate extents (**rotateexts**) is a useful DS8000 feature that can be used to achieve great flexibility and performance with minimal effort. ProtecTIER comes with special requirements where using the **rotateexts** feature does not always make sense.

Because the ProtecTIER product works below the 6 TB LUN size limit and uses repositories that are based on high-performance disks, **rotateexts** does not necessarily provide any performance gains.

The ProtecTIER product does a great job of equally distributing its load to its back-end disks and directly benefits all available resources, even without **rotateexts**. The typical ProtecTIER write pattern does not create hot spots on the back-end disk, so **rotateexts** does not contribute to better I/O performance.

If the repository needs to be grown, the addition of more disks to already existing extent pools, or the addition of another extent pool with all new disks, creates storage that has different performance capabilities than the already existing ones. Adding dedicated arrays with their specific performance characteristics allows the ProtecTIER server to equally distribute all data across all LUNs. So, all back-end LUNs have the same performance characteristics and therefore behave as expected.

Consider the following example. You want to use 300 GB 15 K RPM FC drives for metadata and user data within your DS8000. To reach the wanted performance, you need four 4+4 RAID 10 arrays for metadata. Because you use Fibre Channel drives, go with RAID 5 arrays and configure all user data file systems with 6+1+H RAID 5 or 7+1 RAID 5. With this approach, you do not create RAID 10 arrays with ranks that have a hot spare requirement.

As shown in Figure 12-1, the following example needs some work to be aligned with best practices.

A16	Normal	RAID 5 (6+P+S)	Assigned	R16	4	300	15 Enterprise
A17	Normal	RAID 5 (6+P+S)	Assigned	R17	4	300	15 Enterprise
A18	Normal	RAID 5 (6+P+S)	Assigned	R18	4	300	15 Enterprise
A19	Normal	RAID 5 (6+P+S)	Assigned	R19	4	300	15 Enterprise
A20	Normal	RAID 5 (7+P)	Assigned	R20	4	300	15 Enterprise
A21	Normal	RAID 5 (7+P)	Assigned	R21	4	300	15 Enterprise
A22	Normal	RAID 5 (7+P)	Assigned	R22	4	300	15 Enterprise
A23	Normal	RAID 5 (7+P)	Assigned	R23	4	300	15 Enterprise
A50	Normal	RAID 10 (3*2+2S)	Unassigned		3	300	15 Enterprise
A51	Normal	RAID 10 (4*2)	Unassigned		3	300	15 Enterprise
A52	Normal	RAID 10 (3*2+2S)	Unassigned		3	300	15 Enterprise
A53	Normal	RAID 10 (4*2)	Unassigned		3	300	15 Enterprise
A54	Normal	RAID 5 (7+P)	Unassigned		0	300	15 Enterprise
A55	Normal	RAID 5 (7+P)	Unassigned		0	300	15 Enterprise
A56	Normal	RAID 5 (6+P+S)	Unassigned		1	300	15 Enterprise
A57	Normal	RAID 5 (6+P+S)	Unassigned		1	300	15 Enterprise
A58	Normal	RAID 5 (6+P+S)	Unassigned		1	300	15 Enterprise
A59	Normal	RAID 5 (6+P+S)	Unassigned		1	300	15 Enterprise
A60	Normal	RAID 5 (7+P)	Unassigned		1	300	15 Enterprise
A61	Normal	RAID 5 (7+P)	Unassigned		1	300	15 Enterprise
A62	Normal	RAID 5 (7+P)	Unassigned		1	300	15 Enterprise
A63	Normal	RAID 5 (7+P)	Unassigned		1	300	15 Enterprise

Figure 12-1 DS8000 layout example with bad RAID 10 arrays

Figure 12-2 shows the **dscli** output of the **lsextpool** command and the names that are assigned to extent pools.

```
dscli> lsextpool¶

Name        ID  stgtype rankgrp status availstor (2^30B) %allocated available reserved
numvols¶

==========================================================================================
======¶

TS_RAID10_0 P18 fb           0  below            1054         0     1054     1054       0

TS_RAID10_1 P19 fb           1  below            1054         0     1054     1054       0

TS_RAID5_0  P22 fb           0  below           14494         0    14494        0       0

TS_RAID5_1  P23 fb           1  below           15024         0    15024        0       0
```

Figure 12-2 lsextpool output

Take a closer look at the extpools p18, p19, and p22 in Figure 12-3, and extent pool p23 in Figure 12-4 on page 207.

```
dscli> lsrank -extpool p18¶

ID  Group State  datastate Array RAIDtype extpoolID stgtype¶

==========================================================¶

R20     0 Normal Normal    A20         10 P18       fb¶

R60     0 Normal Normal    A60         10 P18       fb¶

dscli> lsrank -extpool p19¶

ID  Group State  datastate Array RAIDtype extpoolID stgtype¶

==========================================================¶

R21     1 Normal Normal    A21         10 P19       fb¶

R61     1 Normal Normal    A61         10 P19       fb¶

dscli> lsrank -extpool p22¶

ID  Group State  datastate Array RAIDtype extpoolID stgtype¶

==========================================================¶

R16     0 Normal Normal    A16          5 P22       fb¶

R18     0 Normal Normal    A18          5 P22       fb¶

R22     0 Normal Normal    A22          5 P22       fb¶

R50     0 Normal Normal    A50          5 P22       fb¶

R52     0 Normal Normal    A52          5 P22       fb¶

R54     0 Normal Normal    A54          5 P22       fb¶

R56     0 Normal Normal    A56          5 P22       fb¶

R58     0 Normal Normal    A58          5 P22       fb¶

R62     0 Normal Normal    A62          5 P22       fb¶
```

Figure 12-3 Extent pool attributes for p18, p19, and p22

```
dscli> lsrank -extpool p23

ID  Group State  datastate Array RAIDtype extpoolID stgtype

==========================================================

R17     1 Normal Normal    A17          5 P23       fb

R19     1 Normal Normal    A19          5 P23       fb

R23     1 Normal Normal    A23          5 P23       fb

R51     1 Normal Normal    A51          5 P23       fb

R53     1 Normal Normal    A53          5 P23       fb

R55     1 Normal Normal    A55          5 P23       fb

R57     1 Normal Normal    A57          5 P23       fb

R59     1 Normal Normal    A59          5 P23       fb

R63     1 Normal Normal    A63          5 P23       fb
```

Figure 12-4 Extent pool tabulates for p23

Align the respective ranks to dedicated DS8000 cluster nodes by grouping odd and even numbers of resources together in extent pools.

To ensure that you do not use **rotateexts** but keep specific repository LUNs to stick to dedicated 4+4 arrays, use the **chrank** command to reserve ranks and make them unavailable during fixed block volume creation by completing the following steps:

1. Reserve the rank 61 within extent pool p19 to make it unavailable during volume creation (Example 12-1).

   Example 12-1 Reserve rank r61 with extent pool p19

   ```
   dscli> chrank -reserve r61
   CMUC00008I chrank: Rank R61 successfully modified.
   ```

2. Verify the successful execution of the command by running the **lsrank** command (Example 12-2).

   Example 12-2 lsrank command

   ```
   dscli> lsrank -l
   ID  Group State    datastate Array RAIDtype extpoolID extpoolnam  stgtype exts usedexts
   ========================================================================================
   R21     1 Normal   Normal    A21          10 P19       TS_RAID10_1 fb      1054        0
   R61     1 Reserved Normal    A61          10 P19       TS_RAID10_1 fb      1054        0
   ```

3. After verification, you can now create your first of two metadata LUNs within this extent pool (Example 12-3).

   Example 12-3 Create the first of two metadata LUNs

   ```
   dscli> mkfbvol -extpool P19 -cap 1054 -name ProtMETA_#d -volgrp V2 1900
   CMUC00025I mkfbvol: FB volume 1900 successfully created.
   ```

4. After volume creation, verify that the allocated 1054 extents for the newly created fixed block volume 1900 are all placed into rank R21 (Example 12-4).

Example 12-4 Verify allocated extents for new volume in rank r21

```
dscli> lsrank -l
ID  Group State    datastate Array RAIDtype extpoolID extpoolnam  stgtype exts usedexts
==========================================================================================
R21     1 Normal   Normal    A21         10 P19       TS_RAID10_1 fb      1054     1054
R61     1 Reserved Normal    A61         10 P19       TS_RAID10_1 fb      1054        0
```

5. Now, you can release the second rank within your extent pool to allow volume creation on it (Example 12-5).

Example 12-5 Release rank r61

```
dscli> chrank -release r61
CMUC00008I chrank: Rank R61 successfully modified.
```

6. Create the fixed block volume that is used as the metadata LUN (Example 12-6).

Example 12-6 Create metadata LUN

```
dscli> mkfbvol -extpool P19 -cap 1054 -name ProtMETA_#d -volgrp V2 1901
CMUC00025I mkfbvol: FB volume 1901 successfully created.
```

7. After volume creation, verify that the newly allocated extents are all placed in the second rank R61 (Example 12-7).

Example 12-7 Verify that new extents are placed in the second rank r61

```
dscli> lsrank -l
ID  Group State  datastate Array RAIDtype extpoolID extpoolnam  stgtype exts usedexts
==========================================================================================
R21     1 Normal Normal    A21         10 P19       TS_RAID10_1 fb      1054     1054
R61     1 Normal Normal    A61         10 P19       TS_RAID10_1 fb      1054     1054
```

Part 3

Backup management, VTL, OST, and FSI best practices

This part describes ProtecTIER backup management concepts and best practices for specific backup applications in ProtecTIER Virtual Tape Library (VTL), OpenStorage (OST), and the File System Interface (FSI) for CIFS (Windows) and NFS (UNIX) environments.

This part describes the following concepts and topics:

- Backup management introduction
- IBM Tivoli Storage Manager
- Symantec NetBackup and BackupExec
- EMC NetWorker
- HP Data Protector
- CommVault
- IBM i and Backup, Recovery, and Media Services

13

Backup management introduction

This chapter describes the recommended settings that are common to all backup servers, and include information about interoperability, software compatibility, zoning, and more. The subsequent chapters provide information about IBM Tivoli Storage Manager, Symantec NetBackup, and CommVault, and provide an overview and procedural information about VMware and other backup applications. The following chapters in this part focus on best practices for specific backup applications in ProtecTIER Virtual Tape Library (VTL), OpenStorage (OST), and the File System Interface (FSI) for CIFS and NFS environments.

This chapter describes the following topics:

- General recommendations and advice for your backups
- Terminology and vocabulary cross-reference for different backup applications
- Guidelines for your backup application catalog, and remote cloning

13.1 Introduction

Many backup servers have features and settings that are used to optimize performance when writing data to real tape cartridges. In the case of a Virtual Tape Library (VTL) environment, the ProtecTIER repository presents a VTL with virtual drives and cartridges to the backup server, and some settings that are optimized for real tape are not required. This might have a detrimental effect on the ProtecTIER deduplication factor and performance. Check the current settings of your backup server and apply those settings that can be implemented.

13.2 General recommendations

This section provides an overview of general recommendations regarding interoperability, software, backup applications, tape, and SAN zoning recommendations that are common to all backup servers. This section also describes compression, encryption, multiplexing, tape block sizes, and types of data that are targeted for backup.

13.2.1 Interoperability

Check the IBM Interoperability Matrix to ensure that the version of your backup server and operating system are supported for the ProtecTIER product. Also, ensure that your server HBA is compatible with the ProtecTIER product. You can view the matrix at the following website:

http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss

13.2.2 Software compatibility

Ensure that your backup server version platform and operating system version are listed in the supported hardware and software list for the ProtecTIER product. You can view the list at the following website:

http://public.dhe.ibm.com/common/ssi/ecm/en/ivl12348usen/IVL12348USEN.PDF

13.2.3 Software, backup application, and operating system

Ensure that the backup application software is updated to the latest version. This action can impact the overall factoring performance. Also, ensure that the operating system of the backup server is updated to the most recent patch or maintenance level. This action can impact overall HyperFactor performance. For more information, see the following website:

http://www-03.ibm.com/systems/storage/tape/library.html#compatibility

13.2.4 Tape library zoning

The backup server must have a dedicated host bus adapter (HBA) port or ports for the ProtecTIER VTL. This port or ports can be shared with a physical tape library. However, the physical tape library must not be in the same storage area network (SAN) zone as the VTL. When it is not possible to dedicate HBA ports for VTL and physical tape library, have different zones to separate the traffic.

Port sharing: Although sharing a Fibre Channel (FC) port between physical and virtual tape is possible when they are in different SAN zones, you should never share a port with disk attachments. Tape and disk devices require incompatible HBA port settings for reliable operation and optimal performance characteristics. Under stress conditions (high I/O rates for tape, disk, or both) where disk and tape subsystems share a common HBA port, stability problems have been observed.

Other SAN zoning recommendations

The following list provides recommendations that are common to all backup servers:

- Use zones that are based on the worldwide port name (WWPN).
- Use two-member zones, that is, one initiator port and one target port per zone. If this configuration is not possible, consider at least having one single initiator on each zone.
- Do not mix the front-end zoning (zoning between the ProtecTIER server and backup application server) with the back-end zoning (zoning between the ProtecTIER server and the ProtecTIER back-end storage). Including ProtecTIER back-end ports (QLogic HBAs) and front-end ports (Emulex HBAs) in the same SAN zone causes problems.
- Do not include more than one single ProtecTIER front-end port in a SAN zone. A ProtecTIER front-end port must not see other ProtecTIER front-end ports in order to avoid problems.
- For each backup server, create a separate zone for each HBA that accesses ProtecTIER virtual resources.
- Before you create worldwide name (WWN) zones on a SAN switch, you must obtain the WWPN of each port for both your ProtecTIER server and your host computer.
- If you plan to use Control Path Failover (CPF), you can zone your host to all ProtecTIER ports. You have more than one instance of the same robot that is recognized in the OS, but it is ready when the CPF redirects the traffic to the other port without requiring zoning changes.
- If you are not planning to use CPF, zone the host to some ports of the VTL, to balance the backup traffic among the HBAs. Spread drives across the HBAs.

Table 13-1 shows an example of a zone from a system where one backup application server has two HBAs. The workload distribution between the two HBAs is needed and there is no intention to use CPF. Each Tivoli Storage Manager HBA discovers two front-end ports of the VTL. The tape devices in the OS appear only once, and the load is distributed.

Table 13-1 Example of a zone where load distribution is wanted

Initiator	Target
TSM_HBA_Port0	VTL_FrontEnd_port0
TSM_HBA_Port0	VTL_FrontEnd_port2
TSM_HBA_Port1	VTL_FrontEnd_port1
TSM_HBA_Port1	VTL_FrontEnd_port3

13.2.5 Compression

Compression scrambles the data that is sent to the ProtecTIER server, which makes pattern matching difficult. This data scrambling affects data matching rates, even if the same data is sent each time. The ProtecTIER product compresses the data that it sends to the back-end storage disks after the virtual tape drives receive and deduplicate the data. Disable any compression features for the ProtecTIER that are defined on the backup server and clients. Any type of server and client compression, deduplication, or encryption negatively affects the deduplication ratio on the ProtecTIER system.

13.2.6 Encryption

Encryption makes each piece of data that is sent to the ProtecTIER server unique. Encryption affects the data matching rates and the factoring performance. Even if the same data is sent each time, it appears as different data to the deduplication engine. Disable any encryption features in your backup server and client application.

13.2.7 Multiplexing

Do not use the multiplexing feature of any backup application with the ProtecTIER product. Although the ProtecTIER product works with these features, the benefits (disk savings) of the HyperFactor algorithm and compression is reduced. Disable any multiplexing features on the backup server and clients.

13.2.8 Tape block sizes

To optimize the backup server performance, set the block size for data that is sent to the (virtual) tape drives to *256 KB* or greater.

13.2.9 Type of data that is backed up

Another factor that affects performance in a ProtecTIER environment is the *type* of data that is targeted for backup. Some data is well-suited for data deduplication and other data is not. For example, small files (less than 32 KB in size) commonly found in operating systems do not factor well, although the built-in compression might reduce their stored size. For more information about data types, see Chapter 20, “Application considerations and data types” on page 315.

Reevaluate your current backup workloads. Decide which backups are not good candidates for ProtecTIER deduplication.

13.3 General advice for backups

Generally, the preferred method of operation for using the ProtecTIER product is to imitate the procedure that is used with physical cartridges. Implement the time frame mode of operation so that, for every 24-hour cycle, there is a backup window and then a replication window. The user must ensure that there is enough bandwidth (TCP/IP and SAN) and time allotted so that there is no overlap and no replication backlog.

Here is a typical operational flow:

1. Perform regular daily backups to the ProtecTIER system during the defined backup window.
2. After the daily backups are complete, perform a full catalog/DB backup to cartridge to the ProtecTIER repository.
3. Set up the system so that replication starts and is finished before the next backup cycle starts.
4. The user must have a complete and easily recoverable set of their latest daily backups, including the backup application catalog image.
5. If a disaster occurs, the user can revert to the last completed set of backups. So the recovery point objective (RPO) is within the 24-hour window that is typical for the service level agreement (SLA).

13.4 ProtecTIER integration with backup applications

There are three ways the ProtecTIER repository interfaces with backup applications:

- VTL
- OST
- File System Interface (FSI-CIFS and FSI-NFS)

Table 13-2 provides a summary of the backup application support for each type of ProtecTIER interface as of Version 3.3.

Table 13-2 ProtecTIER interfaces and backup applications

Backup application	VTL	OST	FSI-CIFS	FSI-NFS
IBM Tivoli Storage Manager	X		X	X
Symantec Veritas NetBackup (NetBackup)	X		X	X
EMC NetWorker (Legato)	X	X	X	X
CommVault	X		X	[a]
HP Data Protector	X		X	X
Symantec BackupExec	X		X	X

a. Supported as of Version 3.3.1

13.5 Backup application vocabulary cross-reference

Each backup application has its own terminology. Table 13-3 describes and compares the terms that are used by each backup application.

Table 13-3 Backup application vocabulary cross-reference

Term definition	Tivoli Storage Manager	EMC NetWorker	NetBackup	CommVault
The object that is saved in to the backup application, for example, a file or a database table.	Backup	Save set	Image	Backup set
The physical box that connects to the backup devices, such as a tape device.	Tivoli Storage Manager Server or Library Manager	Storage node	Media server	Media agent
The location where the master database of the backups is stored.	Tivoli Storage Manager Server	NetWorker server	-	CommServe
Repository where information is that enables access and decision-making for backups.	Tivoli Storage Manager database	-	Catalog	Database
A system that has data to be backed up, but has no metadata information about the backed-up data.	Client or node	Client node	-	Client
A system that has data to be backed up, and has direct attached storage, typically tape drives.	LAN-free client	Storage node (remote)		LAN-free Media Agent
Application that runs on a client to send data to be backed up.	Tivoli Storage Manager Client	Save	-	Data agent

13.6 Backup application catalog

The backup application catalog/database (DB) is a list of the cartridges that are used for backup and includes the following information:

- The date when the backup was performed
- A list of files that are associated with the backup
- Retention period
- Other backup application-specific information

The backup application supports one catalog or DB per backup server instance. In many cases, the primary and disaster recover (DR) sites have two separate backup servers, each with its own DB or catalog. To efficiently read replicated cartridges at the DR site, the DR site backup server must have access to either the catalog or DB of the primary backup server or an exact copy of it.

There are two basic backup environment topologies:

- Single domain backup environment

  A single domain backup environment shares a catalog across the primary and DR sites. Its entries are visible for both servers always. In a single domain environment, the backup application is fully aware of the whereabouts and status of the cloned cartridges.

- Multiple domain backup environments

  A multiple domain environment requires the user to recover the DR backup server by using a copy of the catalog or DB that matches the replicated repository cartridges that are set before restoring their replicated data from the ProtecTIER system at the DR site.

Single domain backup environment: A single domain backup environment works well with Symantec NetBackup, but does not work with Tivoli Storage Manager.

13.7 Remote cloning of virtual tapes

Remote cloning is the process of using a secondary (DR) site to clone cartridges. You can use ProtecTIER replication to offload tape cloning to your secondary site. Many users replicate their data from the primary site to the secondary (DR) site, and then move it from the disk-based repository on to physical tape cartridges for long-term retention.

One of the advantages of this practice at the secondary site is that it shifts the burden of cloning to physical tape from the production environment to the DR site location. The DR site cloning operation uses the cartridge replicas at the ProtecTIER VTL shelf of the destination. The process imitates the commonly used physical process for the transportation of physical cartridges from the primary site to a DR site.

This feature is effective in single domain backup deployments because in these environments the backup application servers at both sites share the catalog and can be concurrently connected to the ProtecTIER systems. The replication visibility switch control feature is used in these environments. The cartridges to be cloned are moved from the primary repository to the secondary repository and then cloned to physical tapes.

14

IBM Tivoli Storage Manager

This chapter describes recommended settings for IBM Tivoli Storage Manager and includes information about interoperability, software compatibility, zoning, and more.

The ProtecTIER product can be deployed as a Virtual Tape Library (VTL) or File System Interface (FSI) to Tivoli Storage Manager. This chapter describes Tivoli Storage Manager with VTL and with FSI.

Attention: For general VTL considerations for Tivoli Storage Manager servers, see the following website:

`http://www-01.ibm.com/support/docview.wss?uid=swg21425849`

For best practices and configuration of Tivoli Storage Manager in your ProtecTIER FSI environment, see 14.3, “Tivoli Storage Manager: FSI” on page 231.

This chapter describes the following topics:

- Preferred options to enable optimum performance in your ProtecTIER environment
- A technical overview of a typical Tivoli Storage Manager environment that is using the ProtecTIER product
- Best practices when you configure Tivoli Storage Manager to work with the ProtecTIER product

14.1 Tivoli Storage Manager VTL

Combining the advanced capabilities and features of Tivoli Storage Manager with the powerful performance-enhancing and cost reducing capabilities of the ProtecTIER product provide IT organizations with a cost-effective way to improve the performance, reliability, and scalability of data protection.

This chapter describes the settings for Tivoli Storage Manager and includes information about interoperability, software compatibility, zoning, and more.

Attention: The Tivoli Storage Manager parser does not recognize the backed-up file as a Tivoli Storage Manager stream when you use random access mode. Do not use random access mode with Tivoli Storage Manager.

For more information about planning for the Tivoli Storage Manager parser, and estimating the benefits of the Tivoli Storage Manager parser by using the ProcessCSV tool, see Appendix A, “ProtecTIER parsers” on page 447.

14.2 Tivoli Storage Manager: Preferred options

Check the following Tivoli Storage Manager server and client options. If necessary, change the options to enable optimum performance of the ProtecTIER server.

- Use 256 KB I/O for the virtual tape drives, which provides the best factoring ratio. You should configure this setting in the operating system.
- Disable client compression. Keep the default parameter as “**COMPRESSION NO**” in the Tivoli Storage Manager Backup Client option file, or update the Tivoli Storage Manager client node definition in the Tivoli Storage Manager server with the **update node <node_name> compression=no** parameter.
- Set the server option **MOVEBATCHSIZE** to 1000 (the default value).
- Set the server option **MOVESIZETHRESHOLD** to 2048 (the default value).
- When you define the library in the ProtecTIER server, select the TS3500 as the library to be emulated by ProtecTIER (For more information about setting up the virtual library, see 3.3.1, “Creating libraries” on page 39.)
- When you use Windows based Tivoli Storage Manager servers, use IBM Tape Driver, not the Tivoli Storage Manager included drivers, for Windows. Native Windows drivers for the emulated p3000 and DLT7000 drives do not function correctly in this context.
- With the ProtecTIER product, servers can share one virtual library, or you can create a separate virtual library for each Tivoli Storage Manager server.
- Set the Tivoli Storage Manager device class to represent the Ultrium LTO3 tape without compression by using the **FORMAT=ULTRIUM3** parameter.
- Configure the estimated capacity size, in the Tivoli Storage Manager device class, to represent the virtual tape size that is defined in the VTL, by using the **ESTCAPacity** parameter.
- When you define the library in the Tivoli Storage Manager Server, if the Tivoli Storage Manager Server version is Version 6.3 or higher, use the **LIBType=VTL** parameter.

- Use **RELABELSCRatch=yes** in the Tivoli Storage Manager library definition to specify that the server relabels volumes that are deleted and returned to scratch in order to free up the space in the VTL repository. Without relabeling scratch tapes, the allocated space in the ProtecTIER repository is not released.

Tivoli Storage Manager command syntax: To verify the command syntaxes and other Tivoli Storage Manager related topics, see the IBM Tivoli Storage Manager Version 6.4 Information Center at the following website:

http://pic.dhe.ibm.com/infocenter/tsminfo/v6r4/index.jsp

14.2.1 LAN-free backups with the ProtecTIER product

LAN-free backups are simpler with the ProtecTIER product because there are increased tape resources and fewer hardware restrictions. ProtecTIER configured as a VTL has the advantage of presenting greatly increased tape resources to the backup server. So, you are able to perform LAN-free backups to the ProtecTIER server without considering the limitations that are normally applied to these backups, such as tape drive availability.

If you have many LAN-free clients, then it is possible that your LAN-free backup windows are dictated not entirely by business needs but also by hardware availability. With the ProtecTIER product and its maximum of 256 virtual tape drives per ProtecTIER node, you can virtually eliminate any previous hardware restrictions, and schedule your backups as and when they are required by your business needs. The TS7620 Appliance Express can support only up to 64 virtual drives per node.

LUN masking: Enable LUN masking with LAN-free clients. LUN masking reduces the administration of path creation at the Tivoli Storage Manager server. For more information, see 7.3.1, “LUN masking methods and best practices” on page 115.

14.2.2 Data streams

You might be able to reduce your current backup window by taking full advantage of the throughput performance capabilities of the ProtecTIER product. If tape drive availability is limited for concurrent backup operations on your IT storage management (Tivoli Storage Manager) server, you can define a greater number of virtual drives. Reschedule backups to run at the same time to maximize the number of allowable parallel tape operations on ProtecTIER servers.

Important: If you choose to implement this strategy, you might need to increase the value of the **MAXSESSIONS** option on your Tivoli Storage Manager server to specify the maximum number of simultaneous client sessions that can connect to the server. For more information, see the Information Center topic at the following website:

http://pic.dhe.ibm.com/infocenter/tsminfo/v6r3/topic/com.ibm.itsm.srv.ref.doc/r_opt_server_maxsessions.html

You might also need to update the Maximum Number of Mount Points (**MAXNUMMP**) in the Tivoli Storage Manager node registration to specify the maximum number of mount points a node may use on the server or storage agent only for backup operations.

In the Tivoli Storage Manager client option file, you can set the Resource Utilization (**RESOURceutilization**) parameter to specify the level of resources that the Tivoli Storage Manager server and client can use during processing. For more information, see the Information Center topic at the following website:

http://pic.dhe.ibm.com/infocenter/tsminfo/v6r3/topic/com.ibm.itsm.client.doc/r_opt_resourceutilization.html

Also, check the *ProtecTIER and Tivoli Storage Manager Performance Tuning* white paper, found at the following website:

http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102008

14.2.3 Reclamation

Continue to reclaim virtual storage pools that are on the ProtecTIER server. The thresholds for reclamation might need some adjustment until the system reaches a *steady state*. In a steady state, the fluctuating size of the virtual cartridges stabilizes and you can decide what the fixed reclamation limit ought to be.

When you decide how many virtual cartridges to define, consider the current storage pool **reusedelay** value. This value is equal to the number of days that your Tivoli Storage Manager database backups are retained. The same delay period applies to your storage pools that store data on ProtecTIER virtual cartridges. You might need to increase the number of pools to ensure that you always have scratch cartridges available for backup.

Note: There is a Fix Pack that must be applied to certain Tivoli Storage Manager versions so that **REUSEDELAY** and **RELABELSCRACTH** work correctly; otherwise, you receive the following error:

```
IC78915: RELABELSCRATCH PARAMATER DOES NOT WORK ON VTL WITH REUSEDELAY
PARAMETER GREATER THAN 0
```

For more information, go to the following website:

http://www-304.ibm.com/support/docview.wss?uid=swg1IC78915

14.2.4 Collocation

When you use a virtual library, consider implementing collocation (storing your data in two locations) for your primary storage pools. If you begin a restoration when another task (for example, a backup or cartridge reclamation) is using the virtual cartridge, you might not be able to access the data on it immediately.

Collocation means that all of your data for a node or node group is contained on the same set of virtual cartridges. Because you do not have any of the limitations of physical cartridges that are normally associated with this feature (such as media and slot consumption), you can enable the option.

Collocating data with similar expiration characteristics

As much as possible, collocate data with similar expiration characteristics, and then let that data expire. This collocation practice minimizes reclamation, and helps reduce the Tivoli Storage Manager workload. It also reduces the risk of replicated cartridges being out of synchronization because of the timing of the reclamation activity.

14.2.5 Physical tape

Depending on your data protection requirements, it might be necessary to copy the deduplicated data to physical tape. You can do this task by using standard Tivoli Storage Manager copy storage pools that have device classes that direct data to physical libraries and drives.

Tip: Estimate the number of drives that can be used to move data to physical tapes and consider a single stream performance.

14.2.6 Avoiding mount conflicts

To avoid a mount conflict, increase the number of drives (according to your needs) up to 512 per dual-node cluster (256 per node). Depending on your Tivoli Storage Manager version or operating system, these maximum values might change.

Cartridge size: The recommended cartridge size is 100 GB to reduce the reclamation load and to enable concurrent replication.

14.2.7 Multiple streams from the client with resourceutilization parameter

When possible, use multiple streams for the client backup. Try using four or more concurrent streams when you need maximum performance. You can set up multiple streams by modifying the `dsm.opt` (Windows) or `dsm.sys` (UNIX) file on the client and specify the **`resourceutilization`** parameter.

For more information, see the following website:

http://pic.dhe.ibm.com/infocenter/tsminfo/v6r4/topic/com.ibm.itsm.client.doc/r_opt_resourceutilization.html

The option **`RESOURCEUTILIZATION`** increases or decreases the ability of the Tivoli Storage Manager client to create multiple sessions. For Backup or Archive, the value of **`RESOURCEUTILIZATION`** does not directly specify the number of sessions that are created by the client. However, this setting specifies the level of resources that the Tivoli Storage Manager server and client can use during backup or archive processing. The higher the value, the more sessions that the client can start if it deems necessary. The range for the parameter is 1 - 10.

When the option is not set, which is the default, then only two sessions are created on the server. The default **RESOURCEUTILIZATION** level is 0 and it allows up to two sessions running on the server, one for querying the server and one for sending file data. **RESOURCEUTILIZATION=5** permits up to four sessions (two for queries and two for sending data), and **RESOURCEUTILIZATION=10** permits up to eight sessions (four for queries and four for sending data) with the server. The relationship between **RESOURCEUTILIZATION** and the maximum number of sessions that is created is part of an internalized algorithm and, as such, is subject to change.

Table 14-1 lists the relationships between **RESOURCEUTILIZATION** values and the maximum sessions that are created. Producer sessions scan the client system for eligible files. The remaining sessions are consumer sessions and are used for data transfer. The threshold value affects how quickly sessions are created.

Table 14-1 Relationship between the RESOURCEUTILIZATION value and maximum sessions created

RESOURCEUTILIZATION value	Maximum number of sessions	Unique number of producer sessions
1	1	0
2	2	1
3	3	1
4	3	1
5	4	2
6	4	2
7	5	2
8	6	2
9	7	3
10	8	4
Default (0)	2	1

14.2.8 Accommodating increased sessions

Ensure that the **MAXSESSIONS** setting on the Tivoli Storage Manager server can accommodate the increased sessions. The default value for maxsessions is 25. Set this parameter in the Tivoli Storage Manager server options file (Tivoli Storage Manager must be halted and then restarted) or run the **SETOPT** command, as shown in Example 14-1.

Example 14-1 Set MAXSESSIONS parameter

```
tsm: SERVER1>setopt MaxSessions 100

Do you wish to proceed? (Yes (Y)/No (N)) y
ANR2119I The MAXSESSIONS option has been changed in the options file.
```

Also, update the NODE definition on the Tivoli Storage Manager server to allow more than one mount point (**MAXNUMMP**).

For more information, see the following website:

http://pic.dhe.ibm.com/infocenter/tsminfo/v6r4/topic/com.ibm.itsm.srv.ref.doc/r_cmd_node_update.html

Migrating data: Do not expect an effective deduplication when you migrate your existing data from physical tape to the ProtecTIER repository if the data was originally backed up without best practices in place. Use the most current version of the ProtecTIER product so that you implement the appropriate Tivoli Storage Manager parser, which maximizes your overall deduplication factoring ratio.

For more information about how to use Tivoli Storage Manager for Virtual Environments with the ProtecTIER product, see the following website:

http://www.ibm.com/support/docview.wss?uid=swg27021081

14.2.9 Tivoli Storage Manager storage pool selection

When you select storage pools to restore or to retrieve data, the server evaluates the number of volumes that are required for the operation and selects the storage pool with the fewest volumes.

Usually, a VTL that is set up with small logical volumes often has data that is spread out over more volumes than the data in a physical tape library. As a result, the server selects the physical tape storage pool, which has fewer volumes, rather than the faster VTL storage pool.

To force the server to ignore the number of volumes when you select a storage pool to restore or to retrieve data, use the **IGNORENUMVOLSCHECK** Tivoli Storage Manager server option.

Storage pool selection: For more information about the storage pool selection, see the following website:

http://www.ibm.com/support/docview.wss?uid=swg21417248

14.2.10 Technical overview

Figure 14-1 illustrates a typical Tivoli Storage Manager environment that uses the ProtecTIER product. The Tivoli Storage Manager environment is straightforward. The Tivoli Storage Manager servers are connected to storage devices (disk, real tape, or virtual tape), which are used to store data that is backed up from the clients. Every action and backup set that is processed by Tivoli Storage Manager is recorded in the Tivoli Storage Manager database. Without a copy of the Tivoli Storage Manager database, a Tivoli Storage Manager server cannot restore any of the data from the storage devices.

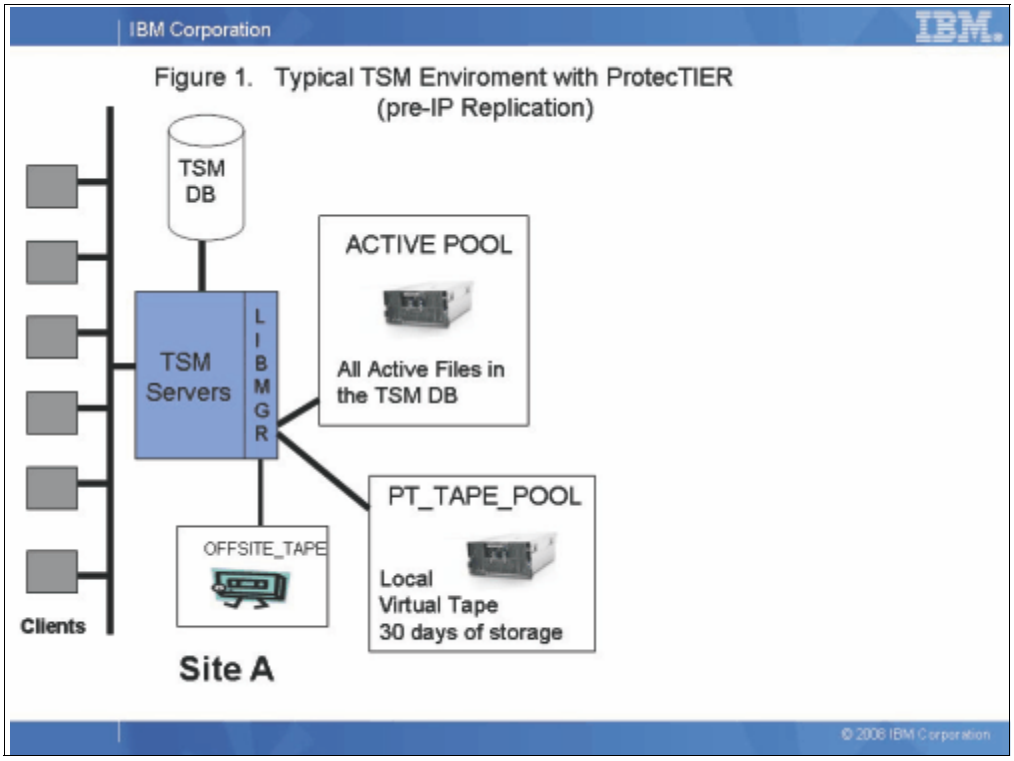

Figure 14-1 Typical Tivoli Storage Manager Environment with ProtecTIER (pre-IP replication)

The ProtecTIER product provides a virtual tape interface to the Tivoli Storage Manager servers and allows the creation of two storage pools:

- The ACTIVE Tivoli Storage Manager pool
- The ONSITE TAPE pool (called PT_TAPE_POOL)

Using the configuration that is shown in Figure 14-1, the user creates a storage pool to create real physical tapes to take offsite (called OFFSITE_TAPE). The user sizes the PT_TAPE_POOL (ProtecTIER system) to store all active client files plus about 30 days worth of inactive client files on virtual tape. The user creates an ACTIVE POOL, which is also hosted on the ProtecTIER system. The ACTIVE POOL pool contains the most recent (active) files that are backed up from all client servers. Client restoration information comes from the ACTIVE Pool.

14.2.11 Advantages of a Tivoli Storage Manager environment with ProtecTIER

The configuration that is shown in Figure 14-1 eliminates the usage of physical tape in the data center and allows faster restoration because the information is coming from the ProtecTIER disk-based virtual tape versus real tape.

Volume definition: When you predefine the volume at Tivoli Storage Manager for backup, this volume is not parsed.

14.2.12 Tivoli Storage Manager version with VTL

This section describes the best practices when you configure Tivoli Storage Manager to work with the ProtecTIER VTL product.

To enable a ProtecTIER server to work with Tivoli Storage Manager, complete the following steps:

1. Ensure that devices (robot and tapes) are recognized by the OS. Example 14-2 shows how the devices appear in an AIX server.

   Output: For the sake of brevity, some of the output in the following examples is shortened.

   Example 14-2 Robot (smc0, smc1) and tapes (rmt0 - rmt63) in the AIX

```
lbsserv38> lsdev -Cc tape
rmt0  Available 02-08-02 IBM 3580 Ultrium Tape Drive (FCP)
rmt1  Available 02-08-02 IBM 3580 Ultrium Tape Drive (FCP)
rmt2  Available 02-08-02 IBM 3580 Ultrium Tape Drive (FCP)
...
rmt61 Available 03-08-02 IBM 3580 Ultrium Tape Drive (FCP)
rmt62 Available 03-08-02 IBM 3580 Ultrium Tape Drive (FCP)
rmt63 Available 03-08-02 IBM 3580 Ultrium Tape Drive (FCP)
smc0  Available 02-08-02 IBM 3584 Library Medium Changer (FCP)
smc1  Available 03-08-02 IBM 3584 Library Medium Changer (FCP)
```

2. In the administrative console (**dsmadmc**), you can define the library by using the VTL library type and **relabelscratch** parameters (Example 14-3).

   Example 14-3 Defining a library in the Tivoli Storage Manager server

```
TSM:SERVER1> define library ptlibrary libtype=vtl relabelscratch=yes shared=yes
ANR8400I Library PTLIBRARY defined.
```

   Library tape VTL: The library type VTL was introduced in Tivoli Storage Manager V6.3 to improve communication between Tivoli Storage Manager and the ProtecTIER VTL.

   If you have a previous version of Tivoli Storage Manager, you can use **libtype=scsi** instead.

3. If you are using Tivoli Storage Manager V6.3 or higher, you can create the path to the library and its subordinate drives by running the **perform libaction** command (for more information, see 14.2.14, “Defining and deleting Tivoli Storage Manager libraries with many drives” on page 230). Example 14-4 shows the **perform libaction** command.

Example 14-4 The perform libaction command

```
tsm: SERVER1>perform libaction ptlibrary action=define device=/dev/smc0
prefix=vtldr
ANR1720I A path from SERVER1 to PTLIBRARY has been defined.
ANR8955I Drive DR00 in library PTLIBRARY with serial number  is updated with
the newly discovered serial number 4641632000.
ANR1720I A path from SERVER1 to PTLIBRARY DR00 has been defined.
```

Tip: To run the **perform libaction** command, the **SANDISCOVERY** parameter needs to be set to on. The SAN discovery function relies on the operating system device names, such as IBMtape0 or IBMchanger0. If you configured customized device names, for example, on Linux with udev, the SAN discovery function does not use those device names. To enable SAN discovery, run the following command:

```
tsm: SERVER1> setopt SANDISCOVERY on
Do you wish to proceed? (Yes (Y)/No (N)) y
ANR2119I The SANDISCOVERY option has been changed in the options file.
```

The Tivoli Storage Manager server setting **SANDISCOVERY** must be disabled for Tivoli Storage Manager CPF/DPF functionality. The ProtecTIER product exports multiple tape drives with the same WWPN, and the SAN discovery feature does not work as expected, so it must be turned off.

The **SANDISCOVERY** setting can be turned on temporarily so that Tivoli Storage Manager can perform the **libaction** command. It can then be turned off when you use CPF/DPF in Tivoli Storage Manager.

4. To display the SAN devices, run the **query san** command, as shown in Example 14-5.

Example 14-5 query san command output

```
TSM:SERVER1>query san

Devicetype  Vendor  Product       Serial Number       Device
-------     -------  ------------  ------------------  -----------
LIBRARY       IBM    03584L32      0046416329990402    /dev/smc0
DRIVE         IBM    ULT3580-TD3   4641632000          /dev/rmt0
DRIVE         IBM    ULT3580-TD3   4641632001          /dev/rmt1
```

5. If you are using a version of Tivoli Storage Manager earlier than Version 6.3, you must manually define the path of the robot and all of its subordinate drives, as shown in Example 14-6.

Drives definition: Define the names for the drive in Tivoli Storage Manager that represent the VTL.

Example 14-6 Manually defining the virtual tapes drives

```
tsm: SERVER1>def path SERVER1 PTLIBRARY srct=SERVER destt=library dev=/dev/smc0
ANR1720I A path from SERVER1 to PTLIBRARY has been defined.
```

```
tsm: SERVER1>q path
Source Name     Source Type     Destination     Destination     On-Line
                                Name            Type
-----------     -----------     -----------     -----------     -------
SERVER1         SERVER          PTLIBRARY       LIBRARY         Yes

tsm: SERVER1>def drive PTLIBRARY drive1
ANR8404I Drive DRIVE1 defined in library PTLIBRARY.

tsm: SERVER1>def path SERVER1 DRIVE1 srct=SERVER destt=drive library=PTLIBRARY
device=/dev/rmt0
ANR1720I A path from SERVER1 to PTLIBRARY DRIVE1 has been defined.
```

Important: The HBA wrapper files that are shipped with the Tivoli Storage Manager server package (except on AIX) provide communication with the virtual library. If AIX:/usr/lib/libhbaapi.a (provided by AIX with the HBAAPI installation) is not correctly configured, the following error might occur:

```
ANR1803W SAN discovery module /opt/tivoli/tsm/server/bin/dsmqsan is not
installed correctly.
ANR1791W HBAAPI wrapper library libHBAAPI.a(shr_64.o) failed to load or is
missing.
ANR1792W HBAAPI vendor library failed to load or is missing.
ANR8396E PERFORM LIBACTION: Library PTLIBRARY is not capable of discovering
the drives that it owns.
```

To resolve this error, ensure that the SAN discovery module can run, has the setuid bit turned on, and is owned by root. The SAN discovery module is called dsmqsan, and must be in the server or storage agent executable directory, as shown here:

```
chown root:system /opt/tivoli/tsm/server/bin/dsmqsan
chmod 4755 /opt/tivoli/tsm/server/bin/dsmqsan
```

6. Label the virtual tapes. Example 14-7 shows the **label libvol** command that is required for label creation.

Example 14-7 The label libvol command

```
tsm: SERVER1>label libvol PTLIBRARY checkin=scratch search=yes
labelsource=barcode
ANS8003I Process number 8 started.
```

Tip: If the **AUTOLabel = yes** parameter is defined in the Tivoli Storage Manager library definition, you can run the **checkin labelsource=barcode** command.

7. Define the device class in the Tivoli Storage Manager Server for the library (Example 14-8).

Example 14-8 Device class configuration

```
tsm: SERVER1>define devclass ptvtldevclass library=ptlibrary devtype=lto
estcapacity=100000M format=ultrium3
ANR2203I Device class PTCLASS defined.
```

8. The remaining tasks vary for each client environment. You must create storage pools by using the device class that is configured for the VTL. Then, you must update the management classes and backup/archive copy groups. After these steps are complete, you can explore the advantages of the ProtecTIER product.

14.2.13 Updating to a VTL library type

You can update an existing small computer system interface (SCSI) library to a VTL library type by running the following command:

```
update library <libname> LIBTYPE=VTL
```

Updating the library to a VTL library type allows ProtecTIER to make more accurate assumptions and skip unnecessary SCSI validations. Depending on the operating system, you might run into limitations, such as 300 - 500 maximum drives.

Setting the `libtype=VTL` also eliminates the restriction of defining only 120 tape drives when you use Tivoli Storage Manager and the ProtecTIER product. This feature is only available with Tivoli Storage Manager V6.3 or later. If you have a previous version of Tivoli Storage Manager, you might want to review APAR IC66166 ‘‘Large number of tape drives can cause volume mounts to perform slowly” at the following website:

http://www.ibm.com/support/docview.wss?uid=swg21425849

Updating the library: The **`UPDATE LIBRARY`** function does not allow mixed media, such as drives with different device types or device generations within the same library (LTO2 and LTO3). Also, this function requires that online paths be defined for servers and storage agents to all drives in the library. If paths are missing or offline, the performance levels degrade to the level of the SCSI library.

14.2.14 Defining and deleting Tivoli Storage Manager libraries with many drives

Tivoli Storage Manager V6.3 introduces the **`PERFORM LIBACTION`** command, which automatically defines the library/path/drive structure in a single command. This action is helpful when you use the ProtecTIER product because it usually contains many drives. When you run the **`PERFORM LIBACTION`** command, you complete the following steps:

1. Create a path to the library.
2. Scan for all drives that belong to the library and define all drives that are found in the library.
3. Define paths to the drives found in the library, respecting the virtual tape drive numbering.

Use Tivoli Storage Manager to define and delete Tivoli Storage Manager libraries (VTL and SCSI library types only) that contain many drives by running the following command:

```
PERFORM LIBACtion <library_name> device=xxx action=define
```

Example 14-9 shows a partial output of this command.

Example 14-9 Partial output of a perform libaction command

```
tsm: SERVER1>perform libaction ptlibrary action=define
device=/dev/smc0 prefix=vtldr
ANR1720I A path from SERVER1 to PTLIBRARY has been defined.
ANR2017I Administrator ADMIN issued command: PERFORM LIBACTION
PTLIBRARY action=define device=/dev/smc0 prefix=dr
```

```
ANR2017I Administrator ADMIN issued command: DEFINE DRIVE
PTLIBRARY DR00
ANR8404I Drive DR00 defined in library PTLIBRARY.
ANR2017I Administrator ADMIN issued command: DEFINE PATH SERVER1
DR00
SRCTYPE=SERVER DESTTYPE=DRIVE LIBRARY=PTLIBRARY
...
ANR2017I Administrator ADMIN issued command: DEFINE PATH SERVER1
DR00
SRCTYPE=SERVER DESTTYPE=DRIVE LIBRARY=PTLIBRARY DEVICE=/dev/rmt32
ANR8955I Drive DR00 in library PTLIBRARY with serial number  is
updated with
the newly discovered serial number 4641632001.
```

Run the **PERFORM LIBACTION** command to set up a single SCSI or virtual tape library (VTL).

If you are setting up or modifying your hardware environment and must create or change many drive definitions, the **PERFORM LIBACTION** command can make this task much simpler. You can define a library and then define all drives and paths to the drives. If you have an existing library that you want to delete, you can delete all existing drives and their paths.

The **PERFORM LIBACTION** command can be used only for SCSI and VTL libraries. If you are defining drives and paths for a library, the **SANDISCOVERY** option must be supported and enabled.

Note: For more information, see the Tivoli Storage Manager Information Center at the following website:

http://publib.boulder.ibm.com/infocenter/tsminfo/v6r4/index.jsp?topic=%2Fcom.ibm.itsm.srv.doc%2Ft_vtl_libaction_perform.html

14.3 Tivoli Storage Manager: FSI

The following section provides steps and best practices for configuring and setting up IBM Tivoli Storage Manager for backup and restore. This section also provides Tivoli Storage Manager parameters and settings for best performance with ProtecTIER FSI. This section describes the configuration steps and parameters to configure Tivoli Storage Manager sequential-access disk (FILE) device classes for usage with ProtecTIER FSI. The required steps and best practices for Tivoli Storage Manager with ProtecTIER FSI-CIFS or Tivoli Storage Manager on a UNIX system with a ProtecTIER FSI-NFS are identical unless otherwise noted.

14.3.1 Setting up backup and restore on Tivoli Storage Manager

Before you set up the new Tivoli Storage Manager device classes for the FSI export, you need to disable the server option directly. On a Linux system, you need to configure the value for the **DIRECTIO** parameter, as shown in Example 14-10, in the file `dsmserv.opt`. Tivoli Storage Manager does try to use direct I/O to files on NFS shares when it is using dataformat=native (the default) on the storage pools.

Example 14-10 Disable direct IO for storage pools on NFS exports

```
DIRECTIO NO
```

To set up back and restore on Tivoli Storage Manager, start the Tivoli Storage Manager server by running the following command on the UNIX host:

```
/opt/tivoli/tsm/server/bin/dsmserv
```

dsmserv starts the Tivoli Storage Manager server. After the server is running, the following steps can be completed from the server in preparation for performing a Tivoli Storage Manager backup:

1. With ProtecTIER FSI, create sequential I/O device classes when you store the Tivoli Storage Manager volumes in an FSI share. A random file device class is not supported. The definition of the device class is the only step that differs between a Windows and UNIX Tivoli Storage Manager configuration, as it contains the operating system-specific file system notations. Therefore, we show two dedicated examples for each operating system type. To create the device class, run the command that is shown in Example 14-11 from the administrative command line of the Tivoli Storage Manager server.

Example 14-11 Define the sequential device class for UNIX with FSI NFS

```
DEFine DEVclass PT1_fsi_devclass DEVType=FILE MOUNTLimit=192 MAXCAPacity=8G
DIRectory=/mnt/puck_tsm2_nfs1,/mnt/puck_tsm2_nfs2
```

Alternatively if you configure a Tivoli Storage Manager server on Windows with a ProtecTIER CIFS share, run the command that is shown in Example 14-12.

Example 14-12 Define a sequential device class for Windows with FSI CIFS

```
DEFine DEVclass PT1_fsi_devclass DEVType=FILE MOUNTLimit=192 MAXCAPacity=8G
DIRectory=\\FSI_IP_alias\share1
```

Important: When you specify a file device class in Tivoli Storage Manager on Windows, do not specify the IP address when you configure the directory parameter. The full path of the Tivoli Storage Manager volumes is stored in the Tivoli Storage Manager database, including the network path with the IP address. Whenever the IP address of the FSI share changes or when you want to switch to a replicated file system on a remote repository, this action is not possible. Therefore, specify a host alias of the FSI interface IP in your domain server or create an entry in the operating system hosts file and specify the alias instead of any IP address.

Note: The MAXCAPacity=8G recommendation is not relevant anymore when using ProtecTIER PGA V3.3.3 or later. From this code level and higher there is no recommendation to limit the backup file size and in fact the bigger the better

The parameter for the **define devclass** command is described with the default values. The recommended values for ProtecTIER FSI, along with additional explanations, are described in Table 14-2.

Table 14-2 Define devclass command parameters, default values, and recommended values

Parameter	Description	Recommended / maximum values	Default value
MOUNTLimit	Number of parallel streams. For VTL, this is the number of parallel mounted virtual cartridges. For FSI, this specifies the number of Tivoli Storage Manager volumes that can be opened simultaneously for read and write operations.	16, with a maximum of 192 parallel streams for TS7650G and a maximum of 64 streams for TS7620 SM2	The default value is 20.
MAXCAPacity	Specifies the maximum file size a Tivoli Storage Manager volume can reach. Based on test results, we recommend a maximum volume size of 8 GB for `deviceType=file`.[a]	8 GB[a]	The default value is 2 GB.
DIRectory	Location of the Tivoli Storage Manager volumes. When you define a file device class, the location is either a CIFS or NFS share. If the server needs to allocate a scratch volume, it creates a file in one of these directories. Furthermore, the distribution of incoming data streams is distributed across the available directory definitions, based on round robin. If one path is unavailable, then Tivoli Storage Manager selects the next available path. If a path is unavailable for read / restore operations, the operation might fail. Therefore, it is the administrators responsibility to implement failover mechanisms, such as bonding.	Specify at least two different file shares, as shown in Example 14-11. Furthermore, the administrator needs to implement fault tolerance mechanisms for the network shares, such as bonding, to prevent path failures.	The default is the current working directory of the server when the command is issued.

a. The MAXCAPacity=8GB recommendation is no longer relevant when using ProtecTIER PGA V3.3.3 or later. From this code level and higher there is no recommendation to limit the backup file size, and in fact the bigger the better

2. To create your ProtecTIER storage pool, run the command that is shown in from the Tivoli Storage Manager server.

Example 14-13 Defining a storage pool

```
DEFine STGpool PT1_stgpool PT1_fsi_devclass POoltype=PRimary
Description="stgpool on ProtecTIER NFS share" REClaim=90 RECLAIMPRocess=16
MAXSCRatch=200 DEDUPlicate=no DataFormat=NATive
```

Important: When you use the **`define stgpool`** command, see Table 14-3 on page 234, for the parameter value descriptions.

Table 14-3 Define the stgpool command parameters, default values, and recommended values

Parameter	Description	Recommended value	Default value
REClaim	Specifies when the server reclaims a volume. This action is based on the percentage of reclaimable space on a volume. Reclamation makes the fragmented space on the volumes usable again by moving any remaining unexpired files from one volume to another volume, thus making the original volume available for reuse. This parameter is optional.	90.	60.
RECLAIMPRocess	Specifies the number of parallel processes to use for reclaiming the volumes in this storage pool. This parameter is optional.	TS7650: 3. TS7620: 3.	1.
MAXSCRatch	Specifies the maximum number of scratch volumes that the server can request for this storage pool. To identify the number of scratch volumes, determine the ProtecTIER file system size that you configured by looking at the size of a scratch volume and reduce the number by 2 to avoid exceeding the file system size: MAXSCRatch= (PT FS size in gigabytes/ 8 GB)[a] - 2. For example: File system size=10 TB MAXCAPacity= 8 GB[a] MAXSCRatch=(10*1024/8) - 2 = 1278	The number of scratch volumes is based on the file system size that you configured on the ProtecTIER system.	
REUsedelay	Specifies the number of days that must elapse after all files are deleted from a volume before the volume can be rewritten or returned to the scratch pool.		The default value is 0, which means that a volume can be rewritten or returned to the scratch pool when all the files are deleted from the volume.
MIGPRocess	Specifies the number of parallel processes to use for migrating the files from the volumes in this storage pool.	TS7650: 3. TS7620: 3.	The default value is 1.

Parameter	Description	Recommended value	Default value
DATAFormat	Specifies the data format to use to back up files to this storage pool and restore files from this storage pool.	Native. Important: Do not use the **DATAFormat=nonblock** option, as it might adversely affect performance and the deduplication ratio.	Native.
DEDUPlicate	If you are using a version earlier than Tivoli Storage Manager V6.2, the **DEDUPlication** option might not exist.	No.	No.

a. The MAXCAPacity=8GB recommendation is not relevant anymore when using ProtecTIER PGA V3.3.3 or later. From this code level and higher there is no recommendation to limit the backup file size, and in fact the bigger the better.

3. To create your policy domain, run the following command:

   ```
   DEFine DOmain domain1 BACKRETention=30 ARCHRETention=365
   ```

4. To create a policy set, run the following command:

   ```
   DEFine POlicyset domain1 policyset1
   ```

5. To create a management class, run the following command:

   ```
   DEFine mgmtclass domain1 policyset1 mgmtclass1
   ```

6. To create a copy group, run the following commands:

   ```
   DEFine COpygroup domain1 policyset1 mgmtclass1 STANDARD Type=Backup FREQuency=0
   DESTination=PT1 VERExists=NOLimit VERDeleted=5 RETExtra=30 RETOnly=60
   DEFine COpygroup domain1 policyset1 mgmtclass1 STANDARD Type=Archive
   DESTination=PT1 RETVer=365 RETInit=CREATion
   ```

7. To set the default management class in the policy, run the following command:

   ```
   ASsign DEFMGmtclass domain1 policyset1 mgmtclass1
   ```

8. To activate the policy, run the following command:

   ```
   ACTivate POlicyset domain1 policyset1
   ```

9. To register a client, run the following command:

   ```
   REGister Node <node name> <node password> PASSExp=0 DOmain=domain1
   COMPression=no ARCHDELete=yes BACKDELete=yes Type=Client MAXNUMMP=80
   DEDUPlication=SERVEROnly
   ```

DEDUPlicate option: If you are using a version earlier than Tivoli Storage Manager Server V6.1 or Tivoli Storage Manager Client V6.2, the **DEDUPlication** option does not exist. Do not add it to the command. Server-side data deduplication was introduced with Tivoli Storage Manager V6.1. Client side data deduplication was introduced with Version 6.2.

10. To set the client option file, open the directory where the Tivoli Storage Manager client is installed and open the `dsm.opt` file. The file options should be set, as shown in Example 14-14.

Example 14-14 Client option file settings

```
NODENAME         <the node name given at the register node command>
TCPSERVERADDRESS <Tivoli Storage Manager server ip>
resourceutilization 10
compression no
deduplication no
```

Encryption: The encryption option should not be activated.

14.3.2 Performing backup and restore on Tivoli Storage Manager

You can perform backup and restore through the Tivoli Storage Manager client GUI or the Tivoli Storage Manager client CLI. We use the Tivoli Storage Manager Client CLI in our example.

To perform the backup and restore, complete the following steps:

1. Change to the location where the Tivoli Storage Manager client is installed (Example 14-15).

Example 14-15 Backup and restore using the CLI

```
Microsoft Windows [Version 6.17600]
Copyright <c> 2009 Micorsoft Corporation. All rights reserved.

C:\Users\Administrator>h:
H:\>cd "Program Files\Tivoli\TSM\baclient"
```

2. To initiate a backup, run the following command:

```
dsmc selective <path>\* -subdir=yes
```

For example:

```
dsmc selective "F:\bu\*" -subdir=yes
```

Subdirectories: If you do not want to back up the subdirectories, do not specify the **-subdir** option.

After you enter this command, you see an output similar to Example 14-16.

Example 14-16 Tivoli Storage Manager CLI login

```
Tivoli Storage Manager
Command Line Backup-Archive Client Interface
Client Version 6, Rlease 3, Level 0.0
client date/time: 04/24/2012, 00:37:40
<c> Copyright by IBM Corporation and other(s) 1990, 2011, All Rights Reserved.

Node Name: LEEDS2
Please enter your user id <Leeds2>:
```

In the user ID field, ensure that this name node is your node name, and confirm the node name by pressing Enter. The following prompt then appears:

```
Please enter password for user id "LEEDS2":
```

Enter your password, and the backup starts.

3. To initiate a restore, run the following command:

```
dsmc restore <path>\* -subdir=yes
```

For example:

```
dsmc restore "F:\bu\*" -subdir=yes
```

Important: In our example, Tivoli Storage Manager restores the data to its original location. It is possible to restore the data to a different location by running the following command:

```
dsmc restore "F:\bu\*" "g:\restore\" -subdir=yes
```

Tivoli Storage Manager then restores the data to `g:\restore\`.

14.3.3 Parameters for best performance with ProtecTIER FSI

Table 14-4 shows the Tivoli Storage Manager parameters and settings for best performance with ProtecTIER FSI.

Table 14-4 Tivoli Storage Manager parameter settings for best performance with ProtecTIER FSI

Value	Parameter	Component	
File	**DEVType**	**Devclass**	**Tivoli Storage Manager Server**
8 GB[a]	**MAXCAPacity**	**Dev class**	
16	**MOUNTLimit**		
Native	**DATAFormat**	**Stg pool**	
<File system size GB> / <8> -2	**MAXSCRatch**		
90	**RECLaim**		
Do not set this option.	**DEDUPlicate**		
0	**REUsedelay**		
No	**COMPression**	**REGister Node**	
16	**MAXNUMMP**		
No	**compression**	**Option file**	**Tivoli Storage Manager Client**
No	**deduplication**		
Do not set this option.	**encryptiontype**		
10	**resource utilization**		

a. The MAXCAPacity=8 GB recommendation is no longer relevant when using ProtecTIER PGA V3.3.3 or later. From this code level and higher there is no recommendation to limit the backup file size, and in fact the bigger the better

Tips:

- You should not use the **`define volume`** command. Let the Tivoli Storage Manager server handle the file creation itself by setting the **`MAXSCRatch`** parameter of **`define stgpool`**.
- Running the **`define volume`** command might affect ProtecTIER performance.
- When you run the **`define stgpool`** command, you should use the **`DATAFormat=native`** parameter. Do not use **`nonblock`**, as it might adversely affect the performance and the deduplication ratio.
- When you work with the ProtecTIER product, you should use the `devclass` file, but not the devclass disk.

15

Symantec NetBackup and BackupExec

This chapter describes the recommended settings and procedural information to integrate the ProtecTIER product in to Symantec NetBackup (NetBackup) environments. The recommended configurations and results can vary, so it is important that you review the configuration with a NetBackup specialist to determine the best configuration for your environment. This chapter also briefly describes Symantec BackupExec in an FSI Environment.

The ProtecTIER product can be deployed as Virtual Tape Library (VTL), OpenStorage (OST), or File System Interface (FSI) to NetBackup. This chapter describes NetBackup with VTL, OST, and FSI.

This chapter describes the following topics:

- A NetBackup overview and description of its three main components
- Guidelines for integrating and configuring NetBackup for optimal performance in ProtectTier deployments for your VTL and OST systems
- Guidelines for integrating and configuring NetBackup for optimal performance in ProtecTIER deployments for your FSI systems
- Basic steps to set up and configure Symantec BackupExec for backups and restores in an FSI environment

15.1 NetBackup overview

Symantec NetBackup is an Open Systems Enterprise backup software solution. Its architecture has the following three main building blocks:

Clients The systems with the data that require backing up.

Media servers The systems that are connected to the backup devices to provide additional storage. Media servers can also increase performance because they distribute the network load.

Master server The system that manages the backups, archives, and restores. Typically, the catalog is also stored on the master server. The catalog is the database that contains information about backup images and configuration information.

In all configurations, at least one media server is required. The media server has access to and manages the storage unit. The master and media servers can be installed on the same hardware. Several media servers can be installed. Each media server controls and manages its own data. NetBackup clients write the backup data to a media server over LAN/IP, but the client and media server can be installed on the same hardware.

In general, a media server uses its own storage unit. A storage unit can be either a disk staging device or a tape storage unit. If a tape storage unit is intended to be shared over several media servers, then an additional license, Shared Storage Option (SSO), is required.

The ProtecTIER product can eliminate or reduce the usage of SSO because it can emulate many virtual tape drives, so sharing might no longer be required.

For detailed product information about Symantec NetBackup, go to the following website:

`http://www.symantec.com/netbackup`

15.2 Recommendations for NetBackup

The ProtecTIER product can be deployed as a VTL, FSI, or OST type when it integrates with NetBackup. Follow the recommendations that are detailed in this section to ensure optimal deduplication and performance.

15.2.1 General recommendations

The following configuration options in NetBackup must be checked and, if necessary, changed to assist with the optimal performance of any ProtecTIER deployment:

- Ensure that you have adequate NetBackup licenses before you perform the implementation.
- Check the ProtecTIER compatibility with NetBackup server hardware and software, operating system, and SAN switches.
- Disable multiplexing, compression, and encryption.
- Select **Allow Multiple Data Streams in policies** to allow a backup job to run in simultaneous streams. It is recommended that the number of streams be set to 16 or less.
- Ensure that your host properties Global Attributes are set appropriately for your configuration.

- Do not mix disks and tapes on a single SCSI bus.
- For each backup server, create a separate zone for each HBA that can access ProtecTIER virtual resources. Use zones that are based on a worldwide port name (WWPN).
- The value of the `NUMBER_DATA_BUFFER` buffer must be at least 32.
- The value of the `SZ_DATA_BUFFER` buffer should be at least 262144 (524288 is recommended). On an AIX system, these buffers can be configured by creating the files on the NetBackup media server (Example 15-1):

Example 15-1 Create the SZ_DATA_BUFFER file

```
– – /usr/openv/netbackup/db/config/SIZE_DATA_BUFFERS –
–/usr/openv/netbackup/db/config/NUMBER_DATA_BUFFERS
```

Best practice: Customers in Linux/UNIX environments that use NetBackup and the PDISC-PRLI (registration request) loop can greatly reduce the amount of noise in the SAN by setting the `AVRD_SCAN_DELAY` entry in the `vm.conf` file to a high value.

This is a parameter that reflects the number of seconds between normal scan cycles. The minimum for `number_of_seconds` is 1. The maximum is 180. A value of zero converts to one second. The default value is 15 seconds. The default low value is used to minimize tape mount time but introduces noise to the SAN. It should also be noted that the benefit of the low value helps in physical tape environments where tape mounts take a long time, but irrelevant in VTL systems, where tape operations are instant

For more information, see the *Symantec NetBackup Administrator's Guide, UNIX and Linux Release 7.5*, found at:

`http://www.symantec.com/business/support/index?page=content&id=DOC5157`

15.3 NetBackup in a VTL environment

The following configuration options in NetBackup must be checked and, if necessary, changed to assist with the optimal performance of ProtecTIER VTL deployments:

- Backup and replication activities can run concurrently; however, a best practice is to use separate time windows for the backup and replication operations.
- Ensure that tape encryption for images that go to the ProtecTIER repository is disabled.
- Use a block size of 512 KB for the best deduplication ratio and performance balance.
- When you create the virtual TS7650G library, select either the data transfer controller (DTC) emulation (for creating a p3000 library) or V-TS3500 (identical to TS3500 emulation). This setting is a requirement from Symantec for NetBackup support.
- When you use Windows Master or Media servers, consider using NetBackup device drivers instead of native windows device drivers, per Symantec recommendations.
- For each backup server, create a separate zone for each HBA that can access ProtecTIER virtual resources. Use zones that are based on a worldwide port name (WWPN).
- If you use more than one HBA on the server to access virtual tape libraries and drives, you might see duplicate virtual tape drive definitions or robotics definitions on your server. To resolve this issue, you can enable zoning, use persistent binding, or ignore the duplicate devices.

15.4 NetBackup in an OST environment

The following configuration options in NetBackup must be checked and, if necessary, changed to assist with the optimal performance of ProtecTIER OST deployments:

- Ensure that you set the Disk Volume Settings high water mark and low water mark appropriately for your configuration. When the high water mark threshold is reached, the disk storage unit is considered full. By default, this setting is 98%. As the high water mark threshold is approached, NetBackup reduces the number of jobs that are sent to the storage unit. No new jobs are assigned to storage units that are considered full.
- Ensure that the Storage Units Maximum concurrent jobs setting is appropriate for your configuration. You must take outbound and inbound duplication jobs in to account. You can avoid unexpected I/O contention in the disk pool by dividing the maximum concurrent jobs count for all the storage units by using the disk pool by the number of volumes in the disk pool.
- Duplication takes longer to complete than the initial backup job and uses a larger amount of bandwidth. Duplication also taxes the NetBackup resource broker. This situation can slow down the rate that all jobs (backups, restores, and duplicates) run because resources are being requested simultaneously. Take extra care to plan the timing of all jobs to prevent delays.

For detailed configuration information, see the *Symantec NetBackup Shared Storage Guide*, found at:

http://www.symantec.com/business/support/index?page=content&id=DOC3659

You can also see the TS7650 Appliance with ProtecTIER Customer Information Center at the following website:

http://pic.dhe.ibm.com/infocenter/ts7650/cust/index.jsp

15.5 NetBackup in an FSI environment

This section provides steps and best practices for configuring and setting up Symantec NetBackup (NetBackup) for backup and restore. This section also provides NetBackup parameters and settings for best performance with ProtecTIER FSI-CIFS and FSI-NFS.

15.5.1 NetBackup in an FSI-CIFS environment

This section provides best practices for configuring and setting up NetBackup in an FSI-CIFS environment for backup and restore. This section also provides NetBackup parameters and settings for best performance.

Setting up for backup and restore

The following sections describe the steps for configuring NetBackup for backup and restore for NetBackup 7.0.1 and higher versions in an FSI-CIFS environment. These sections also describe the steps for configuring backup and restore for a version of NetBackup before Version 7.0.1.

Configuring AdvancedDisk sharing with CIFS for NetBackup V7.0.1 and higher versions

The following section describes how to configure AdvancedDisk sharing on Windows using CIFS. Support for AdvancedDisk sharing on Windows using CIFS is introduced with Symantec NetBackup V7.0.1. For details about the specific additional configuration steps before the disk pool can be configured, see the following website:

http://www.symantec.com/business/support/index?page=content&id=TECH158427

Mount point: The mount point is set as \\FSI_IP\CIFS_name on ProtecTIER FSI by running the following command:

```
nbdevconfig -createdv -stype AdvancedDisk -storage_server windows-ubmu9k3 -dv
\\9.11.109.130\ForNBU
```

To use AdvancedDisk, mount a network drive and assign it a letter. A drive letter must be assigned in order for AdvancedDisk to work.

Creating a disk pool using a shared volume

This section describes the steps to create a disk pool using a shared volume. When you create a disk pool using a shared volume, ensure that you select only the shared volumes.

Complete the following steps:

1. In the NetBackup Administration Console, click **Media and Device Management**.
2. From the list of wizards in the Details pane, click **Configure Disk Pool** and follow the wizard's instructions.
3. Select the type of Disk Pool you are creating. In this example, the type is AdvancedDisk. Click **Next**.
4. Select the storage server for which you want to create the disk pool and click **Next**. The Disk Pool Configuration wizard shows the list mount points that can be used.
5. Select the shared volumes that you want to use. Click **Next**
6. In the Disk Pool Properties window, set **Limit I/O streams**. By setting the Maximum I/O streams option on the disk pool, you can limit the total number of jobs that access the disk pool concurrently, regardless of the job type.

Important: Ensure that you select only the shared volumes.

Storage unit configurations

Storage units that are used with disk type AdvancedDisk are based on the disk pool rather than individual disks. To determine whether storage units exist for the disk pool, open the Administration Console and click **NetBackup Management** → **Storage** → **Storage Units**.

Creating policy and choosing the newly created storage unit

To create a policy, complete the following steps:

1. In the NetBackup Administration Console, expand **NetBackup Management** → **Policies**.
2. Right-click **Policies** and click **New Policy**.

3. From the Change Policy window, select the newly created storage unit from the **Policy storage** drop-down menu (Figure 15-1). **Schedules**, **Clients**, and **Backup Selections** can be selected from the corresponding tabs.

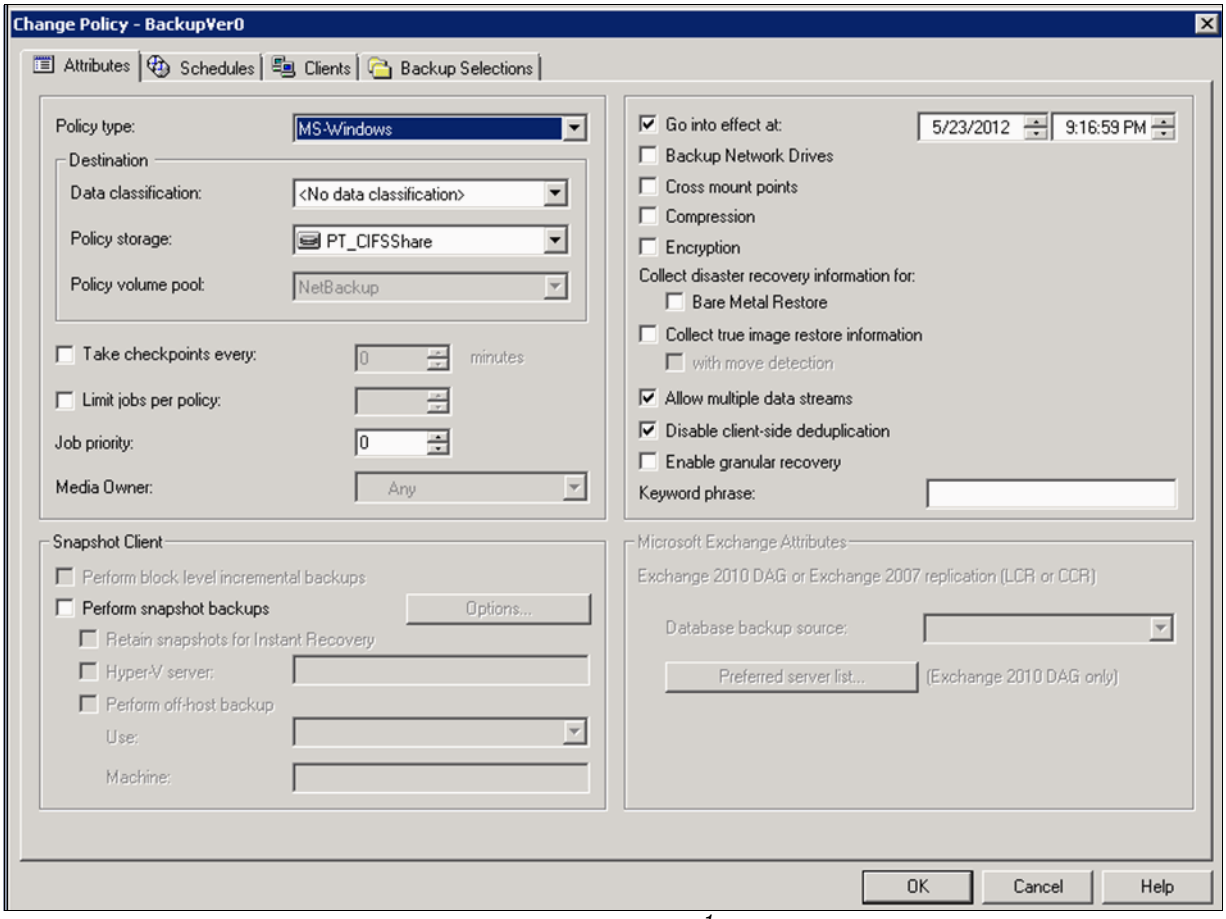

Figure 15-1 Change policy - disable encryption and compression[1]

4. From the Change policy window, specify the date and time that your policy goes into effect by selecting **Go into effect at**, and clear **Compression** and **Encryption**.

> **Note:** Compression, Encryption, and client-side deduplication should be disabled (Figure 15-1 on page 244).

Performing automatic backups or manual backups

This section provides the steps that are needed to perform automatic or manual backups by using NetBackup in your ProtecTIER FSI-CIFS environment. Complete the following steps:

1. In the NetBackup Administration Console, expand **NetBackup Management** → **Policies**.
2. Right-click the policy name in the left pane.
3. To perform a manual backup, click **Manual Backup**.
4. To perform a manual backup and to activate the policy, select the **Go into effect at** check box in the **Policy Attributes** tab. The policy must be active for NetBackup to use the policy.

[1] *Symantec*, Reprinted by Permission.

Important: The *Go into effect at* attribute specifies when the policy can begin to schedule backups. For example, if today is Monday and you enter Wednesday at 12:00 AM, the policy does not run until that time or later. Use the Go into effect at attribute to configure a series of policies in advance of when the policies need to become active.

Performing restores

This section provides the steps that are needed to perform a restore by using NetBackup in your ProtecTIER FSI environment. Complete the following steps:

1. Start the Backup, Archive, and Restore client interface. Click **Select for Restore**, and select the **System State** check box.
2. From the Actions menu, click **Start Restore of Marked Files**.
3. From the Restore Marked Files dialog box, click **Restore everything to its original location** and **Overwrite the existing file**.

   Warning: Do not redirect the System State restore to a different host. System State is computer-specific. Trying to restore it to a different computer can result in an unusable system.

4. Click **Start Restore**.

Parameters for best performance with ProtecTIER FSI-CIFS

This section provides guidelines and parameters for best performance with ProtecTIER FSI-CIFS. We describe network file considerations, and the usage of maximum I/O streams with disk pools for optimum performance. You can adjust the Network Buffer Size, Data Buffer Size, and Number of Data Buffers for performance enhancements.

Multiplexing: For optimum performance in your Protectier FSI-CIFS environment, do not enable multiplexing.

Setting compression, encryption, and client-side deduplication

Compression, encryption, and client-side deduplication should be disabled, as described in “Creating policy and choosing the newly created storage unit” on page 243“.

Network file system considerations

The AdvancedDisk storage implementation presents all mounted file systems as disk volumes to NetBackup, including network file systems (such as NFS and CIFS). For more information about using network file systems, go to the following website:

`http://www.symantec.com/business/support/index?page=content&id=TECH158427`

Using Maximum I/O streams with disk pools

The Maximum Concurrent Jobs setting on the Storage Units menu limits the number of backup or write jobs using each storage unit, but does not limit the number of restore or read activities that might be going on concurrently with the write activity. This situation can cause unexpected I/O contention on the disk pool. For more information, see the following website:

`http://www.symantec.com/business/support/index?page=content&id=TECH158427`

Changing the disk pool I/O streams option

To update the default number of I/O streams per volume, complete the following steps:

1. In the NetBackup Administration Console, expand **Media and Device Management** → **Devices** → **Disk Pools**.
2. Select the disk pool that you want to change in the Details pane.
3. Click **Edit** → **Change**, and the Change Disk Pool window opens, where you can change the value in the Maximum I/O Streams pane.

Setting the maximum jobs per client

To set or change the maximum jobs per client, from the NetBackup Administration Console on the master server, complete the following steps:

1. Expand **NetBackup Management** → **Host Properties** → **Master Server**.
2. Open the host properties of the master server.
3. Click **Global Attributes**.

   The Global Attributes properties apply to the currently selected master servers. The Global Attributes properties affect all operations for all policies and clients. The default values are adequate for most installations, but can be changed.
4. Set **Maximum jobs per client**. The Maximum jobs per client property applies to all clients in all policies.

   This property specifies the maximum number of backup and archive jobs that NetBackup clients can perform concurrently.

15.5.2 NetBackup in an FSI-NFS environment

This section provides best practices for configuring and setting up NetBackup in an FSI-NFS environment for backup and restore. This section also provides NetBackup parameters and settings for best performance

Setting up backup and restore in an NFS environment

This section describes the steps for configuring NetBackup for backup and restore for NetBackup V7.0.1 and later versions.

Creating a disk pool using a shared volume

In this section, we describe the steps to create a disk pool using a shared volume. When you create a disk pool using a shared volume, ensure that you select only the shared volumes. Complete the following steps:

1. In the NetBackup Administration Console, click **Media and Device Management**.
2. From the list of wizards in the Details pane, click **Configure Disk Storage Servers**.
3. Select the type of disk storage you are configuring. The selected disk storage type should be AdvancedDisk. Click **Next**.
4. Select the storage server and media server for which you want to create the disk pool and click **Next**.
5. The wizard guides you through the steps that are required to create a disk pool and a storage unit that uses the newly created disk pool. Click **Next**.
6. Select the type of Disk Pool you are creating. The selected disk pool type should be AdvancedDisk. Click **Next**.
7. Select the storage server for which you want to create the disk pool and click **Next**.

8. The Disk Pool Configuration wizard shows the list of mount points that can be used.
9. Select the FSI-NFS mounted point to add into the disk pool. Click **Next**.

> **Note:** Ensure that you select only the ProtecTIER mounted folder. Preferably, one mount should be defined per disk pool.

10. In the Disk Pool Properties window, set the **Limit I/O streams** parameter to 16. This parameter limits the total number of I/O streams that access the disk pool concurrently, regardless of the job type. The ProtecTIER recommendation is for a total of 16 I/O streams. Click **Next**.
11. The disk pool configuration is completed. Click **Next**.
12. Choose the **Create a storage unit that uses the disk pool** option to configure a storage unit.

Storage unit configurations

Storage units that are used with disk type AdvancedDisk are based on the disk pool rather than individual disks. The Disk Pool Configuration wizard lets you create a storage unit as part of the creation of the disk pool. To determine whether storage units exist for the disk pool, open the Administration Console and click **NetBackup Management** → **Storage** → **Storage Units**.

To create a storage unit, complete the following steps:

1. In the NetBackup Administration Console, select **NetBackup Management** → **Storage** → **Storage Units**.
2. Click **Actions** → **New** → **Storage Unit**.
3. Complete the fields in the New Storage Unit dialog box. We recommend the following settings:
   - The maximum fragment size parameter should be set to 8192 MB (8 GB).
   - The maximum concurrent jobs should be set to 16.

> **Note:** The maximum fragment size recommendation is no longer relevant when using ProtecTIER PGA V3.3.3 or later. From this code level and higher there is no recommendation to limit the backup file size, and in fact the bigger the better.

Creating a policy and choosing the newly created storage unit

To create a policy, complete the following steps:

1. In the NetBackup Administration Console, expand **NetBackup Management Policies**.
2. Right-click **Policies** and click **New Policy**.
3. From the Change Policy window, select the newly created storage unit from the **Policy storage** drop-down menu. **Schedules**, **Clients**, and **Backup Selections** can be selected from the corresponding tabs.
4. Under the Attributes tab, ensure that the following check boxes are *not* selected:
   - Compress
   - Encrypt
5. Under the Attributes tab, ensure that the following check boxes *are* selected:
   - Disable client-side deduplication
   - Allow multiple data streams

6. Under the Schedules tab, double-click every relevant schedule and ensure that the value of **Media Multiplexing** in Attributes tab is set to 1.

Creating configuration files on the media server

These files adjust the responding rate between a media server and ProtecTIER FSI-NFS, and are needed for the ProtecTIER FSI-NFS to work correctly.

To create the configuration files, complete the following steps:

1. Run **`touch /usr/openv/netbackup/db/config/DPS_PROXYNOEXPIRE`**.
2. Run **`echo "1800" > /usr/openv/netbackup/db/config/DPS_PROXYDEFAULTSENDTMO`**.
3. Run **`echo "1800" > /usr/openv/netbackup/db/config/DPS_PROXYDEFAULTRECVTMO`**.
4. Restart nbrmms (NetBackup Remote Manager and Monitor Service) on the media server by running the following commands:

   ```
   pkill nbrmms
   /usr/openv/netbackup/bin/nbrmms
   ```

5. Or, you can stop and restart all services on the MSDP media server by running the following commands:

   ```
   /usr/openv/netbackup/bin/goodies/netbackup stop
   /usr/openv/netbackup/bin/goodies/netbackup start
   ```

Note: The 1800 value should be the only value inside these two files. In general, the allowed values are 10 - 3600. More information about this value can be found at the following link:

http://www.symantec.com/business/support/index?page=content&id=TECH156490

Setting the maximum jobs per client

To set or change the maximum jobs per client, from the NetBackup Administration Console on the master server, complete the following steps:

1. Expand **NetBackup Management** → **Host Properties** → **Master Server**.
2. Open the host properties of the master server.
3. Click **Global Attributes**.

Note: The Global Attributes properties apply to the currently selected master servers and affect all operations for all policies and clients. The default values are adequate for most installations, but can be changed.

4. Set **Maximum jobs per client**. The Maximum jobs per client property applies to all clients in all policies. This property specifies the maximum number of backup and archive jobs that NetBackup clients can perform concurrently.

Tip: For best performance, the recommendation is to set the Maximum jobs per client to 16.

Performing backup

This section provides the steps that are needed to perform automatic or manual backups by using NetBackup in your ProtecTIER FSI-NFS environment. Complete the following steps:

1. In the NetBackup Administration Console, expand **NetBackup Management** → **Policies**.

2. Right-click the policy name in the left pane.
3. To perform a manual backup, click **Manual Backup**.
4. To perform a manual backup and to activate the policy, select the **Go into effect at** check box in the **Policy Attributes** tab. The policy must be active for NetBackup to use the policy.

Important: The *Go into effect at* attribute specifies when the policy can schedule backups. For example, if today is Monday and you enter Wednesday at 12:00 AM, the policy does not run until that time or later. Use the Go into effect at attribute to configure a series of policies in advance of when the policies need to become active.

Performing restore

This section provides the steps that are needed to perform a restore by using NetBackup in your ProtecTIER FSI-NFS environment. Complete the following steps:

1. Start the Backup, Archive, and Restore client interface. Click **Select for Restore**, and select the **System State** check box.
2. From the Actions menu, click **Start Restore of Marked Files**.
3. From the Restore Marked Files dialog box, click **Restore everything to its original location** and **Overwrite the existing file**.

Warning: Do not redirect the System State restore to a different host. System State is computer-specific. Restoring it to a different computer can result in an unusable system.

4. Click **Start Restore**.

Parameters for best performance with ProtecTIER FSI-NFS

Table 15-1 captures all the required and recommended settings for ProtecTIER best practices with NetBackup.

Table 15-1 Recommended settings for ProtecTIER with NetBackup

Component	Parameter	Value
Disk pool definition	Disk pool type	AdvancedDisk
	Limit I/O streams	16[a]
Storage Unit definition	Maximum fragment size	8192 MB (8 GB)[b]
	Maximum concurrent jobs	16[a]
Policy definition	Compress	Disabled (Not Checked)

Component	Parameter	Value
	Encrypt	Disabled (Not Checked)
	Allow multiple data streams	Enabled (Checked)
	Disable client-side deduplication	Enabled (Checked)
	Media Multiplexing	1
Global Attributes	Maximum jobs per client	16[a]

a. To learn about the maximum number of streams that are supported, see 15.2.1, “General recommendations” on page 240.
b. The maximum fragment size recommendation is no longer relevant when using ProtecTIER PGA V3.3.3 or later. From this code level and higher there is no recommendation to limit the backup file size, and in fact the bigger the better

Network file system considerations

The AdvancedDisk storage implementation presents all mounted file systems as disk volumes to NetBackup, including network file systems (such as NFS). For more information about using network file systems, go to the following website:

http://www.symantec.com/business/support/index?page=content&id=TECH158427

15.6 Symantec BackupExec in an FSI environment

This section provides steps and best practices for configuring and setting up Symantec BackupExec (BackupExec) for backup and restore. We also provide BackupExec parameters and settings for best performance with ProtecTIER FSI.

15.6.1 Setting up backup and restore

This section describes the steps that are necessary to create a backup to disk folder. It also describes how to create a job to perform a backup and a job to perform a restore.

Creating a Backup-to-Disk folder

To create a backup to disk folder, complete the following steps:

1. Open the BackupExec GUI and go to **Devices**.
2. Start the Configure Devices Assistant.
3. Create a Backup-to-Disk folder.

4. Configure the parameters of the folder. Figure 15-2 shows the best practice parameters (full path name, 16 GB file size, and maximum concurrent jobs.

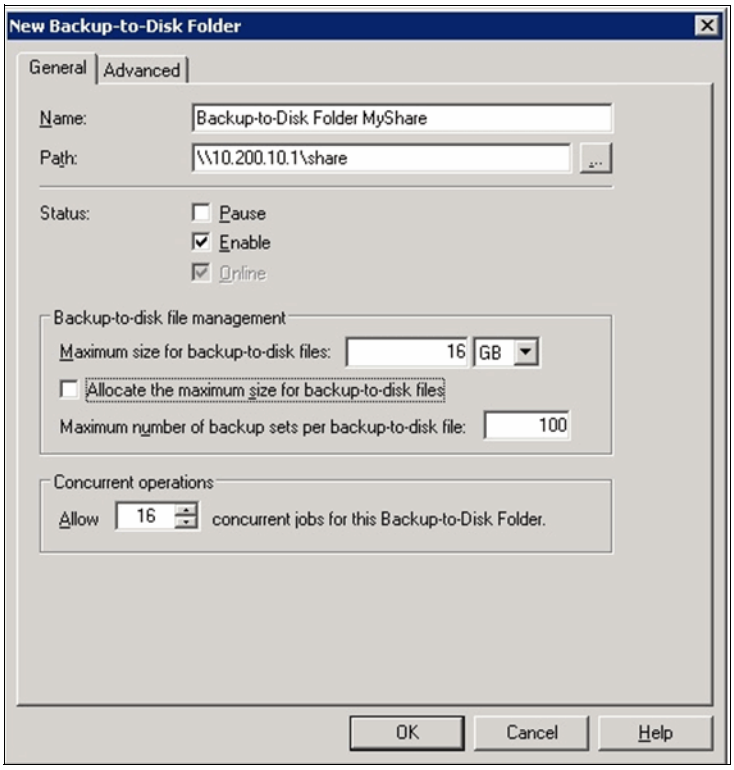

Figure 15-2 Create a backup-to-disk folder

Note: The Maximum size for backup-to-disk recommendation is no longer relevant when using ProtecTIER PGA V3.3.3 or later. From this code level and higher there is no recommendation to limit the backup file size, and in fact the bigger the better.

Creating the backup job to perform backups

To create a backup job to perform backups, complete the following steps:

1. Open the Job Monitor window.
2. Start the backup wizard by clicking the **Backup** icon.
3. Use the Backup wizard recommended parameters. The Backup wizard prompts you with **What do you want to back up?** Click **This computer** and click **Next**.
4. Click **Full backup job** and click **Next**.
5. Choose the time that you want to run this backup at (for example, for a one-time backup, click **Run the full backup now** and click **Next**).

6. Select **Backup-to-disk folder**, and in the drop-down list, click the folder that you configured and then click **Next** (Figure 15-3).

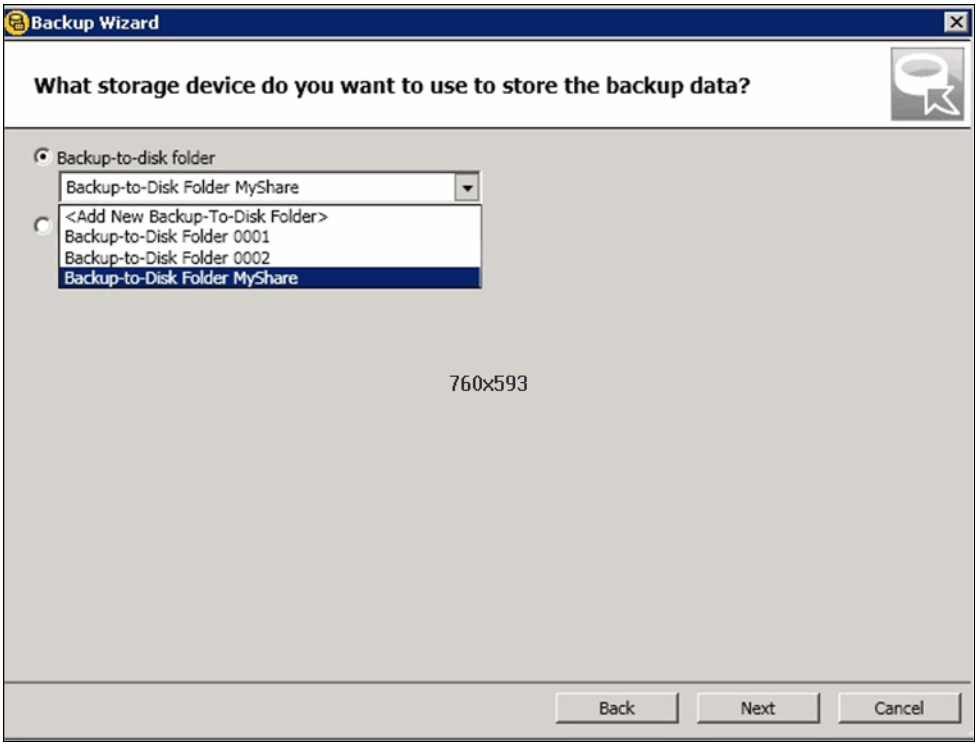

Figure 15-3 Choose backup-to-disk folder from the Backup Wizard window

7. Choose to keep the backups for the retention period you need and click **Next**.
8. Give the backup job a meaningful name and click **Submit**.

Jobs: You can see both the job running (in the Job Monitor view) and the files being written (in the share itself).

Creating the restore job to perform a restore

To create a restore job to perform a restore, complete the following steps:

1. Open the Job Monitor window.
2. Start the restore wizard by clicking the **Restore** icon.
3. Follow the Restore wizard instructions to create the restore job to perform a restore.

16

EMC NetWorker

EMC NetWorker (NetWorker) (formerly known as Legato NetWorker) is one of the enterprise backup applications that provide central backup management for various applications on different operating systems, with different backup methods to different types of storage media.

The ProtecTIER product can be deployed as a Virtual Tape Library (VTL) or a FSI-CIFS share and FSI-NFS export to NetWorker to enhance its data protection ability.

This chapter describes the recommended settings and procedural information that is needed to integrate ProtecTIER VTL and FSI in a NetWorker environment to achieve optimal backup throughput and the factoring ratio of the ProtecTIER system. The recommended configurations and results might vary in different environments. Review the configuration with your NetWorker specialist for the best configuration that fits into your environment.

This chapter describes the following topics:

- Overview
- EMC NetWorker in a VTL environment VTL
- Best practices for setting up and configuring NetWorker for backup and restore in an FSI environment.

16.1 Overview

EMC NetWorker is a centralized and automated backup and recovery product for heterogeneous enterprise data. The NetWorker Server hosts the configuration information and NetWorker databases that track the backups and volumes. It runs on all major operating systems, such as AIX, Linux, Windows, SUN Solaris, and HP-UX. Apart from server software, the NetWorker server always has the NetWorker Storage Node and NetWorker Client installed.

The NetWorker Storage Node is the host that has direct access to tape or disk media. It uses this access to read and write data to storage devices. It sends tracking information only to the NetWorker Server. The NetWorker Client is installed on the customer's servers to generate save sets, and sends them to or retrieves them from the storage node. There are different clients available for the integration of special applications, such as NetWorker for IBM DB2.

A NetWorker datazone consists of one NetWorker Server, several NetWorker Storage Nodes, and several NetWorker Clients. Figure 16-1 shows an illustration of the integration of NetWorker components into a datazone.

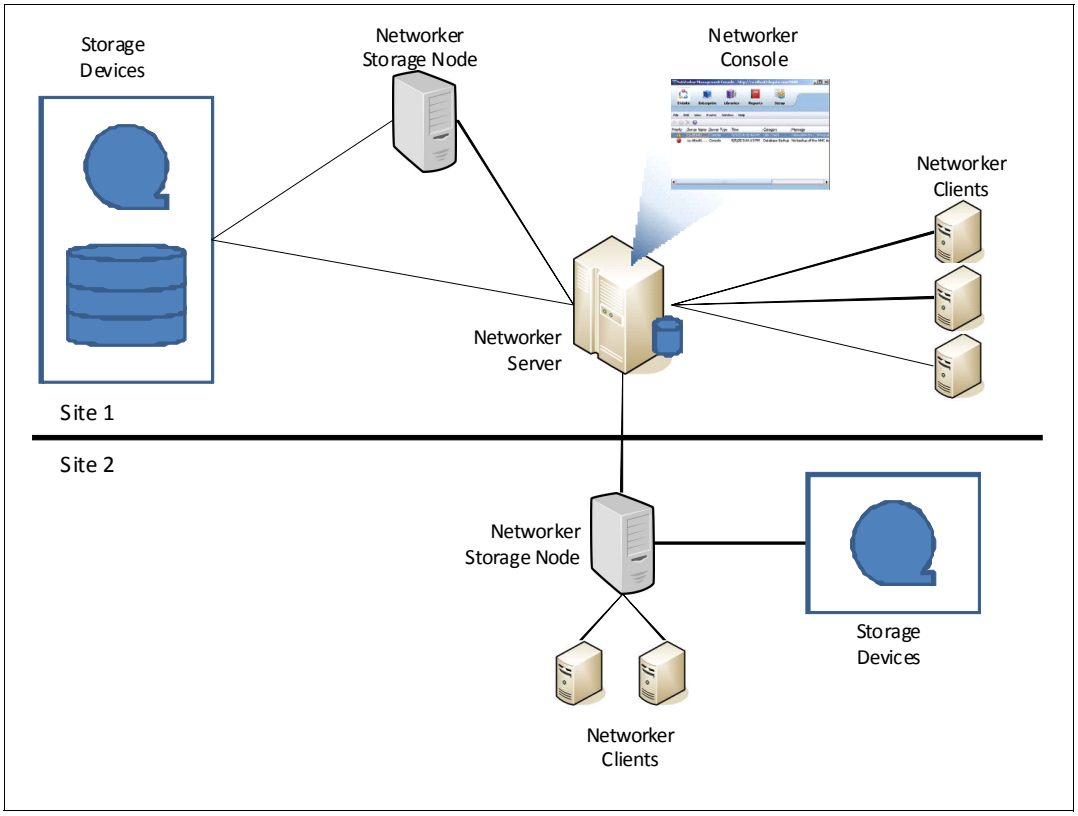

Figure 16-1 NetWorker components in a NetWorker datazone

NetWorker writes data to volumes and identifies the volumes with specified volume IDs. A volume is a physical storage media, such as magnetic tape, an optical disk, or a file system. The volume ID can be changed only when a volume is relabeled. NetWorker organizes the volumes into media pools. A ProtecTIER virtual cartridge or FSI share is seen as a volume in NetWorker and is sorted into specific media pools.

16.2 EMC NetWorker in a VTL environment

The ProtecTIER product can be deployed as a VTL when you integrate it with NetWorker. In traditional tape backups, tape drive sharing can be costly with NetWorker because of the need for an additional license for dynamic drive sharing. With VTL, many virtual tape drives can be created and dedicated to different storage nodes without wasting tape resources. The ProtecTIER product can also be deployed as network-attached storage (NAS) with a Common Internet File System (CIFS) share for the NetWorker Server on Windows or with a Network File System (NFS) export for NetWorker Servers on UNIX systems.

16.2.1 General recommendations

Follow these general recommendations to achieve optimum performance and factoring ratio of the ProtecTIER product with NetWorker:

- Ensure that you have adequate NetWorker licenses before you perform the implementation.
- Check the compatibility of ProtecTIER with the NetWorker Server hardware and software, operating system, and SAN switches.
- Review the configuration with your NetWorker specialist to avoid an adverse performance impact to your environment.
- Disable client compression and encryption.
- Whenever possible, use the same Network Time Protocol (NTP) or time server for the ProtecTIER and NetWorker servers to ease maintenance and troubleshooting tasks.

16.2.2 Recommendation if a ProtecTIER server is used as a VTL

Follow these recommended configurations to achieve optimum performance and the optimum factoring ratio of the ProtecTIER server in your VTL environment:

- Ensure that you have the VTL option enabled or an appropriate number of autochanger and library slot licenses for NetWorker.

  Warning: There are two licenses that are required for EMC NetWorker to work: the *Autochanger Module* (which must be compatible with the number of slots that are created in the ProtecTIER server) and *Virtual Tape Library Capacity* (for the number of terabytes that exist in the ProtecTIER server). Without these two licenses, NetWorker does not work correctly. You might receive the following error:

  ```
  The VTL is not properly licensed. Make sure the VTL is properly enabled and
  authorized (VTL frame enabler needed). See your EMC representative.
  ```

- Consider creating dedicated tape drives for each storage node.
- Use ProtecTIER LUN masking if multiple storage nodes share virtual tape drives. For more information, see 7.3, “LUN masking for VTL systems” on page 114.
- Add virtual tape drives gradually.
- Follow the NetWorker device driver guidelines in Chapter 7, “Host attachment considerations for VTL” on page 103.
- Use persistent binding for virtual tape drives. For more information, see 7.2.7, “Persistent device naming” on page 112.

- Disable Common Device Interface (CDI) on all virtual tape drives.
- Use 512 KB I/O for the virtual tape drives to ensure a good factoring ratio.
- Disable multiplex by setting parallelism to 1 on all virtual drives. You can do so by editing the properties of each virtual tape drive to set the target sessions and maximum sessions value to 1.
- Enable the Auto Media Management for virtual tape libraries to allow space reclamation. See the Media Management section in the *EMC NetWorker Administrator Guide* for details about Auto Media Management. You can find NetWorker documentation at the NetWorker HUB at the following address:

  `http://nsrd.info/docs.html`
- Increase the NetWorker media multiplexor daemon (**`nsrmmd`**) polling interval and **`nsrmmd`** *restart* interval if you see many **`nsrmmd`** restart activities during the backup session. Increase the value gradually, such as **`nsrmmd`** polling interval=6, restart interval=4, followed by polling interval = 9, restart interval = 6, and so on.

16.3 EMC NetWorker in an FSI environment

This section provides steps and best practices for configuring and setting up EMC NetWorker (NetWorker) for backup and restore. It also provides NetWorker parameters and settings and recommendations for best performance with ProtecTIER FSI. The recommendations apply to FSI-CIFS and FSI-NFS unless stated otherwise.

16.3.1 Creating a Windows user for EMC NetWorker

To configure EMC NetWorker (NetWorker) for Windows with ProtecTIER FSI-CIFS, you must create a Windows User for NetWorker.

Complete the steps in the following sections to configure NetWorker and run a simple backup job.

Changing the NetWorker Services' logon properties to the Windows User

Complete the following steps:

1. In the Services window, right-click **NetWorker Remote Exec Service** and click **Properties**.
2. The NetWorker Backup and Recover Server Properties window opens. Click the **Log on** tab, click **This account**, enter the Windows User account that you created, and click **OK**.
3. You receive a message that the changes do not take effect until the service stops and restarts. Click **OK**.
4. Repeat these steps with the NetWorker Backup and Recover Server service.
5. Right-click **NetWorker Remote Exec Service** and click **Restart** to restart the service. This action also restarts the NetWorker Backup and Recover Server service.

Creating a user in NetWorker

Create a user for NetWorker to access your CIFS share. Complete the following steps:

1. Open the NetWorker Management Console.
2. From the Console window, click **Setup**.
3. In the left pane, right-click **Users** and click **New**. The Create User dialog box opens.
4. Enter the user name, the appropriate role information (for example, Console Application Administrator, Console User, and Console Security Administrator), and Password to create a user, and then click **OK**.

16.3.2 Setting up for backup and restore

The following section details the steps for setting up your ProtecTIER FSI-CIFS and FSI-NFS environment to perform backups and restores by using EMC NetWorker.

16.3.3 General configuration recommendations

Here are some general configuration recommendations to configure a NetWorker server for usage with ProtecTIER FSI shares:

- Consider a dedicated network for backup. To review the FSI Guidelines for network preparation, see Chapter 4, “ProtecTIER File System Interface: General introduction” on page 53.
- Ensure that you have the NetWorker DiskBackup option enabled.
- The NetWorker storage node supports only advanced file type devices (AFTD) for FSI shares.
- Use the Universal Naming Convention (UNC) path for device configuration for CIFS shares and references to the FSI mount point on Linux.
- Whenever possible, create different ProtecTIER file systems for different storage nodes.
- NetWorker does not span ongoing saves across multiple AFTD devices. NetWorker suspends all saves being written to AFTDs when the device is full until more space is made available on the device. Plan the AFTDs correctly to have sufficient space that is based on the backup size and retention period.
- Do not share a single FSI share across multiple storage nodes, and create one AFTD in each FSI share on the storage node.
- You can configure NetWorker to split a single save set into multiple sessions in one AFTD. In the AFTD device, each save set is stored as a separate file with a unique saveset ID regardless of its parallelism. Thus, the impact to the deduplication should be minimal.
- When you install the NetWorker server on Windows, create a dedicated user (Workgroup) or add the user (Active Directory) to the Windows Admin and Backup Operator Group of the NetWorker Windows Server. Change the NetWorker Remote Exec server and NetWorker Backup/Restore Server services to be started by the newly created user with administrator rights. The same user must be added as a user of the CIFS shares in the ProtecTIER server with write enabled permission.
- When you install the NetWorker server or storage node on Linux, the read/write permission to the NFS export is granted based on host IP addresses and not on user level granularity.

For more information about AFTD, see Device Operations in *EMC NetWorker Administrator Guide.*[1]

Creating the device

To create a device in the NetWorker Administration window, use the Device Configuration wizard and complete the following steps:

1. In the NetWorker Administration window, right-click **Devices** and click **New Device wizard**.
2. In the Device Configuration wizard, click **Advanced File Type Device (AFTD)** and click **Next**.
3. The Select Storage node window opens. Click **Enter device path**, and in the box that opens, enter the address to the FSI share. Example 16-1 shows the configuration for a FSI-CIFS share by using the UNC path.

   Example 16-1 Create a device on an FSI CIFS share

   ```
   \\10.0.20.134\networker
   ```

   For Linux systems, the path definition looks similar to Example 16-2.

   Example 16-2 Create a device on an FSI NFS share

   ```
   /mnt/fsi_shares/networker
   ```

4. In the next window, click **Backup pool type** and click **Next**.
5. In the next window, set the Number of Target Sessions (number of client's connections in parallel) and click **Next**.
6. The next window should display a summary of the information. Click **Configure**.
7. After the configuration completes, you should see a confirmation window. Click **Finish**.

16.3.4 Setting the information to be backed up

Now that you have set up your device, you must define which information you are backing up to this device. To set up the information to be backed up, complete the following steps:

1. In the NetWorker Administration window, click **Configuration** and click **Groups** → **Default**.
2. Choose the existing group, right-click it, and click **Client Backup Configuration** → **Modify**.
3. In the first window that opens, do not make any changes. Click **Next**.
4. The networker data window opens, and shows the operating system, version, and so on. Verify this information and click **Next**.
5. The next window opens and shows all the directories for all the units that are available. Choose the information that you want to back up and click **Next**.
6. The next window opens and shows the times for the Browse and Retention policy. Make your required changes and click **Next**.
7. The next window opens. You could change the Group for this policy. For this example, do not change any data and click **Next**.

[1] You can find NetWorker documentation at the NetWorker Information Hub at http://nsrd.info/docs.html.

8. In the next window, you could choose another Storage Node for Backup and Restore. For this example, do not change anything and click **Next**.
9. A summary window opens. If all the information is correct, click **Modify**.
10. A confirmation window opens. Click **Finish**.

16.3.5 Setting the time for the backup

To either set the time for your backup or to start the backup now, complete the following steps:

1. In the NetWorker Administration window, click **Configuration**.
2. Click **Groups**, right-click **Default Group**, and then click **Properties**.
3. In the Properties dialog box, you can change the time in the Setup pane. In the drop-down box menu, you can either choose whether the backup starts automatically or at a later time. Choose your option and click **OK**.

16.3.6 Performing a restore

To perform a restore, open the NetWorker User, select the server that backed up the information, click the **Recover** icon, select the information that you want to restore, and click the **Recover** icon to start the restore.

16.3.7 Parameters for best performance with ProtecTIER FSI

This section describes how to get in to diagnostics mode to run diagnostic tests on devices. This section also describes how to turn off compression and encryption in NetWorker.

Setting compression and encryption to none

To set compression and encryption to none, complete the following steps:

1. In the NetWorker Administration window, click **Devices**.
2. Click **View** → **Diagnostic Mode** on the toolbar. The **Devices** window opens.
3. Click **Devices** in the navigation tree (on the left). The Devices detail table appears.
4. Double-click the mounted device from the detail table. The Device Properties window opens.
5. From the Device Properties window, click the **Cloud** tab and, under the Options pan, you find the Compression and Encryption settings.

6. Set Compression and Encryption to **none** (Figure 16-2).

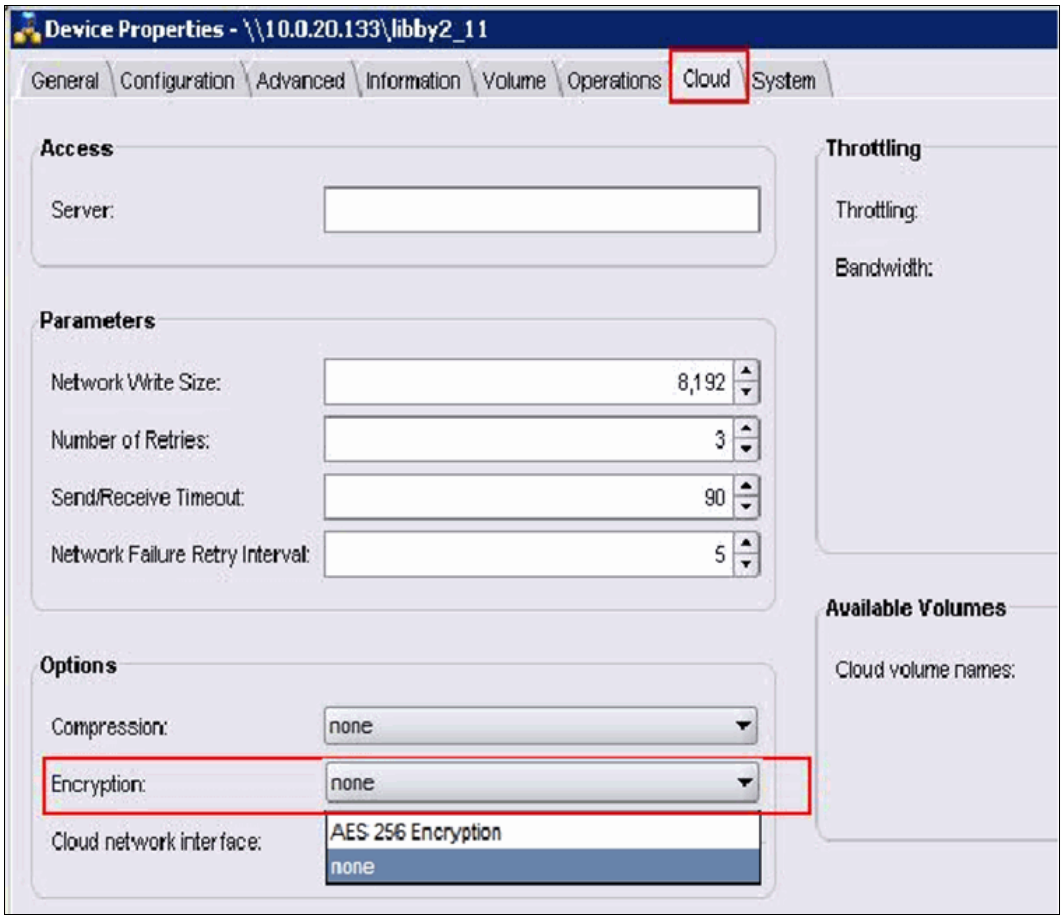

Figure 16-2 Device Properties window - disable encryption

17

HP Data Protector

This chapter describes the settings and parameters that should be modified in HP Data Protector (HPDP) in ProtecTIER environments to enable maximum performance.

Most of the changes and settings are related to the Virtual Tape Library (VTL) mode of the ProtecTIER product, although the File System Interface (FSI-CIFS) mode includes straightforward steps to enable deduplication in your HP Data Protector storage that is provisioned by ProtecTIER.

This chapter describes 17.1, “HP Data Protector with ProtecTIER” on page 262 in a VTL environment, and a brief overview of HP Data ProtecTor with ProtecTIER in an FSI environment.

This chapter describes the following topics:

- HP Data Protector with ProtecTIER
- HP Data Protector in a VTL environment

17.1 HP Data Protector with ProtecTIER

The HP Data Protector (HPDP) supports a wide range of applications, disk and tape devices, data protection tools, and any hypervisor, typically from one centralized console. The HP Data Protector also provides protection for Microsoft Exchange, Microsoft SharePoint, Microsoft SQL Server, Systems Applications and Products (SAP), Oracle, and many more business critical applications in enterprise environments.[1]

The IBM ProtecTIER deduplication solution is compatible with HP Data Protector installed as a backup application. The deployment of ProtecTIER with HPDP is supported by IBM as a VTL or FSI-CIFS.

This section provides basic configuration steps and parameters that are recommended to be enabled in HP Data Protector to achieve good deduplication results, optimal backup and restore performance, and system stability.

17.1.1 HP Data Protector architecture with ProtecTIER

The HPDP infrastructure is based on a *cell* layout, as shown in Figure 17-1. The cell is a network environment, typically in the data center or at a remote branch office, that consists of a *Cell Manager*, *client systems*, and *devices*. The Cell Manager is a central management point where HPDP is installed and operated. It incorporates the Data Protector *Internal Database* (IDB) that maintains the references about all backup files from clients. After the Cell Manager is installed and configured, you might want to add systems to be backed up (protected). These systems become client systems.

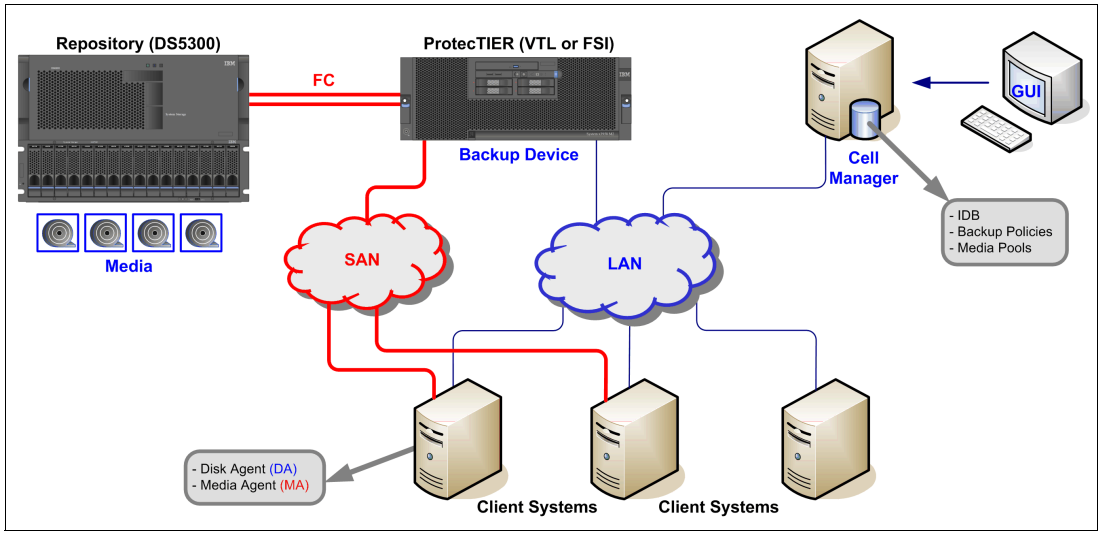

Figure 17-1 The cell diagram of HP Data Protector with ProtecTIER as a VTL or FSI

When client systems send backup files to the HP Data Protector, it stores them on *media* in backup *devices*. In the environments with IBM ProtecTIER deduplication solutions, these media are represented by virtual tape cartridges. Devices are depicted by virtual tape libraries that are provisioned by a ProtecTIER server (VTL mode of the ProtecTIER product). In the case of a FSI-CIFS installation of the ProtecTIER product, the HP Data Protector accesses the CIFS shares from the ProtecTIER server as a media repository for backups.

[1] Source: http://www.hp.com

17.2 HP Data Protector in a VTL environment

The ProtecTIER product can be configured either as VTL or FSI, but not both modes concurrently. This section describes the setup and configuration in a VTL environment for backup and restore.

The Cell Manager is the key system in the cell. It is responsible for the following tasks:

- Operates the cell from a single control point.
- Maintains the IDB records with references to back up data, media IDs, and backup sessions.
- Runs the core HP Data Protector software.
- Runs Session Manager, which triggers backup and restore jobs and records the results in IDB.

On the client systems, there are two core components:

- Disk Agent (DA), also called Backup Agent, which performs a backup of static files in the operating system. An additional utility, called Application Agent, can back up databases online without disruption to the database services.
- Media Agent (MA) offers an interface to the directly attached or SAN-zoned media devices, including virtual tape drives allocated from your ProtecTIER server. To grant control over the attached SCSI tape library to the source systems for direct backup, use General Media Agent.

The following sections provide the key HPDP parameters and configuration steps that are needed to improve your backup and restore performance with the ProtecTIER server as a VTL.

The HP Data Protector offers three types of disk-based backup devices:

- The stand-alone file device is the simplest disk-based backup device. It contains a single slot, to which data can be backed up. The maximum capacity is 2 TB.
- The file jukebox device consists of multiple slots to which you can back up data. Each slot offers a maximum capacity of 2 TB. These volumes must always be created manually.
- The file library device is similar in concept to the file device class in IBM Tivoli Storage Manager. It has multiple slots that are called file depots, which are created and deleted automatically as required. The maximum capacity of each slot is also 2 TB.

For the FSI type of ProtecTIER operation, the File Library device should be used. Select a slot (file depot) capacity 100 - 300 GB as the best practice for optimal performance.

Tip: When an FSI-enabled ProtecTIER server is used together with HP Data Protector, define the maximum slot (file depot) capacity as 100 - 300 GB. The maximum slot capacity size is based on the complexity of your backup infrastructure and an average daily backup volume. In typical medium to large environments, set the file depot size to 200 GB.

After the capacity is configured, do not change it; doing so leads to inefficient space management in the ProtecTIER/HPDP file library.

17.2.1 Enabling the robotic barcode reader

The IBM ProtecTIER Deduplication Gateway in a VTL configuration emulates an IBM System Storage TS3500 Tape Library with Ultrium LTO3 tape drives. It also emulates the numbering of virtual LTO3 tape cartridges that use virtual barcode labels. This feature allows backup applications to perform an inventory of the tape cartridges in the virtual library the same way a real physical tape library with a barcode reader performs an inventory.

To reduce the time that is needed to perform media inventory in the HP Data Protector Backup Device (in our case, VTL), enable robotic barcode reader support. Use these virtual barcodes as a medium label.

Figure 17-2 shows how to enable a robotic barcode reader by selecting devices from the **Environment** folder from the Devices and Media window. Select the **Barcode reader support** check box and the **Use barcode as medium label on initialization** check box.

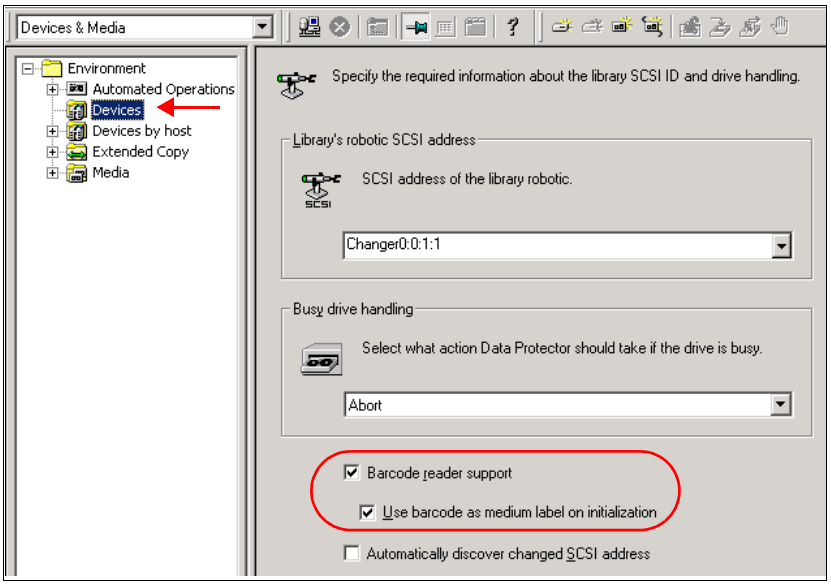

Figure 17-2 Enable robotic barcode reader

On the same window, at the bottom, also enable SCSI reserve/release robotic control to reserve the robotic control only for HP Data Protector operations. This action restricts other systems from working with SCSI reservations on the same target SCSI device, if these systems are zoned to different servers. A typical example is an installation process of a new server, where LAN-free backups are planned and tape zoning is prepared in advance.

17.2.2 Increasing the tape block size

Increase the block size of the tape device from the default value of 64 KB to a minimum of 256 KB and up to 1 MB as the maximum. The ProtecTIER product supports all values up to a 1 MB block size, but a block size of 512 KB provides the best deduplication ratio and performance balance. Increasing the block size reduces the number of virtual tape headers.

Figure 17-3 shows how to increase the tape block size to 512 KB. In this example, select the VTL_LTO3_01 drive from the **Drives** folder. Then, click the **Settings** tab, and you see your Default Media Pool, VTL_TAPE_ONSITE. Click the **Advanced** tab and the Advanced Options window opens. From the Sizes tab, click the **Block size (kB)** menu and set it to 512 KB.

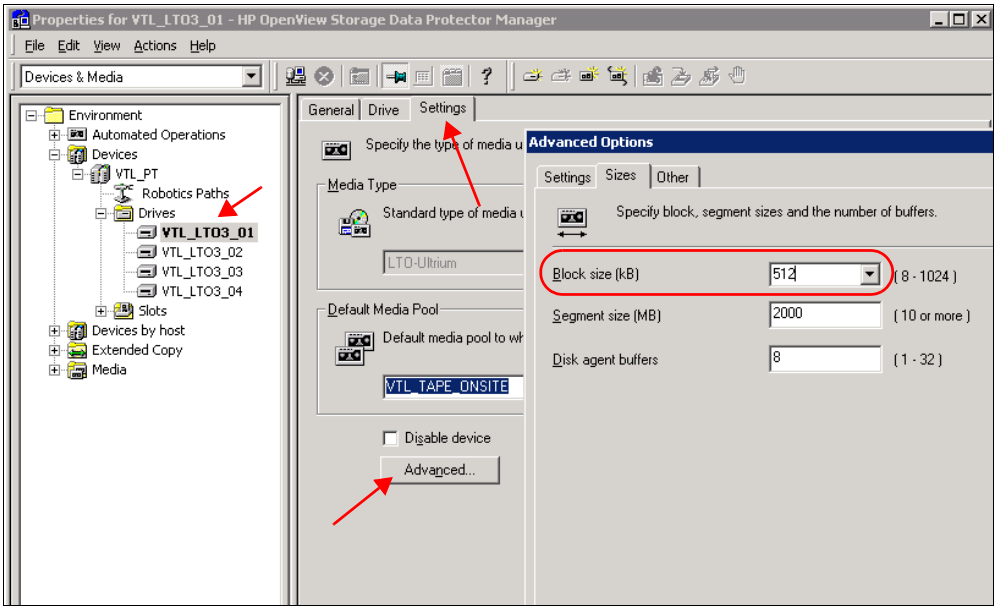

Figure 17-3 Increasing the tape block size

If your Media Agent is running on a Windows server, you must modify the registry.

Important: Before you make any changes in the registry, make a consistent backup.

To modify the registry, complete the following steps:

1. Open the registry editor by running **`regedit.exe`** and navigating to the following path:

   `\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ql2300\Parameters\Device`

   Figure 17-4 shows the registry key.

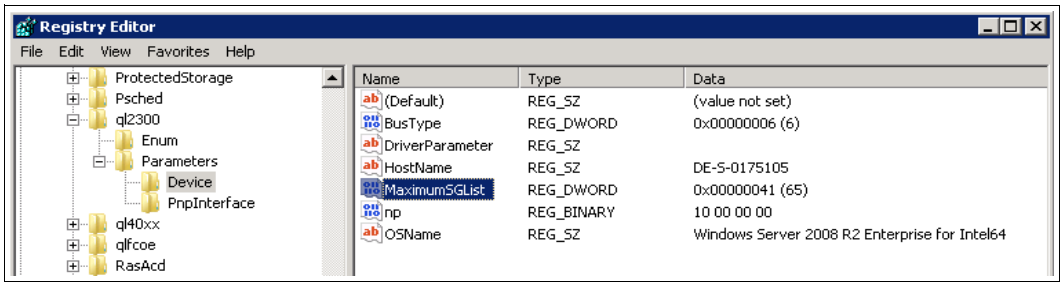

Figure 17-4 Windows registry records of QLogic HBA

Note: The registry value MaximumSGList exists only in Windows 2008 by default. In Windows 2003, you must create it by adding a DWORD of the same name.

2. Modify the value of the tape block size to 512 KB (524288 bytes). Enter this value as a decimal number, not hexadecimal (Figure 17-5). The formula to calculate the “Value data” is as follows:

   “`Value data`**”**`=(wanted Block Size in bytes)/4096+1,`

   in our case:

   “`Value data`**”**`=524288/4096+1=`**`129`**

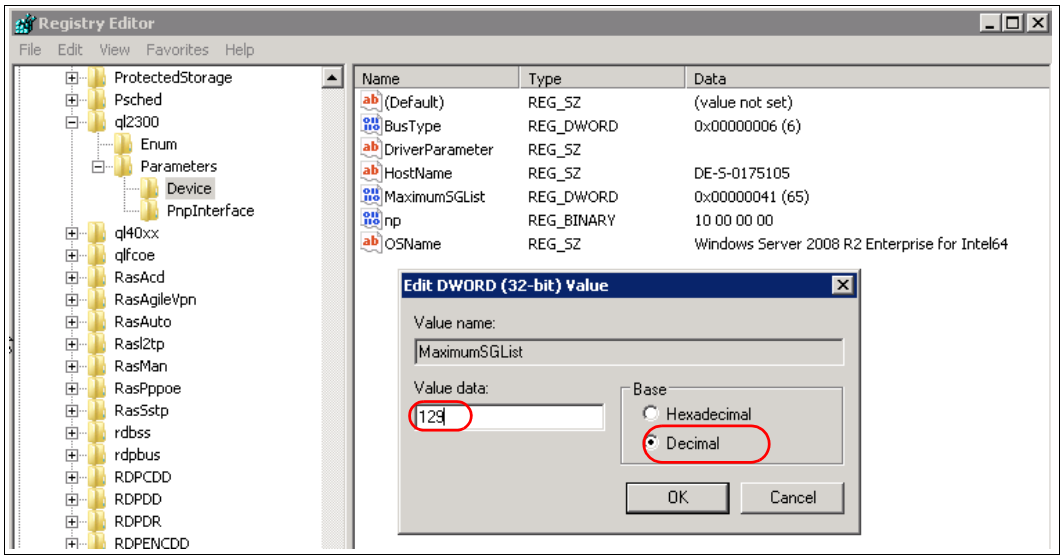

Figure 17-5 Setting a greater tape block size

3. Reboot your server to make the changes effective.

Important: The MaximumSGList parameter is a value 0 - 255. Use this adjustment sparingly, as scatter/gather lists use non-paged memory. Avoid setting it arbitrarily to a high value. For more information, go to the following website:

`http://support.microsoft.com/kb/280793`

To increase the block size above 256 KB for HP-UX, AIX, Linux, or Solaris with the HPDP, add the following variable to the `/opt/omni/.omnirc` file:

`OB2LIMITBLKSIZE=0.`

To allow a block size of 1024 KB, set the **`st_large_recs`** parameter to a non-zero value.

17.2.3 Enabling the lock name

The virtual robotic arm in your VTL is responsible for maintaining records of where all virtual cartridges are. It tracks which tapes are in each storage slot, entry/exit ports, the robotics gripper, and tape drives. These robotic arms are available to all Media Agents where direct backup to tape is needed.

Enable the lock name to prevent a collision when HP Data Protector tries to use the same physical device in several backup sessions at the same time. In this example, select the VTL_LTO3_01 drive from the **Drives** folder. Then, click the **Settings** tab, and you see your Default Media Pool, VTL_TAPE_ONSITE. Next, click the **Advanced** tab and the Advanced Options window opens. Click the **Other** tab and select the **Use Lock Name** check box (Figure 17-6).

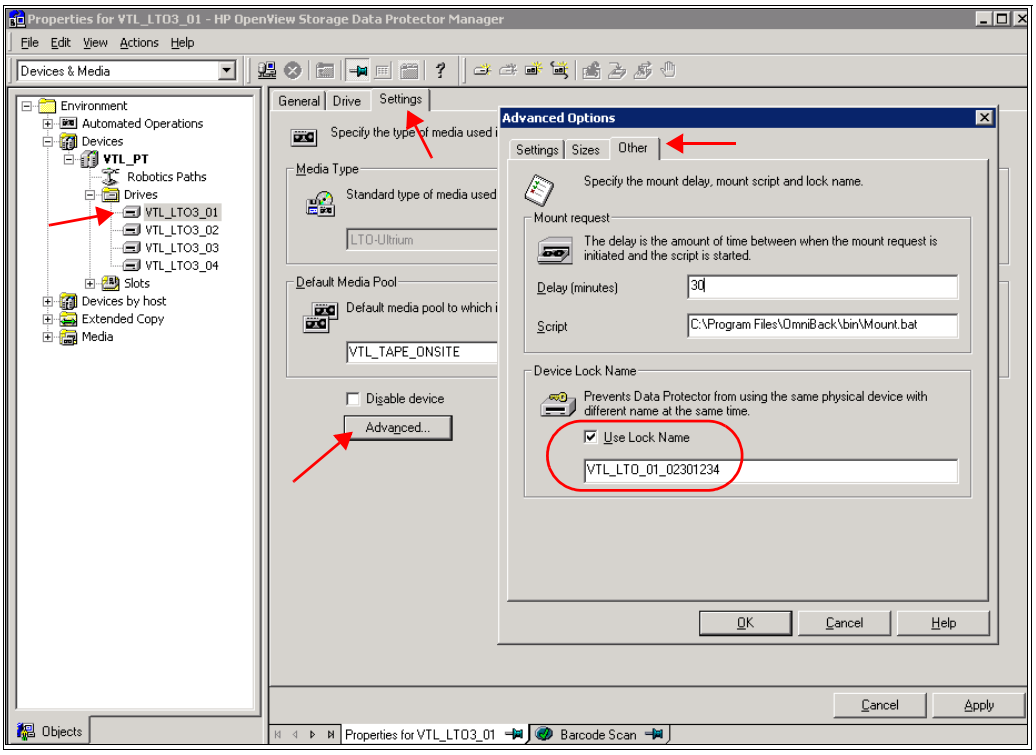

Figure 17-6 Enable lock name

17.2.4 Disabling compression, encryption, and CRC chksum

Compression, encryption, hash-based application deduplication, Cyclic Redundancy Check (CRC) of data integrity, and other processes, impact the ProtecTIER deduplication factoring ratio. Avoid using these techniques when the IBM ProtecTIER solution is implemented as a component of your backup environment.

Although HP Data Protector offers these features, ensure that they are disabled for virtual tape devices that provisioned by the ProtecTIER product. Figure 17-7 shows how to deactivate those options by clearing them. Disable compression, encryption, and CRC by selecting a file system, in this example FS_Daily_Incr. Then, click the **Options** tab, and the Filesystem Options window opens. Clear the **Software compression** and **Encode** check boxes.

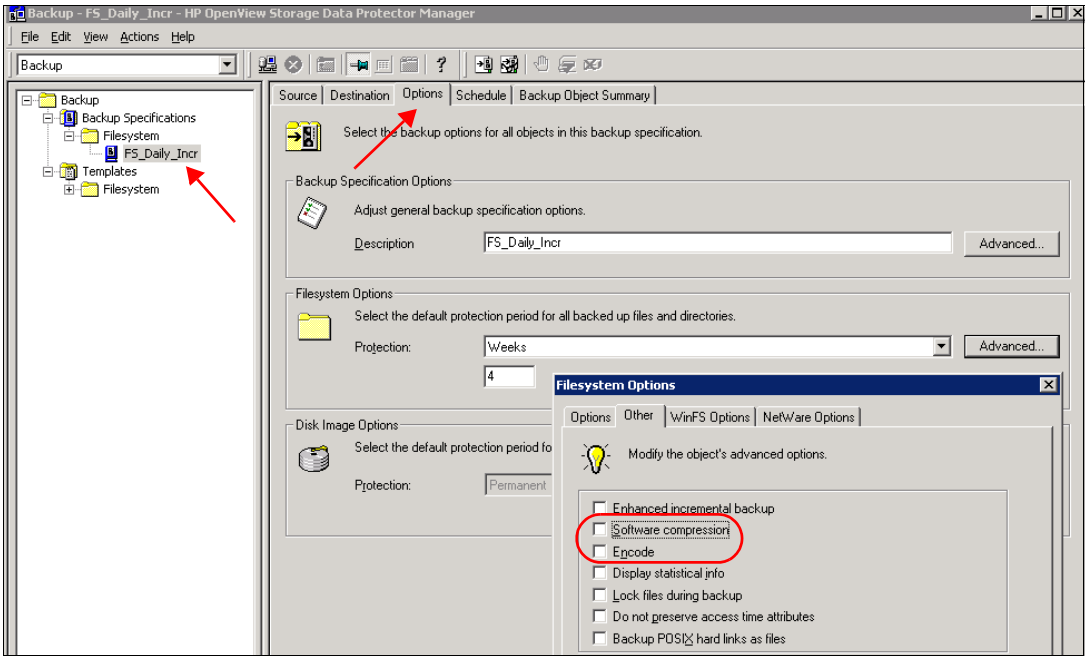

Figure 17-7 Disable compression, encryption, and CRC

You can use HP Data Protector to use multiple Media Devices for different types of backups, or for different client systems (Disk Agents). For certain devices, such as physical tape drives or external disk arrays, it is beneficial to have these storage efficiency techniques enabled. Activate them individually for each backup process to these Media Devices, and ensure that the ProtecTIER virtual tape repository is not affected. This action is exceptionally important when deduplication is enabled.

17.2.5 Hosts multipath support

The TS7650G is shipped with two Emulex dual-port LPe12002 8 Gbps Fibre Channel adapters for the front-end connections to the backup server and, optionally, to the LAN-free client systems (by using a SAN infrastructure). For high availability or performance reasons, you might use more than one HBA on the server to access virtual tape libraries and drives. However, if you implement this configuration, you might see duplicate virtual tape drive definitions or robotics on your server, each accessible through a different HBA.

HP Data Protector supports this multipathing function, although certain settings must be defined. Before you make such changes, ensure that the Datapath and Control Path Failover is supported and enabled by the IBM Tape Device Driver that is installed on a relevant platform.

Hint: Use a dedicated IBM Tape Device Driver for ProtecTIER virtual tape devices. Ensure that multipathing is enabled in the operating system before you enable multipathing in the HP Data Protector Devices and Media Manager.

For examples about how to enable multipathing on a tape device (either robotics or tape drive) in HPDP, see Figure 17-8 and Figure 17-9.

Figure 17-8 shows how to enable multipathing for the robotic path for a tape library. From the Devices & Media drop-down menu, select the VTL_PT library, and select the **MultiPath device** check box.

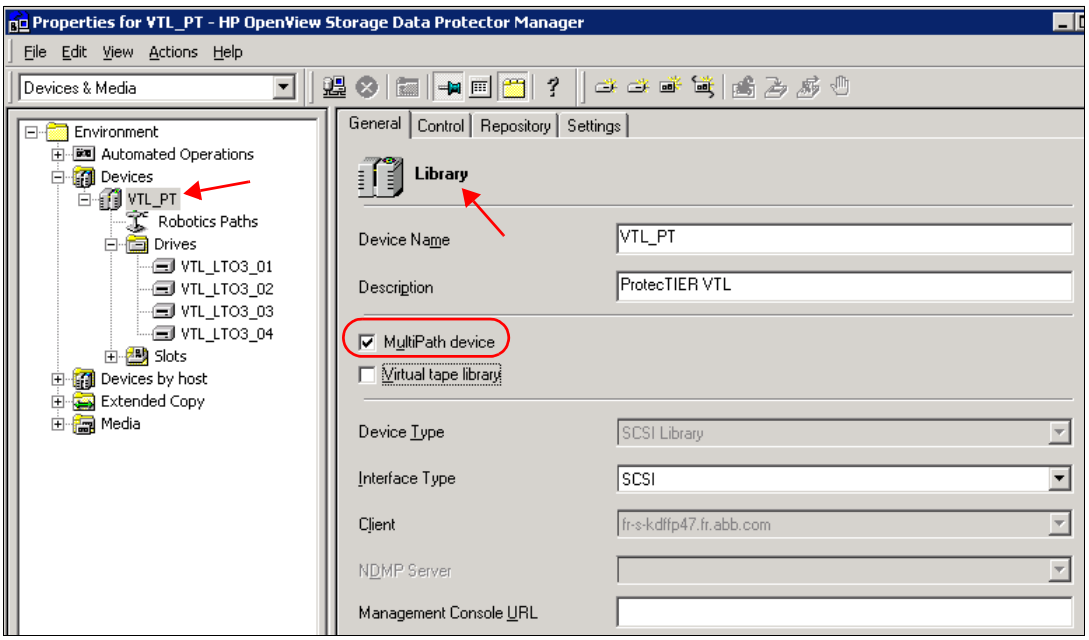

Figure 17-8 Multipath device on library robotics device

Figure 17-9 shows how to enable multipathing for the robotic path for a drive in a tape library. From the Devices & Media drop-down menu, select the drive (in this case, VTL-LTO3_01), and select the **MultiPath device** check box.

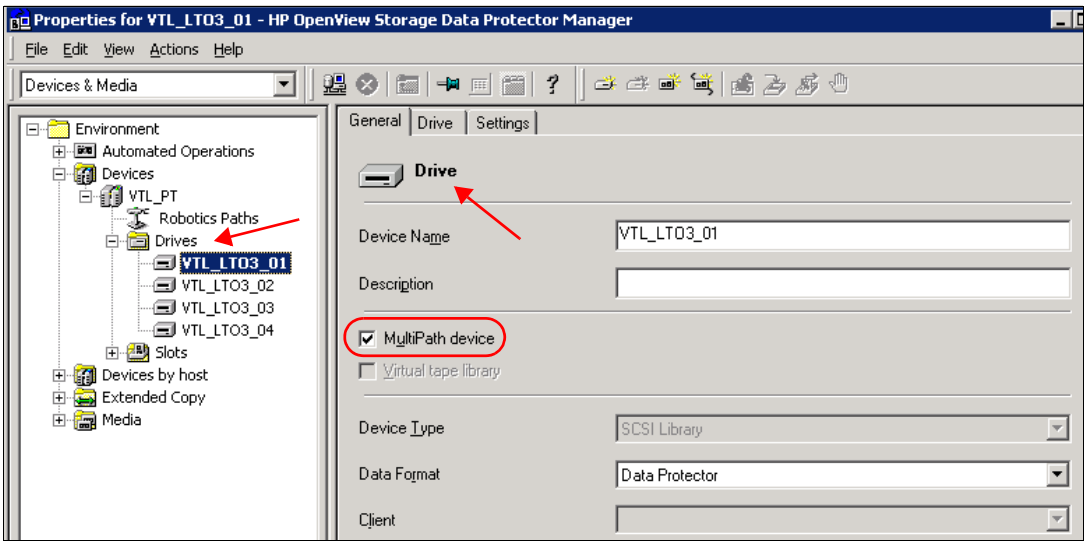

Figure 17-9 Multipath device settings for tape drives

Although the activation of multipathing in the operating system (by using a tape device driver) requires a reboot of the server, the setting of a specific device with multipathing in HP Data Protector is a dynamic process and does not require a reboot. This change is effective immediately.

17.2.6 Load balancing

The load balancing feature specifies how many tape drives are available (minimally) at the start of the backup and how many in total are reserved at the start of this backup to the client system. The Data Protector uses the optimum amount of resources that are required for the tape devices workload. Using inappropriate settings of both parameters can make reserved tape drives unavailable for other backup sessions until the current backup job finishes.

Figure 17-10 provides an example of a setup for load balancing during a backup operation:

1. Select **Backup** from the drop-down menu.
2. For the backup specifications, select the FS_Daily_incr file system.
3. Click the **Destination** tab. Five tape drives are displayed in the Backup pane. Select the VTL_LTO3_04 and VTL_LTO3_05 tape drives, and set the minimum number of tape drives that are available to run backups to 2. Set the maximum number of tape drives to reserve to 5.

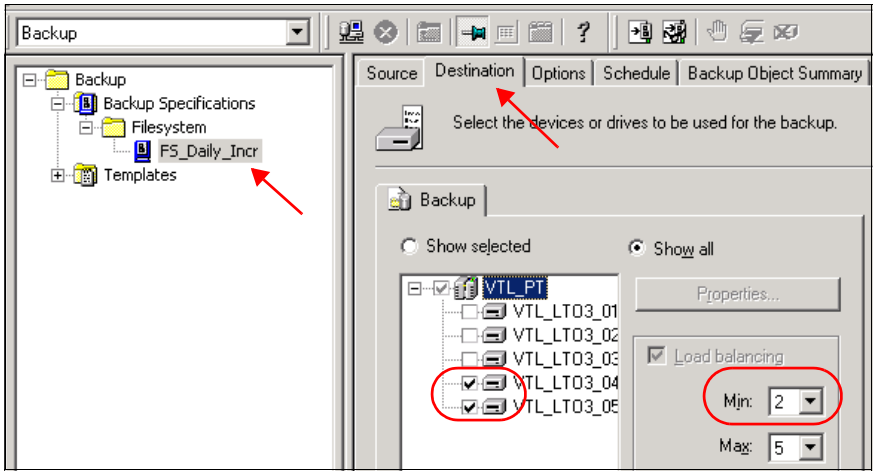

Figure 17-10 Load balancing

In this scenario, HP Data Protector checks whether there is a minimum of two tape drives available for backup; if yes, the backup starts in two parallel streams. During backup, the HPDP can reserve up to five tape drives, making those drives unavailable to other client systems that might start backup later that day. HP Data Protector selects any available drive that you assigned in the backup specification for this load balancing.

Tip: Do not use load balancing when you back up a few large objects, typically databases. In this scenario, HP Data Protector is often not able to balance the load among such devices effectively. Define a dedicated backup policy with disabled load balancing for such objects.

17.2.7 Using a mirroring functionality

You can use the HP Data Protector object mirror devices functionality to write the same data to several media simultaneously during a backup session. In some cases, this feature can replace vaulting or migrating the data between libraries and can decrease the usage of resources and increase performance of backup jobs.

Modify the backup policy settings, as shown in Figure 17-11. The figure shows the usage of a mirror. To add the object mirror functionality:

1. Select the file system; in this example, we select FS_Daily_Incr.
2. Click the **Destination** tab and select the relevant drives. As illustrated in Figure 17-11, from the Mirror 1 tab, we select the VTL_LTO3_01 and VTL_LTO3_02 drives.
3. Click the **Add Mirror** tab.

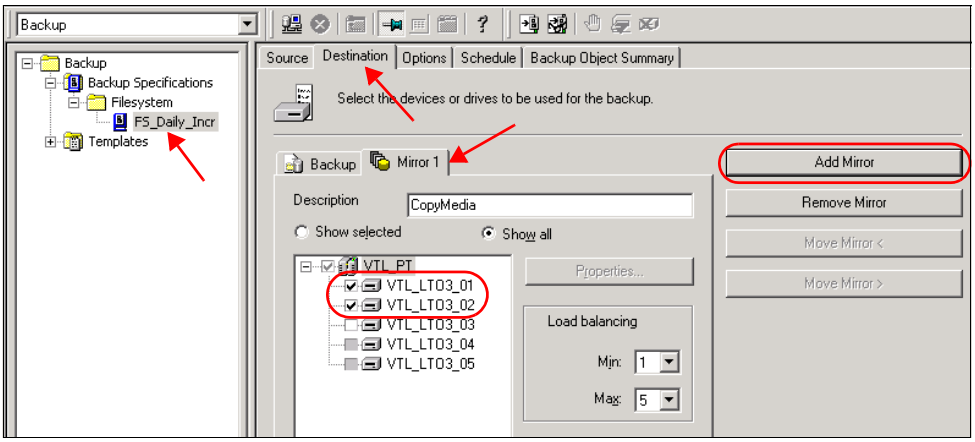

Figure 17-11 Usage of a mirror

17.2.8 Troubleshooting logs

The following information might assist you with preliminary analysis troubleshooting:

- CLI output commands from the Data Protector host server:
  - **`#devbre –dev`**
  - **`#sanconf –list_drivers`**
- Data Protector log files are in the following directories:
  - For Windows systems:

    `<Data_Protector_home>\log`
  - For HP-UX:

    `/var/opt/omni/log and /var/opt/omni/server/log`
  - For other UNIX systems:

    `/usr/omni/log`

For more information that is related to troubleshooting of HP Data Protector and log analysis, go to the following website:

http://support.openview.hp.com/selfsolve/document/KM1120178/binary/DP6.20_Troubleshooting.pdf

The general information and documents library for HP Data Protector is available at the following website:

http://h41112.www4.hp.com/promo/imhub/data_protector/documentation.html

18

IBM i and Backup, Recovery, and Media Services

The ProtecTIER product extends the disk backup solution for IBM i with minimal disk usage by using the ProtecTIER HyperFactor deduplication feature.

The ProtecTIER product is deployed as a Virtual Tape Library (VTL) that emulates tape devices for IBM i. IBM i stores its backup data in ProtecTIER virtual cartridges as though they are physical cartridges. Users have the flexibility to create virtual tape drives and virtual cartridges that are based on their needs. IT centers that require a copy of backup data to be offsite can use ProtecTIER replication to replicate the virtual cartridges to a secondary site.

Because the ProtecTIER product emulates tape devices, you can always share it between IBM i and Open Systems. It is highly recommended to conduct a correct sizing based on your environment before the implementation.

This chapter focuses on using Backup, Recovery, and Media Service (BRMS) to integrate the ProtecTIER product in to an IBM i environment.

This chapter describes the following topics:

- IBM i overview
- Integration of IBM i and ProtecTIER in a VTL environment
- Configuration of BRMS for ProtecTIER
- Deploying ProtecTIER with BRMS for disaster recovery

18.1 IBM i overview

This section describes some IBM i specific features, which you should be familiar with to understand how IBM i backups are performed.

Today, when you purchase IBM i, you purchase IBM POWER® hardware that can run AIX, Linux on Power, or IBM i. You load your operating system in to each logical partition (LPAR); in this example, we load IBM i.

IBM i (formerly known as IBM AS/400®, IBM eServer™ iSeries®, and IBM System i®) is an integrated system that includes hardware, software, security, a database, and other components. IBM i has a flexible architecture where software is independent from hardware, so changing one has little impact on the other one.

IBM i and ProtecTIER: For information about IBM i resources for the ProtecTIER product, see the following website:

http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS4956

18.1.1 Integrated file system

IBM i contains an integrated file system (IFS) that provide a common interface with other file systems. This file system allows applications that are written on other file systems, such as UNIX and Windows, to access data that is stored in IBM i.

18.1.2 Integrated database

DB2 is the integrated relational database in IBM i. DB2 is tightly integrated into the entire system, making it efficient and easy to use by IBM i applications. DB2 is used by various applications, from traditional host-based applications to client/server applications to business intelligence applications.

18.1.3 Object-based architecture

IBM i is an object-based operating system. Unlike most other operating systems where everything is seen as a file, IBM i sees everything as an object. These objects include database files, user profiles, job queues, compiled programs, word-processing documents, and so on. Objects are categorized by type, which allow the users to specify what type of objects are required for a task. The IBM i operating system provides an interface to define operations that can be performed on objects and to provide instructions about the usage of the encapsulated data.

18.1.4 Libraries

IBM i groups objects into libraries. A library is a directory of objects, and it is an object that is used to find other objects in the directory. An object can exist in only one library. A library cannot reference other libraries except for the library that is called QSYS, which contains all the other libraries. Libraries can be associated or referenced by user profiles or applications.

18.1.5 Backup considerations in IBM i

IBM i offers a wide range of backup recovery options. They are intended to help you accomplish the following tasks:

- Make your save operations faster and more efficient.
- Keep your system available for your users.
- Plan and manage your backup and recovery.

Before you implement a backup solution in IBM i, consider the following items:

- Determine the save objects and how often to save them.
  - Consider performing daily saves of the libraries and objects that regularly change, such as application libraries, user profiles, configuration objects, and parts of IFS. The objects that regularly change have a higher possibility of restoration in a shorter period as compared to objects that do not change regularly.
  - Consider performing weekly full system saves. The full system saves provide a baseline copy of your system that can be used for restoring all system objects in the event of a disaster. Alternatively, save all user libraries (*ALLUSR) every week.
- Determine the save window that is based on the following items:
  - Affordable downtime of save objects.
  - Affordable IBM i system downtime for a full system backup.
- The recovery time and availability options.
- Test the developed backup and recovery strategy.

18.2 Integration of IBM i and ProtecTIER in a VTL environment

This section describes the considerations and best practices for configuring IBM i in a ProtecTIER environment.

18.2.1 Backup considerations with ProtecTIER

Using VTL is not necessarily faster than physical tape backup. IBM tape products have been tested and work efficiently with IBM i. IBM i is able to achieve 90% - 100% of tape drive speed in an environment with fewer tape drives. You often require multiple streams in a VTL to achieve the same performance throughput as physical tapes. In this scenario, Backup, Recovery, and Media Service (BRMS) is useful in managing the tape media resources for parallel saves.

In addition to performance throughput, you can use BRMS to share VTL resources across multiple LPARs.

BRMS tracks what you saved, when you saved it, and where it is saved. When you need to do a recovery, BRMS ensures that the correct information is restored from the correct tapes in the correct sequence.

18.2.2 Recommended ProtecTIER and IBM i configuration

Here are some guidelines for configuring ProtecTIER in your IBM i environment:

- Configure parallel or concurrent saves to reduce save windows.
- Use BRMS for automation. Ensure that you are licensed for BRMS.
- The ProtecTIER product supports IBM i from Version 5 Release 4 onward, and both the IOP and IOPless Fibre Channel (FC) adapters are supported. For more information, check the TS7650/TS7650G ISV and Interoperability Matrix for IBM i compatibility with ProtecTIER at the following website:

  http://www-03.ibm.com/systems/storage/tape/resources.html#compatibility
- Create a separate VTL if you are sharing one ProtecTIER system for IBM i and other platforms.
- Do not mix disk and tape on one FC adapter for performance reasons, as a tape workload is always large block while a disk workload is typically small block. The ProtecTIER solution is VTL; it is considered a tape workload when integrated with IBM i.
- When you create virtual tape devices in ProtecTIER Manager, choose IBM TS3500 for your library type and LTO3 (ULT3580-TD3) for the drive type.
- Carefully design the drive attachment and robot attachment on each port (see Figure 18-1 for port assignment). Usually, you would spread the drives across all ports for optimal performance, but in IBM i, there are certain rules for using tape that you might want to consider. We list some of the rules in the next bullet.

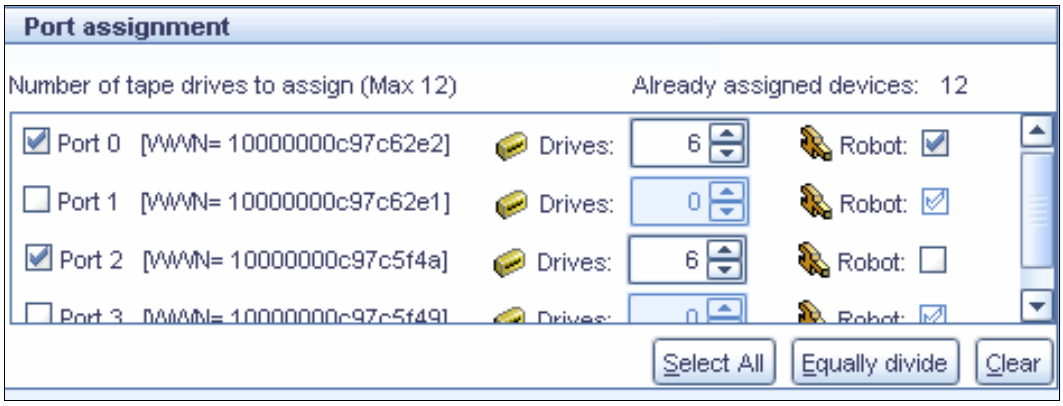

Figure 18-1 Port assignment

- IBM i cannot split the backup evenly across the FC adapters if more than one FC tape adapter is used in one LPAR. Because we advise you to run multiple saves simultaneously for a ProtecTIER deployment, the tape library, tape drive layout, and SAN zoning must be designed carefully. For more information, see SAN design for IBM i Tape at the following website:

  http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS2997

  Here are some of the rules for designing your ProtecTIER environment with IBM i:

  – Multipath is not supported in IBM i for the FC tape adapter. If required, create two SAN zones to the tape devices, one active and one inactive. You must manually activate the inactive SAN zone when the active zone fails.

– There are a maximum number of LUNs (addresses) supported by an FC tape adapter in IBM i. The LUN count includes both drives and control paths. See Table 18-1 for an example about how to count the number of LUNs per FC tape adapter.
  - 16 LUNs on an IOP FC adapter
  - 64 LUNs per port on an IOPless FC adapter

Table 18-1 Example of LUN counts per FC tape adapter

Number of tape libraries	Number of drives per port	Number of connected ProtecTIER ports (robot enabled at all connected ports)	Number of LUN (device) count
1	8	1	9 (1 control path and 8 data paths)
1	8	4	12 (4 control paths and 8 data paths)
16	1	1	32 (16 control paths and 16 data paths)

- When you create the virtual tape cartridges, note the number of tape cartridges, as each cartridge requires a slot.
- Create extra slots for your VTL in case you need some slots for extra tape cartridges, as the library is placed offline when you want to change any of the logical library dimensions.
- Create enough import/export slots to hold all tape cartridges that you eject (replicate) every day.
- Enabling LUN masking on the ProtecTIER server is recommended.
- IBM i has a limitation of a maximum of 15000 elements (places where tape cartridges are stored, that is, slots, drives, picker, import/export slots, and so on). So, the total number of tape cartridges slots, extra slots, and import/export slots of each tape library must not exceed 15000 elements.

Elements: The maximum number of elements was increased to from 5000 to 15000 elements with the following PTFs:

- R611: MF50093, MF55406, MF55407, and MF55408
- R7.1: MF55396, MF55397, MF55409, and MF55393(MF55600)

- Do not use compression or compaction when you save to virtual cartridges on the ProtecTIER server.

18.3 Configuration of BRMS for ProtecTIER

This section describes the basic structure of the Backup, Recovery, and Media Management Services (BRMS) application for IBM i systems. Also provided is BRMS terminology, and recommended configurations of BRMS.

18.3.1 BRMS overview

You can plan, control, and automate the backup, recovery, and media management services for your IBM i systems with BRMS. Figure 18-2 shows how BRMS integrates the tape devices, which can be ProtecTIER virtual tape devices.

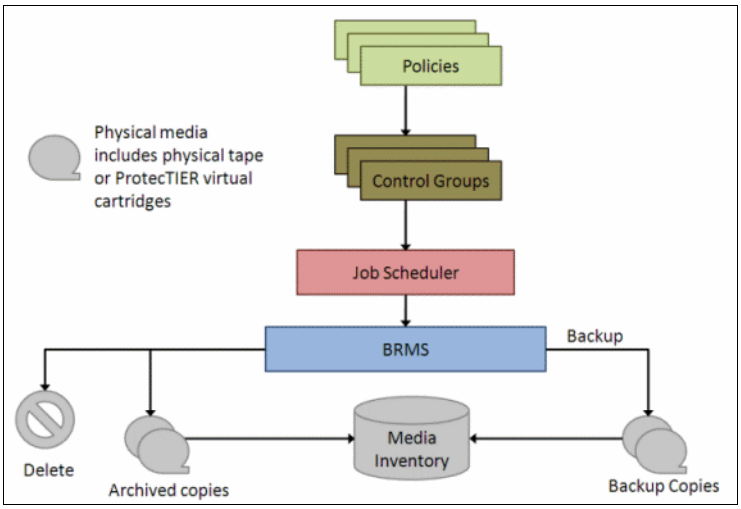

Figure 18-2 Basic structure diagram of BRMS

BRMS terminology

Here is some common BRMS terminology:

Media	A tape cartridge (volume) that holds the saved data.
Media identifier	A name that is given to a media.
Media class	A logical grouping of media with similar physical, logical, or both characteristics.
Control group	A group of libraries, objects, folders, spool files, and other types of data that share common characteristics or a group of items that you want to group it together for backup purposes. It determines which data is to be processed.
Policies	Defines how BRMS operations are processed. There are different types of policies, which include system policy, backup policy, recovery policy, media policy, and move policy. A policy determines how data is to be processed.
IBM i libraries QBRM and QUSRBRM	These libraries contain BRMS-related objects and management information.

18.3.2 Recommended configurations of BRMS

VTL versus physical tape drives: Configuration of BRMS for ProtecTIER virtual tape devices is no different from configuration for physical tape devices. You can find the relevant BRMS version guidelines and steps to configure BRMS tape devices in the IBM BRMS product information section at the following website:

http://www-03.ibm.com/systems/i/support/brms/prodinfo.html

This section focuses on the recommended BRMS configurations that are relevant to ProtecTIER virtual tape devices. Here are some guidelines:

- Use **STRSST** to reset the tape adapter that is attached to the ProtecTIER node to recognize newly created virtual tape devices in ProtecTIER system.
- The ProtecTIER VTL takes the next available TAPMLBxx name by default. Rename it to match the name that you used in ProtecTIER Manager to keep your environment orderly.
- The ProtecTIER product emulates LTO3 tape devices. When you define the BRMS media class in a ProtecTIER environment, create a separate media class for virtual LTO3 volumes from the media class for physical LTO3 volumes. Use a *ULTRIUM3 density.

  Tip: You can make the virtual LTO3 cartridges any size you like, but not the LTO3 size. The cartridges' performance matches whatever your ProtecTIER configuration allows, and is not tied to LTO3 speed.

- Add the virtual cartridges in to BRMS by running **ADDMLMBRM** command and be sure to initialize them to reflect the virtual cartridge's actual label (barcode). Use the actual command instead of navigating there from the **WRKMLMBRM** menu.
- To choose your replication destination with ProtecTIER Manager, complete the following steps:
  - If you choose a remote VTL as a destination, configure the move policy for the selected media so that the media can be moved out from the local VTL to the shelf after a save operation. The replica cartridges can then be moved to a remote VTL and be accessible in remote IBM i.
  - Do not configure a move policy if you want the media to stay in the library after save operations. You can perform a manual move when you want to access it from a remote site.

Configuring a parallel save with BRMS

You can configure a parallel save by specifying the number of parallel device resources in the BRMS control backup control group attributes. An example of the needed specified attributes is shown in Figure 18-3.

```
                    Change Backup Control Group Attributes

Group  . . . . . . . . . . . . . . . . : SAMPLE

Type information, press Enter.

Media policy for:
  Full backups . . . . . . . . . . . . . *BKUPCY     Name, F4 for list
  Incremental backups  . . . . . . . . . *BKUPCY     Name, F4 for list
Backup devices . . . . . . . . . . . . . *BKUPCY     Name, F4 for list



Parallel device resources:
  Minimum resources  . . . . . . . . . . 2           1-32, *NONE, *AVAIL
  Maximum resources  . . . . . . . . . . 4           1-32, *AVAIL, *MIN
Sign off interactive users . . . . . . . *BKUPCY     *YES, *NO, *BKUPCY
Sign off limit . . . . . . . . . . . . . *BKUPCY     0-999 minutes, *BKUPCY
Default weekly activity  . . . . . . . . *BKUPCY     SMTWTFS(F/I), *BKUPCY
Incremental type . . . . . . . . . . . . *BKUPCY     *CUML, *INCR, *BKUPCY
Force full backup days . . . . . . . . . *BKUPCY     0-365, *NOMAX, *BKUPCY

F3=Exit   F4=Prompt   F12=Cancel
```

Figure 18-3 Example of configuring parallel saves with BRMS

18.4 Deploying ProtecTIER with BRMS for disaster recovery

This section describes various scenarios for replication and performing disaster recovery (DR) by using ProtecTIER with BRMS. The first scenario describes how to recover an IBM i production system with BRMS, where BRMS is installed on both the production site and the DR server in one BRMS network. The second scenario describes how to recover an IBM i production system with BRMS installed on the production system, and there is no BRMS installed at the DR site.

18.4.1 BRMS available at the production site and DR site

In this scenario, both the production and DR sites have an IBM i server with a ProtecTIER server connected. BRMS is installed on both the production and DR server in one BRMS network, and the information about the media, devices, and so on, is synchronized between two BRMS systems.

The production IBM i performs backups to a locally attached ProtecTIER system with local BRMS. Replication is established from the production ProtecTIER system to the DR ProtecTIER system.

Replication setup

To set up replication, complete the following steps:

1. Configure the ProtecTIER replication policy with the destination as the VTL of the DR ProtecTIER server.
2. Configure a move policy for these virtual cartridges so that the virtual cartridges can be moved to the shelf after they are ejected from production. Then, the replica cartridges are moved from the shelf to the I/O station (import/export slots) of the selected VTL at the DR site when BRMS movement is run by running **MOVMEDBRM** or **STRMNTBRM MOVMED(*YES)** commands.
3. As soon as the replica cartridges are moved to the VTL at the DR site, the cartridge is on the shelf of the production system and in the IO station on the target.

Disaster at the production site

If there is a disaster at the production site, complete the following steps:

1. Enter DR mode from the DR ProtecTIER server.
2. Generate a ProtecTIER replication statistic (.csv report) that includes statistics for all replica cartridges, including sync time.
3. Review the .csv report to determine whether the cartridges you want to restore have consistent data. If they do not contain consistent data, consider restoring from an earlier set of backups, as described in “Assessing the cartridges’ status and synchronizing with the catalog” on page 420 and “Recovering the data” on page 420.
4. Restore the IBM i system from the consistent replicas of cartridges in the DR ProtecTIER server.
5. You may perform daily saves to the DR ProtecTIER server during the outage of the production system.

Failback

After the production system is running and connected to the production ProtecTIER server, complete the following steps:

1. Create a failback policy to replicate all the cartridges or just the ones that were written during the outage to the production ProtecTIER server.
2. When all the needed virtual cartridges are replicated to the production ProtecTIER server, terminate the failback replication by leaving DR mode from the DR ProtecTIER Manager.

You can now resume the daily saves from the production system.

18.4.2 No BRMS at the DR site

In this scenario, the production IBM i performs backups to a locally attached ProtecTIER system with BRMS. Replication is established from the production ProtecTIER system to the DR ProtecTIER system. There is no BRMS installed at the DR site.

Disaster at the production site

If there is a disaster at the production site, restore the entire system to a server at a DR site by completing the following steps:

1. Enter DR mode from the DR ProtecTIER server.
2. Generate a ProtecTIER replication statistic (.csv report) that includes statistics for all replica cartridges, including sync time.

3. Review the .csv report to determine whether the cartridges you want to restore have consistent data. If they do not contain consistent data, consider restoring from an earlier set of backups, as described in “Assessing the cartridges’ status and synchronizing with the catalog” on page 420 and “Recovering the data” on page 420.
4. After you obtain the list of consistent media, restore the IBM i system with the BRMS recovery report. It might be necessary to complete the following steps:
   a. Restore System i Licensed Internal Code.
   b. Restore the operating system.
   c. Restore the BRMS product and associated libraries QBRM and QUSRBRM.
   d. Restore BRMS-related media information.
   e. Restore user profiles.
   f. Restore System libraries QGPL, QUSRSYS, and QSYS2.
   g. Restore configuration data.
   h. Restore IBM product libraries.
   i. Restore user libraries.
   j. Restore document library.
   k. Restore IFS objects and directories.
   l. Restore spooled files, journal changes, and authorization information.
5. You may perform the daily saves on the recovered IBM i system at the DR site to the DR ProtecTIER server during the outage.

Alternatively, if there is an FC connection from the production system to the DR ProtecTIER server, and depending on the type of failure on the production site, you can establish a connection for the DR ProtecTIER server to the production system and restore the required objects from the DR ProtecTIER server.

Examples of BRMS policies and control groups

In this example, BRMS on a production IBM i and on a DR IBM i are in the BRMS network. Create the source VTL (TS7650SRC1) on the production ProtecTIER system, and create the target VTL (TS7650TGT1) on the DR ProtecTIER system. Both VTLs are known to the BRMS network.

Complete the following steps:

1. After the virtual cartridges are added to BRMS and initialized, set up the BRMS move policy to move the cartridges from the location TS7650SRC1 to TS7650TGT1. The cartridge stays in the TS7650TGT1 until it expires. Give the move policy a name (TS7650) (Figure 18-8 on page 285). The Media policy in Figure 18-4 specifies TS7650.

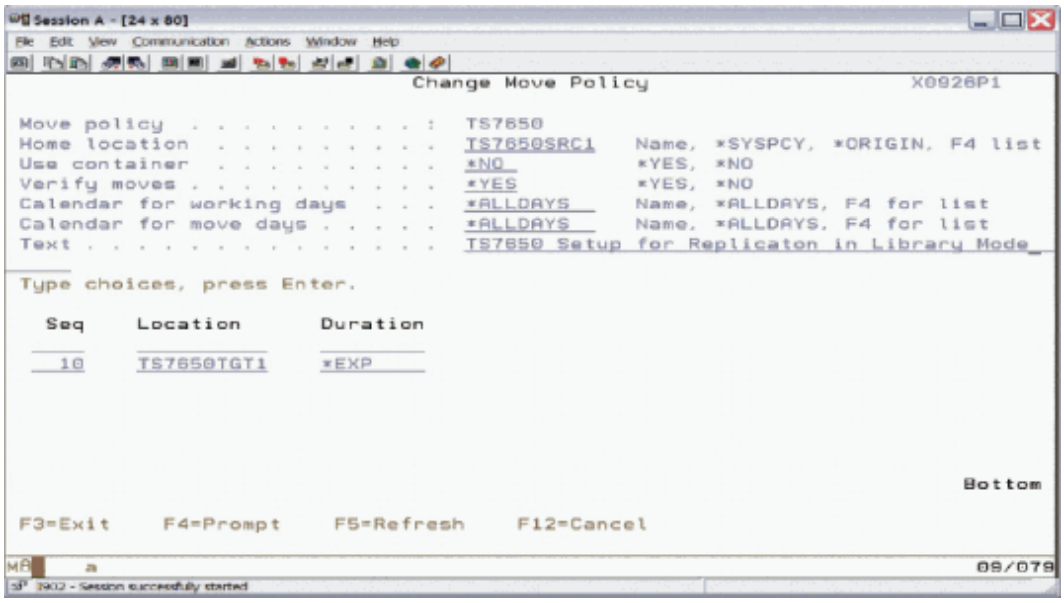

Figure 18-4 Example of a media policy

In Figure 18-8 on page 285, the Home Location specifies the source VTL. Under the Location column, type the name of the target location.

2. Create a BRMS media policy to use this move policy (Figure 18-5).

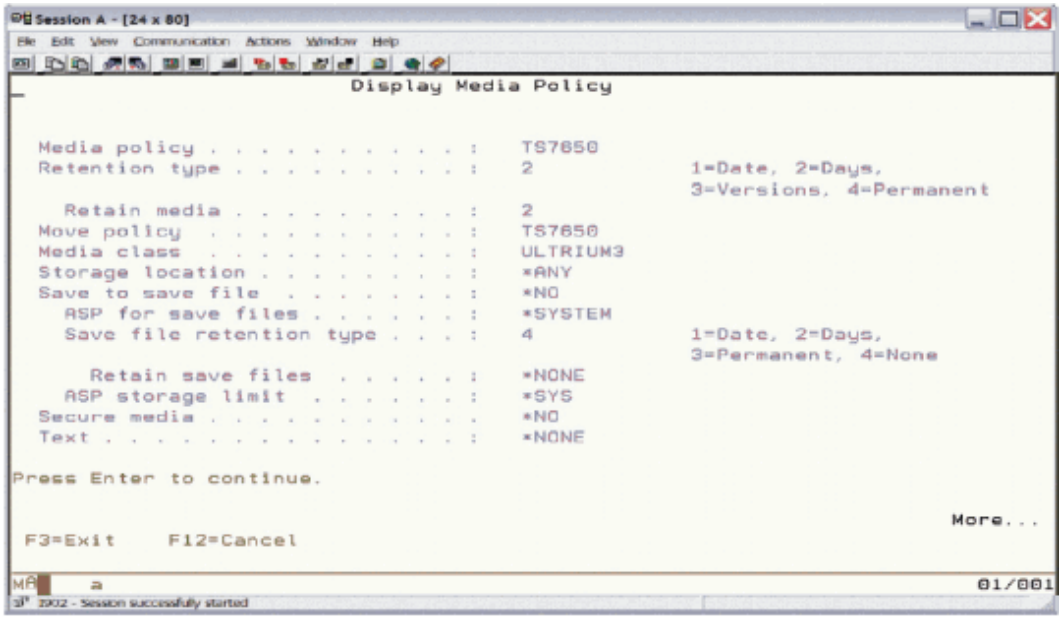

Figure 18-5 Example of change control group attributes

3. Create a BRMS backup control group to use the media policy with the VTL on the production ProtecTIER system (Figure 18-6).

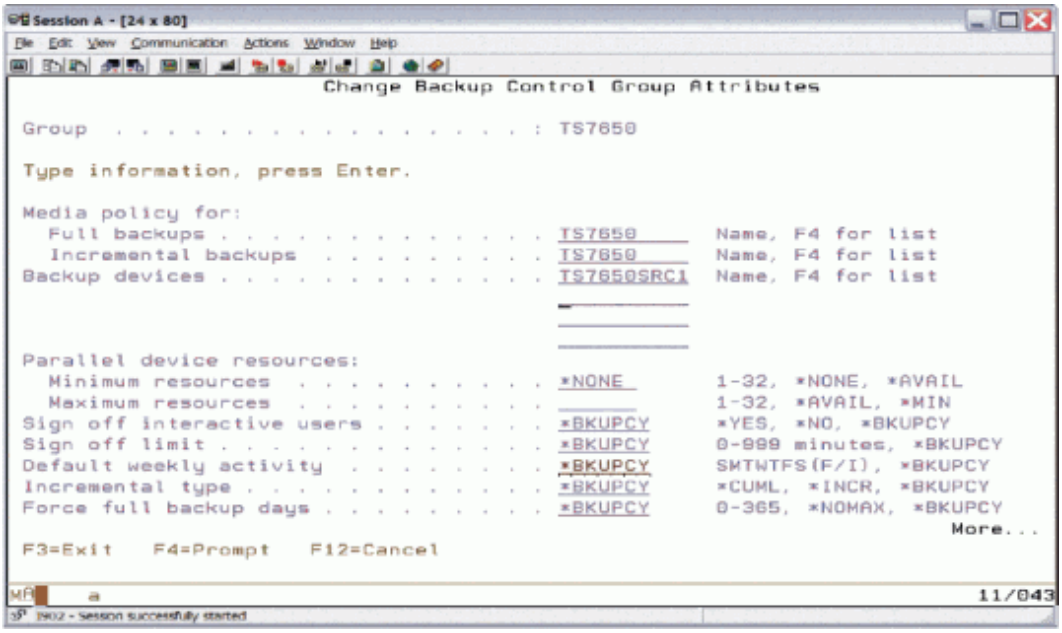

Figure 18-6 Example of verify media move

In Figure 18-6:

- The Full backups field specifies `TS7650`.
- The Incremental backups field specifies `TS7650`.
- The Backup devices field specifies `TS7650SRC1`.

4. On ProtecTIER Manager, configure the replication policy with a destination in the target VTL (TS7650TGT1). The moved media has an `Inserted` status in the DR IBM i system.
5. Verify the media move on the DR IBM i system, which makes the cartridge available (Figure 18-7).

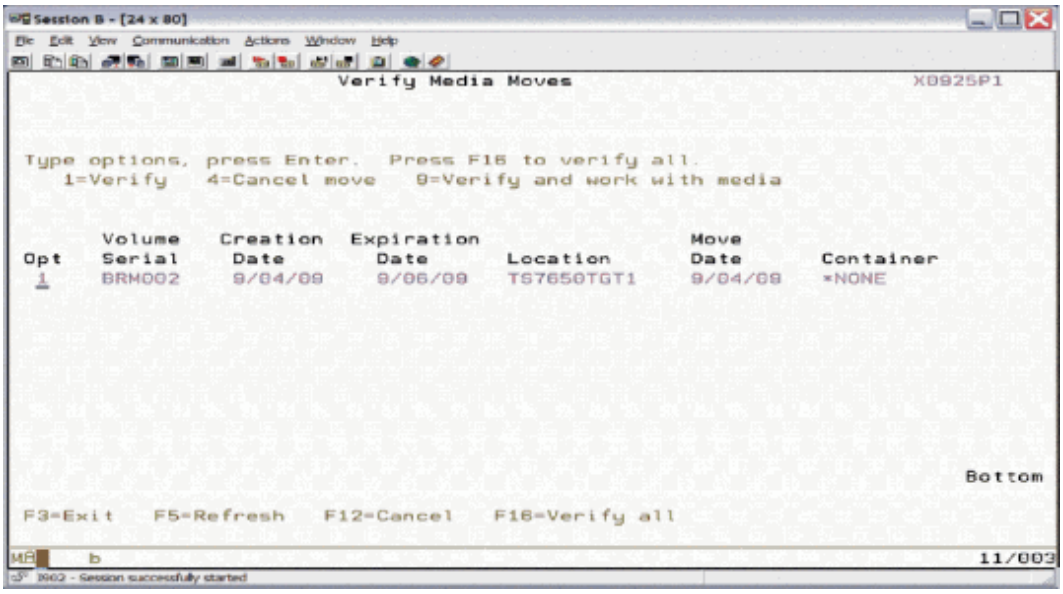

Figure 18-7 Displaying media status

The Location column specifies the name of the target VTL.

6. The cartridge is now available at DR site (Figure 18-8).

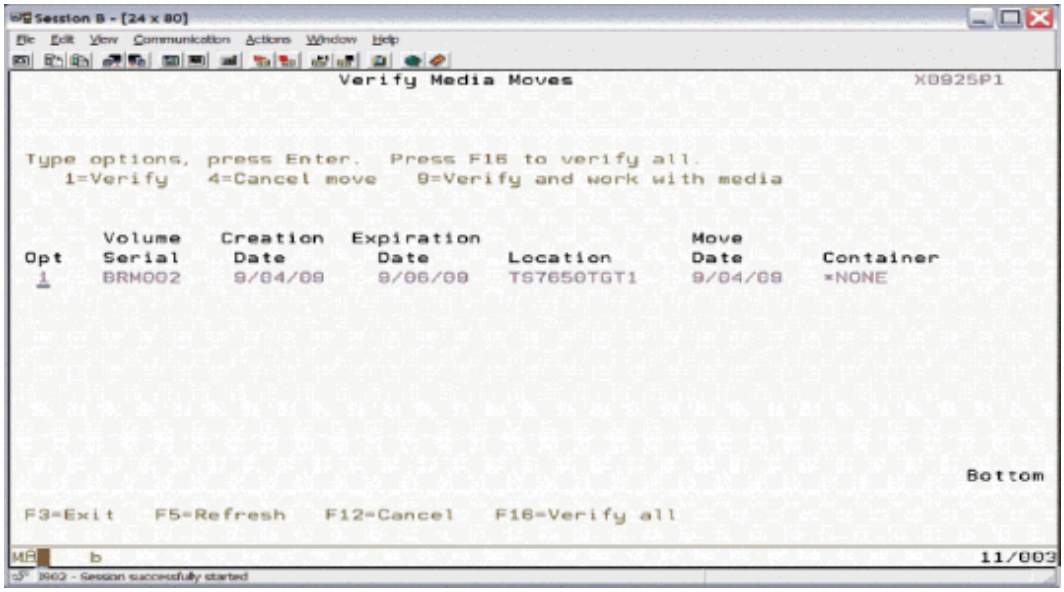

Figure 18-8 Example of change move policy

The Status column for the moved cartridges specifies `Available`.

19

CommVault

This chapter describes all the necessary steps for a successful integration of CommVault Simpana 9 Backup and Recovery software with the IBM ProtecTIER deduplication solutions to get the best factoring ratio and optimal performance.

ProtecTIER can be deployed as Virtual Tape Library (VTL) or File System Interface (FSI-CIFS) to CommVault. This chapter describes CommVault with VTL, CommVault with FSI-CIFS for Windows based servers, and FSI-NFS for UNIX clients.

Attention: For best practices and configuration of CommVault in your ProtecTIER FSI environment, see 19.3, “CommVault FSI” on page 301.

The other versions of CommVault Hitachi Data Protection Suite 7 and 8 are also fully supported by the ProtecTIER product.

This chapter describes the following topics:

- CommVault in a VTL environment
  - CommVault introduction and components
  - CommVault configuration and guidelines in a ProtecTIER environment
  - Alternative data paths
- CommVault in an FSI environment

19.1 CommVault introduction

CommVault Simpana 9 is the latest version of the CommVault enterprise backup and recovery software, which consists of fully integrated modules for backup and recovery, archiving, replication, search, and eDiscovery, all managed from a single user interface.[1]

The key capabilities and benefits of CommVault include the following ones:

Common Technology Platform Incorporates full compatibility across disk and tape products.

Virtual Server Protection Offers the ability to protect virtual machines with a block-based backup approach.

Migration Tools Provides the capability to migrate from different backup software, such a Symantec NetBackup.

Central Management Helps operate the Simpana backup and recovery software from a central management interface.

Gen3 Deduplication Incorporates an integrated and embedded hash-based deduplication solution.

Important: Never enable CommVault integrated deduplication if the IBM ProtecTIER solution is in place. Doing so severely degrades the benefit of HyperFactor deduplication, which offers greater space savings than CommVault hash-based techniques.

Global Reporting Allows you to quickly identify the status of your data protection environment, including the backup results, storage occupancy, and more.

Capacity Licensing Lets you pay for the amount of protected data independently of the complexity of the backup environment.

The CommVault Simpana 9 Backup and Recovery software is a robust suite of data management capabilities that are built on a unified platform. By using it, users simplify their data protection operations.

19.1.1 CommVault components

This section introduces the key components of CommVault Simpana Backup and Recovery software.

CommCell

The CommCell feature provides a set of storage management tools that you can use to move and manage your critical data. You can use these tools to store and retrieve data that is associated with computer systems in your enterprise. The system consists of integrated software modules that can be grouped in a CommCell configuration.

Each CommCell configuration consists of the following key components:

- One or more of the following *Client Agents*:
  - DataAgents that perform backup and recovery operation
  - Archive Management Agents

[1] Source: http://www.commvault.com

  - Quick Recovery agents that create Quick Recovery volumes
  - ContinousDataReplicatior to perform data replication from source to destination clients
- The *Common Technology Engine* (CTE) components, consisting of:
  - One CommServe server
  - One or more MediaAgents
- *Storage Resource Manager* for analyzing and reporting of stored information
- *CommCell Console* for central management and operation of CommCell
- *Content Indexing and Search* engine for easy and fast data discovery

Common Technology Engine

The CTE consists of software modules that provide the necessary tools to manage and administer the Client Agents and also manage the storage media that are associated with the CommCell configuration.

CommServe

The CommServe server ties the CommCell components together; it is the coordinator and administrator of the CommCell components. The CommServe server communicates with all agents in the CommCell to initiate data protection, management, and recovery operations. Similarly, it communicates with MediaAgents when the media subsystem requires management. It maintains a database that contains all the information that relates to the CommCell configuration. In addition, it provides several tools to administer and manage the CommCell components.

MediaAgent

The MediaAgent transfers data between the client computers and the storage media. Each MediaAgent communicates locally or remotely to one or more storage devices, which contain the storage media. The system provides support for various storage devices, including the IBM ProtecTIER virtual tape libraries and virtual tape drives.

CommCell Console

The CommCell Console is the graphical user interface that allows you to control and manage the CommCell element. The CommCell Console can be run in two ways:

- As a stand-alone application, which can be installed directly on any computer that can communicate with the CommServe storage manager.
- As a remote web-based application using Java Web Start, which allows you to remotely access the CommCell Console by using the web browser.

Content Indexing and Search

You can use Content Indexing and Search to search and perform data discovery operations in your CommCell group. You can use this powerful component to search both online and stored data. Administrators, Compliance Officers, and users can use it to search and restore data from several applications, such as FSI, Microsoft Exchange, Microsoft SharePoint, and Lotus Notes, in the CommCell configuration. The search and restore operations can be performed by using either the CommCell Console or the web-based Search Console, which are controlled by a security module.

19.2 CommVault with ProtecTIER VTL

CommVault supports the installation of IBM ProtecTIER deduplication gateways in either a VTL or a FSI-CIFS configuration. However, each of these options requires specific planning and configuration steps that are further explained in the next sections. The general concept of the ProtecTIER product and CommVault integration is presented in Figure 19-1.

Attention: For best practices and configuration of CommVault in your ProtecTIER FSI environment, see 19.3, “CommVault FSI” on page 301.

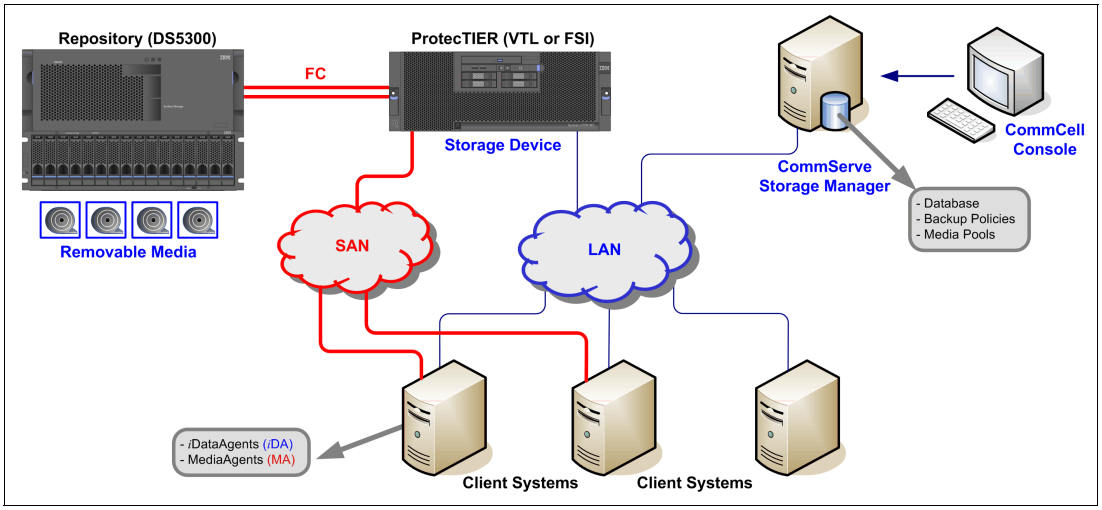

Figure 19-1 Conceptual overview of CommCell with ProtecTIER

The CommServe server is the core component in your CommCell implementation that manages all the data that is backed up from the clients. It has either ProtecTIER virtual tape drives, or file system shares that are configured as a storage repository for data deduplication. In the case of the VTL installation, with the required SAN-based backup, the ProtecTIER virtual tape drives are also zoned to the client systems and configured by using MediaAgents (MA). iDataAgents then transfer the backup data to the CommServe repository, represented by the ProtecTIER deduplication solution.

Important: When you configure VTL on CommServe or Client systems, make sure the latest version of the IBM Tape Device Driver is correctly installed. Data path and control path should be enabled, and enabling the Persistent Naming feature is highly recommended. For guidelines about persistent device name binding, Control Path Failover (CPF), and Data Path Failover (DPF), see Chapter 7, “Host attachment considerations for VTL” on page 103.

19.2.1 CommVault configuration

This section guides you through the initial configuration of the VTL TS3500 with Ultrium LTO3 virtual tape drives, which are emulated by the IBM ProtecTIER deduplication gateway. Furthermore, we provide you with the mandatory parameters and the best practices that must set to achieve the best factoring ratio from your ProtecTIER server.

Initial configuration

Initially, you must set up the physical cabling or SAN zoning, install your CommVault software, and configure your ProtecTIER server with the storage repository. After these tasks are completed, you must perform the initial configuration of the VTL in the CommServe server and on all relevant Clients' MediaAgents.

If you assigned tape drives to more than one ProtecTIER front-end Emulex port, a special procedure is needed to scan the VTL correctly because of a CommVault feature that is called *Exhaustive Detection* and its behavior. To work around Exhaustive Detection, complete the following steps for an existing, automatically detected tape library:

1. Stop all I/O operations and ensure that a backup is not running. On the CommVault Service Control Manager, ensure that all CommVault services are running.
2. Log on to the CommCell Console by using your administrator account and password, click the CommCell Control Panel icon, then click **Library and Drive Configuration** (Figure 19-2).

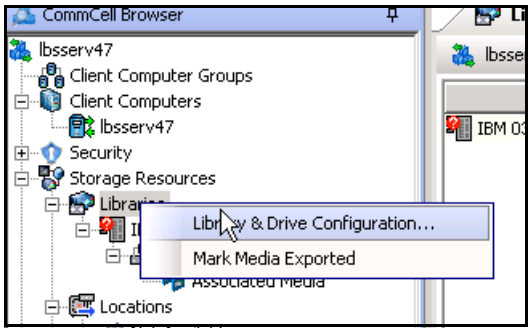

Figure 19-2 Library and drives configuration menu

3. Under the Available MediaAgents pane, double-click the server name. This action moves the server to the right pane. Click **OK** (Figure 19-3).

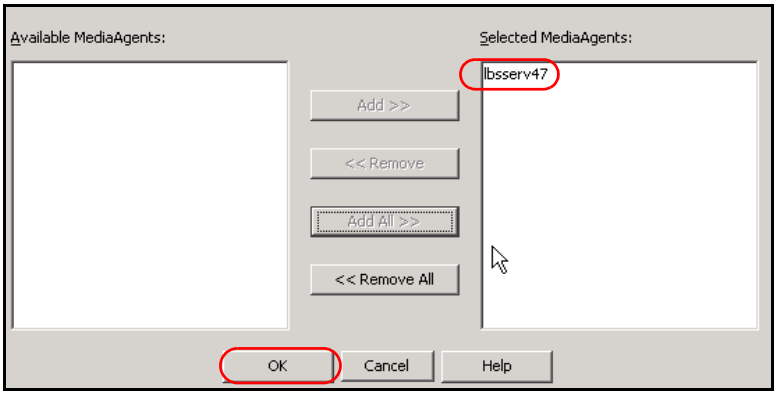

Figure 19-3 Select the MediaAgent

4. In the General tab, under Libraries, right-click the robot, click **Deconfigure**, and confirm the operation. You might be requested to confirm twice (Figure 19-4).

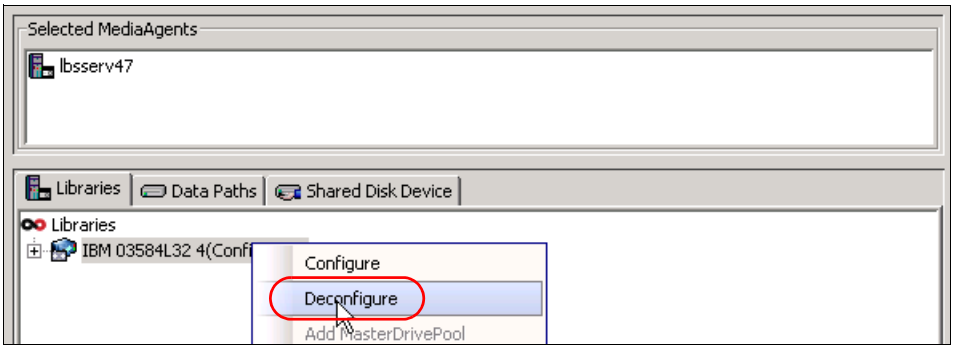

Figure 19-4 Deconfigure the existing tape library

5. As it is not configured, right-click the robot again, click **Delete**, and confirm the operation. The library should now be deleted. Click **Start** and then **Exit**.
6. Now, you must define the library and drives again with the appropriate settings. Within the toolbar icons, click the **Control Panel** and double-click **Library and Drive Configuration**. From the available MediaAgents, double-click the server (lbsserv47) to confirm your selection, and click **OK**.
7. In the General tab, right-click **Libraries** and click **Detect/Config Devices**. The Detect Library window opens; ensure all the check boxes are clear (Figure 19-5).

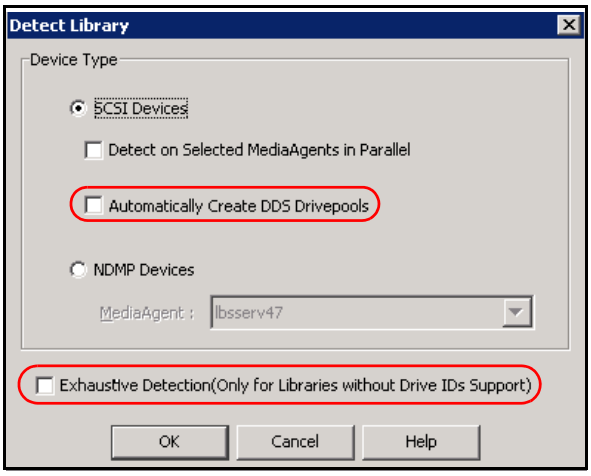

Figure 19-5 Detect new library

8. Verify the correct detection of all devices in the generated log file window (Figure 19-6) and return to the General tab.

Figure 19-6 Detection log window

9. In the General tab, open the current node and navigate to the **Libraries** overview (Figure 19-7). Right-click the unconfigured tape library tag and click **Configure** from the menu.

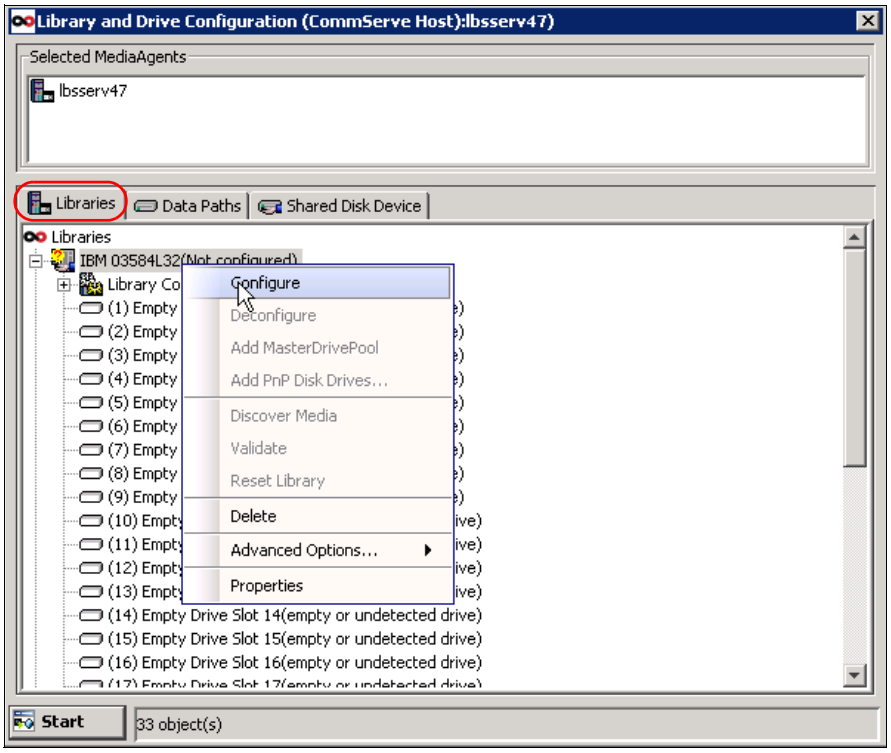

Figure 19-7 Configure the library robotics

The Configuration window opens, where you can select **Library only** or **Library and all Drives** to be configured. Click **Library only** (Figure 19-8). In most cases, the possibility of tape drives configuration is not available yet. Confirm that the library has a barcode reader (Figure 19-8).

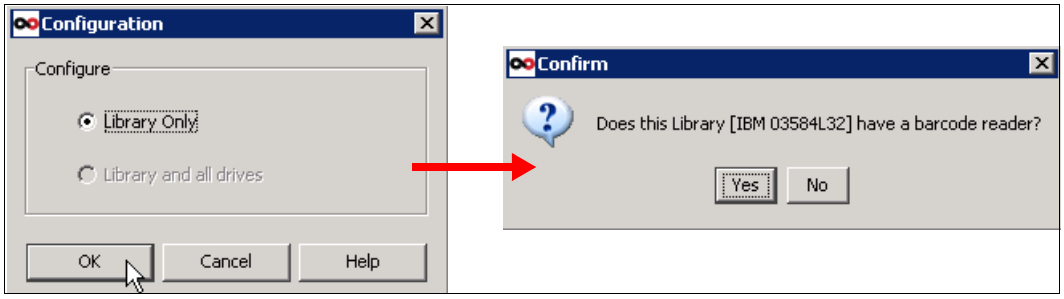

Figure 19-8 Confirm the library robotics configurations

10. Perform a full scan of the detected tape library by using the Full Scan option from the library menu (Figure 19-9).

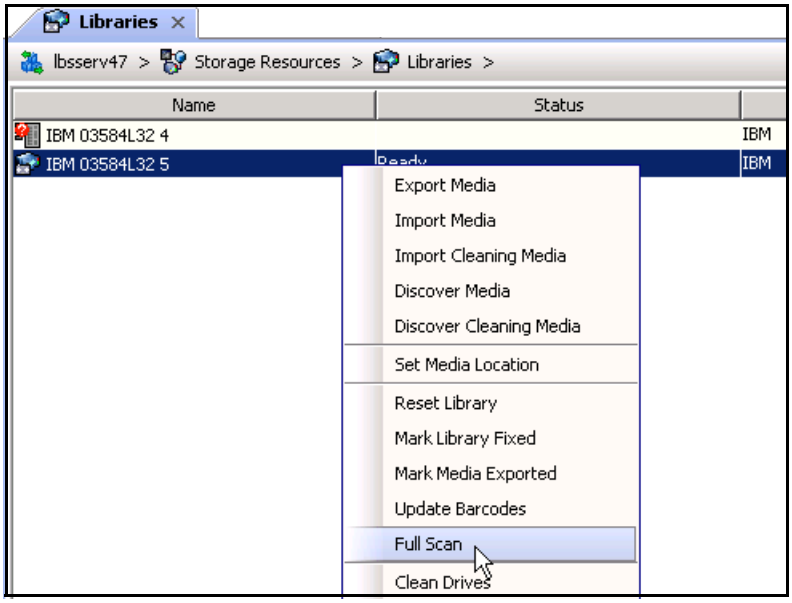

Figure 19-9 Perform a full scan of the library

11. Configure the first detected tape drive from the Libraries tab (Figure 19-10). This action allows you to configure all other tape drives in the future in one step.

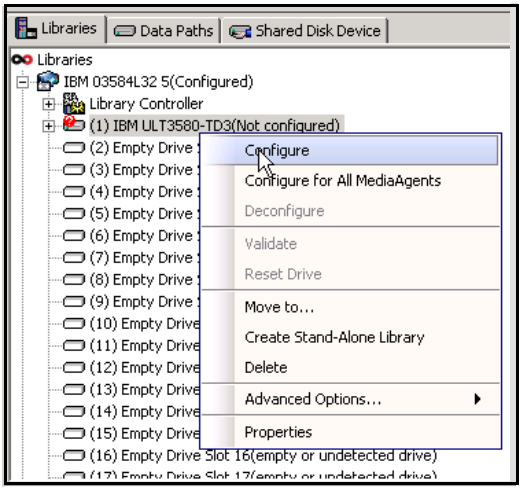

Figure 19-10 Configure first tape drive

12. Ensure that Exhaustive Detection is not used (Figure 19-11).

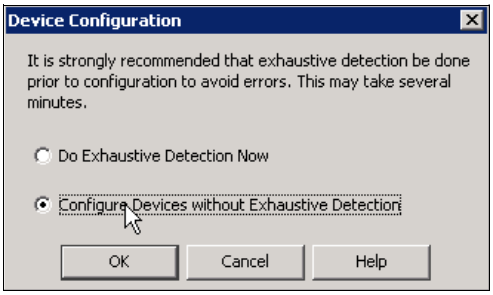

Figure 19-11 Do not use Exhaustive Detection

13. After the first tape drive is correctly configured, configure the remaining tape drives by using Exhaustive Detection. Using Exhaustive Detection saves time in configuring each remaining tape drive (Figure 19-12).

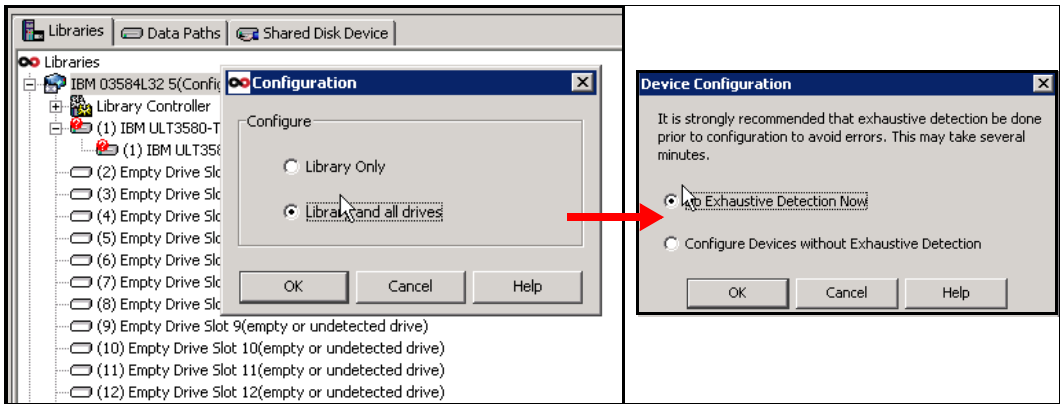

Figure 19-12 Configuration of remaining tape drives

14. When prompted, select the **Ultrium V3** type of media for a Discovery Media Option and verify that all the remaining tape drives are correctly detected and configured (Figure 19-13).

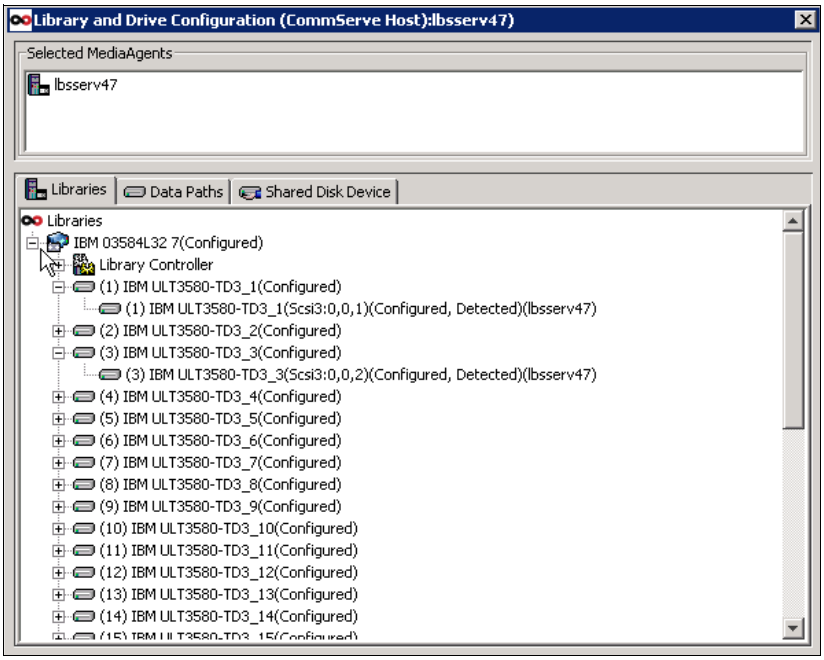

Figure 19-13 Verify the discovered tape drives

Important: If a library is present before the procedure, you must update the existing storage policies to the newly configured tape library. This action is not needed in the case of an initial configuration of the CommVault environment with the ProtecTIER server.

19.2.2 Data multiplexing

Data multiplexing is a CommVault licensed feature that allows multiple backup receivers to combine their backup data streams into one data writer. The data of multiple backup streams are then written to the single media. Figure 19-14 is an overview of how the data multiplexing feature works.

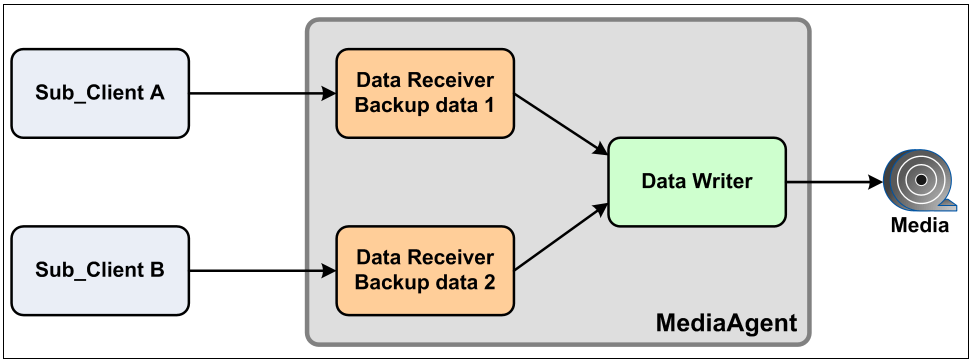

Figure 19-14 The concept of data multiplexing

Data multiplexing negatively impacts the ProtecTIER factoring ratio, so data multiplexing must be disabled. To prevent your subclients from using this feature, data multiplexing must be disabled from the Media tab of Copy Properties dialog box of the primary copy. In the CommCell browser window, navigate to **Storage Policies**, right-click the policy that you need to update, and select **Properties** (Figure 19-15).

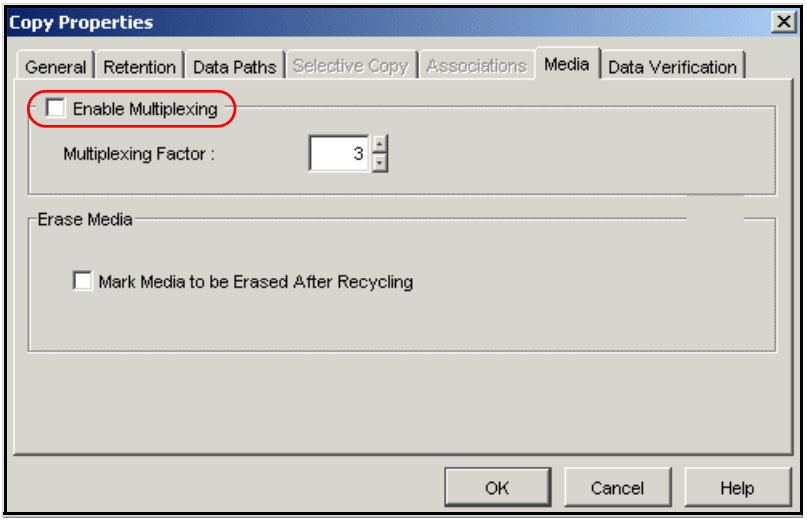

Figure 19-15 Disable data multiplexing

Multistreaming with Oracle jobs

The Oracle iDataAgent applies multiplexing rules for multiple jobs. Also, when you have multiplexing enabled for an Oracle job with multiple streams, all the streams of the job use the available drives sequentially (when one drive is full, it continues to the next). This task is enabled by setting the **JMEnableMultiplexingForOracleAgents** parameter in the CommServe database as follows:

```
Insert into GXGlobalParam values ('JMEnableMultiplexingForOracleAgents','1')
```

This parameter can be used only for Oracle jobs from the CommCell Console or when you initiate the job by using the **qoperation** backup command. In on-demand Oracle jobs, data multiplexing is enabled by default. Disable this feature by using the **QB_NO_MULTIPLEX_STREAM** option.

19.2.3 Hardware compression

You can use the IBM TS3500 Tape Libraries, which are emulated by the ProtecTIER server, to use hardware compression on virtual LTO3 tape drives. This hardware compression affects the factoring ratio, and therefore must be disabled.

In the CommCell browser window, navigate to **Storage Policies**, right-click the policy that you need to update, and select **Properties**. Navigate to the **Data Paths** tab of the Copy Properties dialog box of the primary copy (Figure 19-16). Ensure that the **Hardware Compression** check box is clear.

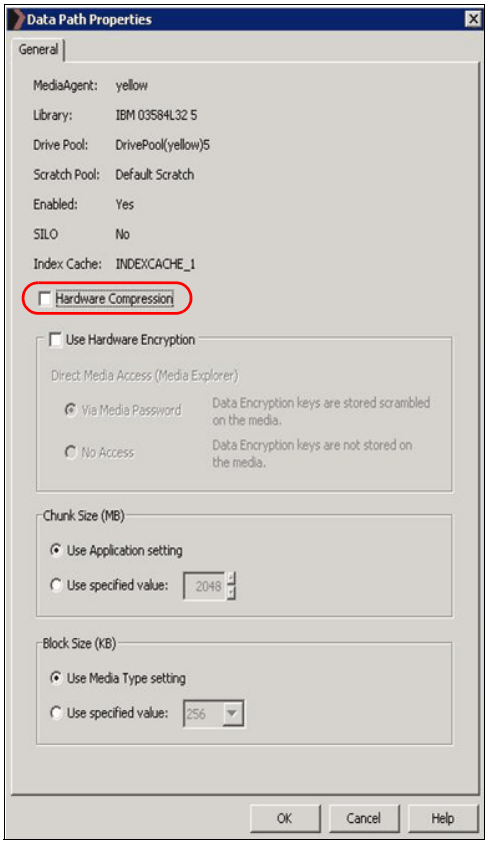

Figure 19-16 Disable hardware compression

19.2.4 Data encryption

The CommVault software supports three types of data encryption:

- Client level encryption of backup data
- Auxiliary copy level of data that is stored in CommServe media
- Hardware encryption on tape drives

None of these types of data encryption should be used with the IBM ProtecTIER solution, or the factoring ratio will be compromised or suppressed. This section briefly describes how to avoid using client encryption. Hardware encryption is not offered by the ProtecTIER emulated VTL, and an auxiliary copy of backup data is typically not stored in the ProtecTIER repository. The auxiliary copy is typically used for physical tapes, which are shipped to the offsite location or local vault regularly.

CommVault is able to use the following type of encryption algorithms (Cipher) with different block sizes:

- Blowfish: 64-bit block size, 128- or 256-bit key length
- Advanced Encryption Standard (AES): 128-bit block size, 128- or 256-bit key length
- Serpent: 128-bit block size, 128- or 256-bit key length
- Twofish: 128-bit block size, 128- or 256-bit key length
- 3-DES (Triple Data Encryption Standard): 64-bit block size, 192-bit key length

To disable client-level encryption feature, right-click the client in the CommServe Console and select its properties. In the Encryption tab, ensure that the **Encrypt Data** check box is clear. This action automatically disables the data encryption for all relevant subclients that belong to the same Client's system (iDataAgents).

19.2.5 Alternative data paths

A data path is a licensed feature that integrates the MediaAgent, Library, Drive Pool, and Scratch Pool features that are used by the storage policy copy to provide backup operations. Each storage policy copy has a unique, single data path by default. For high availability purposes, you can define alternative data paths for each storage policy copy.

The Alternate Data Paths (ADP, also known as GridStor) feature provides the following benefits:

- Automatically redirects a backup stream to an alternative data path, if one of the components in the default data path is not available.
- Alternative data paths can be used to minimize media usage by routing data protection operations from several subclients to the same storage policy and the same media, instead of creating several storage policies by using a different media for each subclient.
- Load balancing (round robin) between alternative data paths provides the mechanism to evenly distribute backup operations between available resources.

If a storage policy is created during the library configuration process, a default data path is created for the primary copy. The default data path is created by using the MediaAgent, Library, Drive Pool, and default Scratch Pool combination for drive pools that are configured within the library. If you create a storage policy, you must specify a Library, MediaAgent, Drive, and Scratch pool combination for the primary copy. In the CommCell browser window, navigate to **Storage Policies**, right-click the policy that you need to update, and select **Properties**. Additional data paths for the primary copy can be defined on the Data Paths tab of the Copy Properties dialog box (Figure 19-17).[2]

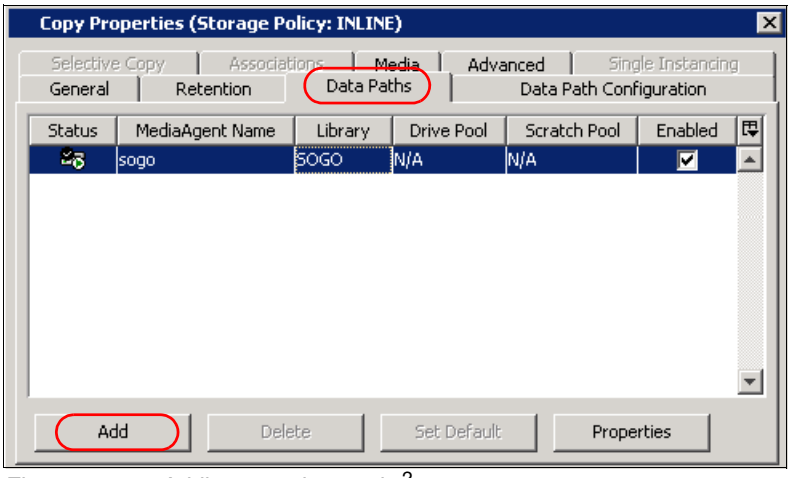

Figure 19-17 Adding new data paths[2]

[2] Source: http://www.documentation.commvault.com

19.3 CommVault FSI

This section provides steps and best practices for configuring and setting up CommVault for backup and restore. It also provides CommVault parameters and settings for best performance with ProtecTIER FSI for Windows based servers through the CIFS protocol (FSI-CIFS) and for UNIX clients through the NFS protocol (FSI-NFS).

19.3.1 Setting up backup and restore in a CIFS environment

This section provides steps and best practices for configuring and setting up CommVault for backup and restore in an FSI-CIFS environment.

Adding a disk library

To add a disk library, complete the following steps:

1. From the Control Panel, double-click **Library and Drive Configuration**.
2. Select the MediaAgents whose devices you want to detect or display from the Available MediaAgents pane.
3. Click **Add >>** to move the MediaAgents to the Selected MediaAgents pane
4. Click **OK**.
5. From the Library and Drive Configuration window, click the **Start** menu, select **Add**, and then click **Disk Library** from the menu.

Adding a shared mount path

To add a shared mount path, complete the following steps:

1. In the Add Mount Path dialog box, click **Network Path**. Complete the Connect As, Password, and Verify Password fields with the information of the CIFS share user account. Write permission is defined when you set CIFS authentication to either Workgroup or Active Directory.
2. Click **Network Path** → **Folder** and enter \\FSI_IP\CIFS_name. In our example, we use \\9.11.109.130\bpgc.

Note: When you specify the network path for the mount point, use host name aliases instead of static IP addresses. By specifying the host name (that is, from `/etc/hosts` or `c:\windows\system32\System32\drivers\etc`), you have flexibility if the IP address of the ProtecTIER FSI changes or you want to switch from the primary copy to a copy of the file system on a remote ProtecTIER system.

Defining the BackupSet and associating it with the storage policy

To define a BackupSet and associate it with a storage policy, complete the following steps:

1. Collapse the **Client Computers** menu in the CommCell Browser
2. Choose the BackupSet you want to use, highlight the subclient, and click **Properties**.
   a. Edit the content, and add the paths that you want to back up.
   b. Select the Storage Policy name, which is automatically created you add a disk library, from the **Storage Policy** drop-down list in the **Data Storage Policy** tab.

Performing a backup

To perform a backup, complete the following steps:

1. Choose the BackupSet, right-click the subclient that you want to back up, and choose **Backup**.
2. Click **Full and Immediate** backup, and then click **OK**.

Performing a restore

To perform a restore, complete the following steps:

1. From the CommCell Browser, navigate to **Client Computers** → **Client** → **File System** → **defaultBackupSet**.
2. Right-click the default subclient and then click **Browse Backup Data**.
3. Expand **defaultBackupSet** and right-click the folder that you want to restore.
4. Click **Restore Current Selected** or **Restore All Selected**.

If you want to restore to the same folder or another CIFS share folder, complete the User Name, Password, and Confirm Password fields with the information of the CIFS share user account with write permission. Write permission is defined when you set CIFS authentication to either Workgroup or Active Directory.

Complete the following steps:

1. From the Restore Options for All Selected Items dialog box, click **Advanced**.
2. The Advanced Restore Options window opens. Click **Advanced Restore Options**, click the **General** tab, and click the **Impersonate User** option.
3. In the User Name and Password boxes, enter a user name and password that has access privileges. In the Confirm Password box, type the password again. Click **OK**.

19.3.2 Parameters for best performance with ProtecTIER FSI-CIFS

This section provides some recommended parameters for the best performance with ProtecTIER FSI.

Clearing hardware compression/using hardware encryption

To disable hardware encryption, complete the following steps:

1. Right-click the storage policy copy and click **Properties**.
2. Click the **Data Paths** tab.
3. Click the data path for which you want to change the hardware compression and then click **Properties**.

4. From the Data Path Properties dialog box, clear the **Hardware Compression** check box to disable hardware compression. Clear the **Use Hardware Encryption** check box to disable hardware encryption (Figure 19-18).

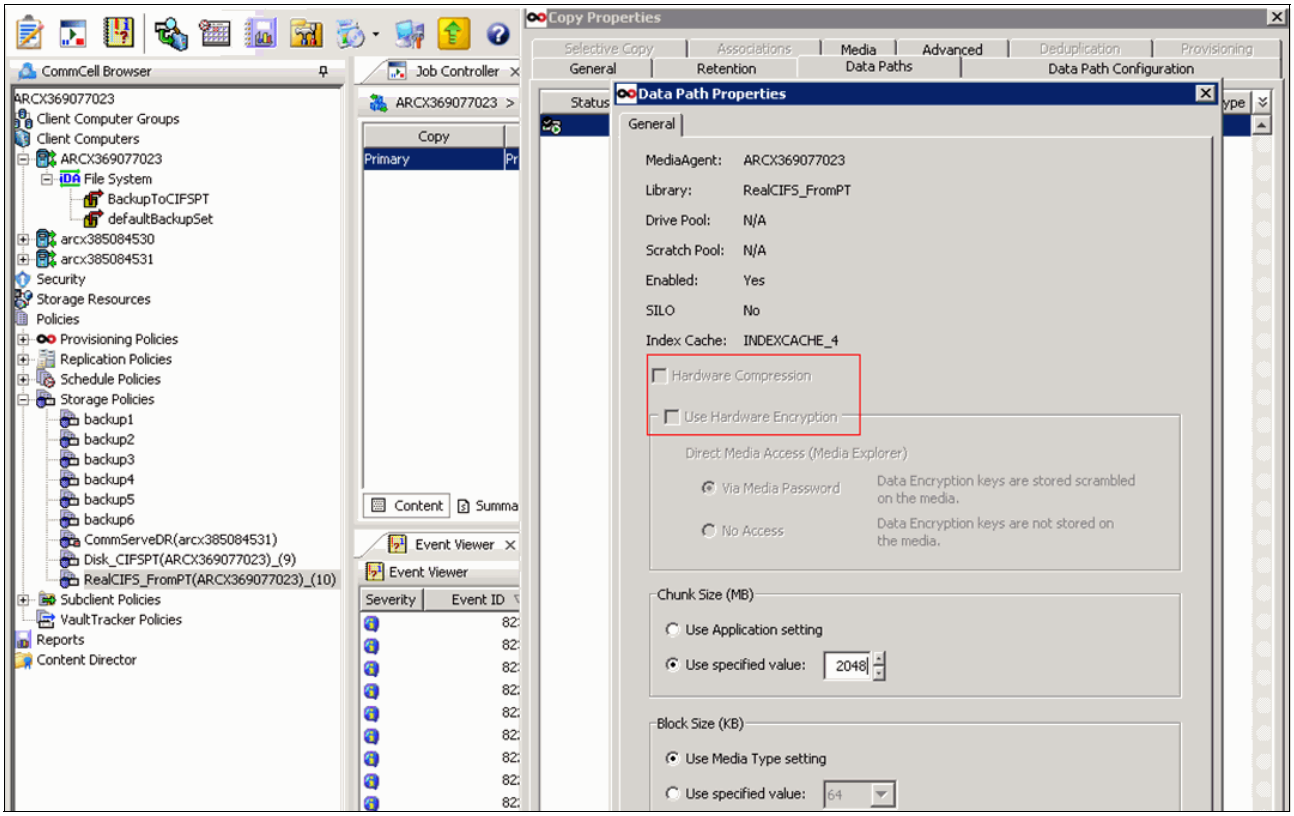

Figure 19-18 Clearing / enabling hardware compression

5. Click **OK** to save your changes.

Disabling data multiplexing

To disable data multiplexing, complete the following steps:

1. In the right pane of the CommCell Browser, right-click the storage policy copy for which you want to configure multiplexing, and select **Properties**.

2. From the Copy Properties (Media) window, clear the **Enable Multiplexing** check box (Figure 19-19).

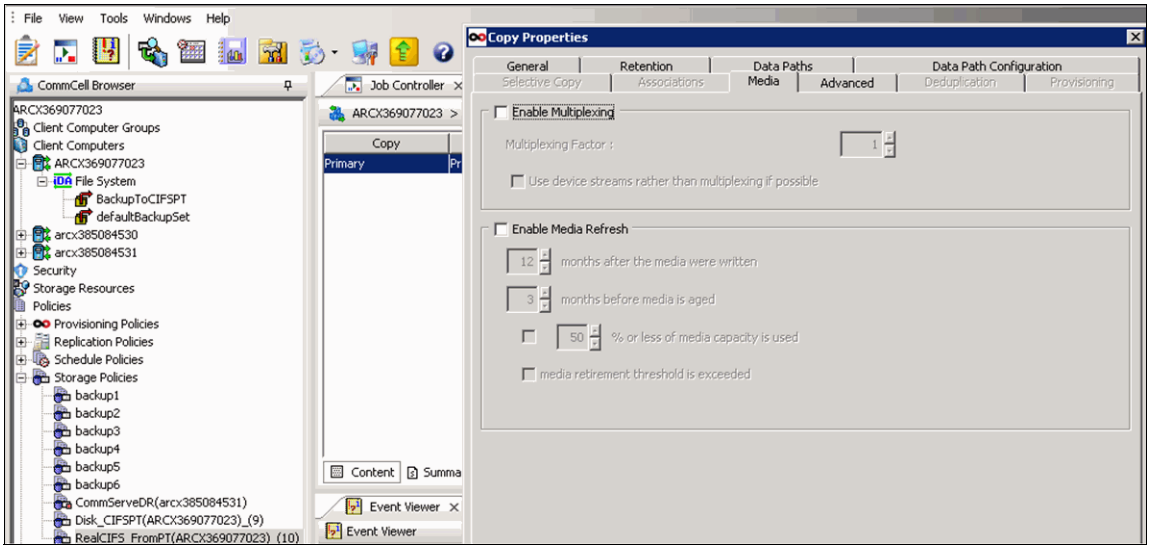

Figure 19-19 Clear Enable Multiplexing

3. Click **OK** to save your changes.

Disabling software compression

To disable software compression, complete the following steps:

1. From the CommCell Browser, right-click the subclient for which you want to disable software compression and then click **Properties**.
2. Click the **Storage Device** tab and under the Data Storage Policy tab, select the storage policy from the Storage Policy list.

3. Click the **Data Transfer Option** tab and select **Off** for the Software Compression option for this subclient (Figure 19-20).

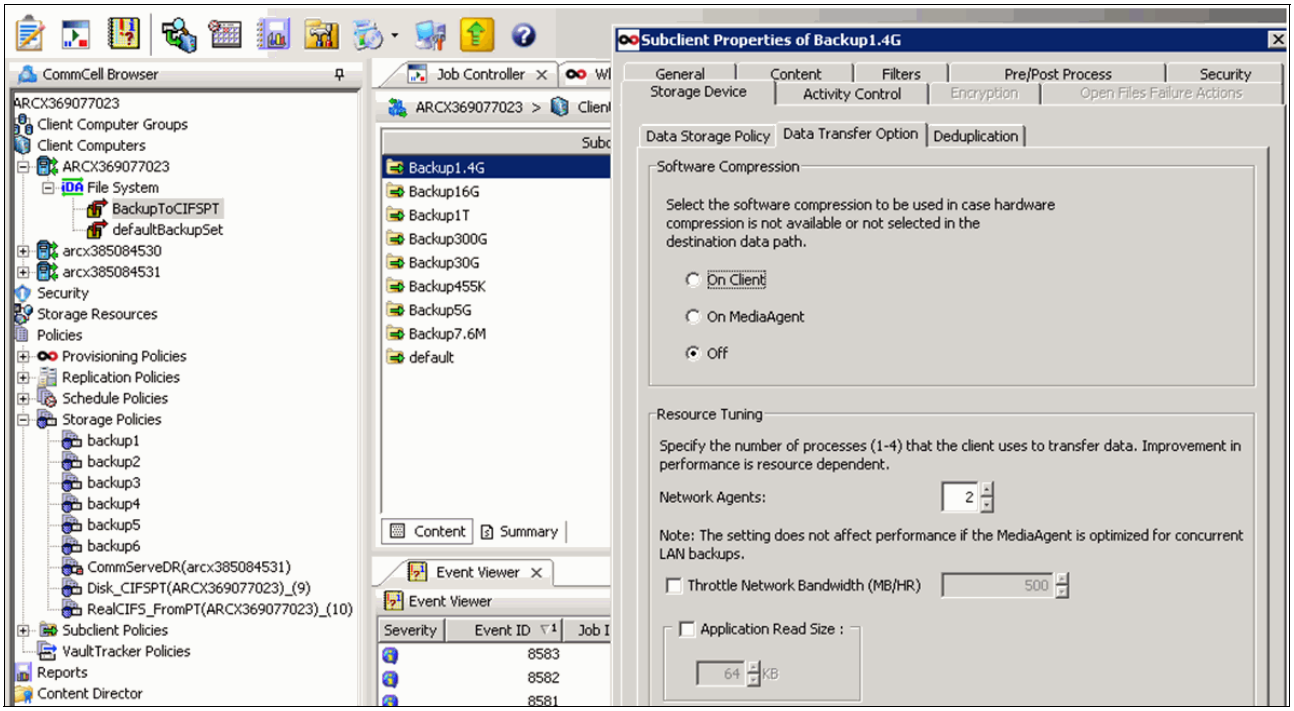

Figure 19-20 Set software compression off

4. Click **OK** to save your changes

Disabling deduplication

To disable deduplication, complete the following steps:

1. From the CommCell browser, right-click the subclient for which you want to disable software compression and then click **Properties**.
2. Click the **Storage Device** tab and, from the Data Storage Policy tab, select the storage policy from the **Storage Policy** list.
3. Click the **Deduplication** tab and clear the **Enable Deduplication** check box for this subclient (Figure 19-21).

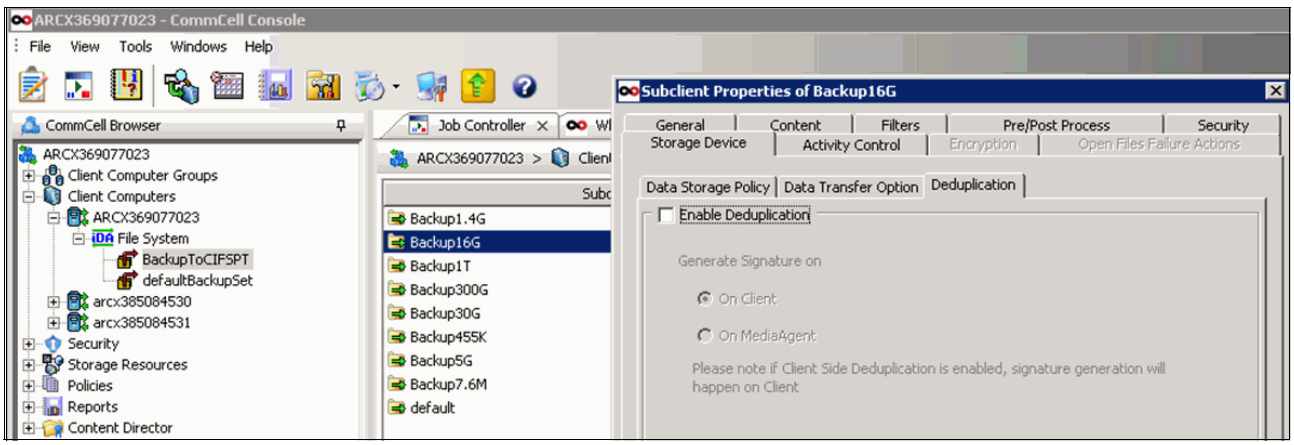

Figure 19-21 Disable deduplication

4. Click **OK** to save your changes.

Configuring multiple streams for backups and changing the maximum number of data streams

You do not have to enable multistreaming if you have multiple mount points that point to the same physical drive. For detailed steps about configuring multistreaming and about changing the maximum number of data streams, go to the following website:

http://documentation.commvault.com/commvault/release_9_0_0/books_online_1/english_us/prod_info/windows.htm?var1=http://documentation.commvault.com/commvault/release_9_0_0/books_online_1/english_us/products/windows/config_adv.htm

You can configure the automatic definition of alternative data paths if multiple MediaAgents share the ProtecTIER VTL. In the CommCell browser window, navigate to **Storage Policies**, right-click the policy that you need to update, and select **Properties**. Navigate to the **Data Path Configuration** tab of your Storage Policy **Copy Properties** window (Figure 19-22).

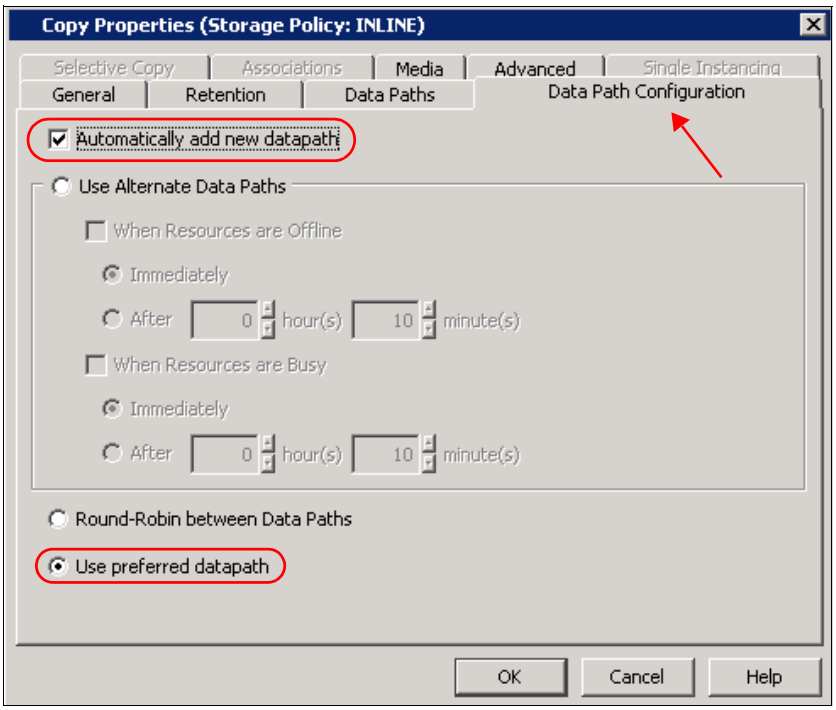

Figure 19-22 Automatically add new data paths

Always use the “Use preferred data path” option when you define new alternative data paths. Avoid using round-robin load balancing between the resources, as it has a negative impact on the factoring ratio. This negative impact occurs because the load balancing is done on tape drive pool level, not on the MediaAgent level.

For more information about configuring alternative data paths, see the CommVault official documentation at the following website:

http://documentation.commvault.com/commvault/release_9_0_0/books_online_1/

19.3.3 Setting up backup and restore in an NFS environment

This section provides steps and best practices for configuring and setting up CommVault for backup and restore in an FSI-NFS environment. It also provides CommVault parameters and settings for best performance with ProtecTIER FSI-NFS

Adding a disk library

To add a disk library, complete the following steps:

1. From the Control Panel, double-click the **Library and Drive Configuration**.
2. Select the MediaAgents whose devices you want to detect or display from the Available MediaAgents pane.
3. Click **Add >>** to move the MediaAgents to the Selected MediaAgents pane and click **OK**.
4. From the Library and Drive Configuration window, click the **Start** menu, select **Add**, and then click **Disk Library** from the menu.
5. Write an alias for the new Disk Library and check the **Automatically create storage policy for the new data paths** option.
6. Click **OK**. The Add Mount Path dialog opens for the Adding a shared mount path option.
7. In the Add Mount Path dialog, leave the **Disk Device** and **Base Folder** options as they are. Confirm that in the MediaAgent box the correct client is selected.
8. Click Local Path Folder and select the /<*mountpoint*> of the NFS export in the new window that appears.
9. Click **OK**.

Defining the BackupSet and associating it with the storage policy

To define a BackupSet and associate it with a storage policy, complete the following steps:

1. From the CommCell Browser, navigate to **Client Computers** → **Client** → **FileSystem**. Right-click **FileSystem** and select **All Tasks** → **Create new Backup Set**.
2. Input the name of the new BackupSet.
3. Select the Storage Policy name, which is automatically created when you add a disk library, from the **Storage Policy** drop-down list, and click **OK**.
4. Choose whether to create a backup schedule or not. For example, choose **Do not schedule**, and click **OK**.
5. Right-click the **BackupSet** created and select **All Tasks** → **Create new Backup Subclient**.
6. Input the name of the new subclient under the General tab of the subclient properties dialog.
7. Add the file path that you want to back up under the Content tab of the subclient properties dialog.
8. Under the **Storage Device** → **Data Storage Policy** tab of the subclient properties dialog, select the Storage Policy name, which is automatically created when you add a disk library, from the **Storage Policy** drop-down list.
9. Click **OK**.

Best practices:

1. Have a total of 16 streams that are used by all BackupSets (one or more) with ProtecTIER FSI-NFS.
2. The total number of streams per one BackupSet is the sum of all the streams started by every subclient that is defined in this BackupSet.
3. By default, each subclient uses one stream. To change this default and configure the number of streams per subclient, click the **General** tab under the subclient properties dialog, select the **Allow multiple data readers within a drive or mount point** option, and input the number of streams that you want in to the Number of Data Readers option.
4. Make sure to configure the correct number of streams per individual subclient, the correct number of subclients per individual BackupSet, and then the correct number of BackupSets to align with the recommendation of total number of streams being 16.
5. If there is a need to run more than 16 streams per BackupSet, then be aware that the maximum number of streams that are supported by each BackupSet is 16. This is a Simpana limitation.

Disabling hardware and software compression and encryption

To disable hardware compression and encryption, complete the following steps:

1. In the CommCell Browser's right pane, click **Policies** → **Storage Policies**. Right-click the storage policy that was created before and click **Properties**.
2. Click the **Data Paths** tab.
3. Click the data path for which you want to change the hardware compression and then click **Properties**.
4. From the Data Path Properties dialog box, clear the **Hardware Compression** check box to disable hardware compression. Clear the **Use Hardware Encryption** check box to disable hardware encryption.
5. Click **OK** to save your changes.

To disable software compression, complete the following steps:

1. From the CommCell Browser, right-click the subclient for which you want to disable software compression and then click **Properties**.
2. Click the **Storage Device** tab and, under the Data Storage Policy tab, select the storage policy from the Storage Policy list.
3. Click the **Data Transfer Option** tab and select **Off** for the Software Compression option for this subclient.
4. Click **OK** to save your changes.

Configuring the chunk size

To configure the chunk size, complete the following steps:

1. In the CommCell Browser's right pane, click **Policies** → **Storage Policies**. Right-click the storage policy that was created before and click **Properties**.
2. Click the **Data Paths** tab.
3. Click the data path that you want to change and then click **Properties**.
4. From the Data Path Properties dialog box, select **Use specified value** option in **Chunk Size** area and set it to 8192 MB.

Note: The Chunk size recommendation is no longer relevant when using ProtecTIER PGA V3.3.3 or later. From this code level and higher, there is no recommendation to limit the backup file size, and in fact the bigger the better

5. Click **OK** to save your changes.

Disabling data multiplexing

To disable data multiplexing, complete the following steps:

1. In the right pane of the CommCell Browser, right-click the storage policy copy for which you want to configure multiplexing, and select **Properties**.
2. From the Copy Properties window, clear the **Enable Multiplexing** check box.
3. Click **OK** to save your changes.

Disabling de-duplication

To disable deduplication, complete the following steps:

1. From the CommCell browser, right-click the subclient for which you want to disable software compression and then click **Properties**.
2. Click the **Storage Device** tab and, from the Data Storage Policy tab, select the storage policy from the Storage Policy list.
3. Click the **Deduplication** tab and clear the **Enable Deduplication** check box for this subclient.
4. Click **OK** to save your changes.

Configuring device streams

To configure device streams, complete the following steps:

1. In the CommCell Browser's right pane, select **Policies Storage Policies**. Right-click the storage policy that was created before and click **Properties**.
2. Click the **General** tab and set **Device Streams** to 16.
3. Click **OK** to save your changes.

Notes:

CommVault requires that the number of streams that are configured in the Storage Policy should be equal to or greater than the specified number of data readers that are defined per subclient. If you need more than 16 streams in your subclients, this setting should be adjusted.

You do not have to enable multistreaming if you have multiple mount points that point to the same physical drive. For detailed steps about configuring multistreaming and about changing the maximum number of data streams, contact CommVault support or review the documentation available at the CommVault website, found at:

http://documentation.commvault.com/commvault/release_9_0_0/books_online_1/english_us/prod_info/windows.htm?var1=http://documentation.commvault.com/commvault/release_9_0_0/books_online_1/english_us/products/windows/config_adv.htm

Configuring semaphores of media agent OS

For detailed steps about configuring the number of Linux Semaphores, contact CommVault support or review the documentation available at the Commvault's website, found at:

http://documentation.commvault.com/commvault/release_9_0_0/books_online_1/english_us/prod_info/linux.htm?var1=http://documentation.commvault.com/commvault/release_9_0_0/books_online_1/english_us/products/linux/config_adv.htm#Configuring_the_Kernel_Parameters

Note: This setting is only required on a Linux Media Agent system.

Performing a backup

To perform a backup, complete the following steps:

1. From the CommCell Browser, click **Client Computers** → **Client** → **FileSystem**.
2. Choose the **BackupSet**, right-click the subclient that you want to back up, and select **Backup**.
3. Click **Full and Immediate backup**, and then click **OK**.

Performing restore

To perform a restore, complete the following steps:

1. From the CommCell Browser, click **Client Computers** → **Client** → **FileSystem** → **defaultBackupSet**.
2. Right-click the default subclient, and then click **Browse Backup Data**.
3. Select **Browsing time** or **Browse the Latest Data backedup**. Click **OK**.
4. Expand **Default Backup Set** and right-click the folder that you want to restore.
5. Click **Restore Current Selected** or **Restore All Selected**.
6. In the window that opens, choose whether you want to overwrite files and the restore destination, either to the same folder or a different destination path.
7. Click **OK**.

19.3.4 Parameters for best performance with ProtecTIER FSI-NFS

Table 19-1 shows all the required and recommended settings for ProtecTIER best practices with CommVault.

Table 19-1 Recommended settings for ProtecTIER with CommVault

Component	Parameter	Value
Data path properties	Hardware Compression	Disabled (Not checked)
	Use Hardware Compression	Disabled (Not checked)
	Chunk Size use specified value	Enabled (Checked) 8192 MB [a] (8 GB)
Copy properties	Enable Multiplexing	Disabled (Not checked)
Data Transfer option	Software Compression	Off
Deduplication tab	Enable Deduplication	Disabled (Not Checked)
General tab of the Storage Policy Properties	Device Streams	16 [b] [c]
Media Agent OS Properties	Number of Linux Semaphores	[d]

a. The Chunk size recommendation is no longer relevant when using ProtecTIER PGA V3.3.3 or later. From this code level and higher there is no recommendation to limit the backup file size, and in fact the bigger the better.

b. The number of streams that are configured in the Storage Policy should be equal to or greater than the specified number of data readers.

c. You do not have to enable multistreaming if you have multiple mount points that point to the same physical drive. For detailed steps about configuring multistreaming and about changing the maximum number of data streams, contact CommVault support or review the documentation available at the CommVault website, found at http://documentation.commvault.com/commvault/release_9_0_0/books_online_1/english_us/prod_info/windows.htm?var1=http://documentation.commvault.com/commvault/release_9_0_0/books_online_1/english_us/products/windows/config_adv.htm.

d. This setting is only required for Linux Media Agent system. For detailed steps about configuring the number of Linux Semaphores, contact CommVault support or review the documentation available at the Commvault's website, found at http://documentation.commvault.com/commvault/release_9_0_0/books_online_1/english_us/prod_info/linux.htm?var1=http://documentation.commvault.com/commvault/release_9_0_0/books_online_1/english_us/products/linux/config_adv.htm#Configuring_the_Kernel_Parameters.

Part 4

Application considerations

This part describes settings and parameters that are modified for optimum deduplication ratios and performance when you work with specific data types. The applications that we focus on are RMAN Oracle, Lotus Domino, Microsoft Exchange, Microsoft SQL Server, DB2, and VMware.

This part describes the following topic:

- Application considerations and data types

20

Application considerations and data types

This chapter described guidelines for the settings and parameters that are modified in your applications, and specific data types for optimal performance and deduplication factoring ratios.

This chapter describes the following topics:

- Lotus Domino
- Microsoft Exchange
- Microsoft SQL Server
- DB2
- Oracle
- SAP
- VMware

Readers of the relevant sections should be familiar with the backup and restore concept of the managed application or data type. Therefore, we do not give detailed steps about how to configure backup applications.

20.1 Lotus Domino

This section describes the settings and parameters that should be modified in Lotus Domino environments to enable the optimal factoring for the ProtecTIER product.

20.1.1 Common server

Lotus Domino employs a common email or application server for several members of the company or network. Clients usually run backup policies from the common server (the Domino server) that stores all the data on the physical or virtual tape.

Domino servers in enterprise or secured environments are configured in *application clusters*. In contrast, the server-based clusters and the shared storage resources that are assigned to the application are available on both (or more) cluster nodes simultaneously. Access to the them is fully controlled by the clustered applications, in our case, Domino servers.

In Domino mail environments, there is typically a configuration of active-active clusters, where each Domino server is always active only on a dedicated node of the dual-node cluster, never fails over, and both Domino applications control the same storage resources. However, only the dedicated portion of the databases (set of mail files) is served at one time by the single node. The common understanding of the application failover to the standby cluster node does not apply in Domino environments. If there is a node failure, the application that is running on a live node takes full management control over all portions (all mail files) of storage resources instantaneously.

From the Lotus Domino functional perspective, there are the following categories of Domino server installations:

- An email server that supports Lotus Notes, IMAP, POP3, SMTP, and WebMail access (IBM iNotes®).
- An application server where the Lotus Notes client provides the application run time.
- A database server that offers Notes Storage Facility.
- A web server that allows Lotus Notes clients to access the data through a web browser.
- A directory server for authentication services (hub/gateway).
- Instant messaging and web conferencing, also known as IBM Sametime®.

This section focuses on email, application, and database servers, which usually hold the most amount of data in Domino server repositories. The other listed features are highly transactional services with small amounts of data, and are therefore not optimal candidates for ProtecTIER deduplication.

20.1.2 Legacy backup and disk space usage

Running the legacy backup commands and using the general recommended methods cause a low factoring ratio of the data even if the change is low. These actions reduce the benefit of using the ProtecTIER solution and disk space usage.

The root cause is the inherent compaction of the database, which reshuffles Notes Storage Format (NSF) files inside the database. Although this function reduces space usage from the perspective of Lotus Domino, it also changes the layout and data pattern of every NSF.

The next time that the ProtecTIER server receives blocks from these databases, they all look unique, so the factoring ratio ends up being low. Ratios of 1:2 - 1:3 are possible in environments with Lotus Domino Version 7. However, running compaction that is based on the Domino **DELETE** operation is a best practice for Lotus Domino, so disabling it is not a solution.

Simply stated, the compaction saves the primary expensive storage on Domino server and increases the performance of mailbox operation, especially in clustered environments, so there is no single client that wants to disable it.

An additional factor to consider regarding deduplication efficiency are the methods that are used to store email attachments. Working documents are compressed and archived by one of the widely available compression tools, and are converted to the file formats zip, rar, tar, 7z, and others, which do not factor optimally. The same is true for media files in email attachments, such as compressed pictures (jpg and gif), movies (mpg, mov, and avi), or music (mp3).

20.1.3 Domino attachments and object service

Deduplication can be improved by using a new feature in the Domino environment that is called *Domino Attachment and Object Service* (DAOS). DAOS removes all the email or application attachments from the Notes Storage Format (NSF) files and stores them separately in the server file system in a *single occurrence* as a Notes Large Object (NLO). Figure 20-1 and Figure 20-2 on page 318 show examples of a dedicated storage repository for NSF and NLO objects. This feature has been available since Domino Version 8.5.

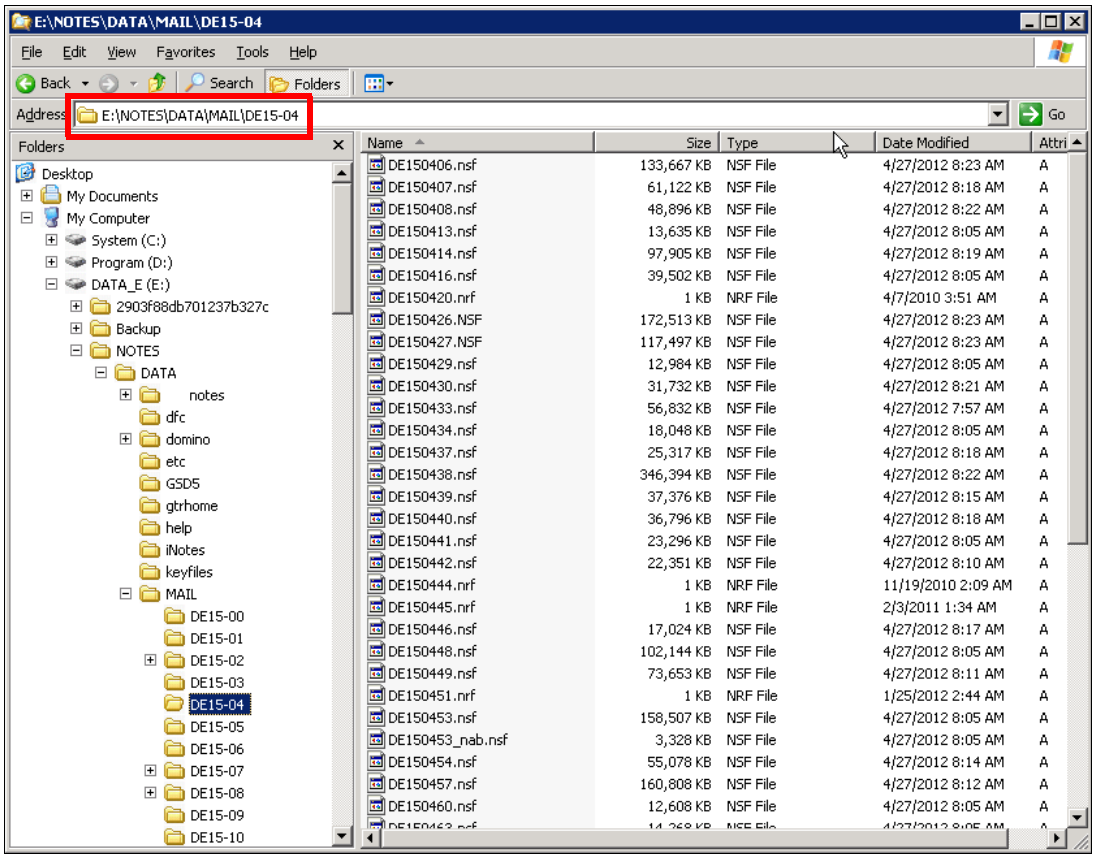

Figure 20-1 The location of mail files in Domino storage

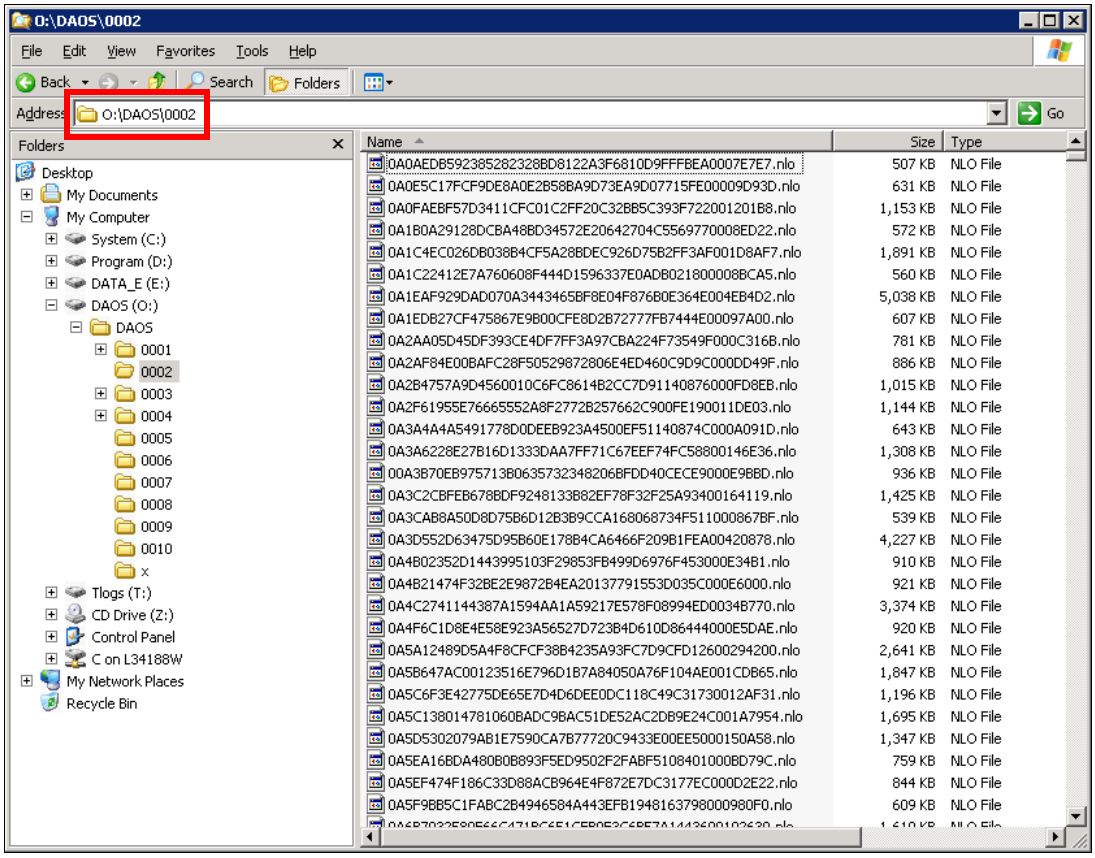

Figure 20-2 The folder structure of DAOS attachments

DAOS divides the database objects into two categories of items:

- .NSF files (database items)
- .NLO files (attachment items)

The NLO holds only one instance of each attachment and the multiple NSF files that contain relevant metadata links, and reuses it. DAOS reduces the impact of using the **DELETE** option because the DAOS layout does not hold each attachment multiple times. This arrangement mitigates the compaction effect and the NLO change is marginal.

Backing up the NLO files in the DAOS repository can be done either while the Domino server is down, or when it is up and running. The backup does not require the usage of any Domino API-based utilities. After the NLO files are initially written, Domino never modifies their contents, so the backup mechanism does not work around file-write activity. NLO files can be backed up as with any other generic files in the file system. Only the NLO files that are complete and not in the process of being written to or renamed must be backed up. Any files that are busy can be skipped until the next backup job runs. Most backup applications automatically skip files that they cannot read because of other activity.

Important: If Domino is running during a backup process (online backup), it is important to first back up all NSF files before you proceed with the NLO backup because the metadata references in the NSFs are related to the newly detached NLO files.

In typical Domino environments, there are dedicated and independent processes that are used to back up NSF and NLO files. The NLO file is just a flat file on the disk, while the NSF file is considered to be in database format, which means different tools are used to back up each of the files (for example, Tivoli Data Protection for Mail, in contrast with the Tivoli Storage Manager Backup/Archive client). It is obvious that operating system flat file backups are not retained in the backup server for the same period as online backups of Domino NSF files.

Important: Ensure that the retention period of NLO file backups is at least the same or longer than the longest retention period used for online backups of NSF files (monthly, quarterly, or yearly backups).

Domino incorporates the feature to keep NLO files on disk for a period after the link to them becomes invalid and the NLO files are not needed anymore. If this DAOS retention period is longer than the backup retention of NSF files in a backup server, the previous statement does not apply. This case is the only exception where backup retention on NLO files does not play a role. A backup copy is never used for the restoration of mail files with detached data, and the relevant NLO file is still retained on the disk by DAOS. However, the minimal backup retention is still needed to protect data against a disaster or file system corruption.

The disk footprint savings with DAOS apply to the backup processing as well. The NLO files represent the static data that used to be in the NSF, and was backed up every cycle even though it had not changed. In a typical mail environment, a large reduction in the NSF footprint, plus a small amount of NLO data, translates almost directly into a reduction in the backup footprint. Not only is the duplicate data eliminated, the mail file data is separated into static and dynamic components. By applying an incremental backup regimen to the static NLO data, only the NLO files that were created since the last backup cycle need to be processed. Those files typically represent a small amount of data compared to the entire set of NLO files.

In the incremental backup case, duplicate NLOs are not backed up again. Thus, the space savings from DAOS are directly proportional to the number of duplicate NLOs seen in the environment, and the backup time savings is the product of the space that is saved and the backup throughput.

The ProtectTIER server greatly benefits from this whole behavior. We have seen a factoring ratio of 3 - 5 times higher than before the DAOS is enabled.

20.1.4 Applying the DAOS solution

Run the DAOS Estimator Tool by running the **`DAOSest`** command in the Domino server console to assess the benefit of using DAOS. Perform this action outside of your office hours, as the estimation procedure impacts the server performance, especially mail servers with hundreds to thousands of users.

The histogram result of the estimation job is shown in Example 20-1, where the top line gives the number of files with the specific size of the detachment threshold (bottom line) and what this result represents in percentage of all attachments (middle line).

Example 20-1 Relative and percentage numbers of attachments of different size

```
================================================================================
|66362 |19744 |33629 |35311 |161416|90946 |18550 | 3458 | 426  |  22  |  0   |
================================================================================
| 0.0% | 0.1% | 0.3% | 0.6% | 7.5% |20.4% |31.5% |25.1% |11.4% | 3.2% | 0.0% |
================================================================================
|  4k  | 8k   | 16k  | 32k  | 64k  | 1MB  | 5MB  | 20MB | 100MB| 1GB  | >1GB |
================================================================================
```

The summary of the estimation process is also shown in different form in Example 20-2 on page 322.

Before anything is done with DAOS, there are some prerequisites that must be addressed. These prerequisites might not all apply in your situation, but it is important to verify them to ensure that any changes that are needed can be accommodated. These items are all requirements for enabling DAOS, and are not optional.

Consultation: Before you implement DAOS, consult with your Domino representative.

Here are the prerequisites:

- Disable SCOS Shared mail.

  The Single Copy Object Store (SCOS) is an older approach to attachment consolidation. This feature is not compatible with DAOS and must be disabled before you enable DAOS.

- Disable NSFDB2.

  The NSFDB2 is a feature that you can use to store NSF data in DB2 running either on the same or a different server. This feature is also not compatible with DAOS and must be disabled on every NSF application that participates in DAOS.

- Upgrade Domino server.

  Although DAOS was introduced in Domino V8.5.0, many important stability and performance improvements were made in subsequent releases. Hence, all new DAOS deployments should use Domino 8.5.3 or later.

- Enable transaction logging.

  The DAOS depends on transaction logging for correct operation. Because DAOS must update several locations simultaneously, it is important that all those updates succeed or fail (and are later rolled back) as a unit.

- Adjust backup/restore processes.

  You must have reliable backup and restore procedures in a production environment to avoid the possibility of data loss. DAOS adds some complexity to the backup and restore process, so you must have a well-established backup and restore foundation for DAOS. Transaction logging introduces some additional features that provide even better recovery options.

- Upgrade `Names.nsf` design.

  The design of the `Names.nsf` file was changed to accommodate DAOS, and the Server document has a new tab that covers the DAOS settings. `Names.nsf` must use the new *pubnames.ntf* template on all Domino servers that are enabled for DAOS.

20.1.5 ProtecTIER considerations

In contrast to the general recommendations for DAOS deployment on Domino servers, here we summarize the best practices or limitations that apply when the ProtecTIER deduplication solution for backup and recovery is in place:

- Disable NLO compression.

  The mail attachments that are represented by NLO files use a certain amount of disk space on the Domino server. Domino administrators tend to enable one of the available compression techniques on the attachments in NFS files. If no attachment compression is enabled on the NSF files, or if Huffman compression is being used, then enabling LZ1 compression can save a significant amount of disk space. Run **`compact -ZU`** to enable LZ1 compression.

  Tip: Avoid using the **`-ZU`** flag during compaction to achieve the best factoring ratio.

- Disable design and data document compression.

  Another Domino space-saving feature is design and data document compression. Enabling these compression forms can also save disk space, but they have a negative impact on deduplication results in the ProtecTIER server. The savings from these features are independent from DAOS and do not achieve the level of savings that you can make with the ProtecTIER solution.

  Tip: Do not compress Design and Data documents.

- Consider the compacting frequency.

  Compacting less frequently is not popular with Domino administrators. However, it does *not* have a significant impact on the performance of a Domino server, mailboxes, or the storage capacity that is used by the Domino server. Complicating factors are backups with retention periods and the database file unique identifier, also called Database Instance Identifier (DBIID). When the DBIID is changed by running a compact job (which is the default action, unless the **`-b`** parameter is not specified), a Domino server always considers this database as eligible for full backup, regardless whether the next backup job is scheduled as an incremental only.

  With a backup schedule that consists of weekly full backups (during the weekend) and daily incremental backup of databases with a changed DBIID (during weekdays), you should perform compact jobs on a weekly basis before the full backup occurs. This setup has a positive effect on the ProtecTIER factoring ratio.

  Tip: Schedule compact jobs less frequently and ensure that they always complete before the next full (or selective) backup of NSF databases. Incremental backup does not back up Domino NSF files, unless the DBIID has changed.

- Compact only selected databases.

  Not all Domino databases must be compacted regularly. If the percentage of *white space* (unused space) in the database is, for example, less than 10% of the mailbox size, consider excluding this database from the compaction. The space savings of such a compact job is negligible, but your factoring ratio decreases. Use the **`-S 10`** option to direct the compact task to databases only with 10% or more of the white space. The database DBIID still changes, unless the **`-b`** option is not used.

Tip: Do not compact databases that use storage space efficiently.

- Disable encryption of attachments.

  When access to server resources is restricted to responsible personnel only and there is minimal or no risk of data exposure, Domino administrators should disable encryption on detached email attachments (NLO files). The enabled encryption has a negative impact on the ProtecTIER factoring ratio, as each block of data that is sent to the ProtecTIER server behaves as a unique block of data. Although the encryption is enabled by default, you can disable it by adding the following parameter to the `Notes.ini` file:

  ```
  DAOS_ENCRYPT_NLO=0
  ```

  The setting cannot be changed retroactively, and the only way to remove encryption from an existing DAOS installation is to completely disable DAOS.

Encryption: The encryption has a negative impact on factoring ratios. Disable it if possible.

- Define the appropriate thresholds of the DAOS process.

  If the attachment is larger than the minimum participation size, it is stored in the DAOS repository. If it is smaller, it is still stored in the NSF, as it would be without the DAOS feature enabled.

  Choosing a size that is too large results in too few attachments being stored in DAOS (low yield), which reduces the savings that DAOS can offer and the ProtecTIER product can benefit from. Conversely, choosing too small of a size can result in a high yield, resulting in an unmanageable number of files in the DAOS repository.

  The statistics in Example 20-2 show the DAOS minimum size versus the number of NLOs and disk space that is required.

Example 20-2 DAOS minimum size versus the number of NLOs and disk space

```
0.0 KB will result in 429864 .nlo files using  180.7 GB
4.0 KB will result in 363502 .nlo files using  136.6 GB
8.0 KB will result in 343758 .nlo files using  130.5 GB
16.0 KB will result in 310129 .nlo files using  128.2 GB
32.0 KB will result in 274818 .nlo files using  119.4 GB
64.0 KB will result in 113402 .nlo files using  110.4 GB
1.0 MB will result in  22456 .nlo files using   85.8 GB
5.0 MB will result in   3906 .nlo files using   47.9 GB
20.0 MB will result in    448 .nlo files using   17.6 GB
100.0 MB will result in     22 .nlo files using    3.8 GB
```

Look for a value that yields about 80 - 90% of the theoretical maximum of the DAOS repository size. Although that value might sound low, it is generally the best trade-off between the DAOS benefits and the resulting number of files.

Hint: Determine the appropriate DAOS size when the attachment is offloaded from the NSF to the NLO file.

20.1.6 Preparing Domino databases for DAOS

The task of preparing Domino databases for DAOS should be performed only once. To accomplish this task, complete the following steps:

1. Depending on your operating system, choose the most appropriate procedure:
   a. If the Domino server is running Windows, click **Start** → **Programs** → **Lotus Applications** → **Lotus Domino Server**.
   b. If the Domino server is running UNIX, enter the following command at the command-line interface (CLI):

      ```
      /opt/lotus/bin/server
      ```
2. Double-click **`nlnotes.exe`**, and go to the workspace window.
3. Browse for the `names.nsf` file. The location is usually **E:\Lotus** → **Domino** → **Data**.
4. Click **Configuration** → **Servers** → **All Server Documents**. Then, open the document that is related to your Domino server.
5. Double-click the page to change it to edit mode, and select the **Transactional Logging** tab.
6. Set the following parameters:
   – Log path: logdir.
   – Logging style: Circular.
   – Maximum log space: 512M.
7. Save your parameters and close the window.
8. Shut down the Domino server by entering the following command at the CLI:

   ```
   exit <password>
   ```
9. Add the following line to the `notes.ini` file:

   ```
   CREATE_R85_DATABASE=1
   ```
10. Start the Domino server again and use the password.

> **Starting time:** For the initial startup sequence after you make these changes, it might take several minutes for the start sequence to run.

11. Complete steps 5 - 7 to edit the server document again. Open the **DAOS** tab. You might need to scroll to the right to see the tab.
12. Update the following parameters:
   – Store Attachments in DAOS: ENABLED.
   – Minimum size: 4096.
   – DAOS base path: daos.
   – Defer deletion: 30 days.
13. Restart the Domino Server by entering the following command at the CLI:

   ```
   restart server [password]
   ```

In the next compaction, you see the DAOS directory that is created in the Domino data directory. Within that directory, you see the following entry:

```
0001/<really_long_name>.nlo
```

20.2 Microsoft Exchange

This section describes the recommended settings for the Microsoft Exchange (Exchange) environment to improve the backup throughput and the factoring ratio of the ProtecTIER server. The examples that are used in this section are based on IBM Tivoli Storage Manager, but the recommended settings apply to most enterprise backup applications. Some of the settings might not be available in other backup applications. Contact the backup application provider for additional information.

20.2.1 Defragmentation

Defragmentation is commonly used in the Microsoft Exchange environment to recover the disk efficiency of fragmented disks. The defragmentation process rearranges the data that is stored on the disk and creates continuous storage space. There are two types of defragmentation processes: online defragmentation and offline defragmentation.

Online defragmentation

The online defragmentation process removes objects that are no longer being used while Exchange databases remain online. Before Microsoft Exchange 2010, the online defragmentation ran as part of daily Mailbox database maintenance, although this Mailbox database maintenance could be scheduled to run at different times.

In Exchange 2010, online defragmentation is separated from the Mailbox database maintenance process and it runs continuously in the background. Additional details about database defragmentation are available in the topic “New Exchange Core Store Functionality” at the following website:

http://technet.microsoft.com/en-us/library/bb125040.aspx#NewESE

Offline defragmentation

Offline defragmentation is a manual process that creates a database file and copies database records without the white space from the original database file to the newly created database file. When the defragmentation process is complete, the original database is removed and the new database file is renamed as the original.

Offline defragmentation is not part of regular Mailbox database maintenance. It can be done only when the Mailbox database is in an offline state, and this action requires much storage space, as both the original database file and newly created database file must coexist on the disk during the defragmentation process.

20.2.2 Recommendations for Microsoft Exchange

Here are the recommended processes and settings of the backup applications to optimize the performance and factoring ratio of the ProtecTIER server:

- Perform a daily full backup instead of daily incremental backups where only transaction logs are backed up.
- Create one backup job for each database (or for each storage group) without multistreaming to keep similar data blocks in the same stream.

- Create concurrent backup jobs if there is more than one database (or more than one storage group) within the Exchange servers to improve overall backup throughput. Example 20-3 shows how to create different backup jobs for different databases in Tivoli Storage Manager. Remember to increase the number of mount points for the client node if multiple databases are housed in one Exchange server.

  Example 20-3 Create multiple backup jobs for different databases with Tivoli Storage Manager

```
TDPEXCC BACKup <Storage Group 01/ Mailbox Database 01> full
TDPEXCC BACKup <Storage Group 02/ Mailbox Database 02> full
TDPEXCC BACKup <Storage Group 03/ Mailbox Database 03> full
```

- Disable compression and encryption within Exchange databases and backup applications.
- If personal archive files (.pst) are backed up, do not enable compression and encryption whenever possible.
- Allow a longer interval for an online defragmentation process, for example, reschedule the daily database maintenance to be on a weekly basis.
- Consider a LAN-free backup that allows data to be sent directly to storage devices.

20.2.3 Microsoft Exchange 2010

Because the online defragmentation is moved out from the daily maintenance process and it runs continuously in background, there is no option to disable or schedule online defragmentation. This situation impacts the factoring ratio. However, we do expect deduplication for Exchange 2010 because Single Instance Storage (SIS) is no longer supported in Exchange 2010. You can find more information about SIS in the Microsoft Exchange team blog at the following website:

http://msexchangeteam.com/archive/2010/02/22/454051.aspx

20.3 Microsoft SQL Server

This section describes the recommended settings for the Microsoft SQL environment to improve the backup throughput and factoring ratio of the ProtecTIER server. The examples that are used in this section are based on Tivoli Storage Manager, but the recommended settings apply to most enterprise backup applications. Some of the settings might not be available in other backup applications. For more information, contact the backup application provider.

20.3.1 Integrating the ProtecTIER server with Microsoft SQL Server backup

A ProtecTIER server can be integrated with traditional backup applications, or with the native SQL server backup utility to back up the Microsoft SQL server.

To back up a Microsoft SQL server with traditional backup applications, such as Tivoli Storage Manager, the ProtecTIER product can be deployed as a Virtual Tape Library (VTL), OpenStorage (OST) (with NetBackup), or the FSI to work in conjunction with the backup applications. To back up the Microsoft SQL Server with the native SQL server backup utility, the ProtecTIER product can be used as CIFS shares through FSI deployment. We describe the recommended ways to integrate the ProtecTIER server with different backup methods in the following sections.

Using native SQL Server backup

You can back up the Microsoft SQL Server with the native SQL server backup and restore utility, where data files can be backed up directly to backup media without using third-party backup applications. The backup media can be disk or tape devices. Most administrators choose to back up to disk instead of to tape devices because the native SQL server backup does not have a tape media management capability like other backup applications. The ProtecTIER product can be deployed as an FSI that provides CIFS shares to a Microsoft SQL server, and the native SQL server backup can use the CIFS share as the destination (Figure 20-3).

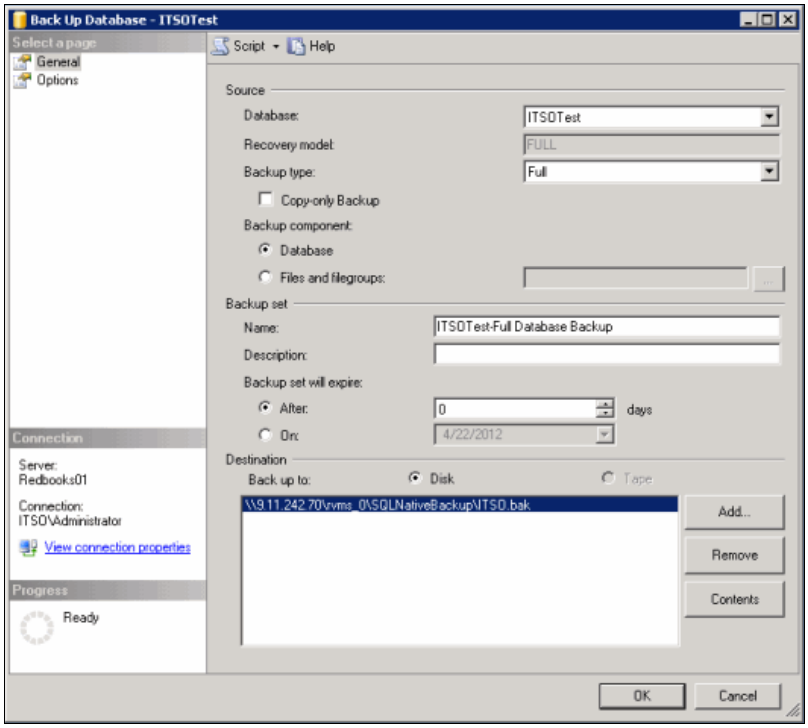

Figure 20-3 Use the ProtecTIER CIFS share as the destination of the SQL native backup

For more information about ProtecTIER FSI, see Chapter 4, “ProtecTIER File System Interface: General introduction” on page 53.

Using third-party backup applications

Most backup applications support Microsoft SQL backup through an SQL backup agent, for example, IBM Tivoli Storage Manager Data Protection for Microsoft SQL Server. Backup applications use the ProtecTIER server as tape devices, file devices, or disk storage as backup destinations during the backup server configuration. For details about how to integrate the ProtecTIER server with different types of back-end storage, see Part 2, “Back-end storage subsystems” on page 123.

20.3.2 Index defragmentation

The Microsoft SQL Server maintains indexes to track table updates, and these indexes can become fragmented over time. Heavily fragmented indexes might affect the database query performance, so Microsoft SQL Server uses a defragmentation feature to reorder the index row in continuous pages.

The defragmentation process rebuilds the indexes by compacting the index pages and reorganizing the index rows. This process results in the indexes being seen as new data blocks in backup streams, which can impact the deduplication process that identifies unique data at block level.

Index defragmentation can adversely affect database workload performance. You should perform index defragmentation only when it is necessary. Before you defragment, see the topic “Microsoft SQL Server 2000 Index Defragmentation Best Practices” at the following website:

http://technet.microsoft.com/en-us/library/cc966523.aspx

20.3.3 Recommendations for Microsoft SQL Server

Here are some suggestions for the Microsoft SQL Server to improve the backup throughput and deduplication ratio of the ProtecTIER server:

- Perform full backups whenever possible.
- When you use the ProtecTIER server as CIFS shares, always use the Universal Network Convention (UNC) path instead of the Windows mapped drive to ensure that the correct CIFS shares are used, and to avoid the Windows connection timeout issue.
- Do not schedule index defragmentation on a regular basis. Perform index defragmentation only when it is necessary.
- Disable compression and encryption within the Microsoft SQL server and backup applications.
- Limit the number of backup streams to one stream for one database backup, or one stream per physical volume, if one single large database is split into multiple physical volumes. Setting the stream to 1 gives the best factoring ratio, but it might impact overall backup performance. Set the number of the stream to the minimal number that does not inhibit the performance by using the following option:

  ```
  STRIPes=1
  ```

- Use a larger buffer size for a better deduplication ratio. Increase the buffer size slowly from the default buffer size of backup application, but do not exceed the amount of buffer that can be handled by the system memory.
- Limit the number of I/O buffers within a backup stream. Ideally, there should be two buffers per stream, with one buffer for reading data from an SQL Server while the other is for sending data to the backup applications, as shown by the following settings:
  - `BUFFer=2`
  - `BUFFERSIze=1024`
  - `SQLBUFFer=0`
  - `SQLBUFFSIze=1024`

20.3.4 LiteSpeed for SQL Server

LiteSpeed for SQL Server is a backup utility that compresses and encrypts the SQL database before the data is stored in backup devices. The factoring ratio is greatly impacted if the data is compressed and encrypted before it reaches the ProtecTIER repository. The ProtecTIER product offers little deduplication benefit if LiteSpeed is used for SQL server backup.

20.4 DB2

This section describes the settings and parameters that should be modified in DB2 environments to enable the maximum performance and optimum factoring for the ProtecTIER server. It also explains why it is possible to combine DB2 compression together with ProtecTIER deduplication.

Updating DB2: Update your DB2 to Version 9.7 Fix Pack 4 or later and use the **`DEDUP_DEVICE`** option for backing up your database. This action results in the best deduplication ratio. DB2 compression types are deduplication friendly.

20.4.1 Combining DB2 compression and ProtecTIER deduplication

DB2 offers multiple options to use compression in conjunction with database rows, database values, or both. Run the **`select tabname,compression from SYSCAT.TABLES`** command to verify the settings for your database.

Table 20-1 shows which types of compression are available.

Table 20-1 DB2 compression types

SYSCAT.TABLES values	Compression type active
R	Row compression is activated if licensed. A row format that supports compression may be used.
V	Value compression is activated. A row format that supports compression is used.
B	Both value and row compression are activated.
N	No compression is activated. A row format that does not support compression is used.

With DB2 compression, data within the database is compressed on a table-row basis. These compressed rows are written to disk as DB2 pages with a default size of 4 K. Changes in a DB2 database with compression enabled affect only the data within these specific DB2 pages; the changes do not affect the entire database because of block based compression, which is different from traditional compression approaches. Also, after changes occur within the database, only the changed pages are recompressed.

Effectively, compressed DB2 pages are not apparent to HyperFactor and a large sequence of compressed pages factors well if they are not changed. There is no general penalty for using compression within DB2. The data change rate affects our deduplication ratio; whether you use compression or not, the behavior is the same.

Remember, even if DB2 compression does have a friendly synergy with ProtecTIER deduplication, the full deduplication potential can be reached only with all sorts of data reduction technology, such as disabled compression.

Warning: Using another form of compression with DB2 database backups, for example, Tivoli Storage Manager compression or the compression feature of another backup software, still impacts your achievable deduplication ratio.

20.4.2 Upgrading the DB2 database to improve deduplication

The most recommended method for improving deduplication is to upgrade the DB2 database to DB2 9.7 Fix Pack 4 or later to be able to use the **DEDUP_DEVICE** option. This special feature improves the DB2 backup process to make it deduplication friendly. Update to DB2 9.7 Fix Pack 4 or later. Only with this version or any later version can you experience the full benefit of the optimized DB2 data handling for deduplication devices.

You can download DB2 Fix Packs for DB2, for Linux, UNIX, and Windows, and IBM DB2 Connect™ products from the following website:

http://www.ibm.com/support/docview.wss?uid=swg27007053

To fully understand the improvements of the **DEDUP_DEVICE** option, look at the default DB2 database backup behavior. When a DB2 backup operation begins, one or more buffer manipulator (db2bm) threads are started. These threads are responsible for accessing data in the database and streaming it to one or more backup buffers. Likewise, one or more media controller (db2med) threads are started and these threads are responsible for writing data in the backup buffers to files on the target backup device. The number of db2bm threads that is used is controlled by the **PARALLELISM** option of the **BACKUP DATABASE** command. The number of db2med threads that is used is controlled by the **OPEN n SESSIONS** option. Finally, a DB2 agent (db2agent) thread is assigned the responsibility of directing communication between the buffer manipulator threads and the media controller threads. This process is shown in Figure 20-4.

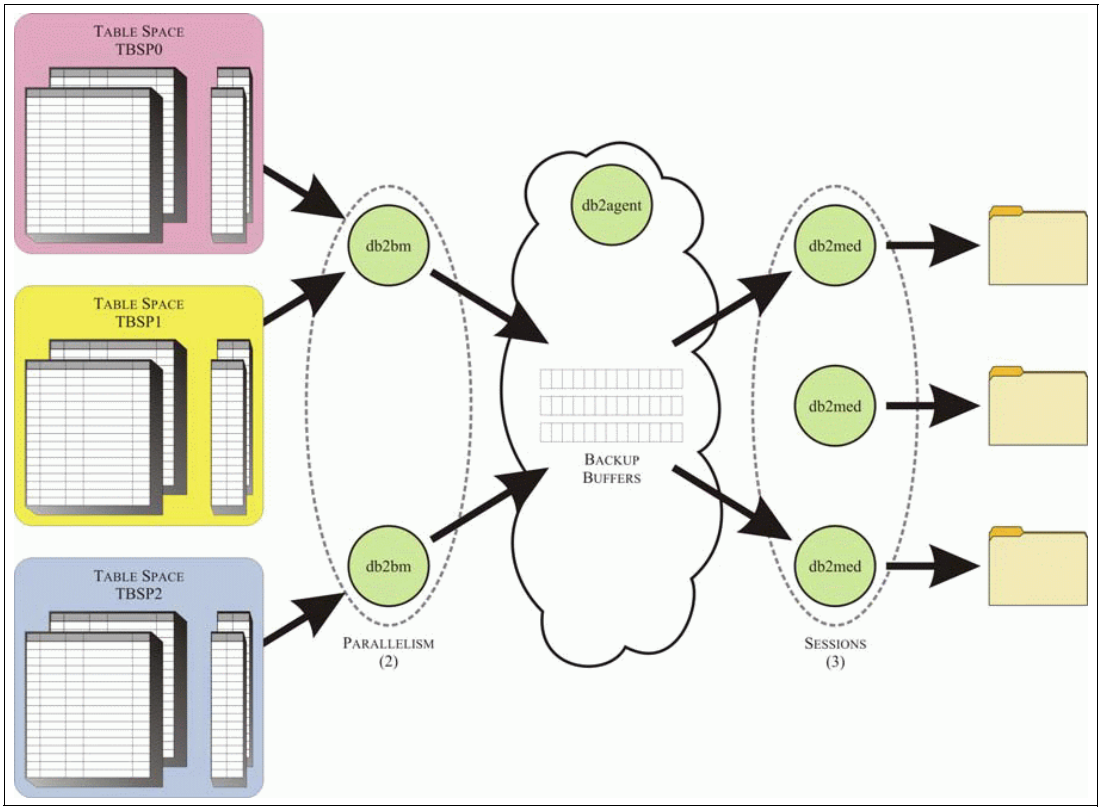

Figure 20-4 DB2 backup process model

Without the **DEDUP_DEVICE** option, data that is retrieved by buffer manipulator (db2bm) threads is read and multiplexed across all of the output streams that are being used by the media controller (db2med) thread. There is no deterministic pattern to the way in which data is placed in the output streams that are used (Figure 20-5). As a result, when the output streams are directed to a deduplication device, the device thrashes in an attempt to identify chunks of data that are already backed up.

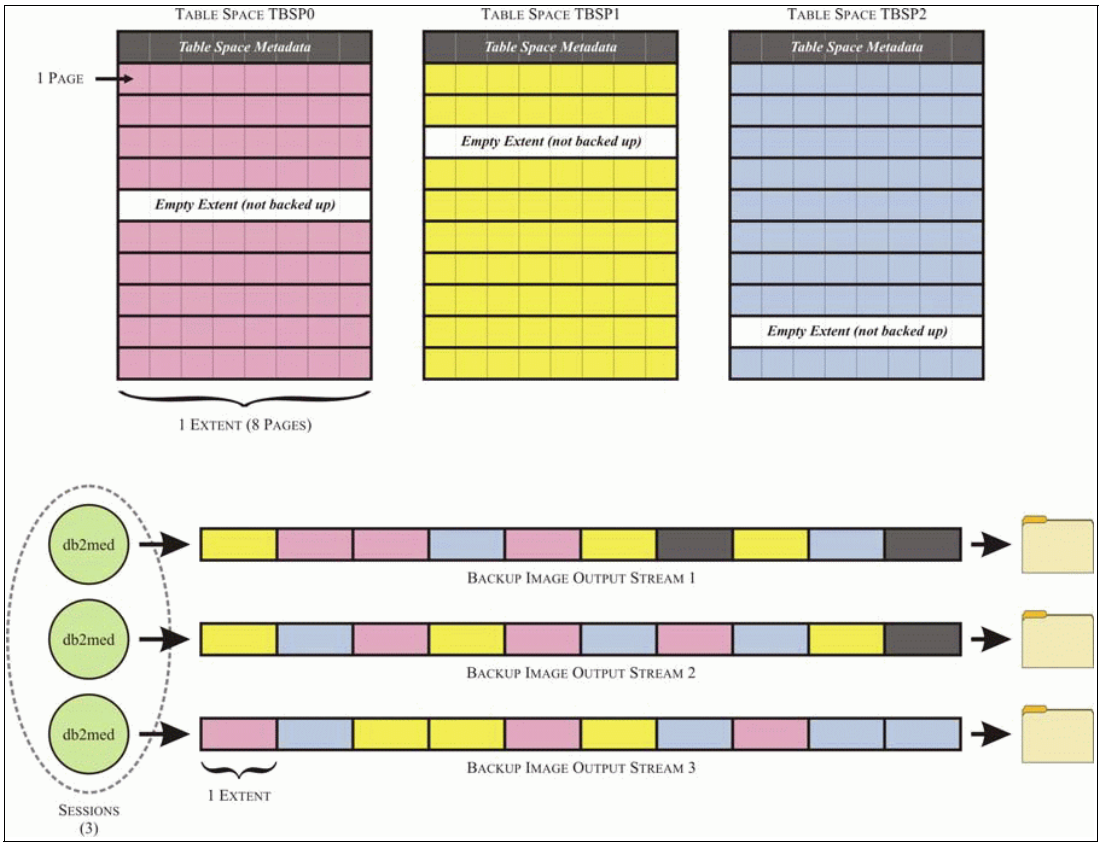

Figure 20-5 Default database backup behavior

20.4.3 DB2 DEDUP_DEVICE setting

When the **DEDUP_DEVICE** option is used with the **BACKUP DATABASE** command, data that is retrieved by buffer manipulator (db2bm) threads is no longer read and multiplexed across the output streams that are being used by the media controller (db2med) threads. Instead, as data is read from a particular table space, all of that table space's data is sent to only one output stream. Furthermore, data for a particular table space is always written in order, from lowest to highest page. As a result, a predictable and deterministic pattern of the data emerges in each output stream, making it easy for a deduplication device to identify chunks of data that are already backed up.

Figure 20-6 illustrates this change in backup behavior when the **DEDUP_DEVICE** option of the **BACKUP DATABASE** command is used.

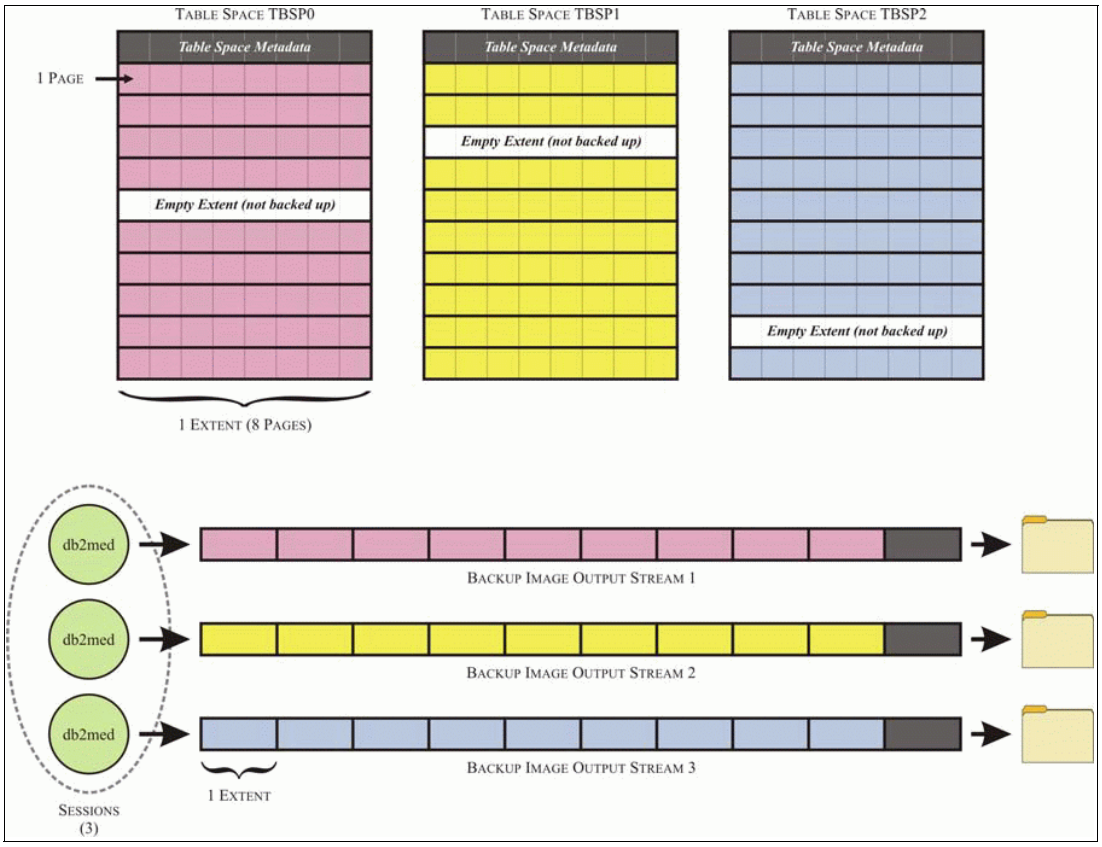

Figure 20-6 Database backup behavior with the DEDUP_DEVICE option

When you use the **DEDUP_DEVICE** option, each table space is backed up to a dedicated tape drive. Using a number of virtual tape drives that is equal or greater than the number of table spaces you want to back up is recommended.

If the database contains table spaces larger than others (above 30% of the entire database size), it prolongs the backup. If this situation affects the backup window, consult your DB2 support to assist you in splitting the larger table spaces and making them smaller. Also, communicate this information to the DB2 database planning staff so that future deployments can directly benefit of the improved deduplication without any drawback.

20.4.4 Example of DEDUP_DEVICE setting

Example 20-4 uses 16 tape drives to back up your 16 table spaces that ideally are of equal size using Tivoli Storage Manager, in parallel, using the **DEDUP_DEVICE** option.

Example 20-4 Multistreamed backup of DB2

```
db2 backup db <database_name> use tsm open 16 sessions dedup_device exclude logs
```

This action results in the best possible deduplication ratio. With DB2 9.7 Fix Pack 4 and later, the DB2 self-tuning capability is able to support this backup command by choosing all tuning values automatically.

20.4.5 Excluding logs from the DB2 database backup

Use the **exclude logs** parameter to avoid backing up your database logs to the same destination as your database. Database logs tend to have a 100% change rate and therefore have a negative impact on your overall HyperFactor ratio. Instead, redirect the archive logs directly to storage with no further active data reduction technology. Using the **include logs** parameter with the DB2 backup command results in archive logs being automatically added to the backup images. This action causes different patterns in the backup streams and reduces deduplication efficiency.

20.4.6 DB2 recommended settings without DEDUP_DEVICE

Backing up to a deduplication device when the **DEDUP_DEVICE** option is not available can still be optimized by applying some rules. The DB2 settings in Table 20-2 provide the best deduplication efficiency for backing up without the **DEDUP_DEVICE** option.

Table 20-2 Recommended DB2 settings

DB2 parameter	Recommended value	Description
sessions/ **OPEN n SESSIONS**	Minimum[a]	Change the value to read the data at the required backup rate.
buffers/ **WITH num-buff BUFFERS**	Parallelism + sessions + 2	The numbers of buffers should be #sessions + #parallelism +2. Also, the following calculation must fit: (num-buffers * buffer-size) < UTIL_HEAP_SZ (UTIL_HEAP_SZ is the database utility heap size).
buffer/ **BUFFER buff-size** (specified in multiples of 4 KB pages)	16384	This value requires much memory. If this value be too much for your environment, use the largest possible **BUFFER** value instead.
parallelism/ **PARALLELISM**	Minimum[a]	Change the value to read the data at the required backup rate.

a. Select the minimum value to allow an acceptable backup window time frame. A value of 1 is the best for deduplication, but it might increase backup times in large multi-table space databases.

BUFFER setting: The large **BUFFER** size of 16384 is the setting with the most impact on your HyperFactor deduplication. The bigger the **BUFFER** value is, the better your deduplication ratio is.

20.4.7 Example of DB2 command using sessions, buffers, and parallelism

Example 20-5 shows an example of a DB2 backup command using four sessions, eight buffers, a buffersize of 16384, and a parallelism of 2.

Example 20-5 Database backup command

```
db2 backup db <databasename> use tsm open 4 sessions with 8 buffers buffer 16384
parallelism 2
```

Tip: Always use the same parameters for restore as you did for backup (number of sessions, buffers, buffer size, and parallelism) to ensure maximum restore performance.

20.5 Oracle

Oracle Recovery Manager (RMAN) is a backup and recovery utility of Oracle databases. The RMAN backs up Oracle databases directly to the disk or to other storage devices using third-party backup applications. Backup applications interface with RMAN to back up Oracle databases with various storage devices, such as tape, file system, or OST (if Symantec NetBackup is used).

The ProtecTIER server can be deployed as a VTL, FSI, or as OST. For more details about how to set up VTL, FSI, and OST, see Chapter 3, “Virtual Tape Library guidelines” on page 37, Chapter 4, “ProtecTIER File System Interface: General introduction” on page 53, and Chapter 6, “OpenStorage guidelines” on page 95.

This section describes the optimal settings and guidelines of RMAN to improve the backup throughput and factoring ratio of the ProtecTIER solution.

20.5.1 Recommendations for RMAN settings

Here are some suggested settings for RMAN:

- Perform a daily full backup whenever possible. Performing a full backup enables the simplest and fastest restoration.
- Enable **ARCHIVELOG** mode for your database. Run **ARCHIVELOG** as often as possible, and to back up the archived logs as soon as possible.
- Disable compression and encryption within Oracle databases and backup applications.
- Disable or minimize multiplexing. Multiplexing enables RMAN to combine data blocks from different files into a single backupset, which impacts the factoring ratio. RMAN multiplexing is affected by the following two parameters:
  - The **FILESPERSET** parameter determines how many files should be included in each backup set. Set **FILESPERSET=1** to send only one file per backupset in each channel (backup stream).
  - The **MAXOPENFILES** parameter defines how many files RMAN can read from the Oracle source simultaneously. Set **MAXOPENFILES=1** to prevent RMAN reading from more than one file at a time.

Example 20-6 shows a calculation of multiplexing with different **FILESPERSET** and **MAXOPENFILES** settings.

Example 20-6 Calculation of multiplexing in RMAN

```
Scenario 1: FILESPERSET=6, MAXOPENFILES=3, number of data files=4
Multiplex = 3 (Limiting by the MAXOPENFILES setting)

Scenario 2: FILESPERSET=2, MAXOPENFILES=3, number of data files=4
Multiplex = 2 (Limiting by the FILESPERSET setting)

Scenario 3: FILESPERSET=8, MAXOPENFILES=4, number of data files=2
Multiplex = 2 (Limiting by the number of data files)

Scenario 4: FILESPERSET=1, MAXOPENFILES=1, number of data files=4
Multiplex = 1 (Limiting by the FILESPERSET and MAXOPENFILES settings)
```

- Increase the number of parallel backup streams to improve backup throughput. Ensure that the number of ProtecTIER virtual tape drives that are available for Oracle backup matches the number of parallel streams that are configured in RMAN. For example, this value is enabled in the definition of the Tivoli Storage Manager client on the Tivoli Storage Manager server by using the **MAXNUMMP=32** parameter. Set **PARALLELISM=32** (up to 64).

Figure 20-7 shows a case study that shows the factoring ratio and a backup throughput result with different multiplexing and parallel channel settings. The result is taken from a case study of Oracle database backup with Tivoli Storage Manager, and a 30-day retention period on ProtecTIER virtual tape. A full backup that is performed on alternate days averages a 5% data change rate between the full backups.

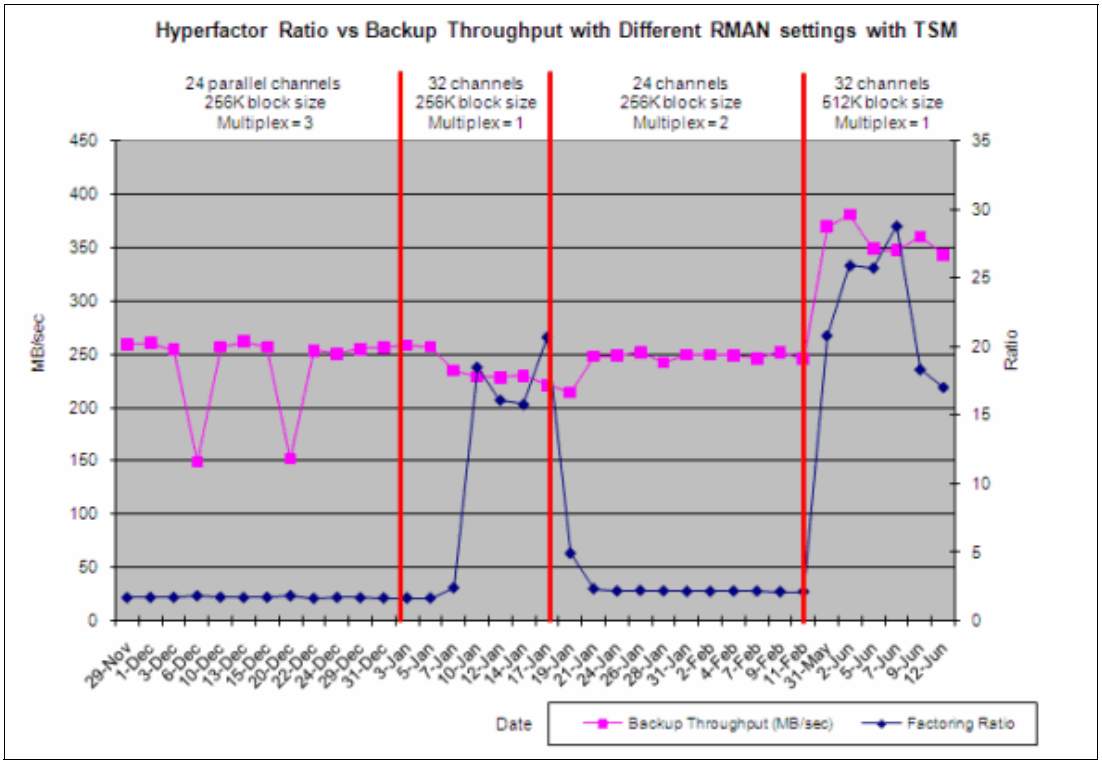

Figure 20-7 Example of multiplexing and parallelism impact on the HyperFactor ratio

Environmental variations: The parameters that are used in this test are not absolute requirements. Different environments might produce different results, depending on the data change rate and backup practices. Fine-tune the RMAN settings in your environment gradually to get the settings that do not inhibit performance.

20.6 SAP

This section describes settings and parameters to be modified for optimum performance when you are working with specific data types, such as SAP integrated with Tivoli Storage Manager.

20.6.1 SAP introduction

SAP is an acronym for *Systems Applications and Products*. SAP provides a common centralized database for all the applications that are running in an organization. The database instance is a mandatory installation component for the installation of an SAP system.

SAP supports the following databases:

- Oracle
- MS SQL Server
- IBM DB2 Universal Database™ for UNIX and Windows
- SAP liveCache technology
- MaxDB
- IBM DB2 Universal Database for z/OS®
- IBM DB2 Universal Database for iSeries
- IBM Informix®

More database and operating system support information can be found in the Product Availability Matrix (PAM) at the SAP Service Marketplace. A login is required to access SAP Service Marketplace, which can be found at the following address:

`http://service.sap.com/pam`

20.6.2 Data protection for SAP

Data protection for the SAP server involves steps to protect all of the software components that are needed by the SAP system to operate. The base components of the SAP server are the operating system, the SAP application server, the database instance, and the data files. Each component requires different data protection techniques.

The SAP system uses the relational database as main storage for all SAP data and meta information. This main storage is the basis for the tight integration of all SAP application modules and ensures consistent data storage. Data in the SAP database is unique for every company, and if the data is lost, it cannot be simply reinstalled in the same manner as an operating system is reinstalled. Therefore, it is important to take special care when you plan the protection of the data that is stored in the SAP database.

Protection of the SAP database

The protection of the SAP database has two parts: protecting the database binary files and configuration files, and protecting data that is stored in the data files.

Database binary files and configuration files are typically backed up as part of the operating system or file system backup. The backup of the database data files and other supporting structures that are associated with SAP data should be performed by a dedicated tool that is designed especially for the database backup. You can use database backup and restore tools to perform backup and restore data in a consistent state.

The backup tools can also perform an online backup of the database and backup of the redo log files just after the log files are archived. A backup of a database creates a copy of the database's data files, control files, and, optionally, log files. It then stores these files on backup media.

A consistent backup, also called an offline backup or cold backup, is a backup of all the data files in the database that is taken when all interim changes are physically written to the data files. With a consistent backup, partial changes from the log files that are not written to the data files are not backed up. If you restore a database from a consistent backup, the database is in a consistent state when the restore operation finishes. In addition:

- For an Oracle database, a consistent backup can be taken only when the database is shut down for the entire duration of the backup procedure.
- For a DB2 Universal Database (UDB), a database must be deactivated, or the instance must be stopped before the backup operation starts.

The database must stay inactive until the backup finishes. This action ensures that there are no data changes on the database at the time the backup is being taken. A consistent backup is always a backup of the entire database; it cannot be a partial or incremental backup.

You can take an offline backup by either using a dedicated database backup tool, such as Oracle Recovery Manager, BR*Tools, the **DB2 BACKUP** command, or by using a non-database backup tool, such as the Tivoli Storage Manager backup archive client. The dedicated database backup tools ensure that all the objects that are required for the successful database restore are included in the backup image. The database backup tool also ensures that the location, time stamp, type, and other information about the backup copy is registered in the repository, such as BR*Tools logs or a DB2 database history file. Using the metadata in the repository, backup tools can perform an automatic restore that is based on the specified time stamp without prompting for the backup images to restore and their location.

IBM offers products for both data protection and data retention and reduction. For example, in the SAP environment, there are Tivoli Storage Manager for Enterprise Resource Planning (ERP) and Tivoli Storage Manager for Database for data protection. For data retention, you can use IBM DB2 CommonStore for SAP. Both solutions can use a Tivoli Storage Manager server for the media manager.

These solutions are illustrated in Figure 20-8.

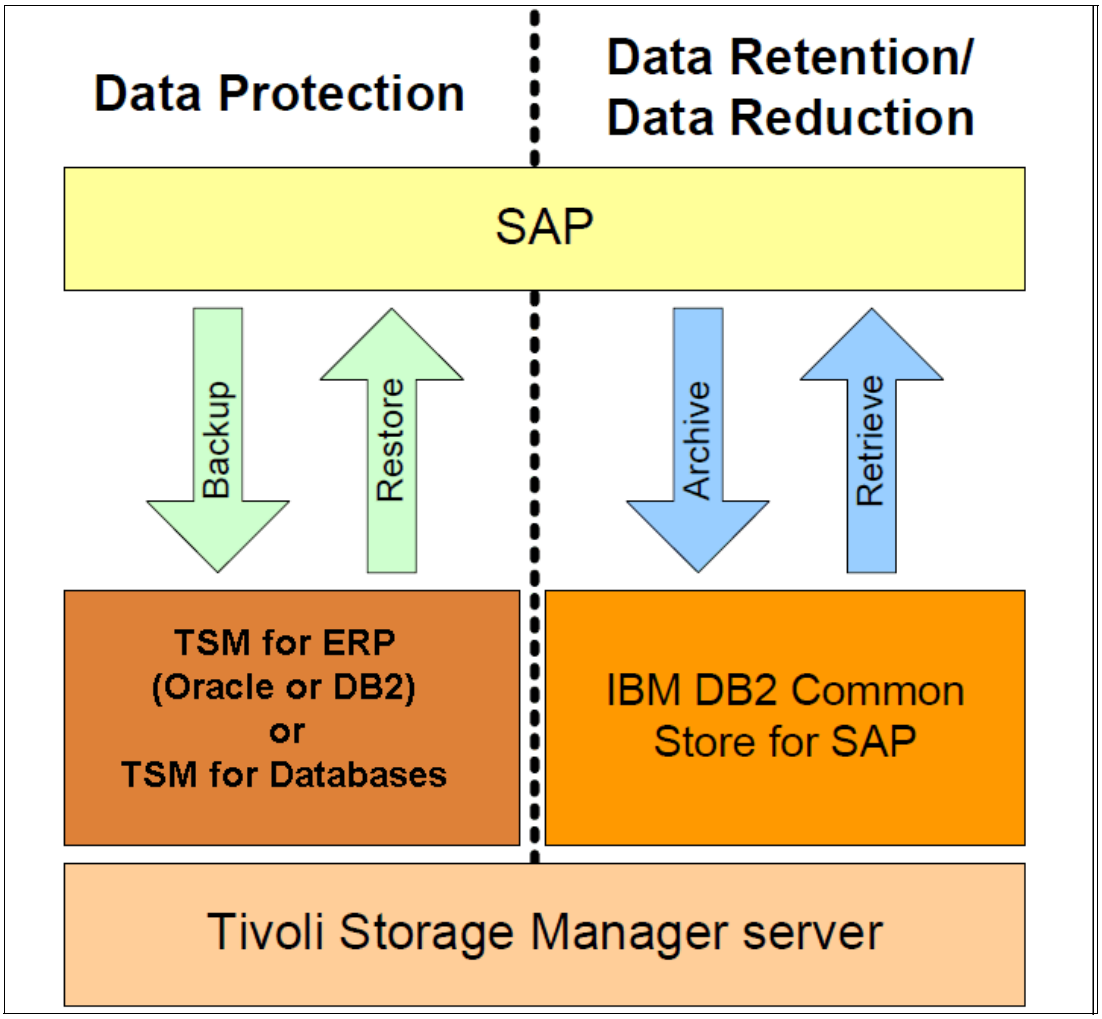

Figure 20-8 Backup and archival products

Tivoli Storage Manager for Enterprise Resource Planning (ERP), formerly known as Tivoli Data Protection for SAP, is a component of Tivoli Storage Manager family that provides a complete backup solution for SAP databases. The current version supports Oracle and DB2 only.

The following features are available for Tivoli Storage Manager for ERP:

- Handles large amounts of data
- Optimized processor usage that reduces the overall time for backup and restore
- Optimized for an SAP environment
- Supports multiple management classes

Additional information: For more information about Tivoli Storage Manager for ERP, go to the following website:

`http://www.ibm.com/software/tivoli/products/storage-mgr-erp/`

20.6.3 Integration of Tivoli Storage Manager for ERP with SAP

Tivoli Storage Manager for ERP is fully integrated in to the SAP environment. The communication between the backup and archive server is performed by an API called *ProLE*. This API is shared with other Tivoli Data Protection products. ProLE runs as a background process and provides communication with the Tivoli Storage Manager server. Figure 20-9 shows a sample architecture of Tivoli Storage Manager for ERP integrated with SAP.

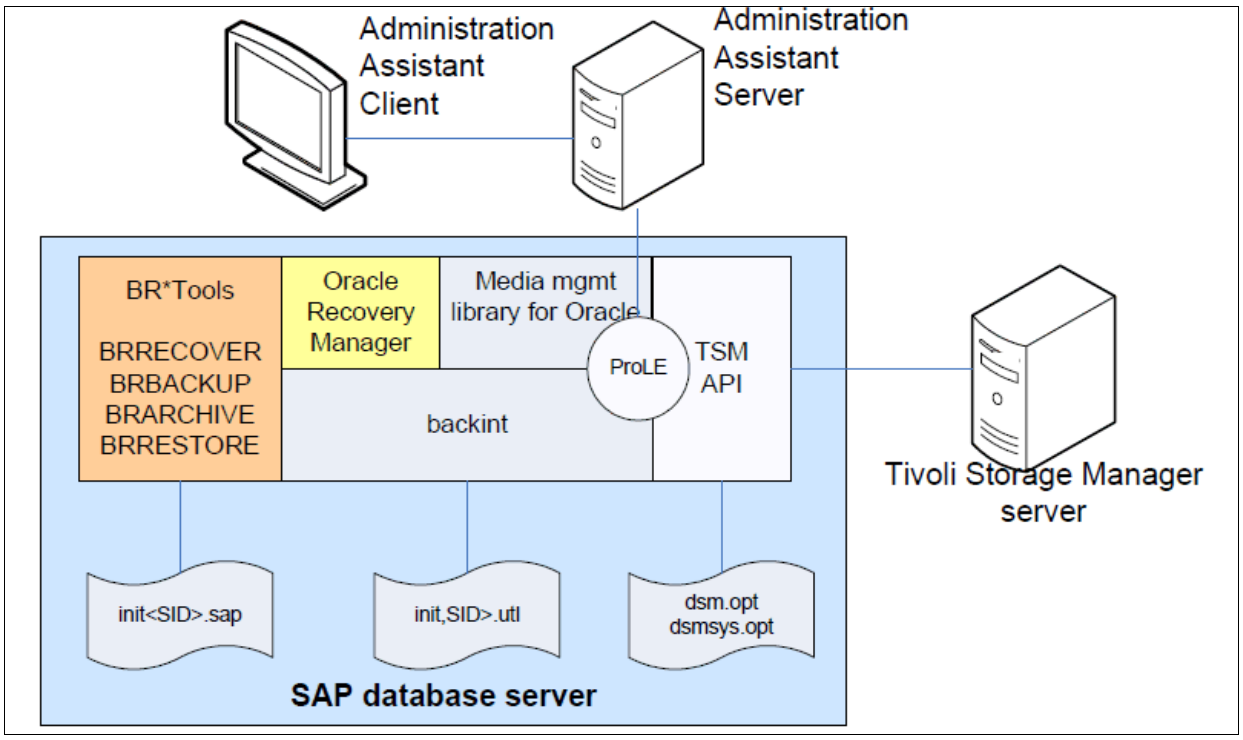

Figure 20-9 Tivoli Storage Manager for ERP sample scenario

Additional information: For more information about Tivoli Storage Manager for ERP, go to the following website:

`http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3/index.jsp?topic=%2Fcom.ibm.itsm.nav.doc%2Ft_protect_dperp.html`

20.6.4 Tivoli Storage Manager for ERP for Oracle database

Tivoli Storage Manager for ERP is a client and server program that manages backups and restores in conjunction with the Tivoli Storage Manager. With Tivoli Storage Manager for ERP, it is possible to handle SAP database backups, and it includes the ability to manage backup storage and processing independently from normal SAP operations. Furthermore, Data Protection for SAP in combination with Tivoli Storage Manager provides reliable, high performance, and repeatable backup and restore processes to manage large volumes of data more efficiently.

For Oracle databases, you have two options to implement a backup using Tivoli Storage Manager:

- Tivoli Storage Manager for ERP using the **BACKINT** interface
- Tivoli Storage Manager for ERP using Oracle Recovery Manager (RMAN)

With the integration, it is possible to follow the ERP backup and restore procedures and to use the integrated SAP database utilities **BRBACKUP**, **BRARCHIVE**, **BRRESTORE**, and **SAPDBA** for backup and restore. Other SAP-related files (executable files) are backed up by using Tivoli Storage Manager standard techniques for file backup and restore, for example, incremental backup, file filtering, and point-in-time recovery.

Tivoli Storage Manager for ERP for Oracle using BACKINT

Using this feature, you can perform the traditional Oracle online backup with automation provided by **BACKINT**. Figure 20-10 shows the data interface between Oracle Databases and Tivoli Storage Manager for ERP for Oracle using the **BACKINT** interface.

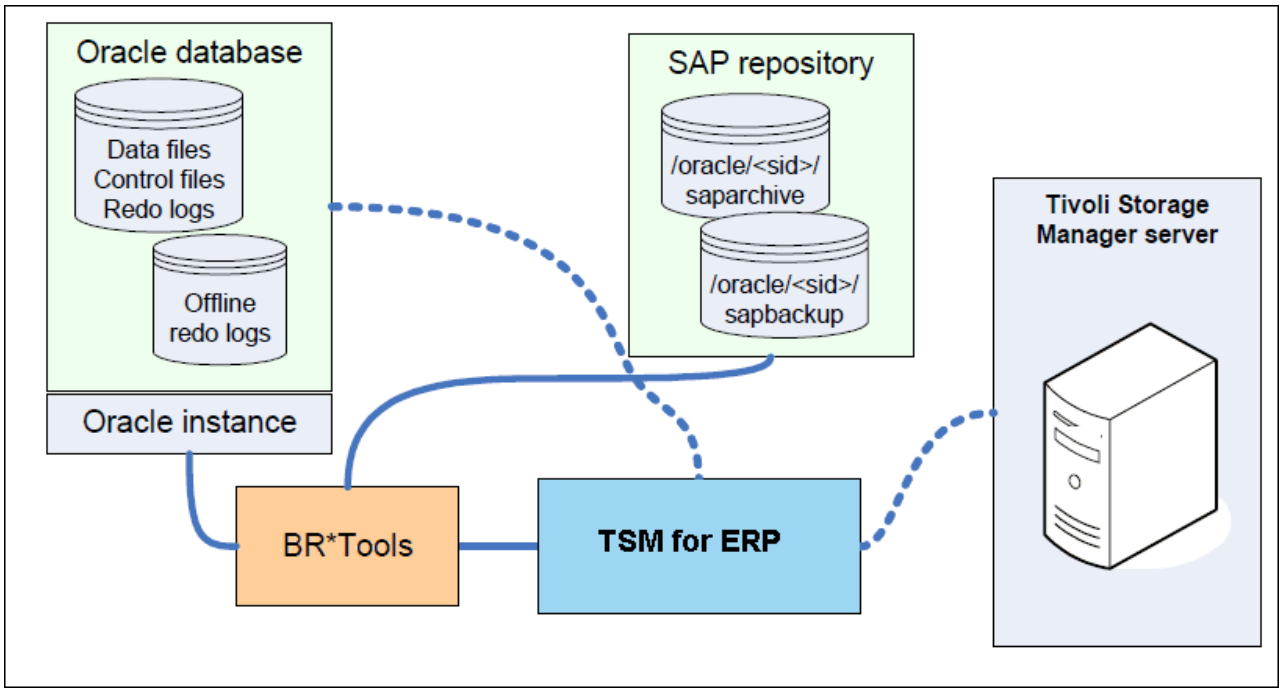

Figure 20-10 Tivoli Storage Manager for ERP for Oracle using BACKINT

The backup proceeds as follows:

1. BR*Tools takes control.
2. **BRBACKUP** calls the Tivoli Storage Manager for ERP by using **BACKINT**.
3. **BACKINT** changes the table spaces to backup mode by running the following command:

   ```
   alter tablespace <tablespace name> begin backup
   ```

4. **BACKINT** using Tivoli Storage Manager for ERP reads all the data files and saves them to Tivoli Storage Manager server.
5. BR*Tools updates the catalog with information about the backed up data file.

Logs: BR*Tools logs are stored in the `/oracle/<SID>/saparch` directory.

BRBACKUP automatically backs up the logs and profiles after every backup operation. In the case of bare metal restore or disaster recovery, logs and profiles must be restored to enable BR*Tools to restore data files. The process can be simplified if the logs and profiles are backed up by a Tivoli Storage Manager backup archive client during the file system backup.

Using this method, the chosen data files are sent to Tivoli Storage Manager one by one. No compression or block checking is performed at this level.

When a database is in backup mode, the amount of redo logs that are written to disk increases because Oracle writes the entire dirty block to the disk, not just the updated data. In some cases, when the backup routine fails for any reason, the data file remains in active backup mode, which can cause some performance impact and additional I/O to the disk.

Tivoli Storage Manager for ERP for Oracle using RMAN

Using this feature, you can take advantage of all the facilities that are provided by RMAN. In general, RMAN is able to perform a backup in less time compared to the traditional backup using **BACKINT** because RMAN sends only used data blocks (in an Oracle data file) to Tivoli Storage Manager. The other interesting feature is block checking, which discovers bad blocks as soon as they occur.

In addition, you can use the Oracle Recovery Manager (RMAN) utility to run some tasks that are not provided by BR*Tools, such as incremental backups, releasing backup versions, and catalog maintenance. Figure 20-11 shows the data interface between Oracle Database and for Oracle for SAP using RMAN.

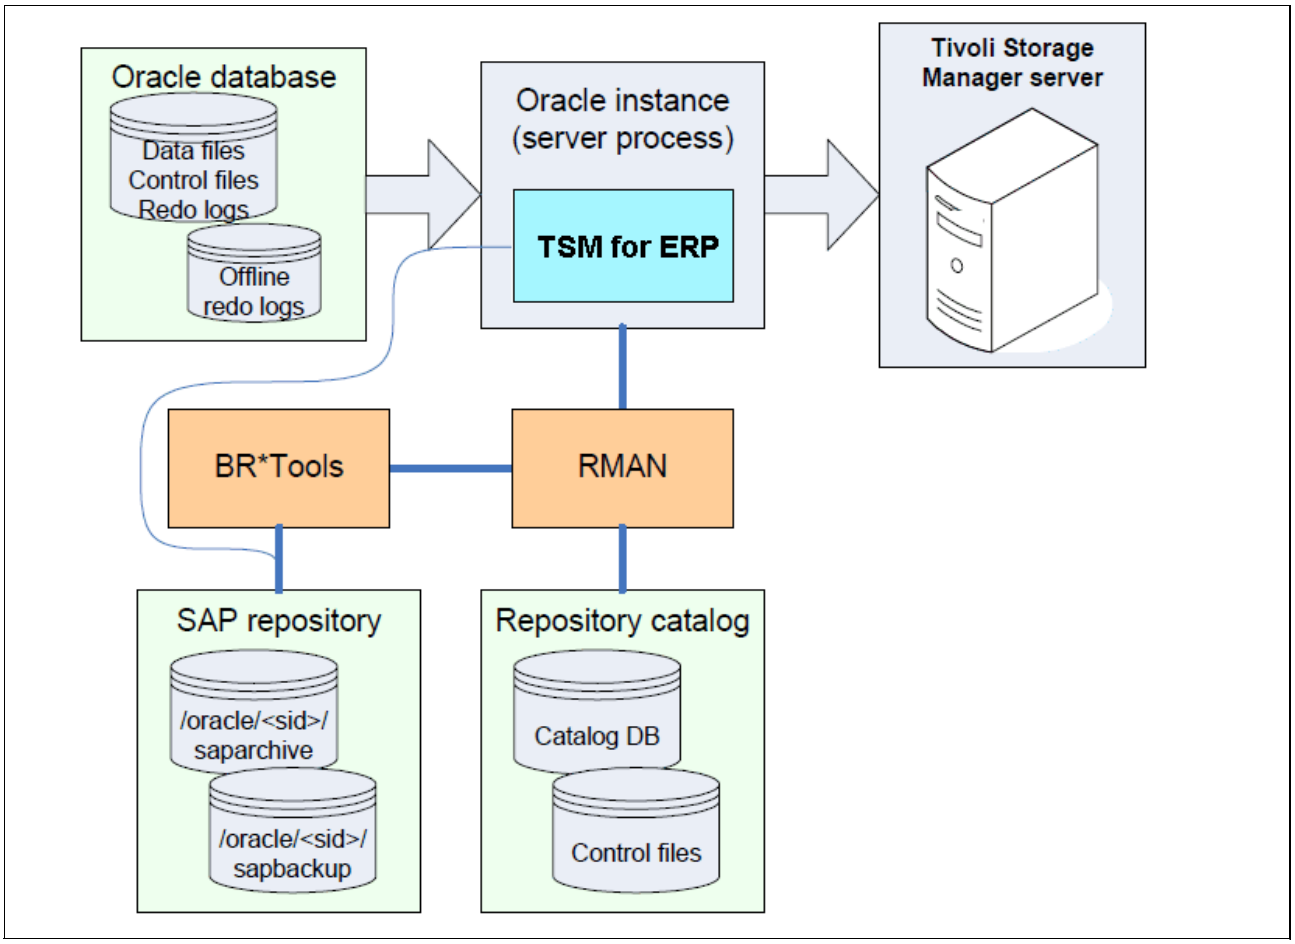

Figure 20-11 Tivoli Storage Manager for ERP for Oracle using RMAN

20.6.5 Tivoli Storage Manager for ERP for DB2

Tivoli Storage Manager for ERP for Oracle for DB2 was created to provide an intelligent interface to manage backup and restore by using Tivoli Storage Manager. It is fully integrated in to the SAP environment. The backup command **`DB2 BACKUP DATABASE`** and the restore command **`DB2 RESTORE DATABASE`** are run at the DB2 CLI, which calls the Tivoli Data Protection for SAP for DBA module.

The backup and restore of the DB2 log files is provided by the BR*Tools commands **`BRARCHIVE`** and **`BRRESTORE`**. In addition, you can use the Tivoli Storage Manager for ERP for DB2 Tools BackOM and the built-in Log Manager. Figure 20-12 shows the data interface between DB2 Databases and Tivoli Storage Manager for ERP for DB2.

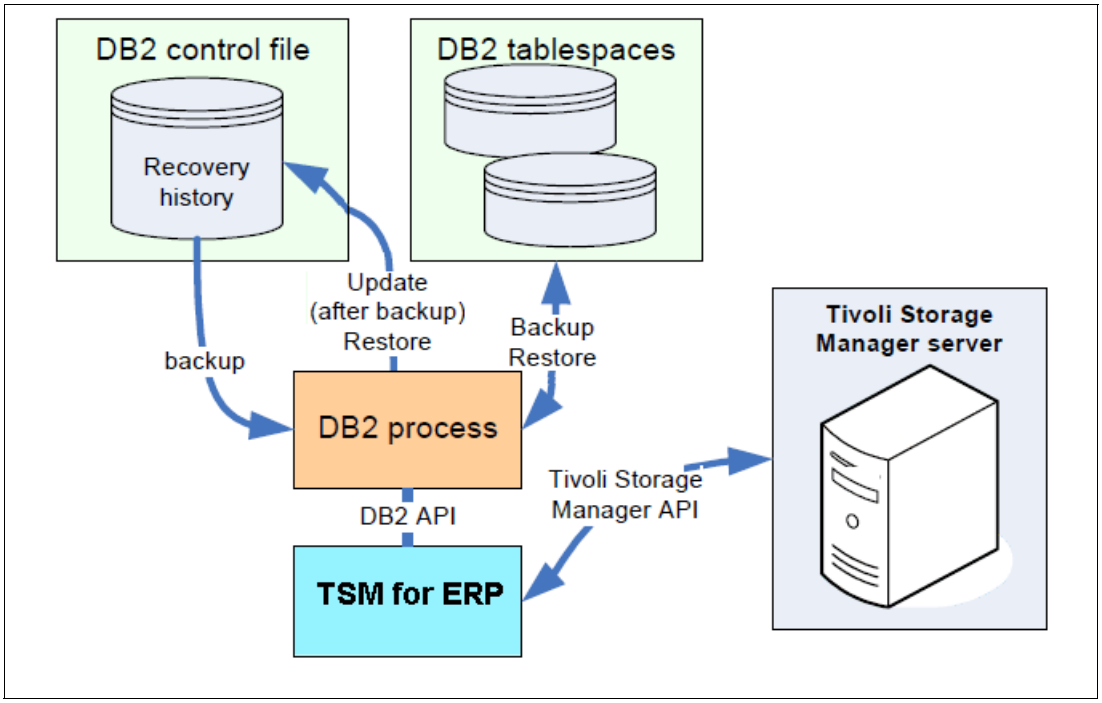

Figure 20-12 Tivoli Storage Manager for ERP for DB2

The archiving of DB2 offline log files is provided by the SAP tool **`BRARCHIVE`**. The retrieval of DB2 offline log files is provided by the SAP tool **`BRRESTORE`** and by the Tivoli Storage Manager for ERP tool BackOM. As of DB2 Version 9.X, offline log files can be archived and retrieved with the DB2 built-in Log Manager.

The DB2 command line processor (CLP) interprets commands for the DB2 database and passes control to a DB2 Server Process. In the case of Tivoli Storage Manager for ERP, the **`LOAD <libraryname>`** option causes DB2 to start the Tivoli Storage Manager for ERP shared library. This process runs during the backup or restore, loads the library dynamically, and communicates with it through the Tivoli Storage Manager API. To start a backup or restore, the DB2 CLP communicates with the DB2 Server Process, providing the server process with the relevant information for processing the database.

Additional information: For more information about backup methodologies for SAP, see *SAP Backup using Tivoli Storage Manager*, SG24-7686.

All the backup solutions that are described in this section can be integrated with advanced backup techniques, such as LAN-free backup, parallel transfer of backup data to and from Tivoli Storage Manager server, or multiplexing.

Reduction of processing time: Implementation of these techniques can reduce backup and restore times, and eliminate the impact of backup data transfers on LAN throughput.

20.6.6 SAP BR*Tools for Oracle using BACKINT

SAP BR*Tools for Oracle is a package of utilities, developed by SAP AG to protect and manage SAP data that is stored in Oracle databases. BR*Tools supports functions for online, offline, partial, or full backups of database (**BRBACKUP**) and backups of archived redo logs (**BRARCHIVE**). It provides functions for database restore and recovery (**BRRECOVER** and **BRRESTORE**).

BR*Tools can be used not only for database recoverability tasks, but it can also serve as a tool for creating homogeneous database copies, and can assist with database migration to different platforms or database versions.

BRBACKUP is the BR*Tools utility that enables online or offline backup of database files (data files, control files, and online redo log files). **BRBACKUP** can be used to back up individual data files, table spaces, or the entire Oracle database. **BRBACKUP** also backs up the BR*Tools configuration profiles and logs that are required for the database's disaster recovery.

The smallest unit that can be saved with **BRBACKUP** is a file. You can use **BRBACKUP** for backing up both files in the database and non-database files and directories. Use the **backup_mode** command from the Initialization Profile `init<DBSID>.sap` file or the command option **brbackup -m|-mode** for this purpose.

Before the offline backup is taken, **BRBACKUP** automatically closes the database and opens it when the backup is accomplished. **BRBACKUP** can also change the status of the table space to be backed up to **BEGIN/END BACKUP**.

You can also instruct **BRBACKUP** to use software compression. The software compression client can enhance the backup, especially if the network is slow.

Compression: If you plan to send data to a ProtecTIER server, do not enable software compression; it might affect the overall deduplication ratio.

The most frequently used **BRBACKUP** function is a full database backup. Example 20-7 shows an example of running **BRBACKUP**.

Example 20-7 Online backup by using the databases BRBACKUP tool

```
$su - cptadm
$BRBACKUP -c -u / -t ONLINE_CONS -m FULL -p /oracle/CPT/102_64/dbs/initCPT.sap
```

You can perform a full backup by running **BRBACKUP** with the following options:

- The mode option (**-mode/-m**) is set to `FULL` or `ALL`.
- You can start a full backup either in online mode (**-type/-t online_cons**) or in offline mode (**-type offline**). In the case of the **online_cons** type, the offline redo log files that are generated during the full backup are also backed up to the same media.

- The backup storage media is defined by the BR*Tools profile file that is specified by the **BRBACKUP** parameter **-profile/-p**.
- The user name and password that is used by **BRBACKUP** to log on the Oracle database system is specified by the parameter **-user/-u**. If you are working as a DBA user that is authenticated to the database by the OS ($OPSuser), you can use "/" as value of this parameter.

The parameter "**-confirm/-c**" stands for an unattended mode, which is mostly used in the backup scripts, so BR*Tools does not prompt you for confirmations.

Archived redo log backup functions

BRARCHIVE provides functions for offline redo log files backup in Oracle databases that run in archiving mode. If archiving is enabled, a database cannot overwrite an active log file until the content is archived. Whenever an active redo log is filled, the database performs a log switch and starts writing to another log file. The full redo log files are archived by Oracle background processes into the `archivelog` directory.

The redo log is the most important database component for a recovery from a crash, media failure, or user failure. Therefore, at least the production databases should be configured in archiving mode. To prevent the `archivelog` directory from filling up, **BRARCHIVE** should be ran periodically to move the offline redo logs from the archive directory to the backup media.

BR*Tools and Tivoli Storage Manager for ERP - Oracle

BR*Tools interacts with Tivoli Storage Manager for ERP - Oracle through the **BACKINT** interface. The communication of BR*Tools and BACKINT occurs as follows:

1. The BR*Tools utility **BRBACKUP** informs Oracle of what data must be backed up and puts the database into the correct backup state (online or offline backup).
2. **BRBACKUP** calls Tivoli Data Protection for ERP using the **BACKINT** interface with a list of all files to be backed up.
3. Tivoli Data Protection for ERP reads all the requested files from the database and reports back to **BRBACKUP**. **BRBACKUP** adds these files to the repository that contains all processed backups.
4. **BACKINT** transfers the data to the Tivoli Storage Manager server by using the Tivoli Storage Manager Client API.
5. The BR*Tools updates the repository that contains information about the status of the files.

BR*Tools configuration

To configure BR*Tools, complete the following steps:

1. The BR*Tools configuration is stored in the `init<SID>.sap` initialization profile file. The configuration file contains parameters that affect the performance of backup and restore functions. The default location of the file is `<ORACLE_HOME>/dbs` (UNIX) or `<ORACLE_HOME>\database` (Windows).
2. Some parameters that are specified in the profile can be overridden if the BR*Tools programs are called with different command options. In the BR*Tools profile, you can specify the backup adapter that is used to transfer data (**cpio**, **BACKINT**, or RMAN).
3. If you set up BR*Tools to use the **BACKINT** adapter, you need to reference the appropriate **BACKINT** profile (*.utl file) in the BR*Tools profile. If you want to instruct BR*Tools to use Oracle RMAN, you must define the RMAN channel parameters in the BR*Tools profile.

4. The configuration profile of Tivoli Storage Manager for ERP is defined in the `init<SID>.utl` file, which is in the same directory as the BR*Tools profile (`init<SID>.sap`). The configuration parameters in the `init<SID>.utl` file include the Tivoli Storage Manager node name and the management classes to be used for data files backup and offline redo logs backup. If the backup retention is going to be controlled by Tivoli Storage Manager for ERP, you can set up the number of backup versions to be kept in this file.
5. The configuration file of the Tivoli Storage Manager API (`dsm.sys`) is, by default, stored in the Tivoli Storage Manager API installation directory (specified by the environmental variable `DSMI_DIR`). The configuration file of the Tivoli Storage Manager API Client defines the network settings (protocol and network address of the Tivoli Storage Manager server) to enable communication between the API client and the Tivoli Storage Manager server.
6. You also specify in this file the authentication type (**`PASSWORDACCESS`**) that the Tivoli Storage Manager API client uses to connect to the Tivoli Storage Manager server. Additionally, if the Storage Agent is operable on the local node in this file, you can instruct the Tivoli Storage Manager API client to use the LAN-free backup (by using the **`LANFREE yes|no`** option).
7. Instruct BR*Tools to use the **BACKINT** interface by setting the **`backup_dev_type`** parameter in the SAP initialization file (`init<SID>.sap`) as follows:

   ```
   backup_dev_type = util_file
   ```

8. Instruct BR*Tools to use the `init<SID>.utl` file (created by the Tivoli Storage Manager for ERP installation wizard) by setting the **`util_par_file`** parameter in the SAP initialization file:

   ```
   util_par_file=<full_path>/init<SID>.utl
   ```

Archiving functions: Tivoli Storage Manager for ERP uses the Tivoli Storage Manager archive functions to transfer data to Tivoli Storage Manager server and vice versa. Thus, the management classes that are assigned to Tivoli Storage Manager for ERP (in the `init<SID>.utl` file) must have an archive copy group defined.

For more information about Tivoli Storage Manager for ERP, go to the following website:

http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3/index.jsp?topic=%2Fcom.ibm.itsm.nav.doc%2Ft_protect_dperp.html

20.6.7 SAP BR*Tools for Oracle using RMAN with Tivoli Storage Manager

To configure BR*Tools for use with the RMAN Tivoli Storage Manager channel, complete the following steps.

On the Tivoli Storage Manager server, complete the following steps:

1. Define a policy domain with two management classes that are used to transfer data and logs. Define an archive management class within each of the management classes. If the retention control is performed at the Tivoli Storage Manager server, specify **`RETVER=<days>`** for each archive copy group. If the retention control is performed at Tivoli Storage Manager for ERP, specify **`RETVER=nolimit`**.
2. Register the Tivoli Storage Manager node with the defined domain. Update the parameter **`MAXNUMMP`** for the Tivoli Storage Manager node to **`MAXNUMMP=2`** (based on the parallelism that is required).

On the client node, complete the following steps:

1. Update or create the DSM.OPT and DSM.SYS files to configure the Tivoli Storage Manager API client. The **PASSWORDACCESS** parameter must be set to “PROMPT” in this configuration.
2. Set up the environment values **DSMI_DIR** and **DSMI_LOG** for the Oracle OS user.
3. Install IBM Tivoli Storage Manager for ERP - Oracle on the Oracle server with SAP installed on it.
4. Configure the client resources for Oracle server in the IBM Tivoli Storage Manager for ERP configuration file (<ORACLE_HOME>\dbs\init<SID>.utl).
5. Check the defined Tivoli Storage Manager node name and Tivoli Storage Manager management classes to be used for the backup of offline redo log files and data files. Ensure that the **SERVER** parameter refers to an existing stanza in the DSM.SYS file. If the retention control is driven by Tivoli Storage Manager for ERP, set the **MAX_VERSIONS** parameter.
6. Switch to the Oracle instance owner and update the Tivoli Storage Manager node password for Oracle by running the following command:

   ```
   backint -p <ORACLE_HOME>\dbs\init<SID>.utl -f password
   ```

7. Ensure that RMAN can access the Tivoli Storage Manager for ERP API. The following links must exist (be created):
   - **ln -s /usr/tivoli/tsm/tdp_r3/ora/libtdp_r3.<ext>**
   - **/usr/lib/libobk.<ext> ln -s /usr/lib/libobk.<ext>**
   - **$ORACLE_HOME/lib/libobk.<ext>**
8. Instruct BR*Tools to use RMAN by setting the **backup_dev_type** and **rman_parms** options in the SAP initialization file (init<SID>.sap) as follows:
   - **backup_dev_type = rman_util**
   - **rman_parms="ENV=(XINT_PROFILE=<ORACLE_HOME>/dbs/init<SID>.utl,PROLE_PORT=<portnumber>,&BR_INFO)"**
9. Instruct BR*Tools to use the init<SID>.utl file for Tivoli Storage Manager specific parameters by setting the **util_par_file** parameter in the SAP initialization file:

   ```
   util_par_file=<path to Tivoli Storage Manager for ERP util file -
   init<SID>.utl>
   ```

20.6.8 SAP BR*Tools for Oracle: Using RMAN to configure DB2 to use Tivoli Storage Manager

To configure DB2 to use Tivoli Storage Manager for ERP, complete the following steps.

On the Tivoli Storage Manager server, complete the following steps:

1. Define a policy domain with two management classes that are used to transfer data and logs. Define an archive copy group for both management classes. If the retention control is performed at the Tivoli Storage Manager server, specify **RETVER=<days>** for each archive copy group. If the retention control is performed at Tivoli Storage Manager for ERP level, specify **RETVER=nolimit**.
2. Register Tivoli Storage Manager node with the defined domain. Update the parameter **MAXNUMMP** for Tivoli Storage Manager node to **MAXNUMMP=2** (based on the parallelism that is required).

On the client node, complete the following steps:

1. Update or create the Tivoli Storage Manager API client option files `DSM.OPT` and `DSM.SYS`. The **`PASSWORDACCESS=GENERATE`** parameter must be set for this configuration.
2. Configure the environment values *DSMI_DIR*, **`DSMI_CONFIG`**, and **`DSMI_LOG`** in the DB2 instance owner user's profile. You must restart the DB2 instance to make the parameters effective for DB2.
3. Install Tivoli Storage Manager for ERP - DB2 on the DB2 UDB server, with SAP already installed. You can use the installation wizard to specify the name of the Tivoli Storage Manager server stanza (in `DSM.SYS`), the Tivoli Storage Manager node name, and the management classes to be used for the backup of data and archived logs.
4. Check the client resource for the Tivoli Storage Manager server in the Tivoli Storage Manager for ERP configuration file `/db2/<SID>/tdp_r3/init<SID>.utl`. Verify that the following environment variables are set correctly in the DB2 owner user's profile:
   - **`XINT_PROFILE`**
   - **`DB2_VENDOR_LIB`**
   - **`TDP_DIR`**
5. Switch to the DB2 instance owner and update the Tivoli Storage Manager client password for DB2 node by running the following command:

   ```
   $/usr/tivoli/tsm/tdp_r3/db264/backom -c password
   ```
6. Restart the DB2 instance.
7. Optionally, you can set up DB2 automatic log management so that the archived logs are sent to Tivoli Storage Manager by using the Tivoli Storage Manager media management library that is provided by Tivoli Storage Manager for ERP. This task can be accomplished by setting the DB2 configuration parameters **`LOGARCHMETH1`** and **`LOGARCHOPT1`** as follows:
   - **`update db cfg for <SID> using LOGARCHMETH1 VENDOR:/usr/tivoli/tsm/tdp_r3/db264/libtdpdb264.a`**
   - **`update db cfg for <SID> using LOGARCHOPT1 /db2/<SID>/tdp_r3/vendor.env`**
8. If you use the direct log backup method that is specified in step 7, you should also specify the **`FAILARCHPATH`** db2 configuration parameter. **`FAILARCHPATH`** points to a directory that is used as a temporary storage for offline logs in case that the Tivoli Storage Manager server is unavailable, which can prevent the DB2 from filling up the log directory. Here is the command syntax:

   ```
   update db cfg for <SID> using FAILARCHPATH <offline log path>
   ```

20.6.9 Best practices for Tivoli Storage Manager for ERP with ProtecTIER

The configuration profile of Tivoli Storage Manager for ERP is defined in the `init<SID>.utl` file, which is in the same directory as the BR*Tools profile (`init<SID>.sap`).

When the ProtecTIER VTL is defined for Tivoli Storage Manager, there are some settings to be done in Tivoli Storage Manager for ERP to optimize this integration.

Set the following settings in the `init<SID>.utl` file:

- Disable multiplexing. The **`MULTIPLEXING`** parameter specifies how many files are read simultaneously and are multiplexed. If a file is multiplexed, it can affect the deduplication ratio. Set **`MULTIPLEXING=1`**.

- Use as many backup sessions in parallel as possible. The **MAX_SESSIONS** parameter defines the number of parallel sessions to be established. The valid range of **MAX_SESSIONS** is 1 - 32. You should also define the **SESSIONS** parameter in each Tivoli Storage Manager stanza in the .utl file to define the maximum number of sessions in that Tivoli Storage Manager server stanza.

  Important: The **MAX_SESSIONS** parameter setting must not exceed the number of tape drives that are available simultaneously to the node in the Tivoli Storage Manager servers to be accessed. This maximum is established by the **MAXNUMMP** parameter settings in the Tivoli Storage Manager node definition.

- Disable compression by configuring **RL_COMPRESSION=NO**. The **RL_COMPRESSION** parameter specifies whether a null block compression of the data should be performed before transmission to Tivoli Storage Manager. Although **RL_COMPRESSION** introduces additional processor load to the SAP server, throughput can be improved when the network is the bottleneck, but it can affect the ProtecTIER deduplication ratio.

On the Tivoli Storage Manager server, complete the following steps:

1. Update the **MAXNUMMP** parameter for the Tivoli Storage Manager node to **MAXNUMMP=*x***, where x should be the number of parallels required. This number should match the **MAXSESSION** parameter that is set in the .utl file. The **MAXNUMMP** parameter specifies the maximum number of mount points a node may use on the server only for operations, such as backup and archive.
2. Update the **COMPression** parameter for the Tivoli Storage Manager node to **COMPression=NO**. This setting specifies that the client node does not compress its files before it sends them to the server for backup and archive.

20.7 VMware

In addition to the now available vStorage APIs, the vStorage APIs for Data Protection (VADP) are also available. VADP replaces the VMware Consolidated Backup (VCB) framework, and offers multiple methods to improve your VMware backup. With the new VADP comes the option to use incremental virtual machine image backups by using the changed block tracking (CBT) feature. In contrast to the full virtual machine image backup, CBT reduces the amount of backed up data because only the changed blocks that are compared to the last full backup are backed up. With CBT enabled, the backup operation backs up only the changed blocks, which results in a high data change rate for the ProtecTIER server, because only new data is backed up. For ProtecTIER deduplication to perform optimally, run at least one full backup per week.

Incremental backups: If you use incremental virtual machine image backups, run at least one full virtual machine image backup per week to optimize your deduplication ratio.

Follow the general best practices and the Tivoli Storage Manager best practices that are described in Chapter 14, “IBM Tivoli Storage Manager” on page 219.

20.7.1 Technical overview

VMware ESX is installed directly on the hardware and does not require any specific operating system. It is a virtualization platform that is used to create the virtual machines (VMs) as a set of configuration and disk files that perform all the functions of a physical machine.

vCenter Server

The vCenter server is a service that acts as a central administration point for ESX hosts that are connected to a network. This service directs actions on the virtual machines and the hosts. The vCenter server is the working core of the vCenter.

Multiple vCenter servers can be joined to a linked mode group, where you can log on to any single vCenter server to view and manage the inventories of all the vCenter server systems in the group.

With vCenter, an administrator can manage every component of a virtual environment. ESX servers, VMs, and extended functions, such Distributed Resource Scheduler (DRS), vMotion, and VM backup, all access the vCenter server by using the vSphere Client GUI.

20.7.2 Settings and tuning for VMware and Tivoli Storage Manager

When you set up VMware for Tivoli Storage Manager, there are guidelines for using the vStorage API, changed block tracking (CBT), and format specification for formats for virtual disk files. This section provides a brief description of those guidelines.

vStorage API

The vStorage application programming interfaces (APIs) for data protection enable backup software to protect system, application, and user data in your virtual machines in a simple and scalable way. These APIs enable backup software to perform the following actions:

- Perform full, differential, and incremental image backup and restore of virtual machines
- Perform file-level backup of virtual machines using supported Windows and Linux operating systems
- Ensure data consistency by using Microsoft Volume Shadow Copy Services (VSS) for virtual machines that run supported Microsoft Windows operating systems

Changed block tracking

Virtual machines that run on ESX/ESXi hosts can track disk sectors that change. This feature is called changed block tracking (CBT). On many file systems, CBT identifies the disk sectors that are altered between two change set IDs. On VMFS partitions, CBT can also identify all the disk sectors in use. CBT is useful when you set up incremental backups.

Virtual disk formats

When you perform certain virtual machine management operations (such as creating a virtual disk, cloning a virtual machine to a template, or migrating a virtual machine), you can specify a format for the virtual disk file. However, you cannot specify the disk format if the disk is on an NFS data store. The NFS server determines the allocation policy for the disk. The disk formats listed in this section are supported.

Thick format

The default virtual disk format. The thick virtual disk does not change its size, and from the beginning occupies the entire data storage space that is provisioned to it. Thick format does not zero the blocks in the allocated space.

Conversion: It is not possible to convert the thick disk format in to thin provisioned format.

Thin provisioned format

Use this format to save storage space. For the thin provisioned format, specify as much data storage space as the disk requires based on the value that you enter for the disk size. However, the thin disk starts small and, at first, uses only as much data storage space as the disk needs for its initial operations.

Thin disk considerations: If a virtual disk supports clustering solutions such as fault tolerance, you cannot make the disk thin. If the thin disk needs more space later, it can grow to its maximum capacity and occupy the entire data storage space that is provisioned to it. Also, you can manually convert the thin disk into thick disk.

20.7.3 Backup solutions

This section describes backup solutions and prerequisites for VMware backup by using the ProtecTIER server and Tivoli Storage Manager.

Full VM backup on ProtecTIER

You can use Tivoli Storage Manager Data Protection (DP) for VMware to back up and restore virtual machine data through SAN-based data movement. There are two different data paths where data movement is possible:

- The first data path is from the VMware data store to the vStorage server through SAN.
- The second data path is from the vStorage server to the ProtecTIER server. It could be SAN-based if Tivoli Storage Manager LAN-free is used.

The backup data path uses the *Tivoli Storage Manager for SAN* feature (LAN-free backup). In this book, LAN-free backup is used on VMware.

The DP for VMware stores virtual machine full backup images (full-VM) as a collection of control and data files. The data files contain the contents of virtual machine disk files, and the control files are small metadata files that are used during full VM restore operations and full VM incremental backups. In most cases, VMs are cloned according to a predetermined template. In other words, there is huge duplication of data. The ProtecTIER solution, in conjunction with Tivoli Storage Manager, deduplicates such data.

Prerequisites to VMware backup using ProtecTIER and Tivoli Storage Manager

Check that the following items are complete before you use VMware to back up your data with the ProtecTIER product and Tivoli Storage Manager:

- ☐ The ProtecTIER repository exists.
- ☐ The ProtecTIER deduplication function is enabled.
- ☐ The Tivoli Storage Manager server is installed with a license.
- ☐ The Tivoli Storage Manager storage agent is installed on a vStorage server, if LAN-free is used.
- ☐ The Tivoli Storage Manager Backup-Archive client is installed on the vStorage server.

Tip: For the best performance, the vStorage server must have separate HBA ports and each port must be connected to the ProtecTIER repository and the disk subsystem that stores the VMware data store.

Prerequisites for ESX

Check that the following items are complete before you use VMware ESX with the ProtecTIER repository and Tivoli Storage Manager:

- ☐ The host must be running ESX/ESXi Version 4.0 or later.
- ☐ The VMware vCenter Server must be Version 4.1.x or later.
- ☐ For incremental backup, the virtual machine that owns the disks to be tracked must be hardware Version 7 or later.
- ☐ CBT must be enabled for the virtual machine. (In the vSphere client, click **Edit** → **Settings** → **Options** → **Advanced/General** → **Configuration Parameters**.)
- ☐ The configuration of the virtual machine (.vmx) file must contain the following entry:

  ```
  ctkEnabled = "TRUE"
  ```

- ☐ For each virtual disk, the .vmx file must contain the following entry:

  ```
  scsix:x.ctkEnabled = "TRUE"
  ```

- ☐ For each virtual disk and snapshot disk, there must be a .ctk file (Example 20-8).

Example 20-8 Both the virtual disk and snapshot disk have an associated .ctk file

```
vmname.vmdk
vmname-flat.vmdk
vmname-ctk.vmdk
vmname-000001.vmdk
vmname-000001-delta.vmdk
vmname-000001-ctk.vmdk
```

VMware topology

Check that the following topology (Figure 20-13) is in place before you use VMware with ProtecTIER and Tivoli Storage Manager.

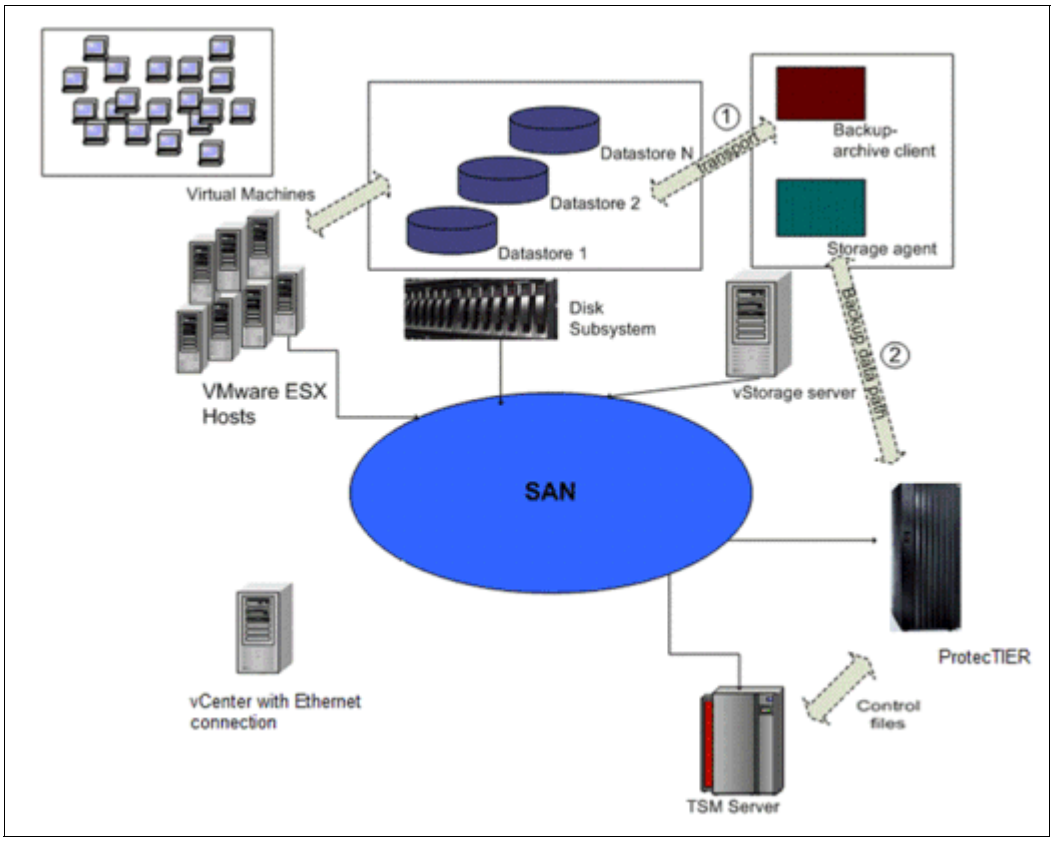

Figure 20-13 Tivoli Storage Manager / VMware topology

Where:

- The Tivoli Storage Manager backup-archive client on the vStorage server must read “GUEST OS data from Disk Subsystem by SAN”.
- The Tivoli Storage Manager Storage agent on the vStorage server must write “GUEST OS data to ProtecTIER by SAN”.
- The Tivoli Storage Manager server writes control data to the ProtecTIER repository through the SAN.

20.7.4 Zoning

This section describes the required fabric zoning for the ProtecTIER repository, the hosts, and Tivoli Storage Manager. For a description of the HBA ports, see Table 20-3, and for a SAN zoning example, see Table 20-4.

Table 20-3 HBA ports

Item	Port
ProtecTIER	PT front-end port_0 PT front-end port_1
Tivoli Storage Manager Storage Agent	Tivoli Storage Manager Storage Agent port_0 Tivoli Storage Manager Storage Agent port_1
Tivoli Storage Manager Server	Tivoli Storage Manager server port_0 Tivoli Storage Manager server port_1
ESX Server	ESX server port_0 ESX server port_1
XIV	XIV_Module4 port_0 XIV_Module5 port_0 XIV_Module6 port_0 XIV_Module4 port_1 XIV_Module5 port_1 XIV_Module6 port_1

Table 20-4 SAN zoning examples

Zone name	Zone members
Zone_XIV_TSM_StorageAgent_0	Tivoli Storage Manager Storage Agent port_0 XIV_Module4 port_0 XIV_Module5 port_0 XIV_Module6 port_0
Zone_XIV_TSM_StorageAgent_1	Tivoli Storage Manager Storage Agent port_1 XIV_Module4 port_1 XIV_Module5 port_1 XIV_Module6 port_1
Zone_PT_TSM_StorageAgent_0	Tivoli Storage Manager Storage Agent port_0 PT front-end port_0
Zone_PT_TSM_StorageAgent_1	Tivoli Storage Manager Storage Agent port_1 PT front-end port_1
Zone_PT_TSM_Server_0	Tivoli Storage Manager server port_0 PT front-end port_0
Zone_PT_TSM_Server_1	Tivoli Storage Manager server port_1 PT front-end port_1
Zone_ESX_XIV_0	Tivoli Storage Manager Storage Agent port_0 XIV Module4 port_0 XIV Module5 port_0 XIV Module6 port_0

Zone name	Zone members
Zone_ESX_XIV_1	Tivoli Storage Manager Storage Agent port_1 XIV Module4 port_1 XIV Module5 port_1 XIV Module6 port_1

20.7.5 Configuring the ProtecTIER server

This section describes all the steps that are required to create and to configure a new VTL in the ProtecTIER server. Also described is the optional procedure to enable and configure LUN masking for the Tivoli Storage Manager server and the Tivoli Storage Manager storage agent (vStorage server in a Tivoli Storage Manager environment).

To create the VTL, complete the following steps:

1. From the **VT** drop-down menu of ProtecTIER Manager, click **VT** → **VT Library** → **Create new library**. The Create new library window opens.
2. Input the name of the library in the VT name field, and press **Next**. The Library type window opens within the Create new library window.
3. Select **IBM TS3500** as the library type to simulate. The Tape model window opens.
4. Select **IBM ULT3580-TD3** as the tape model to simulate. The Port Assignment window opens.
5. Create the robot and the drives (such as one robot and 20 drives). The drives are assigned, crossing all front-end ports to ensure better performance. The Cartridges window opens.
6. Create cartridges that are based on the backup policy (for example, 100 cartridges). The Slots window opens.
7. Create 100 slots and 32 import/export slots by selecting **100** in the **No. of slots** selection box, and **32** in the **Number of import/exports** slots selection box. Click **Next**.

   Important: The creation of the library takes the system offline for a few minutes.

8. (Optional) Enable and configure LUN masking for the Tivoli Storage Manager server and the Tivoli Storage Manager storage agent (vStorage server). If you have multiple backup servers that are connected to the ProtecTIER server, enabling the LUN masking feature is recommended. Enable and configure LUN masking by completing the following steps:
   a. From the expanded list on the left side of the ProtecTIER Manager window, click **VT** → **LUN Masking** → **Enable/Disable LUN masking**. ProtecTIER Manager notifies you that you if you enable the LUN masking feature without configuring LUN masking groups, the devices are hidden from the hosts, and prompts you to confirm whether you want to proceed with this process.
   b. When the Enable/Disable LUN masking dialog box opens, select **Enable LUN masking**, and click **OK**.
   c. From the expanded list on the left side of the ProtecTIER Manager window, click **VT** → **LUN Masking** → **Configure LUN masking groups**. The LUN Masking window opens.
   d. In the “Selected Host Initiators” frame, click **Add**. The Host Initiator Management window opens.

e. Create LUN masking groups for the Tivoli Storage Manager server and the Tivoli Storage Manager storage agent by adding the worldwide port numbers (WWPNs) of the Tivoli Storage Manager server and the Tivoli Storage Manager storage agent in to the list.

f. Select the check box beside each of the added ports, and click **Save Changes**. The LUN Masking window opens.

g. In the Library Mappings frame, click **Add**, and add the library that is called “TSM_VMW” to the library mappings list.

h. Click **Save Changes**.

20.7.6 Installing the tape driver on the Tivoli Storage Manager server and the Tivoli Storage Manager storage agent

This section describes how to install the tape driver on the Tivoli Storage Manager server and the Tivoli Storage Manager storage agent for Windows and Linux based systems.

The Windows Device Manager shows the tape changer as “Unknown Medium Changer” and the tape drives as “IBM ULT3580-TD3 SCSI Sequential Device” (Figure 20-14).

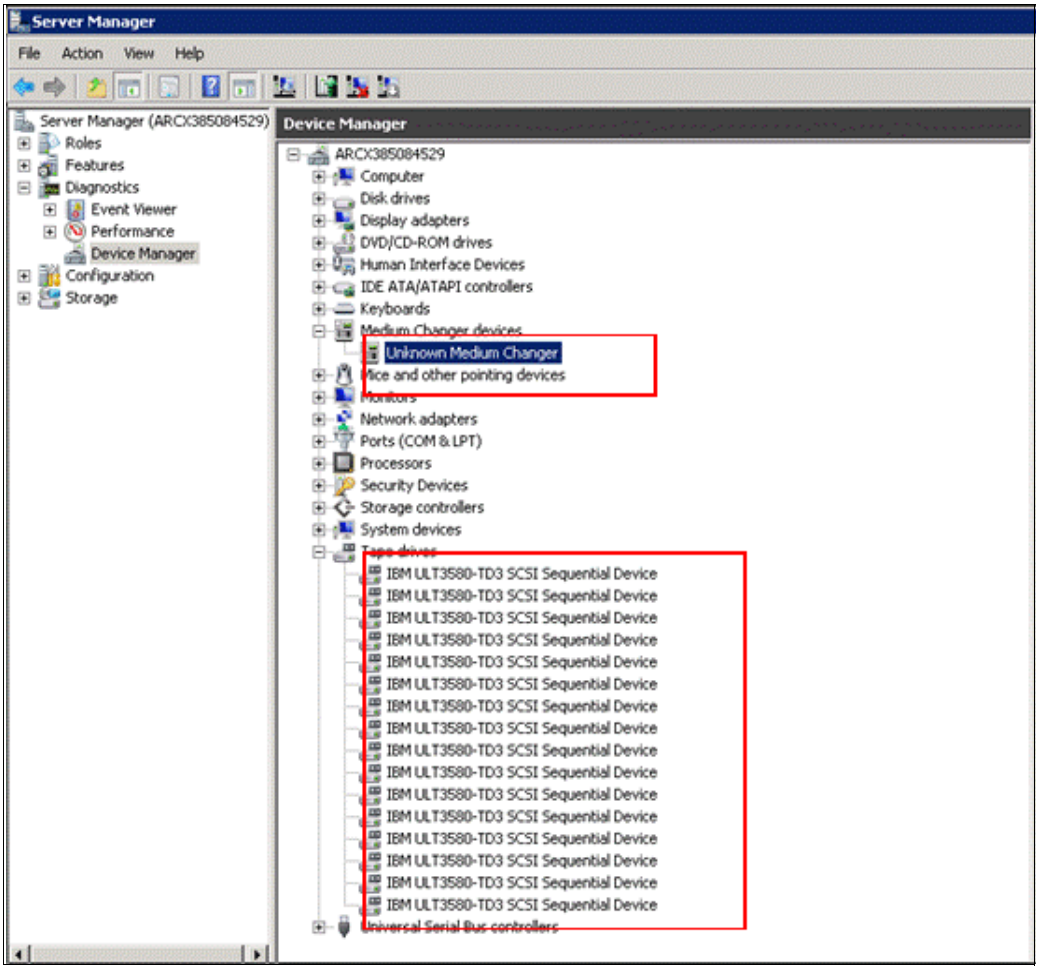

Figure 20-14 Server Manager window

Procedure

To install the tape driver on the Tivoli Storage Manager server and the Tivoli Storage Manager storage agent for Windows and Linux based systems, complete the following steps.

For Windows based systems, complete the following steps:

1. Download the IBM Tape Device Driver from the following website:

   `ftp://ftp.software.ibm.com/storage/devdrvr/FTPSITE_IS_SUNSET.html`

2. Run **`install_exclusive.exe`** (Figure 20-15) to install the IBM Tape Driver for Tivoli Storage Manager. The installation application initiates, and when complete, displays a dialog box that notifies you that the installation was successful.

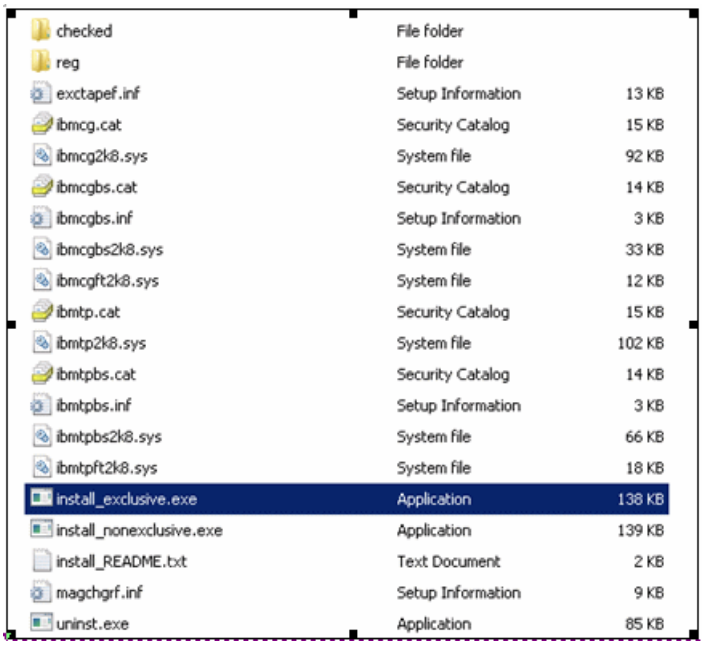

Figure 20-15 Installation program

3. After the installation is complete, the Windows Device Manager shows the tape changer as “IBM 3584 Tape Library”, and the tape drives as “IBM ULTRIUM III 3580 TAPE DRIVE” (Figure 20-16).

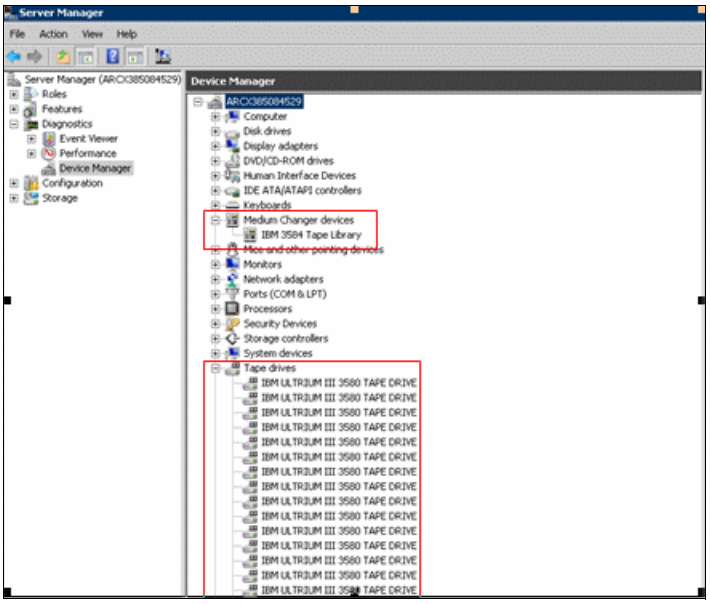

Figure 20-16 Server manager window - showing renamed changer and tape drive

For Linux based systems, download the tape device driver from the following website:

`ftp://ftp.software.ibm.com/storage/devdrvr/FTPSITE_IS_SUNSET.html`

Install the following RPM packages:

- `lin_tape-1.54.0-1`
- `lin_taped-1.54.0-1`

20.7.7 Tivoli Storage Manager storage agent configuration

This section describes the process of configuring the Tivoli Storage Manager storage agent to establish communication with the Tivoli Storage Manager server in a VMware environment. To accomplish this task, complete the following steps:

1. To establish communication for Tivoli Storage Manager server and Tivoli Storage Manager storage agent, run the commands that are shown in Example 20-9 at the CLI of the Tivoli Storage Manager storage agent.

Example 20-9 Commands to establish communication for the Tivoli Storage Manager server and Tivoli Storage Manager storage agent

```
dsmsta.exe setstorageserver
myname=ARCX385084529
mypassword=<user_password>
myhladdress=x.x.110.38
servername=ARCX3650N1332
serverpassword=open1sys
hladdress=x.x.110.65
lladdress=1500
```

2. Disable automatic mounting of volumes on the Tivoli Storage Manager storage agent host by running the following command at the CLI prompt of the Tivoli Storage Manager storage agent:

```
diskpart > automount disable > exit
```

Important: The **diskpart** command is necessary to keep the Tivoli Storage Manager storage agent from damaging the SAN volumes that are used for raw disk mapping (RDM) virtual disks.

3. To enable the Tivoli Storage Manager storage agent to access online GUEST OS data, click **Server manager** → **Storage** → **Disk Management** → **Online** (Figure 20-17).

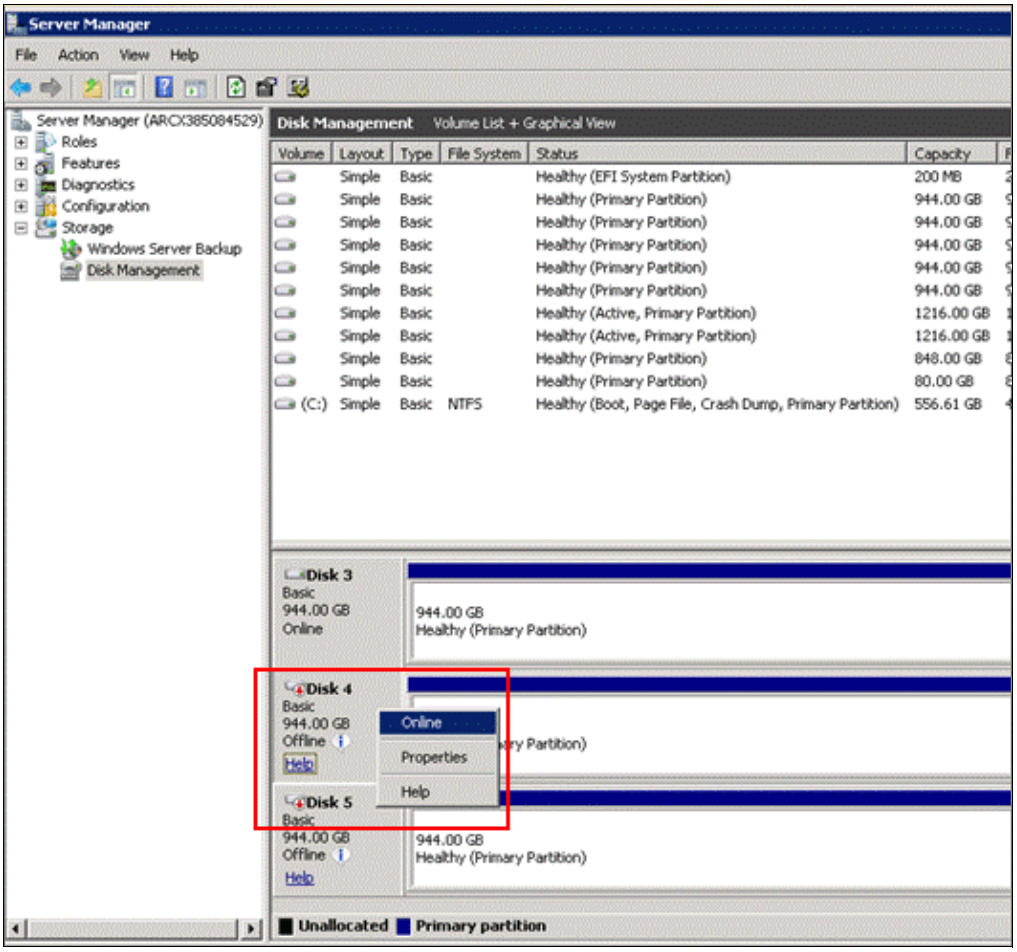

Figure 20-17 Server Manager window that shows the Tivoli Storage Manager storage agent set to online

4. Note the device name and serial number of the Tivoli Storage Manager storage agent (Figure 20-18). You need this information to define the path in the Tivoli Storage Manager server in a later step.

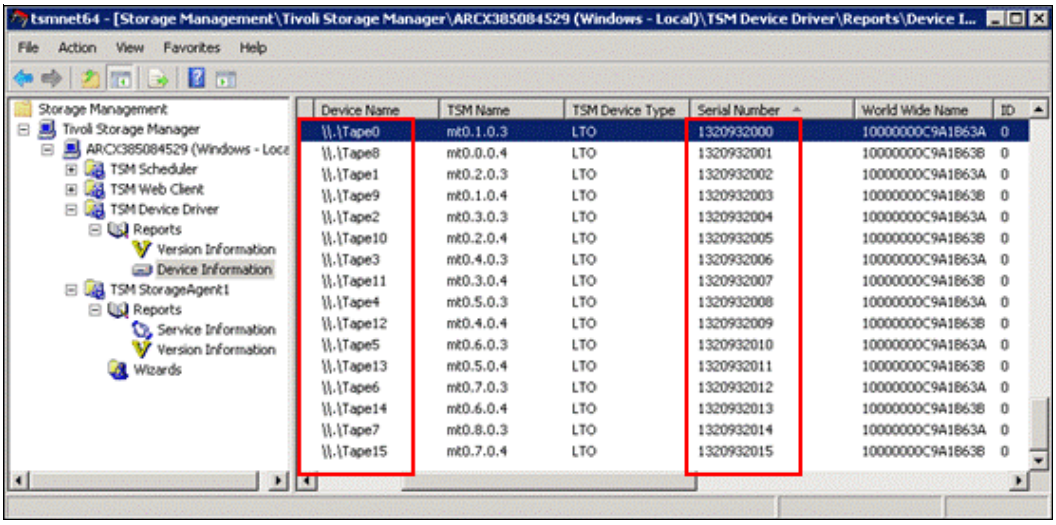

Figure 20-18 Serial numbers of Tivoli Storage Manager storage agent

Persistent naming: The example in Figure 20-18 on page 358 does not show the usage of persistent naming. Follow the guidelines in 7.2.7, “Persistent device naming” on page 112.

If you do not use persistent naming, take a screen capture of the Tivoli Storage Manager management console so that you can have a readily accessible record of the tape device information in the Tivoli Storage Manager storage agent.

20.7.8 Tivoli Storage Manager server configuration

This section describes the procedure to define and to configure the Tivoli Storage Manager server in a VMware environment through the Tivoli Storage Manager server CLI. To accomplish this task, complete the following steps:

1. Define the server for the storage agent by running the following command at the Tivoli Storage Manager server CLI:

   ```
   define server ARCX385084529 SERVERPAssword=admin HLAddress=x.x.110.38
   LLAddress=1500 COMMmethod=TCPIP
   ```

2. Set the server name to Tivoli Storage Manager Server by running the following command:

   ```
   set servername ARCX3650N1332
   ```

3. Set the password by running the following command:

   ```
   set serverpassword admin
   ```

4. Set the Tivoli Storage Manager server IP address by running the following command:

   ```
   set serverhladdress x.xx.xxx.xx
   ```

5. Set the Tivoli Storage Manager server port by running the following command:

   ```
   set serverlladdress 1502
   ```

6. Create the library by running the following command:

   ```
   define library VMW_LIB libtype=scsi autolabel=yes shared=yes RELABELSCRatch=yes
   ```

7. Choose all devices that are related to `tsminst1`.
8. Define a library path from the Tivoli Storage Manager server to the physical OS devices by running the following command:

   ```
   DEFINE PATH ARCX365ON1332 VMW_LIB srctype=server desttype=library
   autodetect=yes device=/dev/IBMchanger0
   ```

9. Define all the drives by running the following commands:
   - **define drive VMW_LIB drive0**
   - **define drive VMW_LIB drive1**
10. Define the drives path from Tivoli Storage Manager server to the physical OS devices by running the following commands:
    - **define path ARCX365ON1332 drive0 srctype=server desttype=drive library=VMW_LIB autodetect=yes device=/dev/IBMtape0**
    - **define path ARCX365ON1332 drive1 srctype=server desttype=drive library=VMW_LIB autodetect=yes device=/dev/IBMtape1**
11. Define the drives path from Tivoli Storage Manager storage agent to the physical OS devices by running the following commands:
    - **define path ARCX385O84529 drive0 srctype=server desttype=drive library=VMW_LIB**

      autodetect=yes device=\\.\Tape0
    - **define path ARCX385O84529 drive1 srctype=server desttype=drive library=VMW_LIB**

      autodetect=yes device=\\.\Tape1

    Important: Ensure that the device on the Tivoli Storage Manager server and the device on the Tivoli Storage Manager storage agent, which are mapped to same drive path, have the same serial number. They are essentially the same device. (For the serial numbers in the Tivoli Storage Manager storage agent, see your notes from step 4 on page 358.)

12. Query the drive and verify that it has a status of `online` by running the following command:

    ```
    query drive
    ```

13. Check in and label the cartridges by running the following command:

    ```
    label LIBVOL VMW_LIB search=yes labelsource=barcode CHECKIN=scratch
    overwrite=yes waitt=0
    ```

14. Define a device class by running the following command:

    ```
    define devclass LTOCLASS3 library=VMW_LIB devtype=lto format=ULTRIUM3C
    ```

15. Define a storage pool by running the following command:

    ```
    define stgpool VMW_POOL LTOCLASS3 pooltype=primary maxscratch=99999
    ```

16. Define a domain by running the following command:

    ```
    DEFine DOmain VMW_DOMAIN BACKRETention=60 ARCHRETention=365
    ```

17. Define a policy set by running the following command:

    ```
    DEFine POlicyset VMW_DOMAIN VMW_POLICY
    ```

18. Define a management class by running the following command:

```
DEFine MGmtclass VMW_DOMAIN VMW_POLICY LTOCLASS3
```

19. Define a copy group by running the following command:

```
DEFine COpygroup VMW_DOMAIN VMW_POLICY LTOCLASS3 DESTination=VMW_POOL
```

20. Define an archive copy group by running the following command:

```
DEFine COpygroup VMW_DOMAIN VMW_POLICY LTOCLASS3 Type=Archive
DESTination=VMW_POOL
```

21. Assign the default management class by running the following command:

```
ASsign DEFMGmtclass VMW_DOMAIN VMW_POLICY LTOCLASS3
```

22. Activate the policy set by running the following command:

```
ACTivate POlicyset VMW_DOMAIN VMW_POLICY
```

23. Register the node for Tivoli Storage Manager BAC by running the following command:

```
register node ARCX385084529 admin passexp=0 userid=admin
domain=VMW_DOMAINcompression=no type=client DATAWritepath=lanfree
DATAReadpath=lanfree hladdress=x.x.110.38 lladdress=1502 archdelete=yes
backdelete=yes maxnummp=999
```

Compression value: Specify the value of compression as `no`. For LAN-free backup, specify the value of **`DATAWritepath`** and **`DATAReadpath`** as `lanfree`.

24. Give the appropriate permissions to the administrator by running the following command:

```
GRant AUTHority admin CLasses=SYstem
```

20.7.9 Tivoli Storage Manager client installation

This section describes the procedure to install and to configure the Tivoli Storage Manager client in a VMware environment by using the IBM Tivoli Storage Management installation program.

To accomplish this task, complete the following steps:

1. From the IBM Tivoli Storage Management installation program (Figure 20-19), select all the components to install, and click **Next**.

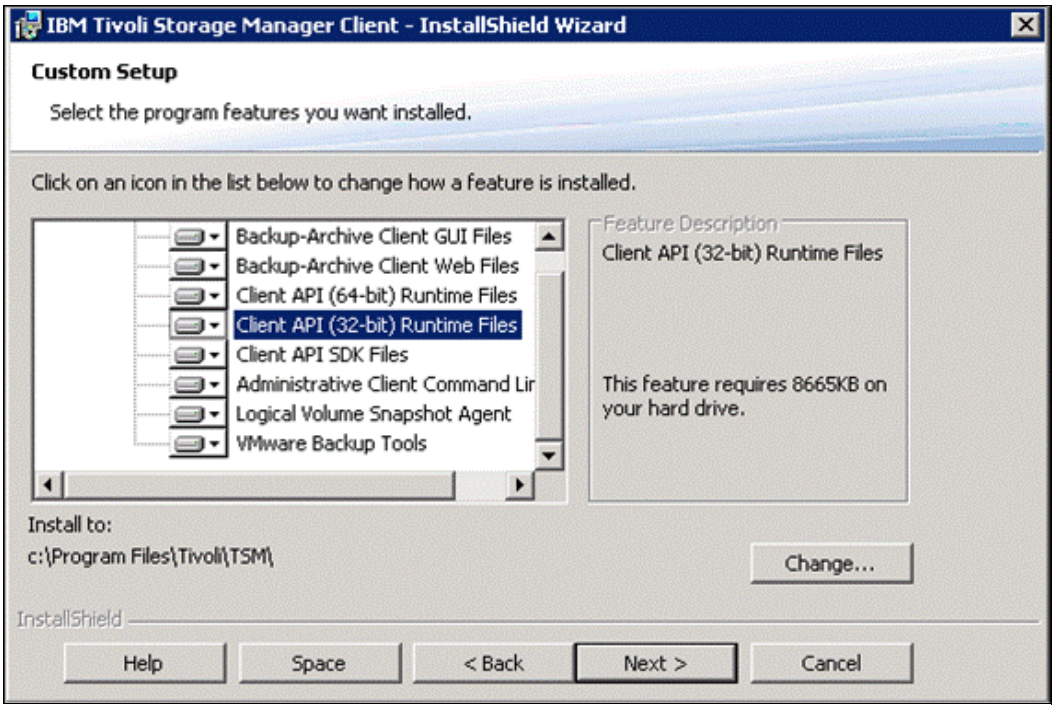

Figure 20-19 Tivoli Storage Manager installation program

Important: Ensure that VMware Backup Tools are installed.

2. When you are prompted to select the type of installation, click **Complete** and finish the installation.

20.7.10 Disabling compression and deduplication on Tivoli Storage Manager

This section describes the procedure to disable compression and deduplication in Tivoli Storage Manager by using the Tivoli Storage Manager GUI.

To accomplish this task, complete the following steps:

1. From the Tivoli Storage Manager GUI drop-down menus, click **Edit** → **Client Preferences** (Figure 20-20). The Client Preferences window opens.

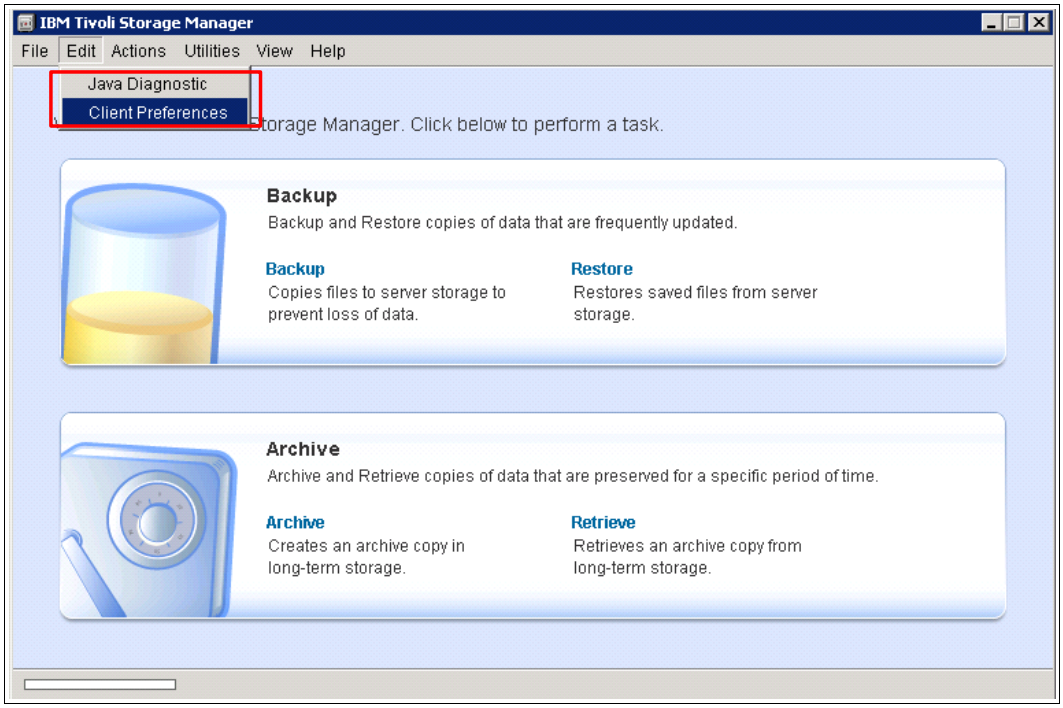

Figure 20-20 Tivoli Storage Manager GUI with Client Preferences selected

2. From the menu at the left, choose **Deduplication** (Figure 20-21). The Deduplication Preferences window opens.

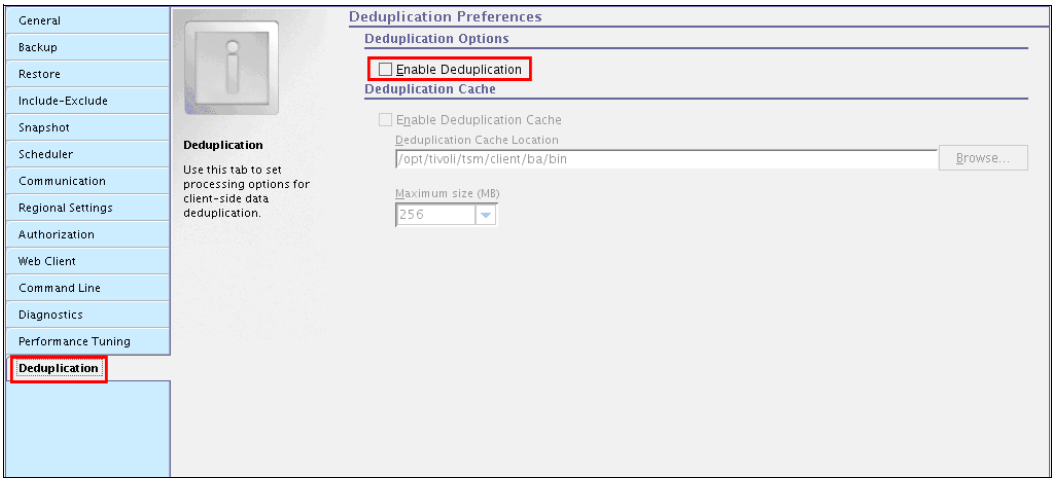

Figure 20-21 Deduplication Preferences window

3. Disable Tivoli Storage Manager deduplication by clearing the **Enable Deduplication** check box.

Note: Tivoli Storage Manager is able to combine both compression and deduplication within itself. The details are explained in Chapter 4, “Introduction to IBM Tivoli Storage Manager deduplication”, in Implementing *Implementing IBM Storage Data Deduplication Solutions*, SG24-7888.

4. From the menu at the left, click **Backup** (Figure 20-22). The Backup Preferences window opens.

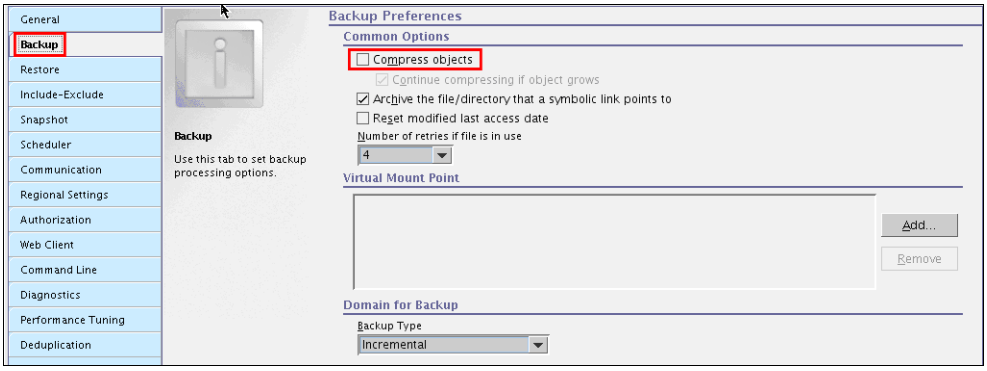

Figure 20-22 Backup Preferences window

5. Disable Tivoli Storage Manager compression by clearing the **Compress Objects** check box.

20.7.11 Configuring a full VM backup through the vStorage API

This section describes the procedure to configure a full VMware backup from the Tivoli Storage Manager GUI. To accomplish this task, complete the following steps:

1. From the Tivoli Storage Manager GUI drop-down menus, click **Edit** → **Client Preferences** The Client Preferences window opens.

2. From the menu at the left, click **VM Backup** (Figure 20-23).

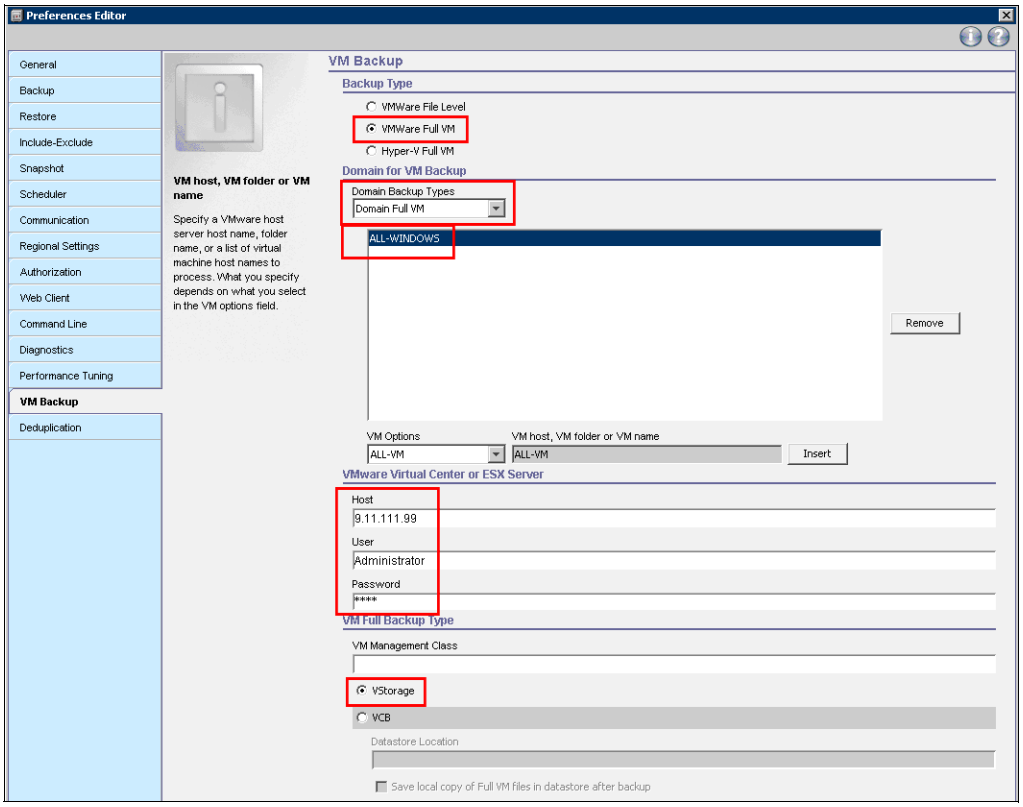

Figure 20-23 VM Backup window

3. Under Backup Type, select **VMware Full VM**.
4. In the Domain for VM Backup selection box, select **Domain Full VM**.
5. In the VM Options selection box, select **All-VM** (If you are using Windows Guest OS, select **ALL-WINDOWS**) and click **Insert**.
6. Enter the IP address, user name, and password of the vCenter.
7. Under VM Management Class, select **VStorage**.

20.7.12 VMware Guest OS backup to ProtecTIER

Use one of the following methods to implement VMware Guest OS backup to the ProtecTIER server.

Backing up VM by using the Tivoli Storage Manager GUI

To accomplish this task, complete the following steps:

1. From the Tivoli Storage Manager GUI drop-down menus, click **Actions** → **Backup VM** (Figure 20-24).

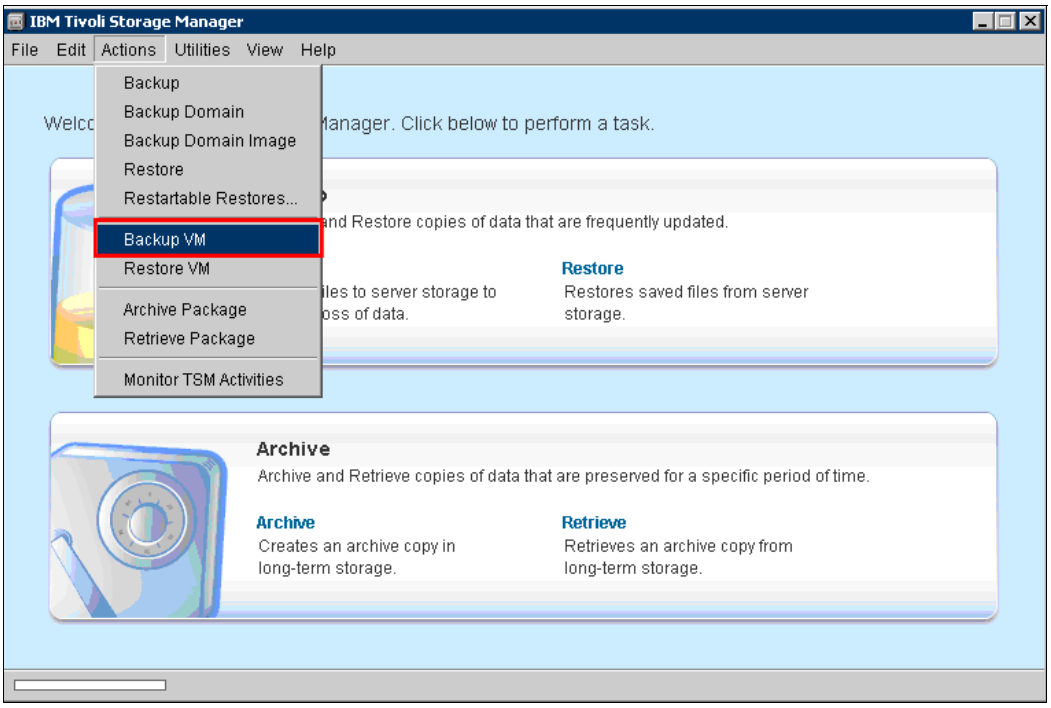

Figure 20-24 Tivoli Storage Manager GUI showing the Backup VM menu cascade

2. The Backup Virtual Machine window opens (Figure 20-25). In the Backup selection box, select the backup type: **VMware Full VM (vStorage)** or **VMware Full VM (incremental)**.

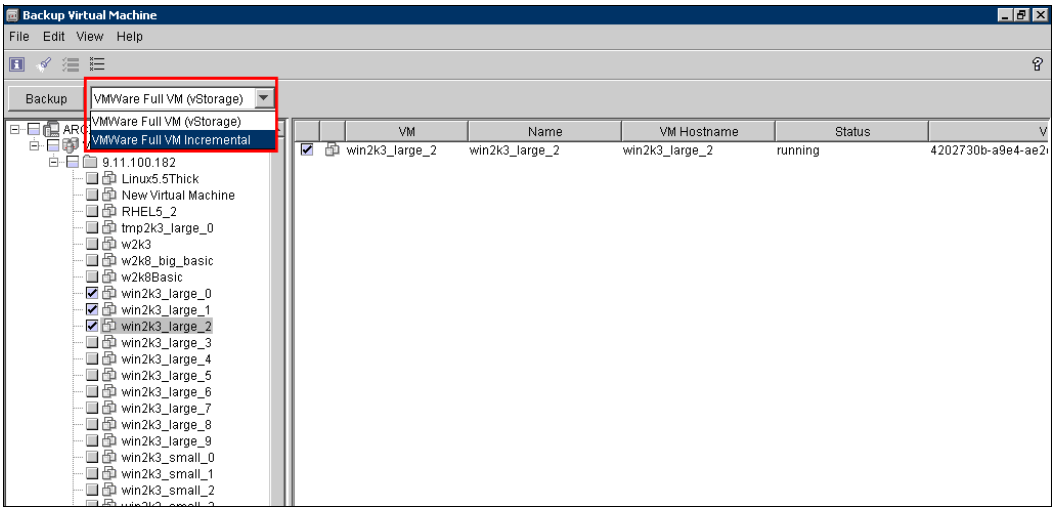

Figure 20-25 Choose the type of backup

Full VM backup using the command-line interface

To accomplish this task, start a full VM backup by using the **mode=full** backup parameter. Tivoli Storage Manager backs up all of the inspected data, and the value of the total destruction ratio is 0%. Run the following command:

```
PS C:\Program Files\Tivoli\tsm\baclient> ./dsmc backup vm win2k3_large_3
-mode=full -vmbackuptype=fullvm
```

The system displays output similar to Example 20-10.

Example 20-10 Output log - full backup

```
Total number of objects inspected:1
Total number of objects backed up:1
Total number of objects updated:0
Total number of objects rebound: 0
Total number of objects deleted: 0
Total number of objects expired: 0
Total number of objects failed: 0
Total number of subfile objects:0
Total number of bytes inspected:15.00 GB
Total number of bytes transferred:15.00 GB
LanFree data bytes:15.00 GB
Data transfer time:842.86 sec
Network data transfer rate:18,661.04 KB/sec
Aggregate data transfer rate:12,541.90 KB/sec
Objects compressed by:0%
Total data reduction ratio:0.00%
Subfile objects reduced by:0%
Elapsed processing time:00:20:54
```

Incremental VM backup by using the command-line interface

To accomplish this task, start an incremental VM backup by using the **mode=incremental** backup parameter. Tivoli Storage Manager backs up only the changed data that is found by VMware CBT, so the value of the data deduction ratio is 99.77%. Enter the following command:

```
PS C:\Program Files\Tivoli\TSM\baclient> ./dsmc backup vm win2k3_large_3
-mode=incremental -vmbackuptype=fullvm
```

The system displays output similar to Example 20-11.

Example 20-11 Output log - full backup

```
Total number of objects inspected:1
Total number of objects backed up:1
Total number of objects updated:0
Total number of objects rebound: 0
Total number of objects deleted: 0
Total number of objects expired: 0
Total number of objects failed: 0
Total number of subfile objects:0
Total number of bytes inspected:15.00 GB
Total number of bytes transferred:38.63 GB
LanFree data bytes:38.63 GB
Data transfer time:2.77 sec
Network data transfer rate:13,461.05 KB/sec
```

```
Aggregate data transfer rate:2,689.30 KB/sec
Objects compressed by:0%
Total data reduction ratio:99.77%
Subfile objects reduced by:0%
Elapsed processing time:00:00:13
```

Restoring VM by using the Tivoli Storage Manager GUI

To accomplish this task, complete the following steps:

1. From the Tivoli Storage Manager GUI drop-down menus, click **Actions** → **Restore VM** (Figure 20-26).

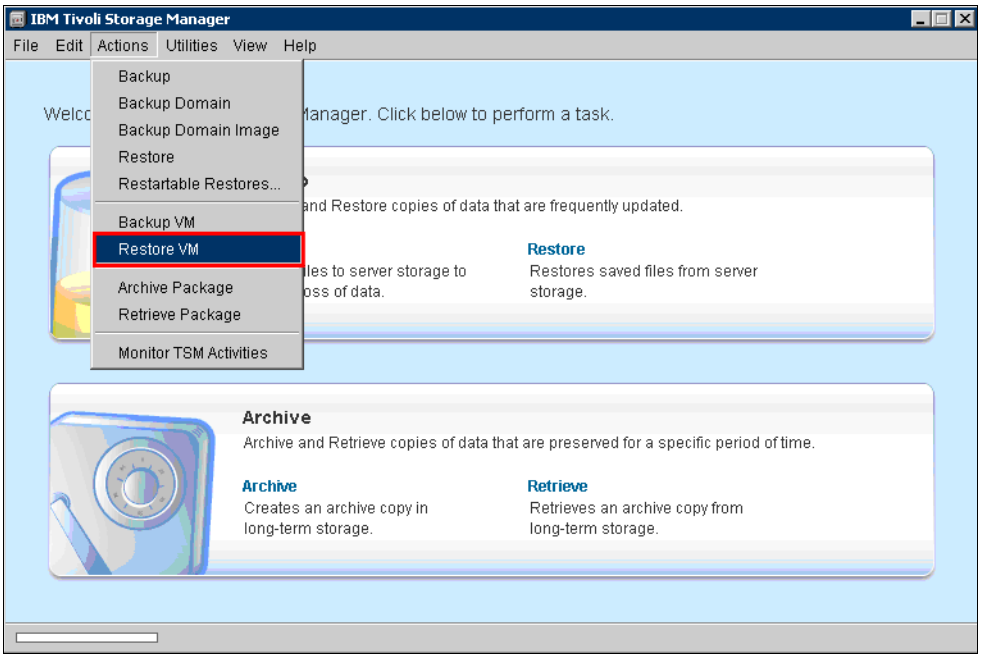

Figure 20-26 Tivoli Storage Manager GUI with Restore VM menu cascade

2. The Restore Virtual Machine window opens (Figure 20-27). Select the version to restored, either full backup (**FULL**) or incremental backup (**INCR**).

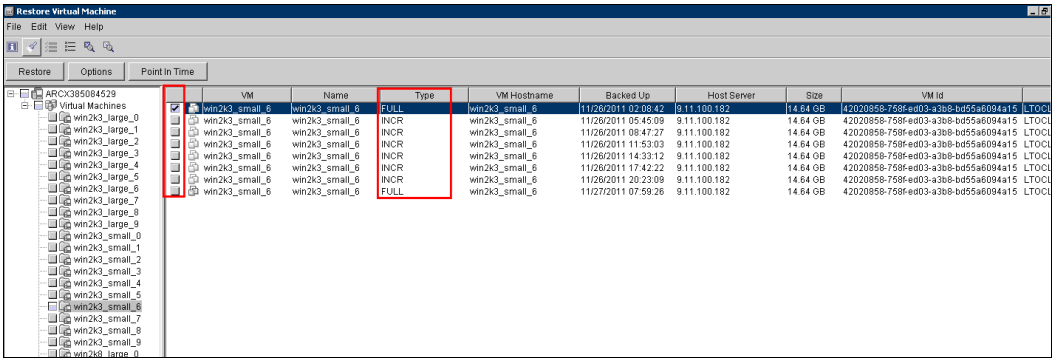

Figure 20-27 Restore Virtual Machine window

3. Tivoli Storage Manager now prompts you to select whether to restore to the original location or to a new location (Figure 20-28). If you choose to restore to a new location, enter the following information, and click **Restore**:
   - Name: The Guest OS name that is managed by vCenter.
   - Datacenter: The name of the data center that stores the new Guest OS.
   - Host: The IP of the ESX server that stores the new Guest OS.
   - Datastore: The name of data store that stores the new Guest OS.

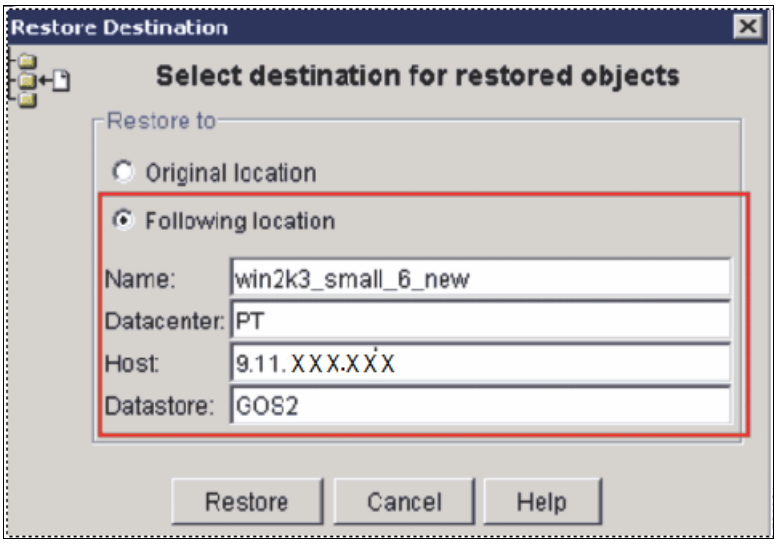

Figure 20-28 Restore Destination window

Part 5

Replication and disaster recovery

ProtecTIER with replication enables virtual tape cartridges to be replicated from multiple primary sites (spokes) to a central secondary location (hub) for enhanced disaster recovery (DR) and business continuity (BC) capabilities. This part describes replication, including such concepts as replication deployment, solution optimization, and the procedures to deploy replication to work with specific backup applications.

This part describes the following topics:

- ProtecTIER replication
- Disaster recovery deployment with backup applications

21

ProtecTIER replication

IT organizations that use a ProtecTIER system with replication can easily expand the coverage of that replication to all of the applications in their environment. You can create replication policies to set rules for replicating data objects across ProtecTIER repositories. This chapter describes the purpose of replication and the enhanced features of the latest ProtecTIER code.

This chapter describes the procedures that are used for replication deployment, including preinstallation steps, creation of the replication grid, and synchronization of the primary and secondary repositories. There is also a section on upgrading the existing system and enabling replication.

In addition, this chapter provides the basic rules and guidance of replication deployment for The ProtecTIER product in environments with OpenStorage (OST), Virtual Tape Library (VTL), and File System Interface (FSI). It also describes the concepts, procedures, and considerations that are related to optimizing replication performance, including the procedures to automate and script the daily operations.

This chapter primarily focuses on best practices for planning, configuration, operation, and testing of ProtecTIER native replication.

The concept, detailed planning, and implementation of native replication is described in *IBM System Storage TS7600 with ProtecTIER Version 3.3*, SG24-7968.

This chapter covers the following topics:

- ProtecTIER IP replication
- Native replication
- Replication policies
- Visibility switching
- Principality
- Replication Manager
- Initial synchronization
- Replication schedules
- Replication backlog
- Replication planning
- Bandwidth validation utility

- Planning ProtecTIER replication
- The backup application database backup
- ProtecTIER Planner tool

21.1 ProtecTIER IP replication

The ProtecTIER IP replication function (Figure 21-1) provides a powerful tool that you can use to design robust disaster recovery architectures. You electronically place backup data into vaults with much less network bandwidth, thus changing the paradigm of how data is taken off-site for safe keeping. The ProtecTIER IP replication feature can eliminate some of the expensive and labor-intensive handling, transport, and securing of the real tapes for disaster recovery purposes.

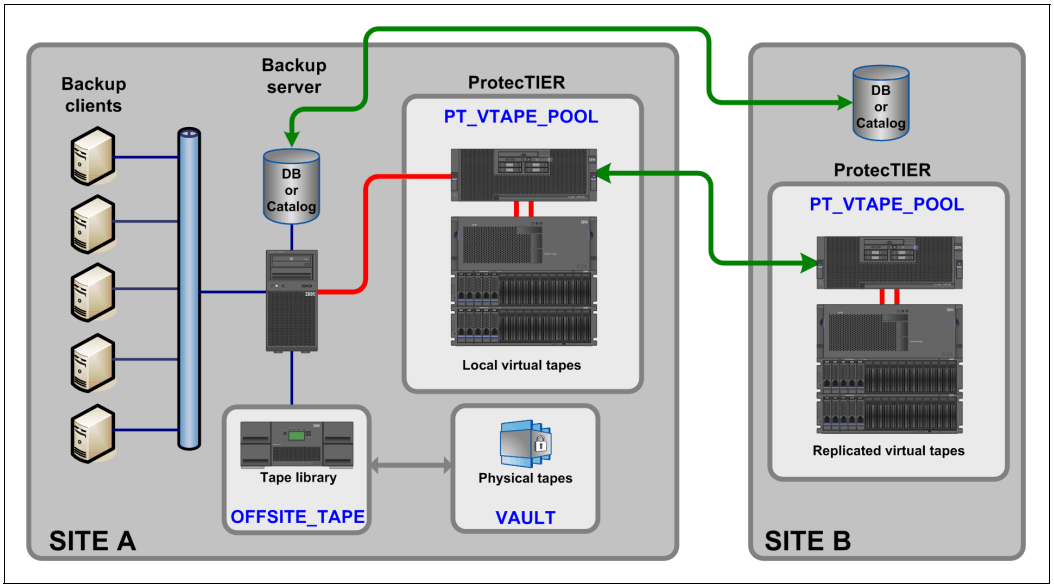

Figure 21-1 IP replication in a backup and recovery environment

Figure 21-1 illustrates how the ProtecTIER IP replication function can be used in a backup and recovery environment. This particular client is using this feature to replicate all of the virtual tapes in ONSITE_VTAPE_POOL off-site. It also backs up all backup application databases or catalogs to virtual tapes. These database backup virtual tapes are also replicated to Site B.

If there is a disaster, you can restore the backup server environment on site B, which is connected to a ProtecTIER VTL. It contains the backup application database (catalog) together with all of the client backup files on virtual tapes in the PT_VTAPE_POOL pool.

21.2 Native replication

ProtecTIER replication enables data replication capability across repositories, among ProtecTIER systems, which are connected to the wide area network (WAN). Because the ProtecTIER product deduplicates data before storing it, only the changes, or unique elements of data, are transferred to the DR site over the replication link. This feature can translate into substantial savings in the bandwidth that is needed for the replication TCP/IP link.

In early versions of ProtecTIER, the repository replication was handled by the disk array subsystems. Starting with Version 2.4, ProtecTIER introduced the function that is known as “native replication”, where the replication of deduplicated data became a function of ProtecTIER. Deduplicated data is replicated to a secondary ProtecTIER system through TCP/IP rather than relying on the back-end disk arrays and their associated infrastructure.

21.2.1 One-to-one replication

The initial replication design consisted of two ProtecTIER systems with one system that is designated as the source and the other system that is designated as the target. The target system (or hub) was dedicated to receiving incoming replicated data and was not eligible to take local backups.

21.2.2 Many-to-one replication

ProtecTIER Version 2.4 expanded the native replication functionality and introduced the many-to-one replication grid. Also known as *spoke and hub*, up to 12 source systems (spokes) can all replicate to a single target ProtecTIER system (hub) simultaneously. The hub system can provide disaster recovery (DR) functionality for one or *more* spokes concurrently, and the hub system can accept and deduplicate local backup data. The hub system cannot replicate outgoing data.

21.2.3 Many-to-many replication

ProtecTIER Version 3.1 built upon existing replication technology and introduced the many-to-many bidirectional replication grid. Up to four systems (all hubs) could accept and deduplicate local backup data, replicate that data to up to three other ProtecTIER systems, and receive incoming replicated data from up to three other ProtecTIER systems.

21.2.4 VTL replication

One-to-One and many-to-one replication implies VTL replication. The ProtecTIER VTL service emulates traditional tape libraries. Your existing backup application can access virtual robots to move virtual cartridges between virtual slots and drives. The backup application perceives that the data is being stored on cartridges while the ProtecTIER product stores data on a deduplicated disk repository. In a VTL replication scenario, data is replicated at the virtual tape cartridge level.

21.2.5 OST replication

Many-to-many replication encompasses VTL replication, but also includes both OST replication and FSI replication. In an OST topology group, up to 12 ProtecTIER OST systems can replicate backup images with multiple target systems, bidirectional replication, and cascading of replicated backup images. ProtecTIER performs the replication of deduplicated data while the replication policies are defined and run by the NetBackup application. Data is replicated at the NetBackup backup image level.

21.2.6 FSI replication

With FSI replication, up to eight ProtecTIER FSI systems can be included in the bidirectional replication group (Figure 21-2). Each FSI system can replicate deduplicated data to as many as three other remote ProtecTIER FSI systems. Data is replicated at the file system level with a maximum of 128 file systems on a single ProtecTIER FSI system.

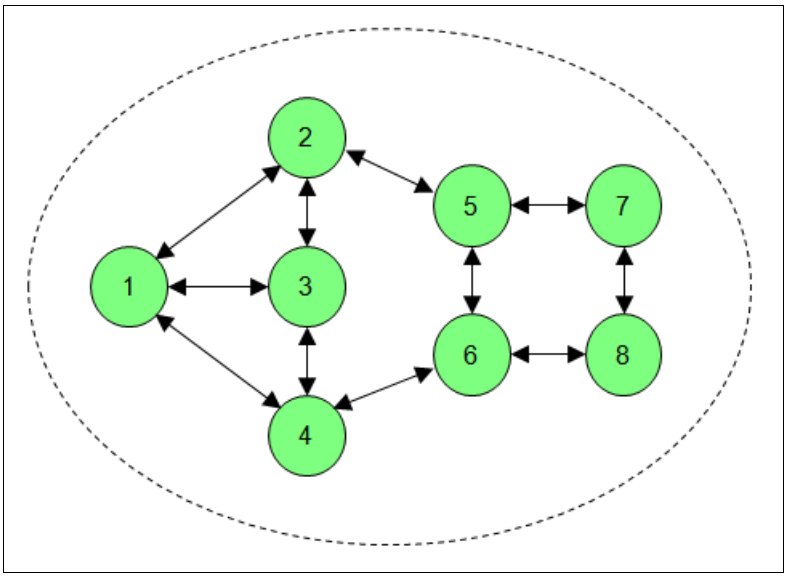

Figure 21-2 FSI replication topology group

21.2.7 Replication grid

A replication grid is a logical set of repositories that can replicate from one repository to other repositories. A ProtecTIER system must be a member of a replication grid before the system creates replication policies.

All ProtecTIER systems are capable of replication. Different models of ProtecTIER systems, such as the TS7650G Gateway, the TS7650 Appliance, and TS7620 Appliance Express models, can be part of the same grid. You can have more than one replication topology group in the same grid. A grid can also contain different types of replication groups, such as groups of VTL, OST, and File System Interface (FSI). A single replication grid can include up to 24 ProtecTIER systems.

Note: A ProtecTIER system can be a member of only one grid. After a ProtecTIER system joins a grid, it is no longer eligible to join any other ProtecTIER replication grid.

21.2.8 Replication topology group

A replication topology group defines the relationship between ProtecTIER systems in a replication grid. A group includes many-to-one, many-to-many, OST groups, and FSI groups. A replication grid can have multiple topology groups of various types, as shown in Figure 21-3.

Note: A ProtecTIER system can be a member of only one topology group at a time. A ProtecTIER system may move from one topology group to another within the grid.

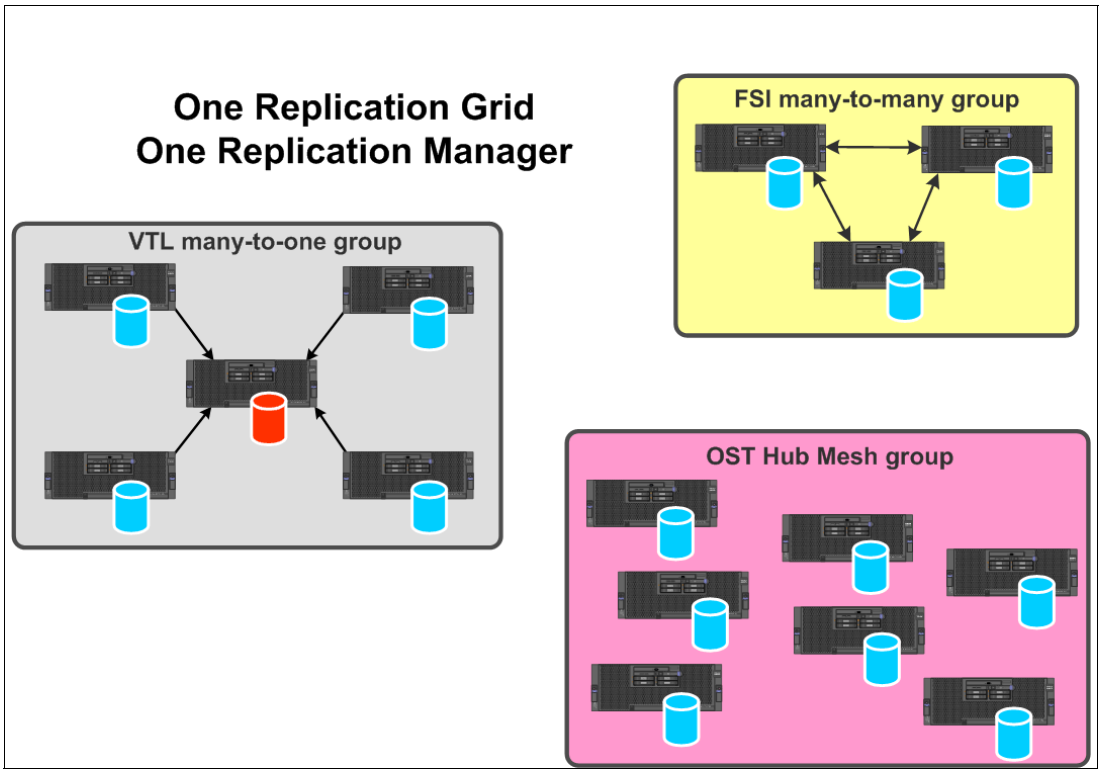

Figure 21-3 Multiple topology groups in a single replication grid

Depending on the version of ProtecTIER, there are three different replication options with different topologies:

- VTL application only:
  - With ProtecTIER V2.3, only as a pair.
  - With Version 2.4 and above, many-to-one (spoke and hub) replication groups are supported.
  - Since Version 3.1, many-to-many, or many-to-one replication groups are supported.
- OST application only:

  As of Version 2.5, the ProtecTIER product supports implicit hubs, where up to 12 hubs can be included in an OST Hub Mesh Group.
- FSI application only:

  With Version 3.2, the many-to-many or many-to-one replication groups are supported.

21.3 Replication policies

The rules for replicating ProtecTIER Data Objects (VTL cartridges, OST backup images, and FSI file systems) are defined in *replication policies*. Replication policies for FSI and VTL Data Objects are defined on the ProtecTIER system. Replication policies are defined on the NetBackup server for OST Data Objects.

When the backup application is writing to a ProtecTIER Data Object (VTL cartridge or FSI file system) that is part of a replication policy, the ProtecTIER software conducts a check on the object to determine its priority and places it in the replication queue.

Data Objects created in the primary site repository are read/write enabled so that the backup application at the primary site has full control of them and their content. Data Objects replicated to the DR site are set in a read-only mode.

In VTL replication, only one cartridge *instance* can be in a library; all replicas are on the virtual shelf in the disaster recovery site repository.

Tip: At any time, you can override the default location of any VTL cartridge and manually move the replica from the virtual shelf to a library in the repository of the disaster recovery site.

Before replication, the *dirty bit* technology system ensures that only unique and new data is transferred. To that end, both the local and secondary sites hold synchronized data for each of their data objects. The destination site then references this synchronized data to determine which data (if any) should be transferred. The replication mechanism has two types of data to transfer:

Metadata Data that describes the actual data and carries all the information about it.

User data The actual backed up data.

Data Objects are marked as *synced* after the data finishes replicating from the primary to the secondary site. So, at the time of synchronization, the local objects and their remote replicas are identical. Before replication starts running, the system ensures that only unique new data is transferred over the TCP/IP link.

Warning: If you delete a Data Object in the source repository, then all the replicas are also deleted in the target repositories.

Network failure: If a network failure occurs during replication, the system continues to try, for up to seven consecutive days, to complete the replication tasks. After seven days, a replication error is logged.

21.4 Visibility switching

Visibility switching is the automated process that transfers the visibility of a VTL cartridge from its master to its replica and vice versa. The visibility switching process is triggered by moving a cartridge to the source library Import/Export (I/E) slot. The cartridge then disappears from the I/E slot and appears at the destination library I/E slot. To move the cartridge back to the source library, the cartridge must be ejected to the shelf from the destination library. The cartridge then disappears from the destination library and reappears at the source I/E slot.

21.5 Principality

Principality is the privilege to write to a cartridge (set it to R/W mode). The principality of each cartridge belongs to only one repository in the grid. By default, the principality belongs to the repository where the cartridge was created.

The cartridge information file includes the principality repository ID field. Principality can be transferred from one repository to another during the failback process if the principality belongs to one of the following repositories:

- The DR repository
- The original primary repository, and this site is the destination for the failback
- The original primary repository with the following exceptions:
  - The original primary repository is out of the replication grid.
  - The target for the failback is a repository that is defined as a replacement repository through the ProtecTIER repository replacement procedure.

21.6 Replication Manager

The ProtecTIER *Replication Manager,* also known as Grid Manager, is the part of the software that is used to remotely manage the replication configuration and activity. From the Replication Manager, you can build and maintain the replication infrastructure and repository relationships.

In most cases, the ProtecTIER Replication Manager is run on one of the ProtecTIER nodes. It is a best practice to designate the DR site system as the Replication Manager.

It is possible to have a ProtecTIER Replication Manager (Grid Manager) installed in a node that is not one of the systems with which it is replicating. It is also possible to have a Replication Manager in a dedicated Linux server, instead of using one of the ProtecTIER nodes. To have a dedicated host as a Replication Manager requires a Request Per Quotation (RPQ), which must be requested by your IBM marketing representative.

The ProtecTIER Replication Manager that is installed on a ProtecTIER node can manage only one grid with up to 24 repositories. If a dedicated server is chosen, *and approved by the RPQ process*, it can manage up to 64 grids with 256 repositories in each grid.

You must activate the Replication Manager function before you can add it to the list of known Grid Managers, as shown in Figure 21-4.

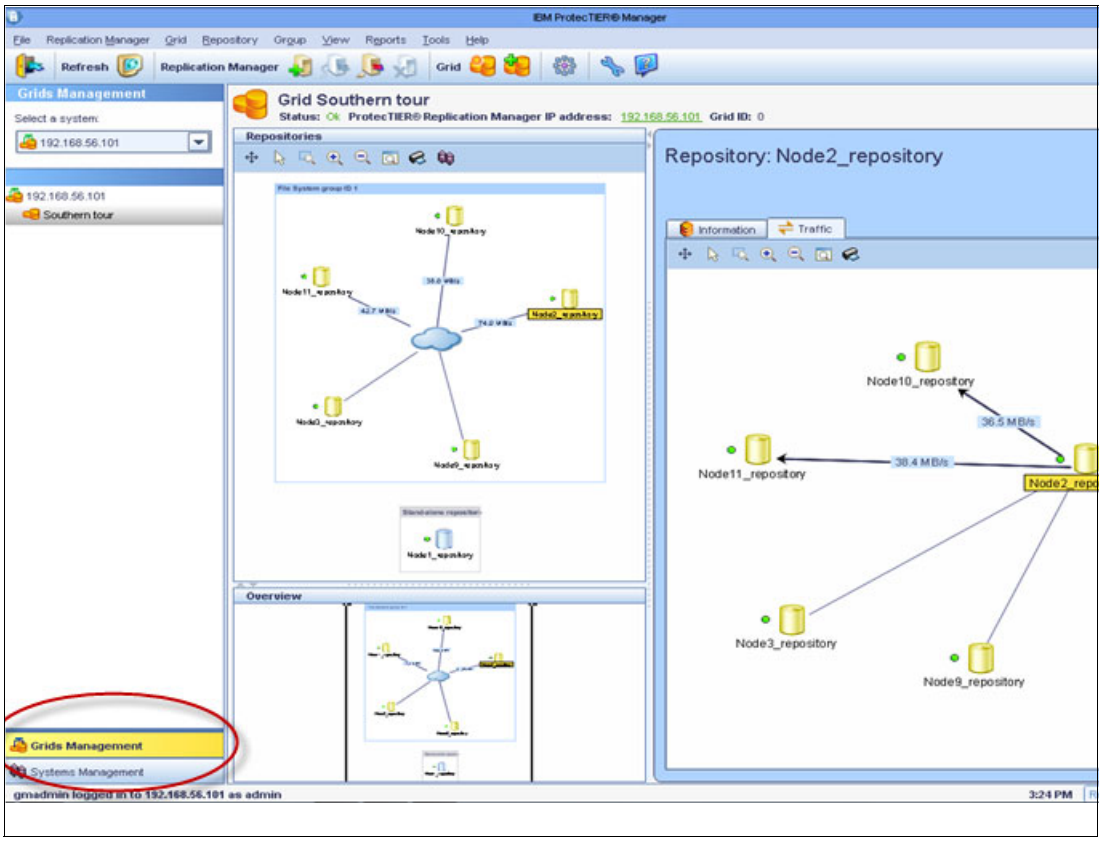

Figure 21-4 Designate ProtecTIER system as a Replication Manager

To designate a ProtecTIER system as a Replication Manager, use the **menu** command, as shown in Figure 21-5.

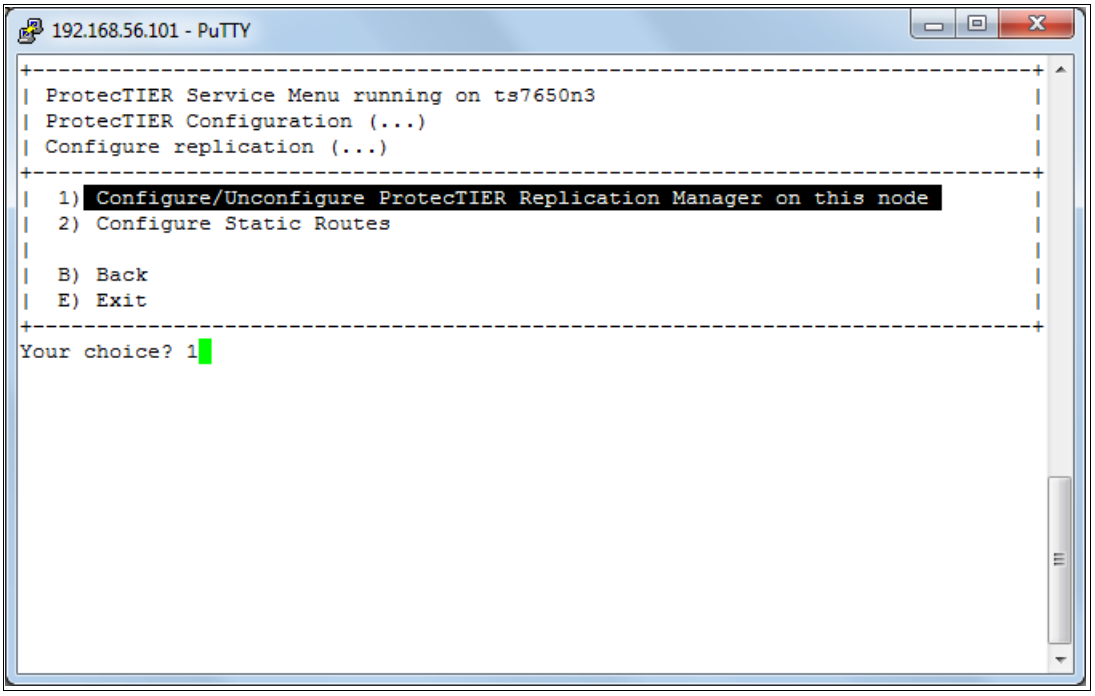

Figure 21-5 Enable Replication Manager function

21.7 Initial synchronization

When a new ProtecTIER system is configured as a replication target (secondary) for an already existing ProtecTIER system (primary), it is necessary to synchronize the primary system with the secondary system.

A deployment of a second ProtecTIER server at a secondary (DR) site has an impact on the planning cycle because the first replication jobs use more bandwidth than required after deduplication takes effect. So, when you prepare for replication deployment, bandwidth is an important consideration. During the planning cycle, the planners and engineers must consider the amount of physical data that will be replicated, the amount of dedicated bandwidth, and the extra time that will be needed for the first several replication runs. It might be necessary to implement a replication policy that allows the first replication job to complete before the next backup activity begins.

Note: For the initial replication, you must allot enough network bandwidth to account for the full nominal size of the data to be replicated.

There are two methods that can be used. Both methods focus on gradually adding workload to the replication policies over time.

Gradual management of policies over time

This is the preferred method, whether you are deploying a new system or adding replication to an existing system. In this method, you add new replication policies over time, and manually ensure that the total daily volume of replicated data remains within the bandwidth limit. Replication policies with a gradual increase are preferred to stay within the available network bandwidth boundaries and within the time frame that is scheduled for replication activity.

Priming the DR repository at a common locality with the primary system

Priming the DR system at a primary site first and then moving it to its DR location has limited practical value, and is not the preferred choice. In a multisite deployment, this method is a poor choice:

- If you take this approach, you must manage the synchronization process again when the systems are placed in to their final location.
- If you are synchronizing a full, partial, or even a newly started repository, the system must have sufficient network bandwidth for primary and secondary systems to synchronize within the available time frame.

21.8 Replication schedules

The ProtecTIER product offers two modes of operation for the replication activity:

- Scheduled replication occurs during a predefined time frame.
- Continuous replication runs constantly.

The mode of operation is configured at the source system. All defined replication policies operate in one of these modes. In most cases, scheduled replication is the best approach. It enables administrators to accurately plan for performance, and to better ensure that SLAs are met. The replication mode of operation is a system-wide option. It affects all polices in the system.

By default, Data Objects are continuously being replicated from the primary (local) site to the repository at the disaster recovery (DR) site. Optionally, a replication schedule can be defined to limit replication activity to specific time slots during the week.

21.8.1 Continuous replication

Continuous replication can run concurrently with the backup operation. Typically, it requires a larger system to enable concurrent operations. This option can affect backup performance because the “read” function is shared between the deduplication processes and the backup operation. The following aspects must be considered when you plan continuous replication:

- Data automatically starts replicating to a DR site repository soon after it is written to the primary ProtecTIER system.
- Replication runs faster (up to 100% of available performance) if the primary system is idle (no backup or restore activity).
- If it is running concurrently, replication is prioritized lower than backup or restore in the ProtecTIER system.

Continuous replication is available or recommended in the following situations:

- A system has consistently lower bandwidth.
- The operation calls for few backup windows that are spread throughout the day.
- Deploying a multisite scenario, especially across multiple time zones.

21.8.2 Scheduled replication

The scheduled replication occurs during a predefined time frame, which is the recommended mode for most applications. This mode imitates the procedure that is used with physical tapes that are being transported to a DR site after backup is completed. This method allows users to keep complete sets of backup data together with a matching backup application catalog or database for every 24 hour period.

With this approach:

- Backups are allowed to finish without performance impact from replication.
- The user defines the start and end of the replication time frame.
- Replication activity begins at the predefined time.
- Replication stops at the end of the time window specified.
  - Each cartridge in transit stops at a consistent point at the end of the time window.
  - Replication does not occur outside of the dedicated time window.

During a replication schedule replication, activity has the same priority as backup and restore activity. If backup and restore activity takes place during the same time frame, they are equally weighted and processed in a first-in-first-out manner. Because the overall system throughput (backup and restore plus replication) can reach the maximum configured rate, the backup duration might vary.

Tip: Because both backup and restore and replication jobs access the same back-end disk repository, contention between these two processes can slow them down. This situation could impact the backup SLA and overall RTO. Therefore, replication tasks should not take place during the time frame that is dedicated for a backup operation.

You must plan for and configure the ProtecTIER system resources to accommodate both types of activities to finish their tasks within the wanted time frames. However, the ProtecTIER system remains available to the backup application throughout the time frame that is dedicated to replication. So if a backup or restore operation is necessary during the replication time frame, the operation can be performed.

A primary benefit of scheduled replication is the ability to strictly classify when the ProtecTIER server uses the network infrastructure, and accurately isolate the usage of the network. This mode of operation is aligned with the backup and DR activity, where users manage a specific backup time frame and schedule cloning or vaulting jobs that follow the backup.

21.8.3 Centralized Replication Schedule Management

Each ProtecTIER system has the optional ability to schedule both incoming and outgoing replication activity by using a weekly schedule that is divided into one-half hour time slots. There is only one schedule for a ProtecTIER system that governs all replication policies on that system. Figure 21-6 shows an overview of this topic.

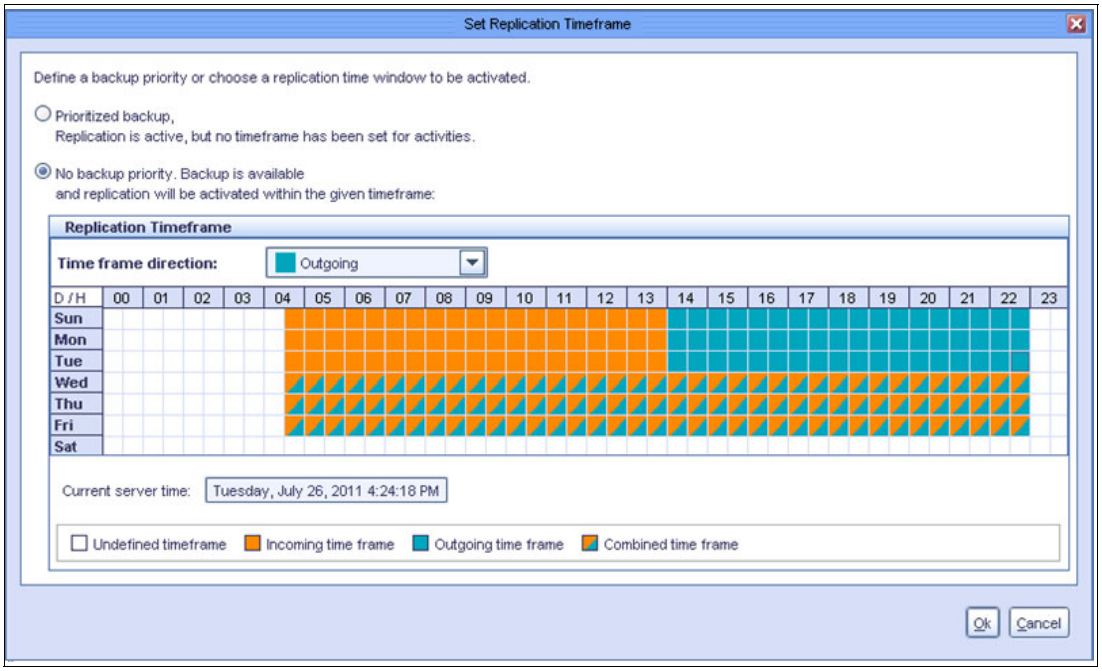

Figure 21-6 Replication schedule

Schedules can be set on both sending (spoke) and receiving (hub) ProtecTIER systems.

Important: Avoid time window conflicts when you define time frames at the hub and at the spokes:

- There is no synchronization mechanism to foresee misalignments, so if you set the hub and spokes to different time slots, replication never runs.
- Ensure that the hub has enough time frame slots to accommodate all of the spokes' combined time frames.

Starting with Version 3.1, ProtecTIER introduced the Centralized Replication Schedule Management function. Using this function, you can view and set replication schedules for all the nodes in a grid and visually check time frame alignment between nodes, as shown in Figure 21-7.

Note: Centralized Schedule Management is available in the Grid Management view of the ProtecTIER Manager GUI

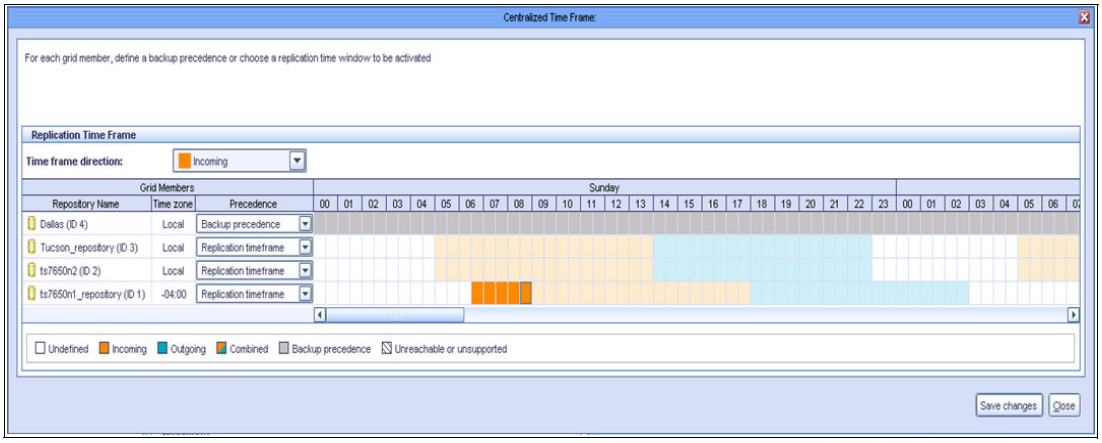

Figure 21-7 Centralized Replication Schedule Management

21.8.4 Replication rate control

There are enhanced system replication throttling and dynamic *system resource allocation* functions for incoming and outgoing replication. ProtecTIER replication offers the following enhanced features and benefits:

- Setting replication performance limits: The nominal performance limit reflects the overall resource consumption of the system. The physical performance limit reflects the network transfer rate of the replication network.
- Enhancements to the replication rate control mechanism: Currently, the replication rate control (RRC) is used when a user does not provide a time frame and the system replicates continuously. The *rate calculation* determines the maximum rate that is possible in both levels of system usage (IDLE and BUSY), and normalizes the rate.

- A new GUI feature that provides an at-a-glance view of the proportion of the repository data, replication data, local backup data, and free space, as shown in Figure 21-8.

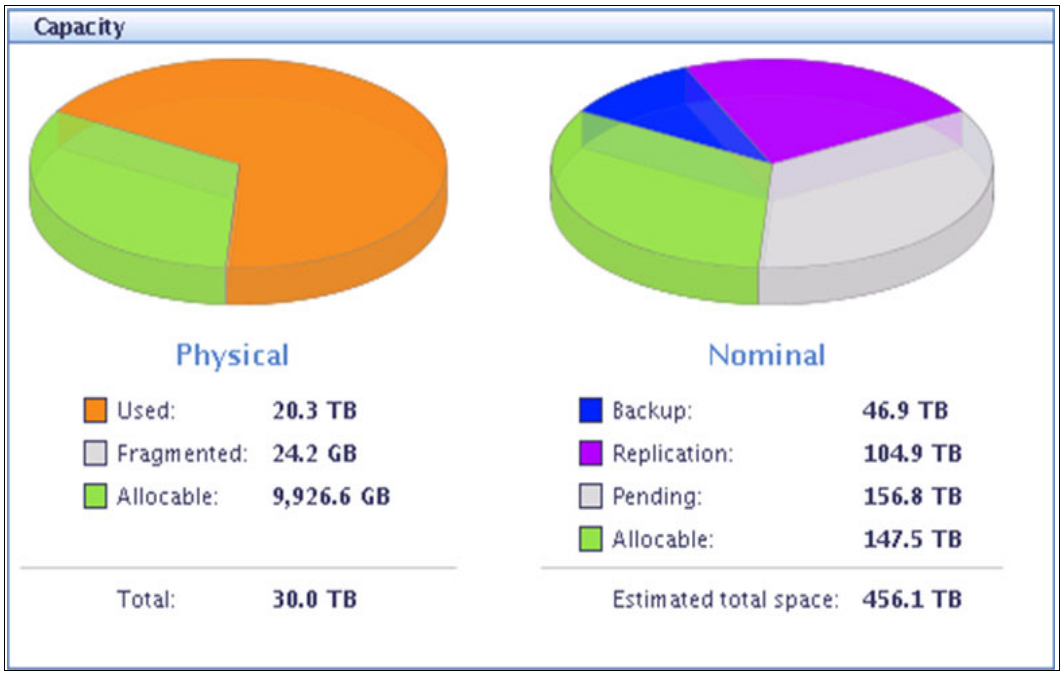

Figure 21-8 Repository usage by category

The nominal and physical throughput (data flow rate) can be limited by setting the replication rate control. The following information must be considered:

- Both the nominal and physical amounts of data that are being processed or transferred.
- The ability to send and receive new unique data between spokes and the hub.
- ProtecTIER validates all the new or updated objects at the target repository before it makes them available for the user. Setting the replication rate control allows the user to limit the nominal and physical throughput (data flow rate of replication). This feature can be used on spokes and on the hub for both sending and receiving. The values that are set for the physical and nominal limits have no explicit influence on one another. The values that are set in the physical throughput might, but do not necessarily, affect those values that are set in the nominal throughput, and vice versa. However, when you use both methods, the physical settings override the nominal ones.

Setting a nominal limit

When you set a nominal limit, you define the maximum ProtecTIER server system resources that can be used to process the replication data. The nominal throughput directly affects the replication data flow and the load on both the source and destination repositories.

By setting a nominal limit for a ProtecTIER system that performs both backup and replication, the replication processing does not compete with the backup operation for system resources. Setting the limit on a source repository ensures that the backup operation realizes the total possible throughput minus the nominal limit set.

For example, on a node with a performance capacity of 500 MBps that performs backup and replication concurrently, the user might set the following limits:

- 300 MB per second when replication is running on its own
- 100 MB per second when replication is running concurrently with a backup

Setting a physical limit

When you set a physical limit, you limit replication network bandwidth consumption by the ProtecTIER server. This limit is intended to be used when the network is shared between the ProtecTIER server and other applications so that all applications can run concurrently. The Physical throughput limit restrains the amount of I/O and resources that the replication processes can use. This limit reduces the total load on the replication networks that are used by the repository and the amount of resources that are needed at the peer repository.

Although this limit can be set at either the spoke or the hub (or both), it is typically set at the spoke. Setting a limit at the hub limits the bandwidth for the entire replication operation, which results in *de facto* limitations on all spokes.

21.8.5 Setting replication rate limits

You can limit the replication rates and throughput (Figure 21-9) in the following situations:

- During a backup or restore operation
- When there is no backup or restore activity
- Within a defined replication time frame

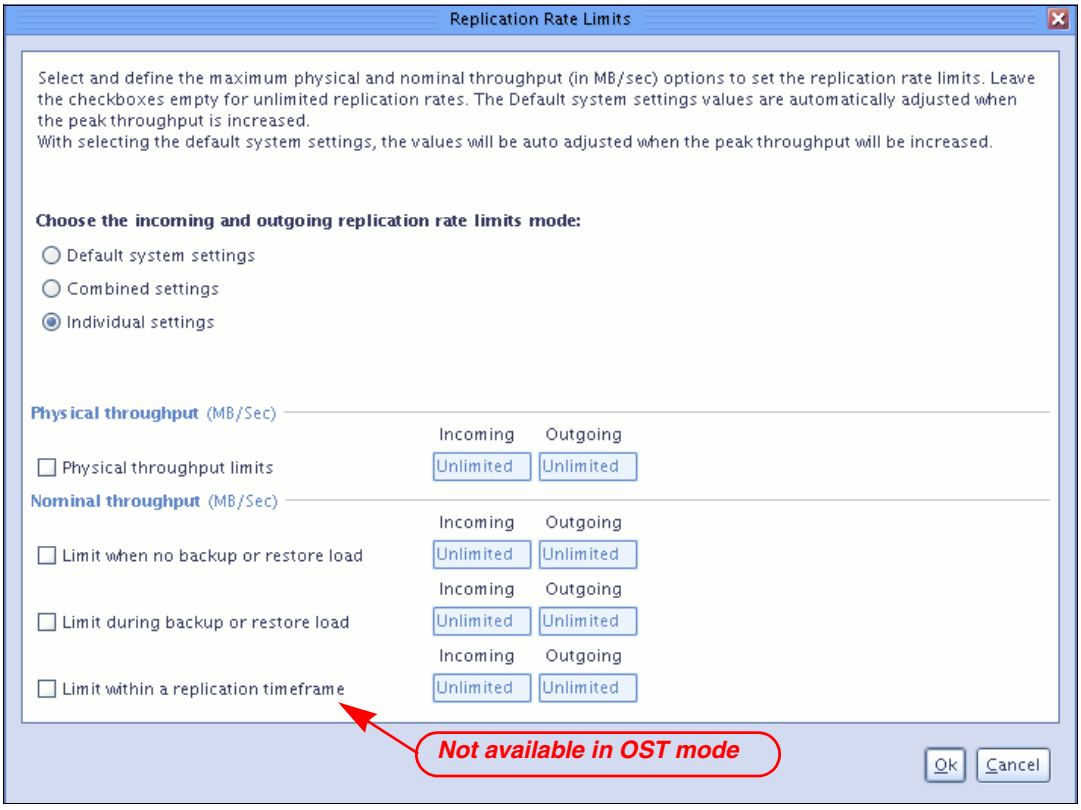

Figure 21-9 Setting replication rate limits

21.8.6 Limiting port bandwidth consumption

Bandwidth throttling (physical limit) controls the speed at which replication operates, where the user can specify a maximum limit for the network usage. The ProtecTIER hardware platform uses two TCP/IP Gigabit Ethernet interfaces per node for sending the actual replicated data traffic across the network. Data transfer occurs at a maximum rate of 175 - 190 MBps per single node and 350 - 380 MBps for a dual-node cluster. By default, there is no configured bandwidth limit. The ProtecTIER server uses as much bandwidth as it can.

If the physical network layer consists of dark fiber or other high-speed network infrastructure, there is typically no reason to limit replication throughput. However, if the ProtecTIER server is running over a smaller network pipe that is shared by other applications, you can restrict the maximum throughput that is used by ProtecTIER replication.

This parameter is adjustable per GigE port on all nodes in the replication grid. It applies only to outgoing data. Set it at the source (sending) system. If the source system is composed of a dual-node cluster, it is important to set the limit at each node.

For example, to hold ProtecTIER replication to a limit of 100 MBps, set each of the four available GbE ports to 25 MBps. Likewise, if the replication traffic is split between two networks with different bandwidth capacities, you can set different limits per port to implement a network-specific cap. By default, the setting per port is Unlimited (Figure 21-10).

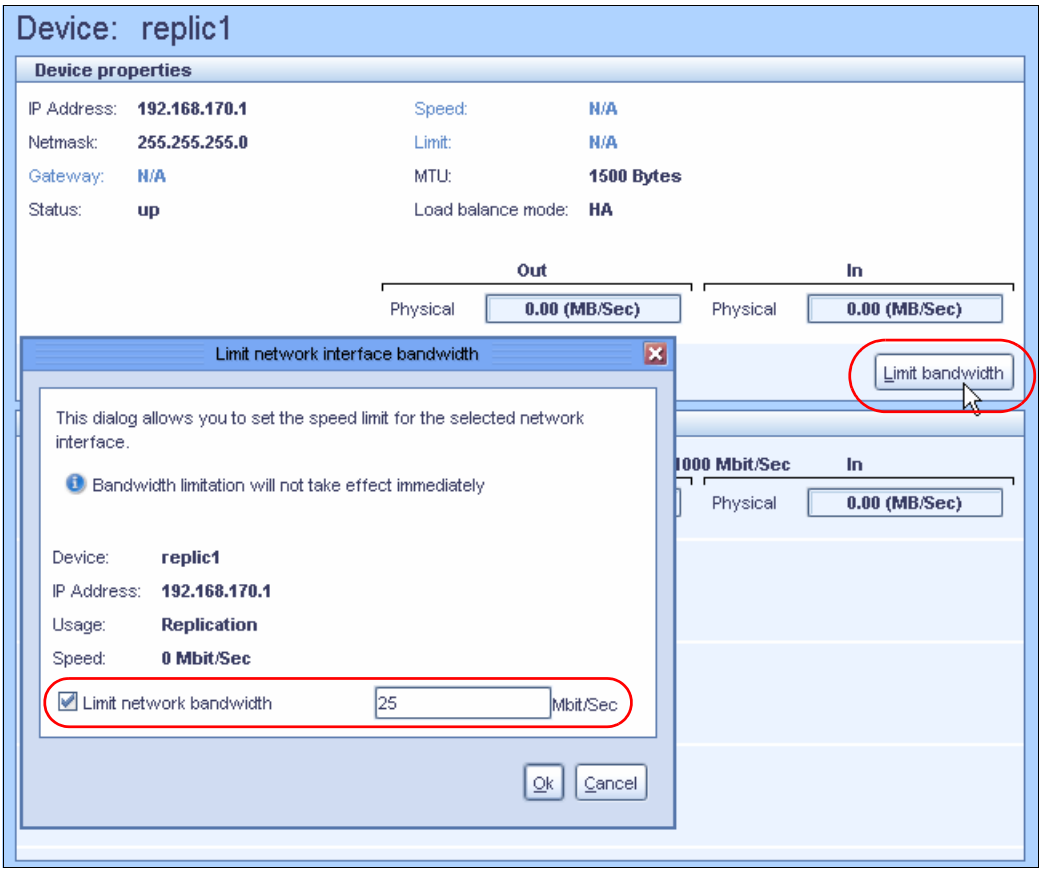

Figure 21-10 Potential modification of the Eth3 interface limit

Changing the bandwidth: If the bandwidth limitation is changed during replication, the change does not take effect immediately. If replication begins after the bandwidth limitation change, the effect is immediate.

21.9 Replication backlog

When replication activity is started, the source system builds a list of new and changed data blocks and sends that list to the receiving system. The receiving system checks the list and determines which data blocks it must synchronize with the source system and then sends requests for the transferal of data blocks. Now, there is a *backlog* of replicated data. The source system monitors and displays the amount of backlog replication data in the ProtecTIER Manager GUI Activities view.

Having a backlog of replicated data is not a problem in and of itself. A potential problem is indicated when the amount of backlog data does not go down over time.

If there is an unscheduled long network or DR site outage, the replication backlog might become too large for the system to catch up. A prolonged replication backlog might be an indication of insufficient available bandwidth that is allocated for the replication operation. In an optimal situation, there should not be any backlog activities that remain (Figure 21-11).

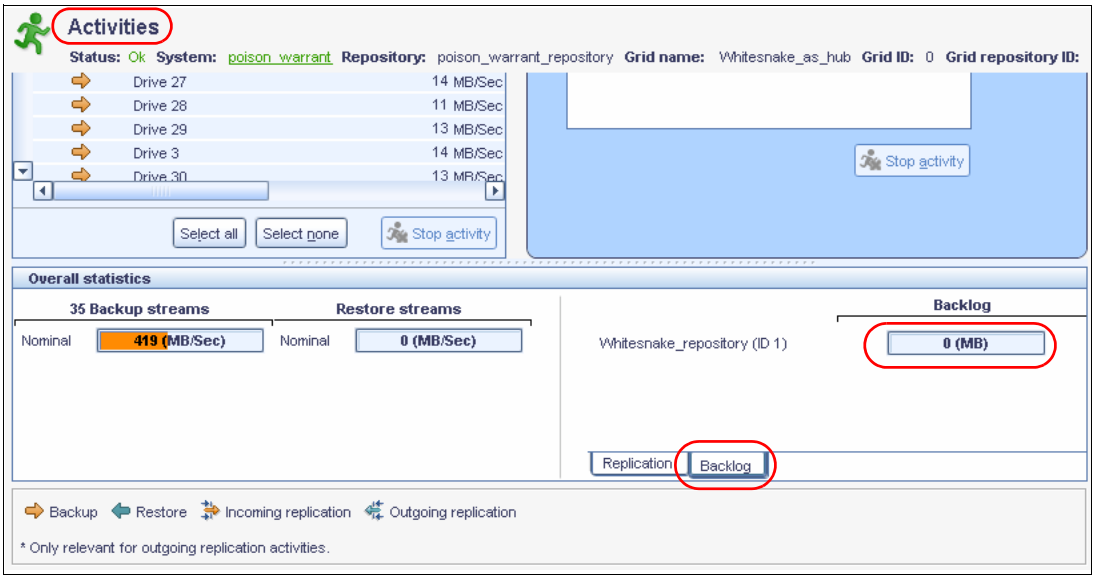

Figure 21-11 Backlog status in replication activities

Use either of these methods to delete the replication backlog:

- From the ProtecTIER Manager replication Policy view, select a specific policy and click **Abort activities**.
- From the ProtecTIER Manager replication Activities view, select a specific activity and click **Abort activities**.

SLAs: For the system to support the organization's set of service level agreements (SLAs), enough bandwidth must be allotted for replication during the replication window so that all the policies are run within the allotted time.

Aborting replication tasks

Aborting replication tasks removes them from the list of pending and running tasks. These tasks are automatically returned to the replication queue if the specific cartridge is in one of the following states:

- Appended
- Ejected from the library
- Selected for manual execution

One way to prevent these replication tasks from rerunning is to mark those cartridges as *read-only* either on the ProtecTIER server or by the backup application. These cartridges are not used for further backups, and therefore do not replicate. New (scratch) sets of cartridges are used for subsequent backups, and do not contain backlog data that does not need to be replicated.

Tip: To resume I/O activity, use different barcodes. Because the earlier data on the older cartridges is replicated before the new data, the backlog is too large to manage. Using a different set of barcodes allows the new data to be replicated, and you can skip replication of the data from the old cartridges.

21.9.1 SNMP alerts for replication backlog

ProtecTIER provides a method for monitoring backlog data and notifying you if backlog data becomes greater than a user-defined threshold setting, as shown in Figure 21-12.

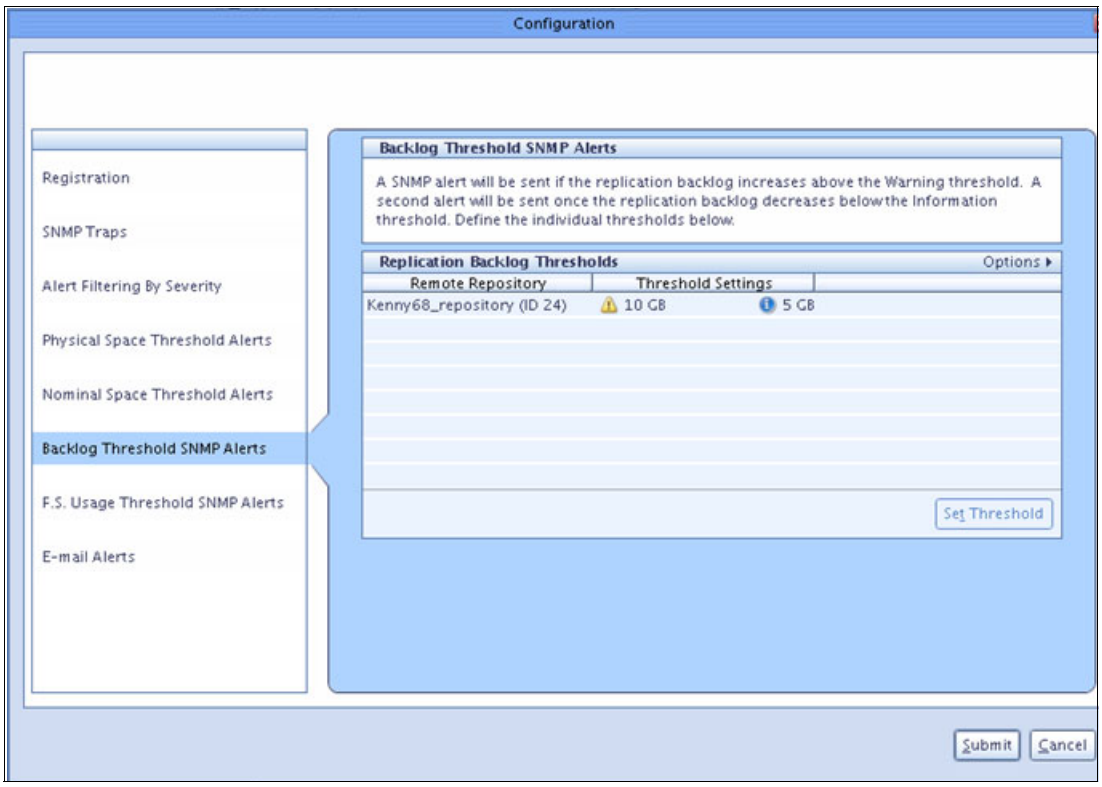

Figure 21-12 SNMP Alerts

21.9.2 Reserving space for local backup data

ProtecTIER can reserve local-backup-only space for the hub repository. You can use this enhancement to exclusively assign a portion of a hub repository for local backups. This enhancement was added to ensure that capacity is reserved only for local backup. Replication cannot be written to this portion of the hub repository. Error notifications display if the repository hub areas that are reserved for local backup or replication are reaching maximum capacity. Figure 21-13 shows the window for this enhancement.

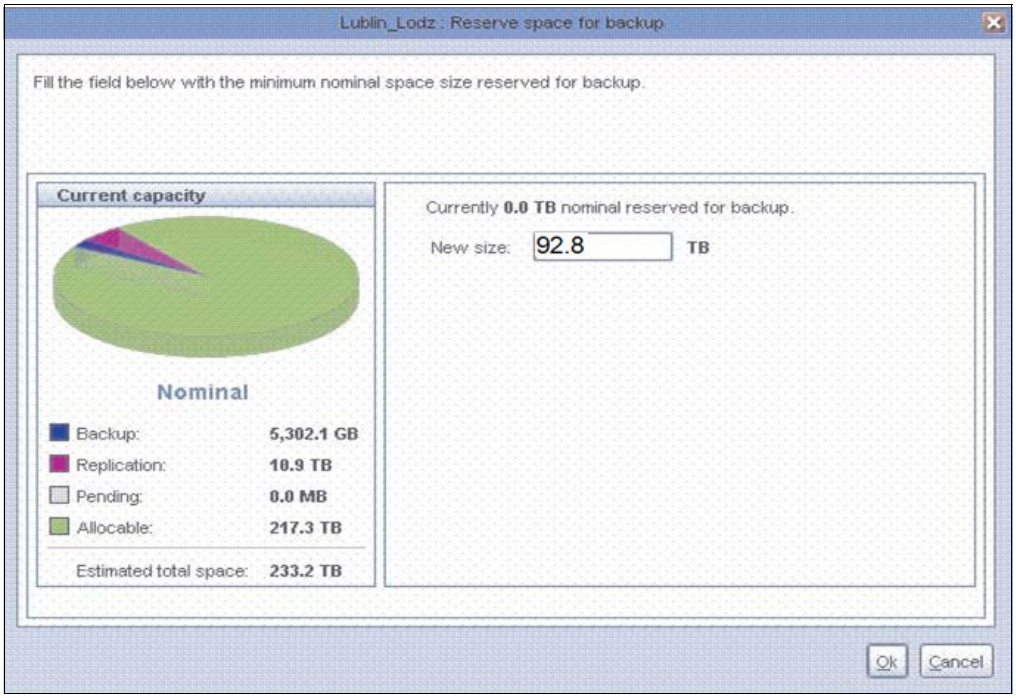

Figure 21-13 Current capacity

21.10 Replication planning

The planning process for ProtecTIER systems with replication deployed requires more input and considerations beyond the individual capacity and performance planning that is needed for a system that is used only as a local VTL, FSI, or OST. When a multiple-site, many-to-one, or many-to-many replication strategy is deployed, the entire configuration, including all spokes and hubs, must be evaluated.

The planning of a many-to-one replication environment is similar to the planning of *a* one-to-one replication strategy. The only difference is that you must combine all replication loads (and potentially a local backup) for the hub. The ProtecTIER Planner tool should be used at the planning stages before any ProtecTIER deployment.

21.10.1 Bandwidth sizing and requirements

The ProtecTIER server replicates only the new or unique deduplicated data. Data that was deduplicated on the primary server is not sent to the DR site. However, the DR site (hub) must synchronize with all of the data on the primary server to ensure 100% data integrity. For example, if there are two cartridges at the primary site, cartridge A and cartridge B, and each contains the same 1 GB of data:

- Replicating Cartridge A transfers 1 GB of physical (which equals nominal) data. The data is new to the DR site repository.
- Replicating Cartridge B transfers 0 GB of physical data to the DR site. Because the same data was transferred with the Cartridge A replication, all of the Cartridge B data exists at the DR site repository. Following the replication action, 1 GB of nominal data is indexed on Cartridge B at the DR site.

Throughput: The maximum replication throughput for any scenario depends on many factors, such as the data type and change rate.

Tip: When you configure a system with unbalanced replication network lines, the total throughput is reduced to the slowest line.

The best practice is for both networks in the replication network topology to be set at the same speed at the source. If need be, set the speed of the port, by using **ethtool**, of the faster network to the same speed as the slower network. Here is an example of the command:

```
ethtool -s eth2 speed 100
```

21.10.2 Replication throughput barriers

The following types of replication data transfer throughput barriers have been identified:

- Physical data-transfer barrier: This barrier results from a ProtecTIER node that has two 1 GB Ethernet ports, where single node system supports up to 190 MBps physical data transfer (two 1 Gbps replication Ethernet ports). A dual-node clustered system supports up to 380 MBps physical data transfer (as it has four 1 Gbps replication Ethernet ports).
- Nominal data barrier: The nominal data barrier results from the maximum processing capability of a given ProtecTIER system (3958-DD5):
  - A single node system supports up to 1540 MBps of nominal data backup ingest, replication, or a combination of these activities.
  - A dual-node clustered system supports sustainable rates of up to 2860 MBps of nominal data backup ingest, replication, or a combination of these activities.

Performance figures: For all calculations in the following sections, we use the performance figures of the new ProtecTIER model 3958-DD5. When you make the estimation with older machines, adjust these variables:

- DD4 maximum system performance (concurrent backup and replication activity):
  - 1,400 MBps single node (nominal) (900 MBps Version 2.5)
  - 2,000 MBps 2-node cluster (nominal) (1500 MBps Version 2.5)
- DD3 maximum system performance (concurrent backup and replication activity):
  - 935 MBps single node (500 MBps Version 2.4)
  - 1,400 MBps 2-node cluster (1000 MBps Version 2.4)

Maximum specifications are based on a TS7650G and a correctly configured back-end disk array. Typical restores are about 15 - 20% faster than backups.

21.10.3 Calculating the replication data transfer

Use the following formula to calculate the replication data transfer. The formula estimates the number of gigabytes of changed data to be sent across the network, and adds 0.5% for control data.

Replication data transfer = daily backup × (Change rate + 0.5%)

Example 21-1 shows an example of this formula.

Example 21-1 Replication of a 6 TB daily backup with change rate of 10%

```
RDT= 6000 GB × (10% + 0.5%)= 630 GB
```

In this scenario, 630 GB of physical data is replicated to the second site, rather than 6 TB of nominal data that would otherwise be transferred without deduplication.

21.10.4 Calculating replication bandwidth

Use this formula to calculate the required replication bandwidth:

Replication bandwidth = replication data transfer ÷ available replication hours

Example 21-2 shows an example of this formula.

Example 21-2 For a replication window of 10 hours

```
replication bandwidth = 630 GB ÷ 10h = 63 GB per hour
```

The WAN bandwidth must be able to transfer an average 63 GB per hour, which represents the requirements for an 18 MBps link between spoke and hub.

Tip: Continuous replication operation (24 hour replication concurrent with a backup operation) is rarely the recommended mode of operation. Add 10% of the required bandwidth for headroom in case of network outages or slowdown periods.

Ports for replication in firewalled environments

In a firewalled user environment, you must open the following TCP ports in order for IP replication to function properly:

- The replication manager uses TCP ports 6202, 3501, and 3503.
- The replication operation between any two repositories uses TCP ports 6520, 6530, 6540, 6550, 3501, and 3503.

ProtecTIER replication does not use any User Datagram Protocol (UDP) ports. In addition to bandwidth performance, two other major factors affect network quality:

- Latency

  Depending upon many factors along the network span, the latency in any WAN varies, but must never exceed 200 ms. If so, it might decrease the system replication throughput. For more information about this topic, contact your network administrator.

- Packet loss

  Packet loss across the network should be 0%. Any other value indicates a major network problem that must be addressed before replication is deployed. For more information about this topic, contact your network administrator.

21.11 Bandwidth validation utility

The **pt_net_perf_util** network testing utility is included as part of the ProtecTIER software package. As a part of the installation process, the installer must ensure that the ProtecTIER nodes at both sites can run this utility concurrently.

Before replication is deployed, the bandwidth validation utility tests and verifies the maximum replication performance between two future ProtecTIER repositories by emulating the network usage patterns of the ProtecTIER replication component. Although this utility does not predict replication performance, it might discover performance bottlenecks.

Tip: It is not necessary to build a repository or configure the ProtecTIER back-end disk to run the **pt_net_perf_util** test tool.

Here are the requirements of the **pt_net_perf_util** utility:

- Red Hat Linux Version 5.6 or later. You must have the following standard external utilities in the current path:
  - **ping**
  - **netstat**
  - **getopt**
  - **echo**
- The utility has two modes of operation: It tests and verifies data flow from the client server to the target server, and vice versa.
  - Data flow from the client server to the target server

    The client server is the ProtecTIER system that transmits the test data. The target server is the ProtecTIER system that receives the data. Based on the data that is sent by the client and received by the target, a script outputs key network parameters. Data flow from the client server to the target server is tested to verify that the target server functions as expected.

– Data flow from the target server to the client server

  It is also important to test data flow in the reverse direction (from target to client) to measure the bandwidth performance during disaster recovery failback. Network bandwidth is not always the same in both directions.

21.11.1 Using the bandwidth validation utility to test the data flow

The target ProtecTIER server must be started and running before the client server. Before you run the utility, shut down all other programs on both the client and target server of the ProtecTIER replication grid.

This procedure, which tests network performance between two servers on a WAN, takes about 25 minutes. The utility performs five foreground tests (tests 1 - 5), and one background test (test 6).

To test the data flow from the client server to the target server, complete the following steps:

1. Start the target server mode of the utility on Server A. Run either of the following commands on the command line:
   – To use the **iperf** external utility, run the following command:

     ```
     /opt/dtc/app/sbin
     ./pt_net_perf_util -s
     ```

   – To use the **nuttcp** external utility, add **-n** to the command:

     ```
     cd /opt/dtc/app/sbin
     ./pt_net_perf_util -s -n
     ```

2. Start the client mode of the utility on Server B. Run either of the following commands (where **-t 300** is the wanted duration of test):

   To use the **iperf** external utility, run the following command:

   ```
   cd /opt/dtc/app/sbin
   ./pt_net_perf_util -c server1 -t 300
   ```

   To use the **nuttcp** external utility, add **-n** to the command:

   ```
   cd /opt/dtc/app/sbin
   ./pt_net_perf_util -c server1 -t 300 -n
   ```

   The utility automatically performs all the tests in sequence. The client server output (Server B) is similar to the output in Example 21-3.

Example 21-3 Sample client server output of pt_net_perf_util

```
*** Latency
PING 10.0.13.194 (10.0.13.194) 56(84) bytes of data.
---10.0.13.194 ping statistics ---120 packets transmitted, 120 received, 0%
packet loss, time 119060ms
rtt min/avg/max/mdev = 57.403/78.491/104.451/9.872 ms
*** throughput -Default TCP
[3] 0.0-120.1 sec 2.41 GBytes 173 Mbits/sec
*** throughput -1 TCP stream(s), 1MB send buffer
[3] 0.0-120.0 sec 2.51 GBytes 180 Mbits/sec
*** throughput -16 TCP stream(s), 1MB send buffer
[SUM] 0.0-121.4 sec 5.91 GBytes 418 Mbits/sec
*** throughput -127 TCP stream(s), 1MB send buffer
[SUM] 0.0-126.1 sec 8.08 GBytes 550 Mbits/sec
```

```
Number of TCP segments sent: 1619061
Number of TCP retransmissions detected: 201038 (12%)
Done.
```

Interpreting the results

The following interpretation is based on the example output that is shown in Example 21-3 on page 393. Actual test results vary.

- Test 1 - Latency: Checks the nominal network link latency and packet loss. As you can see from the script:
  - The average round-trip-time (RTT) was 78.491 ms.
  - There was 0% packet loss.
- Test 2 - Throughput default settings: Checks the maximum TCP throughput by using a single data stream with default TCP settings.
  - The test ran for 120.1 seconds.
  - 2.41 GB of data was transferred, with an average throughput of 173 Mbps.

  Remember: 1 MB = 1,048,576 bytes. 1 MBps = 1,000,000 Bps.

- Test 3 - Throughput of single stream with a 1 MB send buffer: Checks the maximum TCP throughput by using a single data stream with a 1 MB send buffer.
  - The test ran for 120.0 seconds.
  - 2.51 GB of data was transferred, with an average throughput of 180 Mbps.
- Test 4 - Throughput of 16 streams with a 1 MB send buffer: Checks the maximum TCP throughput by using 16 streams with a 1 MB send buffer.
  - The test ran for 121.4 seconds.
  - 5.91 GB of data was transferred, with an average throughput of 418 Mbps.

  The extra streams yielded higher usage of the connection.

  Note: The megabits per second reported in test 4 is the maximum replication performance the system can achieve if the backup environment uses three or fewer cartridges in parallel.

- Test 5 - Throughput of 127 streams with a 1 MB send buffer: Checks the maximum TCP throughput by using 127 streams with a 1 MB send buffer.
  - The test ran for 126.1 seconds.
  - 8.08 GB of data was transferred, with an average throughput of 550 Mbps.

  The megabits per second reported in this test is the maximum replication performance that your system may achieve. If this number is lower than anticipated, contact your network administrator.

  Note: The throughput value that is given by test 5 is the maximum potential physical replication throughput for this system. It is directly affected by the available bandwidth, latency, packet loss, and retransmission rate.

  Tip: Using the TCP protocol, it takes a while for your system to reach its maximum throughput. Run the test with longer testing times (more than 300 seconds).

- Test 6 - TCP retransmissions versus total TCP segments sent: Compares the total TCP transmissions sent with the number of packets that are lost and retransmitted. It also calculates a retransmission rate.
  - A total of 1,619,061 TCP segments were sent during the five tests.
  - 201,038 were lost and retransmitted.
  - The retransmission rate is 12%.

  The retransmission rate imposes a direct penalty on the throughput, as the retransmission of these packets uses bandwidth. The retransmission can be caused by the underlying network (such as packet dropping by an overflowing router). It can also be caused by the TCP layer itself (such as retransmission because of packet reordering). Segment loss can be caused by each of the network layers.

  Hint: TCP retransmission larger than 2% might cause performance degradation and unstable network connectivity. Contact your network administrator to resolve this issue.

21.11.2 Repository replacement

Use the repository replacement function when you want to fail back to a different or rebuilt repository. To accomplish this task, complete the following steps:

1. Cancel the pairing of the original repositories in the replication manager.
2. Take the original primary repository out of the replication grid.

   Important: If a new repository replaces the original one, then the new repository must be installed and join the replication grid.

3. Run the ProtecTIER repository replacement wizard and specify the repository to be replaced and the replacement repository.

After the disaster recovery situation ends and the primary repository is restored or replaced, you can return to normal operation with the production site as the primary site.

Figure 21-14 shows how to leave ProtecTIER DR mode.

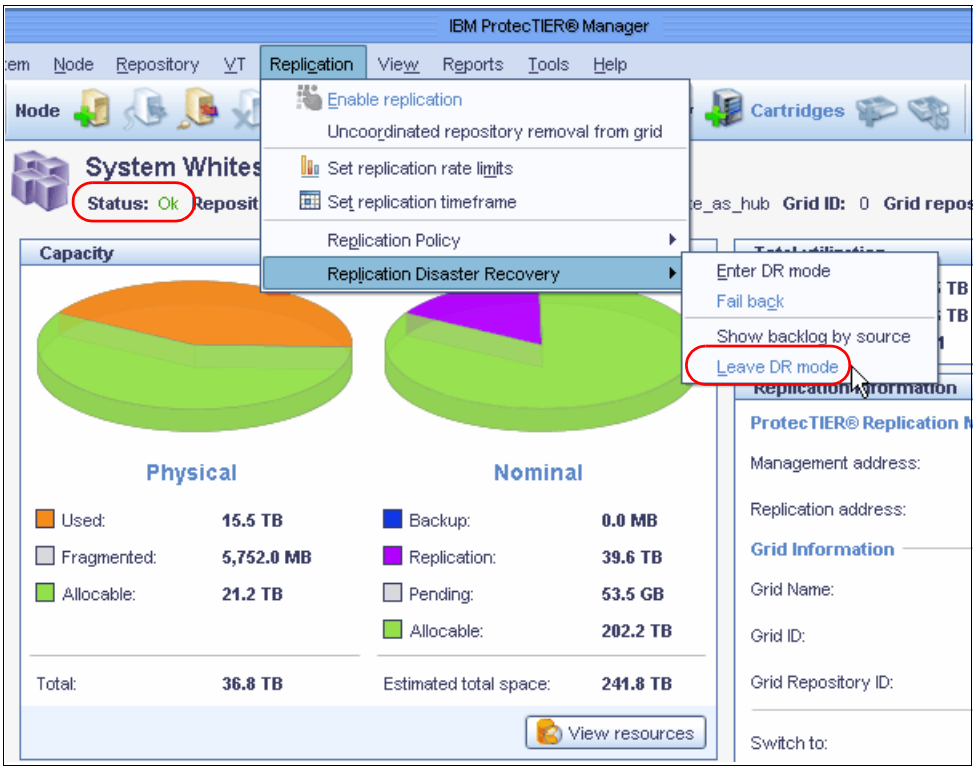

Figure 21-14 Leaving ProtecTIER DR mode

Important: Leaving DR mode should always be preceded by a failback action.

For more information, see *IBM System Storage ProtecTIER User Guide for Enterprise Edition and Appliance Edition*, GC53-1156.

Cartridge ownership takeover

Cartridge ownership takeover enables the local repository, or hub, to take control of cartridges that belong to a deleted repository. Taking ownership of the cartridges on a deleted repository enables the user to write on the cartridges that previously belonged to the replaced (deleted) repository. This process is also known as a change of *principality*.

Cartridge ownership: The repository can take ownership of a cartridge only if the repository is defined on the Replication Manager as the replacement of the deleted repository.

21.12 Planning ProtecTIER replication

In this section, we provide case studies of planning and sizing ProtecTIER Replication. Both many-to-one (Spoke and Hub) replication environment and many-to-many bidirectional replication scenarios are described.

21.12.1 Deployment planning scenario: many-to-many

Figure 21-15 shows a deployment planning scenario for four sites, each with a dual-node gateway. This is an example of building a maximum four node many-to-many VTL configuration with various replication strategies.

VTL, OST, and FSI systems are configured in many-to-many replication groups, so this same sizing strategy applies, but the throughput numbers vary for each type.

At the time of the writing of this book, the maximum listed speed for a dual-node DD5 VTL gateway is 2500 MBps, so all calculations are based on this speed. As ProtecTIER technology improves, the rated performance numbers continue to increase. For the latest published ratings, go to the following website:

http://www-03.ibm.com/systems/storage/tape/ts7650g/index.html

Assume that the following characteristics of the replication grid are present:

- All processes have the same maximum rate (2500 MBps).
- All data exists at all sites.

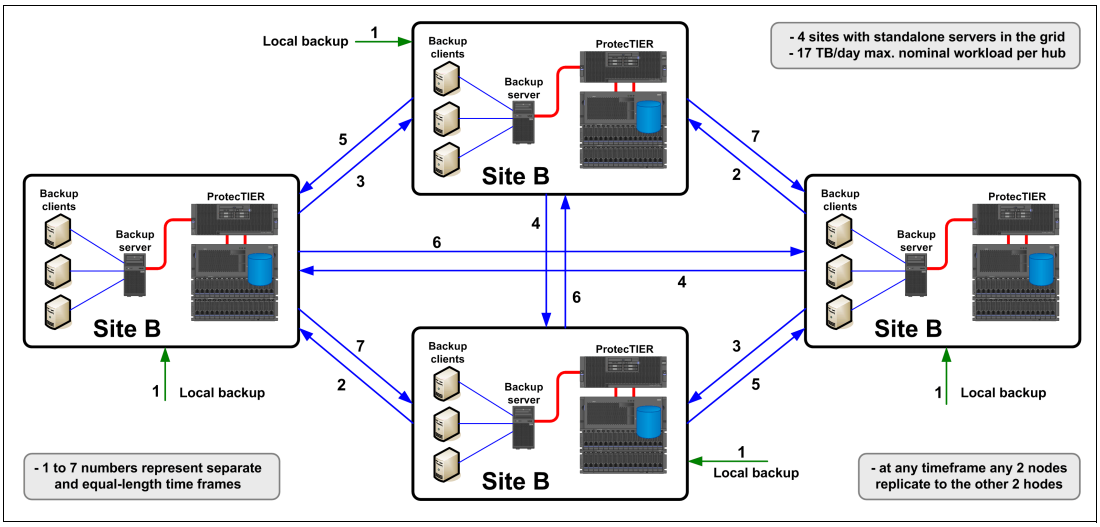

Figure 21-15 Four sites - each with a dual-node gateway hub

Maximum workloads

This section provides scenarios of a maximum backup with no replication, as well as scenarios with maximum workload with one, two, and three replicated copies.

Maximum backup with no replication

With no data replication, a maximum of 24 hours can be used to accept backup data. One 24 hour time slot translates to 216 TB per day for each system. This is not a recommended configuration; it is included here for purposes of the discussion.

Figure 21-16 shows an overview of this scenario.

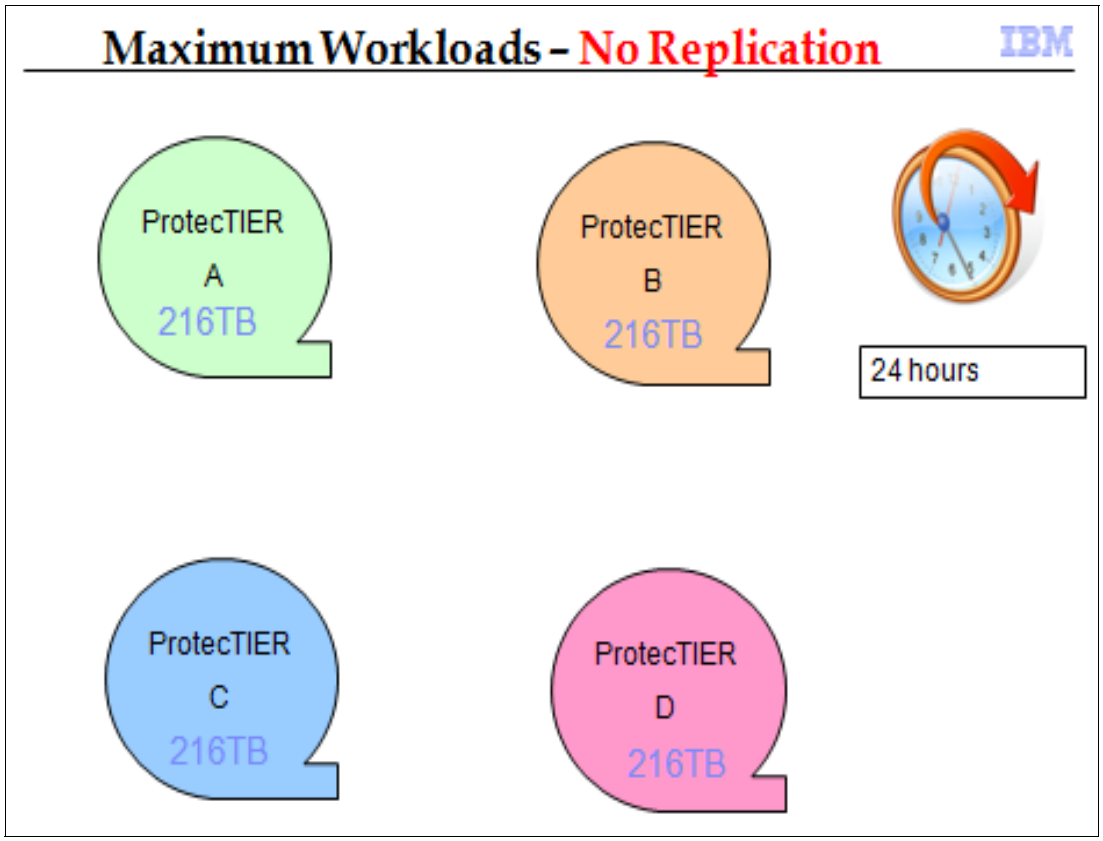

Figure 21-16 Maximum workload with no replication

Maximum workload with one replicated copy

For this example, all four ProtecTIER systems receive and replicate the same maximum amount of data that is possible in a 24 hour period. Because the workloads are equal, we can divide the 24 hour period into three equal time slots:

- One backup process (All four nodes accept backup at the same time.)
- One incoming replication processes
- One outgoing replication processes

With one data replication for each node, a maximum of 8 hours can be used to accept backup data. One 8 hour time slot translates to 72 TB per day for each system.

Figure 21-17 shows an overview of this scenario.

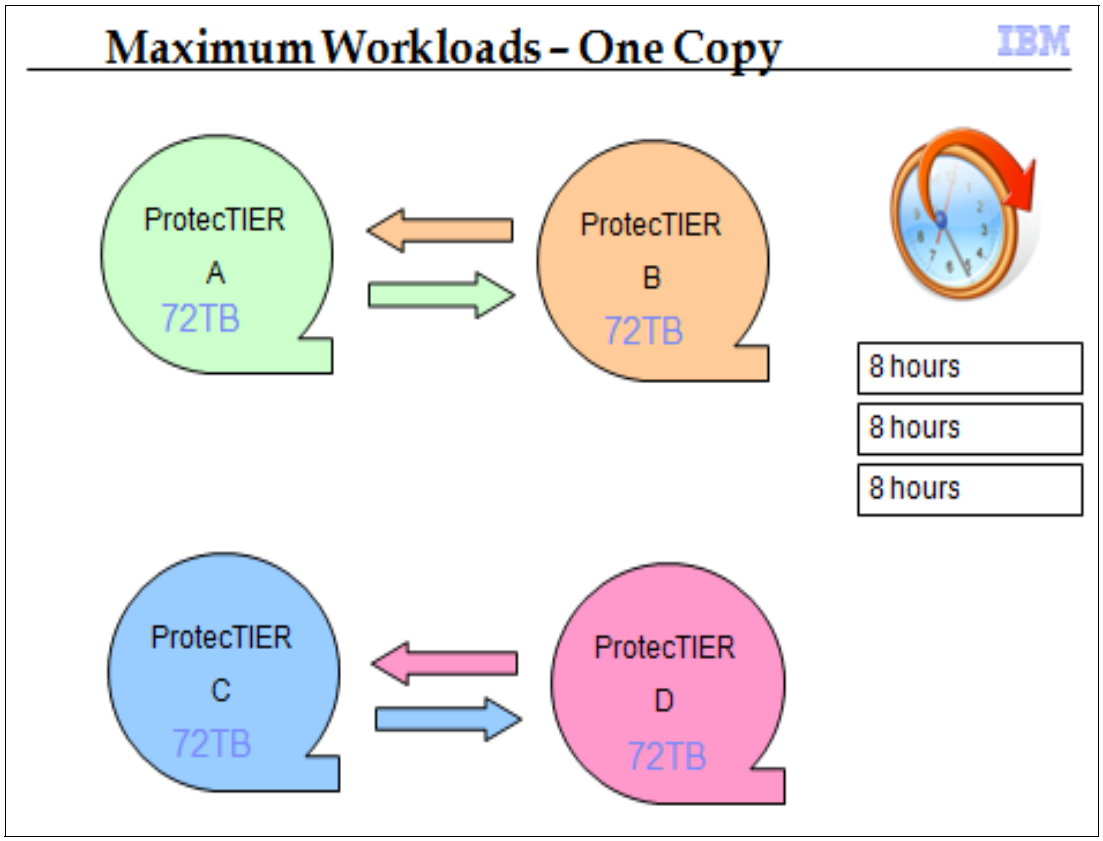

Figure 21-17 Maximum workload with one replicated copy

Two replicated copies

For this example, all four ProtecTIER systems receive and replicate the same maximum amount of data that is possible in a 24 hour period. Because the workloads are equal, we can divide the 24 hour period into five equal time slots:

- One backup process (All four nodes accept backup at the same time.)
- Two incoming replication processes
- Two outgoing replication processes

With two data replications for each node, a maximum of 4.8 hours can be used to accept backup data. One 4.8 hour time slot translates to 43 TB per day for each system.

Figure 21-18 shows an overview of this scenario.

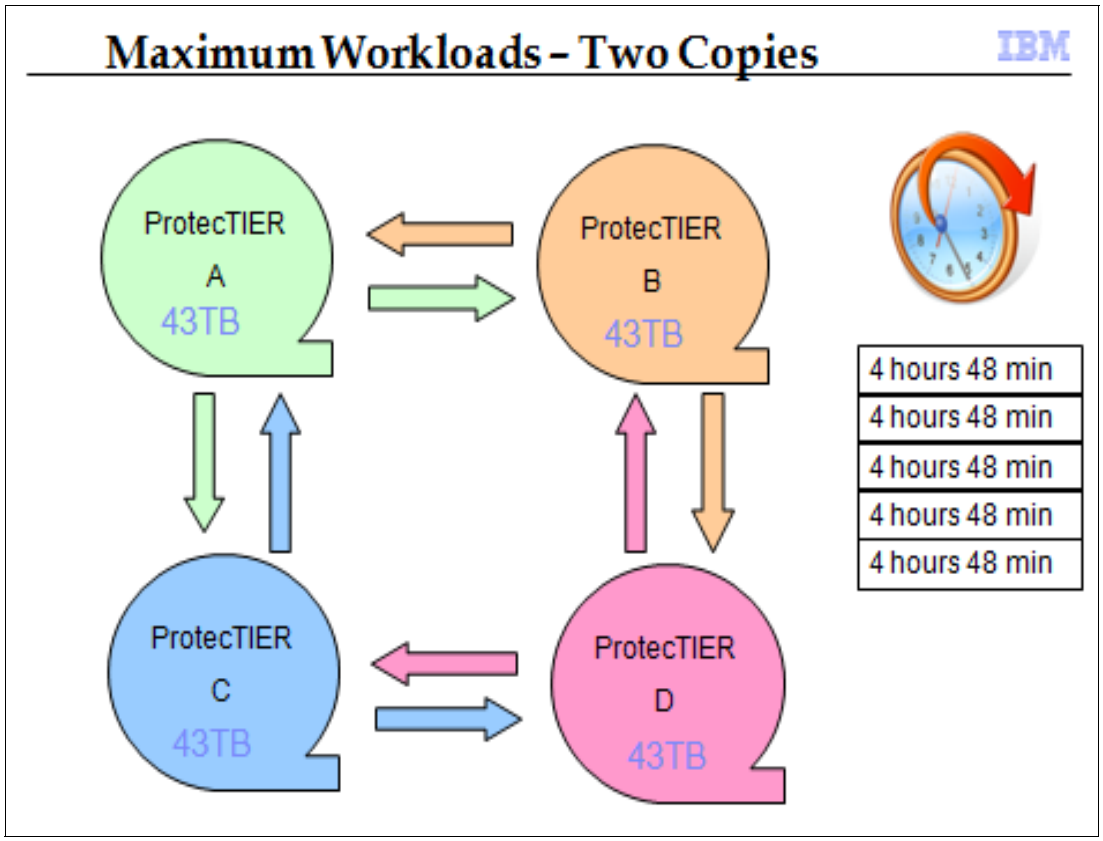

Figure 21-18 Maximum workload with two replicated copies

Three replicated copies

For this example, all four ProtecTIER systems receive and replicate the same maximum amount of data that is possible in a 24 hour period. Because the workloads are equal, we can divide the 24 hour period into seven equal time slots:

- One backup process (All four nodes accept backup at the same time)
- Three incoming replication processes
- Three outgoing replication processes

With three data replications for each node a maximum of 3.4 hours can be used to accept backup data. One 3.4 hour time slot translates to 30.6 TB per day for each system.

Figure 21-19 shows an overview of this scenario.

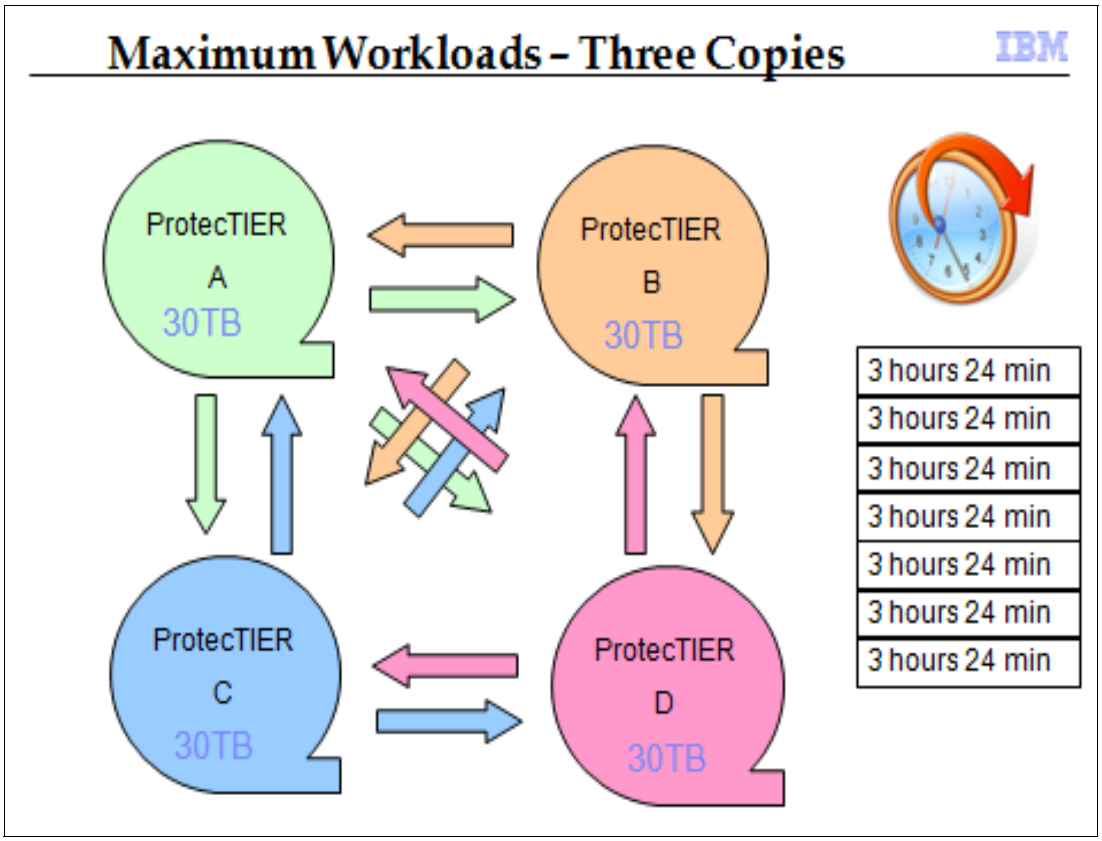

Figure 21-19 Maximum workload with three replicated copies

Table 21-1 depicts activity within different time frames in a 4-way many-to-many replication configuration.

Table 21-1 Example of a 4-way many-to-many replication configuration

Time frame	Activity
1	All systems process backups at 2500 MBps for 3.4 hours (30.6 TB).
2	System C replicates to B at 2500 MBps for 3.4 hours. System D replicates to A at 2500 MBps for 3.4 hours.
3	System C replicates to D at 2500 MBps for 3.4 hours. System A replicates to B at 2500 MBps for 3.4 hours.
4	System C replicates to A at 2500 MBps for 3.4 hours. System B replicates to D at 2500 MBps for 3.4 hours.
5	System B replicates to A at 2500 MBps for 3.4 hours. System D replicates to C at 2500 MBps for 3.4 hours.
6	System D replicates to B at 2500 MBps for 3.4 hours. System A replicates to C at 2500 MBps for 3.4 hours.
7	System B replicates to C at 2500 MBps for 3.4 hours. System A replicates to D at 2500 MBps for 3.4 hours.

21.12.2 Many-to-one replication

This many-to-one planning example uses ProtecTIER TS7650 appliances (Figure 21-20). Although the specific performance numbers would vary, this same process can be followed for SMB appliances and TS7650G Gateway systems when you size and plan a replication scenario.

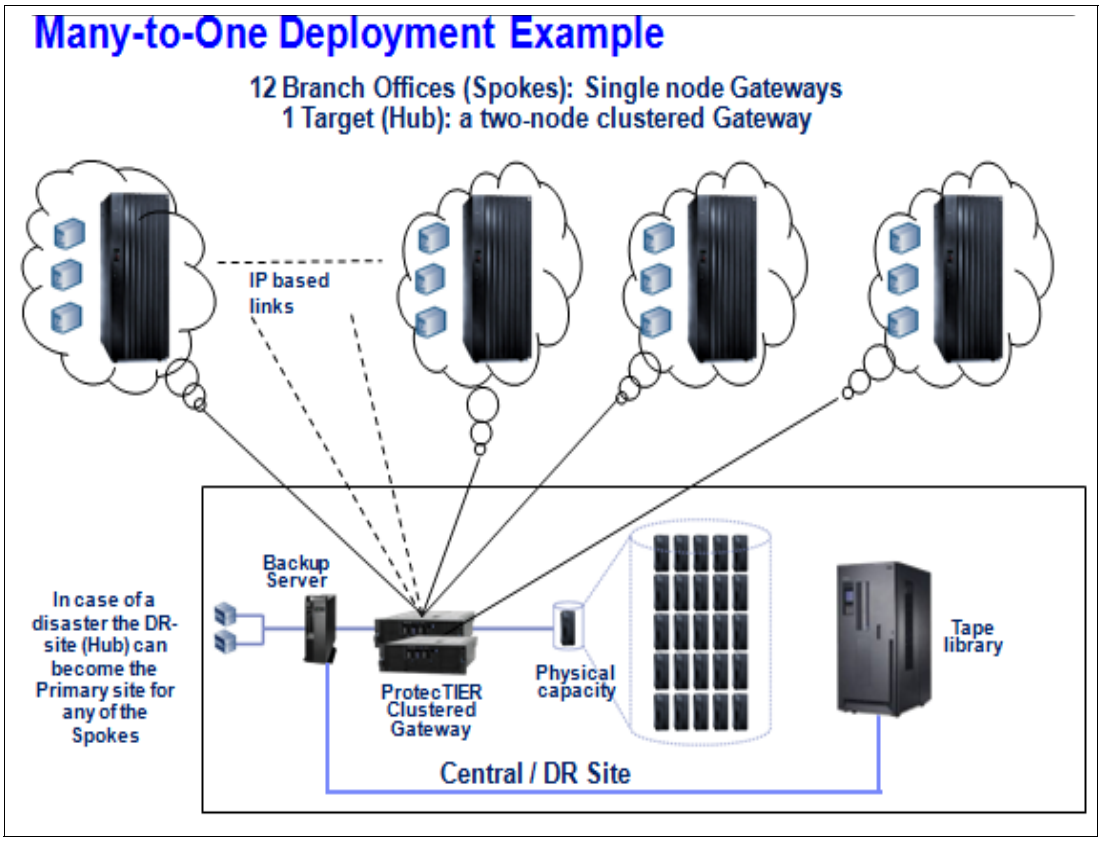

Figure 21-20 Many-to-one replication example

Assumptions

The case modeling is based on the following assumptions:

- A maximum environment of 12 Spoke systems and one hub system.
- Eight hour backup windows (hub and spokes).
- A 16 hour replication window.
- All windows are aligned, meaning that the eight hour backup window is the same actual time at all 13 ProtecTIER systems (hub and spokes).
- Adequate bandwidth between all spokes and hub.
- A 10:1 deduplication ratio throughout the system.
- Data change rate at spokes does not saturate the physical reception capabilities at the hub.

Maximum workloads assumed

Here are the values for the maximum workloads:

- Hub backup:
  - Eight hour backup window.
  - 9 TB per hour (2500 MBps).
  - 72 TB of nominal daily backup at the hub.
  - A 16 hour replication window.
  - 9 TB per hour (2500 MBps) replication performance.
  - 144 TB of nominal data can be replicated from all spokes.
- Spoke backup:
  - Eight hour backup window.
  - 144 TB for all 12 spokes = 12 TB of daily backup data per spoke.
  - 12 TB/eight hours = 1.5 GB per hour or 1500 MBps sustained for eight hours.
  - A spoke could potentially back up 72 TB of nominal data, but can replicate only 12 TB because of configuration constraints.

Sizing the repositories for hub and spokes

This section provides examples for sizing repositories for spokes and hubs. It also provides examples for calculating local backup space and incoming replication space.

Example of spoke repository sizing

In this example, each spoke can process up to 12 TB per day of local backup data, with a 10:1 deduplication ratio.

To size the spoke repository in our example, we complete the following steps:

1. Assuming a 10:1 deduplication ratio, approximately 1200 GB (or 1,200,000 MB) of new data must be replicated to the hub per backup. The total daily space for 27 incremental backups is calculated as follows:

   1200 GB x 27 incrementals = 32.400 GB (or ~32 TB) of physical space (for incrementals)

2. With a backup compression ratio of 2:1, add 6 TB for the first “full” backup (12 TB at 2:1 compression):

   32 TB + 6 TB = 38 TB of physical space for incrementals and full backup

3. Calculate the space that is necessary for spare capacity by multiplying the total physical space that is needed by 10%:

   38 TB x 10% = 3.2 TB of spare capacity

4. Calculate the total physical repository for each spoke by adding the total physical space that is needed and the spare capacity:

   38 TB + 4 TB = 42 TB

Example of hub repository sizing

The hub repository must be sized to handle 27 days of local backups and 27 days of incoming replication from all 12 spokes plus approximately 10% spare capacity.

Local backup space

In this example, the hub system can backup 72 TB in the 8 hour window. The first full backup at the hub requires 36 TB of physical space (72 TB @ 2:1 compression ration). With a 10:1 deduplication ratio, the hub accumulates 7.2 TB of new data for each of the next 27 days.

Here is the calculation for the local backup space:

36 TB +194.4 TB (7.2 TB x 27 days) = 230.4 TB

Incoming replication space

To calculate the incoming replication space in our example, we must complete the following steps:

1. Calculate the hub repository space for a full backup of all 12 spokes at 2:1 compression:

   (12 TB x 12 spokes)/2 = 72 TB of repository space

2. Assuming a 10:1 deduplication ratio, approximately, 1200 GB (1.2 TB) of new data per spoke must be replicated to the hub per backup. Calculate the new data received daily at the hub from all spokes:

   1200 GB x 12 spokes = 14.4TB of new data

3. The total daily space for 27 incremental backups is calculated as follows:

   14.4 TB x 27 incrementals = 388 TB of physical space

4. The total hub repository space that is necessary to accommodate the 27 incremental backups and one full backup is:

   230.4 TB + 388 TB +40 TB (10% spare capacity) = 464 TB for hub repository space

21.13 The backup application database backup

Figure 21-21 illustrates a typical backup and DR environment using the ProtecTIER product.

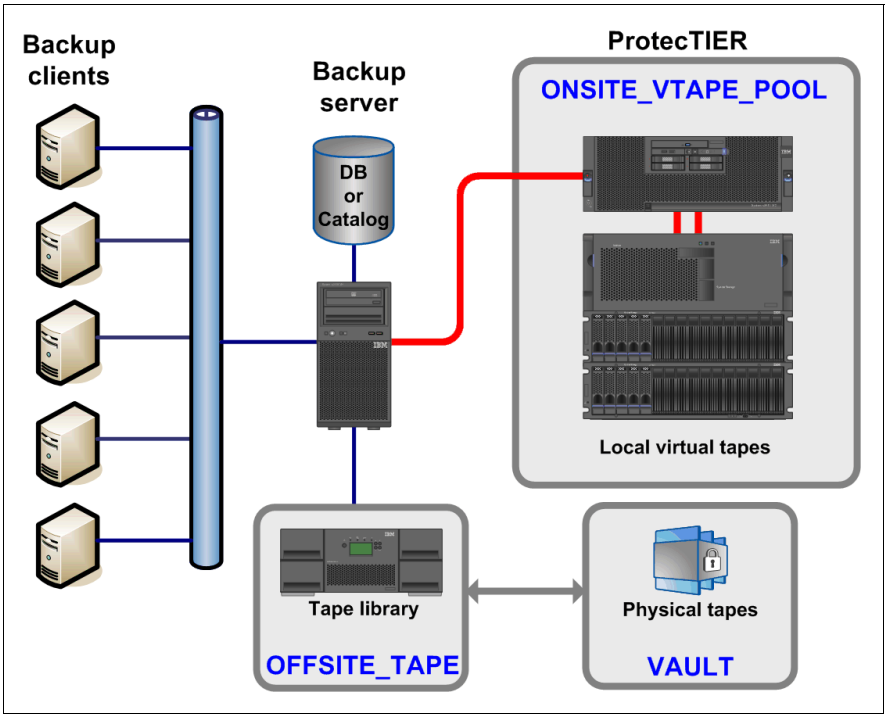

Figure 21-21 Typical backup and DR environment using the ProtecTIER product

The backup application environment is straightforward. The backup application servers are connected to storage devices (disk, real tape, or virtual tape). Every action and backup set that the backup servers process is recorded in the backup application database or catalog. The catalog is at the heart of any recovery operation. Without a valid copy of the database or catalog, restoration of data is difficult, and sometimes even impossible.

The ProtecTIER server provides a virtual tape interface to the backup application server, which you can use to create tapes, as represented by ONSITE_VTAPE_POOL in Figure 21-21 on page 404. The client can also maintain another tape library to create real tapes to take off-site, called OFFSITE_TAPE in Figure 21-21 on page 404.

ONSITE_VTAPE_POOL is where most client recoveries and restores come from. The key advantage of this architecture is that restoration occurs much faster because the data is coming from the ProtecTIER disk-based virtual tape rather than from real tape.

21.14 ProtecTIER Planner tool

Note: The ProtecTIER Planner tool is an IBM internal tool that is available to trained ProtecTIER specialists.

The core component for any capacity sizing and subsequent configuration effort is the ProtecTIER Planner. The primary methodologies to accurately size the required capacity and configure the disk and file system infrastructure depend on the correct usage of this tool.

The primary function of the ProtecTIER Planner enables field engineers to perform key activities when they size and configure the physical disk systems that are connected as back-end storage of the ProtecTIER server. The process starts at a high level, where a general capacity sizing is performed, based on key variables within the organization's environment.

Another aspect of the sizing process is to understand how many metadata and user data file systems are required based on disk technologies and RAID configurations to ensure correct performance. The ProtecTIER Performance Planner aids in this process.

The ProtecTIER Metadata Planner enables the field engineer to understand how many metadata file systems are required to support a repository of a certain size and the size of the metadata file systems.

There are other utilities within the ProtecTIER Planner, such as upgrading from a previous version (with the import of historical user data), and customizing any disk performance information based on unique user scenarios when planning performance.

For more information about capacity planning for a TS7650 Appliance environment, see *IBM System Storage TS7600 with ProtecTIER Version 3.3*, SG24-7968.

The ProtecTIER Planner should be used for capacity and performance planning while manually adapting it to a many-to-many environment at an early stage before the ProtecTIER deployment. The parameters to be considered are the same ones that are used in any replication deployment:

- System workload: Both local backup and replication, including incoming and outgoing
- Network bandwidth: Available and required between the PT nodes on the grid
- Time frames: Backup and replication windows or concurrent operation

22

Disaster recovery deployment with backup applications

This chapter provides best practices, general rules, and setup for disaster recovery considerations for the following backup applications:

- Tivoli Storage Manager
- Symantec NetBackup
- EMC NetWorker
- CommVault

It also provides effective techniques to perform disaster recovery processes and the related failback process after the primary location recovers from the disaster situation.

This chapter covers the following topics:

- Disaster recovery operations
- ProtecTIER replication overview
- Disaster recovery operations with VTL
- Disaster recovery operations with FSI
- Entering ProtecTIER DR mode
- The backup application catalog
- Single domain and multiple domains
- Replication best practices for OST
- Deploying replication with specific backup applications
- Symantec NetBackup deployment with ProtecTIER replication
- EMC NetWorker deployment with ProtecTIER replication
- CommVault

22.1 Disaster recovery operations

Disaster recovery (DR) is the process of recovering production site data at a DR location. Disaster recovery is useful if a disaster occurs or a situation occurs where the production (or primary) site goes offline. The hub, or DR site, can take the place of the production site until the primary site comes back online.

Entering DR mode is done by using the ProtecTIER Manager GUI when you are logged on to the DR ProtecTIER system. When you enter ProtecTIER DR mode, all incoming replication and visibility switching activities from the failed production site to the DR site (hub) are blocked. The DR node remains eligible for local backup and restore activity.

When the primary site is rebuilt or replaced, you can then use the ProtecTIER Failback function to return the systems to their normal state and resume the backup and replication operations. When the primary site comes back online, previously replicated and newly created objects can be moved to the main production site by using the failback process. The primary site then resumes its roll as a production site.

22.2 ProtecTIER replication overview

ProtecTIER replication enables virtual tape cartridges (VTL mode), file system volumes (FSI mode), or NetBackup images (OST mode) to be replicated from the primary site to secondary or multiple sites. This capability is intended to enhance disaster recovery (DR) and business continuity (BC).

Users can select some or all of their data objects to be replicated to one or more sites. ProtecTIER replication provides DR capability and enhanced availability in accessing backed up data for restore operations across sites. Replication can run in parallel with backup and restore activities or in dedicated scheduled time frames.

Data transfer is started based on trigger points in either on-demand continuous replication or replication schedules. Data verification and validation is done at the DR site to ensure the integrity of the transferred data before you make the virtual cartridge or tape available.

In summary, replication functions as follows:

- Replication is a process that transfers logical objects such as cartridges from one ProtecTIER repository to another repository.
- The replication function allows ProtecTIER deployment to be distributed across sites. Each site can have a single or clustered ProtecTIER environment.
- Each ProtecTIER environment has at least one ProtecTIER server.
- The ProtecTIER server that is a part of a replication grid has two dedicated replication ports that are used for replication.
- Replication ports are connected to the customer's WAN. By default, each ProtecTIER server is configured on two replication subnets.
- The replication groups are configured and managed in the ProtecTIER Grid Manager.

22.2.1 Replication data transfer

When the replication action is started either manually or based on a policy, the source (primary) ProtecTIER system carries out the following procedures:

- Initiates the sync-cartridge function between its own (source, that is, primary site) repository and the destination (DR site) repository
- Reads the unique replication data units on requests from the remote ProtecTIER system based on what it is missing
- Sends the unique data elements, by using TCP protocol, over the WAN to the remote (DR) site

At the same time, the destination ProtecTIER system performs the following handshake actions in this order:

1. Calculates the relevant cartridges' sync point from where the replication must start.
2. Receives many data units concurrently as part of the replication action.
3. Verifies cyclic redundancy check (CRC) for all replicated data before it becomes available as part of the data object.

After the CRC check is successful, the system moves each of the verified data elements in to scope and makes it available at the DR site.

22.3 Disaster recovery operations with VTL

This section describes disaster recover operations in a Virtual Tape Library (VTL) environment. It describes managing cartridges with replication by using basic disaster recovery and visibility switching, cartridge replication requirements, and importing/exporting slots allocation and searching in a VTL.

22.3.1 Managing cartridges after replication

When the replication jobs complete, the handling of the replicated cartridges is determined by the replication policy. The cartridge is always replicated to the ProtecTIER shelf at the target site. When the source cartridge is ejected from the source system virtual library, one of two things happens: It is left off the shelf or it is automatically moved to the import export slot of the target virtual library. The automated cartridge movement is known as *visibility switching*.

These two options are commonly referred to as follows:

- Basic disaster recovery (DR)
- Visibility switching

In the following sections, we provide a brief overview of both options.

Basic DR

The replicated cartridge stays on the target shelf after replication completes. Basic DR is similar to disaster recovery where cartridges are kept on a physical shelf at the DR site or at a remote storage facility. When the source site fails, physical cartridges can be rapidly transferred to the DR location and imported in to the standby library. The same concept exists with the ProtecTIER product when you use the basic DR mode. Cartridges are on the virtual shelf of the target system and are ready to be imported to an existing or new virtual library.

Visibility switching

Visibility switching is similar to a warm backup site practice. Physical cartridges are shipped from the source location to the DR site and stored in the physical slots of the standby library. When a disaster is declared, the necessary cartridges are immediately available for recovery at the DR site.

Added value of visibility switching

In a virtual world, the difference in processor usage between the previously described modes is minimal. Importing cartridges from the virtual shelf to a library is fast and requires little effort. It can be done from anywhere (with network access to the system). Storing cartridges on the virtual shelf does not make the DR system less reactive for recovery. As a result, the RTO that a ProtecTIER replication-based DR solution offers represents an improvement over a physical tape-based solution.

The advantage of visibility switching is more versatile management for backup applications that support cross-site distributed tape library management of their catalogs. Backup applications that can manage multiple sites through a universal catalog or database can use *automated cartridge movement*. Automated cartridge movement easily moves cartridges from site to site without using any other interface other than the backup application itself. With a replication policy that is configured to use visibility switching, when the backup application ejects a cartridge from a source library the cartridge then appears at an import and export slot of a designated DR site library (pending completion of replication). Likewise, cartridges can be moved back to the source site library by using the reverse process. Having control of the cartridge movement through the backup application simplifies the process of cloning cartridges to physical tape at the target site.

By eliminating the usage of *physical tape* from the replication process, a few more steps by the ProtecTIER server are required if visibility switching is not used. Because backup applications cannot handle the same barcode at multiple sites, cartridges are visible in only one library at a time. Therefore, the ProtecTIER server does not permit a cartridge to be visible in two libraries even if the data exists in both locations. To create a physical copy at the DR site without visibility switching, the administrator must perform the following actions:

- Import the replica into the target library after replication completes.
- Import the source cartridge back into the primary library when the clone job completes.

Some well-known backup applications that support single domain are Symantec NetBackup, EMC NetWorker, and IBM System i BRMS.

For applications that do not support a single domain, the automated visibility switching mechanism is of no real value. Each library is managed by a separate entity with no shared knowledge of the replicated volume's content and whereabouts. In these environments, the local backup server must proceed through a recent backup catalog or database that describes the content of the associated data volumes. Every set of replicated data cartridges that is imported into a target library needs to be preceded by a recent catalog or database update.

After the backup server at the target site is updated, the creation of physical tapes requires the moving of the cartridges. The cartridges are moved from the shelf to the library and exported back to the shelf when you are finished. This part of the tape creation procedure is the same as it is in a single domain environment.

22.3.2 Cartridge replication requirements

- A ProtecTIER server does not replicate a cartridge if the target instance is in a library.
- Cartridges should be exported back to the shelf after cloning.

- A cartridge that remains in the target library after a cloning prevents further replication.
- Before you move cartridges into a source library, verify that a cartridge is not left in the target library.

22.3.3 Importing/exporting slots allocation in VTL

When you create a virtual library, the default is to create eight import/export slots. A low import/export slot count can affect your DR process. For this reason, allocate *more than eight* import/export slots for DR purposes (the maximum value is 1022 per library and 4096 per repository). Allocation of more import/export slots is also important if visibility switching is used and in cases of a DR strategy with a heavy cartridge ejection requirement is implemented.

Import/export slots: Create enough import/export slots so that all of your DR cartridges are processed in one single step.

For example, if you need to move 32 cartridges for your initial DR operation, create at least 32 import/export slots in your virtual library. This configuration reduces the DR complexity.

22.3.4 Import/export slots searching

When you move replicated cartridges from the virtual shelf to a library, the backup application needs to scan and find the new cartridges in the import/export slots. Scanning for the new cartridges is a time consuming action. Alternatively, you can create more empty slots (for example, when you create the library, choose X cartridges and X+100 slots). The slot locations are already known to the backup application and therefore reduce the scan time that is required to scan for new cartridges.

22.4 Disaster recovery operations with FSI

You can use the ProtecTIER product to configure replication policies to replicate file system's directories and all objects that are contained in these directories recursively to remote ProtecTIER repositories without any disruption to the operation of the file system as a target for backup. It is possible to define up to 64 source directories and up to three remote ProtecTIER destinations per replication policy. The replicated data in the remote destination can be easily used to restore data in the case of a disaster recovery, or in the case of a disaster recovery test (without any interruption to the backup and replication procedures).

It is important to allow the ProtecTIER product to supervise all the changes that are made to a directory, or to a set of directories, that are defined within a replication policy. Therefore, you should not disable a replication policy unless this policy is no longer considered relevant. If there is a scheduled maintenance of the network that is used for replication, it is possible (though not mandatory) to *suspend* the replication to a specific destination. *Suspend* allows the ProtecTIER product to continue supervising all of the changes made; it does not attempt to send the replication data through the network for the time that is defined by the suspend operation (the suspend operation is limited in time).

In the case where a policy was disabled for some reason, a new Replication Destination Directory (RDD) must be defined to re-enable the policy. The ProtecTIER product does not need to replicate all of the data from scratch if the old RDD was not deleted. It needs to create only the new structure and metadata within the new RDD. Therefore, you should not delete the old RDD until a new cycle of replication to the new RDD is complete.

In summary:

- File system replication is a major feature of the ProtecTIER FSI project. It provides an inherent disaster recovery ability for ProtecTIER FSI customers.
- As opposed to VTL replication, FSI must replicate file's data and metadata, and the directory tree structure.
- FSI replication enables directory replication capability across repositories.
- FSI replication provides high availability in accessing backed up data across sites.
- FSI structural changes are replicated one by one to maintain the order of changes.
- Data replication is separated from structural replication to prevent delaying the order of changes.
- Upon a write operation, a file closure creates two replication triggers, one for data and one for structure.
- Structural changes create a structural trigger to replicate the change.
- Renaming a directory that is included in a replication policy is not allowed and results in a client error.
- Manual replication uses FSI snapshot to efficiently perform the replication.
- FSI replication policies are defined on a set of directories in the source repository and target directories in the RDD.

22.4.1 Replication Destination Directory

The RDD is a specially designated directory under the root of a file system at the DR site.

- The RDD is allocated and flagged upon policy creation to ensure that only one replication policy replicates to the RDD.
- After an RDD is allocated by a source policy, it cannot be used by any other policy.
- A replication policy is first created with empty RDDs.
- The source directories may not be empty when you create a policy, so part of them might not be synchronized at the destination until changes to it are applied.
- On FSI replication, a destination file or directory either exists in its synchronized state or does not exist there at all.
- Under no circumstances should a file or directory exist in the destination in an unsynchronized state. This a crucial assumption of the design.
- Replication can be ran manually, either on a policy as a whole, or a specific directory in it.

22.4.2 ProtecTIER FSI cloning

Cloning in FSI lets you perform disaster recovery tests at a replication destination repository while allowing replication of the original data from the source repository to the destination repository to continue. To perform a disaster recovery (DR) test, a space-efficient clone of a read-only replication destination directory can be created on the destination repository that is read/write accessible. All existing data in the repository remains accessible during cloning. If backup data is written to the clone, the relevant modified section is written as new data on the repository.

To clone a replication destination directory, complete the following steps:

- From the Systems Management view, click **Replication** → **Replication Directory** → **Clone a replication directory**. The Clone a replication directory window opens.
- In the Source directory pane, select the file system from which to clone the replication directory.
- Click **Browse** to search for and select the directory name of the replication directory you want to clone.
- In the Target directory pane, select the file system to which to clone the replication directory.
- Click **Browse** to select the new directory path and enter the new directory name to which the contents of the replication directory are to be copied.
- Click **Clone**. A message displays that it might take time for all of the directory contents to be copied.
- Click **Yes** to confirm and begin the operation. The Cloning progress window displays the source and destination paths, the total number of files to be cloned, the start time of the operation, and the number of files that are cloned per second.
- Click to hide the Cloning progress window.

22.4.3 Replication best practices for FSI

Here we describe some best practices regarding network replication in combination with ProtecTIER File System Interface (FSI).

File system naming

Choose meaningful names when you create file systems and network shares on the ProtecTIER server. Consider using the same name for the file systems that you create and the shares that you create in these file systems. Consider some obvious scheme to be able to re-create the relationship between file shares and file systems.

Backup server separation

Each backup server output should point to a separate directory that is configured in the ProtecTIER FSI. Because network replication policies are established between the source and destination at the directory level, creating a directory per backup server ensures consistency of individual backup servers. You can use this configuration to individually configure network replication for only a subset of your backup servers.

22.5 Entering ProtecTIER DR mode

Keep the process for entering DR mode simple to avoid potential complications and human errors when you handle an emergency situation. For VTL, the ProtecTIER DR wizard in ProtecTIER Manager at the DR site facilitates the initiation of DR mode.

Figure 22-1 shows an example of entering ProtecTIER DR mode.

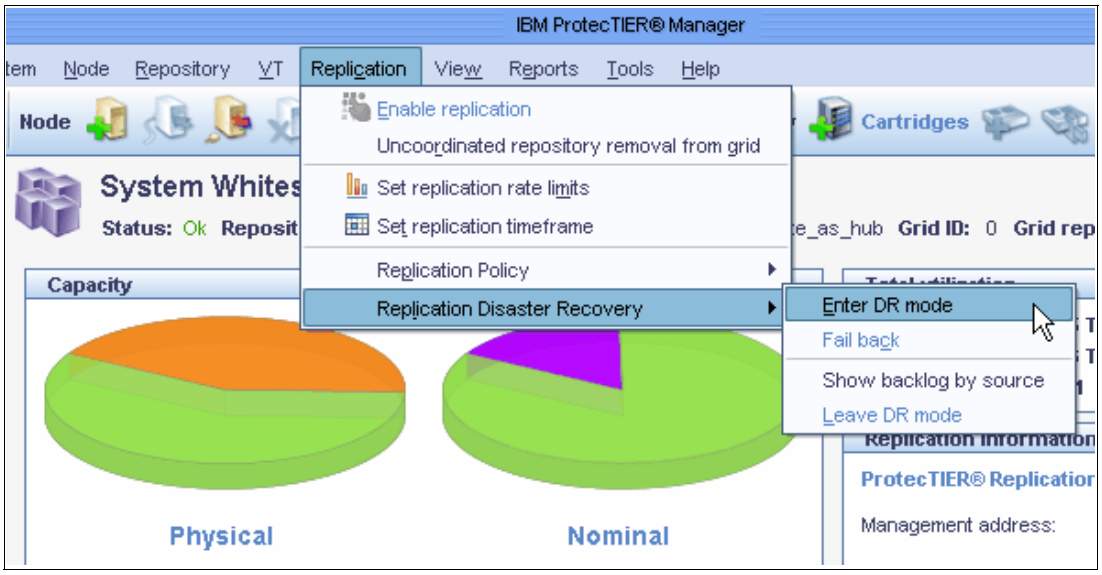

Figure 22-1 Entering ProtecTIER DR mode

As shown in Figure 22-1, to enter the mode, complete the following steps:

1. From the Systems Management view, click **Replication** → **Replication Disaster Recovery** → **Enter DR Mode**.
2. Choose the repository that you want to enter DR mode and confirm the selection.

An automatic procedure is run that blocks incoming replication to the DR site from the source selected.

Important: Entering and activating ProtecTIER DR mode allows you to have all the latest data visible at the remote site. However, the backup application must be restored manually or by using prepared semi-automatic procedures and scripts. You must be familiar with this backup application recovery process.

22.5.1 Working at the disaster recovery site

By default, replicated cartridges at the DR (hub) repository that were created at a primary site (spoke) are in read-only mode. If it is necessary to run local backups at the DR site, then use new locally created cartridges.

22.5.2 Inventory command options for a disaster recovery scenario

The inventory command options are used to filter cartridges in a ProtecTIER repository by using various criteria. Also, these options can be used to move cartridges that match a certain criteria.

Disaster recovery with the ptcli

The ProtecTIER ptcli provides a command-line interface (CLI) to manage your ProtecTIER servers and replication grid. Use the ptcli to perform the following tasks:

- Determine the consistency point at the hub when all cartridges from a specific spoke were fully synchronized (replicated).
- Automatically move cartridges in bulk from a shelf to a library and back.
- Determine the cartridges that were created at the hub during DR mode and then move them to the shelf so they can be replicated back to the spoke.

When you use the ptcli commands outside of DR mode:

- The ptcli snapshot might take up to 15 minutes to be created and populated.
- The snapshot is a static one, so it reflects the status of all cartridges only at the point in time it was taken.

A DR scenario and example procedure

Assume that a disaster occurs at a spoke when replication is running and a DR condition for the specific spoke is declared at the hub. You would complete the following steps:

1. Refer to the DR site (the hub) to determine when the last full backup occurred and recover the data from the DR to the primary site.
2. Use ptcli to produce a list of cartridges that were synchronized at the time of the disaster and a list of which cartridges were not synchronized at the time the last full backup at the spoke was completed.
3. Decide which cartridges to use at the DR site and use ptcli to move them (all or some) from the shelf to the library.
4. Save your results in a .csv file by using the **`-output`** command switch.

Regarding the output file:

- The output .csv file with results can be used as an input to a CLI **`move`** command. The .csv file is editable, and you can remove rows (each line represents a cartridge).
- You can also create your own barcodes file and use them as an input source for a **`move`** command.

22.5.3 Commonly used disaster recovery queries

Note: For more information about using the ptcli CLI, see Appendix A, “Command Line Interface”, of *IBM System Storage TS7600 with ProtecTIER Version 3.1*, SG24-7968.

Here are some of the commonly used disaster recovery queries that you can run:

- To create a snapshot of the current replication status, run the following command:

```
./ptcli InventoryRefresh --ip xxx.xxx.xxx.xxx -login file
```

Important: Creating a snapshot must be done before you run any other queries. Before you begin to filter the cartridges or to move cartridges by using the CLI, you must create a snapshot of the cartridges by running **`InventoryRefresh`**. This snapshot captures the current properties of all cartridges. Any filter or move operation is run based on the content of the snapshot. Moving a cartridge before you create a snapshot might fail if the snapshot is not up-to-date (not taken directly before the move).

- To list all in-sync cartridges, run the following command:

  ```
  ./ptcli InventoryFilter --ip xxx.xxx.xxx.xxx --querytype replica --query
  “in_sync = true” –login file –output /tmp/not_dirty_carts
  ```

- To list all unsynchronized cartridges marked with dirty bit, run the following command:

  ```
  ./ptcli InventoryFilter --ip xxx.xxx.xxx.xxx --querytype replica --query
  "in_sync = false" -login file -output /tmp/dirty_carts
  ```

- To list cartridges synchronized with the destination at a certain time range on the source, run the following command:

  ```
  ./ptcli InventoryFilter --ip xxx.xxx.xxx.xxx --querytype replica --query
  "source_time_for_last_sync_point > datetime(?2009-11-13 08:00:00')" -login file
  ```

- To list all cartridges that are replicated to repository 18 in grid 1, run the following command:

  ```
  ./ptcli InventoryFilter --ip xxx.xxx.xxx.xxx --querytype origin --query
  "destination_repository_id = 18 and destination_grid_id = 1" -login file
  ```

- To list all cartridges in barcode range BR0300 and BR0330, run the following command:

  ```
  ./ptcli InventoryFilter --ip xxx.xxx.xxx.xxx --querytype all --query "barcode >
  BR0300 and barcode < BR0330" -login file -output barcodes_file
  ```

- To move all synchronized cartridges to the shelf, run the following command:

  ```
  ./ptcli InventoryMoveFilter --ip xxx.xxx.xxx.xxx --querytype replica --query
  "in_sync = true" -destination shelf -login file
  ```

22.5.4 Returning to normal operations

This section explains the meaning of failback and failback policy, and provides the steps for using the ProtecTIER Manager failback wizard.

Failback

Failback is the procedure for returning visibility to the primary site and replicating any new data that is created at the DR site to the original (or replaced) primary repository. You use failback to return to the original working mode of the primary site.

Failback: The failback procedure can be initiated only while in DR mode. It is mandatory to perform the failback process before you leave DR mode at the remote site.

Failback policy

The DR site uses a *one-time replication policy* during failback that places a hold on all cartridges that are replicated to the primary site. This special policy also transfers the ownership of the relevant cartridges (created at the DR site when the primary site was down) to the primary repository or to its replacement in case the original primary repository is no longer functioning.

Figure 22-2 shows an example of the ProtecTIER Manager failback wizard.

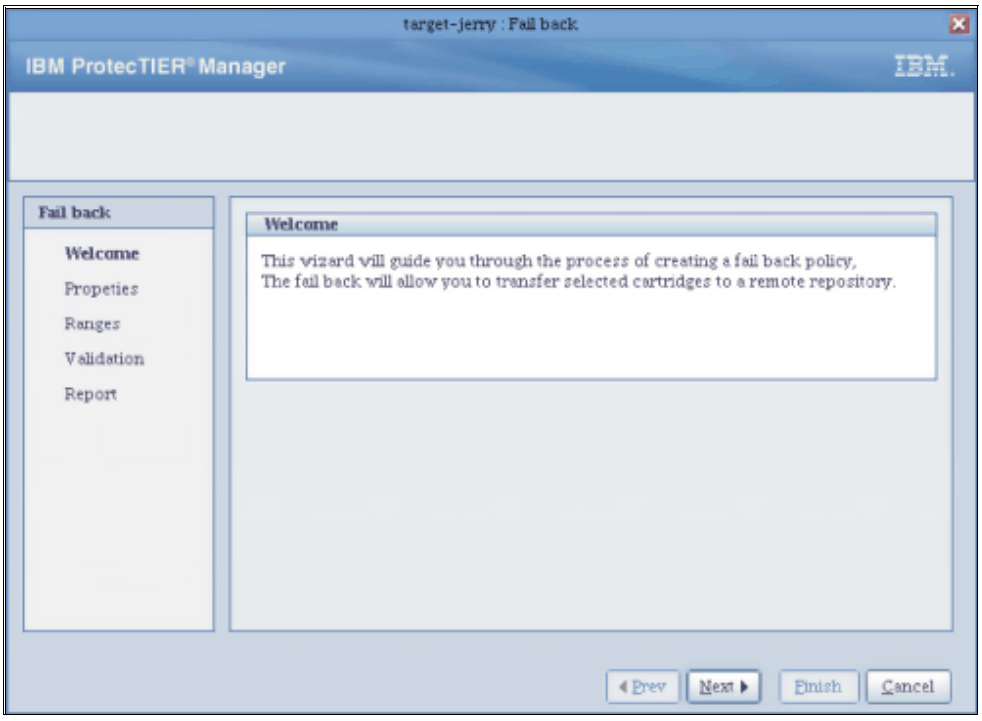

Figure 22-2 ProtecTIER Manager failback wizard

ProtecTIER Manager failback wizard

To initiate the failback process by using the ProtecTIER Manager failback wizard, complete the following steps:

1. Cartridges must be ejected out of the library to the shelf at the DR site. These cartridges include both the original replicated cartridges and any new cartridges that are created during DR operations.
2. Define a policy with all the cartridges that must be transferred to the primary site. This policy can be run only manually. The system log ignores runtime events for this policy.
3. Approve the execution of the policy. (This approval is a manual execution of the policy in the VTL.)
4. Close the failback wizard. The system provides information about the number of pending, running, and completed cartridges. ProtecTIER Manager presents this information to the user to indicate that the failback process is complete.
5. Delete the policy after the failback task completes.

22.6 The backup application catalog

Figure 22-3 illustrates a typical backup and DR environment using the ProtecTIER product. The backup application environment is straightforward. The backup application servers are connected to storage devices (disk, real tape, or virtual tape). Every action and backup set that the backup servers process is recorded in the backup application database or catalog. The catalog is at the heart of any recovery operation. Without a valid copy of the database or catalog, restoration of data is difficult, sometimes even impossible.

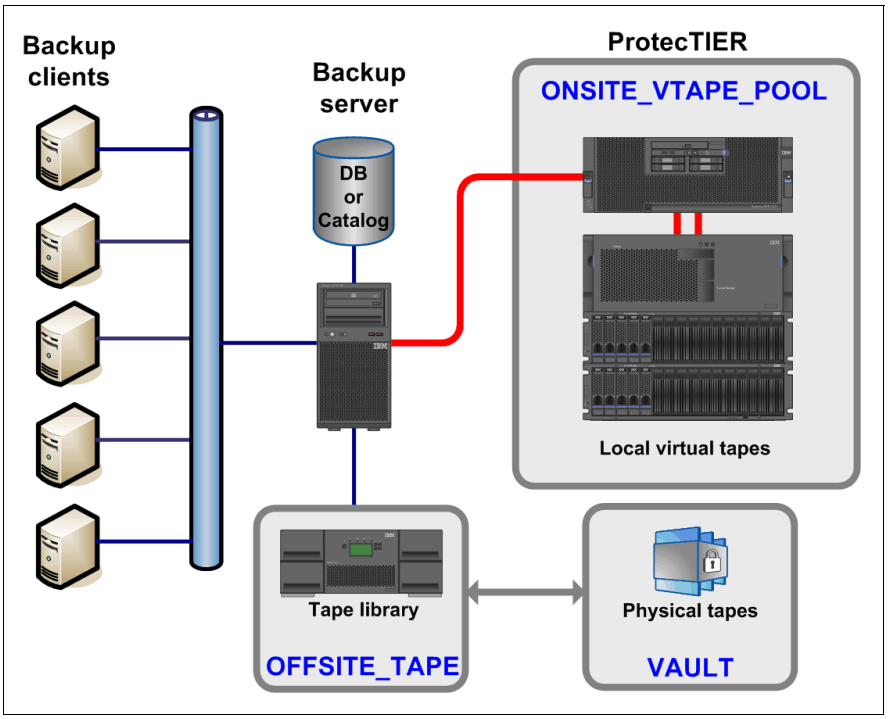

Figure 22-3 Typical backup and DR environment using the ProtecTIER product

The ProtecTIER server provides a virtual tape interface to the backup application server, which you can use to create of tapes, represented by ONSITE_VTAPE_POOL in Figure 22-3. The client can also maintain another tape library to create real tapes to take off-site, called OFFSITE_TAPE in Figure 22-3.

ONSITE_VTAPE_POOL is where most client recoveries and restores come from. The key advantage of this architecture is that restoration occurs much faster because the data is coming from the ProtecTIER disk-based virtual tape rather than from real tape.

Your Tivoli Storage Manager database or the backup server catalog in general should be backed up only after all backup sets are complete for a specific backup server. The database or catalog should be backed up in the ProtecTIER FSI-specific directory for that backup server.

Important: ProtecTIER FSI cannot be used to store the active version of the catalog that is constantly being updated. Do not place your Tivoli Storage Manager database or the backup server catalog directly in to the FSI. Only backups of supported applications are supposed to be writing to ProtecTIER FSI. For more details about Tivoli Storage Manager setup and best practices in an FSI environment, see 14.3, “Tivoli Storage Manager: FSI” on page 231.

22.6.1 ProtecTIER replication with Tivoli Storage Manager

The IP replication function of ProtecTIER provides a powerful tool that you can use to design a robust DR architecture. Because of data deduplication, you can now electronically place backup data into a vault while using less bandwidth. The ProtecTIER IP replication functionality can be used in a Tivoli Storage Manager environment, as shown in Figure 22-4.

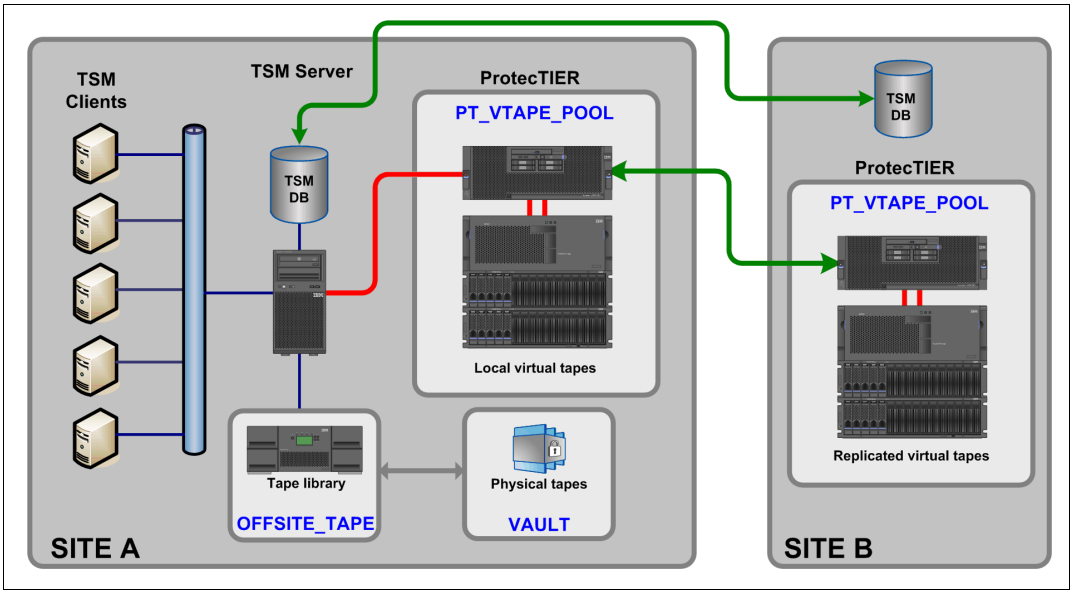

Figure 22-4 ProtecTIER replication and Tivoli Storage Manager

In Figure 22-4, the user chooses to replicate all of the virtual tapes in the PT_VTAPE_POOL off-site. The user also backs up their Tivoli Storage Manager database to virtual tapes, which are also replicated to Site B. If there is a disaster, you can restore the Tivoli Storage Manager server in Site B, which is connected to a ProtecTIER VTL. The ProtecTIER VTL contains the Tivoli Storage Manager database on virtual tape and all of the client ACTIVE files on virtual tapes.

22.6.2 Recovering the backup application catalog

There are several ways to obtain a copy of the catalog at the DR site:

- From a catalog backup on a virtual cartridge that is replicated to the DR site.
- From disk-based replication, or by other means.

If the catalog is backed up to a virtual cartridge, check if the DR site on which this cartridge appears is In-Sync with the primary site. If the cartridge is not In-Sync, you need to compare the cartridge's last synchronization time with the time of the last full backup.

To recover the backup application catalog from a backup on a virtual cartridge, you must work with the replicated cartridges on the hub to get an updated copy of the catalog back to the DR site. From the Systems Management window, select the **Replica** properties view on the Cartridges tab and use the following guidelines for each cartridge before you run the procedure for recovering the catalog.

Important: The procedure for recovering the selected catalog backup depends on the backup application and should be documented in the backup application's official documentation.

If the cartridge is replicated, either a red X or a green check mark appears in the In-Sync column. If the In-Sync column has a green check mark, then nothing further needs to be verified and this cartridge is valid for recovery.

If the cartridge is not marked In-Sync, refer to the Last sync time column. This column displays the last time each cartridge's data was fully replicated to the DR site. The cartridge marked with the most recent last sync time date should be used to recover the backup application catalog.

The sync time is updated not only when replication for this cartridge is finished, but also *during* replication.

Assessing the cartridges' status and synchronizing with the catalog

After the DR site backup application server is recovered, the user must review the status of the replicated cartridges. This review ensures that their replication is consistent with the backup catalog or database. This section explains the process for assessing cartridges' status on the DR site and synchronizing the backup application catalog with the cartridges.

Before you run a restore for disaster recovery, you must verify that the list of associated cartridges are marked as In-Sync with the primary site. Otherwise, an earlier full backup image must be used for recovery. The easiest way to determine the time of the last full backup is if you have a specific time each day where your replication backlog is zero (that is, there is no pending data to replicate and backups are not running). If not, you can assess the cartridges by recovering the backup application catalog and scanning it to find the last full backup where its associated cartridges completed replication.

Recovering the data

After the data is recovered, scan the backup application catalog and search for the full backup image you want to recover:

- Get the start and end backup time of the full backup image.
- View the list of cartridges that are associated with this full backup.

Run the **`PTCLI inventory filter`** command to filter the cartridges according to the following properties:

- In-Sync
- Last update time
- Last sync time

All the cartridges that are marked as In-Sync are valid for recovery. For those cartridges not marked as In-Sync, compare the last update time, which represents the last time that the replica was updated and the last sync point destination time. If the last update time is less than or equal to the last sync point destination time, the replica cartridge has consistent point-in-time. Otherwise, the cartridge is incomplete, or in transit. If the cartridge has consistent point-in-time, ensure that this time stamp is larger than the full backup image end time. This time stamp indicates that the cartridge contains all the required data for this recovery operation. Otherwise, the user must use a previous full backup image for recovery.

You might have a case where the cartridge sync point is after the backup start time, but before the end of the backup. This situation might happen in cases where replication is working in parallel to the backup. If the backup has many cartridges, the first cartridges might finish replicating before the backup ends and get a synchronization point earlier than the backup end time.

If the Last sync time flag on one (or more) of the cartridges indicates a time later than the backup start time, but earlier than the backup complete time, those cartridges need further inspection. Scan the backup application catalog for each of those cartridges and get the backup start time and the backup complete time.

If the Last sync time flag on all the cartridges indicates a time later than the backup complete time, your backup image was fully replicated.

Attention: When you process the cartridge list to find a complete set of DR tapes, you must track the date and time discrepancies. Compare the date and time values of the source master backup server and the source ProtecTIER system. The destination environment might be in a different time zone or might be set to the incorrect date and time.

Use the *source* date and time, rather than the *destination* sync tim, when you compare cartridge states to the backup catalog or database. The destination sync time should be used only to determine which cartridges are complete.

There could be a time difference between the source backup server and the source ProtecTIER server. Your administrator should be aware of the discrepancy, measure it regularly, and communicate the delta to your DR administrator or operators.

For example, if the backup server is two hours behind, a cartridge might have a sync time that precedes its backup complete time. If there is uncertainty about the time differences, compare the nominal size of the cartridge to the catalog or database value as an additional (not a substitute) layer of verification.

22.6.3 Tivoli Storage Manager reclamation and housekeeping

When you use Tivoli Storage Manager and ProtecTIER replication, configure your Tivoli Storage Manager reclamations to keep the Tivoli Storage Manager reclamation process from running during the ProtecTIER replication window. For maximum efficiency, choose dedicated windows for ProtecTIER replication and backup server housekeeping.

22.7 Single domain and multiple domains

Backup application software can be set up in many topologies. From the ProtecTIER replication standpoint, there are two general setup methods for these environments:

- Single domain
- Multiple domain

The backup application catalog or database has an entry for each cartridge that is used for backup:

- Date when the backup was performed
- List of files that are associated with the backup
- Retention period
- Other backup application-specific information

The backup application supports one catalog or database per backup server instance. In many cases, the primary and remote (secondary DR) sites have two separate backup servers, each with its own database or catalog. To efficiently read replicated cartridges at the remote site, the remote backup server needs access to the actual catalog or database of the primary backup server, or an exact copy of it.

22.7.1 Single domain environment

In a single domain environment, the same backup application catalog (or database) is shared across the separate primary and secondary sites. In these environments, the catalog is always updated in real time on the locations of the cartridges (physical and virtual). For example, this type of environment is more commonly used with Symantec NetBackup (NetBackup) and does not work with most deployments of Tivoli Storage Manager.

22.7.2 Multiple domain environment

A multiple domain approach is more widely used. The backup application does not share a catalog between the primary (local) and secondary (remote DR) sites. This scenario is the most common with Tivoli Storage Manager environments. In this type of deployment, each of the backup servers in both the primary and the secondary locations has its own backup catalog.

22.8 Replication best practices for OST

This section describes the OST NetBackup application programming interface (API) capabilities. It also describes the typical considerations and best practices for replication in an OST configuration.

The OST interface is a NetBackup API that supports communications between NetBackup systems and OST-enabled ProtecTIER systems. The OST interface provides the following capabilities:

- The NetBackup software directs ProtecTIER systems about when to create, copy, delete, or restore backup images.
- Backup images can be replicated to up to 12 different ProtecTIER systems.
- Workloads and performance are balanced between NetBackup media servers and ProtecTIER systems.
- Detailed statistics about the state of the ProtecTIER systems are available through the API.

With OST, a ProtecTIER server can be integrated with NetBackup to provide backup-to-disk without the need to emulate tape. ProtecTIER uses a plug-in that is installed on an OST-enabled media server. With this plug-in, a ProtecTIER server can implement a communication protocol that supports data transfer and control between the backup server and the ProtecTIER server.

Here are the typical considerations and guidelines for a replication in an OST configuration:

- Before any replication policies can be defined, a connection must be created between the repositories. The connections are created by defining the replication group in the ProtecTIER Grid Manager GUI.

- Replication connections for OST are bidirectional. A maximum of 12 repositories can replicate to each other at any particular time with up to 256 repositories existing in the entire grid.
- In an OST environment, replication policies and activities are managed by NetBackup through the duplication feature. ProtecTIER Replication Manager is only used to manage the replication grid, to create application groups, and to define the replication connections within the replication grid.
- Duplication takes longer to complete than the initial backup job and uses a larger amount of bandwidth. Duplication also taxes the NetBackup resource broker. This situation can slow down the rate that all jobs (backups, restores, and duplicates) run because resources are being requested simultaneously. Pay extra attention to plan the timing of all jobs to prevent delays.

22.8.1 The OpenStorage operating environment

There are two major components that make up the OST operating environment and they communicate through a Internet Protocol network:

- The storage server
- The plug-in

The *storage server* is an entity that runs on the ProtecTIER servers. It uses the major internal functionality of the ProtecTIER platform (such as HyperFactor, clustering, and replication).

The *plug-in* is a shared library (that is, a stateless software component) that is installed on the NetBackup machine. It is dynamically linked to the NetBackup application for data transfer to the ProtecTIER storage server emulation.

22.8.2 Automation of daily operation

The automation process of moving cartridges between sites and performing clone-to-physical-tape operations at the secondary site is more suitable in single-domain backup application environments. Some of the major backup applications, such as NetBackup, EMC NetWorker, and BRMS, allow for this type of environment.

Here is an example of a possible automation opportunity within a NetBackup backup application environment:

- Vault profile

  You can create a vault profile for ejecting cartridges from a library by running the following command:

  ```
  vltrun <vault profile name>
  ```

- You can inject a cartridge at the DR site to the library by running the following command:

  ```
  vltinject <vault profile name>
  ```

- Barcodes
  - You can eject cartridges from a library by running the following command:

    ```
    vmchange –res –multi_eject -ml <barcodeA:barcodeB:..:barcodeZ>
    ```

  - You can inject cartridges to a library by running the following command:

    ```
    vmchange –res –multi_inject -ml <barcodeA:barcodeB:..:barcodeZ>
    ```

- Inventory command

  This command scans the import/export slots and injects all available cartridges to the library:

  ```
  vmupdate –rt <robot type>-m<robot #> -empty_map
  ```

Scripting the inject/eject commands

Vault and inject/eject commands can be scripted to run periodically on the backup application host. The script triggers an automatic cartridge movement from the import/export slot to the library whenever the relevant cartridge is in the import/export slot. This process ensures free import/export slots. Example 22-1 shows an example of this script.

Example 22-1 Example script of inject/eject commands

```
#!/bin/csh
while (1)
vltinject myVault
sleep 600
end
```

Scripting the inventory command

Scripting the inventory command is not recommended because it scans the robot and therefore might take a long time to complete on libraries with many cartridges.

22.8.3 Gauging the replication completion status

You can gauge the replication completion status through the ProtecTIER Manager GUI or the ptcli command-line tool.

To use the GUI, click your Shelf (1), choose **Replica properties** from the Shelf cartridges (2), and sort the display list by the In-Sync status (3), as shown in Figure 22-5.

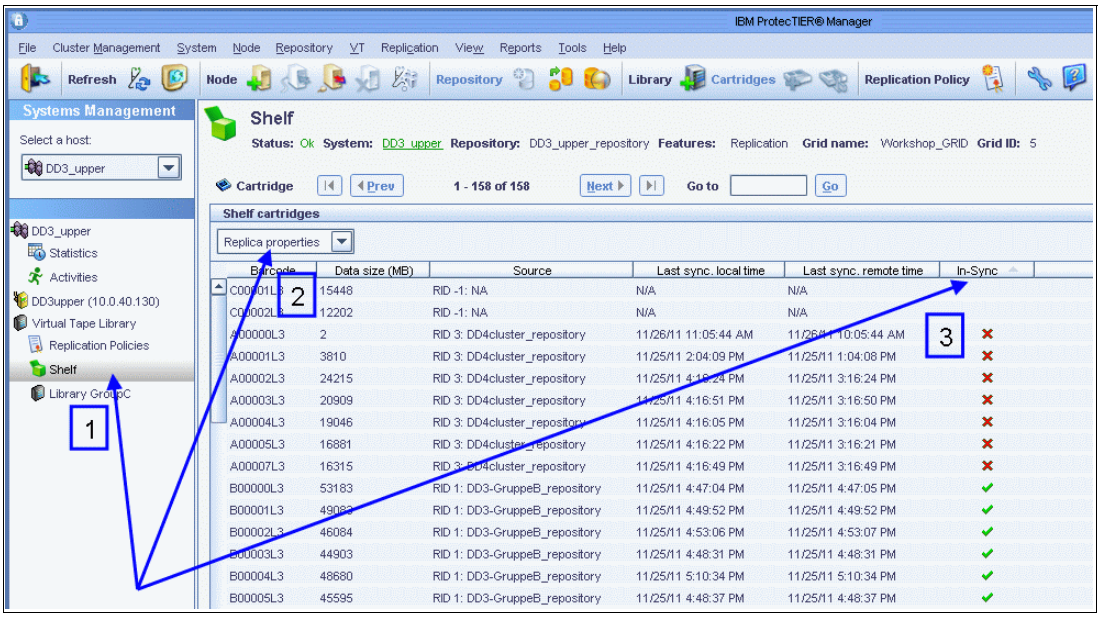

Figure 22-5 Gauging the replication status with the ProtecTIER Manager GUI

Should you prefer to gauge the replication status with the ptcli, complete the following steps. The ptcli command line is located within the `/opt/dtc/ptcli` directory on each ProtecTIER node. These commands must query the replication target machine (your hub machine) to successfully work.

1. To use the ProtecTIER ptcli, first create a ptcli profile by running the following command:

   ```
   ./ptcli –p ptadmin
   ```

   Using the ptcli tool: The ptcli tool requires usage of the ptadmin user. The ptadmin user is needed to run some of the commands to gather the information for the gauging of the replication completion status. The ProtecTIER product still obeys the single-logon rule.

   You cannot run these commands if there as an active ptadmin user already logged on. There is a **`force`** option for the **`ptcli`** command that you can use to end an already existing active session. Be careful when you use the **`force`** option, as issuing these commands within a script could affect an active admin session.

2. Generate a fresh copy of the inventory statistics by running the following command:

   ```
   ./ptcli InventoryRefresh --login ptadmin --ip 10.0.40.120 –force
   ```

3. Query for all out-of-sync cartridges by running the following command:

   ```
   ./ptcli InventoryFilter --login ptadmin --ip 10.0.40.120 --querytype replica
   --query in_sync=false --output /tmp/dirty_carts.csv
   ```

The output that is generated in the `/tmp/dirty_carts.csv` file lists all cartridges with pending replication tasks. Should the command in step 3 create no output, this situation indicates that the hub is synchronized with all spokes.

22.9 Deploying replication with specific backup applications

In ProtecTIER, VTL mode is generally the preferred method of replication to simulate the procedure that is used with physical cartridges. Implement the time-frame mode of operation so that for every 24 hour cycle, there is a backup window, and then a replication window. The user should ensure that there is enough bandwidth and time allotted so that there is no overlap and no replication backlog. Here is a typical operational flow:

- Perform regular daily backups to the ProtecTIER system during the defined backup window.
- The system should be set up so that replication starts and finishes before the next backup cycle starts.
- The user should have a complete and easily recoverable set of their latest daily backup, including the backup-application catalog image.
- If there is a disaster, the user can revert to the last completed set of backups, so the RPO is within the 24 hour window, which is typical for a service-level agreement (SLA).

22.9.1 Recovery point objective

When you design a ProtecTIER replication environment, one of the most important questions to consider is “What is the recovery point objective (RPO)?” How much lag time is acceptable for a backup that is written to virtual tape in Site A, to be replicated to Site B?

Tape-based DR

The RPO for tape-based DR is typically 24 hours. For example, consider a typical user case in which backups begin at 6 p.m. on Monday evening and the tape courier picks up the box of physical tapes at 10 a.m. Tuesday morning for transport to the vault. Therefore, on a typical day, there can be a 14 hour delay between the time the first backup begins and when the data is safely offsite.

However, if a disaster occurs before the courier arrives, the customer recovers the applications from the Sunday replication workload, which is a day behind, providing a 24 hour RPO.

ProtecTIER replication

With ProtecTIER replication, it is possible to get the backups offsite almost immediately (provided there is enough bandwidth). Because the ProtecTIER product is always working within the backup application, the RPO typically remains 24 hours.

22.9.2 Tivoli Storage Manager

This section provides an introduction to the various replication scenarios for when a disaster occurs at different stages of the replication process. Tivoli Storage Manager is a typical representation of a multiple domain environment, so the valid and most recent Tivoli Storage Manager server database backup must be available at the DR site.

Important: The valid and most recent Tivoli Storage Manager server database backup is a critical component of the recovery process at the disaster site.

As of Tivoli Storage Manager V6.1 ensure that you always have the file with all the records of Tivoli Storage Manager database backups (`volhistory`). If you do not have a copy of the volume history file, you cannot recover the Tivoli Storage Manager server database from your database backup.

Scenario 1: Replication complete

Tivoli Storage Manager allows the data replication to begin in a scheduled time frame after the backup cycle is completed, as shown in Figure 22-6.

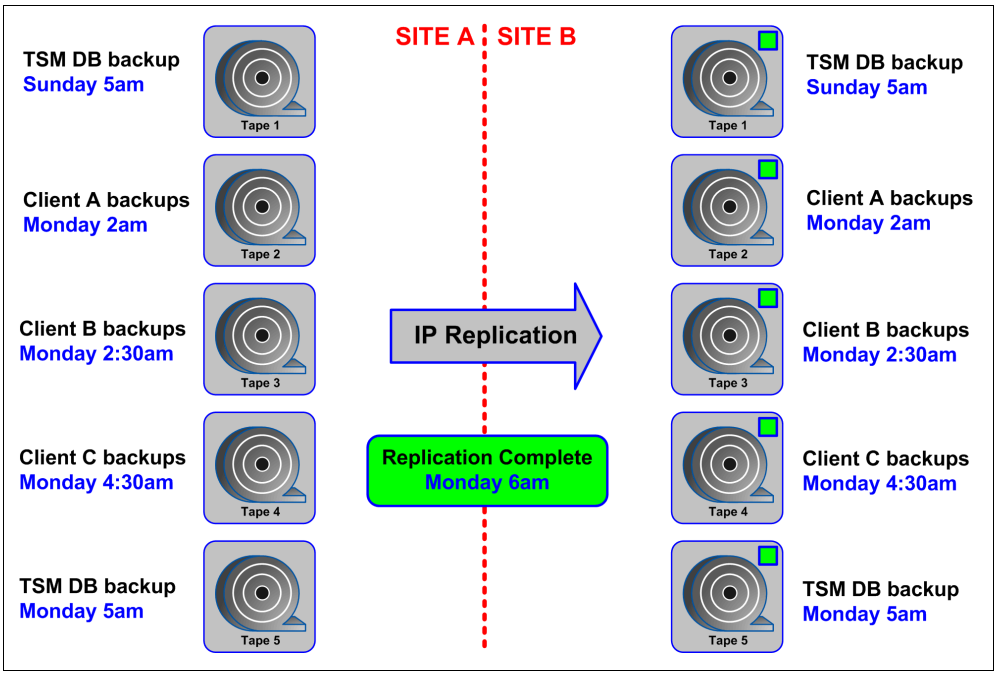

Figure 22-6 Scenario 1 - replication complete

For example, assume that the replication window ended at 6 a.m. and the disaster occurred at 6:15 a.m., 15 minutes after the replication cycle completes. (The status of all virtual cartridges is 100% replicated when the disaster event occurred.) The user starts the Tivoli Storage Manager server by using the Tivoli Storage Manager database backup that occurred on Monday 5 a.m. to Tape 5. The Tivoli Storage Manager database has knowledge of all the prior tapes, and restoration immediately begins from tapes 2 - 4, which hold the latest client backups.

Scenario 2: Replication incomplete

Suppose that the disaster event occurred at 5:15 a.m. before the nightly replication cycle completes. This scenario means that some of the backups from last night are not 100% replicated to Site B. For example, assume that the replication cycle ended at 6 a.m., but the disaster event occurred at 5:15 a.m. Monday morning. Tapes 4 and 5 were not replicated when the link went down.

In this case, because the Tivoli Storage Manager database is not entirely replicated, the user must restore from the last available Tivoli Storage Manager database backup on Tape 1 from Sunday at 5 a.m. Because Tapes 2 and 3 were also created after Sunday 5 a.m., they are not in the Tivoli Storage Manager database and cannot be used. Thus, the information that is on tapes 2 and 3 is lost.

Scenario 3: Auditing tapes and fixing the Tivoli Storage Manager database

Another possibility is when the most recent Tivoli Storage Manager database virtual tape is replicated, but not all the associated client backup tapes completed their replication before the disaster event occurred. For example, the disaster occurred 30 minutes before the anticipated replication cycle completion. In this case, tapes 1, 2, 3, and 5 replicated, but tape 4, because of the size of the backup data set that is stored on it, did not finish replication before the disaster.

In this case, the Tivoli Storage Manager server is restored by using the Tivoli Storage Manager database that is stored on Tape 5. However, because Tape 4 did not complete the replication, it must be audited to fix Tivoli Storage Manager database records that are related to the missing data on tape. Here is the Tivoli Storage Manager command for auditing Tape 4:

```
audit volume Tape4 fix=yes
```

The audit and fix must be performed for every tape that was not fully replicated when the disaster occurred. If the number of tapes that need this corrective action exceed the reasonable limits to run each job manually, you can consider auditing the pool that contains the inconsistent volumes. Use the following command to let Tivoli Storage Manager server audit every tape volume that belongs to the affected pool:

```
audit stgpool PT_VTAPE_POOL fix=yes
```

Auditing hundreds of tape cartridges can take a significant amount of time to complete (healthy cartridges are also audited). You must find a compromise between the time that is processed by auditing all the tapes at once, and running the audit job for each inconsistent tape separately.

Reclamation considerations

Another important item to consider when you use ProtecTIER IP replication in a Tivoli Storage Manager environment is the efficiency that reclamation causes. Reclamation is the Tivoli Storage Manager process that frees up space on tapes, and returns empty tapes to the scratch pool. Reclamation is accomplished by deleting records in the Tivoli Storage Manager database that are related to the expired data on tapes and moving the remaining valid data to other tapes, to efficiently use tape resources.

All reclamation-related data movement is recorded in the Tivoli Storage Manager database. To restore the accurate tape environment in the event of a disaster, all the database records related to the data movement must be replicated. If the disaster occurs when reclamation is running and the database backup that is available for restore does not reflect the status of data, an audit and fix of the volumes might be required. However, this audit process fixes only the inconsistent records in Tivoli Storage Manager database, but does not make the latest primary backup version of files available at the disaster site, unless they were replicated.

You can suppress the effect of the reclamation by using the **REUSEDELAY** parameter (valid for sequential-access storage pool, which includes virtual tapes). The **REUSEDELAY** parameter specifies the number of days that must elapse before a volume can be reused or returned to scratch status after all files are expired, deleted, or moved to another volume.

Tip: Set the **REUSEDELAY** period to the same value as the retention period of your Tivoli Storage Manager database backups. Only then you can safely restore Tivoli Storage Manager to any available point in time.

Pending state

When you delay the reuse of such volumes and they no longer contain any files, they enter the pending state. Volumes remain in the pending state as specified by the **REUSEDELAY** parameter.

Delaying the reuse of volumes can be helpful for disaster recovery under certain conditions. When files are expired, deleted, or moved from a volume, they are not physically erased from the tape. Only the database references to these files are removed. Thus, the files still exist on a sequential volume if the volume is not immediately reused. This prevents a situation where the reclamation ran after the last Tivoli Storage Manager database backup to virtual tape and reclaimed tapes are replicated. If the disaster occurs at a point after the reclamation ran, you are forced to restore from the less current database backup. If so, some files might not be recoverable because the server cannot find them on replicated volumes. However, the files might exist on volumes that are in the pending state. You might be able to use the volumes in the pending state to recover data by completing the following steps:

1. Restore the database to a point in time before file expiration.
2. Use a primary or copy volume that is rewritten and contains the expired file at the time of the database backup.

If you back up your primary storage pool, set the **REUSEDELAY** parameter for the primary storage pool to 0 to efficiently reuse primary scratch volumes. For the copy storage pool, you must delay the reuse of volumes while you keep your oldest database backup.

Disabling automatic reclamation

You might need to disable automatic reclamation by changing the reclaim parameter of the sequential-access attributes by running the following command:

```
update stgpool <STG POOL NAME> reclaim=100
```

Consider disabling reclamation altogether if the user deployed the ProtecTIER product with small virtual tape cartridge sizes (for example, 50 GB or less).There is statistically less data to reclaim on a smaller volume than on a larger one. Everything on a smaller volume is likely to expire together, causing the tape to go immediately to the scratch pool.

In summary:

- Ensure that catalog and database backups are performed to virtual tape and replicated along with the daily workload each day. Perform a database backup and replicate the backup data at the end of each backup cycle.
- Create a separate tape pool for database backups.
- Consider adjusting the Tivoli Storage Manager reclamation processing to ensure that actions are in sync with the replicated database, which includes setting **REUSEDELAY** to provide a 2 or 3 day delay in the reuse of tapes that are reclaimed.
- Consider deploying more network bandwidth to decrease the backlog. Database synchronization becomes a more serious issue if the backlog exceeds one backup cycle (typically 12 - 24 hours).

Determining what is available for restoration at the DR site

To determine what is available for restoration at the DR site, you must answer two questions:

- Which database copy at the DR site is optimal for recovery
- What cartridges at the DR site are valid (replication completed)

Which database copy is optimal

Before you run a restore for DR, verify that the list of associated cartridges are replicated to the DR site. Otherwise, an earlier full backup image must be used for recovery. The time of the last full backup is the specific time each day where your replication backlog is zero (there is no pending data to replicate).

If you do not know the specific time each day when your backlog is zero, assess the cartridges by recovering the backup application catalog and scanning it to find the last full backup where its associated cartridges completed replication. There are two primary methods to obtain a copy of the catalog at the DR site:

- From a catalog backup on a virtual cartridge that is replicated to the DR site
- From disk-based replication

If the catalog is backed up to a virtual cartridge, then use the cartridge view of the library in ProtecTIER Manager to examine each of the cartridges that are used for catalog backup and find the most recent sync dates that are marked on the cartridges. If there are multiple backup copies, find the latest backup that completed replication. To recover the backup application catalog from a backup on a virtual cartridge, you must work with the replicated cartridges to get an updated copy of the catalog to the DR site.

Each cartridge contains a time stamp of a last sync time, which displays the last time that the cartridge's data was fully replicated to the DR site. The sync time is updated during the replication process, not just when the replication for this cartridge is finished. The cartridge marked with the most recent last sync time must be used to recover the backup application catalog.

Which cartridges at the DR site are valid for restore

When the Tivoli Storage Manager server at the DR site is recovered, review the status of the replicated cartridges to ensure that their replication is consistent with the Tivoli Storage Manager database. Use the available ProtecTIER Manager system Replicated Cartridges Status Report.

To eject a cartridge from a library, run the following command:

```
CHECKOUT libvolume
```

For example:

```
TSM:TUSCSON1>CHECKOUT LIBVOL <name of library> REMOVE=BULK FORCE=yes
CHECKLABEL=YES VOLLIST=<volume1,volume2,...volume3>
```

To inject or import (insert) a cartridge to a library, run the following command:

```
CHECKIN libvol
```

For example:

```
TSM:TUCSON1>CHECKIN libvol <name of library> search=BULK checklabel=barcode
status=SCRATCH WAITTIME=0
```

22.10 Symantec NetBackup deployment with ProtecTIER replication

This section describes the usage of the IBM ProtecTIER IP replication system in a NetBackup environment and describes the ramifications of possible scenarios that are related to DR.

Here are possible automation options within a NetBackup backup application environment:

- Vault profile
  - You can create a vault profile for ejecting cartridges from a library by running the following command:

    ```
    vltrun <vault profile name>
    ```

  - You can inject a cartridge in to the library at the DR site by running the following command:

    ```
    vltinject <vault profile name>
    ```

- Barcodes
  - You can eject cartridges from a library by running the following command:

    ```
    vmchange –res –multi_eject -ml <barcodeA:barcodeB:..:barcodeZ>
    ```

  - You can inject cartridges to a library by entering the following command:

    ```
    vmchange –res –multi_inject -ml <barcodeA:barcodeB:..:barcodeZ>
    ```

- Inventory command

  This command scans the import/export slots and injects all available cartridges to the library:

  ```
  vmupdate –rt <robot type>-m<robot #> -empty_map
  ```

22.10.1 Scripting the inject/eject commands

Vault and inject/eject commands can be scripted to run periodically on the backup application host (see Example 22-2). These commands trigger automatic cartridge movement from the import/export slot to the library whenever the relevant cartridge is in the import/export slot. This process ensures free import/export slots.

Example 22-2 Example script of inject/eject commands

```
#!/bin/csh
while (1)
vltinject myVault
sleep 600
end
```

22.10.2 Scripting the inventory commands

Scripting the inventory command is not recommended because it scans the robot and therefore might take a long time to complete on libraries with many cartridges.

22.10.3 Setting up NetBackup for backup and restore

NetBackup deployments typically use a schema of weekly full backups and daily incremental backups. There are two types of incremental backups:

- Cumulative: Backs up everything since the last full backup.
- Differential: Backs up everything since the last backup.

22.10.4 Setting up NetBackup for disaster recovery

When you set up NetBackup for disaster recovery, consider a number of key issues:

- NetBackup architecture: Does the NetBackup domain span across the primary and DR sites or are they two separate domains? This is a key step to understand and has strong implications in DR.
- Classification of clients (RTO): When a company plans for DR, each server is given a recovery time objective (RTO) that is dependent on the importance of its application and the associated data to the business. Servers with short RTOs (typically less than 24 hours) generally do not use backup systems for DR. These servers typically use clustering, volume mirroring, or some form of data replication to maintain business continuity. Servers with RTO greater than 24 hours tend to use tapes for DR. Servers are then prioritized into recovery bands of 24, 36, 48, or 72 hours, depending on business requirements.

  DR and production servers: Typically, only production servers are set up for DR. Test and development servers are out of scope for DR. However, the ProtecTIER product makes DR protection affordable for all applications in any environment.

- Classification of Clients (RPO): Running alongside RTO is a recovery point objective (RPO). RPO is the point in time to which the server must be recovered. For most servers that use a tape DR scenario, the RPO is the point of the last complete backup before the disaster. For example, if a disaster strikes at 9:00 a.m., the RPO is the previous nightly backup.

Single domain

To cater to these disaster recovery requirements, configure NetBackup with a single domain that spans both sites (NetBackup Clustered).

In this configuration, the master uses host-based replication to mirror the NetBackup databases and a clustering product to manage host failover. If there is a DR event, the NetBackup master's operations can seamlessly fail over to the DR site. As the NetBackup databases are replicated, all of the backup information is known at the DR site and therefore data restores can begin immediately.

22.10.5 Cross-site backups

There are two possible options for cross-site backup scenarios:

- Connect clients from one site through the IP protocol to media servers on the DR site.

  All backups are then in the DR site library and ready for restore. The primary downside is that large IP pipes are required and backups are limited to the speed of the cross-site network.

- Stretched Tape SAN.

  A local client backs up to a local media server, which then sends the data across the SAN to the DR site. Backups are in the DR site library and are ready for restore. The disadvantage of this option is the large SAN pipes that are required, and the back ups are limited to the speed of a cross-site SAN.

Downside of both options

Because normal backups are now in the DR library, regular restores are slower because the data must come from a DR library.

Here are some options and possibilities to partially eliminate this negative effect:

- Turn multiplexing off. To achieve the best restore performance (to meet RTOs), NetBackup must be configured without multiplexing.
- Dedicate individual volume pools of RTO tiers or clients. For optimum restore times (and with sufficient media in libraries), implement individual volume pools per client. In this case, there is no contention between media when you do restores. In the physical tape environments, where the number of tape drives is limited, this configuration is often impractical.
- Systems in current production can implement cross-site backups with client backups going to dedicated volume pools, but this is limited to 30 clients with low RTOs. With separate volume pools, you need separate backup policies per client.

If the NetBackup configuration at the DR site is not in the same domain as the primary site, then a different strategy is required. Because the DR site has no knowledge of the backups, tapes, and so on, that are used by the primary site, you must first get a copy of the NetBackup catalog from the primary site and load it in to the NetBackup master on the DR site. This task can either be done through disk replication or tape backup.

NetBackup catalog backups: NetBackup catalog backups are different from regular backups and need special handling to restore. Not having the catalog available at the DR site means that every tape must be imported to build the catalog, which is impractical and is not considered a viable option. With the catalog in place at the DR site, the tapes can be loaded into the library, the library inventoried, and restores can commence in a short time frame.

22.10.6 ProtecTIER disaster recovery with Symantec NetBackup

Discuss the following key concepts with the NetBackup architects and senior administrators within the user organization:

- In normal operation, back up to a local VTL.
  - Backing up to a local VTL provides quick backups and quick restores.
  - Because VTL replication is at the cartridge level, and only the deduplicated data is transferred, it reduces the bandwidth that is needed, as compared to traditional cross-site replication/backups.
- Split servers for DR in to their RTO classifications.

  Have servers for DR (usually production servers) split into their RTO classifications and plan for separate volume pools and backup policies. For servers with low RTO requirements, consider individual volume pools and backup policies.

- Turn off multiplexing (MPX) for all backups that require DR.

  Multiplexing is accomplished at either the storage unit level or backup policy level. Disable MPX for all backups that go to the ProtecTIER VTL.

- Use large fragment sizes.

  Fragment sizes are configured at the storage unit level. Large fragment sizes improve the restore performance of whole file systems.

- Disable storage checkpoints.

  Storage checkpoints have an adverse effect on the deduplication ratios.

- Disable software compression.

  Software compression might reduce the efficiency of the ProtecTIER deduplication and affect its factoring ratio.

22.10.7 Single domain versus two separate domains

After your architects and administrators understand the basic concepts, they need to decide whether to have one domain that spans both sites or two separate domains.

Single domain approach

With a single domain approach, the same NetBackup catalog is shared across sites and is always updated with the whereabouts of all cartridges.

Replicating cartridges

ProtecTIER replicates cartridges per the policies that are set by the user. Cartridges are copied on to a virtual shelf at the DR site.

Moving cartridges

Cartridges can also be moved by using the replication policy with the visibility switch option so that they are visible to the NetBackup application at the DR site (although the actual data is available to ProtecTIER on both sites). Moving cartridges includes the following functions:

Ejecting cartridges Eject (export) a cartridge from a primary library.

Injecting cartridges Inject (import) a cartridge in to the inventory at the DR site library.

This operation can be set manually or by using the NetBackup vault. Either way, it can be automated from within the NetBackup environment.

Separate (multiple) domains approach

If a separate (multiple) domains approach is used, the items that are listed in this section apply.

Replicating cartridges

ProtecTIER replicates cartridges per the policies that are set by the user. Cartridges are copied into a virtual shelf at the DR site.

Perform a catalog backup to virtual tape at the end of the backup window. Replicate it at the end of each replication cycle to the DR site. This approach ensures that at the end of every day (assuming a 24 hour backup and replication cycle) that the DR site holds a full set of replicated cartridges with a matching NetBackup catalog.

Disaster recovery

If a disaster occurs:

- Your first step is to get the NetBackup catalog restored on the DR site's NetBackup server by restoring the cartridges that contain the catalog.
- Your second step is to inject the cartridges from the DR shelf at the DR site ProtecTIER server into the library and perform an inventory.

After the NetBackup server is up and running with the DR repository, restores and local backup operations can resume at the DR site. After the disaster situation is cleared and the primary site is back online, the user should use the ProtecTIER failback procedure to move their main operation back to the primary site.

22.10.8 Disaster recovery scenarios

ProtecTIER replication reduces cross-site backup traffic because it replicates only deduplicated data. ProtecTIER replication also improves ease of operation (by enabling simple inject and inventory actions), and if there is a disaster or DR test, makes recovery easy to plan and implement. Deploying the ProtecTIER product in a NetBackup environment makes the business more secure and reduces the burden of NetBackup architects and administrators.

Single domain environment

In a single domain configuration, there are two possible disaster situations.

Clustered configurations with all operations complete

In a clustered configuration where all backups and all replication operations are complete and disaster occurs, no NetBackup recovery action is necessary. The NetBackup catalog at the DR site is up to date. In this case, take the following actions:

- Within the ProtecTIER server, move the cartridges from the virtual shelf to the import slots.
- Within NetBackup, the library inventory must be refreshed. Choose the option to import tapes.

After the inventory operation is complete, restores and local backups at the DR site can resume.

Clustered configurations with backup complete but replication incomplete

If the configuration is clustered and all backups are complete but the replication operation is incomplete, then the NetBackup catalog database at the DR site is up to date. However, because replication is not complete, roll back to the previous nightly catalog and cartridges set (RPO of one day). After the inventory operation is complete, restores and local backups at the DR site can resume.

Important: When you are working in a single domain NetBackup clustered environment *and* using the visibility switch option within the ProtecTIER server to move cartridges from the primary site directly into a DR site library, the catalog is always up to date.

Multiple domain environment

The following scenarios might occur in a multiple domain environment with NetBackup and the ProtecTIER product.

Stand-alone configurations with all operations complete

If you are working with a stand-alone configuration, and all backups and all replication operations are complete when the disaster occurs, then the catalog database at the DR site is *not* up to date. The NetBackup catalog recovery action is necessary.

Identify the latest backup catalog tape and load (import) it in to the ProtecTIER library at the DR site. After the library is inventoried, begin a standard NetBackup catalog recovery operation. After the recovery operation is complete, restores and local backups at the DR site can resume.

Stand-alone configurations with backup complete but replication incomplete

If you are working with a stand-alone configuration and all backups are complete but the replication operation is incomplete when the disaster occurs, then the catalog database at the DR site is *not* up to date. The NetBackup catalog recovery action is necessary.

Find the NetBackup backup catalog tape from the previous night and load (import) it into the ProtecTIER library at the DR site. After the library is inventoried, begin a standard NetBackup catalog recovery operation. After the recovery operation is complete, restores and local backups at the DR site can resume.

22.10.9 Determining what is available for restore at the disaster recovery site

First, determine which NetBackup database copy at the DR site is valid. Before you run a restore for disaster recovery, verify that the list of associated cartridges is replicated to the DR site. Otherwise, an earlier full backup image (usually the backup from the previous night) must be used for recovery.

Tip: The easiest way to determine the time of the last full backup is if you have a specific time each day where the replication backlog is zero (there is no pending data to replicate).

If not, then assess the cartridges by recovering the backup application catalog and scanning it to find the last full backup where its associated cartridges completed replication.

Ensuring that a copy of the catalog is available at the DR site

The best practice for ensuring that a copy of the catalog is available at the DR site is to use the ProtecTIER native replication function. Each day, the catalog should be backed up on a virtual cartridge after the daily backup workload completes so that the catalog is replicated to the DR site at the end of each replication cycle.

If the catalog is backed up to a virtual cartridge, use the cartridge view of the library in ProtecTIER Manager to examine each of the cartridges that are used for catalog backup to find the most recent sync dates marked on the cartridges. If there are multiple backup copies, then find the latest backup that finished replication. To recover the backup application catalog from a backup on a virtual cartridge, work with the replicated cartridges to get an updated copy of the catalog to the DR site:

- Each cartridge has a *Last Sync* time that displays the last time the cartridge's data was fully replicated to the DR site. (The sync time is updated during the replication and not only when the replication for this cartridge is finished.)
- The cartridge marked with the most recent Last Sync Time should be used to recover the backup application catalog.

22.10.10 Eject and inject commands from NetBackup software

Although the process can be manually scripted to enable automation, the easiest way of using the NetBackup commands for automating this process is by using the vault service within the NetBackup software.

Ejecting a cartridge from a library

Ejecting cartridges from a library can be initiated through the NetBackup GUI or by using the NetBackup **vault** option. If you are using the **vault** command, first run your vault policy by running the following command:

```
/usr/openv/netbackup/bin/vltrun<vault policy name>
```

At the end of the backup, eject the cartridge by running the following command:

```
/usr/openv/netbackup/bin/vltinject<vault policy name>
```

Inserting a cartridge to a library

Inserting a cartridge can be automated through commands. Update the media manager volume by running the following command:

```
/usr/openv/volmgr/bin/vmupdate -rt dlt –r
```

Recovering a master server from an existing database copy

You can recover the NetBackup catalog by using either the online method or the offline method.

- Online catalog backup

  To use the offline catalog backup method, from the NetBackup Help menu, click **Online, hot catalog backup method**.

- Offline cold catalog backup

  To use the offline cold catalog backup method, from the NetBackup Help menu, click **Offline, cold catalog backup method**.

For more information, consult the official NetBackup application documentation.

22.11 EMC NetWorker deployment with ProtecTIER replication

This section describes how to integrate the ProtecTIER product with the EMC NetWorker (NetWorker) environment for DR by using IP replication. You can find information about integrating the ProtecTIER product with NetWorker through VTL or FSI in Chapter 16, “EMC NetWorker” on page 253, and general best practices of ProtecTIER replication in Chapter 21, “ProtecTIER replication” on page 371.

There are several benefits to deploying ProtecTIER replication in the NetWorker environment:

- It reduces the administrative impact of managing the physical transportation of tape media.
- ProtecTIER replicates only the data changes to a secondary site so minimal network bandwidth is required. This situation is especially beneficial to customers who previously refrained from implementing electronic vaulting because of the network cost.

- In addition to DR protection, ProtecTIER can be used with NetWorker to clone data to physical tape at a secondary site for third copy protection or long-term archive purposes.

Here are some important NetWorker terms that we use in our description:

Datazone A group of computers that are administered by a NetWorker server. There is only one NetWorker server in a datazone with several storage nodes and several NetWorker clients. A datazone can be single site or it can span across multiple sites with remote storage nodes, protecting NetWorker clients at different sites, and communicating the backup information to the NetWorker server through a LAN or WAN.

Clone pool A collection of clone volumes that are used only for cloning purposes. NetWorker cloning copies data from backup volumes in backup pools to clone volumes in clone pools. You can use ProtecTIER virtual cartridges or AFTD (FSI share) as a backup volume and clone it to a clone volume. The clone volumes can be virtual cartridges, AFTDs, or physical cartridges.

Bootstrap A save set that is essential for NetWorker DR. A bootstrap is composed of the media database, the resource database, and NetWorker client file indexes. You should back up the bootstrap at least once a day to local storage devices on the NetWorker server. The bootstrap backup should be included in your daily NetWorker server backup routine or whenever changes are made to the NetWorker server. At least two copies of bootstrap backup should be created where one copy is kept onsite for operational recovery and another copy is sent off-site for DR purposes. With ProtecTIER, the bootstrap can be backed up to a virtual cartridge or an FSI share on a ProtecTIER system. It can also be replicated to a ProtecTIER system at a secondary site while using minimum TCP/IP bandwidth. This practice eases the system administration management of handling physical tapes for bootstrap backups.

22.11.1 Cloning physical tapes with ProtecTIER replication

This section describes cloning data at a secondary site. The ProtecTIER product replicates backup volumes on virtual cartridges on ProtecTIER server A to ProtecTIER server B, then NetWorker clones the replica virtual cartridges in ProtecTIER server B to physical cartridges at a secondary site. For cloning requirement and configuration details, see *EMC NetWorker Administration Guide.*[1]

LTO3 media: The ProtecTIER product emulates LTO3 tape media in NetWorker, but the size of the ProtecTIER LTO3 media might vary from physical LTO3 tape media. Ensure that you have sufficient physical tape in the clone pool. NetWorker supports cloning from one tape media type to another tape media type, and from multiple tape media to single tape media.

[1] You can find NetWorker documentation at the NetWorker Information Hub at http://nsrd.info/docs.html.

Figure 22-7 shows a single datazone deployment where the NetWorker server is at a primary site with VTL on ProtecTIER server A attached to it. On a secondary site, the remote storage node has a VTL and a physical tape library that is attached to it. Data is backed up to virtual cartridges of ProtecTIER server A at the primary site by the NetWorker server and replicated to ProtecTIER server B at the secondary site with ProtecTIER replication. Then, the remote storage node clones the data from replica virtual cartridges in ProtecTIER server B to physical cartridges. We use this diagram for the rest of the description in this section.

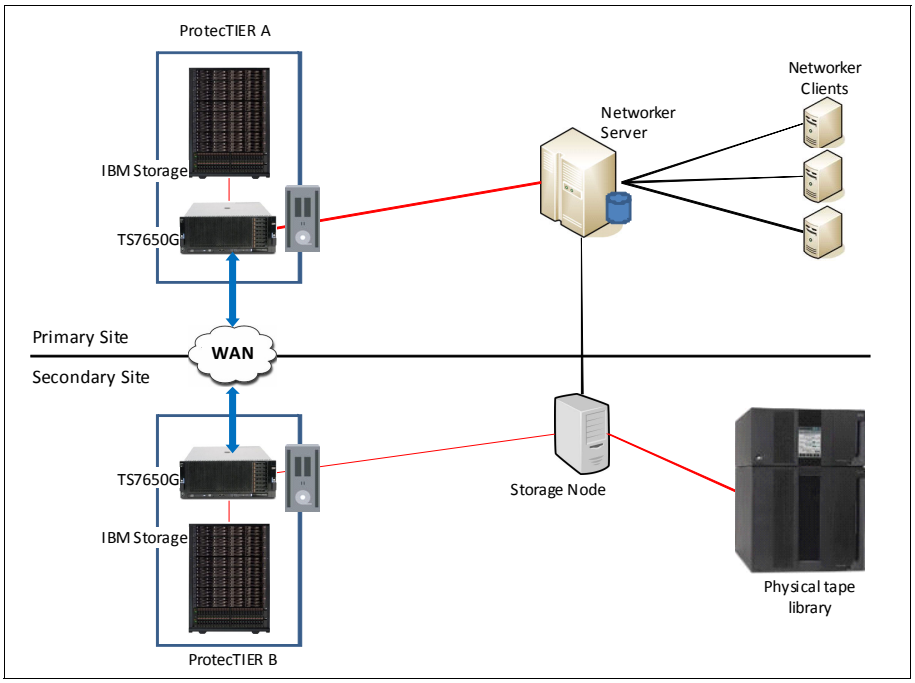

Figure 22-7 How NetWorker clones data to physical tape at a secondary site

Ensure that the following actions and prerequisites are completed and present before cloning:

- You must have a storage node at the secondary site that belongs to the same datazone.
- A VTL must be created on ProtecTIER server B and attached to the remote storage node.
- Define a cloning policy in NetWorker server.

- Define a replication policy at ProtecTIER server A with a target destination to a VTL of ProtecTIER serve B (Figure 22-8). By default, replication takes place after a backup is started and it runs to completion. This can make it difficult to determine the replication completion time. Use a replication schedule so that the cloning process can be scheduled upon replication completion.

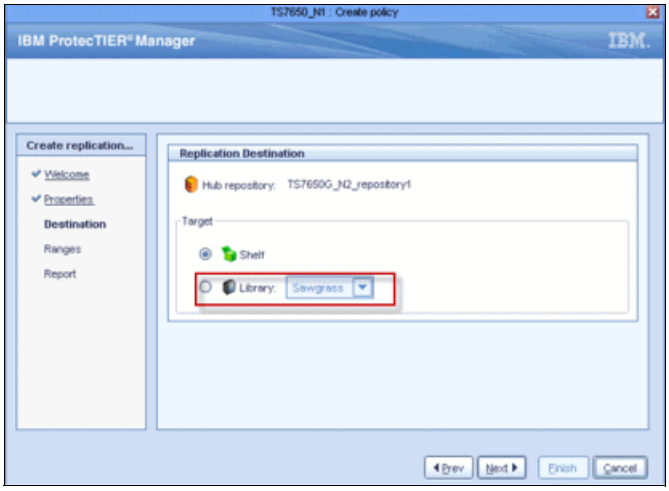

Figure 22-8 Set replication destination

The cloning process with the ProtecTIER product

The following steps detail the cloning process with ProtecTIER in an EMC NetWorker environment:

1. The NetWorker server writes save sets to virtual cartridges in ProtecTIER server A. ProtecTIER server A then replicates the virtual cartridges to the secondary site either immediately or during the configured replication schedule.
2. Upon completion of the replication, virtual cartridges are ejected (exported) from the source VTL and moved to the virtual shelf of ProtecTIER server A. This step preserves the unique volume ID in a NetWorker datazone when the replica cartridges (with identical volume ID as source tapes) are deposited (imported) in to the target VTL at the secondary site.
3. After source virtual cartridges are exported to the virtual shelf, ProtecTIER server B moves the replica virtual cartridges to the virtual import/export slots of the target VTL. Ensure that you have sufficient virtual import/export slots in the target VTL. For more information about how to allocate virtual import/export slots, see Chapter 21, “ProtecTIER replication” on page 371.
4. Run a library inventory in the remote storage node. The replica virtual cartridges are now accessible by the remote storage node and are ready for cloning.
5. Start cloning from replica virtual cartridges to physical cartridges.
6. Upon completion of the cloning process, eject the replica virtual cartridges to the virtual shelf of ProtecTIER server B so that the source virtual cartridges can be returned to the source VTL for next backup and replication jobs.

This process can be automated with NetWorker scripts so that each step takes place in sequence.

Data deduplication: Data is rehydrated when it is cloned to physical tape. This process is transparent to the NetWorker server and storage node. Thus NetWorker can restore the data from physical tapes without the presence of a ProtecTIER system.

22.11.2 Disaster recovery with ProtecTIER replication

Figure 22-9 shows a DR configuration where backup activity is performed at the primary site to the VTL of ProtecTIER server A. ProtecTIER server A replicates virtual cartridges to ProtecTIER server B at the secondary site. There is a standby server at the secondary site that is used for NetWorker server recovery during a disaster event.

For this description, a DR is a situation where the NetWorker server and the ProtecTIER system at the primary site are not available. During the disaster, you recover the NetWorker server on the secondary site by using the replicated cartridges of ProtecTIER server B.

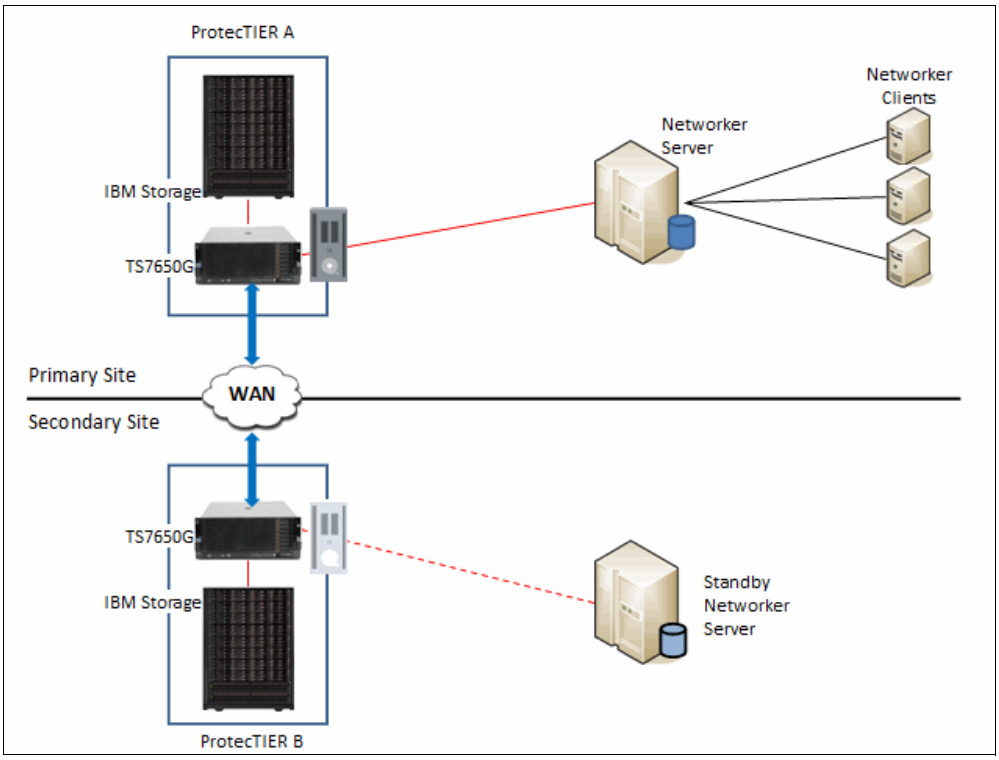

Figure 22-9 Disaster recovery with ProtecTIER replication

For more information about the NetWorker DR procedure, see *EMC NetWorker Disaster Recovery Guide*. For more information about the NetWorker CLI commands that are used in this description, see *EMC Commands Reference Guide*.

Disaster recovery preparation

Here are some suggestions about DR preparation:

- Configure EMC NetWorker based on the recommendations that are described in Chapter 16, “EMC NetWorker” on page 253.
- Send the bootstrap and the client file indexes to the same media pool. Create a separate media pool from a normal backup data pool for the bootstrap and the client file indexes backup.

- Write the bootstrap and client file indexes to the ProtecTIER server and replicate it over to the secondary site with ProtecTIER replication with high priority.
- During a disaster, the replica virtual cartridges on the secondary site are always in read-only mode. If write enabled cartridges are required, create virtual cartridges with a different barcode naming convention (volume ID). This action eliminates the possibility of barcode conflict (identical volume ID) during failback.

The NetWorker requirements for DR include:

- A good backup of the bootstrap. It should include the NetWorker media database, the resource database, and client file indexes that must be available at the secondary site during DR.
- An accurate record of the NetWorker server hardware, software, network, device, and media components at the secondary site.

The ProtecTIER system requirements for a secondary site include:

- ProtecTIER Manager must be installed and able to access ProtecTIER server B at the secondary site.
- If the Replication Manager is not hosted at the secondary site, see *IBM System Storage TS7600 with ProtecTIER User Guide*[2] to recover the ProtecTIER Replication Manager. From the ProtecTIER Manager GUI, enter DR mode at ProtecTIER server B.
- If there is no VTL created on ProtecTIER server B at the secondary site, then create a VTL with sufficient import/export slots.
- Ensure that the replica cartridge of the bootstrap is in a fully synchronized state.

Procedure to recover NetWorker Server

To recover an EMC NetWorker server, complete the following steps:

1. Replace the server with compatible hardware and operating system.
2. Configure the server with the identical network configuration, including the host name.
3. Install the server with the same version of NetWorker with patch levels that are equivalent to the original location.
4. If the links for the `/nsr` directory or any of its subdirectories, except for `/nsr/res`, are missing, then re-create these links.
5. Configure the NetWorker server storage device to detect the ProtecTIER VTL and tape drives by running the **`jbconfig`** command.
6. Make the replica cartridges or volumes so that the bootstrap is available to the NetWorker server. From ProtecTIER Manager, move the virtual cartridges from the virtual shelf to the import/export slot of the target VTL and import it into the VTL by running the following command:

   ```
   nsrib -b <volume name>
   ```

7. From the NetWorker server command line, inventory the VTL by running the following command:

   ```
   nsrjb –lnv – S <the slot with first volume> –f <pathname of first drive>
   ```

8. If you do not know the bootstrap save set ID and volume ID, see *EMC NetWorker Disaster Recovery Guide* for the steps to recover the server bootstrap.

[2] You can find ProtecTIER documentation at the IBM Support Portal at http://www.ibm.com/support/.

9. Recover the bootstrap save set by running the **`mmrecov -N`** command, which prevents overwriting the data in the tape. NetWorker scans the volumes for the appropriate save set and recovers the NetWorker media database and resource database.
10. After the media database and resource database are recovered, you must stop the NetWorker services and then rename the NetWorker resource directory (`/nsr/res`). The directory is recovered to an alternative location, as the NetWorker service is running during the bootstrap recovery process.
11. Restart NetWorker services and start recovering all the client file indexes from the NetWorker server.
12. The NetWorker server is now ready for client and storage node recovery. Perform a test backup and recovery with the standby ProtecTIER VTL.
13. Begin normal operations at the DR site.

Recovering client data

After the NetWorker Server is recovered at the secondary site, you can restore the client data. Use the NetWorker CLI or management console to generate the list of volumes to be recovered. Based on the list of volumes, review the status of the volumes (replicated virtual cartridges) to ensure that all the data is replicated to the secondary site.

After you determine the status of the recoverable volumes, restore the save set through NetWorker.

Resuming NetWorker server operation at the primary site

If new client data backups are done at the secondary site during the disaster, the updated bootstrap and client file indexes must be recovered at the primary site before you replicate the new virtual cartridges to the primary site.

Before you leave DR mode, complete the following steps:

1. Eject all virtual cartridges from the VTL in ProtecTIER server B with NetWorker and then perform a library inventory.
2. In ProtecTIER Manager, create a failback policy to replicate all new virtual cartridges that are created at ProtecTIER server B during the disaster to a primary site.
3. Start failback replication and wait until failback replication completes.
4. At the primary site, move the virtual cartridges from the virtual shelf to a VTL on ProtecTIER server A.
5. Start the NetWorker volume deposit process and perform a library inventory.

You can now resume normal backup and restoration operation.

22.12 CommVault

This section introduces the best practices for managing operations by using the CommVault backup application to replicate cartridges from the primary/spoke site to a DR/hub site. Figure 22-10 shows a typical CommVault DR scenario.

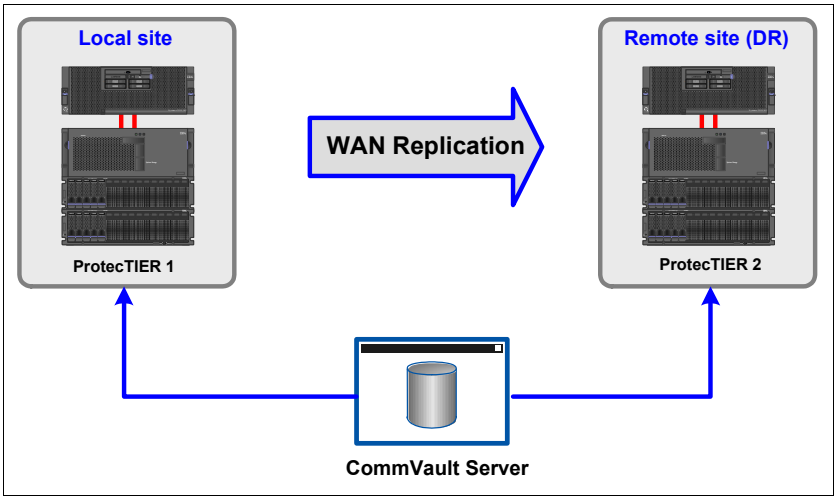

Figure 22-10 Typical CommVault DR scenario

22.12.1 Prerequisites

Here are the prerequisites for running the CommVault backup application:

- Verify that both the primary/spoke and the DR/hub ProtecTIER servers are connected, and are identified by the CommVault backup application.
- Verify that replication activities are running and finishing successfully by checking the active jobs log. To check this log, click **PT Manager** → **Repositories View** → **Replication's Activities** and select the active jobs.
- Verify that a visibility switch is defined on all replication policies. For example, the replication destination is the DR library. To check the visibility, click **PT Manager** → **Repositories View** → **Replication's Policies** and select the policy's name.

22.12.2 Running the CommVault backup operation

To run the CommVault backup operation, complete the following steps:

1. Eject media from a library through the CommVault Graphic User Interface (GUI):
   a. From the CommVault GUI, click **Storage Resources** → **Libraries** → **Local's site library** → **Media By Groups** → **Assigned**.
   b. Select the required cartridges and click **Export**.
   c. Define the new outside storage location. Click **OK**.
   d. Verify the change:
      - The CommVault GUI shows that the selected media moved to the new outside storage location.
      - The ProtecTIER GUI shows that the selected cartridges moved from the local library to the shelf.

- The ProtecTIER GUI on the DR site shows that the library cartridges are in the import/export slots.

2. Import (insert) media into a library through the CommVault GUI:
   a. From the CommVault GUI, click **Storage Resources** → **Libraries** → **Remote's site library** → **Import Media**.
   b. Click **Continue**.
   c. Verify the change:
      - The CommVault GUI shows that the selected media moved to the DR site library. You can access this view by clicking **Storage Resources** → **Libraries** → **Remote's site library** → **Media by groups** → **Assigned**.
      - The ProtecTIER GUI shows that the selected cartridges moved from the local library to the shelf (subject to visibility).
3. Eject (export) and import a cartridge through CommVault CLI:
   a. In a Windows environment, run the following command:

      ```
      <CV installation path> Simpana\Base
      ```

   b. Log in to the CommServe (the CommVault server) by running the following command:

      ```
      qlogin.exe
      ```

      Note: The user and password are the same as for the login to the CommVault console.

   c. Run the following command:

      ```
      qmedia.exe export -b <barcode(s)> -el <exportlocation>
      ```

      For example:

      ```
      qmedia.exe export –b XZXZX1L3 –el Comm_Shelf
      ```

   d. Verify the change:
      - The CommVault GUI shows that the selected cartridge exported successfully.
      - The ProtecTIER GUI shows that the selected cartridges moved from the local library into the shelf (subject to visibility).
   e. Run the following command:

      ```
      qlibrary.exe import -l <library>
      ```

   f. Verify the change:
      - The CommVault GUI shows the selected cartridge in the Assigned media window at the DR site.
      - The ProtecTIER GUI shows that the selected cartridges are located inside the DR site library.

22.12.3 CommVault resources

To obtain more technical details about CommVault, see the following documents:

- CommVault DR strategy and settings:

  http://www.commvault.com/solutions-disaster-recovery.html#t-1

- CommVault recovery, restoration, and retrieval using Data Agents:

  http://documentation.commvault.com/commvault/release_8_0_0/books_online_1/english_us/getting_started/getting_started.htm

- CommVault CLI for configuration and command usage:

  http://documentation.commvault.com/commvault/release_7_0_0/books_online_1/english_us/features/cli/cli.htm

22.12.4 Disaster recovery operations with OpenStorage

Data Replication in an OST environment is configured and ran within the NetBackup application. As a result, all DR operations are handled by using the NetBackup software. For more information about using ProtecTIER OST in DR operations, see Chapter 15, “Symantec NetBackup and BackupExec” on page 239.

A

ProtecTIER parsers

This appendix describes the ProtecTIER parsers that are used with various backup applications to improve deduplication ratios. It describes terminology, explains how metadata from backup applications hinders deduplication and what the ProtecTIER parser does to reverse this impact, and the supported parsers in ProtecTIER.

This appendix helps you understand what workloads benefit from ProtecTIER parsers, and what are the causes of low deduplication ratios. It also reviews several sample environments and describes whether they benefit from parsers.

This appendix also describes using the **`analyze_sessions`** utility to monitor a ProtecTIER parser.

The ProtecTIER parser

The ProtecTIER product excels when it finds large *matches* that can be deduplicated. Some common backup applications add metadata, also known as *backup application headers*, to the backup stream for various purposes. This metadata interrupts the large matches, which hinders deduplication. The ProtecTIER parsers separate the metadata from the backup stream dynamically, leaving the users' data to deduplicate without interruptions. When the data is restored, the ProtecTIER product adds the metadata back in to the data stream so that the backup application can use it.

The following sections describe terminology, why the ProtecTIER product needs a parser, and the causes of fragmented data, which hinder the matches and ProtecTIER performance and deduplication ratios. We also review several sample environments and describe whether they benefit from parsers.

Terminology

The users' servers send data to the ProtecTIER server, which deduplicates the data. On arrival at the ProtecTIER server, the data is first passed through the deduplication engine, which searches for the data in the repository. Some data is not found (the *new data*) and some is found (the *old data*).

The *change rate* is the ratio of new data to total data in the backup, that is, the percentage of data that the deduplication engine did not find.

The ProtecTIER server examines the old data and might decide that some of it cannot be deduplicated efficiently because doing so would introduce fragmentation and impact the restore performance. The percentage of data in these cases is called *old data not factored*, that is, the percent of data that could not be stored effectively.

In terms of deduplication, both the change rate and old data not factored represent data that was written to the ProtecTIER disk without being deduplicated. Their sum is termed the *system change rate*, that is, the percentage of data that could not be deduplicated because the data was either new to the repository or because deduplicating the data would cause unacceptable fragmentation.

Note: System Change Rate = Change Rate + Old Data Not Factored.

Explanation of how metadata from the backup application hinders deduplication

Some backup applications insert metadata into the backup stream for various control purposes by prefixing each block of user data with a small header that includes items such as the sequence number, session identifier, and cartridge barcode.

Figure A-1 shows how a backup application adjusts the sequence of user data by inserting its headers (the backup application metadata) at regular intervals in the tape image.

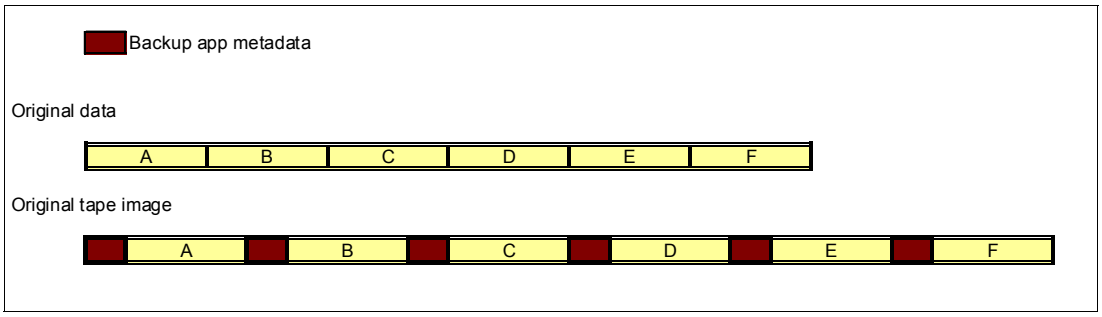

Figure A-1 Backup applications inserting metadata at regular intervals

The headers are always at regular intervals regarding the start of the tape cartridge. Figure A-2 shows how this spacing affects the tape image if the user data changes its configuration, perhaps because of an insertion of data (for example, the A-tag in Figure A-2). The block B in the original tape image is unchanged, but in the modified tape image, this block is split between B1 and B2, separated by a header.

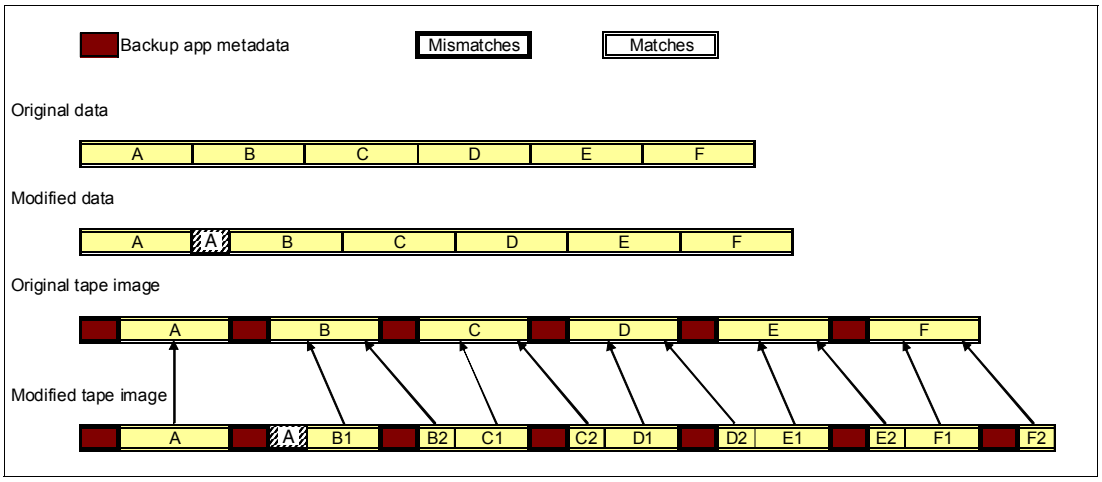

Figure A-2 Top - changes to data; bottom - fragmented data

The deduplication engine finds the B1 and B2 data blocks, but the backup application header between them interrupts their original sequence. The ProtecTIER server needs one pointer for B1 (pointing to part of the original data in B). A second pointer points to where the data of the backup application header is stored (not shown in Figure A-2). A third pointer is needed to point to the location that stores the data of segment B2.

The sequence of backup application headers introduces artificial change. Because the data in the backup image is shifted, it has a ripple effect all the way to the end of the backup or the end of the virtual cartridge. This effect multiplies the number of pointers that are needed to store the new data. Each of these pointers point to smaller data segments than the previous backup. Each generation of the backup potentially amplifies the fragmentation. The cost of adding a pointer to these fragments of old data eventually becomes so high, both in terms of extra I/Os and extra space, that the ProtecTIER server decides not to add the pointers. This old data is merged with adjacent new data and stored again as though it were new. This is how fragmentation leads to old data not factored.

ProtecTIER parser functionality

A ProtecTIER parser reverses the impact of the backup application headers by extracting them from the data stream. The backup data then deduplicates much better against previous backups and other data that is in the repository. When a user asks to restore the data, the ProtecTIER product reinserts the metadata in to the backup stream before it returns it to the backup application.

A ProtecTIER parser examines the start of each backup stream to determine whether the backup application is one that it needs to parse. If the backup application is not identified as needing a parser, the rest of the cartridge is processed normally. If the backup application needs parsing, the parser infrastructure sets up a mechanism in the data path that extracts the backup application headers from the user data. The backup application headers are compressed and stored separately. The remaining user data is now free of the backup application headers and the ripple effects that they cause, and is passed on to the deduplication engine.

The usage of a parser introduces a performance impact that is below 3%.

Deduplication ratios can increase by as much as the value of old data not factored, depending on the characteristics of the data, because the benefit is achieved by avoiding fragmentation. Section “Estimating the benefit of a parser” on page 452 explains how to estimate the expected improvement on your system.

ProtecTIER parsers: Support

Three backup applications, CommVault, Legato, and Tivoli Storage Manager, add metadata to the backup stream, so backups using these applications benefit from ProtecTIER parsers.

For VTL models of the ProtecTIER product, the following ProtecTIER parsers are available:

- CommVault: Parser available since IBM acquired ProtecTIER.
- Legato parser
  - ProtecTIER V2.3.4 and V2.4.1 introduced a parser for Legato 7.4 / 7.5 / 7.6 (May / June 2010)
  - ProtecTIER V2.4.7 and V 2.5.5 introduced a parser for Legato 7.6.1.6 (August 2011)
- Tivoli Storage Manager: ProtecTIER V3.1.4 introduced a parser for Tivoli Storage Manager V5.5 and later (October 2011).

For OpenStorage (OST) models of the ProtecTIER, product, the parsers are not relevant because OST is a Symantec NetBackup (NetBackup) function, and NetBackup does not insert regularly spaced metadata into the backup stream.

For CIFS models of ProtecTIER servers, the following ProtecTIER parsers are available:

- Tivoli Storage Manager

  Important: When you use the Tivoli Storage Manager parser with a CIFS model of the ProtecTIER product, use the **`DATAFormat = NATive`** option with the disk storage pool definitions. This value is the default value and is the format that the CIFS ProtecTIER Tivoli Storage Manager parser recognizes. Do not use the **`DATAFormat-NONblock`** option.

- CommVault

A Legato parser was not written for CIFS because there are Legato settings that make it unnecessary. Legato users should choose *Advanced File* rather than *File* disk storage when they set up the CIFS share on a ProtecTIER server. This setting does not require a parser.

Backup applications and parsers: These are the only backup applications that add metadata and the only ones that need parsers.

Sometimes, though rarely, new versions of backup applications introduce changes to the format of headers that they insert in the tape stream. This situation might cause the ProtecTIER parser to miss headers. Although this situation does not risk backup data in any way, it can cause deduplication to drop as the old data not factored increases. ProtecTIER Quality Assurance monitors the new versions of these backup applications, so check with ProtecTIER Support before you upgrade backup applications. Support can advise whether there is a change in the efficiency of the parser at that level or possibly provide a revised version of the ProtecTIER product if the revised parser is available.

What workloads benefit from the ProtecTIER parsers

The deduplication ratio that is achieved by a backup is affected by two key factors:

1. The change rate of the data (*change rate*): This rate depends on how much the data in the backup changes from day to day and some of the parameters that are set as part of the backup (for example, encryption, compression, and multiplexing).
2. The amount of data that is not factored (old data not factored): This rate can be high in environments where changes are interspersed throughout the data, causing fragmentation, such as might happen with a backup application that inserts metadata.

The ProtecTIER parsers are designed to reduce the amount of old data not factored and thus increase the deduplication ratio in environments where the backup application inserts metadata in the backup stream. A system with a larger percentage of data in this category (for example, 15%) benefits more from adding a ProtecTIER parser than a system with a smaller amount of data in this category (for example, 3%).

High change rate: If the system also has a high change rate, the benefit from the parser might be less noticeable.

Workloads that achieve lower deduplication ratios because of a high change rate should look at other avenues to increase the deduplication rate.

Background information: Causes of low deduplication ratios

There are two bases to good deduplication:

- The first and most important is multiple copies of the data. If there are not multiple copies of the data, there is no deduplication.
- The second is the similarity between (successive) copies. If the copies are similar (low change rate), then they deduplicate well.

Poor deduplication is caused by not enough copies or too much change between copies. However, sometimes successive copies are similar but the changes are small and evenly distributed. In this case, the storage subsystem cannot effectively store just the changes, but must rewrite some of the common data as though it were new. This phenomenon is measured by old data not factored.

The objective of parsers in the ProtecTIER environment is to focus on the old data not factored that is caused by the backup application headers in the backup stream. By removing these headers and storing them separately, the ProtecTIER parsers remove a cause of small and evenly distributed change that interferes with the underlying user data. This action prevents some of the fragmentation in the data and the extra pointers that are needed to track the data.

Often, the number of copies of the data (*retention*) is set by company policy. The old data not factored is controlled by characteristics of the actual and imposed changes. If the deduplication ratio is still low, then reducing the change rate is the best place to concentrate your efforts. There are many causes of a high change rate. Some common causes of high change rate are:

- Multiplexed backups: Some backup applications intermix backups from multiple sources in to the same (virtual) tape image. This situation causes the deduplication engine to try to search too many different sources for common data. A similar phenomenon occurs if the order of files in the backup does not remain stable.
- Backup of compressed or encrypted data: Sometimes unchanged compressed files deduplicate by using their previous backup as a reference, but in general, compressed or encrypted data does not deduplicate well.
- Cross-deduplication between files with similar content but a different layout: This situation can happen when different VM images are backed up. Each VM has similar files in it, but their layout in each VM's file system is different.
- Files with common data but high internal change: Some applications occasionally reorganize, defragment, or reindex their data. Backups that are made immediately after these operations are not likely to deduplicate well.

In extreme cases, the ProtecTIER product has a set of parameters that control the deduplication engine and the storage subsystem. ProtecTIER Level 3 support can recommend parameter changes to help in cases of poor deduplication.

Estimating the benefit of a parser

This section provides guidance for estimating the benefit of a parser.

Consider a hypothetical system with a change rate of 20% and old data not factored of 13.3% before you implement the parser. This situation means that the system change rate is 33.3% and the user sees one-third of the data change each backup, or two-thirds of the data deduplicated. Such a system could reach a deduplication ratio of 3:1 (before compression) if sufficient backup copies are retained. If the parser reduces the old data not factored to 5%, then the system change rate drops to 25%, and the user sees one-quarter of the data change each backup. The system could now reach a 4:1 deduplication ratio (before compression) if sufficient backup copies are retained. If compression achieves a 2:1 ratio, then this system would improve from 6:1 to 8:1.

ProtecTIER Support can review a system and provide the figures for change rate and old data not factored to help with this calculation.

Environments that benefit from parsers

The following four cases illustrate scenarios in two application environments: NetBackup and Tivoli Storage Manager. The cases that are provided demonstrate environments where a parser would be helpful in improving deduplication ratio and sometimes performance. Each case has a graph that shows the change rate and old data not factored values. The graphs were created by using data from a ProtecTIER server and a spreadsheet graphing tool.

Case 1: Low change rate and low old data not factored

This case shows a customer who is using NetBackup. The workload has a low change rate and low old data not factored values (Figure B-3). This workload does not need a parser.

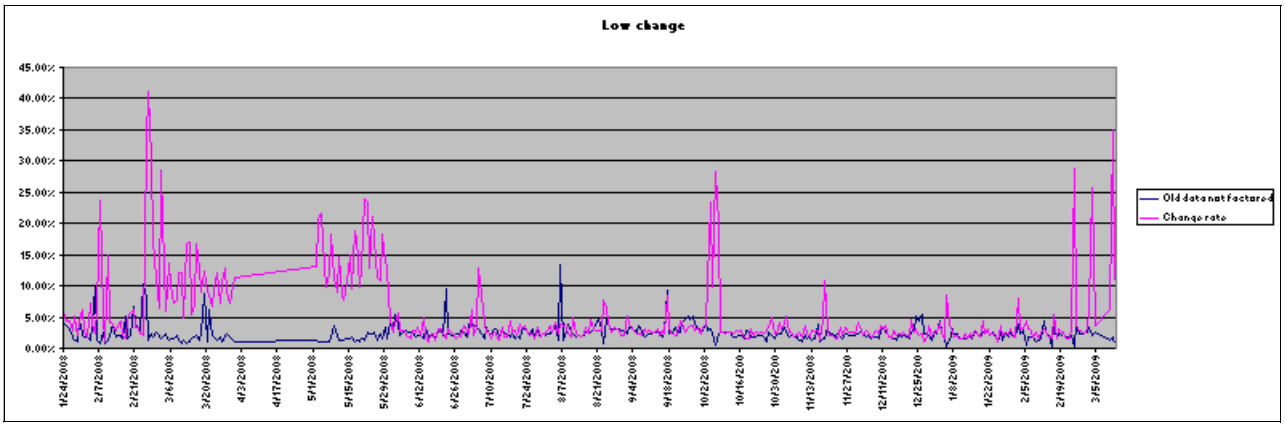

Figure A-3 Case 1 - low change rate and low old data not factored

Case 2: Moderate change rate and high old data not factored

Figure B-4 shows a case where a parser is effective. The backup application is Tivoli Storage Manager. The change rate is moderate, but the old data not factored is high.

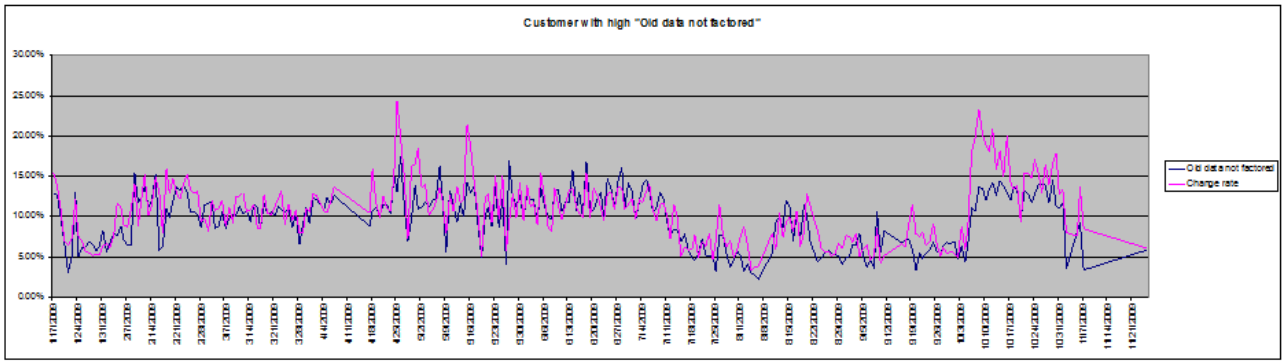

Figure A-4 Case 2 - moderate change rate and high old data not factored

Case 3: High change rate, moderate to low old data not factored

In Figure A-5, the change rate is high and the old data not factored is moderate to low. A parser might help, but the benefit that it offers is marginalized by the high change rate. In this case, the best action might be to look for causes of the high change rate in the environment and try to reduce that value.

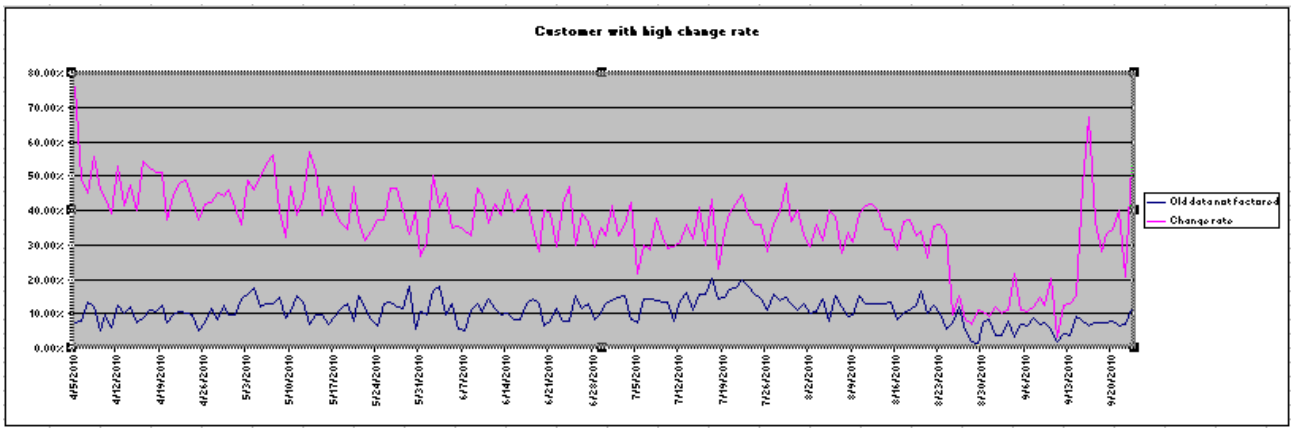

Figure A-5 Case 3 - high change rate, moderate to low old data not factored

Case 4: High change rate, low old data not factored

Taking case 3 to an extreme, Figure A-6 shows a ProtecTIER installation that reports a high change rate and a low old data not factored. The Tivoli Storage Manager parser does not help here at all, and the best action is to look for causes of the high change rate in the environment.

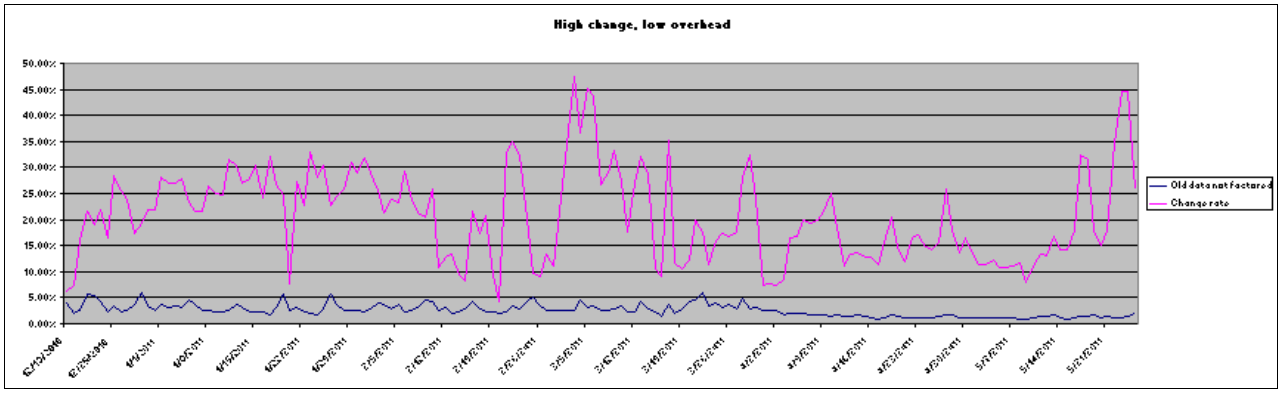

Figure A-6 Case 4 - high change rate, low old data not factored

Experience from one user site

Figure A-7 tracks the change in old data not factored of a ProtecTIER installation. It shows about four weeks of history while the site was still running ProtecTIER V2.5.7 and about six weeks after upgrading to Version 3.1.9. Over a period of about a week after the upgrade, the old data not factored dropped fairly rapidly from around 6% to around 3%.

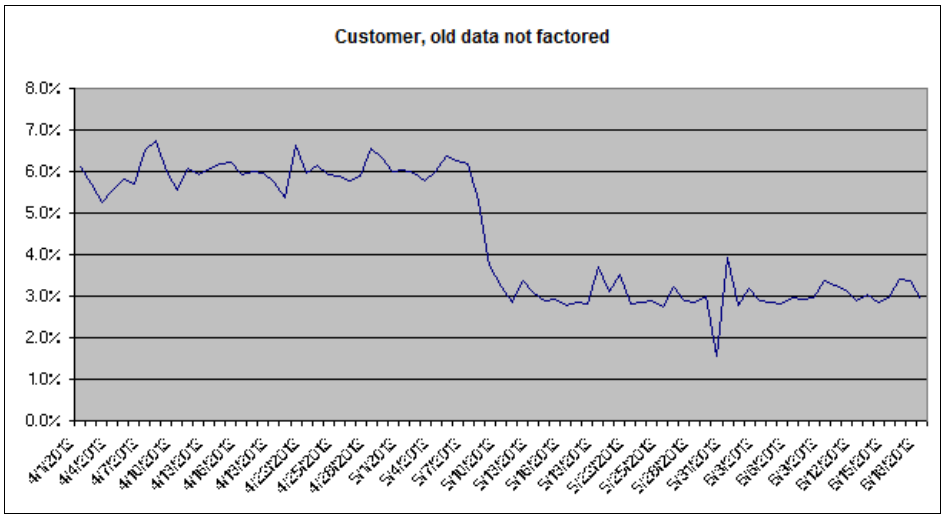

Figure A-7 Change in old data not factored at a user site after upgrading from ProtecTIER V2.5.7 to Version 3.1.9

Not every site upgrading to a ProtecTIER release with a parser sees such an immediate improvement in the old data not factored. The actual improvement depends on many factors outside of the ProtecTIER product.

Using analyze_sessions to monitor the benefit of a ProtecTIER parser

This section explains how to monitor changes in the deduplication ratio at a fine-grain level, whether it is at a per cartridge, per session, daily, or even hourly level. You can use the information in this section to contrast deduplication before you install a ProtecTIER version that supports a parser with the deduplication after the installation.

The **`analyze_sessions`** utility assists you with understanding the deduplication of your workloads. It examines the deduplication logs of the ProtecTIER cluster and produces a report with the deduplication split (by default) into sessions. The definition of a session is a period of system ingest activity. A session ends when there is a short break in the I/O, during which the ProtecTIER product is idle. There can therefore be one session that covers a multiple hour period, or many sessions within an hour.

The **`analyze_sessions`** utility can be found in the `/opt/dtc/app/utils` directory on a ProtecTIER server. This utility examines the deduplication logs in their default location in the repository of the ProtecTIER server and generates a report on the deduplication that is achieved by the ProtecTIER cluster. The report is in CSV format and is placed in the `/pt_work` directory. The report is best viewed by using a spreadsheet program, which can also be used to plot the deduplication ratios of the sessions and their trends.

There are several parameters that can be used with **`analyze_sessions`**:

`-n <number>` Report only <number> of months back.

`-s <start date>` Start reporting only from date (date format is YYYY-MM-DD-HH-MM).

`-e <end date>` Report only up until the end date.

-sd <directory>	The directory that contains the deduplication logs (if it is not the default location).
-min	Report only overall statistics.
-daily	Include a daily summary.
-hourly	Include an hourly summary.
-d <minutes>	The number of minutes of idle time that defines a session boundary.
-i <file of carts>	The file that contains a list of cartridge barcodes. Only these barcodes are reported.
-c	Used with **-i**. Create a separate file for each cartridge.
-l	Used with **-i**. Add a line to the report for each cartridge.
-o <output>	Specify an output file name other than the default.

Figure A-8 shows the default output of the **analyze_sessions** utility. The important fields are the system change rate (the effect of only the deduplication, not the compression), and the compressedBytesCount (the actual amount of data that is written to the repository by this session). The deduplication of a particular session (regarding all the data already in the repository) is column C divided by column F, and includes the effects of deduplication and compression of the data.

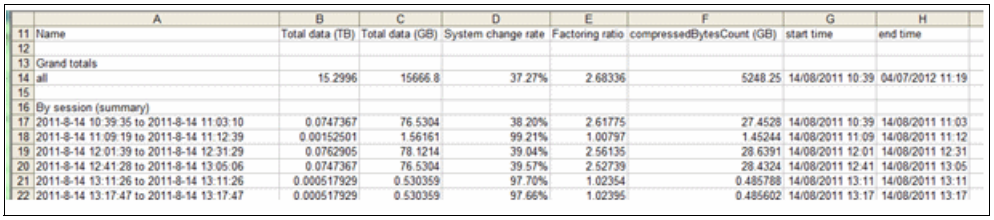

	A	B	C	D	E	F	G	H
11	Name	Total data (TB)	Total data (GB)	System change rate	Factoring ratio	compressedBytesCount (GB)	start time	end time
12								
13	Grand totals							
14	all	15.2996	15666.8	37.27%	2.68336	5248.25	14/08/2011 10:39	04/07/2012 11:19
15								
16	By session (summary)							
17	2011-8-14 10:39:35 to 2011-8-14 11:03:10	0.0747367	76.5304	38.20%	2.61775	27.4528	14/08/2011 10:39	14/08/2011 11:03
18	2011-8-14 11:09:19 to 2011-8-14 11:12:39	0.00152501	1.56161	99.21%	1.00797	1.45244	14/08/2011 11:09	14/08/2011 11:12
19	2011-8-14 12:01:39 to 2011-8-14 12:31:29	0.0762905	78.1214	39.04%	2.56135	28.6391	14/08/2011 12:01	14/08/2011 12:31
20	2011-8-14 12:41:28 to 2011-8-14 13:05:06	0.0747367	76.5304	39.57%	2.52739	28.4324	14/08/2011 12:41	14/08/2011 13:05
21	2011-8-14 13:11:26 to 2011-8-14 13:11:26	0.000517929	0.530359	97.70%	1.02354	0.485788	14/08/2011 13:11	14/08/2011 13:11
22	2011-8-14 13:17:47 to 2011-8-14 13:17:47	0.000517929	0.530359	97.66%	1.02395	0.485602	14/08/2011 13:17	14/08/2011 13:17

Figure A-8 Output from the analyze_sessions command

Planning for the Tivoli Storage Manager parser

When a new parser is added to a ProtecTIER system, there might be a slight performance degradation (less than 3%).

Initially, there might be a temporary and slight decrease in deduplication ratio and an increase in the space that is used in the repository because the newly parsed user backups might not match as well with the existing unparsed backups in the repository. After a few generations, a new steady state is reached, where parsed user backups deduplicate by using other parsed user backups as their reference for deduplication. Because there are no interfering headers, old data not factored is expected to decline, causing the deduplication ratio to reach an improved level. Consequently, used space in the repository drops. There might also be improved performance because the ProtecTIER product processes data more efficiently when the deduplication level is higher.

For environments where the backup window is tight or the repository is full, customers should work with ProtecTIER Support before they load the first ProtecTIER release that contains the applicable parser. ProtecTIER Support has tools to analyze problem reports that give a historic breakdown of actual change rate and old data not factored. These figures can help you understand both the startup impact and how effective the parser is in improving deduplication ratio over time.

B

Entry-level and midrange disks

This appendix provides the best practices and guidelines for the IBM System Storage DS3000, IBM System Storage DS4000, and IBM System Storage DS5000 family of disk storage subsystems in your ProtecTIER environment. It also includes the procedure to check and adjust the Automated Volume Transfer (ATV) settings, zoning, cabling guidelines, and RAID configuration for these entry-level and midrange disk subsystems.

This appendix describes the following topics:

- General considerations and best practices for attaching entry-level or midrange storage subsystems to the ProtecTIER server
- Considerations that are specific to the attachment and configuration of the DS3000 family

Important: As of ProtecTIER V3.3.0, the entry-level and midrange family of disk storage subsystems for DS3000, DS4000, and DS5000 are not supported.

DS3000, DS4000, and DS5000 storage systems that are attached to ProtecTIER systems that run earlier releases are still supported. We provide the best practices and guidelines for those systems

The list of supported entry-level and midrange disk storage subsystems can be found in the TS7650/TS7650G ISV and Interoperability Matrix, found at:

`http://www.ibm.com/systems/storage/tape/resources.html`

General considerations

This section briefly summarizes the common best practices that are used when you attach entry-level or midrange disk storage subsystems to the ProtecTIER server. It then describes the detailed steps and considerations in each section that is dedicated to a specific device.

Using the latest version of software

When you install a new or redeploy an existing storage subsystem, ensure that all the firmware levels are up-to-date and at the same version on all installed controllers. For the latest version of firmware and to verify all the upgrade steps, consult your IBM technical representative or visit the IBM Fix Central at:

http://www.ibm.com/support/fixcentral/options

Always use the latest supported version of the IBM DS Storage Manager that is valid for your operating system. Even though the important enhancements were made in Version 10.60, you should use the latest version (at the time of writing, Version 10.77) that is available for the Windows platform. You can find the latest version at:

http://www.ibm.com/support/fixcentral/swg/selectFixes?parent=Mid-range+disk+systems&product=ibm/Storage_Disk/DS3950&release=DS_SM_v10.77.x5.28&platform=Windows+64-bit,+x86&function=all

Hint: Upgrade the firmware of your DS storage subsystem and DS Storage Manager to the most recent version that supports your configuration.

Also, the IBM DS Storage Manager is the common management and operational software that supports all types of entry-level and midrange disk storage subsystems, including the IBM System Storage DS3500 family (IBM Systems Storage DS3950, DS4000, and DS5000).

General settings on DS storage subsystems

When you use entry-level or midrange disk storage subsystems as a back-end storage repository, ProtecTIER benefits from the parameters that are listed in Table B-1. These parameters might increase the performance and stability of the system, but they do not change the factoring ratio.

Table B-1 DS storage subsystem settings

Option on DSx000	Value for ProtecTIER
Media scan frequency (in days)	30
Default logical drive	Unnamed
Cache block size (in KB)	32
Start cache flushing at (in percentage)	50%
Stop cache flushing at (in percentage)	50%
Segment size	128 KB
Flush write cache after (in seconds)	10.00
Write cache without batteries	Disabled

Option on DSx000	Value for ProtecTIER
Write cache with mirroring	Enabled
Read cache	Enabled
Write cache	Enabled
Enable background media scan	Enabled
Modification priority	High
Pre-read redundancy check	Disabled
Media scan with redundancy	Enabled

Disabling Automatic Volume Transfer

When you attach one of the DS storage subsystems, you must disable Automatic Volume Transfer (AVT) in your DS Storage Manager software. The AVT is a built-in feature of controller firmware that allows logical drive-level failover rather than controller-level failover. It provides redundant I/O paths with a multipath driver that is installed on the host system.

The AVT feature is automatically disabled or enabled depending on the type of host ports on the host partition to which you mapped the logical drives. It is disabled by default for Microsoft Windows, IBM AIX, and Sun Solaris operating systems. It is enabled by default for Linux (and therefore for the ProtecTIER server), Novell NetWare, and HP-UX operating systems.

The Automatic Volume Transfer (AVT) is also known as Auto Disk Transfer (ADT).

Important: Disable AVT on your DS storage subsystem. The procedure requires a restart of all controllers in your storage subsystem.

This step is not needed with DS Storage Manager V10.77 and higher when operating the DS3500 family, where AVT/ADT is disabled for Linux host type by default.

Procedure

There are two options that are available to disable AVT:

- Using the Storage Manager command **`SMcli.exe`**
- Directly in the DS Storage Manager GUI

The initial steps in the following section are used for both options. You must ensure that no I/O operations are in process when you use this procedure, so you should shut down all ProtecTIER controllers that are connected to the relevant DS storage controller. Stop the services and power off all ProtecTIER nodes that are connected to the DS storage by completing the following steps:

1. Shut down the IBM TS3000 System Console (TSSC):
   a. Log in to the TSSC with the user name “service” and password “service”.
   b. Right-click the TSSC blue desktop. The TSSC menu opens.
   c. Select **Shutdown Console**. A dialog box opens with this question:

      “Shutdown the machine?”

   d. Click **OK** to start the shutdown process. The shutdown process is complete when you see this message:

"It is now safe to turn off the TS3000 System Console. Hold down the power button for 5 seconds to power off".

e. Press the power button (it is not necessary to hold the button down). When the power off process is complete, the power LED on the front panel of the ProtecTIER server flashes steadily, which indicates that the server is in standby mode.

2. If you have not already done so, log in to the ProtecTIER server by attaching a keyboard and monitor and logging in with user name "ptconfig" and password "ptconfig".
3. From the ProtecTIER Service menu, click **Manage ProtecTIER services (...)**. The system displays the Manage ProtecTIER services menu.
4. From the Manage ProtecTIER Services menu, click **Stop all services**. When the services are stopped, the system shows the following messages:

```
Stopping ptrasd                          [ Done ]
Stopping vtfd                            [ Done ]
Stopping ptcluster                       [ Done ]
```

5. Press Enter to return to the ProtecTIER Service Menu, then type `E` to exit.
6. Log in to the server with the user name "ptadmin" and password "ptadmin".
7. At the command-line prompt, type **`poweroff`** and press Enter. When the power off process is complete, the Power LED on the front of the server panel flashes steadily, which indicates that the server is in standby mode.

Because the TSSC console and ProtecTIER servers are down, there is no I/O activity that is seen on the attached storage controllers. You can continue to safely disable AVT. Example B-1 shows the host types that are defined in the DS Storage Manager and the default settings of AVT.

Example B-1 Default settings of AVT for each host type

HOST TYPE	ADT STATUS	INDEX
AIX	Disabled	6
AIX-ADT/AVT	Enabled	4
DEFAULT	Disabled	0
HP-UX	Enabled	7
HPXTPGS	Disabled	17
IBM TS SAN VCE	Enabled	12
Irix	Disabled	10
LNXCLVMWARE	**Disabled**	**13**
Linux	**Enabled**	**5**
Linux_DMP	Disabled	18
NetWare Failover	Enabled	11
Solaris (with Veritas DMP)	Enabled	14
Solaris (with or without MPXIO)	Disabled	8
Unused1	Enabled	1
Windows 2000/Server 2003/Server 2008 Clustered	Disabled	3
Windows 2000/Server 2003/Server 2008 Clustered (supports DMP)	Enabled	15
Windows 2000/Server 2003/Server 2008 Non-Clustered	Disabled	2
Windows 2000/Server 2003/Server 2008 Non-Clustered (supports DMP)	Enabled	9

Now we describe how to complete the rest of the procedure by using the two options.

If you want to use the Storage Manager command `SMcli.exe`, log in to each of the controllers by using **`SMcli.exe`** and complete the following steps:

1. Use the following command to disable AVT on Linux host types:

```
SMcli.exe ip-addr-for-A set controller [a] HostNVSRAMBYTE [5,0x24]=0x00
```

Variable 5 stands for Linux hosts and 24 refers to the AVT value.

2. Reboot the controller by running the following command:

   ```
   SMcli.exe ip-addr-for-A reset controller [a]
   ```

3. Repeat these steps for controller B:

   - **SMcli.exe ip-addr-for-B set controller [b] HostNVSRAMBYTE [5,0x24]=0x00**
   - **SMcli.exe ip-addr-for-B reset controller [b]**

4. After you reboot both controllers, ensure that the settings are in effect by running the following commands:

   - **SMcli.exe ip-addr-for-A show controller [a] HostNVSRAMBYTE [5,0x24]**
   - **SMcli.exe ip-addr-for-B show controller [b] HostNVSRAMBYTE [5,0x24]**

AVT settings: Although changing the AVT settings for a Linux host type works and does the same job, it is a preferred practice to change the host type of your ProtecTIER to LXNCLVMWARE, which disables the AVT by default. This method prevents the uncontrolled reset of AVT for a Linux host to “enabled” when you upgrade your DS Storage subsystem firmware.

If you want to change the host type from your DS storage subsystem by using DS Storage Manager, complete the following steps. The preferred host type for your ProtecTIER servers is “Linux Cluster/VMware”, which is represented by the expression LNXCLVMWARE. In this case, instead of disabling AVT for Linux host type, change the host type of ProtecTIER to the one that has AVT disabled by default.

1. In your DS Storage Manager, navigate to your profile to determine whether you must change your current settings. These settings are at the bottom, near the end of the Mappings profile (Figure B-1).

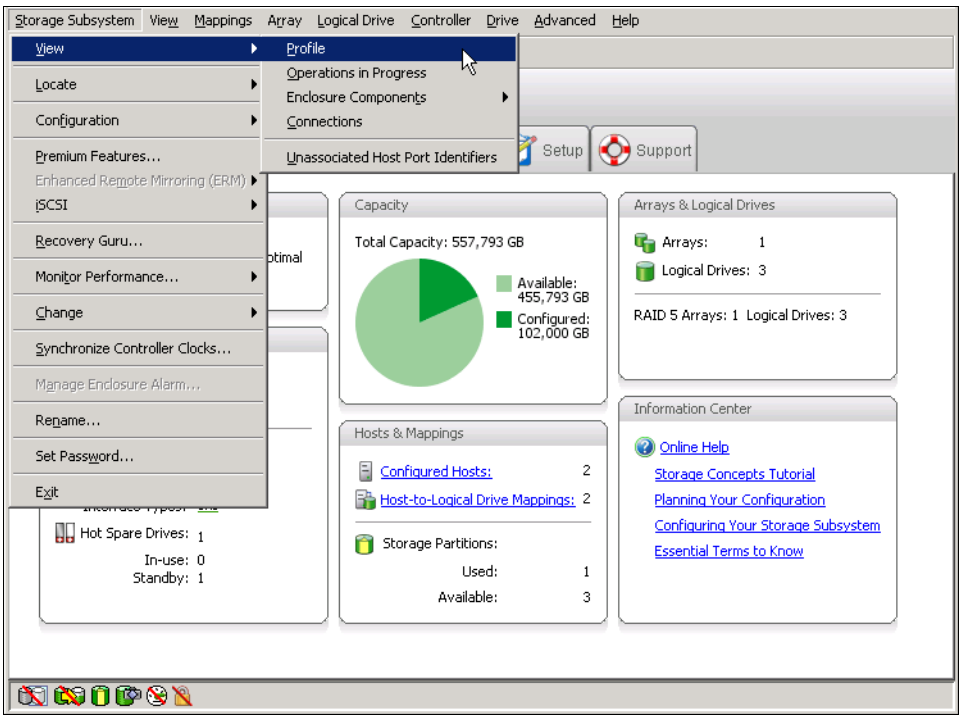

Figure B-1 Profile view in DS Storage Manager

2. Determine the host type. If your host type is set at the subsystem level, modify the settings for the default host type (Figure B-2). If your host type is set at the host group level, modify the settings for the default group.

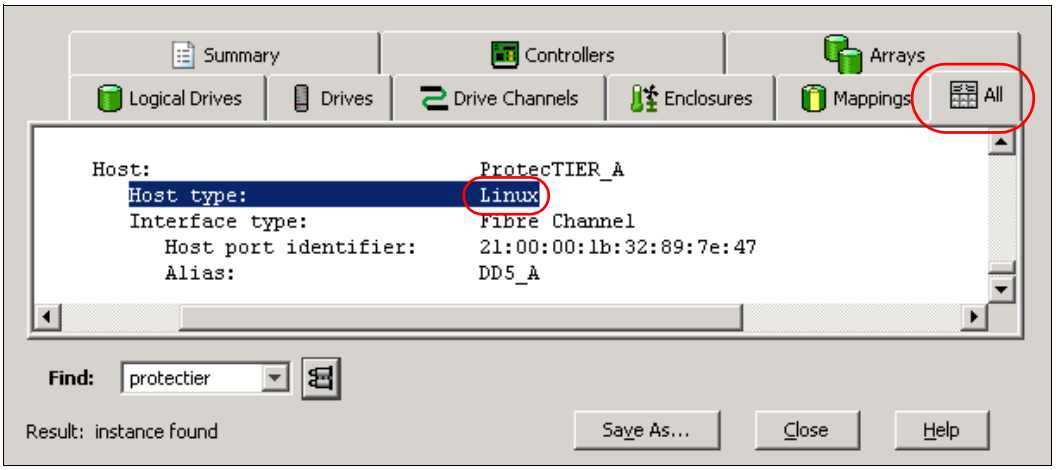

Figure B-2 Determine the host type of your ProtecTIER

3. Scroll down to the Host Type Definitions section of the profile. The ADT STATUS column shows if AVT is enabled or disabled (Figure B-3).

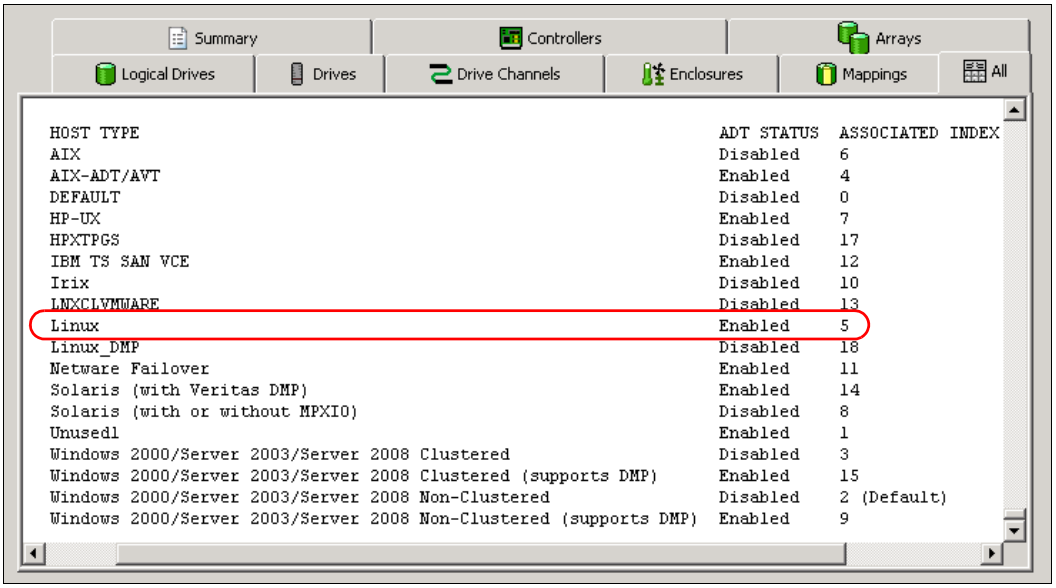

Figure B-3 General definition of an AVT that is different for different host types

4. Close the profile window and return to the Storage Manager window. The host type of your ProtecTIER node must be changed to a different host type that has the AVT disabled by default. The best matching profile is the LNXCLVMWARE.

– If the host type is set by using the default host type, click **Storage Subsystem** → **Change** → **Default Host Type**, and select **LNXCLVMWARE** in the Change Default Host Type window (Figure B-4). The LNXCLVMWARE host type now has AVT disabled.

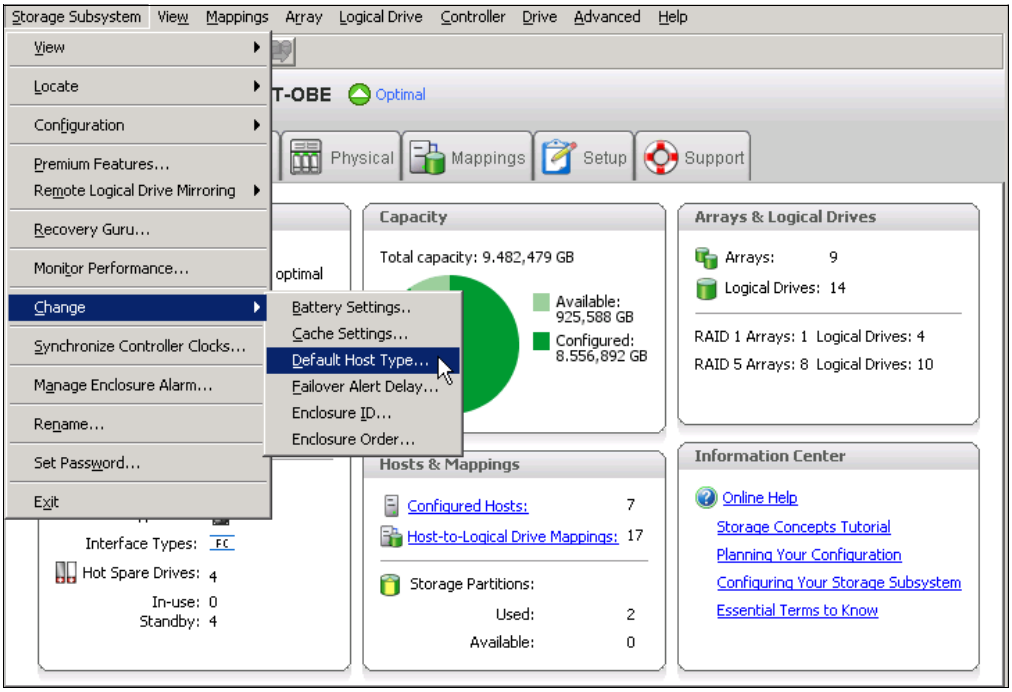

Figure B-4 Changing the default host type

– If the host type is set at the host group or host port level, click the **Mappings** tab in the Storage Manager window. Locate and highlight your host name, and click **Change Host Type** (Figure B-5).

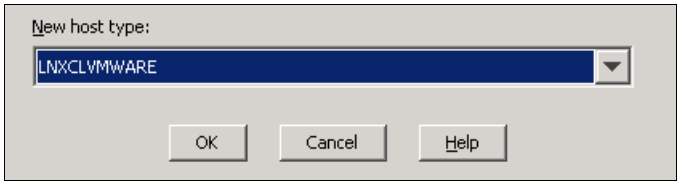

Figure B-5 Selecting the LNXCLVMWARE host type in DS Storage Manager

– In newer versions of Storage Manager, click the **Mappings** tab in the Storage Manager window, locate and highlight your host name, and click **Change** → **Host Operating System**.

Important: As of storage controller firmware Version 7.60.40, the host type LNXCLVMWARE is no longer available (replaced by LNXCLUSTER and VMWARE). When you upgrade the controller firmware, your host type is automatically changed to LNXCLUSTER. Because this host type still matches your requirements, no further action is needed, as shown in the following output:

```
HOST TYPE                                                    ADT STATUS  INDEX
LNXAVT                                                       Enabled     5
LNXCLUSTER                                                   Disabled    13
Linux                                                        Disabled    18
VMWARE                                                       Disabled    16
```

Direct cabling without SAN

This section provides guidelines about how to directly connect your ProtecTIER stand-alone server or dual-node cluster to the back-end disk storage subsystems without an underlying SAN network. This scenario is recommended when your ProtecTIER environment consists of a single back-end DS storage repository and you do not plan to extend it by attaching more storage devices. It is a best practice because it eliminates the factor of SAN management (upgrading SAN switch code, SAN failures, attention on correct zoning, and so on). The concept is the same for all entry-level and midrange DS storage subsystems from IBM.

Tip: Keep your infrastructure as simple as possible; ensure that there is always a single active path and one redundant path only. More redundant paths increase the complexity of the solution and are not needed.

Stand-alone ProtecTIER with dual controller DS

In this configuration, you need only two shortwave, single mode fiber cables, one for each storage controller. The schema of the cabling is shown in Figure B-6.

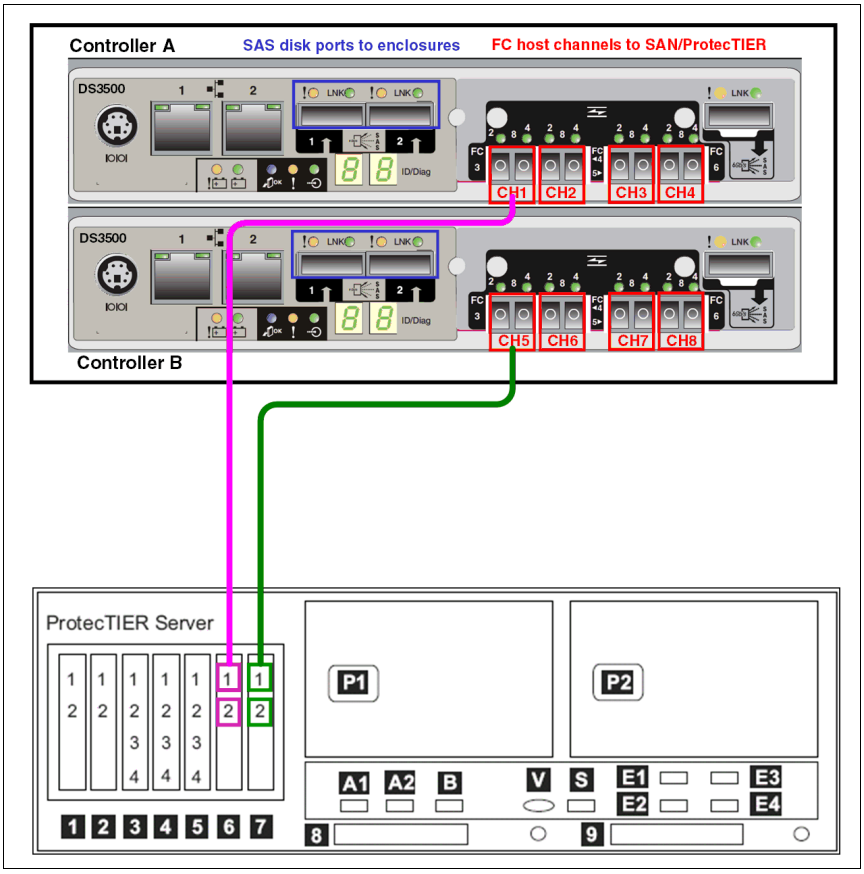

Figure B-6 Direct cabling of single ProtecTIER to DS storage controllers

Only the host channels CH1 (Controller A) and CH5 (Controller B) are used for direct connection of ProtecTIER QLogic HBAs to each of the storage subsystem controllers. This installation provides exactly one active and one standby path to your repository, depending on which storage controller holds the control over associated LUNs.

Tip: If there is future deployment of more storage subsystems, use the available QLogic ports on your ProtecTIER server in slot 6, port 2 and slot 7, port 2.

A solution with three or more storage subsystems requires a SAN network.

ProtecTIER cluster with dual controller DS

ProtecTIER clusters can also be connected directly to your back-end storage system without using an existing SAN network. The conceptual diagram is shown in Figure B-7.

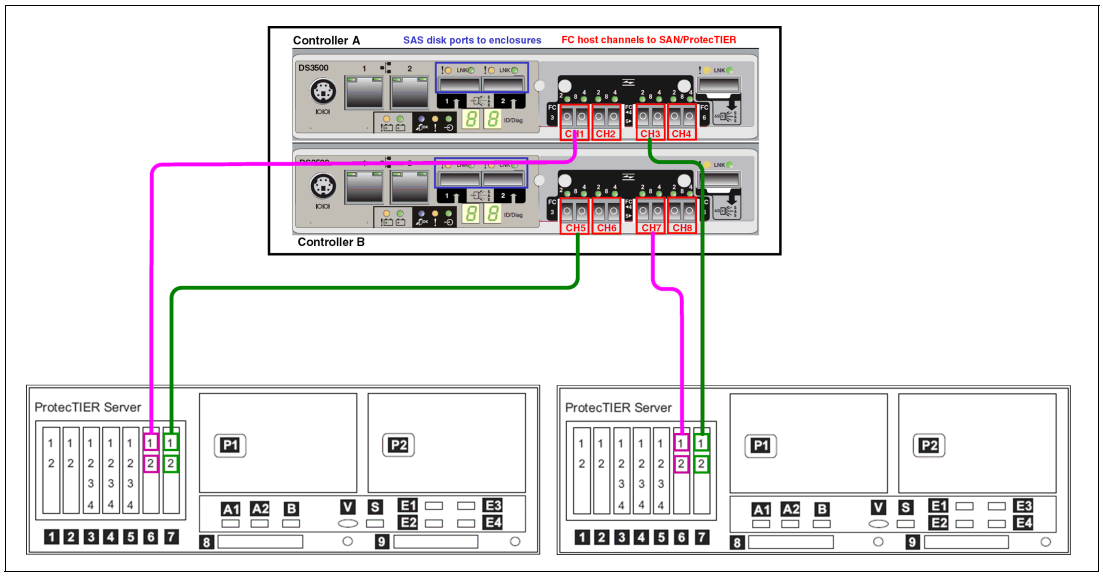

Figure B-7 ProtecTIER cluster that is directly cabled to the DS storage subsystem

The diagram interprets the same rule as in the case of a stand-alone ProtecTIER server. Keep the infrastructure simple by enabling just one active and one redundant path to each component in the network.

Cabling and zoning in SAN

The following section describes examples of best practices when you do SAN zoning of ProtecTIER servers and DS storage controllers as stand-alone servers or in dual-node clusters.

Assume, as a general best practice for system high availability, that all back-end entry-level or midrange storage subsystems in the ProtecTIER environments are implemented with dual controllers, each equipped with four Fibre Channel ports per controller. This configuration is available and supported on all models of the DS3000, DS4000, and DS5000 family of storage subsystems. However, as the Fibre Channel (FC) HBAs can be an optional feature for certain models, it must be ordered explicitly. For example, the DS3500 systems (models DS3512 and DS3524) are in the default configuration (delivered with two 6 Gbps SAS ports). The 4-port FC host adapter is available on demand as a replacement.

With this assumption, the SAN zoning of your back-end storage and ProtecTIER nodes is similar for all types of entry-level and midrange disk storage subsystems.

Reminder: Always connect a dual-controller disk storage subsystem as back-end storage. Avoid using storage subsystems that are equipped with only a single controller.

When you cable and zone your ProtecTIER servers with disk storage subsystems at the back end, always simulate the deployment scenario with direct cabling. This setup is valid either in a stand-alone or cluster configuration of ProtecTIER. Explicitly, it means to follow these rules:

- Use worldwide port name (WWPN) zoning, also called soft zoning. Avoid using hard zoning type, where the ProtecTIER HBA or storage subsystem HBA is bound to a specific port on the SAN switch. This setup helps resolve a malfunctioning port quickly by just replugging the FC cable.
- Ensure that there is always one active path and one redundant path only to each component in the SAN network, especially to your ProtecTIER servers. More redundant paths increase the complexity of the solution and might have significant impact on performance during path failover.
- Define zones with one initiator and one target only. Even if a single port of the ProtecTIER QLogic HBA needs access to the multiple channels on different storage subsystems, do not put them into one zone. Use two dedicated zones.
- Isolate the ProtecTIER traffic from other systems or applications. For example, on IBM b-type SAN switches, dedicate all that ports that are managed by a single Application-Specific Integrated Circuit (ASIC) to the ProtecTIER servers, back-end storage subsystems, and, optionally, backup server hosts. On Cisco switches, this dedication does not bring any further benefit because of the different switching architecture (crossbar fabric matrix).
- In a scenario with ProtecTIER clusters, it is important to use the dual-fabric SAN network. It eliminates the risk of a single switch failure, which leads to the suppression of high availability.

To understand the naming convention of ports that is used for zoning, see Figure B-8, which shows a DS3500 with dual controllers and how the ports are represented in your zoning.

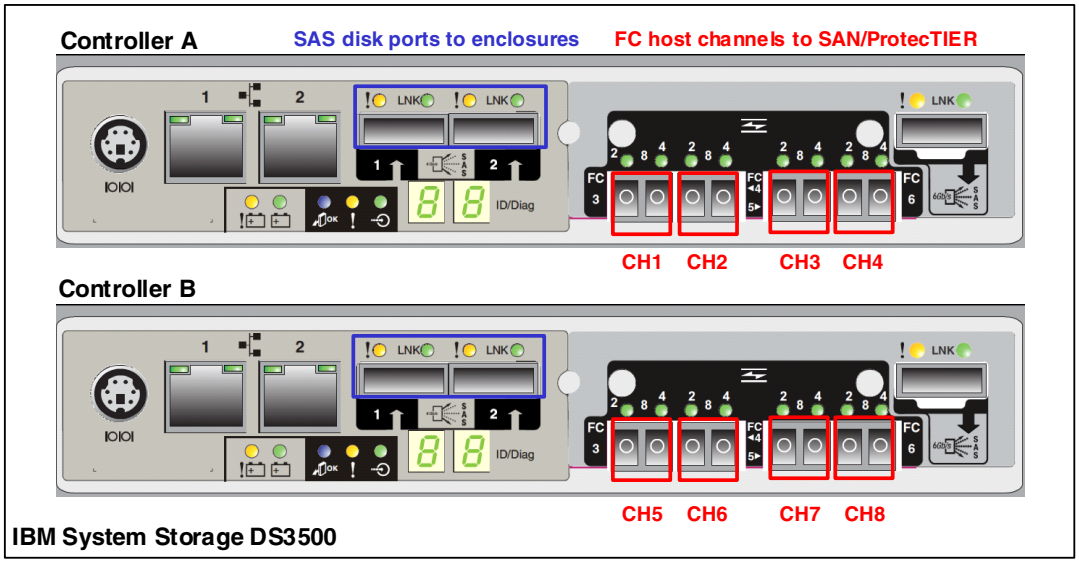

Figure B-8 Naming convention that is used for zoning of HBA channels on DS3500

This section introduces the four most frequently used deployments:

- Stand-alone ProtecTIER server that is connected to a single SAN switch
- Stand-alone ProtecTIER server that is connected to a dual-fabric SAN network
- ProtecTIER cluster that is connected to a dual-fabric SAN network
- ProtecTIER cluster with two back-end storage servers in a dual-fabric SAN

There are more available scenarios, such a connecting a ProtecTIER cluster to a single SAN switch, but those scenarios are not considered best practices. The factor of single points of failure plays a key role. We do not cover the scenario of three or more back-end storage devices, as we are only concerned with attaching disk subsystems to the same set of SAN switches and working with zoning.

Stand-alone ProtecTIER and a single SAN switch

In this scenario, the ProtecTIER stand-alone server is connected to the back-end disk storage subsystem through a single SAN fabric switch. Although the best practice for such an implementation is to interconnect these two components directly, some data centers cannot use this option. Sharing the storage subsystem is not considered a best practice because of distance limitations. This concept is shown in Figure B-9.

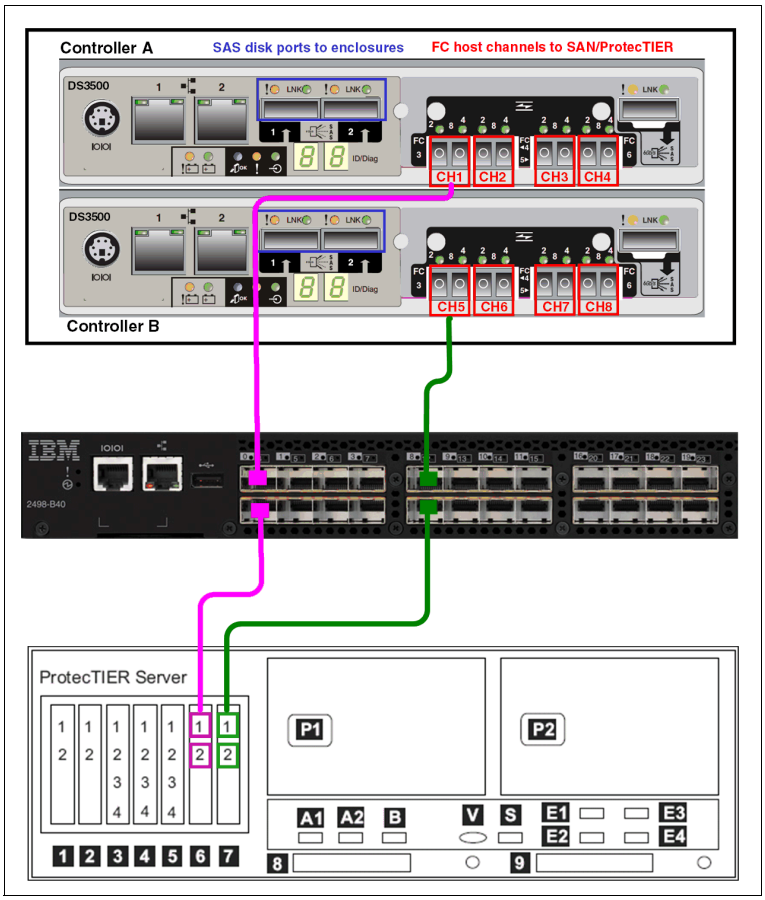

Figure B-9 Single SAN network with stand-alone ProtecTIER

Two SAN zones must be defined on your switches, as shown in Table B-2.

Table B-2 SAN zoning definition

Effective configuration	
zone_1	WWPN of ProtecTIER Node 1 HBA 6 Port 1 WWPN (S6P1) WWPN of DSx000 #1 Controller “A” Channel 1 (A CH1)
zone_2	WWPN of ProtecTIER Node 1 HBA 7 Port 1 WWPN (S7P1) WWPN of DSx000 #1 Controller “B” Channel 5 (B CH5)

You should replace this configuration by implementing direct cabling between ProtecTIER and the back-end storage subsystem when possible.

Stand-alone ProtecTIER in a dual-fabric SAN

This situation is similar to the configuration that is described in “Stand-alone ProtecTIER and a single SAN switch” on page 468, but instead of a single SAN switch, use a dual-fabric SAN with two independent SAN switches. The cabling concept is shown in Figure B-10.

In the figure, the magenta color represents the primary path, and the dark green color represents the redundant, inactive paths. Orange connections are not used (not zoned) to the ProtecTIER server, but we recommend connecting each controller of the storage subsystems to the dedicated SAN switch.

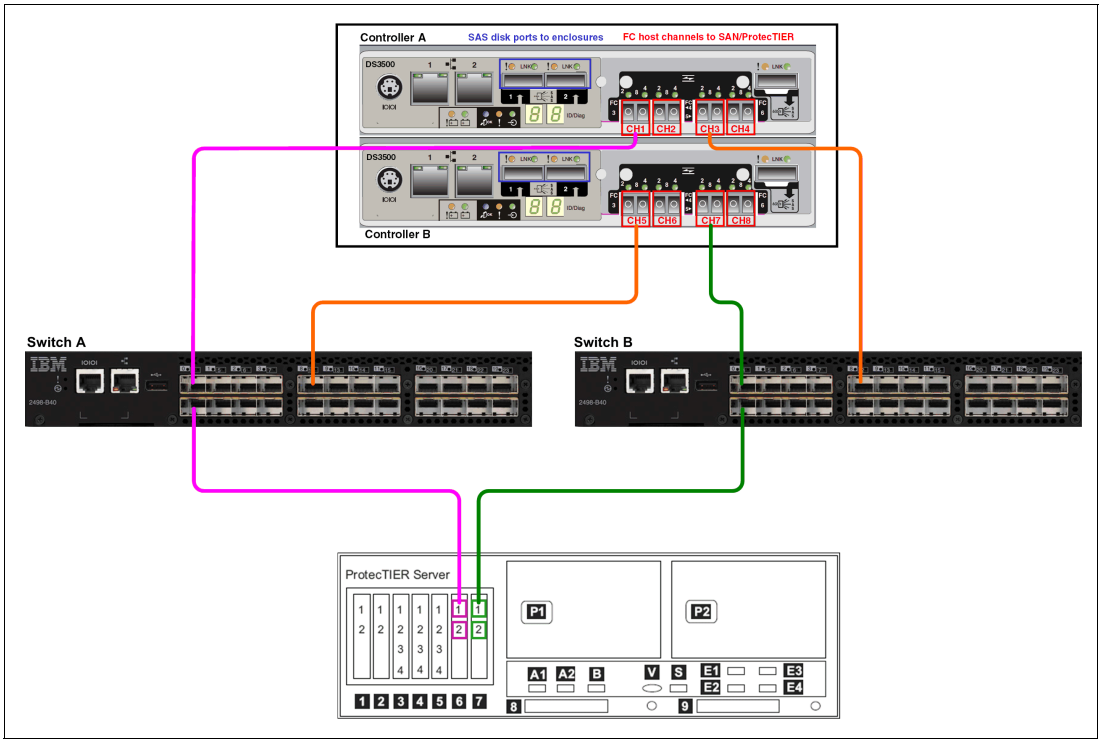

Figure B-10 Stand-alone ProtecTIER server in dual-fabric SAN

The zoning for each SAN switch is shown in Table B-3.

Table B-3 SAN zones that are defined on each switch

Effective configuration	
switch_A_zone_1	WWPN of ProtecTIER Node 1 HBA 6 Port 1 WWPN (S6P1) WWPN of DSx000 #1 Controller “A” Channel 1 (A CH1)
switch_B_zone_1	WWPN of ProtecTIER Node 1 HBA 7 Port 1 WWPN (S7P1) WWPN of DSx000 #1 Controller “B” Channel 7 (B CH7)

You should not define the additional redundant paths by adding the remaining two combinations of SAN zones by using orange connections. The setup that is shown in Table B-4 increases complexity, and does not provide any additional benefits:

Table B-4 SAN zoning - do not define additional redundant paths

Ineffective configuration	
Switch A	WWPN of ProtecTIER Node 1 HBA 6 Port 1 WWPN (S6P1) WWPN of DSx000 #1 Controller “B” Channel 5 (B CH5)
Switch B	WWPN of ProtecTIER Node 1 HBA 7 Port 1 WWPN (S7P1) WWPN of DSx000 #1 Controller “A” Channel 3 (A CH3)

ProtecTIER cluster in a dual-fabric SAN

This deployment is the most common one of IBM deduplication solutions. The dual-node ProtecTIER cluster increases the backup and restore performance. The connection through two independent fabric SAN switches provides the necessary level of high availability and storage path redundancy.

Four zones need to be defined, with two in each fabric, as shown in Table B-5.

Table B-5 Dual-node ProtecTIER cluster with single back-end storage

Effective configuration	
switch_A_zone_1	WWPN of ProtecTIER Node 1 HBA 6 Port 1 WWPN (S6P1) WWPN of DSx000 #1 Controller “A” Channel 1 (A CH1)
switch_A_zone_2	WWPN of ProtecTIER Node 2 HBA 6 Port 1 WWPN (S6P1) WWPN of DSx000 #1 Controller “A” Channel 1 (A CH1)
switch_B_zone_1	WWPN of ProtecTIER Node 1 HBA 7 Port 1 WWPN (S7P1) WWPN of DSx000 #1 Controller “B” Channel 7 (B CH7)
switch_B_zone_2	WWPN of ProtecTIER Node 2 HBA 7 Port 1 WWPN (S7P1) WWPN of DSx000 #1 Controller “B” Channel 7 (B CH7)

Figure B-11 shows the Fibre Channel cabling.

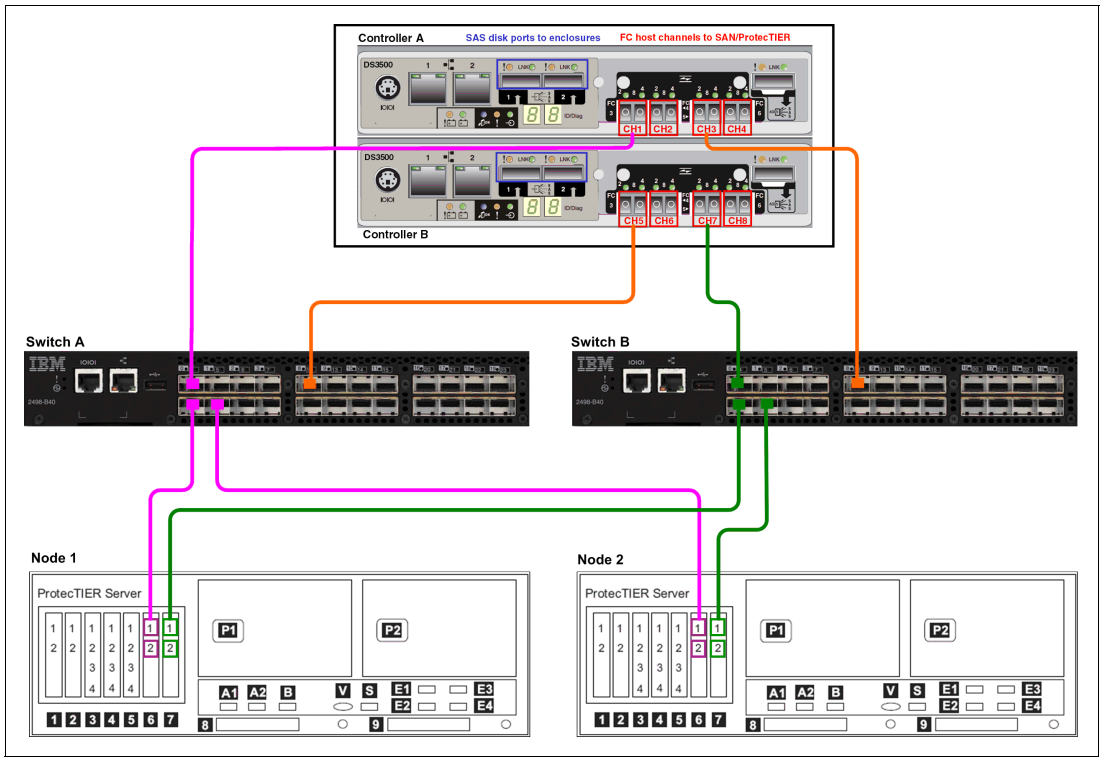

Figure B-11 Cabling of ProtecTIER cluster in a dual fabric

Also in this configuration, do not define SAN zones for the redundant storage paths that are depicted by the orange connectors in the figure. These zones are not required for a ProtecTIER server and their definition increases the complexity.

In one scenario, the definition of SAN zones is not needed, and two zones on each switch can be merged into a single zone, represented by three ports that are marked with magenta color on switch A and three ports that are highlighted by dark green on switch B. Although this setup can be done, one of the ProtecTIER rules is that there are SAN zones with a single target and a single initiator only. Therefore, two zones are needed in each fabric.

ProtecTIER cluster with more back-end storage subsystems

This scenario provides an example of SAN zones where your dual-node ProtecTIER cluster uses two dedicated DS storage subsystems, which are mapped by using a dual SAN fabric. For production environments, you should use a dual-fabric SAN network with independent SAN switches to eliminate a single point of failure.

The cabling diagram is shown in Figure B-12. The magenta connectors are primary paths from each ProtecTIER node to “Controller A” of every storage subsystem in the repository chain, and inactive paths are marked dark green. The orange connectors represent the connectivity of each storage controller to all fabrics in the SAN. For the ProtecTIER solutions, these paths are not enabled by SAN zoning to ProtecTIER nodes.

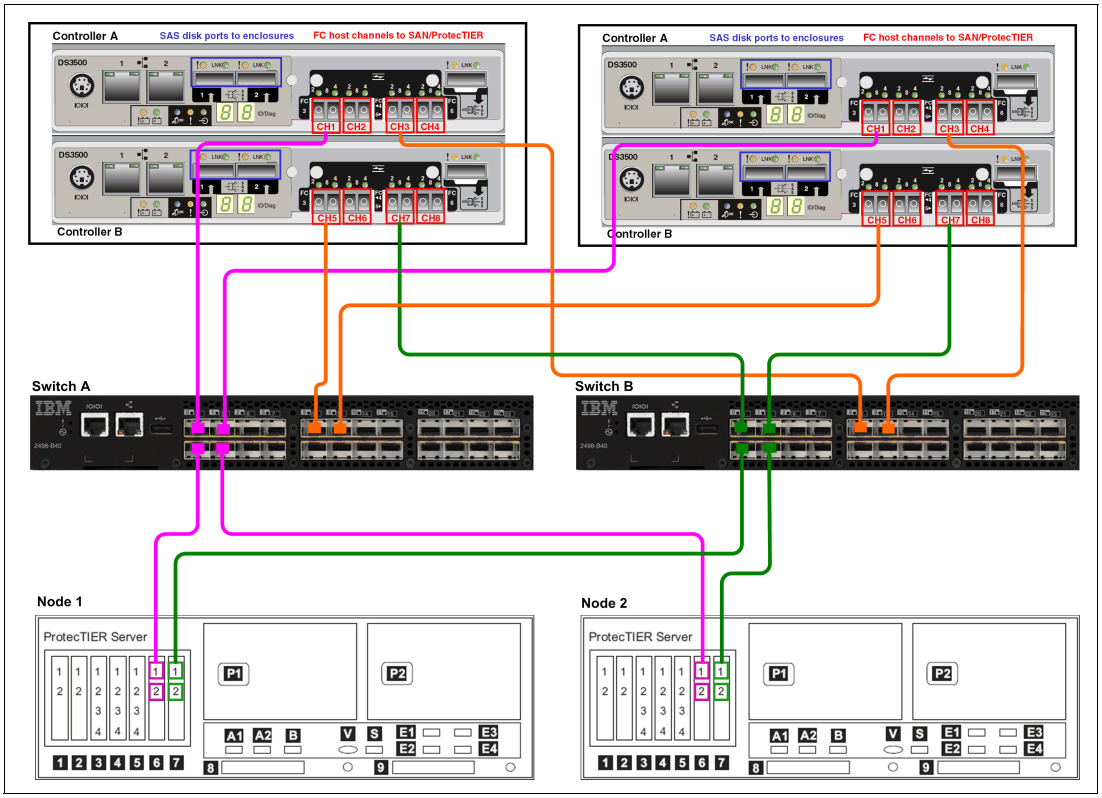

Figure B-12 ProtecTIER cluster with two disk storage subsystems in a repository

Also in this configuration, do not cable ports 2 of each ProtecTIER QLogic HBA. Doing so increases the complexity of the structured cabling and it is not needed. It does not improve the redundancy of the ProtecTIER cluster. These ports are cabled only when you attach two back-end storage subsystems directly, without using a SAN network.

In addition, ensure that each zone that is defined on both SAN switches includes one initiator and one target only. Do not combine multiple targets on storage controllers with a single ProtecTIER initiator. The list of eight required zones is shown in Table B-6. Do not zone Channel 3 of Controllers A and Channel 5 of Controllers B.

Table B-6 Dual-node ProtecTIER cluster with two back-end storage subsystems

Effective configuration	
switch_A_zone_1	WWPN of ProtecTIER Node 1 HBA 6 Port 1 WWPN (1_S6P1) WWPN of DSx000 #1 Controller “A” Channel 1 (1_A CH1)
switch_A_zone_2	WWPN of ProtecTIER Node 1 HBA 6 Port 1 WWPN (1_S6P1) WWPN of DSx000 #2 Controller “A” Channel 1 (2_A CH1)
switch_A_zone_3	WWPN of ProtecTIER Node 2 HBA 6 Port 1 WWPN (2_S6P1) WWPN of DSx000 #1 Controller “A” Channel 1 (1_A CH1)
switch_A_zone_4	WWPN of ProtecTIER Node 2 HBA 6 Port 1 WWPN (2_S6P1) WWPN of DSx000 #2 Controller “A” Channel 1 (2_A CH1)

Effective configuration	
switch_B_zone_1	WWPN of ProtecTIER Node 1 HBA 7 Port 1 WWPN (1_S7P1) WWPN of DSx000 #1 Controller “B” Channel 7 (1_B CH7)
switch_B_zone_2	WWPN of ProtecTIER Node 1 HBA 7 Port 1 WWPN (1_S7P1) WWPN of DSx000 #2 Controller “B” Channel 7 (2_B CH7)
switch_B_zone_3	WWPN of ProtecTIER Node 2 HBA 7 Port 1 WWPN (2_S7P1) WWPN of DSx000 #1 Controller “B” Channel 7 (1_B CH7)
switch_B_zone_4	WWPN of ProtecTIER Node 2 HBA 7 Port 1 WWPN (2_S7P1) WWPN of DSx000 #2 Controller “B” Channel 7 (2_B CH7)

The same analogy applies when you connect a third and next back-end storage. Each zone that is set on both SAN switches is extended by two additional zones (four in total), similar to the even lines in the Table B-6 on page 472. The number #2 of the second DS storage subsystem is replaced by number 3, 4, and so on. Each *n* DS storage subsystem adds the following two WWPNs to each zone:

```
WWPN of DSx000 #n Controller "A" Channel 1 (n_A CH1)
WWPN of DSx000 #n Controller "B" Channel 7 (n_B CH7)
```

The DS3000 series

The DS3000 series consists of two products: the DS3500 and the DS3950. Both of these products are a good fit for the entry-level to midrange SAN and direct-attach market space, including the ProtecTIER environments. With the common Storage Manager that is shared by these DS3000 storage systems and the DS5000 storage systems, there is a smooth link in to the DS5000 series storage systems, with remote mirroring and copy services features being shared by these two platforms. The DS3500 and the DS3950 offer robust functionality, exceptional reliability, and availability with the common ease of shared storage management.

Note: As of ProtecTIER Version 3.3.0, the entry-level and midrange disk storage subsystems for DS3500 are no longer supported. However, DS3000, DS4000, and DS5000 storage that is attached to ProtecTIER systems that are running earlier releases continue to be supported. This section provides performance and capacity recommendations for the DS3500

The DS3500 series storage subsystems support up to two redundant RAID controllers in either a 12 or 24 drive configuration. The models for the storage servers are DS3512 and DS3524. There are also two models of drive expansion chassis (a 12 and a 24 drive) that can be attached to either of the storage subsystems. The models for these chassis are EXP3512 and EXP3524. However, there are certain limitations and recommendations about how to interconnect them with ProtecTIER servers.

EXP3500 attachment

The EXP3512 and EXP3524 expansion subsystems allow for the growth of the DS3500 storage subsystem up to the 96 drive maximum, by adding either the 12 or 24 drive chassis to the storage server's serial attached SCSI (SAS) drive expansion port. Any mix of the expansion models can be added, up to the maximum allowed drive count. The EXP3512 and EXP3524 differ from the DS3512 and DS3524 in that in place of the controller module, they are equipped with an Environmental Services Module (ESM). As with the DS3500 controllers, the expansions can be optionally upgraded with a second ESM module for redundant paths. Each ESM has a 6 Gbps SAS connection that provides 600 MBps throughput.

With the EXP3512 and EXP3524, only one of the two IN ports are used on each ESM to connect expansions together into a cascaded loop configuration. As shown in Figure B-13, the cabling scheme that is used for connecting these expansions follows what is known as a *top down, bottom up* method. This method provides the expansion loops with redundant paths to the enclosures, and in the event of one expansion encountering a catastrophic failure, the others are still able to continue to run. With a correct RAID layout, this configuration can provide uninterrupted operations.

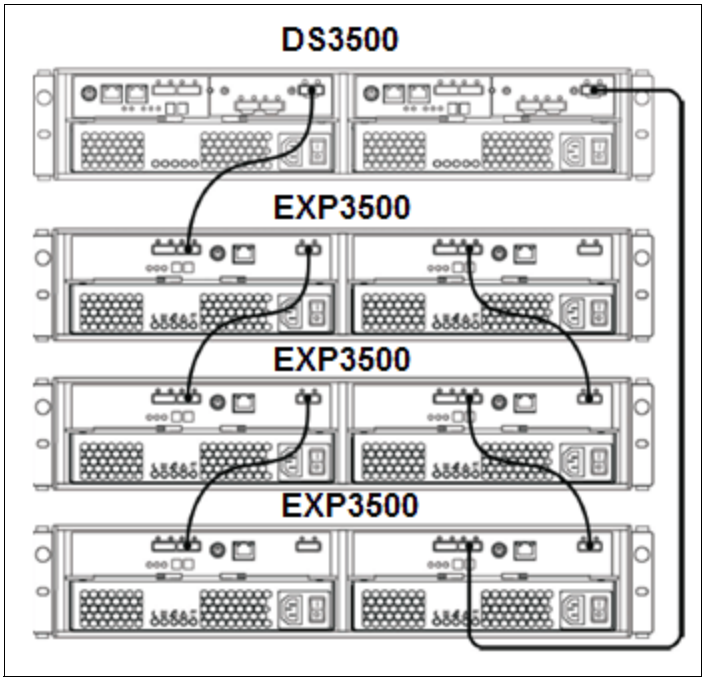

Figure B-13 EXP3500 expansions cascaded loop

Best practice: For optimal performance, do not attach more than three expansion units to the base dual-controller module. Instead, expand your ProtecTIER solution by adding an extra set of dual-controller modules DS3000 with their own EXP3500 units. Build the ProtecTIER repository by combining multiple DS3000 units.

A single DS3500 can in theory offer up to 500 MBps backup and restore performance, but practically size your solution for about 400 MBps, based on your hardware configuration. You should consider using the FC/SAS 10 K/15 K RPM disk drives, cache memory upgrade, licensed features, and so on.

You should always follow these two basic connection rules when you attach EXP3500 enclosures to a DS3500 storage subsystem:

- Connect the drive expansion SAS port on the DS3500 controller to the SAS IN port on the EXP3500.
- Connect the SAS OUT port on the EXP3500 to the SAS IN port on the next EXP3500.

Adding an ESM: The EXP3500 ships with one ESM installed. If you are going to attach the EXP3500 to a dual controller DS3500 storage subsystem, then you must install a second ESM in the EXP3500 expansion enclosure to provide redundant drive paths to the EXP3500 enclosures. Connecting an EXP3500 with one ESM installed to a DS3500, with two controllers that are installed, is not a supported configuration.

If the SAS cabling on the drive side is incorrect, it is detected by the DS3500 controller firmware. The Storage Manager application alerts you about the condition by logging a critical Major Event Log event. In addition, the Storage Manager Recovery Guru points to the mis-wired condition and advises you of the steps to correct the problem. A mis-wired condition is reported only if the wrong SAS connections results in a non-working configuration. It is also possible to attach the EXP3500 enclosures in a technically correct manner that is not optimal. Such configurations do not produce a mis-wire event. Therefore, you should carefully follow the recommendations that are documented in this chapter for best results.

An example of a non-optimal, but technically correct configuration, would be to connect both the left and right side ESMs in a top-to-bottom order. Although this configuration works and does not result in a mis-wire condition, it does not provide redundancy. If the entire EXP3500 fails, all the EXP3500 enclosures beneath it lose access to the DS3500 subsystem.

Cache upgrade

Each DS3000 storage controller is equipped with a 1 GB memory cache by default. Consider performing a cache memory upgrade to 2 GB per controller, especially if there are more than 48 disks in a repository (more than one EXP3500 expansion unit).

When you use 10 K RPM SAS for metadata or SATA disks for user data, always upgrade your cache to 2 GB per controller of every DS3000 unit in the repository chain.

Turbo Performance feature

If you plan for a high performance system, it is *strongly recommended* that you enable the Turbo Performance feature. Turbo Performance is an optional premium feature that, when enabled, boosts the performance of a storage system across both throughput and IOPS workloads, allowing you to take full advantage of DS3500 performance capabilities with bandwidth-intensive applications. The DS3500 Turbo Performance feature offers the following capabilities:

- Scalability to midrange performance and features that start at entry-level prices
- Efficiencies to help reduce annual energy expenditures and environmental footprint
- Simplicity that does not sacrifice control with the perfect combination of robustness and ease of use

Figure B-14 shows performance with and without the Turbo Performance premium feature enabled.

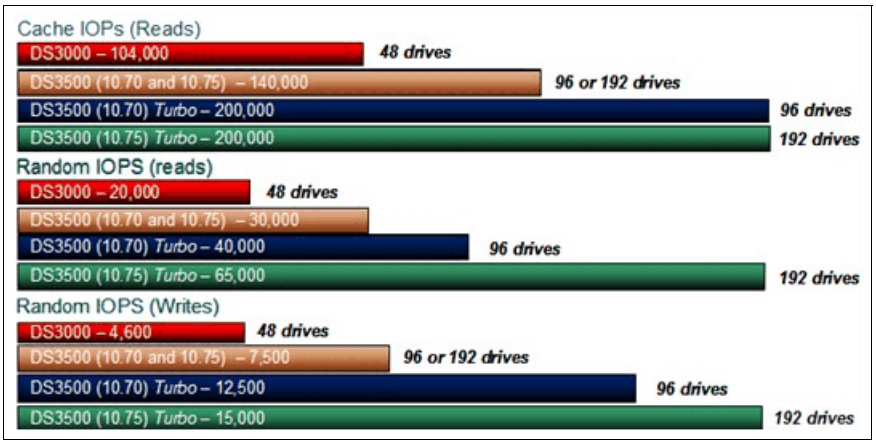

Figure B-14 The Turbo Performance feature characteristics

The Turbo Performance feature can enhance the dual-controller version of the DS3500 with a 30% improvement in IOPS and up to a 100% improvement in throughput. It is not supported on DS3500 with a single storage controller. This feature is not available on DS3950.

When you enable the Turbo Performance feature, consider the number of disks. When the system performance is limited by the low number of disks, the number of IOPS that is processed by disk controller becomes low. Relatively, the effect of the Turbo Performance feature becomes irrelevant to the license cost. In this case, consider deploying more DS3500 systems.

Media scan

Media scan is a background process that checks the physical disks for defects by reading the raw data from the disk and writing it back, which detects possible problems that are caused by bad sectors of the physical disks before they disrupt normal data reads or writes. This process is sometimes known as *data scrubbing*.

Media scan continuously runs in the background, using spare cycles to complete its work. The default media scan is for a scan every 30 days, that is, the maximum time media scan has to complete the task. During the scan process, the DS3500 system calculates how much longer the scan process will take to complete, and adjusts the priority of the scan to ensure that the scan completes within the time allocated. After the media scan completes, it starts over and resets its time for completion to the current setting. This media scan setting can be reduced, however. If the setting is too low, priority is given to media scan over host activity to ensure that the scan completes in the allocated time. This scan can impact performance, but improve data integrity.

A media scan can be considered a surface scan of the hard disk drives, and a redundancy check scans the blocks of a RAID 3, 5, or 6 logical drive and compares it against the redundancy data. In the case of a RAID 1 logical drive, the redundancy scan compares blocks between copies on mirrored drives.

We have seen no effect on I/O with a 30 day setting unless the processor is used in excess of 95%. The length of time that it takes to scan the LUNs depends on the capacity of all the LUNs on the system and the usage of the controller.

C

Networking

This appendix describes general information about bonding for the backup server side. Because we described the concepts of bonding in Chapter 5, “Networking essentials” on page 79, this appendix provides only a basic overview of potential bonding options of different host operating systems. You must ensure that all involved communication partners within the data flow of your backup are configured to allow bonding to work. Reach out to your networking staff and work together with them on implementing the network solution that best suits your needs.

This appendix covers the following topics:

- Bonding on Linux and UNIX machines
- Teaming on Microsoft based machines
- Link aggregation and similar technologies on the LAN switches

Bonding on Linux and UNIX machines

This section describes bonding types on Linux machines and on UNIX machines.

Bonding on Linux machines

This section describes bonding types on Linux machines. In Linux machines, the following bonding types exist:

Mode 0	Sets a round robin policy for fault tolerance and load balancing.
Mode 1	Sets an active-backup policy for fault tolerance.
Mode 2	Sets an exclusive-OR (XOR) policy for fault tolerance and load balancing.
Mode 3	Sets a broadcast policy for fault tolerance. All transmissions are sent on all subordinate interfaces.
Mode 4	Sets an IEEE 802.3ad dynamic link aggregation policy.
Mode 5	Sets a Transmit Load Balancing (TLB) policy for fault tolerance and load balancing.
Mode 6	Sets an Active Load Balancing (ALB) policy for fault tolerance and load balancing. Includes transmit and receive load balancing for IPV4 traffic. Receive load balancing is achieved through Address Resolution Protocol (ARP) negotiation.

Modes 2 and 4 use a default transmit hash policy of Layer 2 (source MAC destination MAC)%N (number of subordinates). The hash policy can be modified to Layer 3 + 4, where both the source and destination IP and port are considered.

Bonding on UNIX machines

This section describes bonding types on UNIX machines. These modes are supported in IBM AIX:

Default hash mode	The traditional IBM AIX type. The adapter selection algorithm uses the last byte of the destination IP address (for TCP/IP traffic) or MAC address (for ARP and other non-IP traffic).
Standard or 802.3ad	src_dst_port hash mode. The outgoing adapter path is selected by an algorithm by using the combined source and destination TCP or UDP port values.
Standard or 802.3ad src_port	The adapter selection algorithm uses the source TCP or UDP port value.
Standard or 802.3ad dst_port	The outgoing adapter path is selected by the algorithm by using the destination system port value.
Round-robin	Outgoing traffic is spread evenly across all of the adapter ports in the Etherchannel.

Teaming on Microsoft based machines

On Microsoft based machines, the teaming methods are defined by the Network Interface Card (NIC) vendor. This section describes the teaming modes for Broadcom and Intel NICs. The teaming modes that are described are only the ones that are used for load sharing.

Broadcom NICs

Broadcom supports the following balancing modes:

Smart Load Balancing (SLB) In this method, both transmit and receive load balancing are enabled, based on source and destination L3/L4 IP addresses and TCP/UDP port numbers.

Generic trunking In this switch-assisted teaming mode, the LAN switch to which the server is attached also must be configured for one of the aggregation methods. As is the case for SLB, the IP/TCP/UDP source and destination addresses load balance the transmit traffic from the server.

Link aggregation (IEEE 802.3ad LACP) Link aggregation is similar to generic trunking except that it uses the Link Aggregation Control Protocol (LACP) to negotiate the ports that make up the team.

Intel NICs

Intel supports the following balancing modes:

Adaptive Load Balancing (ALB) This method allows transmission over 2 - 8 ports to multiple destination addresses, along with fault tolerance. In this method, transmit is done through 2 - 8 adapters in load balancing, while the team receives packets only through the main adapter. This method works on Layer 3 and 4.

Receive Load Balancing (RLB) This method, which can be configured only with ALB, adds the receive load balancing feature to it, and is also based on Layer 3 and 4. This method is switch-less.

Virtual Machine Load Balancing (VMLB) Provides transmit and receive traffic load balancing across virtual machines that are bound to the team interface, and fault tolerance in the event of switch port, cable, or adapter failure. This teaming type is a switch-less method.

IEEE 802.3ad In this method, the standard supports static and dynamic modes. Intel supports both modes, and must be configured with a LAN switch that supports the 802.3ad standard or Cisco Etherchannel technology.

Link aggregation and similar technologies on the LAN switches

For LAN switches, configurations are provided for Cisco switches. Some of the features are present in Cisco IOS Release 12.2(33)SB and higher. The 802.3ad is a market-wide standard, supported by all common vendors. In Cisco switches, both Layer 2 and Layer 3, two methods are available:

- Etherchannel – GEC/FEC (Giga/Fast Ethernet ports), and PaGP
- 802.3ad Link Aggregation and LACP control protocol

In switches from other vendors (Avaya, Juniper, 3Com, Hewlett-Packard, and others), only 802.3ad is used. Etherchannel is a Cisco proprietary technology. In Cisco versions up to Release 15.0(1)S, mechanisms for load balancing Ethernet service instances over member links in a port channel do not account for the service instances traffic loads, which can lead to unequal distribution of traffic over member links.

In IOS Release 15.0(1)S, a new feature was introduced: the 802.3ad Link Aggregation with Weighted Load Balancing feature (802.3ad LAG with WLB). You can use it to assign weights to service instances to efficiently distribute traffic flow across active member links in a port channel.

The LAG with WLB feature supports both LACP (active or passive mode) and manual (mode on) Etherchannel bundling. A weighted load balancing configuration does not affect the selection of active member links in the Etherchannel. As member links become active or inactive, a load-balancing algorithm adjusts the distribution of Ethernet service instances to use the active member links.

D

Managing cartridge sizes with ProtecTIER

This appendix provides general information about managing cartridge sizes with ProtecTIER. With ProtecTIER, the total amount of space that is available in your repository is variable. These dynamics can make managing cartridges complex. This appendix is intended to help administrators to plan and manage their ProtecTIER cartridges and repository.

This appendix covers the following topics:

- Effects of dynamic cartridge sizes
- The mechanism behind fluctuating cartridge sizes

Effects of dynamic cartridge sizes

In a long running ProtecTIER environment, the HyperFactor ration tends to stabilize around a certain value. If you then encounter changes in your environment, these changes have an unexpected impact on the HyperFactor ratio. When the HyperFactor ratio changes, the amount of free space that is available in your repository is recalculated. A higher HyperFactor ratio results in more free space being available. A lower HyperFactor ratio results in less free space being available.

After such events, you might find it difficult to determine how many scratch tapes or free space are available. Also, you might see virtual cartridges that are marked as full before such an event is expected.

The mechanism behind fluctuating cartridge sizes

All virtual cartridges get an equal amount of space up to either the LIMIT size that is configured when the virtual cartridge was created (maximum cartridge size), or the calculated amount of Nominal Space / Number of Carts.

After a tape is full, an early warning is sent to the backup application, and the tape is removed from the calculation of Number of Carts. This action allows the system to adjust the size of the remaining (non-full) tapes to use the nominal space effectively.

If the nominal space is not large enough to hold the total of Number of Carts * the LIMIT size of the cart, then those tapes are marked full before they reach their size (that is, there is no thin provisioning)

Knowing this behavior can help an administrator know when it is best to add more cartridges and whether it is best to limit the size of the carts or leave them unlimited.

IBM recommends to always define a fixed cartridge size. Values of 100 GB for a virtual cartridge are reasonable. If you use 100 GB virtual cartridges, the database/catalog of the backup application will not have an unreasonably high number of cartridges, and therefore be unmanageable. Also, the housekeeping jobs on 100 GB virtual cartridges can be ran reasonably early to free up unused space and make it available to ProtecTIER again (as relabel scratch space). Deciding on a limit for the virtual cartridge size allows you to reserve space or divide space in libraries. With this strategy, you can prevent one of your virtual libraries from using up all the space in your repository, causing the other libraries to run out of space.

For example, if you have two libraries, one with 500 carts and the other with 1000 carts, your usage is 33% for the smaller library, and 66% for the larger library.

The decision about how many carts you have should be based on the nominal space in the repository. The number of carts should be either calculated for a wanted approximate size (unlimited) or calculated so that there is enough space for all carts to reach the limit that is set, while leaving a little room for fluctuation (the factoring ratio).

Here are some cases where having the wrong cartridge size has led to problems:

- You might think that because you have low scratch tapes that adding scratch tapes at a limited size brings more space to the backup application. This situation works if there is a reserve of nominal space in the repository. However, if the factoring ratio is lower then planned, and less space results, adding cartridges results in an even smaller cartridge size.

Important: Adding more virtual cartridges does *not* increase the available free space in your repository. If you face an out-of-space condition, you must expand your ProtecTIER repository by adding more physical disk to the back end.

- An insufficient idle time for background jobs results in a build-up of delete and defragmentation data, which reduces allocatable space for backups. Adding cartridges in this scenario can reduce the size of the cartridges, which has a negative impact.
- An extra library with 1000 extra empty tapes results in the production library running out of room, even though the repository showed plenty of nominal space. Adding tapes in this case results in even smaller tapes as well.
- Collocation of backups with many partial tapes might make the repository appear to have much space, but the partial tapes consume more space than expected because of the size and usage. In this case, limiting the tapes to a smaller size could allow the tapes to store data more efficiently.
- Finally, if the repository has more space available when you multiply the total number of tapes by the amount of space, then it is a good plan to add more limited cartridges (This situation will not happen if you do not have any limited size tapes.)

If you have a mixture of tape sizes and types, the management of these tapes becomes complicated, especially if you have a large variance in the sizes. If you want to “unlimit” all of your existing tapes, contact IBM Support and request help with running the support utility to “unlimit” the cartridge sizes.

Glossary

3958 DD1 This is the original server, which has been available since August 2008. This server is based on the IBM System x3850 M2 Type 7141. When it i used as a server in the TS7650G, its machine type and model are 3958 DD1. Use this machine type and model for service purposes.

3958 DD3 This is a higher performance server, which has been available since March 2009. This server is based on the IBM System x3850 M2 Type 7233. When used as a server in the TS7650G, its machine type and model are 3958 DD3. Use this machine type and model for service purposes.

3958 DD4 This is a newer, higher performance server, which has been available since December 2010. This server is based on the IBM System x3850 X5 Type 7145-AC1. When used as a server in the TS7650G, its machine type and model are 3958 DD4. Use this machine type and model for service purposes.

3958 DD5 This is a newer, higher performance server, which has been available since May 2012. This server is based on the IBM System x 7145 model. When used as a server in the TS7650G, its machine type and model are 3958 DD5. Use this machine type and model for service purposes.

asynchronously parallel system A system in which the backed up data does not return to the host (and out to file) in the same order each time.

Backup, Recovery, and Media Services for IBM i (BRMS) Helps you implement a disciplined approach to managing your backups, and provides you with an orderly way to retrieve lost or damaged data. BRMS also enables you to track all of your backup media from creation to expiration

BRMS policies A set of defaults that is commonly used (for example, device or media class). Generally used defaults are in the BRMS system policy. Backup-related defaults are in the BRMS backup policy.

chown The `chown` command (abbreviation for **ch**ange **own**er) is used on UNIX based systems to change the owner of a file. In most implementations, it can be run by only the superuser to prevent users from simply changing the ownership of files randomly.

compaction (data compaction) The reduction of the number of data elements, bandwidth, cost, and time for the generation, transmission, and storage of data without loss of information by eliminating unnecessary redundancy.

Common Internet System (CIFS) ProtecTIER emulates Windows file system behavior and presents a virtualized hierarchy of file systems, directories, and files to Windows CIFS clients. When configured for FSI-CIFS, ProtecTIER emulates a network-attached storage (NAS) backup target that can use both HyperFactor and ProtecTIER native replication bandwidth reduction techniques for storing and replicating deduplicated data.

concurrent saves and restores The ability to save or restore different objects from a single library or directory to multiple backup devices or different libraries or directories to multiple backup devices at the same time from different jobs.

control group A group of items (for example, libraries or stream files) to back up, and the attributes that are associated with how to back them up.

CSV file (.csv) Comma-separated **v**alue file, sometimes called comma-delimited. This type of file is a specially formatted plain text file that stores spreadsheet or basic database-style information in a simple format, with one record on each line, and each field within that record separated by a comma. CSV files are used by ProtecTIER Manager as a simple way to transfer a large volume of database information between programs. This type of file can be imported into most spreadsheet programs.

deduplication A data compression technique in which redundant data is eliminated. The technique improves storage usage and can also be applied to network data transferals to reduce the number of bytes that must be sent across a link**.**

direct attachment Refers to a digital storage system that is directly attached to a server or workstation, without a storage network in between.

dirty bit A dirty bit is a flag that indicates whether an attribute must be updated. This situation occurs when a bit in a memory cache or virtual memory is changed by a processor but is not updated in storage.

dirty bit technology The ProtecTIER system uses a “dirty-bit” feature/technology and cartridges are marked as in-sync after the data finishes replicating from the primary to the secondary site, so that at the time of synchronization, the local cartridges and their DR site replicas are identical.

disaster recovery (DR) The process of recovering production site data at a DR location. Disaster recovery is useful if a disaster occurs or a situation occurs where the production (or primary) site goes offline.

disk controller The disk controller for the TS7650 Appliance is IBM Feature Code 3708: 4.8 TB Fibre Channel Disk Controller. Use this feature code for service purposes.

factoring ratio The ratio of nominal capacity to physical capacity in the ProtecTIER repository. For example, if you have 100 TB of user data (nominal capacity) and it is stored on 10 TB of physical capacity, your factoring ratio is 10:1.

File System Interface (FSI) The File System Interface (FSI) presents ProtecTIER as a network-attached storage backup and recovery target that can use the HyperFactor algorithm and ProtecTIER native replication bandwidth reduction techniques for storing and replicating deduplicated data. The FSI configuration option allows ProtecTIER to present disk repository storage as a virtualized hierarchy of file systems.

disk expansion unit The disk expansion unit for the TS7650 Appliance is IBM Feature Code 3707: 4.8 TB Fibre Channel Disk Expansion Unit. Use this feature code for service purposes.

failback The procedure for replicating updated cartridges, new or old, from the DR site to the original (or restored) production site to bring it up to date in case it was down, or lost and rebuilt.

hub The hub (target server) is connected, through your Ethernet replication network, to one or more spokes (source servers). The hub stores backup images of the data repositories, file system configurations, and other system settings that are on the spokes. If there is a spoke failure, the stored image can be easily retrieved from the hub and transmitted to the spoke. This action restores the spoke to its previous configuration with minimal data loss

IBM Tivoli Assist On-site (AOS) IBM Tivoli Assist On-site (AOS) is a web-based tool that enables a remote support representative from IBM to view or control the management node desktop. More information can be found at the Tivoli Assist On-site website at `http://www.ibm.com/support/assistonsite`.

IP address Internet Protocol address. A numerical label that is assigned to each device that participates in a computer network.

load throttling Load throttling is a process that helps avoid dangerous overload situations. Load throttling limits the number of permitted incoming connections, allowing resources to be allocated to all processes.

logical partition (LPAR) A division of a computer's processors, memory, and storage into multiple sets of resources so that each set of resources can be operated independently with its own operating system instance and applications. The number of LPARs that can be created depends on the system's processor model and resources that are available.

logical unit number (LUN) A number that is used to identify a logical unit that is a device that is addressed by Fibre Channel. A LUN may be used with any device that supports read/write operations, such as a tape drive, but is most often used to refer to a logical disk that is created on a SAN.

LUN masking An authorization process that makes a LUN available to some hosts and unavailable to other hosts. LUN masking is used in the ProtecTIER product as a precaution against servers corrupting disks that belong to other servers. By masking (hiding) LUNs from a specific server (or servers), you effectively tell those servers that the LUN does not exist, and those servers cannot corrupt the disks in the LUN.

media A tape cartridge (volume) that holds saved data.

media class A logical grouping of media with similar physical, logical, or both of these characteristics (for example, density).

media identifier A name that is given to a physical piece of media.

Network File System (NFS) ProtecTIER emulates UNIX file system behavior and presents a virtualized hierarchy of file systems, directories, and files to UNIX based clients using the NFS protocol. When configured for FSI-NFS, ProtecTIER emulates a network-attached storage (NAS) backup target that can use both HyperFactor and ProtecTIER Native Replication bandwidth reduction techniques for storing and replicating deduplicated data.

nominal capacity of the repository The physical space and expected factoring ratio.

nominal data The original amount of backed-up data before you apply the ProtecTIER deduplication factor.

OpenStorage (OST) Allows the ProtecTIER product to be integrated with NetBackup to provide backup-to-disk without using a Virtual Tape Library (VTL) emulation.

parallel saves and restores The ability to save or restore a single object or library or directory across multiple backup devices from the same job.

peak throughput The maximum of the ProtecTIER server capabilities.

principality The privilege to write to a cartridge (set to read/write mode). The principality of each cartridge belongs to only one repository in the grid. By default, the principality belongs to the repository where the cartridge was created.

ptcli ProtecTIER command-line interface.

ProtecTIER When used by itself, this expression points to the IBM patented deduplication solution based on HyperFactor. Depending on the context, it can mean the family of products, a specific device, or just the deduplication engine.

RAID 10 Many storage controllers allow RAID levels to be nested. The elements of a RAID may be either individual drives or RAIDs themselves. Thus, a RAID 10 (or RAID 1+0) is a configuration in which multiple drives are first combined into multiple RAID arrays. Each RAID 1 array is treated as a single drive. These arrays are then combined into a single RAID 0 array.

RAID 6 The RAID 6 architecture supports block-level striping with double distributed parity. It provides fault tolerance of two drive failures; the array continues to operate with up to two failed drives. This situation makes larger RAID groups more practical, especially for high-availability systems. This situation becomes increasingly important as large-capacity drives lengthen the time that is needed to recover from the failure of a single drive. Single-parity RAID levels are vulnerable to data loss until the failed drive is replaced and its data rebuilt. Double parity gives more time to rebuild the array without the data being at risk if another drive fails before the rebuild is complete.

reclamation The Tivoli Storage Manager process that frees up space on tapes, and returns empty tapes to the scratch pool. Reclamation is accomplished by deleting expired data from tapes and moving any unexpired data to other tapes to more efficiently use tape space.

recovery point objective (RPO) How much lag time is acceptable for a backup that is written to virtual tape in Site A to be replicated to Site B.

redundant array of independent disks (RAID) A storage technology that combines multiple disk drive components into a logical unit. Data is distributed across the drives in one of several ways (RAID levels). The physical drives are said to be in a RAID array, which is accessed by the operating system as one single drive. The different schemes or architectures are named by the word RAID followed by a number (for example, RAID 0 or RAID 1).

redundant host connection The duplication of connections, or two or more paths that connect two objects in a network. The intention of redundancy is to increase the reliability of the system, or to provide a backup or failsafe if one of the connections fails.

remote cloning The process of using a secondary (DR) site to clone cartridges. ProtecTIER replication enables users to offload tape cloning to their secondary site.

replication A process that transfers logical objects, such as cartridges, from one ProtecTIER repository to another one.

replication grid A set of repositories that shares a common ID and can potentially transmit and receive logical objects through replication.

replication grid ID A number 0 - 63 that identifies a replication grid within an organization.

replication grid member A repository that is a member in a replication grid.

Replication Manager The utility in the ProtecTIER Manager GUI through which you can set replication policies, define replication time frame windows, delete replication activities, and much more.

replication pairs Two repositories within a replication grid that replicate from one to another.

replication policy A policy made up of rules that define a set of objects (for example, VTL cartridges) from a source repository to be replicated to a target repository.

replication rate control (RRC) A built-in resource-governing mechanism. RRC gives precedence to backup and restore requests and throttles down replication traffic whenever backup and restore activity increases above an idle state. When the backup and restore workload drops below that idle threshold, RRC returns to the default priority. The RRC determines the maximum replication rate for both system states, IDLE and BUSY, based on the performance limits set by the user.

replication time frame A scheduled period for replication to take place for all policies.

replication window The time frame during which replication runs.

repository A warehouse to store data for safekeeping.

repository unique ID (RID) A number that uniquely identifies the repository. The RID is created from the replication grid ID and the repository internal ID in the grid.

SAN fabric The hardware that connects workstations and servers to storage devices in a SAN is referred to as a fabric. The SAN fabric enables any server to connect to any storage device through Fibre Channel switching.

shelf A container of VTL cartridges within a ProtecTIER repository.

SLA Service level agreement.

spoke The spokes are the servers that process and store the information that is generated during daily business operations. The stored information is then replicated to a hub, according to a user-defined replication policy.

storage area network (SAN) A dedicated network that provides access to consolidated, block-level data storage. SANs make storage devices, such as disk arrays and tape libraries, accessible to servers so that the devices appear like locally attached devices to the operating system.

storage checkpoint Storage checkpoints are pointers added to the backup stream so that if the backup fails, a rerun of the backup starts from the last storage checkpoint, rather than the beginning of the stream.

system console The system console is an IBM TS3000 System Console (TSSC). It is the console used with a keyboard to issue commands to the server through the CLI, and monitors the functions of the server. The TS3000 combined with the keyboard is commonly referred to as the KVM (keyboard, video, monitor).

TS7600 When used alone, this term signifies the IBM family of virtualization solutions that operate on the ProtecTIER platform.

Virtual Tape Library (VTL) The ProtecTIER VTL service that emulates traditional tape libraries.

visibility switching The automated process that transfers the visibility of a VTL cartridge from its master to its replica and vice versa.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- *DS8800 Performance Monitoring and Tuning*, SG24-8013
- *IBM System Storage DS8870 Architecture and Implementation*, SG24-8085
- *IBM System Storage Solutions Handbook*, SG24-5250
- *IBM System Storage Tape Library Guide for Open Systems*, SG24-5946
- *IBM System Storage TS7600 with ProtecTIER Version 3.3*, SG24-7968
- *IBM Tivoli Storage Manager Implementation Guide*, SG24-5416
- *IBM XIV Storage System: Host Attachment and Interoperability*, SG24-7904
- *Implementing IBM Storage Data Deduplication Solutions*, SG24-7888
- *Implementing the IBM Storwize V7000 V6.3*, SG24-7938
- *Implementing the IBM System Storage SAN Volume Controller V6.3*, SG24-7933
- *Introduction to IBM Real-time Compression Appliances*, SG24-7953
- *SAP Backup using Tivoli Storage Manager*, SG24-7686

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Other publications

These publications are also relevant as further information sources.

Publications common to the TS7650 Appliance and TS7650G

- *IBM System Storage Labeling Instructions for the TS7650 ProtecTIER Deduplication Appliance and TS7650G ProtecTIER Deduplication Gateway*, 95P8942
- *IBM System Storage TS7600 with ProtecTIER Installation Instructions for the RAS Package, BIOS, and Firmware updates following a FRU replacement for models 3958 DD1, 3958 DD3, and 3958 AP1*, 46X6059
- *IBM System Storage TS7600 - ProtecTIER User's Guide for FSI Systems, v.3.3,* GA32-2235

- *IBM System Storage TS7600 - ProtecTIER User's Guide for OpenStorage Systems,* v.3.3, GA32-2234
- *IBM System Storage TS7600 - ProtecTIER User's Guide for VTL Systems, V3.3*, GA32-0922
- I*BM System Storage TS7650 ProtecTIER 3.3 Deduplication Appliance and TS7650G Deduplication Gateway Introduction and Planning Guide* (3958-AP1 and 3958-DD5), GA32-0918
- *IBM System Storage TS7650 ProtecTIER Deduplication Solutions ProtecTIER version 3.3 TS7650 and TS7650G Problem Determination and Service Guide*, GA32-0923
- *IBM System Storage TS7650 with ProtecTIER for the TS7650 and TS7650G Software Upgrade Guide for the TS7650 ProtecTIER 3.3 Deduplication Appliance and TS7650G ProtecTIER Deduplication Gateway*, SC27-3643
- *IBM System Storage TS7650G ProtecTIER Deduplication Gateway for ProtecTIER 3.3 Installation Roadmap Guide*, GA32-0921

TS7650 Appliance publications

- *IBM System Storage TS7650 ProtecTIER 3.2 Deduplication Appliance Installation Roadmap Guide [3958 AP1]*, GA32-0920
- *IBM System x3850 M2 and System x3950 M2 Type 7141 and 7233 User's Guide*:

  http://content.etilize.com/User-Manual/1013913138.pdf
- *Start Here for IBM System Storage TS7650 with ProtecTIER 3.3*, GI13-1831-03

TS7610 Appliance Express and TS7620 Appliance Express publications

- *IBM System Storage TS7610 ProtecTIER Deduplication Appliance Express V3.3 Increasing Capacity on the 3959 SM1 from 4 TB to 6 TB (Feature Code 9314),* SC27-3642
- *IBM System Storage TS7610 and TS7620 ProtecTIER 3.3 Deduplication Appliance Express Service Guide*, GA32-0915
- *IBM System Storage TS7610 and TS7620 ProtecTIER Deduplication Appliance Express - ProtecTIER Maintenance Guide,* v3.3, GA32-2232
- *IBM System Storage TS7610 and TS7620 ProtecTIER Deduplication Appliance Express - ProtecTIER User's Guide for FSI, v3.3,* GA32-2231
- *IBM System Storage TS7610 and TS7620 ProtecTIER Deduplication Appliance Express - ProtecTIER User's Guide for OpenStorage Systems,* v3.3, GA32-2230
- *IBM System Storage TS7610 and TS7620 ProtecTIER Deduplication Appliance Express - ProtecTIER User's Guide for VTL Systems, v3.3*, GA32-0916
- *IBM System Storage TS7610 or TS7620 ProtecTIER V3.3 Deduplication Appliance Express V3.3 Software Upgrade Guide*, SC27-3641
- *IBM System Storage TS7620 ProtecTIER 3.3 Deduplication Appliance Express Introduction and Planning Guide*, GA32-0913
- *IBM System Storage TS7620 ProtecTIER 3.3 Deduplication Appliance Express Installation and Setup Guide for VTL, OpenStorage, and FSI systems,* GA32-0914

- *IBM System Storage TS7620 ProtecTIER Deduplication Appliance Express V3.3 Increasing Capacity on the 3959 SM2 from 6 TB to 12 TB (Feature Code 9317)*, GA32-2222
- *IBM System Storage TS7620 ProtecTIER V3.3 Deduplication Appliance Express - Feature Code 9345 (3959 EXP), Field Installation of Expansion Drawer*, SC27-5413

Integrated Management Module and Remote Supervisor Adaptor publications

- *Integrated Management Module User's Guide*:

  http://download.boulder.ibm.com/ibmdl/pub/systems/support/system_x_pdf/00d2490.pdf
- *Remote Supervisor Adapter II Slimline and Remote Supervisor Adapter II User's Guide*:

  http://download.boulder.ibm.com/ibmdl/pub/systems/support/system_x_pdf/43w7827.pdf

Online resources

These websites are also relevant as further information sources:

- Compatibility matrix:

  http://www-03.ibm.com/systems/storage/tape/library.html#compatibility
- IBM i resources for ProtecTIER and tape

  http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS4956
- IBM Interoperability Guide:

  http://www.ibm.com/systems/support/storage/config/ssic/index.jsp
- IBM SAN Volume Controller (SVC):

  http://publib.boulder.ibm.com/infocenter/svc/ic/index.jsp
- IBM Storwize V7000:

  http://publib.boulder.ibm.com/infocenter/storwize/ic/index.jsp
- IBM Storwize V7000 Unified Storage:

  http://publib.boulder.ibm.com/infocenter/storwize/unified_ic/index.jsp
- IBM Support website:

  http://www.ibm.com/systems/support/storage/config/ssic/index.jsp
- IBM Tivoli Storage Manager V6.3 Information Center:

  http://pic.dhe.ibm.com/infocenter/tsminfo/v6r3/index.jsp
- IBM XIV:

  http://www.xivstorage.com/
- List of supported Fibre Channel switches:

  http://www.ibm.com/systems/support/storage/config/ssic/index.jsp
- Red Hat Linux:

  https://www.redhat.com/wapps/store/allProducts.html

- Symantec NetBackup:
  - http://www.symantec.com/enterprise/products/overview.jsp?pcid=1018&pvid=2_1
  - http://www.symantec.com/enterprise/support/index.jsp
- TS7610 and TS7620 V3.3 Customer Information Center:

  http://publib.boulder.ibm.com/infocenter/ts7610/cust/index.jsp
- TS7650 Combined Customer Information Center:

  http://publib.boulder.ibm.com/infocenter/ts7650/cust/index.jsp
- TS7650 V3.3 Customer Information Center:

  http://pic.dhe.ibm.com/infocenter/ts7650/cust/index.jsp
- TS7650G supported disk arrays:

  http://www.ibm.com/systems/support/storage/config/ssic/index.jsp
- TS7650G supported operating systems:

  http://www.ibm.com/systems/support/storage/config/ssic/index.jsp

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Redbooks

IBM ProtecTIER Implementation and Best Practices Guide

(1.0” spine)
0.875”<->1.498”
460 <-> 788 pages

IBM ProtecTIER Implementation and Best Practices Guide

Implement scenarios for backup recovery applications with ProtecTIER

Optimize data deduplication and storage efficiency

Exploit the File System Interface using CIFS and NFS

This IBM Redbooks publication provides best practice guidance for planning, installing, and configuring the IBM TS7600 ProtecTIER family of products. This guide provides all the latest best practices for using ProtecTIER Software Version 3.3 and the revolutionary and patented IBM HyperFactor deduplication engine, along with other data storage efficiency techniques, such as compression and defragmentation.

The IBM System Storage TS7650G ProtecTIER Deduplication Gateway and the IBM System Storage TS7620 ProtecTIER Deduplication Appliance Express are disk-based data storage systems that are configured for three available interfaces:

- The Virtual Tape Library (VTL) interface is the foundation of ProtecTIER and emulates traditional automated tape libraries.
- The Symantec NetBackup OpenStorage (OST) API can be integrated with Symantec NetBackup to provide backup-to-disk without having to emulate traditional tape libraries.
- The newly available File System Interface (FSI) supports Common Internet File System (CIFS) and Network File System (NFS) as backup targets.

For your existing ProtecTIER solution, this guide provides best practices and suggestions to boost the performance and the effectiveness of the data deduplication with regards to your application platforms for your VTL, OST, and FSI systems.

When you build a ProtecTIER data deduplication environment, this guide helps your IT architects and solution designers plan for the best option and scenario for data deduplication for their environments. This guide helps you optimize your deduplication ratio, while reducing the hardware, power and cooling, and management costs.

This guide provides expertise that was gained from the IBM ProtecTIER Field Technical Sales Support (FTSS/CSS) Group, development, and Quality Assurance teams.

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

SG24-8025-01 ISBN 0738438081