


# IBM System Storage TS7650, TS7650G, and TS7610

Many to one native replication

Benefits of deduplication

Version 2.5 introduced



Alex Osuna  
Eva Balogh  
DaeSung Kim  
Kai Lanzmich  
Karen Orlando  
Siew Chin Wong

**Redbooks**





International Technical Support Organization

**IBM System Storage TS7650, TS7650G, and TS7610**

August 2011

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

**Fourth Edition (August 2011)**

This edition applies to IBM ProtecTIER Version 2.5

© Copyright International Business Machines Corporation 2010, 2011. All rights reserved.

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

# Contents

<b>Notices</b> .....	xiii
Trademarks .....	xiv
<b>Preface</b> .....	xv
The team who wrote this book .....	xv
Now you can become a published author, too! .....	xvii
Comments welcome .....	xvii
Stay connected to IBM Redbooks .....	xvii
<b>Summary of changes</b> .....	xix
August 2011, Fourth Edition .....	xix
August 2010, Third Edition .....	xix
March 2010, Second Edition .....	xx

## **Part 1. Introduction and architecture** ..... 1

<b>Chapter 1. Concepts of data deduplication</b> .....	3
1.1 Data deduplication .....	4
1.2 Types of data deduplication .....	5
1.2.1 Hash-based data deduplication .....	5
1.2.2 Content aware .....	7
1.2.3 HyperFactor, deduplication, and bandwidth savings .....	7
1.3 Data deduplication processing .....	9
1.3.1 Inline method .....	9
1.3.2 Post-processing method .....	9
1.4 Components of a data deduplication system .....	9
1.4.1 Server .....	10
1.4.2 Data deduplication software .....	10
1.4.3 Disk array .....	10
1.5 Benefits of data deduplication .....	10
1.5.1 Reduction of storage requirements .....	10
1.5.2 Reduction of environmental costs .....	10
<b>Chapter 2. IBM System Storage ProtecTIER architecture</b> .....	11
2.1 General overview of ProtecTIER .....	12
2.1.1 Types of ProtecTIER models .....	12
2.1.2 ProtecTIER service modes .....	13
2.1.3 ProtecTIER Manager .....	13
2.1.4 ProtecTIER Replication Manager .....	13
2.2 Terms and definitions .....	14
2.3 TS7650G ProtecTIER Deduplication Gateway .....	17
2.3.1 TS7650G Gateway models .....	17
2.3.2 TS7650G ProtecTIER Deduplication Gateway (3958-DD4) .....	18
2.3.3 Disk array .....	20
2.3.4 Deployment .....	20
2.4 TS7650 ProtecTIER Deduplication Appliance .....	26
2.4.1 TS7650 Deduplication Appliance .....	26
2.4.2 TS7650 ProtecTIER Deduplication Appliance features .....	26
2.4.3 Available models .....	27

2.4.4	Deployment	27
2.5	TS7610 ProtecTIER Deduplication SMB Appliance	29
2.5.1	TS7610 hardware components	30
2.5.2	Deployment	31
2.6	ProtecTIER virtual tape library	32
2.6.1	ProtecTIER VTL concepts	33
2.6.2	Steady state	34
2.7	ProtecTIER OpenStorage	36
2.8	Data deduplication	37
2.8.1	Virtual tape library concept	38
2.8.2	HyperFactor	40
2.8.3	ProtecTIER data ingest flow	41
2.9	ProtecTIER native replication	42
2.9.1	Replication licensing	43
2.9.2	Hardware	43
2.9.3	Replication grid	45
2.10	ProtecTIER Manager	46
2.10.1	Virtual libraries	48
2.10.2	Virtual drives	49
2.10.3	Virtual cartridges	49
2.11	IBM TS3000 System Console	50
2.12	Operating system	50
<b>Part 2.</b>	<b>Planning for data deduplication and replication</b>	<b>51</b>
<b>Chapter 3.</b>	<b>Planning for deduplication and replication</b>	<b>53</b>
3.1	Planning for deduplication	54
3.1.1	Sizing inputs	56
3.2	APTARE overview	61
3.2.1	APTARE StorageConsole Backup Manager	61
3.2.2	APTARE architecture	62
3.2.3	Managed backup environment security	63
3.2.4	APTARE setup	64
3.2.5	APTARE reports	68
3.3	Throughput considerations	68
3.3.1	Attached host systems	69
3.3.2	SAN connectivity	69
3.3.3	Disk array	70
3.3.4	Data type	72
3.3.5	Local repository sizing	75
3.3.6	Factoring ratio considerations	77
3.3.7	Storage sizing	80
3.4	Replication considerations	89
3.4.1	Supported replication configurations	89
3.4.2	Virtual shelf	90
3.4.3	Replication parallelism scheme	91
3.4.4	Initial synchronization considerations	91
3.4.5	Bandwidth sizing and requirements	92
3.4.6	Network bandwidth sizing tips	93
3.4.7	Bandwidth validation tool	95
3.4.8	Repository sizing for replication with performance	100
3.5	Planning for OpenStorage	102
3.5.1	Replication for OST	102

3.5.2	Deployment planning guidelines example. . . . .	105
3.5.3	Replication policy . . . . .	108
3.6	Choosing a replication mode of operation: Time frame versus continuous. . . . .	109
3.6.1	Operation . . . . .	109
3.6.2	Best practices for choosing the mode of operation. . . . .	112
3.6.3	Remote repository. . . . .	112
3.7	Tips for using the visibility switch control feature . . . . .	113
3.8	Planning for OST. . . . .	115
3.9	Planning for cartridges . . . . .	115
3.9.1	Capacity management in the traditional tape library paradigm. . . . .	116
3.9.2	ProtectTIER systems versus traditional tape libraries . . . . .	116
3.9.3	How the TS7600 with ProtectTIER manages changes in nominal capacity . . . . .	117
3.9.4	Managing capacity fluctuations. . . . .	117
3.9.5	Capacity management implications for the TS7650, TS7610, or TS7650G . . . . .	118
3.9.6	Capacity management implications: Initialization phase. . . . .	118
3.9.7	Management of IBM System Storage TS7600 with ProtectTIER. . . . .	118
3.9.8	Capacity management implications: Adding new data sets to an existing IBM System Storage TS7600 with ProtectTIER. . . . .	119
3.9.9	Space reclamation and steady state. . . . .	119
3.9.10	Summary of TS7600 with ProtectTIER capacity management . . . . .	119
<b>Chapter 4.</b>	<b>Hardware planning for IBM System Storage ProtectTIER . . . . .</b>	<b>121</b>
4.1	General overview of the TS7610, TS7650, and TS7650G . . . . .	122
4.2	Hardware and software components for 3959-SM1, 3958-AP1, and 3958-DD4 . . . . .	124
4.2.1	The 3958-SM1 server features . . . . .	124
4.2.2	3959-SM1 server characteristics. . . . .	125
4.2.3	3958-AP1 server features. . . . .	125
4.2.4	3958-AP1 server characteristics. . . . .	126
4.2.5	3958-DD4 server features. . . . .	127
4.2.6	3958-DD4 server characteristics. . . . .	129
4.2.7	TS3000 System Console . . . . .	131
4.3	3959-SM1, 3958-AP1, and 3958-DD4 feature codes . . . . .	131
4.3.1	Features codes for 3959-SM1 . . . . .	131
4.3.2	Features codes for 3958-AP1 server . . . . .	132
4.3.3	Feature codes for 3958-DD4 server . . . . .	135
4.4	IBM System Storage TS7600 with ProtectTIER software . . . . .	138
4.4.1	5639-XXB ProtectTIER Enterprise Edition (EE) V2.5 Base Software . . . . .	138
4.4.2	5639-XXP ProtectTIER Appliance Edition (AE) V2.5 Software . . . . .	139
4.4.3	ProtectTIER Manager V2.5 console software . . . . .	140
4.5	Feature codes for Red Hat Linux . . . . .	140
4.5.1	ProtectTIER Enterprise Edition V2.5 . . . . .	140
4.6	3958-DD4 server configuration options. . . . .	141
4.6.1	Single node configuration . . . . .	141
4.6.2	Two-node cluster configuration. . . . .	147
4.6.3	Host attachment ports configuration for VTL and OST. . . . .	154
4.7	Usage considerations . . . . .	156
4.7.1	Virtual tape libraries and drives. . . . .	156
4.7.2	Fibre Channel ports and host assignment considerations . . . . .	156
4.7.3	Firewall environments: Ports assignments in ProtectTIER Replication Manager . . . . .	157
4.8	Installation planning . . . . .	158
4.8.1	Installation worksheets . . . . .	158
4.8.2	Supported backup server operating environments . . . . .	158
4.8.3	Planning the ProtectTIER installation. . . . .	158

4.8.4	Installation tasks . . . . .	160
4.8.5	Host attachment considerations . . . . .	162
4.8.6	SAN configuration . . . . .	162
<b>Part 3.</b>	<b>Implementing and administering the IBM System Storage TS7650G and TS7650 servers . .</b>	<b>169</b>
<b>Chapter 5.</b>	<b>IBM System Storage TS7600 with ProtecTIER initial setup . . . . .</b>	<b>171</b>
5.1	Enabling ProtecTIER SNMP support . . . . .	172
5.1.1	Defining the IP address . . . . .	172
5.1.2	IBM MIB definition file . . . . .	173
5.1.3	SNMP compatibility . . . . .	173
5.2	Installing ProtecTIER Manager . . . . .	174
5.2.1	Prerequisites . . . . .	175
5.2.2	Installing on Windows . . . . .	175
5.2.3	Installing ProtecTIER Manager on Linux . . . . .	179
5.3	Getting started . . . . .	190
5.3.1	Adding nodes to ProtecTIER Manager . . . . .	190
5.3.2	Logging in and out . . . . .	192
5.3.3	TS7610 start message . . . . .	194
5.3.4	Saving and printing data . . . . .	195
5.3.5	Refreshing ProtecTIER Manager . . . . .	196
5.3.6	Renaming the system . . . . .	198
5.4	Initial setup for the TS7610 model . . . . .	199
5.4.1	TS7610 configuration wizard setup . . . . .	200
5.5	Creating a ProtecTIER repository for TS7650G . . . . .	208
5.5.1	Creating file systems . . . . .	209
5.5.2	fsCreate parameters . . . . .	215
5.5.3	Creating the repository . . . . .	216
5.5.4	Renaming the repository . . . . .	225
5.5.5	Deleting existing repository and file systems . . . . .	227
5.6	Setting up the virtual library and cartridges . . . . .	230
5.6.1	Creating libraries . . . . .	230
5.7	Working with OpenStorage using ProtecTIER . . . . .	240
5.7.1	The OpenStorage operating environment . . . . .	240
5.7.2	Installing the ProtecTIER storage appliance . . . . .	241
5.7.3	Configuring ProtecTIER to work with the OpenStorage environment . . . . .	241
5.7.4	Configuring a storage server . . . . .	242
5.7.5	Modifying the storage server credentials . . . . .	244
5.7.6	Deleting a storage server . . . . .	245
5.7.7	Configuring a logical storage unit . . . . .	246
5.7.8	Adding an logical storage unit . . . . .	247
5.7.9	Modifying an logical storage unit configuration . . . . .	249
5.7.10	Managing the logical storage unit configuration . . . . .	250
5.8	Setting up native replication . . . . .	252
5.8.1	Replication throughput control . . . . .	252
5.8.2	ProtecTIER Replication Manager . . . . .	253
5.8.3	Adding replication to an existing production system . . . . .	254
5.8.4	Planning for a new installation . . . . .	256
5.8.5	Activating the ProtecTIER Replication Manager . . . . .	256
5.8.6	Deactivating the ProtecTIER Replication Manager . . . . .	257
5.8.7	Adding the repository to the grid . . . . .	259
5.8.8	Setting up the replication policies . . . . .	264
5.8.9	Enabling and disabling a policy . . . . .	272



5.8.10	Running a policy . . . . .	272
5.9	Setup replication on OpenStorage systems . . . . .	273
<b>Chapter 6. Host implementation for virtual tape libraries . . . . .</b>		<b>275</b>
6.1	Connecting hosts to ProtecTIER systems . . . . .	276
6.1.1	What bandwidth you need for ProtecTIER . . . . .	276
6.1.2	Multiple paths to tape drives . . . . .	277
6.1.3	Tape and disk on the same HBA . . . . .	279
6.1.4	SAN zoning . . . . .	281
6.1.5	LUN Masking for VTL systems . . . . .	283
6.1.6	Deduplication rate consideration . . . . .	296
6.1.7	Path failover . . . . .	296
6.2	Installing and configuring the device driver in OS . . . . .	297
6.2.1	Getting device drivers . . . . .	297
6.2.2	Installing IBM tape device drivers for AIX . . . . .	298
6.2.3	Installing IBM tape device drivers for Windows 2003 and 2008 . . . . .	309
6.2.4	Installing the IBM lin_tape driver for Red Hat Linux . . . . .	314
6.2.5	Open source device driver: lin_tape . . . . .	327
6.3	Setting up ProtecTIER on IBM i . . . . .	330
6.3.1	Prerequisites and test environment . . . . .	330
6.3.2	Interoperability . . . . .	331
6.3.3	Connecting the ProtecTIER FC ports to IBM i Fibre Channel adapters . . . . .	331
6.3.4	Creating a VTL for IBM i . . . . .	334
6.3.5	Managing the virtual devices in IBM i . . . . .	339
<b>Chapter 7. Backup and restore applications . . . . .</b>		<b>343</b>
7.1	Considerations for all backup servers . . . . .	344
7.1.1	General considerations . . . . .	344
7.1.2	Additional factors . . . . .	347
7.1.3	Assessing cartridge status and synchronizing with the catalog . . . . .	347
7.2	IBM Tivoli Storage Manager . . . . .	350
7.2.1	ProtecTIER IP replication in the Tivoli Storage Manager environment . . . . .	352
7.2.2	Implementing a virtual tape library . . . . .	352
7.2.3	The ProtecTIER virtual tape library definition . . . . .	353
7.2.4	Defining the virtual tape library to AIX with IBM Tivoli Storage Manager . . . . .	355
7.2.5	ProtecTIER and Tivoli Storage Manager attachment on Windows 2008 . . . . .	366
7.2.6	Tivoli Storage Manager database replication status . . . . .	377
7.2.7	Reclamation considerations . . . . .	380
7.2.8	Summary . . . . .	381
7.3	Determining what is available for restore at the DR site . . . . .	382
7.3.1	Which database copy at the remote site is valid . . . . .	382
7.3.2	Determining which cartridges at the remote are valid for restore . . . . .	382
7.3.3	Configuration changes . . . . .	383
7.3.4	LAN-free backup to disk with ProtecTIER . . . . .	385
7.3.5	Moving data between real tape libraries and ProtecTIER systems . . . . .	386
7.4	Symantec NetBackup . . . . .	389
7.4.1	Setting up NetBackup for backup and restore implementation . . . . .	390
7.4.2	Before you start . . . . .	390
7.4.3	Setting up NetBackup for disaster recovery . . . . .	390
7.4.4	Disaster recovery scenarios . . . . .	393
7.4.5	Determining what is available for restoration at the disaster recovery site . . . . .	394
7.4.6	NBU procedure to recover a master server from an existing DB copy . . . . .	395
7.4.7	Configuration changes . . . . .	396

7.5	EMC NetWorker	396
7.5.1	EMC Legato NetWorker and ProtecTIER Replication in a LAN/WAN	397
7.5.2	Implementation steps	397
7.5.3	Replicating NetWorker database (bootstrap) backups	398
7.5.4	Legato disaster recover procedures at the DR site	398
7.5.5	Determining which cartridges at the remote site are valid for restore	399
7.6	Backup, Recovery, and Media Services	400
7.6.1	Advantages of ProtecTIER replication for an IBM i data center	400
7.6.2	Setting up BRMS for backups to ProtecTIER	401
7.7	Data types	408
7.7.1	Oracle database data	408
<b>Chapter 8. IBM System Storage ProtecTIER native replication operation</b>		<b>409</b>
8.1	How replication works	410
8.1.1	Replication features	411
8.1.2	Typical deployment	413
8.1.3	ProtecTIER native replication Management Interface	415
8.2	Normal operation concepts	416
8.2.1	Replication	416
8.2.2	Replication data transfer	416
8.2.3	Visibility switch control	417
8.2.4	Single domain and multiple domain backup application environments	418
8.3	Many to one replication	420
8.3.1	ProtecTIER with Virtual Tape Library	420
8.3.2	ProtecTIER with OpenStorage	421
8.3.3	Replication management	424
<b>Chapter 9. IBM System Storage ProtecTIER with Symantec OpenStorage</b>		<b>433</b>
9.1	OpenStorage setup	434
9.1.1	OpenStorage elements	434
9.1.2	The OpenStorage operating environment	436
9.1.3	OpenStorage network configuration	436
9.2	Configuring OpenStorage with ProtecTIER Manager	438
9.2.1	Configuring the storage server	439
9.2.2	Configuring a logical storage unit	441
9.3	The ProtecTIER OpenStorage plug-in	446
9.3.1	Installing the OpenStorage plug-in on a NetBackup media server	446
9.3.2	Configuring the OpenStorage plug-in on a NetBackup media server	448
9.4	Configuring ProtecTIER OpenStorage on NetBackup	452
9.5	Replication settings for OpenStorage	466
9.5.1	Working with application groups in OpenStorage	466
9.6	Native replication with OpenStorage	471
9.7	The replication grid	472
<b>Chapter 10. Managing IBM System Storage ProtecTIER systems</b>		<b>473</b>
10.1	Managing nodes in ProtecTIER Manager	474
10.1.1	Adding a node and adding node subnetworks	474
10.1.2	Removing a node from ProtecTIER Manager	475
10.2	Managing repositories	476
10.2.1	Planning an expansion of the repository	476
10.2.2	Expanding existing file systems	476
10.2.3	Expanding the repository	481
10.2.4	Deleting the repository	485
10.3	Managing virtual libraries and cartridges	487

10.3.1	Editing library parameters . . . . .	487
10.3.2	Reassigning devices . . . . .	494
10.3.3	Adding cartridges . . . . .	500
10.3.4	Deleting cartridges . . . . .	504
10.3.5	Switching cartridges to read-only mode . . . . .	507
10.3.6	Renaming libraries . . . . .	509
10.3.7	Deleting libraries . . . . .	510
10.4	Viewing the alerts and event log windows . . . . .	513
10.4.1	Access alerts . . . . .	513
10.4.2	Access events . . . . .	515
10.4.3	Grid Manager log . . . . .	516
10.5	Wizard error messages . . . . .	517
10.6	Generating a service report . . . . .	517
10.6.1	Creating a problem report in Systems Management . . . . .	518
10.6.2	Creating a problem report in Grids Management . . . . .	521
10.7	Adding and removing cluster members using the ProtecTIER Manager . . . . .	524
10.7.1	Adding a cluster member . . . . .	524
10.7.2	Removing a cluster member . . . . .	528
10.8	Common maintenance tasks . . . . .	535
10.8.1	Starting and stopping the server . . . . .	535
10.8.2	Rebooting a node . . . . .	540
10.8.3	Disabling defragmentation . . . . .	542
10.8.4	Disabling compression . . . . .	543
10.8.5	Changing the HyperFactor mode . . . . .	544
10.8.6	Modifying the trace buffer . . . . .	544
10.8.7	Resetting drives . . . . .	547
10.8.8	Resetting robot . . . . .	547
10.8.9	Unloading and moving cartridges . . . . .	548
10.8.10	Modifying port attributes . . . . .	552
10.8.11	Checking the system . . . . .	555
10.8.12	Cartridge metadata verification . . . . .	557
10.8.13	Cartridge integrity verification . . . . .	561
10.9	Exchanging tape cartridges using the shelf . . . . .	565
10.10	Replication bandwidth throttling . . . . .	570
10.11	Automation of daily operations . . . . .	571
10.12	Updating the ProtecTIER Manager . . . . .	573
<b>Chapter 11</b>	<b>Native replication and disaster recovery . . . . .</b>	<b>575</b>
11.1	Moving to ProtecTIER DR mode . . . . .	576
11.2	Failback for all data at the DR site . . . . .	579
11.2.1	Splitting the failback according to priorities . . . . .	580
11.2.2	Creating a failback policy . . . . .	580
11.3	Recovery management . . . . .	586
11.3.1	Dirty bit attribute . . . . .	586
11.3.2	ProtecTIER command-line interface . . . . .	587
11.4	Disaster recovery with TS7650 OpenStorage and Symantec NetBackup . . . . .	590
11.4.1	NetBackup background . . . . .	590
11.4.2	Disaster recovery scenarios . . . . .	594
11.5	Completing failback and leaving DR mode . . . . .	596
11.6	Principality . . . . .	598
11.6.1	Taking over principality . . . . .	599
11.7	Repository replacement . . . . .	600
11.7.1	Replacing a destroyed VTL repository . . . . .	601

11.8 Restoring the ProtecTIER Replication Manager . . . . .	602
11.8.1 Restoring from default backup . . . . .	603
11.8.2 Restoring from file . . . . .	604
11.8.3 Restoring from IP address . . . . .	604
11.9 Returning to normal operations . . . . .	605
11.10 Flushing the replication backlog after a long link outage/unavailability . . . . .	606
11.10.1 Replication modes of operation: Visibility switch versus basic DR . . . . .	607
11.10.2 Use cases to demonstrate features of replication/DR operation . . . . .	609

**Chapter 12. Monitoring and reporting of the IBM System Storage TS7600 with ProtecTIER . . . . .**

<b>ProtecTIER . . . . .</b>	619
12.1 ProtecTIER Manager user management . . . . .	620
12.2 Monitoring ProtecTIER . . . . .	623
12.2.1 The status line . . . . .	624
12.2.2 The Navigation pane . . . . .	626
12.2.3 The Systems window . . . . .	628
12.2.4 The Nodes window . . . . .	633
12.2.5 The Repository window . . . . .	637
12.3 Monitoring the ProtecTIER virtual tape libraries service . . . . .	642
12.3.1 The Library window . . . . .	642
12.3.2 The General tab . . . . .	643
12.3.3 The Drives tab . . . . .	646
12.3.4 The Cartridges tab . . . . .	647
12.3.5 The Slots tab . . . . .	648
12.3.6 The Import/Exports tab . . . . .	649
12.3.7 Monitoring the shelf . . . . .	650
12.4 Reporting on ProtecTIER activity . . . . .	650
12.4.1 The ANALYZE_SESSIONS utility . . . . .	651
12.4.2 Long-term statistical data report . . . . .	654
12.5 Monitoring replication policies and activities . . . . .	657
12.5.1 Replication Policies window . . . . .	658
12.5.2 Replication Activities window . . . . .	660
12.5.3 Cartridge replication status information . . . . .	662
12.5.4 Cartridge verification after replication . . . . .	664
12.5.5 Replication long-term statistics . . . . .	665
12.5.6 Remote Cartridges report . . . . .	668
12.6 Network replication performance validation . . . . .	670
12.6.1 Interpreting the test results . . . . .	672
12.7 Managing and reporting TS7650 and TS7650G using the command-line interface . . . . .	674
12.7.1 Command-line interface . . . . .	674

**Part 4. Appendixes . . . . . 685**

<b>Appendix A. Installation and implementation checklists . . . . .</b>	687
Customer installation responsibilities . . . . .	688
Customer information work sheet . . . . .	688
Customer network settings work sheets . . . . .	690
TSSC network IP scheme . . . . .	690
IP address worksheet . . . . .	691
Customer IP addresses . . . . .	692
Customer and replication IP addresses . . . . .	692
Host names and DNS settings . . . . .	693
Replication settings worksheet . . . . .	695

<b>Appendix B. Western Telmatic Inc. Network Power Switch</b> .....	697
<b>Related publications</b> .....	703
IBM Redbooks .....	703
Other publications .....	703
Online resources .....	704
Help from IBM .....	704
<b>Index</b> .....	705



# Notices

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

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

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

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

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

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

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

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

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

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


## COPYRIGHT LICENSE:

This information contains sample application programs in source language, which 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](http://www.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 5L™	Informix®	System i®
AIX®	Lotus Notes®	System p®
DB2®	Lotus®	System Storage®
Diligent®	Notes®	System x®
Domino®	POWER5™	Tivoli®
DS4000®	ProtectTIER®	TotalStorage®
DS8000®	Redbooks®	VTF®
eServer™	Redbooks (logo)  ®	XIV®
HyperFactor®	RS/6000®	xSeries®
IBM®	Symphony™	

The following terms are trademarks of other companies:

Intel Xeon, Intel, Pentium, 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.

LTO, 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.

NetApp, and the NetApp logo are trademarks or registered trademarks of NetApp, Inc. in the U.S. and other countries.

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.

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

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



# Preface

The revolutionary and patented inline deduplication technology of IBM® ProtecTIER® enables you to easily harness the power of deduplication to retain data longer, protect data more efficiently and reliably, and save money by reducing energy, floor space, and maintenance requirements.

IBM ProtecTIER offers advanced native replication technology to automate the electronic replication of backup data between multiple data center locations for enhanced disaster recovery and business continuity. By eliminating the need to physically transport actual tape cartridges, data can be recovered faster and more reliably, enabling systems to get back online quicker in the event of a disaster or major system outage. ProtecTIER also lowers the total cost of ownership for backup and recovery by eliminating the costs associated with moving and storing physical tape cartridges.

This IBM Redbooks® publication covers these revolutionary technologies, as well as assist you with planning, installation, and administration. This publication is intended for system programmers, storage administrators, hardware and software planners, and other IT personnel involved in planning, implementing, and using the IBM deduplication solution, as well as anyone seeking detailed technical information about the IBM System Storage® TS7600 with ProtecTIER.

## 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, San Jose Center.

**Alex Osuna** is a project leader at the International Technical Support Organization, Tucson Center. He writes extensively about IBM System Storage N series and Tape. Before joining the ITSO 6 years ago, Alex worked in the IBM Tivoli® Western Region as a Systems Engineer. Alex has over 32 years in the I/T industry, focused mainly on storage. He has 30 years with IBM and holds certifications from IBM, Red Hat, Microsoft, and the Open Group.

**Eva Balogh** is an STG Lab Based Services consultant in the CEE/MEA region working at the IBM System Storage DS8000® Storage Server manufacturing in Vác, Hungary. She has been working there since 1997 as a contractor and joined IBM in 2000. In her current position, she holds ProtecTIER workshops, delivering implementation services for customers in the region for ProtecTIER and providing consultation in the same field. Her area of expertise is ProtecTIER for open systems. She has experience with the HP Data Protector 6.11 backup application. She also delivers encryption services for the DS8000 by using Tivoli Key Lifecycle Management (TKLM) and Secure Data Overwrite (data shredding) services on the DS8000 and ESS. She has been certified as a Project Management Professional by the Project Management Institute.

**Daesung Kim** is an IT Specialist from Korea working at the IBM Global Technology Services Organization. He provides second level support as a Product Field Engineer at the Technical Support Group. He joined IBM in 2006 and worked for many years as an SSR for an industry customer. In his current role as a second level support engineer in Seoul, Korea, he supports Tape Library, midrange storage, and NAS hardware products. He has experience with Linux, Windows, Sun Solaris, and IBM AIX®. This is his first IBM Redbooks publication.

**Kai Lanzmich** is a certified specialist for IBM RMSS products from Germany. He works as an IBM region designated specialist. He joined IBM 10 years ago and started working in different customer projects. Since 2004, he works exclusively on IBM tape products. From 2004 to 2006, he worked in the TSCC HW SSG in the IBM RMSS Storage Front office in Mainz. In 2007, he moved to Munich and started working as an RMSS region specialist for southern Germany. In his job role, he maintains installations, works with critical customer situations, and all other matters concerning IBM tape products in his region. He also works with the IBM RMSS product engineers. He has experience with AIX, Linux, and Windows, and backup software, such as Tivoli Storage Manager.

**Karen Orlando** is a Project Leader at the International Technical Support Organization, Tucson Arizona Center. Karen has over 25 years in the IT industry with extensive experience in open systems, product test, 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.

**Siew Chin Wong** is currently with the IBM Systems & Technology Group as a Storage Specialist handling disk systems, tape systems, storage virtualization and SAN infrastructure. She has more than 4 years of experience working with IBM customers in the ASEAN region, mainly as a technical consultant providing storage solutions, proofs of concept, and benchmarking. She has more than 10 years of professional experience in the IT industry, started as a Systems Engineer with a Systems Integrator (SI), has over 6 years of experience in open systems, specializing in UNIX and Linux, and has managed UNIX and Linux High Performance Computing Clusters for 4 years in a HPC environment. She joined IBM in 2006.

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

Dave R. Accordino  
**IBM Systems & Technology Group, Systems Hardware Development  
Tucson Arizona, USA**

Karen E. Orlando  
**International Technical Support Organization  
Tucson Arizona, USA**

Michael C. Barton  
**IBM Systems & Technology Group, Storage Platform  
Beaverton, USA**

Thanks to the authors of the previous editions of this book.

- ▶ Authors of the first edition, IBM System Storage TS7650, TS7650G, and TS7610, published in November 2008, were:
  - Babette Haeusser
  - Alessio Bagnaresi
  - Michael Bocian
  - Rik Foote
  - Abbe Woodcock
  
- ▶ Authors of the second edition, IBM System Storage TS7650, TS7650G, and TS7610, published in March 2010, were:
  - Alex Osuna, Lothar Weinert, Erwin Zwemmer, Xu X Yan and Reimar Pflieger
  
- ▶ Authors of the third edition, IBM System Storage TS7650, TS7650G, and TS7610, published in August 2010, were:
  - Alex Osuna, Erwin Zwemmer, Lothar Weinert, Reimar Pflieger, and Xu X Yan.

## 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](http://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](http://ibm.com/redbooks)

- ▶ Send your comments in an email to:

[redbooks@us.ibm.com](mailto: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>



# Summary of changes

This section describes the technical changes made in this edition of the book and in previous editions. This edition might also include minor corrections and editorial changes that are not identified.

Summary of Changes  
for SG24-7652-03  
for IBM System Storage TS7650, TS7650G, and TS7610  
as created or updated on September 6, 2011.

## August 2011, Fourth Edition

This revision reflects the addition, deletion, or modification of new and changed information described below.

### New information

- ▶ ProtecTIER V2.5 introduces support for Symantec OpenStorage (OST) interface.
- ▶ For the 2.5 release (new installation or upgrade), you can use the following states:
  - VTL application only
  - OST application only

### Changed information

- ▶ New DD4 and AP1 server (x3850 X5):
  - Additional memory.
  - Faster bus.
  - Faster CPUs.
  - Same number of PCI expansion slots.
- ▶ New network cards: FC Network
  - Emulex LPe12002 - 8 Gb FC HBA (front end)
  - Qlogic QLE2562 - 8 Gb FC HBA (back-end disk)
- ▶ The Maintenance Race Controller MRC now determines the rate of delete and defrag tasks according to the maximum limit and the calculated resources percentage.
  - The user can modify the maximum rate for delete and defrag.
  - Configuration is done online.

## August 2010, Third Edition

This revision included some changes to recommendations with regards to compression, which can be found in the following chapters:

- ▶ Chapter 1, “Concepts of data deduplication” on page 3
- ▶ Chapter 3, “Planning for deduplication and replication” on page 53
- ▶ Chapter 7, “Backup and restore applications” on page 343

## March 2010, Second Edition

This revision reflects the addition, deletion, or modification of new and changed information described below.

### **New information**

The V2.3 Enterprise and Appliance Edition ProtecTIER software provide the following functions:

- ▶ Network Replication Software
- ▶ 2nd NIC MES (DD1 only)
- ▶ ProtecTIER Replication Manager console software

The ProtecTIER Replication Manager console software must be installed on one of the 3958-AP, 3958-DD3, or 3958-DD1 servers in a replication pair. Alternatively, customers can submit an RPQ request to put the ProtecTIER Replication Manager console software on a separate piece of hardware.



# Part 1

## Introduction and architecture

In this part, we introduce the basic concepts of data replication and native replication and describe the architecture and components of the IBM solution for data deduplication and native replication.







# Concepts of data deduplication

In this chapter, we describe general data deduplication concepts and type, data deduplication methods, system components, and benefits of data deduplication.

## 1.1 Data deduplication

Data deduplication is a technology that is used to reduce the amount of space required to store data on disk. It is achieved by storing a single copy of data that is backed up repetitively. Data deduplication can provide greater data reduction than previous technologies, such as Lempel-Ziv (LZ) compression and differencing, which is used for differential backups.

Figure 1-1 illustrates the basic components of data deduplication for IBM System Storage 7600 ProtecTIER servers.

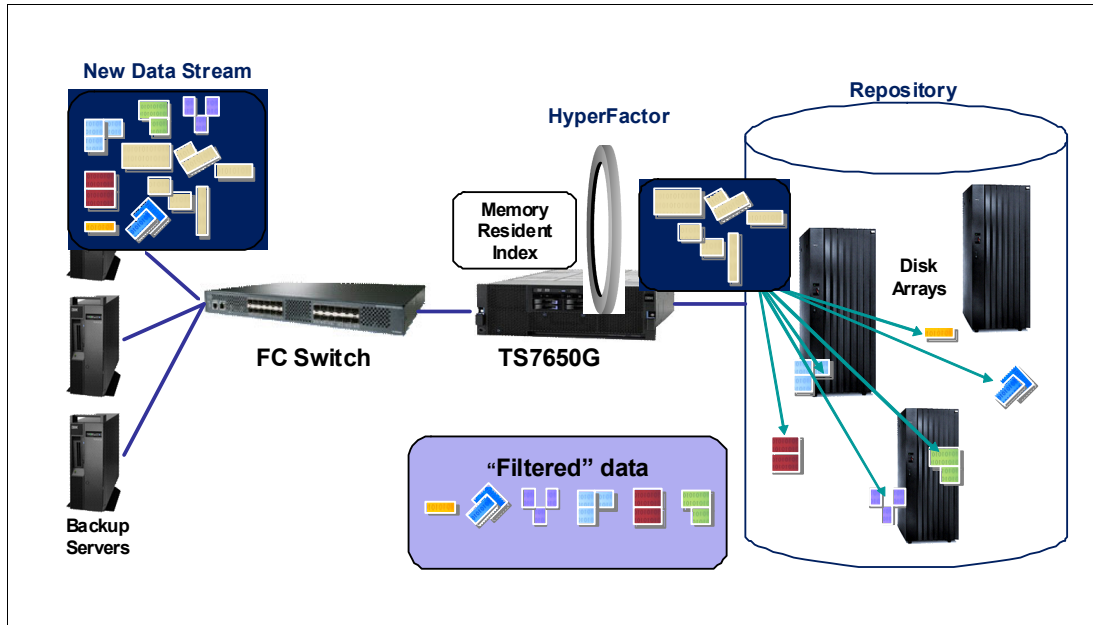


Figure 1-1 Basic concepts of data deduplication

With data deduplication, data is read by the data deduplication product while it looks for duplicate data. Different data deduplication products use different methods of breaking up the data into elements, but each product uses a technique (see 1.2, “Types of data deduplication” on page 5) to create a signature or identifier for each data element. Whether using inline or post processing data deduplication (see 1.3, “Data deduplication processing” on page 9), data element signature values are compared to identify duplicate data. After the duplicate data is identified, one copy of each element is retained, pointers are created for the duplicate items, and the duplicate items are not stored.

The effectiveness of data deduplication is dependent upon many variables, including the data rate of data change, the number of backups, and the data retention period. 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 to 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 to 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, let us assume that you take weekly full and daily differential incremental backups. Let us also assume that your data change rate for the full backups is 15% and your daily incrementals is 30%. After 30 days, your deduplication ratio might be around 6 to 1, but if you kept your backups up to 180 days, your deduplication ratio might have increased to 10 to 1.

In the examples above, and in the remainder of this book, we discuss the deduplication ratio as being the total backup data received divided by the amount of disk space used to store it.

Data deduplication can reduce your storage requirements, 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 like IBM Tivoli Storage Manager Incremental Forever, Oracle RMAN, or Light Speed, can reduce the deduplication ratio. Compressed, encrypted, or otherwise scrambled workloads typically do not benefit from deduplication.

## 1.2 Types of data deduplication

Many vendors offer products that perform deduplication. Various methods are used for deduplicating data. Three methods frequently used are:

- ▶ Hash based
- ▶ Content aware
- ▶ IBM HyperFactor®

### 1.2.1 Hash-based data deduplication

Hash-based data deduplication methods use a hashing algorithm to identify *chunks* of data. Commonly used algorithms are Secure Hash Algorithm 1 (SHA-1) and Message-Digest Algorithm 5 (MD5). When data is processed by a hashing algorithm, a hash is created that represents the data. A hash is a bit string (128 bits for MD5 and 160 bits for SHA-1) that represents the data processed. If you process the same data through the hashing algorithm multiple times, the same hash is created each time.

Here are some commonly used hash codes:

- ▶ MD5: 16-byte long hash
  - # echo “The Quick Brown Fox Jumps Over the Lazy Dog” | md5sum  
9d56076597de1aeb532727f7f681bcb0
  - # echo “The Quick Brown Fox Dumps Over the Lazy Dog” | md5sum  
5800fccb352352308b02d442170b039d
- ▶ SHA-1: 20-byte long hash
  - # echo “The Quick Brown Fox Jumps Over the Lazy Dog” | sha1sum  
F68f38ee07e310fd263c9c491273d81963fbff35
  - # echo “The Quick Brown Fox Dumps Over the Lazy Dog” | sha1sum  
d4e6aa9ab83076e8b8a21930cc1fb8b5e5ba2335

Hash-based deduplication breaks data into chunks, either fixed or variable length, depending on the product, and processes the chunk with the hashing algorithm to create a hash. If the hash already exists, the data is deemed to be a duplicate and is not stored. If the hash does not exist, then the data is stored and the hash index is updated with the new hash.

In Figure 1-2, data chunks A, B, C, D, and E are processed by the hash algorithm and create hashes  $A_h$ ,  $B_h$ ,  $C_h$ ,  $D_h$ , and  $E_h$ . For the purposes of this example, we assume that this is all new data. Later, chunks A, B, C, D, and F are processed. F generates a new hash,  $F_h$ . Because A, B, C, and D generated the same hash, the data is presumed to be the same data, so it is not stored again. Because F generates a new hash, the new hash and new data are stored.

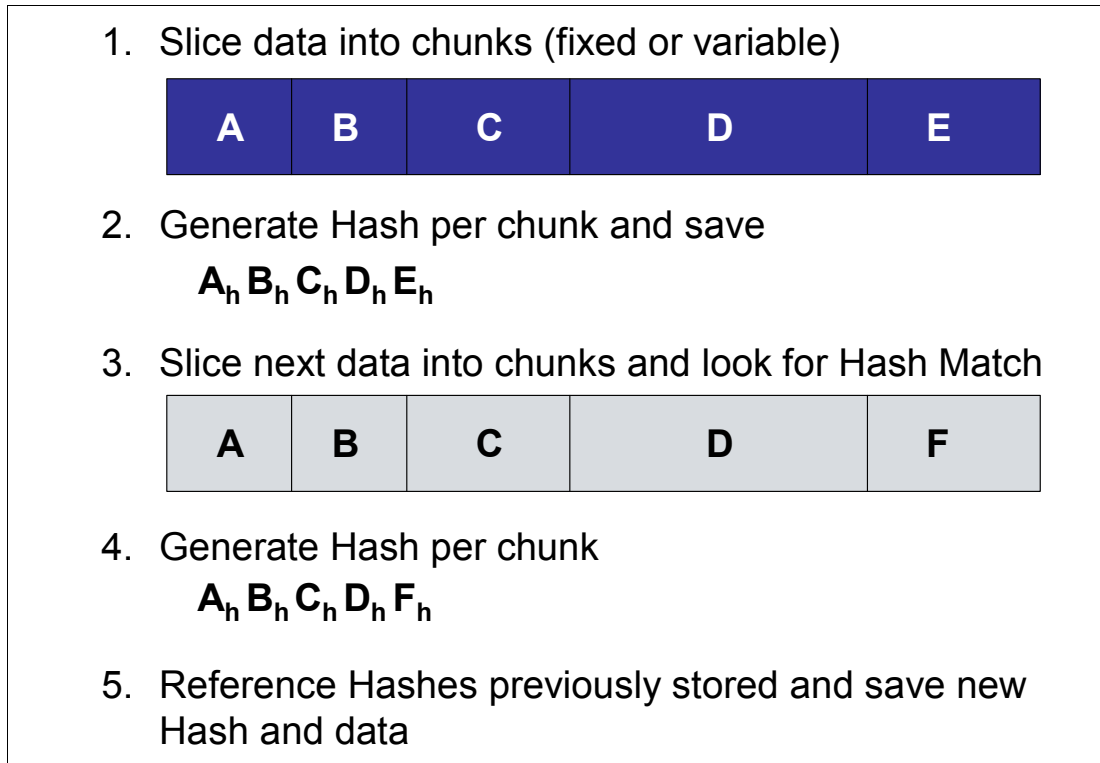


Figure 1-2 Hash-based deduplication

Hash-based deduplication must store all hashes in an index that can be large and might not fit in memory, and consequently must be stored on disk. Querying the index to identify hash matches can be time consuming, which can impact performance. The size of the index might also impact scalability, as the index space is required to increase. Assuming an 8 KB data chunk, processing 10 TB of data might require 1,250,000,000 accesses to an index.

There is some concern in the industry that two chunks of different data could create the same hash, causing a hash collision, and, furthermore, there is no way of determining that the data has been corrupted by a hash collision. With a hash collision, you could inadvertently lose data, as the deduplication process does not save new data because it assumes that because the hashes match, the data has already been stored. Opinions vary on the level of exposure to hash collisions. Products using hash-based deduplication can mitigate the potential problem by employing techniques such as processing the data with both the SHA-1 and MD5 algorithms for consistency or doing a byte comparison on data. When reviewing deduplication products, you should discuss this topic with the product vendor.

## 1.2.2 Content aware

Content aware deduplication methods are aware of the structure or common patterns of data used by applications. It assumes that the best candidate to deduplicate against is an object with the same properties, such as a file name. When a file match is found, a bit-by-bit comparison is performed to determine whether data has changed and saves the changed data.

In Figure 1-3, with content-aware deduplication, the system looks through the backup data to find the fully qualified names and then looks for previous versions of the same file. If a match is found, the system performs a bit-by-bit comparison and updates the reference points as needed.

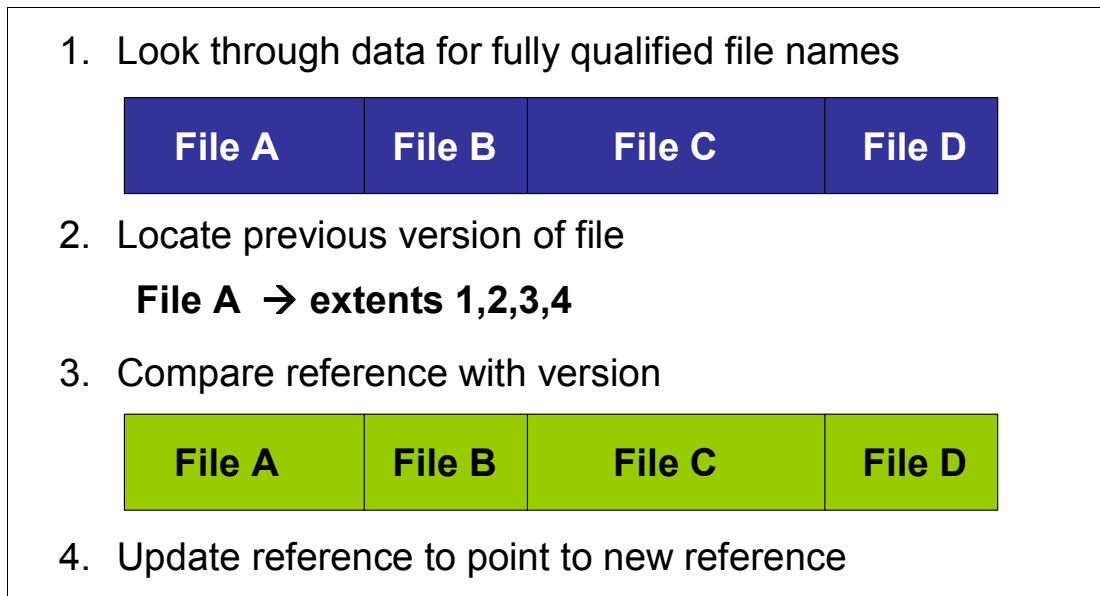


Figure 1-3 Content-aware data deduplication

A content-aware data deduplication product must know the data structure of the backup applications that it supports. It must also keep track of potentially millions of file names, which reduces the possibility of a Memory Resident Index for quick reference. If a backup application changes the data structure, the content-aware data deduplication product must be updated to reflect that change.

The index is file size dependent. An average file size of 1 MB would require 10,000,000 accesses to an index to process 10 TB of data.

## 1.2.3 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's 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 have previously been stored by ProtecTIER. Over time, HyperFactor can increase the usable capacity of a given 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 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 ProtecTIER native replication technology, HyperFactor finds any similar data elements that have already been stored. This search is extremely quick using 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 is able to surpass the reduction ratios attainable by any other data reduction method. HyperFactor can reduce any duplicate data, regardless of its location or how recently it was stored. Unlike hash-based techniques, HyperFactor finds duplicate data without needing exact matches of chunks of data. When new data is received, HyperFactor checks to see whether similar data has already been stored. If similar data has already been stored, then only the difference between the new data and previously stored data must be retained. Not only is this an effective technique of finding duplicate data, but it performs well.

In Figure 1-4, with HyperFactor deduplication, when new data is received, HyperFactor looks for data similarities and checks those similarities in the Memory Resident Index. When similarity matches are found, a binary differential is performed on the 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.

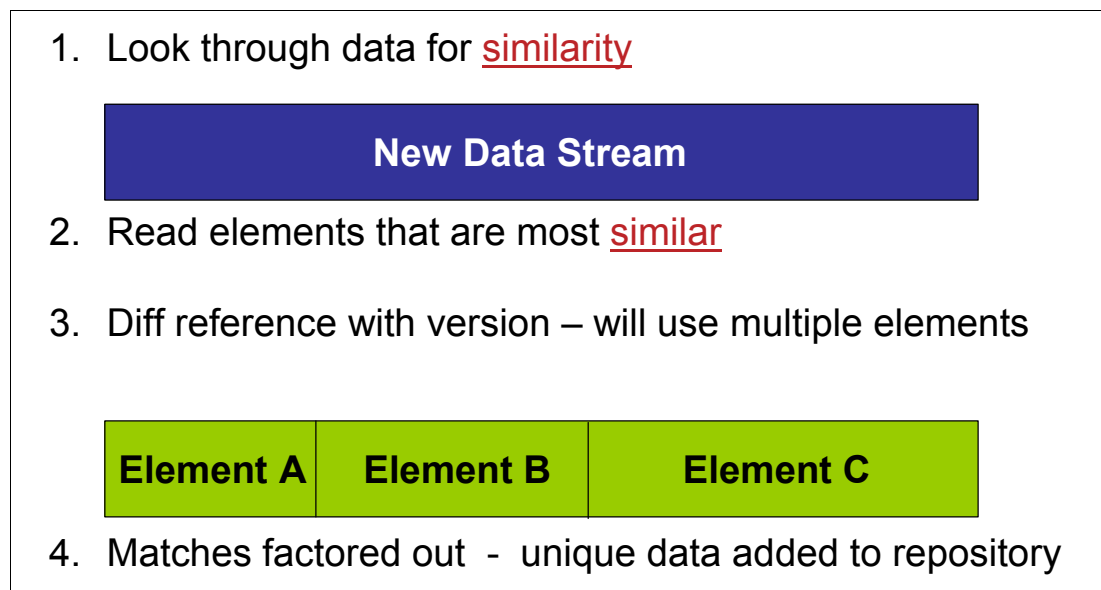


Figure 1-4 HyperFactor data deduplication

HyperFactor data deduplication uses a 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 to 1, you could store 12 PB on that 1 PB of disk array. With the Memory Resident Index, HyperFactor can identify potentially duplicate data quickly for large amounts of data and does this 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 faster than the write rate to the system, because there is no risk of fragmentation, and no access to the index or heavy computation is required during a restore activity. It just requires you to open metadata files and fetch the data according to the pointers that they contain.

## 1.3 Data deduplication processing

Data deduplication can either be performed while the data is being backed up to the storage media (inline) or after the data has been written to the storage media (post processing).

### 1.3.1 Inline method

Inline data deduplication is not dependent on the type of data deduplication used. An advantage of inline data deduplication is that the data is only processed once and there is no additional processing after the backup window. Inline data deduplication requires less disk storage because the native data is not stored prior to data deduplication. Depending on the implementation, a disadvantage of inline data deduplication is that the data deduplication processing could slow down the backup data stream. Algorithms used for data deduplication can be CPU intensive and data deduplication might require additional read or write access if the index is disk based.

### 1.3.2 Post-processing method

With a post-processing data deduplication method, data is backed up first, and after the backup window has completed, the data deduplication is performed. The advantage of this method is that the original backup stream is not slowed and your backup window is not impacted.

There are disadvantages to the post-processing method:

- ▶ Increased input/output (I/O) to the storage device. Because the data is written during the backup, reads to identify duplicate data and the pointers must be updated if there is duplicate data. Overall, the data deduplication cycle will likely be longer than if performed inline.
- ▶ More disk is required than with an inline method, as all the data must be stored prior to deduplication.

If the post-processing data deduplication period extends too much, you could encounter a situation where your data deduplication process has not completed before the start of your next backup window.

## 1.4 Components of a data deduplication system

Data deduplication systems' implementations vary. You can have a hash-based data deduplication system that uses the post-processing method and another hash-based data deduplication system that uses the inline method. Some systems might integrate all the components required for data deduplications and others might need to be integrated by you.

Regardless of the type and method of data deduplication or the packaging, a system has three required components. Those components are covered in the following three sections.

## 1.4.1 Server

Every data deduplication system must have a server, with an operating system, on which the data deduplication software runs. The data deduplication software might be integrated with a virtual tape library or similar software and run on a single server or a separate data deduplication server might be required. The server might be packaged and integrated with the data deduplication software or you might need to acquire the server and the operating system separately.

## 1.4.2 Data deduplication software

The data deduplication software performs the deduplication process. Available products might integrate the data deduplication software with a virtual tape library application or the data deduplication software might run separately.

## 1.4.3 Disk array

Data processed by data deduplication software is stored on a disk array. A vendor might package and integrate the disk array with the server and the data deduplication software or you might need to acquire the disk array separately.

**Note:** ProtecTIER does not require additional compression to be effective. ProtecTIER performs compression, by default, after the deduplication process. Do not allocate disk arrays for ProtecTIER VTLs that perform their own compression.

# 1.5 Benefits of data deduplication

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, as 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.

## 1.5.1 Reduction of storage requirements

With the amount of data that corporations are required to maintain for compliance with government regulations and for normal business operations, data deduplication can reduce the amount of disk storage 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 data required might have been shipped offsite.

## 1.5.2 Reduction of environmental costs

If data deduplication reduces your disk storage requirements, the environmental costs for running and cooling the disk storage are also reduced.





# IBM System Storage ProtecTIER architecture

In this chapter, we introduce the IBM System Storage TS7600 ProtecTIER Deduplication Solution and discuss the following topics:

- ▶ IBM System Storage ProtecTIER General Overview
- ▶ IBM System Storage ProtecTIER Manager and Replication Manager
- ▶ IBM System Storage TS7650G ProtecTIER Deduplication Gateway technology and product requirements
- ▶ IBM System Storage TS7650 ProtecTIER Deduplication Appliance technology and product requirements
- ▶ IBM System Storage TS7610 ProtecTIER Deduplication SMB Appliance technology and product requirements
- ▶ IBM System Storage ProtecTIER V2.5 Software with Data Deduplication and HyperFactor

## 2.1 General overview of ProtecTIER

This section describes the general hardware and software components of a ProtecTIER system.

### 2.1.1 Types of ProtecTIER models

There are three different models that you can choose from to implement the ProtecTIER deduplication system:

- ▶ TS7650G ProtecTIER Deduplication Gateway
  - It includes only the server installed with the ProtecTIER code.
  - This is the most scalable model of the three. It can start from a small disk subsystem and grow up to 1 PB of repository size.
  - You can have single node Gateway or two-node cluster Gateways.
  - You can choose to attach to a new external disk subsystems or existing disk subsystems in your environment. For the list of supported disk systems, refer to the document found at:

[ftp://service.boulder.ibm.com/storage/tape/ts7650\\_support\\_matrix.pdf](ftp://service.boulder.ibm.com/storage/tape/ts7650_support_matrix.pdf)

For more information, refer to 2.3, “TS7650G ProtecTIER Deduplication Gateway” on page 17.

- ▶ TS7650 ProtecTIER Deduplication Appliance
  - It includes the server installed with the ProtecTIER code and disk subsystems for the repository.
  - The size of the repository of the appliance is 7 TB, 18 TB, or 36 TB.
  - You can have single node Appliance or two-node cluster Appliances in a 36 TB capacity.

For more information, refer to 2.4, “TS7650 ProtecTIER Deduplication Appliance” on page 26.

- ▶ TS7610 ProtecTIER Deduplication SMB Appliance
  - A customer instatiable entry level package that includes the appliance installed with the ProtecTIER code and internal disks for the repository.
  - The size of the repository of the SMB appliance is 4.0 TB or 5.4 TB.

For more information, refer to 2.5, “TS7610 ProtecTIER Deduplication SMB Appliance” on page 29.

Table 2-1 gives a comparison of the maximum limits of the different ProtecTIER models.

Table 2-1 Comparison of maximum limits of the ProtecTIER models

Models	Enterprise Gateway Edition		Appliance Edition		SMB Appliance Edition	
	Repository	Library	Repository	Library	Repository	Library
Maximum repository size	1 PB	1 to 16	36 TB	1 to 12	5.4 TB	1 to 4
Virtual libraries per repository	1 to 16	N/A	1 to 12	N/A	1 to 4	N/A

Models	Enterprise Gateway Edition		Appliance Edition		SMB Appliance Edition	
Virtual Drives (single node)	256	256	256	256	64	64
Virtual Drives (two-node cluster)	512	512	512	512	N/A	N/A
Virtual Cartridges	512 K	62 K	128 K	62 K	8 K	8 K
Virtual Slots	512 K	62 K	128 K	62 K	64 K	64 K
Virtual Import/Export Slots	4096	1022	4096	1022	1022	1022

## 2.1.2 ProtecTIER service modes

ProtecTIER is a disk-based data storage system. It uses data deduplication technology to store data to disk arrays. All the models of the ProtecTIER deduplication system can be configured to operate in one of the two modes:

- ▶ ProtecTIER Virtual Tape Library (VTL):

The ProtecTIER VTL service emulates traditional tape libraries, which enables you to make a transition from a traditional tape backup to disk backup without having to replace your entire backup environment. For more information, refer to 2.6, “ProtecTIER virtual tape library” on page 32.

- ▶ Symantec OpenStorage (OST):

OST implements a storage server emulation that can be integrated with NetBackup to provide the means for backup-to-disk without having to emulate traditional tape libraries. For more information, refer to 2.7, “ProtecTIER OpenStorage” on page 36.

**Note:** A single ProtecTIER system cannot be configured to run both modes at the same time.

## 2.1.3 ProtecTIER Manager

The ProtecTIER Manager (PT Manager) application is a graphical user interface (GUI) you use to view, configure, manage, and troubleshoot the operation of the IBM TS7600 ProtecTIER family of products. It can be installed on a Windows or Linux based PC on your Ethernet network. The PT Manager does not need to be located near the ProtecTIER system, as long as Ethernet connectivity to the ProtecTIER system is allowed.

For more information about PT Manager, refer to 2.10, “ProtecTIER Manager” on page 46.

## 2.1.4 ProtecTIER Replication Manager

Sometimes referred to as *Grid Manager*, the ProtecTIER Replication Manager should be able to recognize all the members of the entire network that the ProtecTIER Replication Manager handles on both replication subnets. The ProtecTIER Replication Manager is enabled separately from the ProtecTIER Manager on the customer's ProtecTIER server. The ProtecTIER Replication Manager manages the configuration of multiple replication grids in an organization. An agent on every node in each ProtecTIER server interacts with the server and maintains a table of its grid members.

Native replication lets you replicate data objects between ProtecTIER repositories. In order for a logical set of repositories to replicate from one to another, you must create a replication grid. The replication grid is remotely created and managed by the Replication Manager.

Each ProtecTIER Replication Manager has a unique identity. A repository, after it has joined a replication manager, cannot join a replication grid managed by a different replication manager, even if it has left the grid. This setup prevents data collision.

The ProtecTIER Replication Manager is a server that remotely manages the replication grids within an organization. The ProtecTIER Manager connects to the ProtecTIER Replication Manager using the IP address of the ProtecTIER Replication Manager server. The ProtecTIER Replication Manager can be installed on a dedicated host (which requires a special request for support from IBM), or on a ProtecTIER node.

If the ProtecTIER Replication Manager is installed on a ProtecTIER node, it can manage up to one single grid with 24 repositories. If the ProtecTIER Replication Manager is installed on a dedicated server, it can manage up to 64 grids with 256 repositories in each grid.

**Note:** Consider only using one grid.

The PT Replication Manager's responsibility is to:

- ▶ Manage the repositories in the replication grid
- ▶ Maintain the replication IP addresses of all repositories
- ▶ Update the repositories leaving and joining the grid
- ▶ Provide high-level monitoring and statistics of traffic in the replication grid

For more information about replication, refer to Chapter 11, "Native replication and disaster recovery" on page 575.

## 2.2 Terms and definitions

The following terms are used in this and later chapters:

<b>Front end</b>	The connection between the ProtecTIER system and the backup server is referred to as a front-end connection.
<b>Back end</b>	The connection between the ProtecTIER system and the disk array is referred to as a back-end connection.
<b>Node</b>	A single ProtecTIER server (TS7650G, TS7650, or TS7610) is viewed as a node from the ProtecTIER Manager software. You can have a single node or two-node clusters.
<b>Metadata</b>	Metadata is the data used to keep track of the data about your backup data, including where it is stored on the disk.
<b>User data</b>	User data is the backup files and data sets stored on the virtual tape library. It is the data that you are storing on disk.

<b>Metadata file system</b>	The metadata file system stores all aspects of the data that is backed up and cataloged, but not the data itself, whether it requires new disk space or not. It is critical that the performance of the metadata file system be optimal. Therefore, in general, Strongly consider using RAID 10 RAID groups (3+3 up to 8+8 disks) on Fibre Channel disk for the metadata file systems.
<b>User data file system</b>	The user data file system stores the actual data that is backed up or referenced by new generations of the data. The user data file system is stored on a RAID 5 configuration.
<b>Nominal capacity</b>	The amount of user data that ProtecTIER is managing.
<b>Physical capacity</b>	The physical capacity used in the array.
<b>Factoring ratio</b>	The factoring ratio refers to the ratio of nominal capacity to physical capacity. 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 to 1.
<b>Repository</b>	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: <ul style="list-style-type: none"> <li>– Metadata</li> <li>– User data</li> </ul>
<b>Disk array</b>	The disk array attaches to the IBM System Storage TS7600 family through back-end connections and holds the repository or cache of factored backup data.
<b>Data retention period</b>	The period of time (usually measured in days) that defines how long customers keep their disk-based backups online. This period of time typically ranges from a period of 30 to 90 days, but can be less or longer.
<b>Change rate</b>	The percentage of change in data from one backup set to the next. The daily change rate is the percentage of change from one day's backup cycle to the next. For example, if the daily change rate is 10%, it means that only 10% of the backed-up data changes from one day to the next.
<b>Disaster recovery</b>	Disaster recovery (DR) is the process of recovering production site data from a remote location. It includes a way to indicate to a remote repository that the production site has gone down and the actual data recovery process.
<b>Disaster recovery test</b>	A simulation of the disaster recovery process.
<b>Failover</b>	The process of failover is undertaken when continued operations at the source location is no longer possible. A disaster is declared.
<b>Failback</b>	A process that is initiated from the remote site when the source site is now able to continue production operations and therefore back up processes. The process ensures that the paired repositories are re-synchronized using the least amount of bandwidth and maintaining the most recent copies of backups. After the failback process, the operational norms prior to the execution of a DR resume.

<b>One-time policy</b>	A special policy that can only be executed manually and is deleted after it is completed.
<b>Principality/ownership</b>	An attribute indicating the repository in which an individual cartridge can be updated or written on by a backup application. A cartridge at its principal repository can be read/write (R/W) or read-only (R/O). At other sites it is R/O. A cartridge can have principality/ownership turned on for only one site.
<b>Replication</b>	A process that transfers logical objects like cartridges from one ProtecTIER repository to another one. The replication function allows ProtecTIER deployment to be distributed across sites. Each site has a single or clustered ProtecTIER environment. The ProtecTIER server that is a part of the replication grid has two dedicated replication ports. The port number depends on the ProtecTIER model. Replication ports are connected to the customer's WAN and are configured on two subnets as the default.
<b>Replication grid</b>	A set of repositories that share a common ID and can potentially transmit and receive logical objects through replication. A replication grid defines a set of ProtecTIER repositories and actions between them and is configured using the ProtecTIER Replication Manager. The ProtecTIER Replication Manager is a software component that is enabled on a ProtecTIER server.
<b>Replication grid ID</b>	A number from 0 to 63 that identifies a replication grid within an organization.
<b>Replication grid member</b>	A repository that is a member in a replication grid.
<b>Hub</b>	A replication and backup target. It may receive replication from up to 12 spokes in a many-to-one replication system setup environment. In a many-to-many replication system setup environment, up to 4 hubs repositories can replicate to each other; in this case, the repositories are hub and spoke at same time. In a OST mesh grid, it may be configured to replicate to another system and act as a source concurrently.
<b>Spoke</b>	A backup source that can only replicate to a single hub in a many-to-one replication system setup environment. Spokes are not supported in a many-to-many replication group.
<b>Replication pairs</b>	Two repositories within a replication grid that replicate from one to another.
<b>Replication policy</b>	A policy made up of rules that define a set of objects from a source repository to be replicated to a target repository.
<b>Replication time frame</b>	A scheduled period of time for replication to take place for all policies.
<b>Repository unique ID (RID)</b>	A number that uniquely identifies the repository. The RID is created from the replication grid ID and the repository internal ID in the grid.
<b>Shelf</b>	A container of virtual tape library (VTL) cartridges within a ProtecTIER repository. This is analogous to a shelf or a rack where physical tapes are kept when outside an automated tape library.

<b>TS7600</b>	When used alone, this term signifies the IBM family of virtualization solutions that operate on the ProtecTIER platform.
<b>Visibility</b>	Whether an application backup server can see or has visibility to a VTL cartridge. This construct is unique to VTL technology and it is the means by which ProtecTIER ensures that a tape is only in one place at any given time. This is analogous to a physical piece of media.
<b>Visibility switching</b>	The automated process that transfers the visibility of a VTL cartridge from its master to its replica and vice versa (source or local to destination or remote). The visibility switching process is triggered by moving a cartridge to the source library import/export (I/E) slot. The cartridge will become invisible to the source backup application and appear at the destination library I/E slot. To gain visibility back at the source library, you need to move the cartridge to the I/E slot, where it will then be placed on the shelf at the destination location and placed in the I/E slot at the source location. The cartridge will then disappear from the destination library and reappear at the source I/E slot.

## 2.3 TS7650G ProtecTIER Deduplication Gateway

The IBM System Storage TS7650G is a preconfigured virtualization solution of IBM systems, and the IBM revolutionary ProtecTIER data deduplication software designed to improve backup and recovery operations. The solution is available in a single node or two-node cluster configurations designed to meet the disk-based data protection needs of a wide variety of IT environments' enterprise data centers organizations.

The IBM System Storage TS7650G, composed of 3958-DD4 hardware combined with IBM System Storage ProtecTIER Enterprise Edition V2.5 software, offers the acquired data deduplication technology.

**Note:** ProtecTIER software is ordered with the TS7650G order, but it is shipped separately.

### 2.3.1 TS7650G Gateway models

There are a few models of the IBM virtualization solution from the TS7600 family that does not include a disk storage repository, allowing the customer to choose from a variety of storage options. The TS7650G consists of the 3958-DD4, 3958-DD3 or 3958-DD1 models, which are three types of servers used as the ProtecTIER Gateway. The 3958-DD1 and 3958-DD3 are the oldest models.

► **3958-DD4:**

This is the latest higher performance server available since November 2010. This server is based on the IBM System x3850 Type 7145 server, and when ordered as the ProtecTIER TS7650G, the machine type and model are 3958-DD4. Use this machine type and model for service purposes.

▶ 3958-DD3:

This is the second generation server 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. This model is no longer available for ordering.

▶ 3958-DD1:

This is the original server introduced in 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. Use this machine type and model for service purposes. This model is no longer available for ordering.

### 2.3.2 TS7650G ProtecTIER Deduplication Gateway (3958-DD4)

The TS7650G, composed of the 3958-DD4 hardware combined with IBM System Storage ProtecTIER Enterprise Edition software, is designed to address the data protection needs of enterprise data centers. The solution offers high performance, high capacity, scalability, and a choice of disk-based targets for backup and archive data. The TS7650G offers:

- ▶ Inline data deduplication powered by HyperFactor technology.
- ▶ A multicore virtualization and deduplication engine.
- ▶ Clustering support for higher performance and high-availability.
- ▶ Fibre Channel ports for host and server connectivity.
- ▶ Performance up to 2000 MBps or more sustained inline deduplication (two-node clusters, depending on the configuration of attached disks).
- ▶ Virtual tape emulation of up to 16 virtual tape libraries per single node or two-node cluster configuration and up to 512 virtual tape drives per two-node cluster or 256 virtual tape drives per single node.
- ▶ Emulation of the IBM TS3500 Tape Library with IBM LTO Ultrium 3 tape drives.
- ▶ Emulation of the Quantum P3000 Tape Library with DLT7000 tape drives or IBM LTO Ultrium 2 tape drives.
- ▶ Emulation of the IBM DTC VTF® 0100 virtual tape libraries with DLT7000 tape drives or IBM LTO Ultrium 2 tape drives.
- ▶ Scales up to 1 PB of physical storage over 25 PB of user data (depending on your deduplication ratio).
- ▶ The system console is a TS3000 System Console (TSSC).

The TS7650G is an enterprise-class data protection platform designed to quickly and safely protect business information while reducing the amount of space required to store it.

Deploying the TS7650G can help organizations more efficiently protect their corporate data on disk-based storage while helping them manage the exponential growth of new data through reduction and elimination of duplicate data in their backups.

The use of disk cache as a deduplicated storage pool with existing backup applications offers a potentially lower total cost of ownership than conventional disk. To facilitate backup applications that are designed for use with tape storage, the TS7650G emulates a traditional tape library unit, so deployment does not force any changes to a customer's existing data protection processes.



Figure 2-1 shows the components of the TS7650G ProtecTIER Deduplication Gateway DD4. Table 2-2 lists the slot assignment, ports and connections. Table 2-3 lists the new capabilities of the 3850 X5 Gateway. More information about these components follows in the next sections.




Figure 2-1 TS7650G ProtecTIER Deduplication Gateway DD4 rear view: Generic

Table 2-2 3958 DD4 server rear view: Slot assignments, ports, and connections

Slot, port or connection	VTL use	OpenStorage use
Slot 1	Emulex HBA for front-end attachment	Blank
Slot 2	Emulex HBA for front-end attachment	Blank
Slot 3	Blank	Intel Pro Quad-port gigabit Ethernet adapter (Port 1 = Eth8, Port 2 = Eth9, Port 3 = Eth10, Port 4 = Eth11)
Slot 4	Blank	Intel Pro Quad-port gigabit Ethernet adapter (Port 1 = Eth4, Port 2 = Eth5, Port 3 = Eth6, Port 4 = Eth7)
Slot 5	Intel Pro Quad-port gigabit Ethernet adapter (Port 1 = Eth0, Port 2 = Eth1, Port 3 = Eth2, Port 4 = Eth3)	Intel Pro Quad-port gigabit Ethernet adapter (Port 1 = Eth0, Port 2 = Eth1, Port 3 = Eth2, Port 4 = Eth3)
Slot 6	Qlogic HBA for back-end attached disks	Qlogic HBA for back-end attached disks
Slot 7	Qlogic HBA for back-end attached disks	Qlogic HBA for back-end attached disks

Table 2-3 Capabilities of the new 3850 X5 Gateway

Component	3958-DD1	3958-DD3	3958-DD4
IBM System	x3850 M2	x3850 M2	x3850 X5
CPU	4x quad core (16 cores)	4x hex cores (24 cores)	4x octo cores (32 cores)
RAM	32 GB	32 GB	64 GB

**Note:** For existing customers who are using the 3958-DD1 or 3958-DD3 models, there is a new option to increase the memory of the ProtecTIER system from 32 GB to 64 GB. This upgrade is only necessary if you want to have a repository size greater than 512 TB with replication enabled and using more than 128 virtual drives per node.

In a clustered configuration, the upgrade procedure should allow for one node to remain online, providing customer data access while the peer node in the cluster undergoes the upgrade procedure.

You can upgrade from an existing 3958-DD3 single node to a clustered gateway with a 3958-DD4 server. However, to support an upgrade of an existing 3958-DD1 single node to a clustered gateway, you need special approval from IBM (RPQ).

### 2.3.3 Disk array

A critical hardware component in a TS7650G implementation is the disk array that holds the ProtecTIER repository. The repository holds the ProtecTIER deduped data.

The types of supported disk arrays can be found at:

[ftp://service.boulder.ibm.com/storage/tape/ts7650\\_support\\_matrix.pdf](ftp://service.boulder.ibm.com/storage/tape/ts7650_support_matrix.pdf)

**Note:**

- ▶ The disk array is not packaged as part of a TS7650G order from IBM. You may order a new disk subsystem from IBM or another vendor (if supported) or you may use an existing disk subsystem at your site that is listed in the referenced above support matrix.
- ▶ ProtecTIER does not require additional compression to be effective. ProtecTIER performs compression, by default, after the deduplication process. Do not allocate disk arrays for ProtecTIER VTLs that perform their own compression.

### 2.3.4 Deployment

A TS7650G can be deployed in a single node or a two-node cluster. The two-node cluster provides increased performance, provides hardware redundancy, and provides a high availability solution that can allow you to continue to have access to the all virtual data stored even if one node fail, with some manual intervention.

## Single node configuration

A single node TS7650G (Figure 2-2) contains a single repository. Within that repository, you can define:

- ▶ Up to 16 virtual libraries
- ▶ Up to 256 virtual tape drives
- ▶ Up to 512,000 virtual cartridges




Figure 2-2 TS7650G single node configuration

A single node configuration has four front-end ports for attaching backup applications and provides up to 1400 MBps throughput for deduplication (depend on the configuration of attached disks) and up to 1000 MBps throughput for replication. Figure 2-2 shows an example of single node configuration.

## Two-node cluster configuration

A two-node cluster TS7650G configuration has two IBM System x5 and provides better performance, up to 2000 MBps (depend on the configuration of attached disks), than a single node system. A two-node cluster configuration has a single and same repository dimensions as a single node configuration. Within that repository, you can define:

- ▶ Up to 16 virtual libraries
- ▶ Up to 512 virtual tape drives (256 per node)
- ▶ Up to 512,000 virtual cartridges

A two-node cluster configuration has an additional network power switch that is used to automatically control the power to a failed node. It also has two Ethernet switches, which are used to connect the TS7650G servers with the network power switch, as shown in Figure 2-3. The Ethernet switches also connect the TS7650G servers with the ProtecTIER Manager workstation and the IBM TotalStorage® Service Console (TSSC) TS3000.

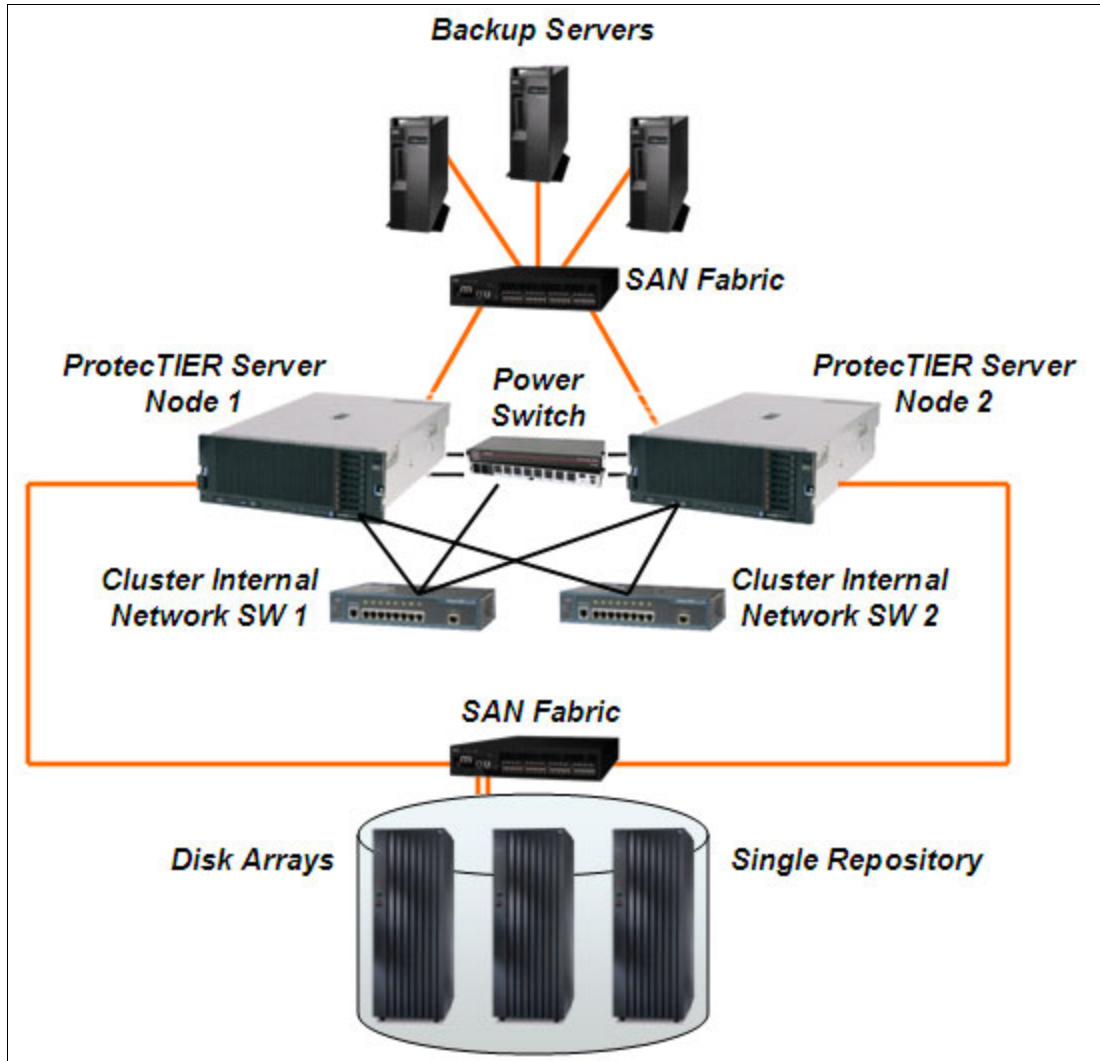


Figure 2-3 TS7650G two-node cluster configuration

The two-node cluster configuration was designed for better performance, but having redundant nodes might provide improved availability for a single node outage. Because both nodes share a single repository, they have access to all virtual resources defined in the repository. As appropriate for your operations, you can have a single virtual library or multiple virtual libraries. Within a virtual library, for example, the robot is assigned to one node in the two-node cluster, but virtual tape drives contained in a library can be defined to either node to provide load balancing.

Figure 2-4 shows a TS7650G two-node cluster configuration with a single virtual tape library. For this example, we assume a single backup application on a single node with a requirement for six virtual tape drives. The virtual tape library consists of a virtual tape library robot and six virtual tape drives. To provide load balancing, the virtual robot and three virtual tape drives are assigned to server 1 and the other three virtual tape drives are assigned to server 2.




Figure 2-4 TS7650G two-node cluster with single virtual tape library

While this configuration can provide load balancing, it does have a single point of failure. If server 1 fails, the virtual tape library robot is unavailable and all access to the virtual tape library is lost. All jobs will fail. Using the ProtecTIER Manager, you can reassign the virtual tape library robot and virtual tape drives to server 2, perform the appropriate reconfiguration in your backup application, and restart the backup application tasks or server, as appropriate for your backup application.

If server 2 fails, instead of server 1, you would lose three virtual tape drives and the backup application tasks that were running on that server, but you would be able to continue to run on server 1 on with the three virtual tape drives available. You can reassign the virtual tape drives to server 1. You must restart the tasks that were running on server 2.

Figure 2-5 shows a TS7650G two-node cluster configuration with two virtual tape libraries. For this example, we assume a single backup application on a single server with a requirement for six virtual tape drives. Each virtual library consists of a virtual tape library robot and six virtual tape drives. The virtual robot and six virtual tape drives from virtual tape library A are assigned to server 1 and the virtual robot and six virtual tape drives from virtual tape library B are assigned to server 2. The backup application actively uses both virtual tape libraries, but only uses three of the virtual tape drives in each virtual library during normal operations.

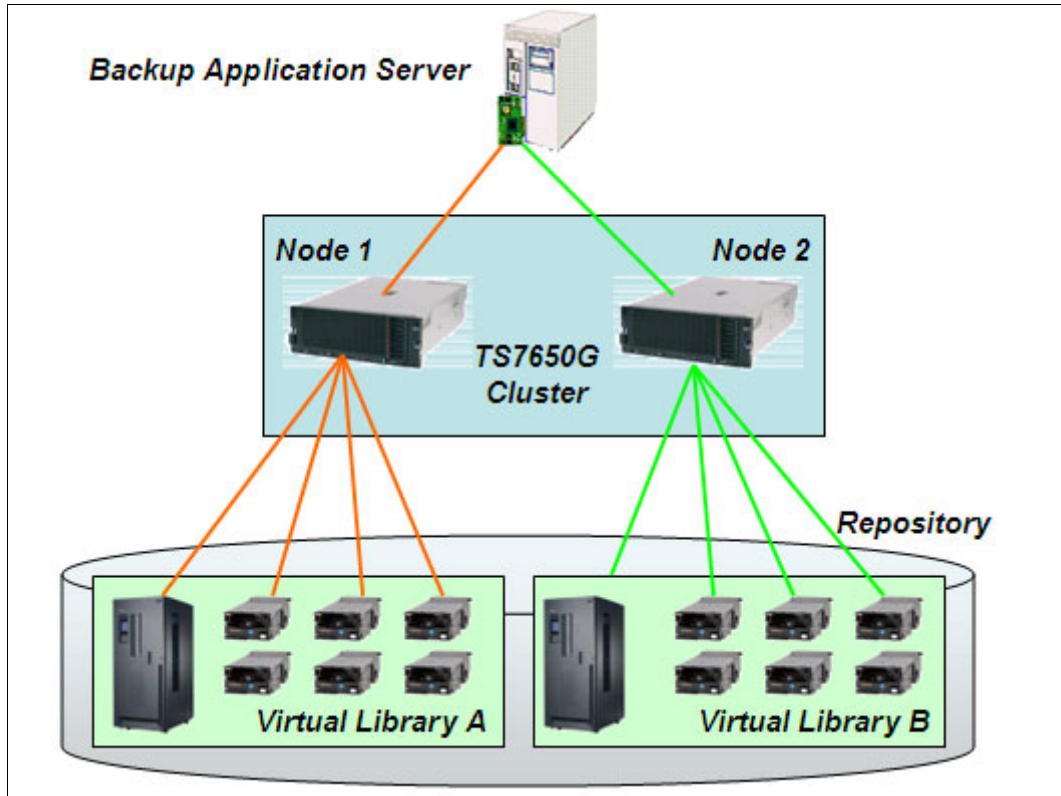


Figure 2-5 TS7650G two-node cluster configuration with two virtual tape libraries

This configuration does not have a single point of failure because there is one virtual robot per TS7650G server. If server 1 fails, server 2 can continue to mount virtual tapes and run backup application jobs (or vice versa). However, you do not have access to the virtual cartridges in virtual tape library A. All tasks running on server 1 fail.

Using the ProtectTIER Manager, you can reassign the virtual tape library robot and virtual tape drives in virtual library A to server 2, perform the appropriate reconfiguration in your backup application, and restart the backup application tasks or server, as appropriate for your backup application. Alternatively, if you do not need to access the virtual cartridges in library A, the backup application can continue using server 2 and virtual tape library B with six virtual tape drives by using the extra three drives that were defined but previously unused.

If all the prerequisites are available and you are able to enable control path failover (CPF), you can assign multiple robots to one tape library. Using TS3500 tape library emulation and an IBM tape device driver on your host, you can enable CPF, which means there will be redundant paths to the robot. In case there is a failure on Server 1, you still have a path to the robot on Server 2 and paths to 3 drives, and you will be able to access all cartridges. See Figure 2-6.

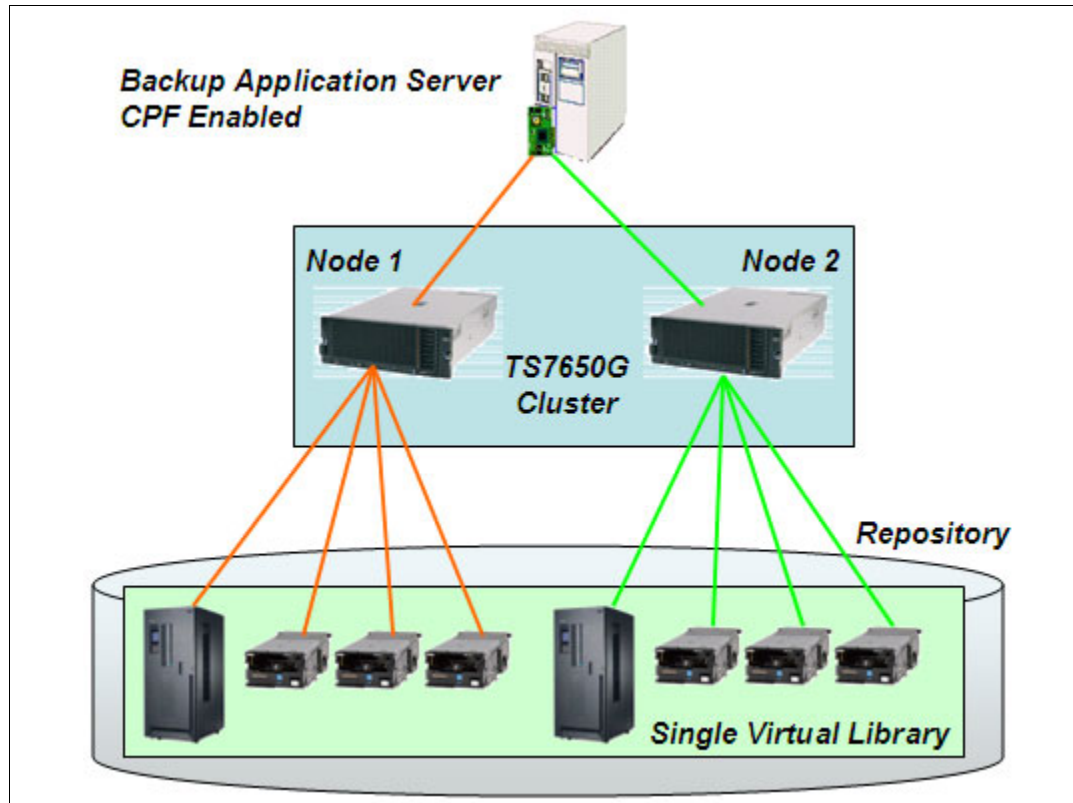


Figure 2-6 TS7650G two-node cluster configuration with one virtual tape library with CPF enabled

This configuration does not have a single point of failure, as there is one virtual robot per TS7650G server. If server 1 fails, server 2 can continue to mount virtual tapes and run backup application jobs (or vice versa), and still be able to access all the virtual cartridges in the virtual tape library. There is no need to reassign the virtual tape library robot or virtual tape drives to server 2. However, if there are more jobs running in the queue than the number of drives available, some of the jobs will have to wait until that one of the drives becomes available.

**Note:** Beside the CPF enabled by IBM tape device driver, Symantec NetBackup 6.5.2 onwards also support multiple paths for robots or tape libraries to achieve highly available tape libraries.

As you can see, with a two-node cluster configuration, you have the possibility to continue operating due to the redundancy of the nodes, but you do not have a failover or high availability function at the backup server.

## 2.4 TS7650 ProtecTIER Deduplication Appliance

The IBM System Storage TS7650 ProtecTIER Deduplication Appliance is a preconfigured virtualization solution of IBM storage, IBM systems, and the revolutionary IBM ProtecTIER data deduplication software designed to improve backup and recovery operations. This solution is not just a bundle of components, but a truly integrated solution that makes it easy to harness the power of deduplication without making radical changes to the existing environment. The solution is available in four configurations designed to meet the disk-based data protection needs of a wide variety of organizations, from mid-sized IT environments to enterprise data centers.

### 2.4.1 TS7650 Deduplication Appliance

These are terms for the IBM self-contained virtualization solution from the IBM System Storage TS7600 family that includes a disk storage repository. The TS7650 Appliance consists of the following components:

<b>Server</b>	The 3958-AP1 server is based on an IBM System x3850 X5 Type 7145. When used as a server in the TS7650 Appliance, its machine type and model are 3958-AP1. Use this machine type and model for service purposes.
<b>System console</b>	The system console is a TS3000 System Console (TSSC). This book uses the terms <i>system console</i> and <i>TSSC</i> interchangeably.
<b>Disk controller</b>	The disk controller is an IBM System Storage DS5020 Express. When used as a disk controller in the TS7650 Appliance, its machine type and model are 1814-70H. Use this machine type and model for service purposes.
<b>Disk expansion module</b>	The disk expansion module is an IBM System Storage DS4000® EXP810 Storage Expansion Unit. When used as a disk expansion module in the TS7650 Appliance, its machine type and the IBM disk expansion module are an IBM System Storage DS4000 EXP810 Storage Expansion Unit. Each expansion unit supports sixteen 450 GB (15k rpm), 4 Gbps, Fibre Channel disk drive modules (Dams).

**Note:** All components are mounted in an shipped IBM 36U Rack unit.

### 2.4.2 TS7650 ProtecTIER Deduplication Appliance features

The IBM System Storage TS7650 Appliance offers many features that can create savings in physical storage, processing, and network bandwidth. Some of the main features are:

- ▶ IBM ProtecTIER with patented HyperFactor
- ▶ Data deduplication technology
- ▶ ProtecTIER native replication technology
- ▶ IBM System x® server for enterprise-level performance
- ▶ IBM Storage Controller with highly reliable Fibre Channel drives
- ▶ Rack, cables, and other components needed to provide a complete solution
- ▶ Up to 500 MBps or more inline data deduplication performance
- ▶ Up to 25 times or more storage capacity reduction
- ▶ Emulation of up to 12 virtual libraries, 256 virtual drives, and 128,000 virtual cartridges



- ▶ Preconfigured with a DS5020 for storage repositories
- ▶ Simplified configuration and deployment

### 2.4.3 Available models

IBM System Storage TS7650 Appliance is available in the four configurations provided in Table 2-4.

Table 2-4 Available models

Size of repository	7 TB configuration	18 TB configuration	36 TB configuration	36 TB two-node configuration
Nominal capacity <sup>a</sup>	175 TB	450 TB	900 TB	900 TB
Maximum throughput	Up to 100 MBps	Up to 250 MBps	Up to 500 MBps	Up to 500 MBps

a. Based on a deduplication ratio of 25:1.

### 2.4.4 Deployment

A TS7650 Appliance can be deployed in a single node or two-node cluster. The two-node cluster provides an increased number of virtual drives and provides hardware redundancy that can allow you to continue to have access to the virtual data after a node failure with some manual intervention.

#### Single node configuration

A single node TS7650 Appliance (Figure 2-7) contains a single repository. Within that repository, you can define:

- ▶ Up to 12 virtual libraries per two-node cluster
- ▶ Up to 256 virtual tape drives
- ▶ Up to 128,000 virtual cartridges

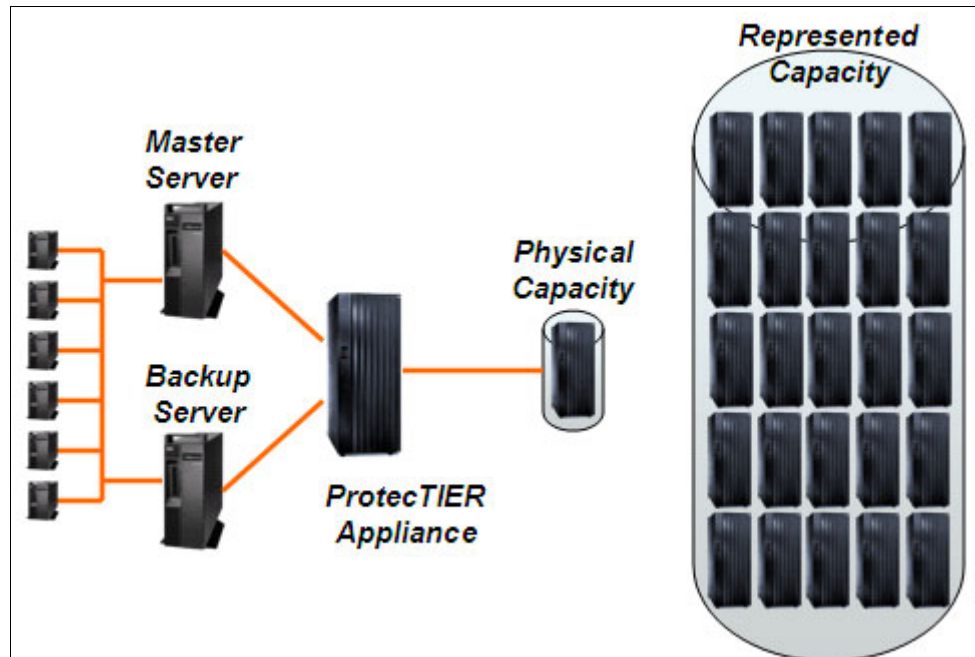


Figure 2-7 Single node appliance configuration

## Two-node cluster configuration

A two-node cluster TS7650 Appliance configuration (Figure 2-8) has two IBM System x3850 X5 Type 7145. A two-node cluster configuration has the same repository dimensions as a single node configuration. A two-node cluster TS7650 Appliance contains a single repository. Within that repository, you can define:

- ▶ Up to 12 virtual libraries
- ▶ Up to 512 virtual tape drives (256 per node)
- ▶ Up to 128,000 virtual cartridges per node




Figure 2-8 Two-node cluster appliance configuration

A two-node cluster TS7650 Appliance and Gateway configuration (Figure 2-9) has an additional Network Power Switch that is used to automatically control the power to a failed node. It also has two Ethernet switches, which are used to connect the TS7650 servers with the network power switch. The Ethernet switches also connect the TS7650 servers with the ProtecTIER Manager workstation and the TotalStorage Service Console (TSSC) TS3000.

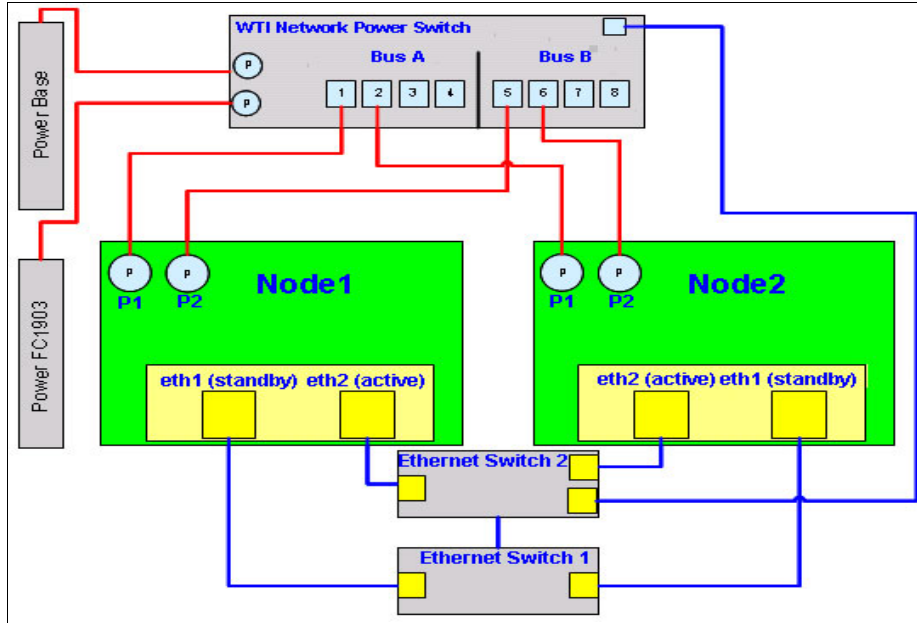


Figure 2-9 Power switch control

Because both nodes share a single repository, they have access to all virtual resources defined in the repository. As appropriate for your operations, you can have a single virtual tape library or multiple virtual tape libraries. In a virtual tape library, for example, the robot is assigned to one node in the two-node cluster, but virtual tape drives contained in a library can be defined to either node to provide load balancing.

## 2.5 TS7610 ProtecTIER Deduplication SMB Appliance

The IBM System Storage TS7610 ProtecTIER Deduplication Appliance Express (TS7610 Appliance) is a customer installable data deduplication and replication solution that compresses and stores backup data within the TS7610. It can be configured to run either as a Virtual Tape Library or a OpenStorage system.

**Note:** The TS7610 Appliance is available in a stand-alone one node configuration only.

Do not connect a 3959-SM1 server directly to a public optical network. The customer must use an additional connectivity device between an 3959-SM1 optical Host Bus Adapter and an external public network. Use a device such as a patch panel, a router, or a switch. You do not need an additional connectivity device for optical fibre connectivity that does not pass through a public network.

The TS7610 Appliance uses IBM technology and software to emulate a TS3500 Tape Library unit containing IBM Ultrium LTO3 tape drives. This emulation capability is designed to help customers achieve the following operational and throughput efficiencies:

- ▶ Backup window reduction
- ▶ Restoration time reduction
- ▶ Direct access to the information
- ▶ Data sharing and resource virtualization optimization
- ▶ Operational efficiency
- ▶ Improved sharing of virtual tape libraries across applications and servers

### 2.5.1 TS7610 hardware components

The TS7610 Appliance consists of the left and right side rails, support tray, and cable management assembly, an external USB DVD drive, and the 3959-SM1 server (Figure 2-10 and Figure 2-11). The shelf rail kit occupies an additional 1U rack space (3U in total frame space).

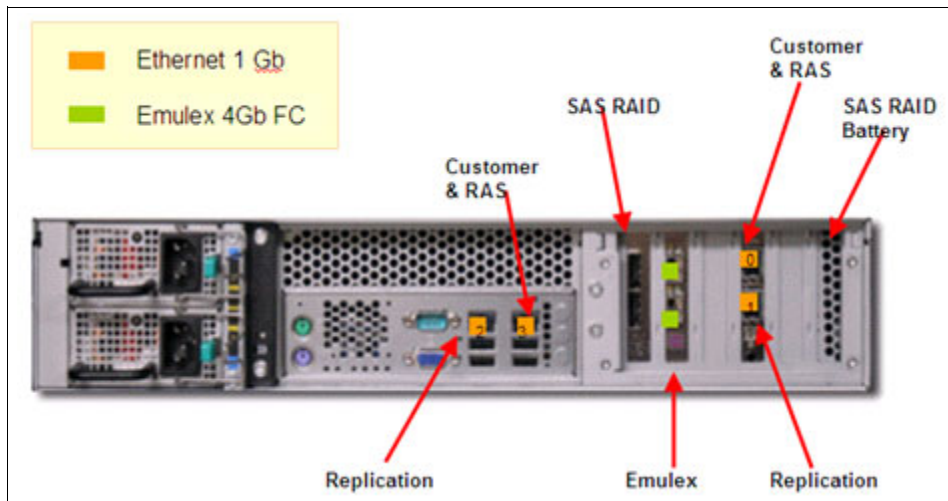


Figure 2-10 Back view of 3959-SM1 server for Virtual Tape Library




Figure 2-11 Back view of 3959-SM1 server for OpenStorage

The TS7610 Appliance provides two different capacity models:

- ▶ Small TS7610 Appliance repository binary capacity of 4.0 TB  
This configuration is ideal for customers with 500 GB or less incremental backups per day and have average data growth.
- ▶ Medium TS7610 Appliance repository binary capacity of 5.4 TB  
This configuration is ideal for customers with 1 TB or less incremental backups per day and 1 TB to 3 TB full backups each week with average to rapid data growth.

A TS7610 small model of 4.0 TB capacity can be upgraded to a 5.4 TB model by a software upgrade; no additional hardware is required.

The 3959-SM1 Appliance server has the following components:

- ▶ One quad core 2.33 GHz Intel XEON processor
- ▶ Twelve 1 TB SATA hard disk drives (HDDs)
- ▶ Two power supplies
- ▶ RAID battery backup unit
- ▶ Six 4 GB dual in-line memory modules (DIMMs)
- ▶ Dual-port 4 Gb Emulex Fibre Channel host bus adapter (for VTL)
- ▶ Dual-port 1 Gb Ethernet adapter (for OST)

Also, you can see all hardware components using ProtecTIER Manager GUI, as shown in Figure 2-12, by choosing the resources at the hardware resources column.

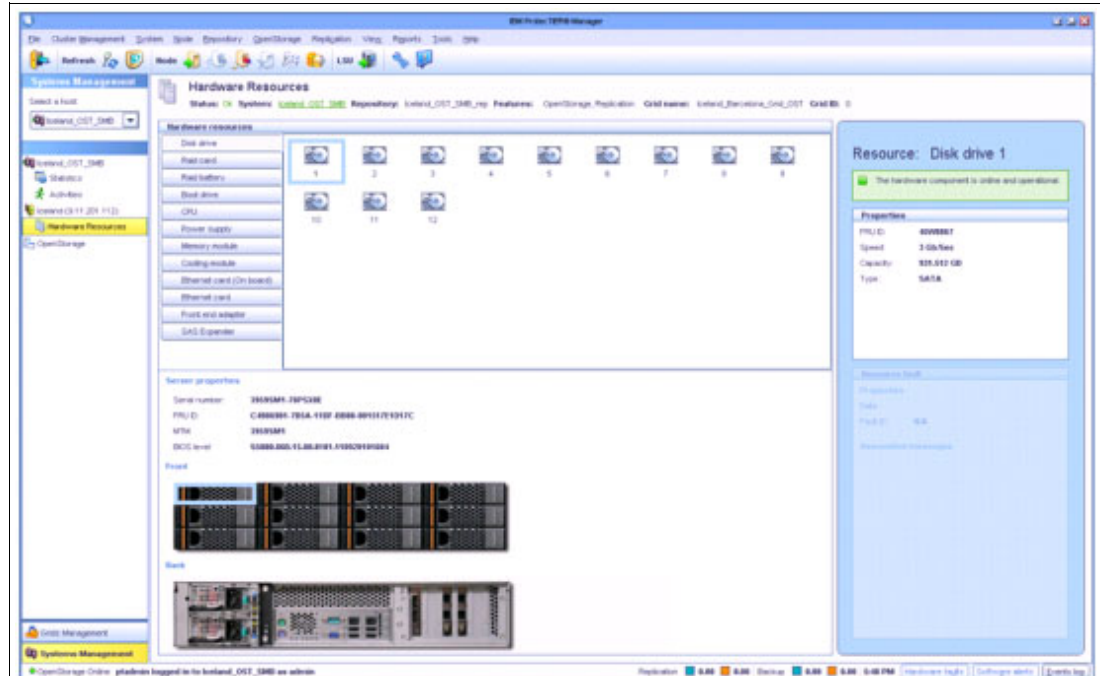


Figure 2-12 ProtecTIER Manager hardware resource view

## 2.5.2 Deployment

The TS7610 Appliance system is a one node system. A one node system uses one server to transfer data from the backup server to the storage fabric, as illustrated in Figure 2-13.




Figure 2-13 A typical TS7610 one node ProtecTIER system deployment

Where:

1. Backup client server (supplied by customer)
2. Main backup server (supplied by customer)
3. Fibre Channel cables connection (not included in ProtecTIER package)
4. TS7610 SMB Appliance

Within the TS7610 appliance, you can define:

- ▶ Up to 4 x TS3500 virtual tape libraries
- ▶ Up to 64 x LTO 3 virtual tape drives
- ▶ Up to 8,000 virtual cartridges
- ▶ Up to 1,022 import/export slots

The TS7610 VTL appliance will be shipped with the software preconfigured as follows:

- ▶ A TS3500 virtual library.
- ▶ 16 LTO3 virtual drives, balanced evenly across both FC ports.
- ▶ 200 GB cartridge size.
- ▶ 16 virtual import export slots.
- ▶ The Small model will have 400 virtual slots and 400 virtual cartridges.
- ▶ The Medium model will have 540 virtual slots and 540 virtual cartridges.

You can modify the configuration according to your own requirements when you need to add more tape libraries, tape drives, cartridges, and slots.

## 2.6 ProtecTIER virtual tape library

A ProtecTIER virtual tape library (VTL) (Figure 2-14) emulates traditional tape libraries. By emulating tape libraries, ProtecTIER enables you to transition 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 ProtecTIER actually stores data on a deduplicated disk repository on the storage fabric or direct attached storage.

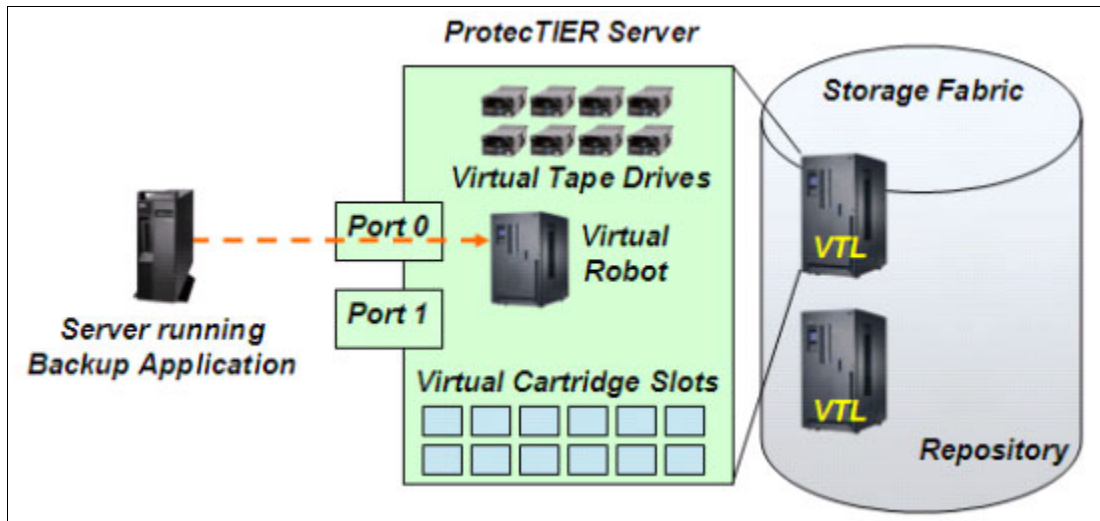


Figure 2-14 ProtecTIER virtual tape library

## 2.6.1 ProtecTIER VTL concepts

After the data is ingested, the ProtecTIER VTL functions like a traditional VTL with the addition of the deduplication processing requirements.

When duplicate data is identified by ProtecTIER, it updates a reference count in the database. ProtecTIER uses the reference count to determine when a data segment can be overwritten (deleted). As shown in Figure 2-15, sector 3 represents a segment that occurs four times within the virtual cartridges in the repository. In the lower left corner is a representation of the reference table showing that sector 3 is referenced four times.

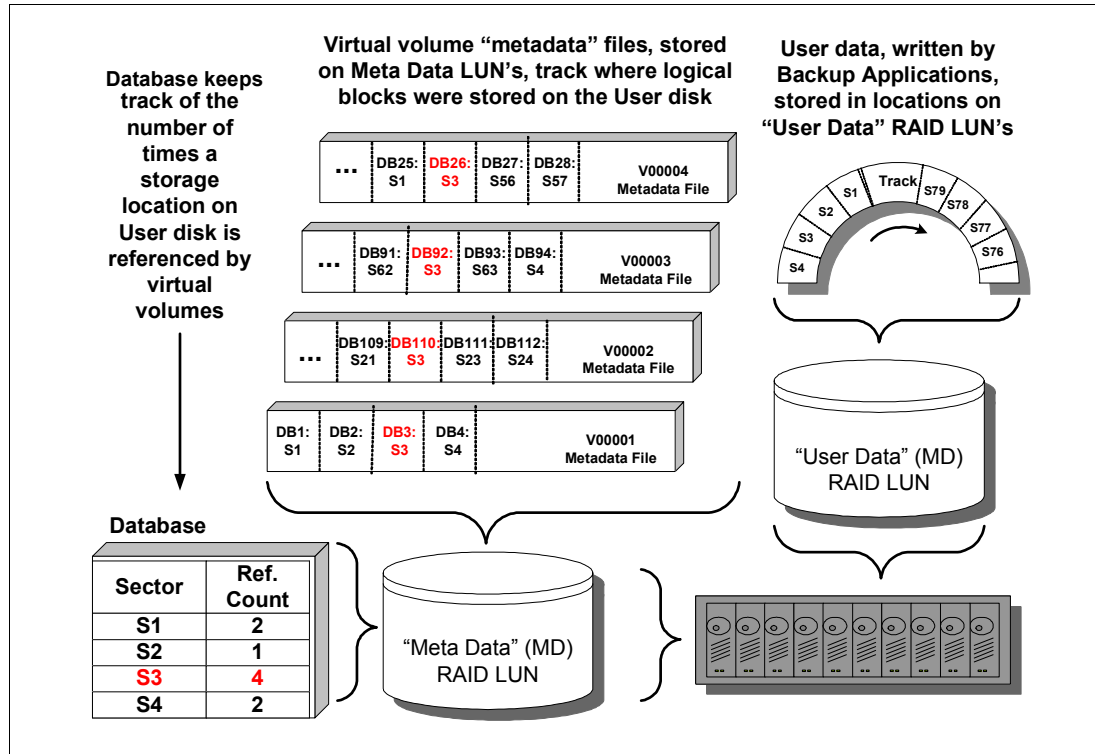


Figure 2-15 ProtecTIER virtual tape library concepts

ProtecTIER uses the metadata files to read back the virtual cartridges. When a virtual cartridge is overwritten or deleted, the reference count for each segment of user data on a virtual cartridge is decremented. After the reference count reaches zero, the space occupied by that user data segment is released.

ProtecTIER uses the Global File System (GFS) to allow multiple references to the same data.

## 2.6.2 Steady state

After ProtecTIER has been operational for some period of time, it reaches a *steady state*. The point at which a steady state is reached varies based on the size of the cache and the frequency of backups. To understand the concept, it might help to think in terms of a real physical library. In those terms, a physical library reaches a steady state when all cartridges have been used, but enough cartridges become scratch every night to provide the media required for the next day's backup.

In that context, if you allow ProtecTIER to decide how much data fits on every cartridge, in theory, if you accurately predicted the factoring ratio, when you fill the last available scratch cartridge you have consumed the last of the usable space in the repository.



Until all available scratch cartridges are used, ProtecTIER only performs two types of input/output (I/O) to the RAID. While doing backup, it is performing random reads to the user data disk to *prove* the duplicate data (90%) and it is performing writes of new user data (10%). As that backup data is being written to user data disk, ProtecTIER is also doing roughly 90% random write I/O to the metadata LUNs, as it updates the virtual cartridge metadata files, to record where the user data is stored for each virtual cartridge.

After you fill your last virtual scratch tape, and you use all the space in the repository, then you are positioned to enter a steady state. At that point, the next time that the backup application performs a backup, the data must be written to a virtual cartridge that was previously used and filled. When that virtual cartridge is mounted and positioned at load point, and writing begins, all of the metadata files associated with the prior use of that virtual cartridges must be processed.

The first step of that processing reads the contents of each of the old metadata files, finds every reference to the user data on disk, and decrements the reference count for each storage block identified. After all references in a metadata file are processed, the metadata file is deleted. Each time that the reference count of a storage block goes to zero, that storage block gets returned to the pool of *free blocks* and becomes usable as free space.

Not all units of *free space* are usable in ProtecTIER. The smallest unit of usable, or *allocatable space*, in the repository is 1 MB. A storage block is 16 K. As storage blocks are freed as a result of an overwrite of a virtual cartridge, some of the space freed will be in amounts that are less than 1 MB of contiguous blocks, and the system must defragment the file system. ProtecTIER keeps one block group free for defragmentation. The *active blocks* in a single block group are copied to contiguous space in the free block groups, essentially *defragging* the block group. The block group from which the data was copied becomes the new free block group.

All of this processing occurs in the background and can occur while new data is written to new metadata files associated with the virtual cartridge being overwritten.

After the system enters a steady state, ProtecTIER is managing four types of I/O activity to the disk: the two types of I/O performed as all virtual cartridges were filled (standard backup activity, as described previously), plus the new additional work of:

- ▶ Reading and deleting old metadata files
- ▶ Moving data from fragmented block groups to free space in block groups

To prevent the defragmentation operations from impacting performance, they are only allowed a maximum of 15% of the total input/outputs operations per second (IOPS) of the system.

From a performance standpoint, when the ProtecTIER system first begins ingesting data, ProtecTIER is not matching new data to existing data and no fragmentation is occurring. This enables high performance. After a steady state is achieved, performance stabilizes.

Figure 2-16 is a conceptualization of the performance from initial implementation through a steady state. The change in performance and the time to reach a steady state depends on the size of the repository and the operational characteristics of your data. For more information about performance, refer to Chapter 3, “Planning for deduplication and replication” on page 53.

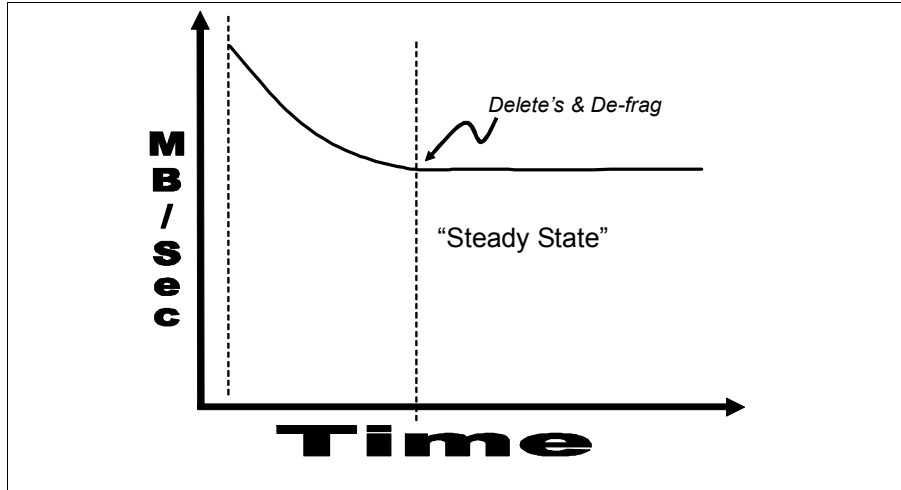


Figure 2-16 ProtecTIER performance from initial implementation through a steady state

## 2.7 ProtecTIER OpenStorage

With OpenStorage, ProtecTIER can be integrated with NetBackup to provide the means for backup-to-disk without having to emulate traditional tape libraries. Using a *plug-in* that is installed on an OST-enabled media server, ProtecTIER can implement a communication protocol that supports data transfer and control between the backup server and the ProtecTIER server. Therefore, to support the plug-in, ProtecTIER implements a storage server emulation (Figure 2-17).

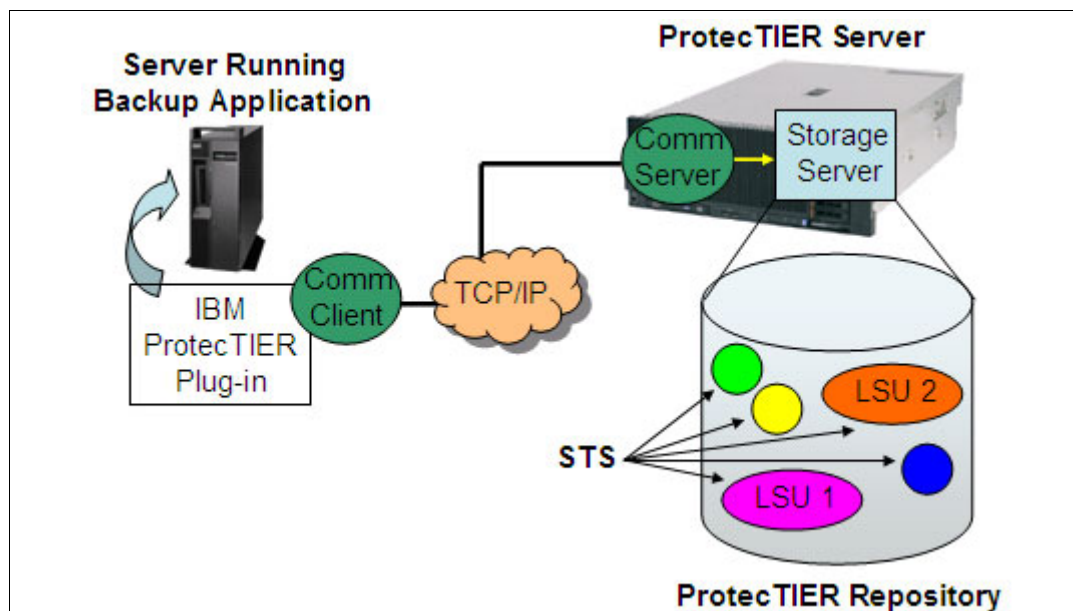


Figure 2-17 OpenStorage emulation

The OpenStorage interface is a NetBackup appliance programming interface (API) supporting communications between NetBackup systems and OpenStorage-enabled ProtecTIER systems. The OpenStorage interface provides the following capabilities:

- ▶ NetBackup software directs ProtecTIER systems when to create, copy, delete, or restore backup images.
- ▶ Backup images may be replicated to up to 12 different ProtecTIER systems.
- ▶ Workloads and performance are balanced between NetBackup media servers and ProtecTIER systems.
- ▶ Detailed statistics on the state of the ProtecTIER systems are available through API.

**Note:** OpenStorage supports many-to-many bidirectional replication with up to 12 systems in a hub mesh group.

For more information, refer to Chapter 9, “IBM System Storage ProtecTIER with Symantec OpenStorage” on page 433.

## 2.8 Data deduplication

Data deduplication solutions from IBM employ an advanced form of data compression that identifies and eliminates redundant data across the data landscape, making it possible to significantly reduce the amount of data that must be protected. This in turn dramatically increases the effective capacity of existing disk storage so that far less physical disk is required to protect the data, beyond the direct resource savings associated with needing less disk space, which can be in the hundreds of thousands or millions of dollars (Figure 2-18).




Figure 2-18 Disk savings with ProtecTIER

The benefits of data deduplication include:

- ▶ Greater productivity that comes from being able to perform more frequent backups with the same amount of physical disk
- ▶ Increased efficiency because of the greater likelihood that data will be able to be restored from disk rather than a slower medium
- ▶ Reduced energy consumption that results from reducing the amount of disk in operation

Data deduplication uses algorithms to compare data and eliminate duplicated or redundant data. Standard compression only works at the file level, while IBM Real-time Compression appliances and database real-time compression both work at the subfile level. Data deduplication can be applied at the subfile level to reduce data size across a much broader landscape.

HyperFactor technology uses a pattern algorithm that can reduce the amount of space required for storage in the backup environment by up to a factor of 25 times, based on evidence from existing implementations. The capacity expansion that results from data deduplication is often expressed as a ratio, essentially the ratio of nominal data to the physical storage used. A 10:1 ratio, for example, means that 10 times more nominal data is being managed than the physical space required to store it. Capacity savings of 18:1 and greater have been reported from data deduplication, up to 25:1 in the case of IBM solutions.

## 2.8.1 Virtual tape library concept

A virtual tape library is a software application that emulates physical tape while storing your data on a disk subsystem. Some vendors only sell the software, and you are required to assemble your own system by procuring the supported components (server, disk subsystem, and switches) and then you must integrate those components. Other vendors might provide an appliance in which all the components are included and integrated. Still other vendors might provide a solution somewhere between the two solutions.

Regardless of the procurement and integration method, all VTLs provide the basic function of emulating physical tape. To perform this function, the VTL software must accept the normal tape commands issued by a backup application and provide the appropriate responses back to the backup application. The VTL software must also store your data on the random access disk subsystem that allows it to retrieve the data and present it to the backup application as a sequential stream.

The basic concept of VTL data storage is shown in Figure 2-19.

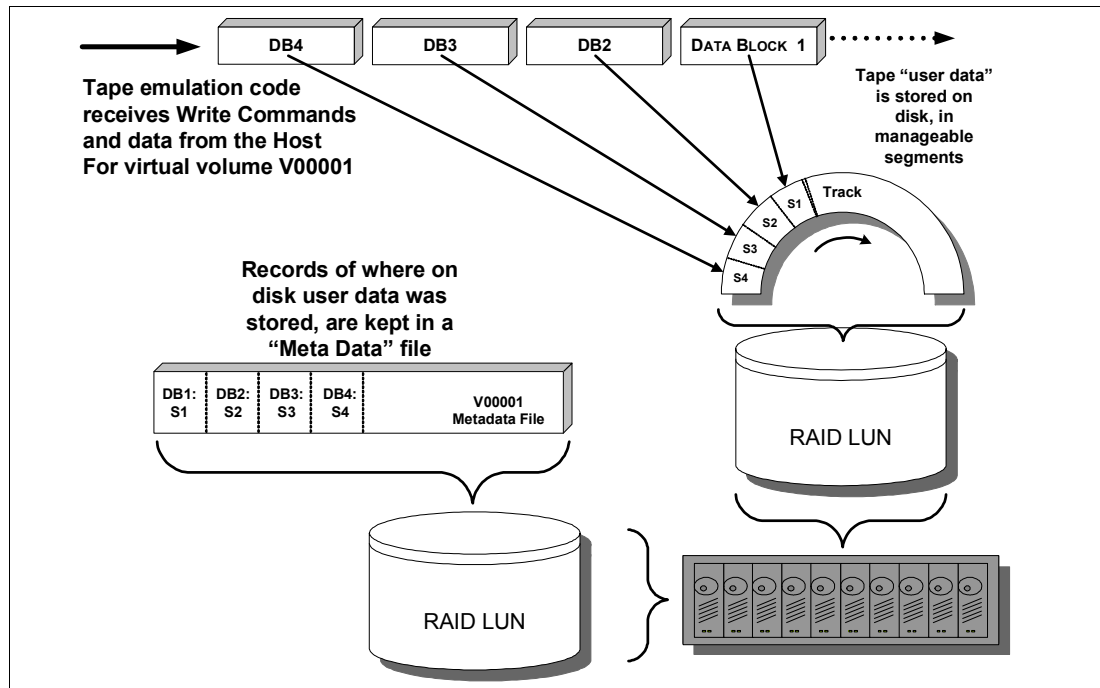


Figure 2-19 Virtual tape library concept

When data is written from the backup application to the VTL, the VTL receives the write commands and the data, in this case, virtual volume V00001. The data is stored on the disk in segments and the VTL uses a metadata file, or database, to keep track of each segment of data. When the VTL receives a read command from the backup application, it uses the metadata file to retrieve the data and present it to the backup application as sequential data.

One concept of a virtual tape library that frequently causes confusion is the allocation of space for virtual cartridges after the backup application has *expired* a cartridge or marked it as *scratch*. When expiring tape cartridges in a backup application, a virtual tape library emulates the operation of a physical tape library, which means that when the backup application expires a tape cartridge, it is simply an indication within the backup application that the tape cartridge is now available for scratch use. The backup application does not issue any commands to the tape library to modify the expired tape cartridge. In a physical library, the physical tape cartridge remains in the library and the data written on the tape cartridge stays on the tape cartridge until the backup application mounts the cartridge and performs a scratch write (write from beginning of tape (BOT)).

A virtual tape library functions in the same way. When the backup application expires a virtual tape cartridge, the virtual tape cartridge stays in the library and the data contained on the virtual tape cartridge remains until the backup application performs a scratch write (write from BOT). When viewing the space utilization for the virtual library, the expired virtual tape cartridges continue to use allocated space until the scratch write is performed. If the utilized space on your virtual tape library is high but you have sufficient virtual tape cartridges in expired or scratch status in the backup application, you might still be able to work within your installed capacity. If the utilized space on your virtual tape library is high and you do not have sufficient virtual tape cartridges in expired status in your backup application, you might need to investigate increasing your capacity.

## 2.8.2 HyperFactor

ProtectTIER performs the basic VTL function, but has the added benefit of data deduplication provided by the IBM HyperFactor technology. HyperFactor was designed from the top-down to overcome the known limitations of hash-based deduplication.

Hash-based data deduplication is an all or nothing proposition. Either a chunk of data is identical or it is not. Changing a single byte renders a chunk completely different, even though almost all of the data remains the same. Thus, systems that look for exact chunk matches have many *missed opportunities* for data factoring because the new data is often closely related to previously stored data.

HyperFactor takes a different approach and therefore reduces the phenomenon of missed factoring opportunities. Rather than depending on exact chunk matches, HyperFactor finds matches between *similar data*. When new data is received by ProtectTIER, HyperFactor finds any similar data elements that have already been stored (Figure 2-20). This search is extremely quick, using a small and efficient memory-resident index. After the similar data elements are found, HyperFactor can then compare the new data to the similar data to identify and store only the byte-level changes.

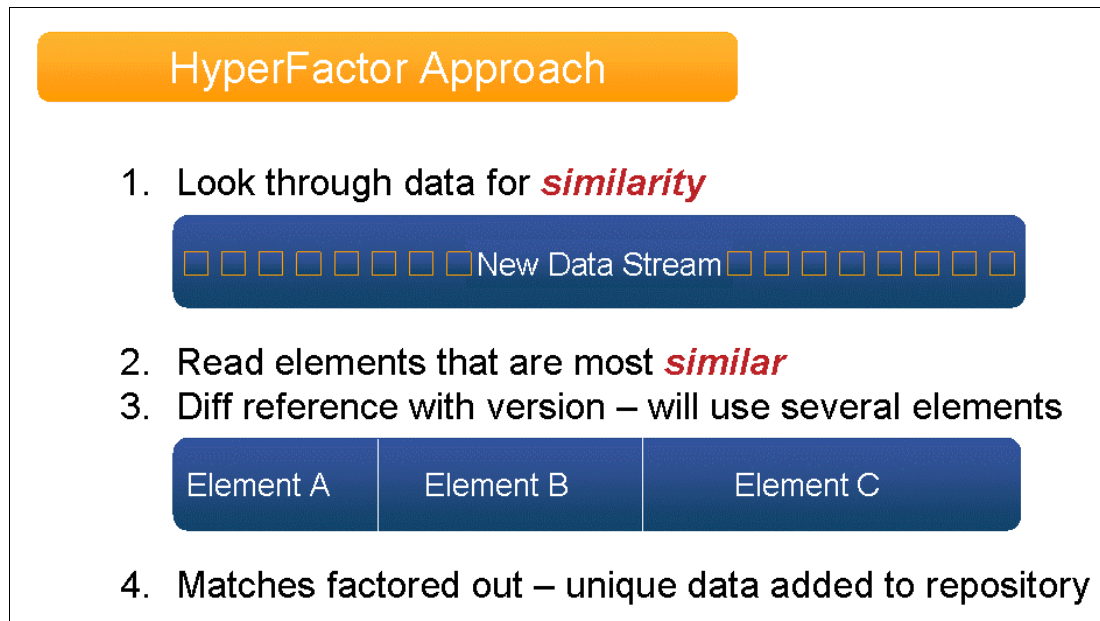


Figure 2-20 HyperFactor approach

With this approach, HyperFactor is able to surpass the reduction ratios attainable by any other data reduction method. HyperFactor can reduce any duplicate data, regardless of its location or how recently it was stored. Unlike hash-based techniques, HyperFactor finds duplicate data without needing exact matches of chunks of data. When new data is received, HyperFactor checks to see whether similar data has already been stored. If similar data has already been stored, then only the difference between the new data and previously stored data must be retained. Not only is this an effective technique of finding duplicate data, but it performs well.

High performance is enabled by using a limited amount of memory and computation processing to store the index, requiring fewer items and much less memory than with hash-based data deduplication. HyperFactor requires only 4 GB of memory to index 1 PB of data, leading to a fast and efficient system that can search its index without disk I/O or large amounts of RAM.

The core technology in ProtecTIER-HyperFactor is a series of algorithms that factor, or deduplicate, data efficiently. In each new backup, HyperFactor finds the data in common with previous backups. This common data in the new backup is effectively *filtered out* and pointers are used to reference existing data in the repository. The net effect is that the entire content of the new backup is stored in the space required to only store the small fraction of it that is truly new data.

Over time, the effect of HyperFactor is a system-wide *factoring ratio*. In essence, the factoring ratio is the ratio of nominal data (the sum of all user data backup streams) to the physical storage used. With ProtecTIER, this factoring ratio can grow to 25:1 or more, meaning that 25 times more nominal data is managed by ProtecTIER than the physical space required to store it.

The factoring ratio of your data depends heavily on two key variables:

1. Data retention period: The period of time (usually measured in days) that defines how long customers will keep their disk-based backups online. This period of time typically ranges from a period of 30 to 90 days, but can be less or much longer.
2. Data change rate: The rate at which the data received from the backup application changes from backup to backup. This measurement has the most relevance when *like* backup policies are compared (Data change rates might range from 1% to >25%, but are difficult to directly observe).

### 2.8.3 ProtecTIER data ingest flow

ProtecTIER performs the deduplication process on data ingest, which means data deduplication is performed inline. The data flow is shown in Figure 2-21.

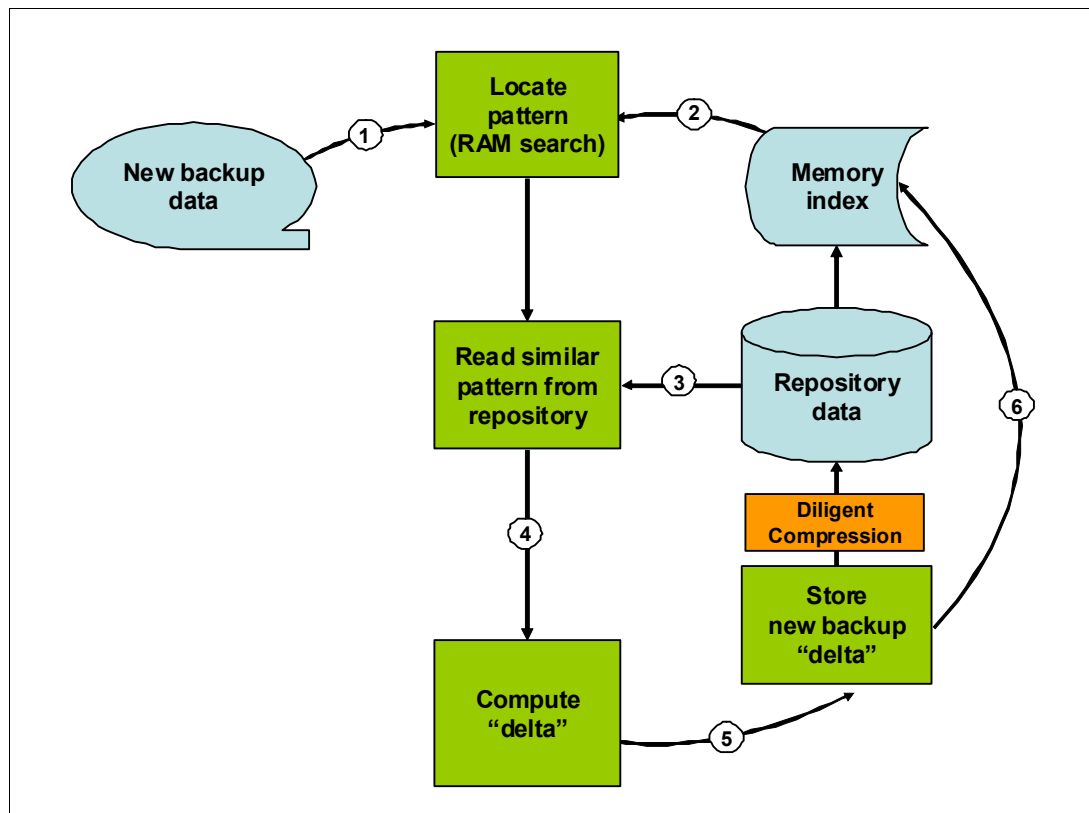


Figure 2-21 ProtecTIER software data ingest flow

The data flow is as follows:

1. A new data stream is sent to the ProtecTIER server, where it is first received and analyzed by HyperFactor.
2. For each data element in the new data stream, HyperFactor searches the Memory Resident Index in ProtecTIER to locate the data in the repository that is most similar to the data element.
3. The similar data from the repository is read.
4. A binary differential between the new data element and the data from the repository is performed, resulting in the delta difference.
5. The delta from step 4 is now written to the disk repository after being processed with the IBM Diligent® Compression algorithm, which is similar to Lempel-Ziv-Haruyasu (LZH). With LZH compression, additional size reduction might be achieved for delta data. Some size reduction might be accomplished for new data (such as the initial backup of unique new data) through this compression.
6. The Memory Resident Index is updated with the location of the new data that has been added. The Memory Resident Index is written to the metadata file system frequently.

After the duplicate data is identified, the Memory Resident Index is not needed to read the data, which eliminates the concern that the Memory Resident Index could be corrupted or lost and therefore access to the data might be compromised. Because the Memory Resident Index is only used for data deduplication on data ingest, data accessibility remains if the index is lost. The Memory Resident Index is restored from the metadata file system, if needed. If the Memory Resident Index was lost and restored, any index updates for deduplication that occurred in the window between the last index save and the index loss would be unavailable and new data could not be compared for any similarities developed during that short window. The only impact from this would be a slight, probably unmeasurable, reduction in the overall deduplication ratio.

## 2.9 ProtecTIER native replication

Prior to the availability of ProtecTIER native replication, ProtecTIER users relied, in most cases, on disk array based replication products. This meant that every changed disk block was replicated, even when it was only due to *data scrubbing* inside the ProtecTIER system. This situation eliminated a large portion of the ProtecTIER deduplication effect on the replication, and resulted in only a small amount of reduction in the amount of network bandwidth required compared with a first generation of VTL replication with no deduplication. Now, with ProtecTIER's native replication, only the changed elements of data (unique data) are being transferred over the network, which results in a dramatic reduction in the bandwidth requirements.

There are two types of native replication:

- ▶ Many to one replication: One hub can receive until there are 12 spokes of replication.
- ▶ Many to Many replication: In a four-way group of ProtecTIER, you can replicate from one to another in a bidirectional fashion.

For more information about replication, refer to Chapter 11, “Native replication and disaster recovery” on page 575.



## 2.9.1 Replication licensing

Replication licensing is done by capacity and varies slightly within the TS7600 family.

### Replication licensing for the TS7650G

Replication is authorized by capacity corresponding to the total terabytes (TB) managed by the TS7650G, which entitles the customer to use the replication functionality of ProtecTIER. It is required for every TS7650G that replicates data to another TS7600 family or receives replicated data from another TS7600 family. Terabyte capacity ordering is by cumulative tier. Starting with tier 1 (1 - 12 TB) until tier 5 (more than 250 TB), order the quantity from each successive tier required to reach the total TB managed by the TS7650G, filling each tier in succession.

### Replication licensing for the TS7650 Appliance

Replication is authorized by capacity corresponding to the total terabytes managed by the TS7650 Appliance, which entitles the customer to use the replication functionality of ProtecTIER. It is required for every TS7650 Appliance that replicates data to another TS7600 family or receives replicated data from another TS7600 family. Terabyte capacity ordering is by cumulative tie. Starting with tier 1 (1 - 12 TB), order the quantity from each successive tier required to reach the total TB managed by the TS7650 Appliance, filling each tier in succession. Valid capacity values for the TS7650 Appliance are 7 TB, 18 TB, and 36 TB.

### Replication licensing for the TS7610 SMB Appliance

Replication is authorized by capacity corresponding to the total terabytes managed by the TS7610 SMB Appliance, which entitles the customer to use the replication functionality of ProtecTIER. It is required for every TS7610 SMB Appliance that replicates data to another TS7600 family or receives replicated data from another TS7600 family. Terabyte capacity ordering is by total appliance capacity, which can be 4.0 TB or 5.4 TiB.

## 2.9.2 Hardware

No additional hardware is required for the 3958-DD3, 3958-DD4 Gateways, 3958-AP1 Appliance, or 3959-SM1 Appliance. For the 3958-DD1 Gateway (oldest model), an additional network card is required.

1. The replication grid is the container object for the ProtecTier one node or two-node cluster systems.
2. Each ProtecTIER two-node cluster has its own repository.

3. The ProtecTIER two-node clusters are connected through an IP-based WAN and exchange data (Figure 2-22).




Figure 2-22 Physical Replication Grid

There are several additions/changes for accommodating IP replication as follows:

- ▶ Each site has a single node or two-node cluster ProtecTIER system.
- ▶ The TS7610 (SM1) has two dedicated Ethernet ports for replication, one onboard and the other on an additional network interface card.
- ▶ The TS7650G (DD3, DD4) and TS7650 (AP1) in a stand-alone configuration have two dedicated Ethernet ports for replication. The TS7650G (DD3, DD4) and TS7650 (AP1) in a cluster configuration have four dedicated Ethernet ports for replication.
- ▶ Replication ports are connected to the customer's WAN.
- ▶ Replication ports are configured on two different subnets as the default.

**Note:** For the 3958-DD1 Gateway (oldest model), an additional network card is required for replication. Contact your IBM representative.

For more information about Ethernet ports and IP Address, refer to Appendix A, "Installation and implementation checklists" on page 687.

## 2.9.3 Replication grid

A replication grid is a set of repositories that share a common ID and can potentially transmit and receive logical objects through replication. A replication grid defines a set of ProtecTIER repositories and actions between them and is configured using the ProtecTIER Replication Manager. ProtecTIER Replication Manager is a software component that comes with new orders of ProtecTIER V2.5. ProtecTIER Replication Manager (Figure 2-23) is activated on a ProtecTIER node.

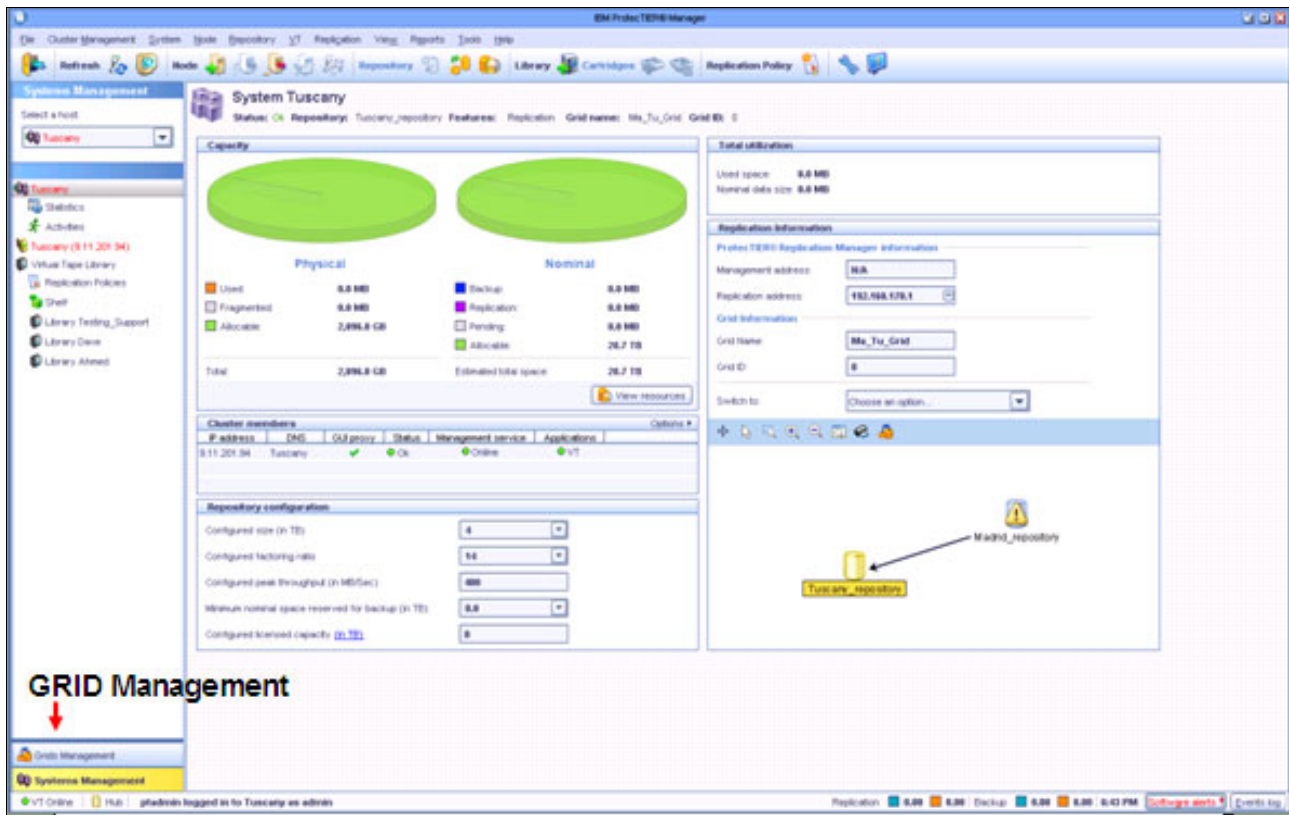


Figure 2-23 ProtecTIER Replication Manager

Figure 2-24 shows an example of two ProtecTIER two-node clusters with a link between them. Each system is composed of a repository that holds VTL and a shelf. A ProtecTier system holds a CM Virtual Tape Facility Daemon (VTFD), and VTL services. Notice in this diagram that there is a virtual shelf. One of the main uses of the virtual shelf is to receive replicated data from the source repository.

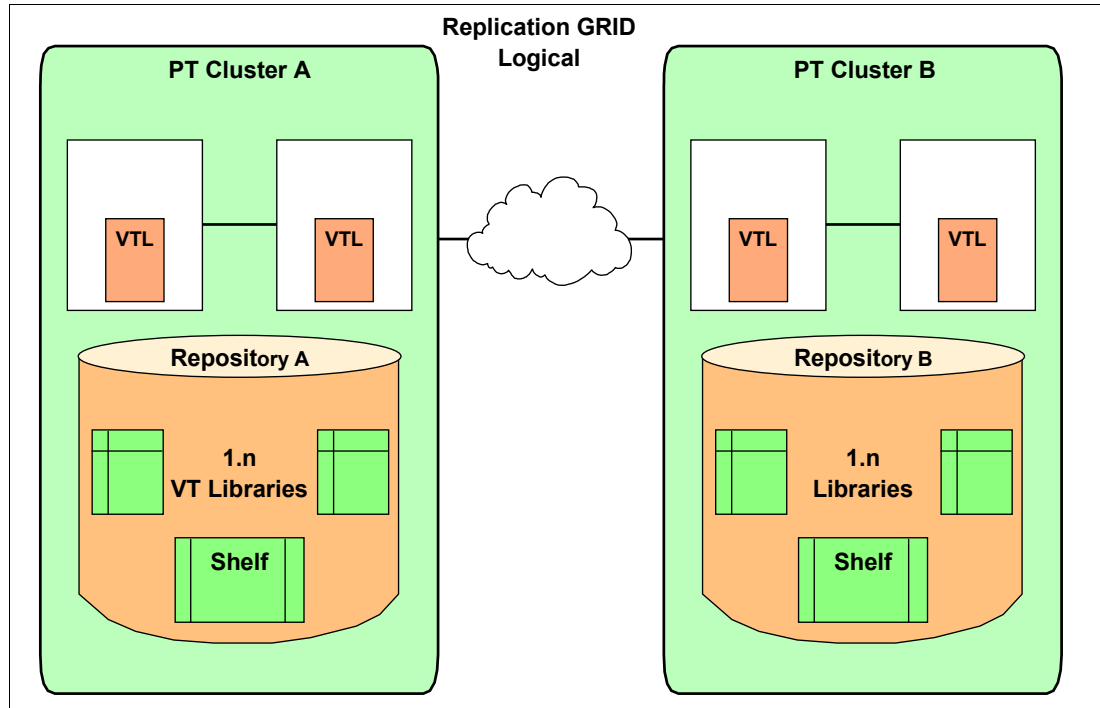


Figure 2-24 Replication grid topology example: Replication pair with two-node cluster systems

Replication has been designed to minimize its impact on backup and restore activities. Policies can be created to allow the user the option of selecting which tapes they want to replicate. Up to 254 policies per repository can be defined.

For more information about replication, refer to Chapter 11, “Native replication and disaster recovery” on page 575.

## 2.10 ProtecTIER Manager

The ProtecTIER Manager (PT Manager) application is a graphical user interface (GUI) you use to view, configure, manage, and troubleshoot the operation of the IBM TS7600 family of products. It can be installed on a Windows or Linux based PC on your Ethernet network. The PT Manager does not need to be located near the ProtecTIER system, as long as Ethernet connectivity to the ProtecTIER system is allowed.

The workstation for the installation of the PT Manager is not included in the ProtecTIER system. You can install the PT Manager GUI on any workstation that meets the following requirements:

- ▶ x86 (Pentium or higher) microprocessor
- ▶ 256 MB memory at minimum
- ▶ 1.2 GB of disk space

- ▶ Access to the ProtecTIER service node's IP address (port 3501 and 3503 are open on the firewall)
- ▶ Keyboard, mouse, and CD-ROM drive
- ▶ Screen resolution of 1024 x 768 or higher
- ▶ 24-bit color or higher
- ▶ Operating environments supported:
  - Windows 32/64 bit (2003, XP, or 7)
  - Linux Red Hat 32/64 bit (Enterprise 4 or higher)

**Note:** If you are planning to run ProtecTIER Manager on a UNIX system, configure your graphics card and X Window System. This is done either manually or by using the Xconfigurator utility. For instructions, refer to the appropriate Linux documentation.

**Note:** The console hardware is not included in a TS7600 family order. You may order the console hardware that meets the specifications from IBM or an OEM vendor.

Figure 2-25 shows one Virtual Tape Library, which is the only virtual library for the repository on the single node system called Hungary. You can see that there is one virtual tape library defined in the repository with 24 virtual drives. You can view more detail about the library by selecting the **Drives**, **Cartridges**, **Slots**, or **Import/Exports** tabs at the top of the view.

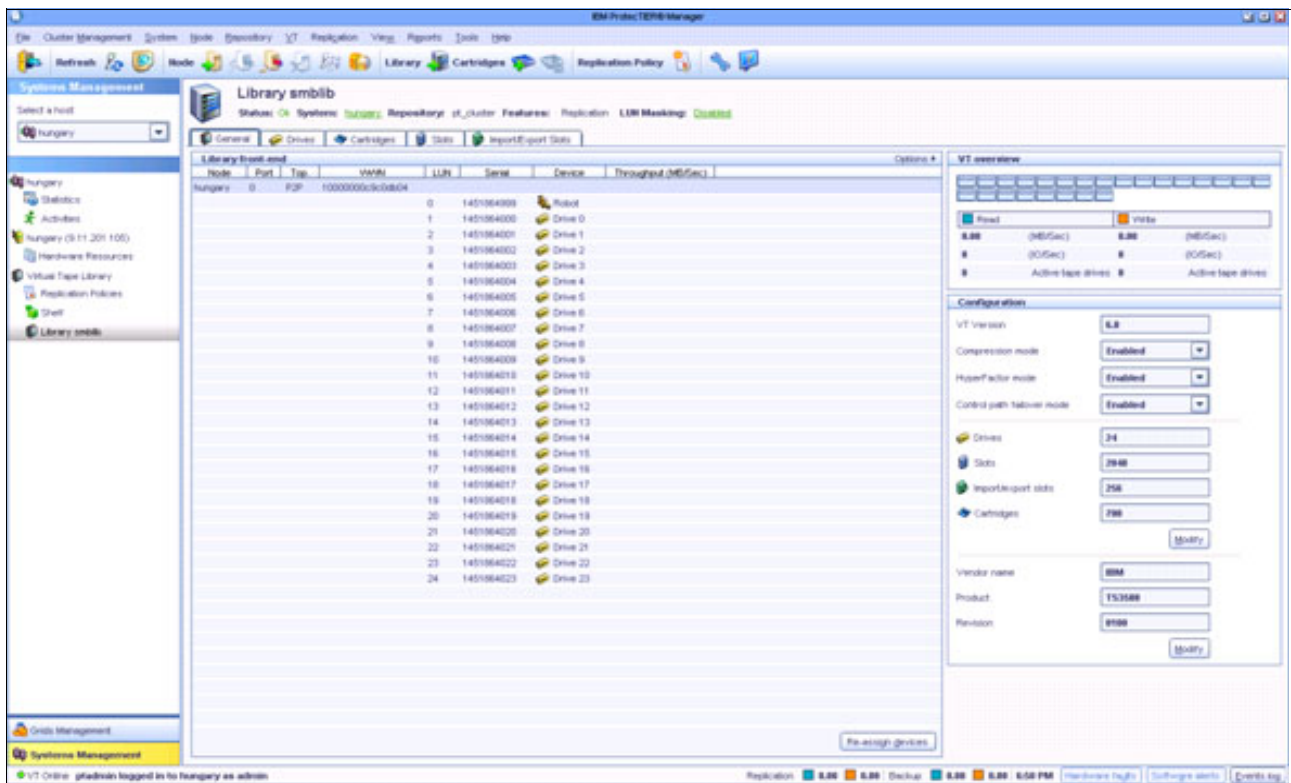


Figure 2-25 ProtecTIER Manager Library view

You may modify your configuration using the functions described in the following sections according your requirements.

## 2.10.1 Virtual libraries

For virtual libraries, you have the options shown in Figure 2-26.

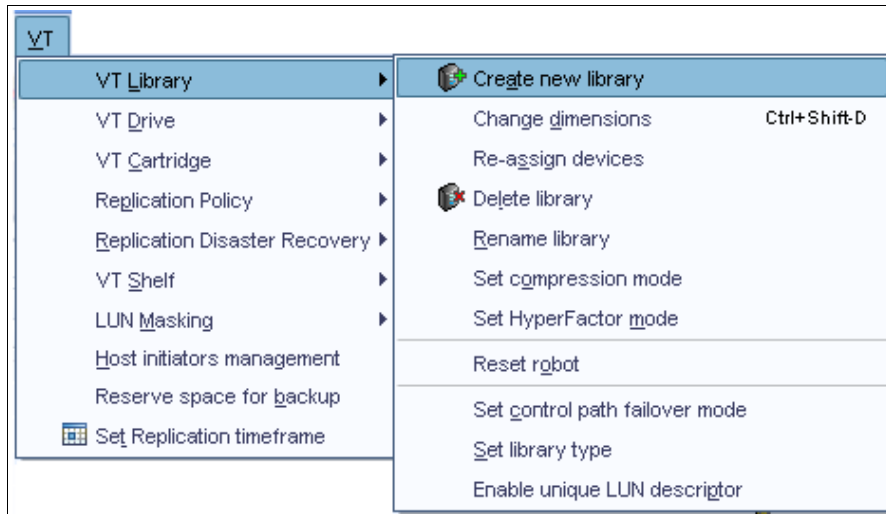


Figure 2-26 Virtual Library options

Where:

**Create new library**

Creates a new virtual library.

**Change dimensions**

Changes the number of drives, internal slots, or import/export slots. It is a disruptive task.

**Re-assign devices**

Changes the assignment of a virtual medium changer (robot) or a virtual tape drive front port assignment.

**Delete library**

Deletes an existing virtual library, including all virtual tape drives and virtual cartridges.

**Rename library**

Renames an existing library.

**Set compression mode**

Sets compression on or off. By default, compression is enabled on the library, and the ProtecTIER system compresses data.

**Set HyperFactor mode**

Sets HyperFactor mode on (default) or off. This must be used only at the direction of IBM Support. HyperFactor is the IBM technology that deduplicates data inline as it is received from the backup application.

**Reset robot**

Resets the status of the virtual medium changer.

**Set control path failover mode**

Enables control path failover (CPF) to ensure redundant paths to a robot.

**Set library type**

Sets library type to a supported virtual library.

**Enable Unique LUN descriptor**

Support for HP V11.3 hosts. This feature enables a "unique LUN descriptor".

## 2.10.2 Virtual drives

For virtual drives, you have the options shown in Figure 2-27.

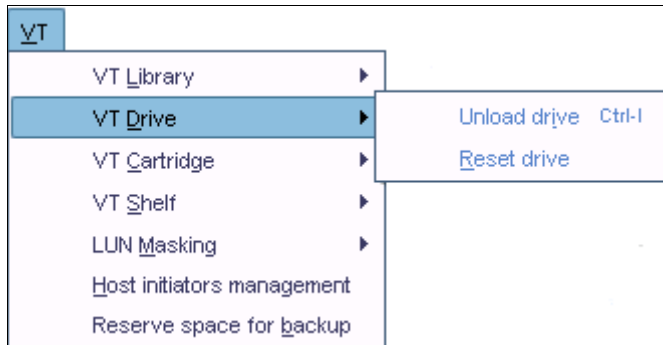


Figure 2-27 Virtual Drive options

Where:

**Unload drive** Unloads a virtual cartridge to allow the cartridge to change visibility from the virtual drive to another location.

**Reset drive** Resets the status of the drive.

## 2.10.3 Virtual cartridges

For virtual cartridges, the options shown in Figure 2-28 are available.

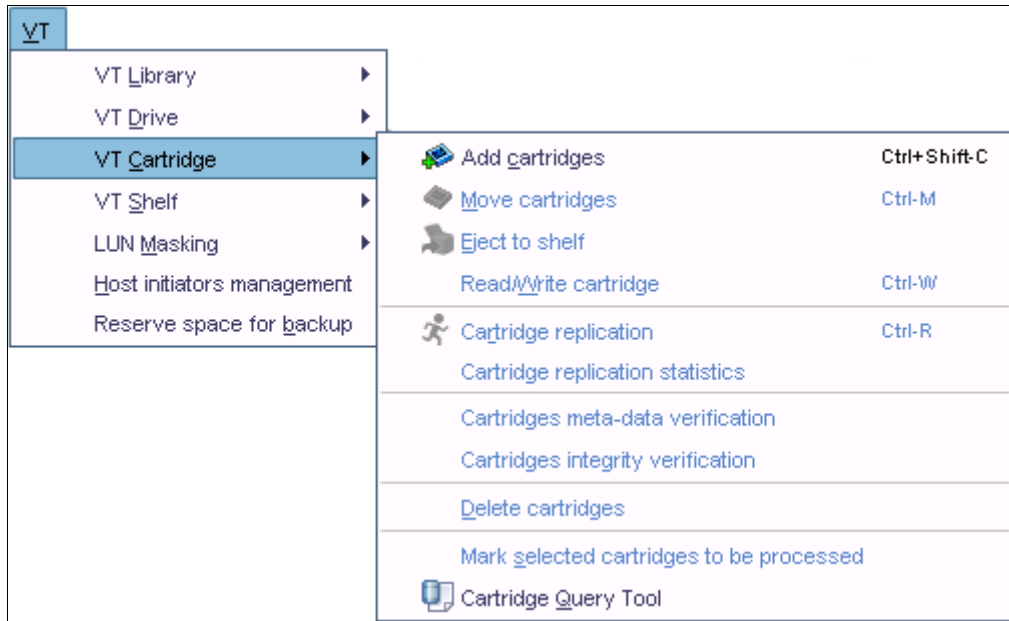


Figure 2-28 Virtual Cartridge options

Where:

**Add cartridges** Creates additional virtual cartridges.

**Move cartridges** Changes the visibility of a cartridge to a virtual drive, slot, import/export slot, or shelf (through Export slots).

	This function is normally only used at the direction of IBM Support.
<b>Eject to shelf</b>	Ejects specific cartridges to a shelf to make it available to other libraries.
<b>Read/write cartridge</b>	Sets a cartridge to read/write or read-only. This function is normally only used at the direction of IBM Support.
<b>Cartridge replication</b>	In case of a cartridge that already belongs to a policy replication, using this command forces the cartridge to be replicated immediately.
<b>Cartridge replication statistics</b>	Statistics about the percentage of a cartridge replication. Show the total amount of the cartridge that already was replicated.
<b>Cartridges meta-data verification</b>	Verifies the metadata for a cartridge. This function is normally only used at the direction of IBM Support.
<b>Cartridges integrity verification</b>	Verifies the data on the cartridge. Note that depending on the size of the cartridge and the amount of data, this task can take a significant amount of time. This function is only available on the Slots view.
<b>Cartridges ownership takeover</b>	Changing the principality of cartridges to the actual library.
<b>Delete cartridges</b>	Deleting cartridge. Data loss has to be entered to ensure that this action has been done on purpose.

ProtectTIER Manager can be run from the management console workstation or from your personal workstation.

For more information about the ProtectTIER Manager functions and windows, refer to Chapter 10, “Managing IBM System Storage ProtecTIER systems” on page 473.

## 2.11 IBM TS3000 System Console

An IBM TS3000 System Console (TS3000) (formerly the IBM TotalStorage Service Console) is required for the installation, service, and maintenance of the TS7600 family. You may use an existing TS3000 if it is within line of sight of the TS7600 family. The TS3000 provides a single point of service for gathering logs and remote support. An analog phone line is required for the auto-call feature.

**Note:** If you need a TS3000, you can order it with a member of the TS7600 family, and then they are shipped together.

## 2.12 Operating system

The required operating system used with the ProtecTIER is Red Hat Enterprise Linux 5.4 Advanced Platform. IBM System Storage ProtecTIER Enterprise Edition V2.5 software and Red Hat Enterprise Linux 5.4 Advanced Platform (RHELAP) x86\_64 are loaded on the IBM System Storage TS7600 family of products with ProtecTIER initial order.





## Part 2

# Planning for data deduplication and replication

In this part, we discuss configuration and sizing considerations and give you detailed planning information to prepare for a smooth implementation of the TS7650 and TS7650G in your environment.





## Planning for deduplication and replication

In this chapter, we provide an overview of planning, sizing, and configuring the IBM System Storage TS7600 with the ProtecTIER family of products. We consider your requirements in terms of architecture, scalability, and performance.

### 3.1 Planning for deduplication

HyperFactor is the core technology of ProtecTIER and consists of a series of algorithms that factor, or deduplicate, data efficiently (Figure 3-1). In each new backup, HyperFactor finds the data in common with previous backups. This common data in the new backup is effectively *filtered out* and pointers are used to reference existing data in the repository (see 2.6, “ProtecTIER virtual tape library” on page 32 for more details about the deduplication process). The entire contents of the new backup are stored and only new data is required to be stored, that is, a small fraction of the entire amount of the new backup data.

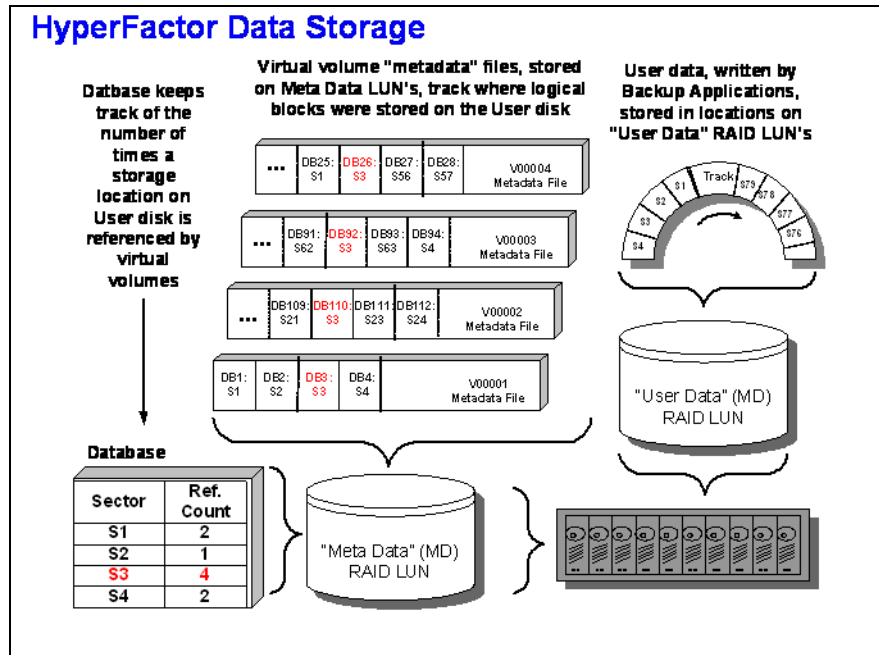


Figure 3-1 HyperFactor

The capacity of the ProtecTIER Repository is composed of the factored backup streams and the metadata that describes the factored backup streams, so it is essential to have the proper amount of back-end disk capacity as part of the ProtecTIER System configuration. The capacity reduction effect of deduplication is expressed as a deduplication ratio or factoring ratio. In essence, the deduplication ratio is the ratio of nominal data (the sum of all user data backup streams) to the physical storage used (including all user data, metadata, and spare capacity, that is, the total amount of disk storage for which the user pays).

To figure out the factoring ratio, consider having a detailed picture of your backup/recovery environment. In this chapter, we help you identify your goals and objectives and consider all the variables, and we suggest useful best practices. This chapter covers the following topics:

- Sizing inputs

Section 3.1.1, “Sizing inputs” on page 56 describes a deep analysis of your requirements, workloads, backup application, and data center topology. This data is required to be able to proceed for the next step of sizing, which is capacity sizing.

- Capacity sizing

We provide an overview of the process that is used to estimate your physical storage requirements for your current environment and your scalability needs. The capacity sizing provides a required throughput, which will be used for the following step, which is performance sizing.

► Performance sizing

Section 3.6.3, “Remote repository” on page 112 helps you understand how many metadata and user data file systems are required, based on disk technologies and RAID configurations, to ensure proper performance in terms of backup/recovery throughput.

The methods described in this chapter provide the basis for assessing and understanding how the IBM System Storage TS7600 with ProtecTIER can be fully integrated in your environment. Figure 5-2 shows a flowchart of the sizing process. Although each step is described in the following sections, IBM technical personnel or Business Partner personnel will perform the sizing process before the pre-sales and pre-installation phases.

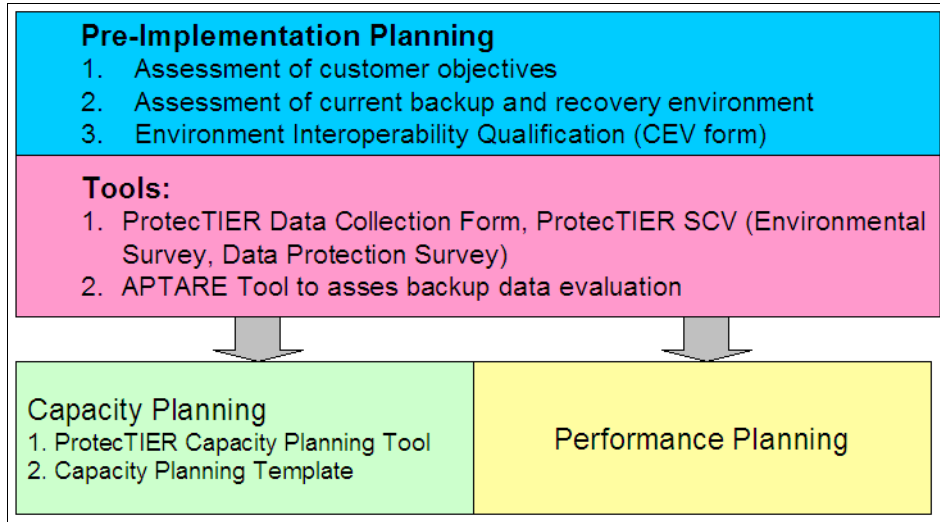


Figure 3-2 Flowchart for sizing a ProtecTIER solution

You must understand how the product itself calculates the deduplication ratio and presents and displays the data. ProtecTIER calculates the deduplication ratio by comparing the nominal data sent to the system to the physical capacity used to store the data. This information is displayed through the ProtecTIER Manager GUI and through other ProtecTIER utilities.

Table 3-1 describes general customer profiles for open systems, giving a high level overview and suggestion for a possible ProtecTIER product. Because appliances have a predefined capacity and performance, you have to make sure that your needs are covered with them.

**Note:** In the case of appliances, upgrades can be done only within the same appliance family, either TS7610 or TS7650.

Table 3-1 General customer profiles for open systems

ProtecTIER solution	Physical capacity	Performance	General profile
TS7610 Small Appliance	4 TB	Up to 80 MBps	
TS7610 Medium Appliance	5.4 TB	Up to 80 MBps	
TS7650 Appliance	7 TB	Up to 150 MBps	

ProtecTIER solution	Physical capacity	Performance	General profile
TS7650 Appliance	18 TB	Up to 250 MBps	
TS7650 Appliance	36 TB	Up to 500 MBps	
TS7650G-DD3 single node Gateway	Up to 1 PB	Up to 900 MBps	
TS7650G-DD3 dual node Gateway		Up to 1400 MBps	
TS7650G-DD4 single node Gateway		Up to 1400 MBps	
TS7650G-DD4 dual node Gateway		Up to 2000 MBps	

The performance listed in the table is based on testing performed under laboratory conditions. For the TS7650, the DS4000 disk subsystem, and for the TS7650G, the new V7000 disk subsystem was tested. The performance of the different products can show different values in different customer environments.

### 3.1.1 Sizing inputs

In this section, we discuss all the information required to assess your current environment. It is important to understand your environment because the ProtecTIER system sizing depends on many factors directly related to backup and restore policies.

#### Understanding the requirements

A key point to remember when you want to evaluate a virtual tape library (VTL) with deduplication and replication technology is to make sure that you understand your requirements and how you will be using the technology. For example, if you must support both open systems and IBM mainframe tape processing from a single VTL, your options might be more limited than if you were looking at just open systems support.

This analysis starts with the environment and must answer the following general questions about your requirements:

- ▶ Why do you want to introduce a VTL solution with data deduplication and replication in your environment? For example, are you going to use the solution to support disk-to-disk-to-tape (D2D2T) backups, to support disk-to-disk (D2D) with disk replication, or for archiving?
- ▶ What is the impact of a deduplication and replication solution on your current backup/recovery environment?
- ▶ Are you struggling with data proliferation?
- ▶ Are your applications good candidates for data deduplication or native replication?
- ▶ How many media servers do you need to attach and with how much storage capacity?
- ▶ What host interface technologies are you going to use (Fibre Channel, iSCSI, NAS, and so on)?
- ▶ Where are the potential bottlenecks in your current backup and recovery environment? How can a deduplication and replication solution fix them?

- ▶ What are your current and projected performance requirements? Performance *should* include today's backup requirements and peak throughput needs, and should also factor in the data growth and the implications of future performance needs.
- ▶ What is your estimated annual data growth rate?
- ▶ What are your expected capacity savings with a deduplication solution?
- ▶ What are the possible changes that you plan to make in your current backup architecture? For example, you might have a current backup environment similar to the following one:
  - LAN-free backup to the physical tape library for databases.
  - Disk-to-disk-to-tape (D2D2T) for file servers, web servers, mail servers, and application servers.
  - A disaster recovery solution based on remote tape vaulting by a truck.

You want to change your environment to this configuration:

- LAN-free backup to the virtual tape library for databases.
- Disk to virtual tape library for file servers, web servers, mail servers, and application servers. For file servers with small files, you might choose to perform NDMP image backups to VTL or have a backup application copy its disk storage pool to VTL.
- Disaster recovery solutions are based on remote virtual tape vaulting through replication. By greatly reducing the amount of data that is stored through the factoring process, only a fraction of the original data must be replicated to protect against disaster. With the reduction in the amount of data, the required bandwidth and disk storage is greatly minimized. As a result, IBM System Storage TS7600 with ProtecTIER provides recovery from online disks and recovery might be fast, reliable, and manageable.

After the requirements of the environment are well understood, the capabilities of a given solution must be assessed. This assessment might have two stages:

- ▶ An evaluation of the characteristics of the solution itself
- ▶ Actual testing of the system in a live environment

## Understanding the existing environment

This section describes the information necessary to understand your current backup and recovery infrastructure. An environment questionnaire is used to collect this information and assess a pre-implementation plan. After it is completed, the questionnaire may be used to determine how ProtecTIER can fit into your environment.

These questions will be asked by an IBM representative or a Business Partner representative and will help evaluate three major areas:

- ▶ **Customer environment:** Used to gather the characteristics of the backup application, the backup tiers, and the backup architecture. Because the TS7650G has an ISV compatibility matrix that specifies the supported backup software and version, it is essential to verify that your current backup software and version matches the ProtecTIER system requirements. If not, an infrastructure and economy impact analysis is required.
- ▶ **Existing backup infrastructure:** Used to examine how you back up your data, the current amount of data, the estimated annual data growth rate, and the type of technology used, for example, Ultrium Linear Tape Open (LTO) or Digital Linear Tape (DLT).
- ▶ **Disaster recovery and high availability:** Used to determine whether a ProtecTIER solution can be integrated into your current disaster recovery environment.
- ▶ **Network Replication many to many:** There is deployment of data to different data centers for safety reasons using WAN to locations in different countries.

## Environment configuration survey

In this section, we describe other important information that plays a role in deduplication and replication sizing. Other environmental information affects sizing. Your IBM or Business Partner representative requires the following information to complete sizing:

- ▶ Total amount of backup jobs per night and at peak (if you run full backup during the weekends)
- ▶ Your backup window
- ▶ When data must leave the primary site and be replicated to the DR site
- ▶ The length of time that you would like to retain backup data on disk at the local site and at the DR site
- ▶ The profile of applications that are being backed
- ▶ The analysis of the data streams generated by performing the different backups to assess the possibility of factoring
- ▶ Other unique elements of requirements

## Operating systems and hardware inventory for compatibility check

You must verify that all elements of your operating environment that will interact with the IBM System Storage TS7600 with ProtecTIER are qualified by IBM or the Business Partner to work effectively. For a complete list of supported systems, go to the following address:

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

## Media server qualification

For each media server that will connect to the IBM System Storage TS7600 with ProtecTIER, provide the data shown in Table 3-2.

Table 3-2 Characteristics of the deployed media server

Item	Media server 1	Media server 2	Media server 3
Media server operating system and version			
Backup software and version			
FC HBA model			
HBA firmware version			
HBA driver version			
Connected to ProtecTIER through loop or switch			



### Front-end fabric connectivity

For each switch that will be connected to the front-end ports (media server facing) of IBM System Storage TS7600 with ProtecTIER, provide the information shown in Table 3-3.

Table 3-3 Characteristics of the fabrics

Switch characteristic	Switch 1	Switch 2
Switch model		
Switch release		

### Storage

For each storage array connected to the TS7650Gs, provide the information shown in Table 3-4 after an accurate assessment of the disk sizing (see “Disk configurations for TS7650G” on page 81).

Table 3-4 Characteristics of the disk arrays

Item	Disk array 1	Disk array 2
Disk array make and model.		
Disk capacity on implementation.		
Number of hard disk drives (HDDs).		
Size of HDDs.		
HDDs revolutions per minute (RPM).		
Number of controllers.		
Controller cache size.		
The connection between the TS7650G and disk array is a loop or switch topology.		

### Data protection survey

Accurate capacity planning considers the behavior of each data type. A data type can be a file system, an operating system, databases, and so on. The size of one full backup is usually equal to the size of the online disk capacity. This assumption can be different if a file system that includes the operating system on one hard drive is backed up. While the operating system is running on the same hard drive as the file system, the operating system might block certain files for read access, which causes this file not to be backed up at this point in time. The *data protection survey* is a worksheet that IBM technical personnel use during capacity sizing that provides information about the number of versions, frequency, and retention for each backup. We assume that the retention of a weekly backup is associated with the retention of its incremental backup.

Important information for this survey includes:

- ▶ All the workloads that you back up in your environment.
- ▶ How much capacity is used for your current full backups to physical tape.
- ▶ How much capacity is used for the current *daily* backups to physical tape, including differentials, incrementals, and cumulative backups.
- ▶ The rate at which the data received from the backup application changes from backup to backup. This measurement has most relevance when *like* backup policies are compared. (Data change rates might range from 1% to >25%, but are difficult to observe directly.)
- ▶ How often the full backups are performed.
- ▶ How many cycles of full backups are kept.
- ▶ The relationship of how many daily backups are performed in between full backups, including differential, incremental, and cumulative backups.
- ▶ How many cycles of full and incremental, differential, or cumulative backups are kept.
- ▶ Whether a monthly full backup is kept for longer periods than the regular weekly full backup.
- ▶ How much data do you like to replicate to another location?
- ▶ What operating systems are backed up and what is the path, for example, C:\ for Windows systems or root for AIX or Linux based systems.

The information in Table 3-5 is relevant for IBM Tivoli Storage Manager users with an *incremental forever* policy for some of their data. IBM Tivoli Storage Manager uses a different paradigm that is known as *Progressive Incremental*. In this case, an initial full backup is taken, then all future backups (even those on the weekends) are considered incremental, so there is no full backup on a weekly basis. IBM Tivoli Storage Manager uses a much more sophisticated and intelligent way to perform backups, because only new or changed files are backed up. IBM Tivoli Storage Manager is empowered to do this type of backup because of its relational database that tracks each individual file and knows exactly what your computer's state was on each day. When a restore is required, just the version of the file needed is restored.

In addition to the information listed above, provide the information in Table 3-5 if you are using IBM Tivoli Storage Manager as your backup application.

Table 3-5 IBM Tivoli Storage Manager backup policies

Data type	Primary disk size held by this data type (or all)	Policy: Full+ Incremental or Incremental Forever	Versions data exists	Version data deleted (optional)	Retain only versions (optional)
Databases					
Small files					
IBM Lotus® Domino®					
Flat files					
Data type 5					
OS type					

A good way to get a real overview of the customers backup data is by using *APTARE*. This software tool can be installed in the customer environment to collect and analyze the different data types, which provides a measure about the amount of data backed up and what different types of data are stored to better assess the data for deduplication. All sources, types of data for backup, and the estimated amount of data for the different host systems to back up must be estimated and listed.

## 3.2 APTARE overview

The APTARE Data Collector Tool is intended to help ProtecTIER sales and technical teams enhance the collection of current backup information from an environment for use in ProtecTIER Capacity Planning. The tool extracts current and historical backup information from the backup servers (existing backup job details and tape information), which enables the improvement of pre-sale sizing/planning activities. After collecting the data, the sales and technical team will be able to better size a ProtecTIER solution to your environmental.

The APTARE StorageConsole Data Collector is a centralized and remotely managed data collection mechanism. This Java application is responsible for interfacing with backup servers to gather information related to storage backup/recovery.

The Data Collector continuously collects data and sends this data, using an http or https connection, to the APTARE Data Receiver. This Data Receiver runs on the APTARE Portal Server and stores the data it receives in the Reporting Database.

The Data Collector obtains all of its monitoring rules from a Data Collector Configuration File. Passwords configured for policies in the Data Collector are encrypted prior to insertion into the database. They are decrypted in local Data Collector memory immediately prior to use. This configuration file resides in the Reporting Database in XML format. When the Data Collector first starts, it downloads this file from the Reporting Database. The Data Collector uses this file to determine the list of backup servers that are to be monitored and included in its data collection process.

APTARE provides storage management solutions for an end-to-end view of your storage environment, enabling you to increase data reliability and reduce costs. APTARE delivers Data Protection Management, a backup manager for reporting on your backup environment, which provides an enterprise-wide view of your backup environment through a single portal.

APTARE can be deployed securely in a variety of operational environments, including a private LAN, virtual private network (VPN), corporate intranet, and even the public Internet. The security features for APTARE communications (for example, between web client and server) can be customized for your operational environment. APTARE can be configured to provide an end-to-end security context between the user's browser and the APTARE Portal Server, which provides security, privacy, and user authentication, and no repudiation.

### 3.2.1 APTARE StorageConsole Backup Manager

The APTARE StorageConsole Backup Manager provides comprehensive reporting of major backup software environments. The Backup Manager Data Collector interfaces with each of the supported backup and recovery software systems to extract metadata about the underlying backup and recovery environment, such as backup job details and tape inventory information. The Backup Manager Data Collector can run on any stand-alone server, the Portal Server, or any backup server for all backup solutions. In a Veritas NetBackup environment, the Backup Manager Data Collector must reside on each Veritas NetBackup Master Server.

### 3.2.2 APTARE architecture

APTARE has a hub and spoke architecture with two main components:

- ▶ **Portal:** A centralized location (the underlying database is embedded and includes licenses and so on for Oracle 10g).
- ▶ **Data Collector:** For APTARE Backup Manager, the Data Collector interfaces with the underlying Backup Vendor Products (IBM Tivoli Storage Manager, NetBackup, Legato, HP Data Protector, or Backup Exec) using the standard published interfaces to that backup product. Likewise, for APTARE Capacity Manager, the Data Collector uses storage array mechanisms specific to each vendor's storage system (IBM, HDS, EMC, and NetApp). In addition, host resources data is acquired. The data is parsed and packaged into Java objects, serialized into an HTTP or HTTPS data stream, compressed and sent over a network through port 80 (http) or port 443 (https) and inserted into the underlying Aparte Portal database.

Figure 3-3 shows a view of the APTARE architecture.

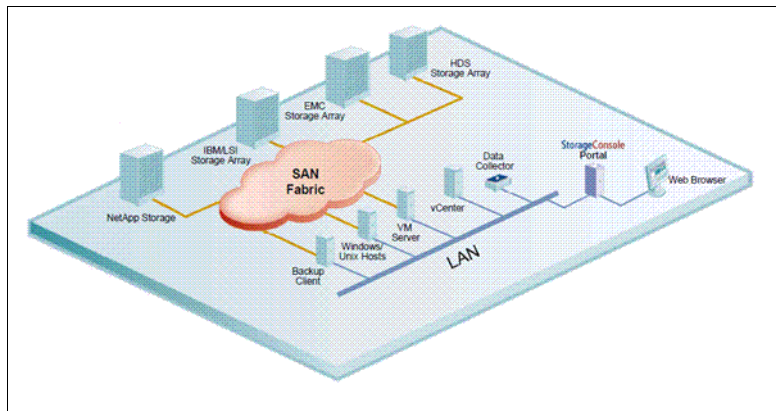


Figure 3-3 APTARE StorageConsole architecture

Table 3-6 shows the supported operating systems for Data Collectors.

Table 3-6 Supported operating systems for data collectors

IBM AIX V5.3 or later (JRE 1.6 required)
HP-UX 11i (JRE 1.6 required)
CentOS 4.0 and 5.0
Red Hat Enterprise Linux 4 & 5
Solaris x86, Solaris 8, 9, and 10 (SPARC)
SUSE Linux Enterprise 9 and 10
Windows 2003 and 2008

**Note:** For a more detailed configuration of the APTARE portal interface, go to the following address:

<http://www.aptare.com>

### 3.2.3 Managed backup environment security

We now discuss some of the major backup applications and the environmental security associated with them.

#### IBM Tivoli Storage Manager

For each Tivoli Storage Manager instance, the Backup Manager Data Collector establishes connections to the database using the **dsmdmc** command. The Data Collector Configuration file contains all the connection information for each Tivoli Storage Manager Instance, including such parameters as the Tivoli Storage Manager user name and password for login, the Tivoli Storage Manager Instance name, the IP address of the Tivoli Storage Manager Host Server, and the Tivoli Storage Manager Port. The Backup Manager Data Collector passes various QUERY and SELECT commands through **dsmdmc** to obtain its information from each separate Tivoli Storage Manager Instance. The information is then sent by way of http(s) to the Portal.

#### Veritas NetBackup

The NetBackup Agent is lightweight and nondisruptive. Essentially, the only task it performs directly is to parse, serialize, compress, and transmit data returned from the NetBackup commands. All Agent configuration information is maintained in local configuration files on the master server. All communication between the master server and the portal is initiated by the Agent. The portal will never try to make a connection with the master server, thus protecting the security of the master server.

#### Symantec Backup Exec

For each Backup Exec server, the Backup Manager Data Collector establishes connections to the Backup Exec database. The connection information for each Backup Exec server is retrieved from the Portal or from a locally stored, encrypted file. This connection information includes parameters such as the administrator user name, domain name and password, and server host name or IP address. The Data Collector passes database commands through TCP/IP to obtain its information from each Backup Exec server. The information is stored in the Portal database, enabling a global view of all of the backup servers and clients.

#### EMC Legato NetWorker

For each Legato Server, the Backup Manager Data Collector establishes connections to the database using the **nsradmin** command. The connection information for each Legato NetWorker server is retrieved from the Portal or from a locally stored encrypted file. This connection information includes parameters such as the administrator user name, domain name and password, and server host name or IP address. The Data Collector uses the command-line utilities such as **mminfo**, **nsradmin**, and **nsrinfo** to obtain its information from each Legato Server. The Data Collector also uses **ssh** to connect to remote Legato Servers to retrieve log file details. This information is stored in the Portal database, enabling a global view of all of the backup servers and clients.

#### HP Data Protector

The Backup Manager Data Collector establishes connections to the HPDP Cell Manager Server and collects data using the **omnicellinfo**, **omnicc**, **omnirpt**, and **omnimm** commands. The Data Collector Configuration file contains all the connection information.

### 3.2.4 APTARE setup

This section contains a brief description of a Tivoli Storage Manager setup. For more information about other operational system setup, contact your IBM Service representative.

To install APTARE Data Collector tool at Tivoli Storage Manager, you must download the tool from the following address:

<http://www.aptare.com/agent80/install.htm>

Choose the right platform and follow the instructions.

**Note:** Data Collector can be installed on the Tivoli Storage Manager server or another server that has the dsmdmc setup.

#### Preinstallation questionnaire

The following information must be acquired before installing Data Collector:

- ▶ Customer information
- ▶ Tivoli Storage Manager host server(s)
  - Host name
  - IP address
- ▶ Tivoli Storage Manager Instances
  - Instance name and server port number
  - User name / password
  - The home directory where dsmdmc is installed
- ▶ Internet access through port 443 to aptareagent.lotus.com

You must download the installer onto the data collector server(s).

#### Portal setup

The following steps must be completed:

1. Create a server group under “eval”. Use the customer’s name (no spaces).
2. Add the Tivoli Storage Manager host server(s) under the new server group.
3. Create a domain. This is *important*. Use the customer’s name (no spaces).
4. Create a new user.
5. Create a data collector policy. Do not forget to select the customer’s domain.
6. Create a Tivoli Storage Manager policy. Do not forget to select the customer’s domain.

**Note:** Make sure these steps are done prior to the data collector installation.

Figure 3-4 on page 65 to Figure 3-9 on page 67 shows the steps of the portal setup.

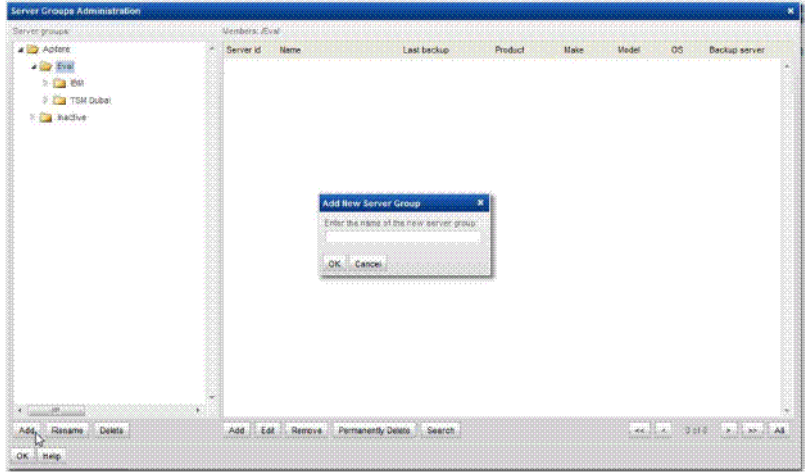


Figure 3-4 Create a server group

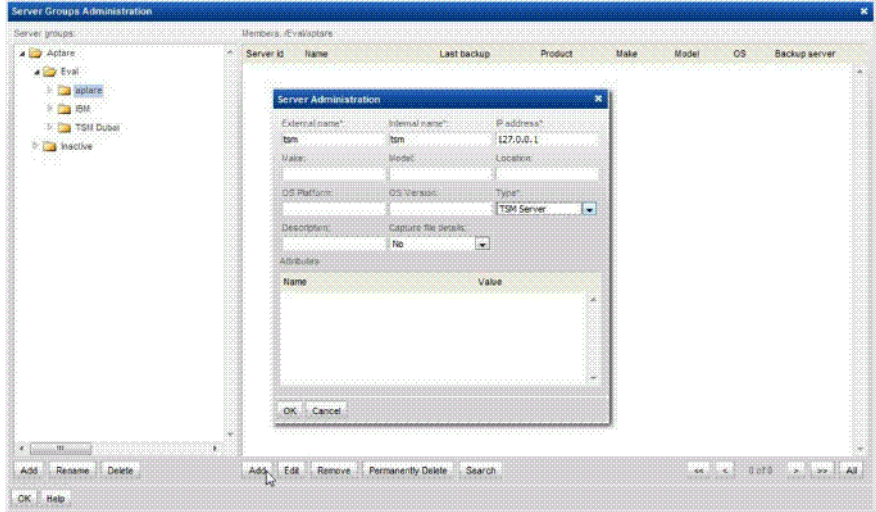


Figure 3-5 Add a Tivoli Storage Manager host server

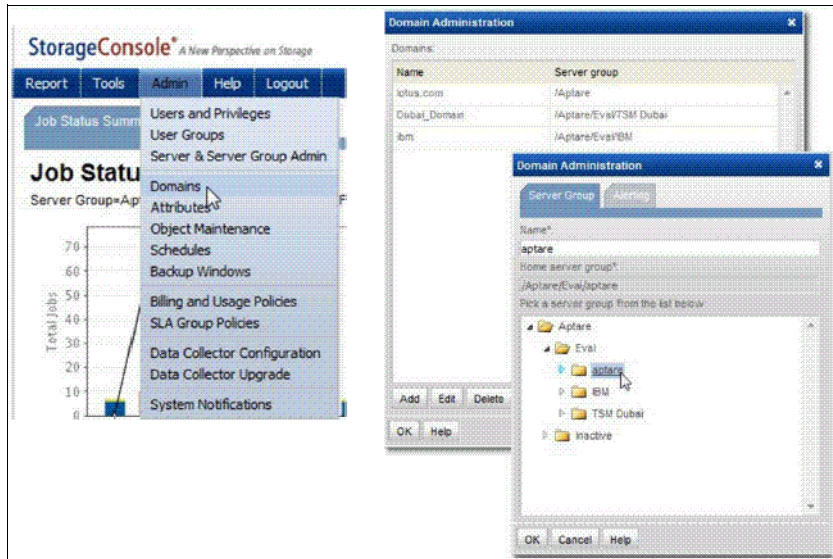


Figure 3-6 Create a domain

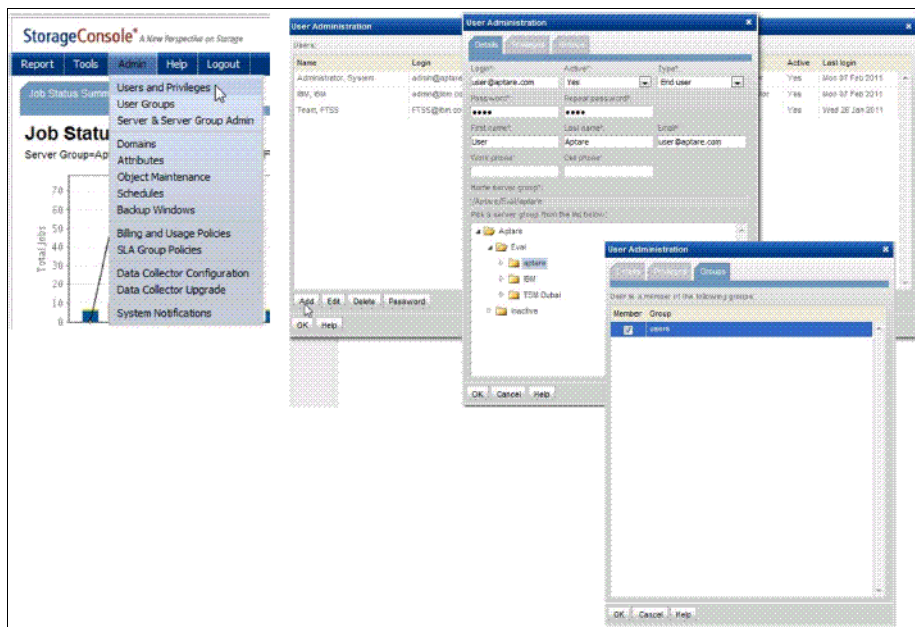


Figure 3-7 Create a new user



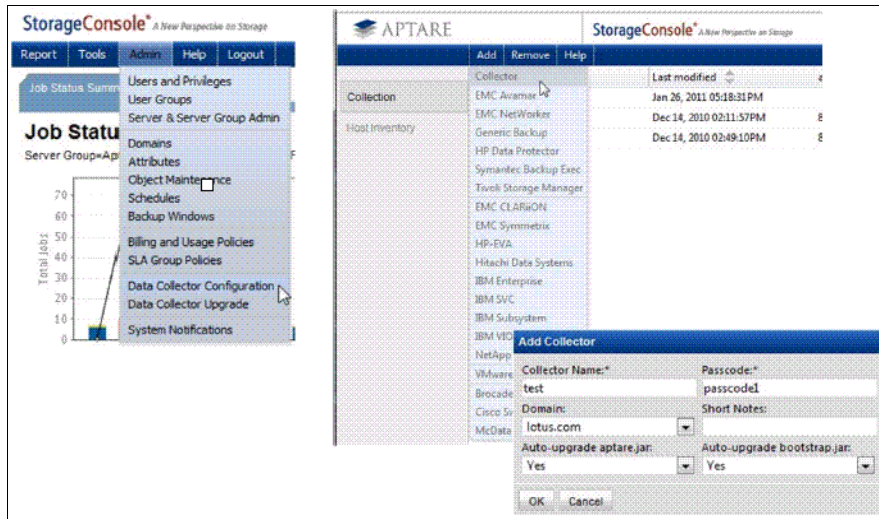


Figure 3-8 Create a data collector policy

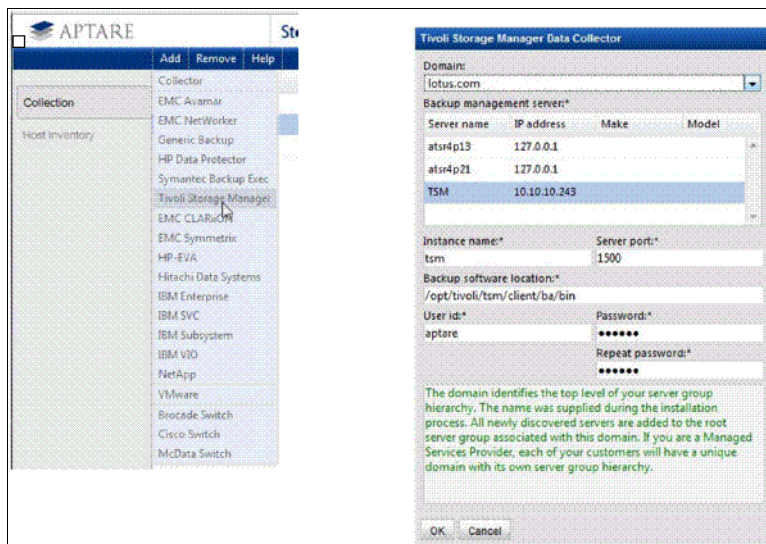


Figure 3-9 Create a Tivoli Storage Manager policy

## Data collector installation

To install the data collector, complete the following steps:

1. Log onto the data collector server as the admin or root user.
2. Depending on your operating system, run the following executables:
  - UNIX: `sh ./aptareagentinstall.bin`
  - Windows: `aptareagentinstall.exe`
3. Follow the instructions and complete the following entries:
  - Data collector name: Use the one from the portal setup.
  - Passcode: Use the one from the portal setup.
  - URL: Use `https://aptareagent.lotus.com`.
4. Execute `checkinstall.sh|bat`.
5. Start the data collector.

## Data collector uninstallation

To uninstall the data collector, complete the following steps:

1. Log onto the data collector server as the admin or root user.
2. Depending on your operating system, stop the data collector by completing the following tasks
  - UNIX:
    - i. Run:

```
find /etc -name aptare_agent -print | xargs rm.
```
    - ii. Run:

```
rm -rf /opt/aptare/installlogs /opt/aptare/jre  
/opt/aptare/lib/opt/aptare/mbs /opt/aptare/UninstallerData
```
  - Windows: Select **Start** → **Programs** → **APTARE StorageConsole Agent** → **Un-install APTARE StorageConsole Agent**.

## Portal side uninstallation

To uninstall the Portal, complete the following steps:

1. Delete the Tivoli Storage Manager policy, data collector, customer domain, and user.
2. Move the server group to “inactive”.

## 3.2.5 APTARE reports

After the data collector runs, you will be able to download all the reports that are set. Those reports can be sent to your IBM representative to be used as input data for a ProtecTIER Planner tool. Those reports help the IBM representative understand the backup requirements and develop a better customer solution by using the ProtecTIER Planner tool.

The ProtecTIER Planner tool is used by ProtecTIER specialists (FTSS, ATS, and BP) to calculate the required capacity and performance of a TS7600 solution.

This tool is an interactive modeling tool using MS Access that allows the input of the key variables required to determine the TS7600 capacity.

**Note:** All information in this section about APTARE can be used by an IBM technical representative. For ProtecTIER Planner, an IBM technical representative *must* be contacted.

## 3.3 Throughput considerations

The IBM System Storage TS7600 with ProtecTIER is a virtual tape library with enterprise scale in-band factoring, which means that all data reduction occurs in real time as the backups are running. The in-band factoring approach has many advantages, but also requires the appropriate level of hardware and proper configurations to achieve optimal performance. Properly configured and based on DD4 hardware, a single Enterprise-level IBM System Storage ProtecTIER node is capable of achieving sustained throughput rates of up to 1400 MBps in live production environments. Using a two-node clustered configuration, the IBM System Storage TS7600-DD4 with ProtecTIER can achieve sustained throughput rates of up to 2000 MBps. The actual performance that any given environment achieves depends on several variables that we cover in this section.

The purpose of this section is to discuss performance considerations that can impact throughput performance, measured in megabytes per second (MBps), when testing and deploying the IBM System Storage ProtecTIER System.

The following components play a role in the overall system throughput that ProtecTIER can achieve:

- ▶ Attached host systems
- ▶ SAN connectivity
- ▶ Disk array
- ▶ Data type (also called backup policy)
- ▶ Replication activities

For each component, we list the best practices for optimal performance.

### 3.3.1 Attached host systems

From the host system, the backup data is delivered to the attached backup device. Depending on the host hardware and configuration of internal or external disks, the performance can vary. Consider the following items:

- ▶ Different hosts can have different basic hardware performance. Consider the type of the host that performs a backup, for example, a PC does not perform like a server.
- ▶ Consider what disk type the host is reading the data to be stored to the backend device. If the host internal disk drive is a SCSI standard, the achievable performance might be low, while a modern SATA delivers the data to be stored much faster.
- ▶ The external attached types of disk subsystems to a host should be considered too, depending on how many hosts might share the external host disk subsystem where the data should be read from for backup.

### 3.3.2 SAN connectivity

For the best SAN connectivity:

- ▶ Make sure that the fabric switches are up to the latest firmware revision of their operating system (contact the manufacturer or reseller).
- ▶ IBM System Storage TS7600 with ProtecTIER front-end ports should not be in a zone with any other IBM System Storage TS7600 with ProtecTIER front-end ports.
- ▶ If possible, dedicated Host Bus Adapter (HBA) ports in the backup server should be zoned to single IBM System Storage TS7600 with ProtecTIER and its front-end ports.
- ▶ Ensure that the IBM System Storage TS7600 with ProtecTIER front-end ports are in a separate zone than the IBM System Storage TS7600 with ProtecTIER back-end ports to avoid reservation conflicts on the disk subsystem acting as the repository.
- ▶ Ensure that Inter-Switch Links (ISL) between switches, connected to an IBM System Storage TS7600 with ProtecTIER ports and backup servers or storage arrays, are not oversubscribed and provide the needed bandwidth for optimal operation.
- ▶ Use 8 Gbps HBAs for the TS7650G back-end and front-end connections.

### 3.3.3 Disk array

A critical hardware component in a ProtecTIER implementation is the disk array that holds the ProtecTIER Repository (Figure 3-10). The repository is the physical disk that holds the ProtecTIER *HyperFactored* data. There are two types of file systems that make up the ProtecTIER Repository:

- ▶ Metadata
- ▶ User data

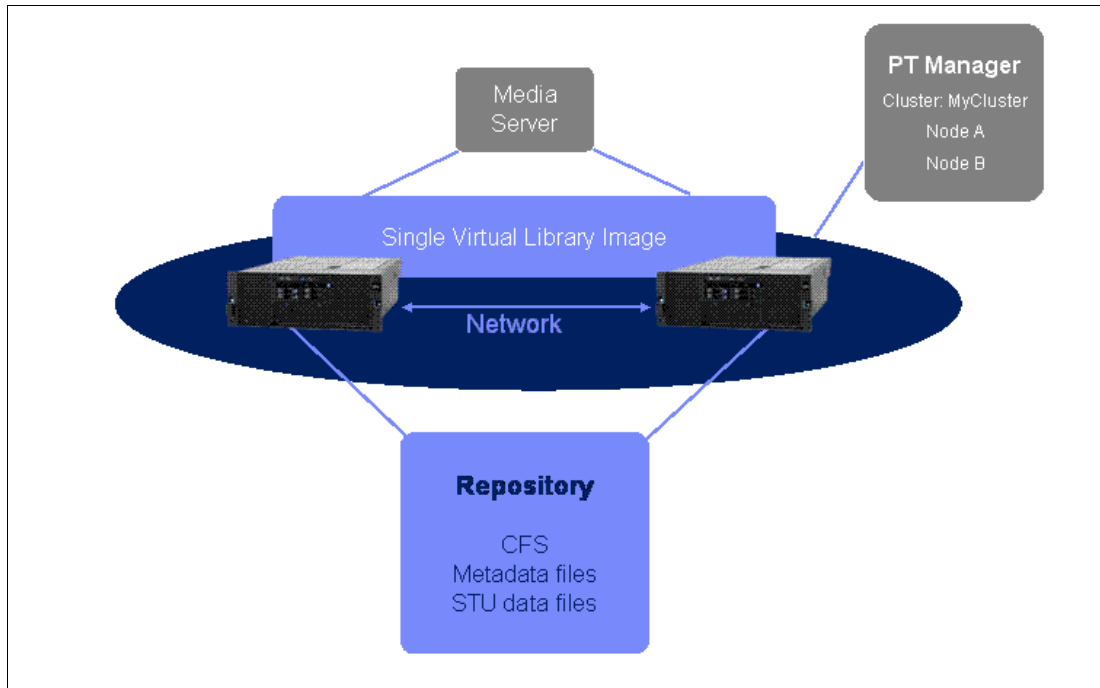


Figure 3-10 Repository metadata and user data

Metadata file systems store all aspects of the data that is backed up and cataloged, but not the data itself, whether it requires new disk space or not, which means that every already available pattern of data must be referenced as well as every new pattern of data stored it to the repository. These metadata references are small in size, which can impact the amount of I/O to the storage subsystem.

Figure 3-11 shows the relationship of the metadata file system size and the amount of spindles grouped to a RAID array in regards to the factoring ratio and performance.




Figure 3-11 Metadata file system size and factoring ratio impact performance

The graph implies a linear behavior. The higher the factoring ratio, the more references are stored in a shorter period of time. Therefore, more bandwidth for the disk subsystem is needed, which can be provided by providing the metadata raid arrays with more spindles at the initial setup to provide optimal physical performance capabilities.

It is critical that the performance of the metadata file system be optimal. In general, we recommend RAID 10 RAID groups (4+4, 6+6, or 8+8 disks) for the metadata file systems. To improve performance, more than one metadata RAID group could be defined at the initial setup, depending on the customer performance needs.

The user data file systems store the actual odds of data backed up or referenced by new generations of the data. To provide the needed bandwidth (the user data RAID groups represent the repository of the TS7650), the amount of disks used depends on the needed size of the repository to hold all the customers backed up data.

Figure 3-12 shows how the performance is impacted by the amount of disk spindles arranged in a RAID array and the factoring ratio. Compared to Figure 3-11 on page 70 the factoring ratio inversely impacts the performance for user data raid arrays. The lower the factoring ratio, the more I/O there is for the user data file systems. Therefore, the user data raid groups should provide more performance than can be provided by more parallel disk spindles.

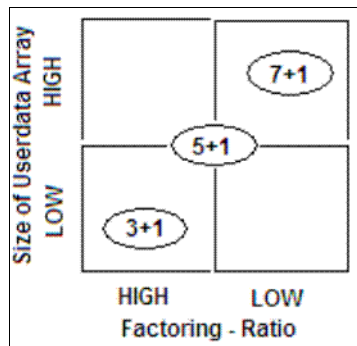


Figure 3-12 User data file system size and factoring ratio impact performance

The ProtectTIER software uses all the available RAID groups randomly to optimize performance based on the available RAID array groups that hold the internal file systems of the repository. If the system design is not balanced for performance and factoring, either the design of the metadata or user data may limit performance.

The expansion of a repository at a later point in time than the initial setup and configuration must not imply changes in performance. Only the metadata file systems can be extended later. The user data file systems cannot be extended; only RAID groups for user data can be added.

**Note:** The configuration of the disk array is the variable that has the *greatest impact* on overall system performance.

Tuning the array for the unique ProtecTIER I/O pattern is critical. ProtecTIER is *random read* oriented. Eighty to ninety percent of I/O in a typical ProtecTIER system environment is random reads. Therefore, any storage array deployed with ProtecTIER should be optimized for this I/O pattern. In all cases, the disk array manufacturer (or reseller) should be consulted to determine the best tuning and configuration parameters for the particular array being deployed. The user data LUNs should be tuned for a random-read intensive workload, while the metadata LUNs should be tuned for a random-write intensive workload. In 3.3.7, “Storage sizing” on page 80, we describe implementation considerations related to the disk array configuration.

### 3.3.4 Data type

The other factor that affects performance in a ProtecTIER system environment is the data that is being targeted for backup. Some data, such as databases and Lotus Domino email, is highly compressible and also factors quite well. Other data, such as video or seismic data, cannot be compressed or factored well. Also, the various backup applications have features, such as encryption or multiplexing, that can also affect the ProtecTIER factoring ratio and performance. The type of data that ProtecTIER systems are *HyperFactoring* can affect both the factoring ratio and system performance. This section also includes general system testing and benchmarking considerations. See Table 3-7 for a sample composition of data types.

Table 3-7 Sample composition of data types

Data type	Example
Production Database	IBM DB2®
Data Warehouse	IBM Informix®
Email	IBM Lotus Notes®
File & Print Services	Windows or UNIX

Here is a list of software, settings, and tools known to impact the structure of a backup data stream for deduplication. Use the list to evaluate your environment and to discuss your requirements with your IBM technical sales representatives.

- ▶ DB2 compression
- ▶ Tivoli Storage Manager Incremental Forever
- ▶ Lotus Notes compression/DAOS
- ▶ Client/Server multiplexing
- ▶ Client-side compression
- ▶ Oracle RMAN
- ▶ SQL LightSpeed
- ▶ VMWARE vRanger

**Tip:** Accurate planning assesses the capacity and performance requirements for each data types separately.

Consider the following items:

- ▶ If multiple backup streams of the same data set are used for testing, a single copy of the data should be backed up first to populate the repository.

- ▶ If encryption features of the backup application are turned on, the factoring ratio and performance of these data sets will degrade drastically.

**Note:** Always encrypt last. Deduplicating encrypted data is ineffective. Compressing encrypted data can decrease security. Drive-level encryption has no performance impact, and ensures that encryption is the last action performed.

- ▶ Backup application or database backup programs should disable compression. Compression is common with SQL database backups using LiteSpeed. Data should be sent to the ProtecTIER servers uncompressed, or the factoring ratio for this data will be low. ProtecTIER can manage multiple VTLs, each with its own configuration. For compressed data streams, create a new ProtecTIER VTL with compression turned off. Compressing data a second time can cause data to expand, so compressed data should be segregated in ProtecTIER whenever possible.
- ▶ Multiplexing features of the backup application or database backup tool (RMAN or similar) should be disabled. The backup application should send only one stream to each virtual tape drive. Because ProtecTIER systems allow up to 256 virtual tape drives per node for the IBM System Storage TS7600 with ProtecTIER, the system can process many streams.

For example, before IBM System Storage TS7650, TS7610, and TS7650G with ProtecTIER, RMAN sent nine streams (three file streams to three real tape drives each). With the IBM System Storage TS7650 and TS7650G, the parameters should be adjusted to send one file stream each to nine virtual tape drives, which does not affect the database backup speed (it might improve). Another kind of multiplexing could be generated by backing up a client's internal hard drive that holds, besides the file systems for the different users, the operating system to operate the hardware. All operating systems have different sections where data changes more or less. Infrequent changes in the system configuration may occur; more often, changes in the user data area are to be expected and continuous change can be seen in the system and application logs. While a system is continuously operating, the operating system prevents its internal files from being read by the backup application. Because the backup of such a system always starts at the root of the file system tree and continues to the next file to back up, a file might be skipped because read access is blocked by the running operating system. This situation changes the content of the data stream to the backup device. A backup of such a system internal hard drive while it is operating is expected to lead to a reduced factoring ratio.

- ▶ In evaluation or test scenarios, do not create big files that repeat the same data. In addition, try to avoid copying the same files over and over to create multi-stream tests (for example, using the same 100 MB to create a 40 GB file and copy this file several times to allow multi-stream tests). Although this increases the factoring ratio, it reduces performance, as all data is referenced to the same blocks.
- ▶ For testing, do not use **dd** commands on a directory that you duplicated many times. Similar data hits to the same RAID group always will use the same RAID group to read.
- ▶ You can check the storage array read performance capability by using a **dd** command to read from disk to /dev/null on UNIX, for example:
 

```
dd if=/vol1/backup.txt of=/dev/null
```
- ▶ **vmstat** and **iostat** are two system commands that can be helpful when checking the performance of the ProtecTIER systems:
  - If **vmstat** shows that the CPU is waiting, this means that the disk is not fast enough to service the I/O queues.
  - If **vmstat** shows that the CPU system and user is reaching 90 - 95% capacity, then we might be on the edge of the CPU's capability.

- In optimal conditions, **vmstat** should show the CPU utilization parameters of CPU idle 0, CPU system and user are high, but not exceeding 85%, and CPU wait is 0.
- **vmstat** can show you whether you are in the READ or WRITE state. If mostly READ, it means that you are already in the factoring process, and if mostly WRITE, then you are writing your new data or baseline.
- **iostat** shows you the activity on each LUN (for example, /dev/sdX) (Figure 3-13). If while doing the tests not all of your /dev/sd devices are active in READ, then probably not all of the RAID and disk spindles are participating. You have optimized your system if, during backup, all of your RAID group are active.

Device:	tps	Blk_read/s	Blk_wrtn/s	Blk_read	Blk_wrtn
sda	3.50	1.96	124.92	1356690	86443930
sdb	9.32	11.61	103.49	8034030	71613632
sdc	0.37	6.29	3.86	4354682	2670701
sdd	0.51	5.71	62.61	3951658	43330396
sde	0.00	0.00	0.00	1373	0
sdf	0.00	0.00	0.00	1373	0
sdg	0.36	6.20	3.80	4293030	2630861
sdh	0.00	0.00	0.00	1373	0
sdi	0.00	0.00	0.00	1373	0
sdj	0.36	6.20	3.80	4293030	2630829

Figure 3-13 Example IOSTAT

- ▶ The first backup baseline is important. At the first backup, all available RAID groups must be configured and reachable. To establish the system baseline correctly, use the same data sources used for backup. The system is designed for large backup streams. The baseline is the minimum amount of references needed to find similar data. If the data at the first time is different than the data received on later backups, the system needs to change the references that can impact performance and the factoring ratio.
- ▶ Small files do not factor as well as larger files. For best system performance, in test or production, *at least* 24 data streams of backup data sets should be run at the same time, which takes advantage of the fact that ProtecTIER systems can process 24 storage units at the same time. The storage unit is one of the four allocation entities that ProtecTIER systems use for abstracting the physical disk to provide contiguous logical space when actual physical space is fragmented.

Other options to improve factoring for small files (less than 32 KB) would be to:

- ▶ For files residing on NAS boxes, perform NDMP image backups and send the backups to ProtecTIER.
- ▶ File level backups should first back up to the backup application Disk Storage Pools, and then Disk Storage Pool can be copied to ProtecTIER.

**Tip:** To outline the best possible performance, use as many as possible virtual tape drives for the backups. IBM System Storage TS7600 with ProtecTIER is optimized for a large number of virtual drives and heavy workloads of backup data.



### 3.3.5 Local repository sizing

In this section, we discuss the importance of capacity sizing. This process requires a pre-sales engagement with IBM technical personnel. The use of IBM internal tools and questionnaires must be used to correctly size the IBM System Storage TS7600 with ProtecTIER through a comprehensive discussion with you about your data protection environment, requirements, and business objectives.

In regards to the performance of the repository sizing, the maximum throughputs of which the configurations are capable are kept in mind. Our example includes the replication activities of the system to be sized, which means the used disk subsystem representing the repository is shared for local backup and replicated data. Remember to save enough space for local backup by using the ProtecTIER manager. With ProtecTIER Version 3.1, replication activities can impact the system's performance for local backup when running in parallel, and the attached disk subsystem cannot provide the I/O bandwidth required for this workload. Using IBM Storwize V7000 as a repository sized for the I/O load to handle this impact only provides a minimum negligible performance value, as verified in lab assessments. A different disk subsystems hardware architecture could be used, which can result in different values of performance.

The disk subsystem of choice *must* be able to process I/O access for local backups, replication data received, and replication data sent, in addition to the I/O generated by possible data expiration caused by the backup application.

The specifications below are based on a realistic customer workload, assuming properly configured back-end disk arrays for the repository.

You can perform repository sizing through the planner tool for pre-sales and also by using the ProtecTIER Manager Create repository planning option. Replication repository sizing is not described in this section; refer to 3.5.2, "Deployment planning guidelines example" on page 105 for more details about that topic.

Note that:

- ▶ A single node DD4 server is capable of 1400 MBps of I/O activity for backup and while replicating as a target system.
- ▶ A two-node DD4 cluster is capable of 2000 MBps of I/O activity for backup and while replicating as a target system.

The following examples demonstrate the calculation of the required performance of the ProtecTIER system under different scenarios, that is, scheduled mode and continuous mode of replication operation:

- ▶ Example A (scheduled replication mode of operation)
  - There is backup activity running for 10 hours a day at a 500 MBps ingest rate.
  - The spokes replication activity is running in a separate time slot of 12 hours a day at 500 MBps to receive data.
  - The repository should support a sustained rate of 500 MBps.
- ▶ Example B (Replication activity runs in a continuous mode of operation.)
  - The backup activity is running 24 x 7 at an ingest rate of 1300 MBps.
  - Replication runs concurrently and in parallel with the backup activity at 600 MBps.
  - The repository should support a sustained rate of 1900 MBps.

- ▶ Example C (Replication activity runs in a continuous mode of operation on a many-to-many configuration; during local backup, it is sending and receiving data from replication.)
  - The backup activity is running 24 x 7 at a ingest rate of 1100 MBps.
  - Concurrently, the hub is replicating data to a spoke at a ingest rate of 400 MBps.
  - Concurrently, the hub is receiving replication data from a spoke at a ingest rate of 500 MBps.
  - The repository should support a sustained rate of 2000 MBps.

The correct sizing of the ProtecTIER Repository is important, because it:

- ▶ Enables you to select the correct disk product
- ▶ Enables you to purchase the correct amount of disk
- ▶ Keeps backup operations running smoothly

Capacity sizing might rely on estimates and forecasts. It appears that business growth develops in an unforeseen positive way, causing the amount of backup data to grow higher in later operations, as assumed in your planning. Therefore, there is always a margin of error in the estimates and you should plan for additional capacity to compensate for this margin. Adding 10% to your estimated capacity is a best practice.

Table 3-8 shows the capacity sizing terminology that we use in our examples.

Table 3-8 Capacity sizing terms

Term	Definition
Nominal capacity	The amount of user data that the ProtecTIER system is protecting.
Physical capacity	The physical capacity used in the array.
Factoring ratio	The ratio of nominal to physical capacity.
Data change rate	The rate at which data received from the backup application changes from backup to backup. This measurement is more relevant when <i>like</i> policies are compared. Data change rates can be from 1 - 25% and are best calculated with tools like the IBM TPC for Data product.
Data Retention period	The period in time (usually in days) that defines how long customers keep their disk-based backups online. Retention periods on average are 30 - 90 days, but can be longer depending on business and government regulations.

All the information required in the data protection survey is fundamental to sizing the solution, and the data change rate is the most important variable in determining the size of the ProtecTIER repository. The data change rate can be an estimate or can be measured through a site assessment. The estimated or measured data change rate and all the information gathered by the data protection survey provides an estimate or measurement for the *factoring ratio*, which is defined as the ratio of nominal capacity (the sum of all user data backup streams) to the physical capacity used (including all user data, metadata, and spare capacity, that is, the total amount of disk storage for which the user pays). In the following section, we discuss the formulas and worksheets necessary to calculate the factoring ratio and the corresponding physical capacity to purchase for sustaining a certain amount of nominal capacity.

### 3.3.6 Factoring ratio considerations

In this section, we discuss the factoring ratio in more detail, focusing on the parameters on which it depends. With ProtecTIER systems, the factoring ratio can grow to 25:1 or more, depending on these parameters. The factoring ratio depends *heavily* on two key variables:

- ▶ The data retention period: This is the period of time (usually measured in days) that defines how long you are going to keep your disk-based backups online. This period of time typically ranges from a period of 30 to 90 days, but can be much longer. This value is required when you compile the data protection survey, as discussed in “Data protection survey” on page 59.

**Note:** A longer retention period yields a higher factoring ratio because the data change rate decreases and therefore less new data comes in the repository, reducing the physical capacity required.

- ▶ The data change rate: This is the rate at which the data received from the backup application changes from backup to backup. This measurement has the most relevance when *like* backup policies are compared. (Data change rates might range from 1% to > 25%, but are difficult to directly observe.) The data change rate can be directly measured through an onsite assessment or can be estimated. Therefore, the more accurate the data change rate is, the more accurate the estimate will be about the sizing of the ProtecTIER Repository (see “Calculating deduplication factoring ratios” on page 78). Note that the factoring ratio is roughly the inverse of the data change rate. The data change rate is required when IBM conducts the data protection survey as part of the sizing process during an onsite, pre-sales engagement, as discussed in “Data protection survey” on page 59, and when you calculate the estimated factoring ratio, as described in “Calculating deduplication factoring ratios” on page 78.

In production environments, the ProtecTIER Repository will be a blend of many backup policies (data types) that protect many different application and data environments. Each backup policy has two variables that primarily influence the realized factoring ratio (and subsequent physical storage requirements for the ProtecTIER Repository):

- ▶ The data change rate
- ▶ The data retention period

The values of these variables differ across the various backup policies and associated data sets.

**Note:** Each policy can be said to have its own unique factoring ratio and nominal and physical storage capacities. If possible, use different VTLs or different ranges of barcodes assigned to the different policies to ease tasks for possible optimization.

The key task in capacity planning is to determine the physical storage required for *all* data types used in the analysis, which is done by first determining the nominal and physical storage capacities required for each data type and totaling these values up for all data types. After a total nominal and total physical storage capacity is calculated, a system-level factoring ratio can be calculated for the overall repository. Therefore, a weighted average change rate is calculated based on percentage estimates of each type of backup policy.

Capacity planning is both an art and a science. When sizing the ProtecTIER Repository capacity, it is important to build in some extra capacity. This allows for a margin of error and adds a buffer for scenarios that require more capacity, for example:

- ▶ You add more backup policies to your environment.
- ▶ Your backup policies grow (corporate data growth).
- ▶ Changes in the IT environment in terms of new applications.
- ▶ Amount of customers increase/change unexpected, if you provide backup external services.

The size of this buffer or *padding* will vary from situation to situation.

**Note:** Adding 10% to the physical storage calculations is a best practice.

If you can appreciate the importance of this margin, and given the value in disk savings that ProtecTIER systems provides, the incremental cost of the disk is easily justified.

### Calculating deduplication factoring ratios

The formulas listed below let you calculate the estimated factoring ratio for each data type (also called the backup policy). The required input is:

- ▶ Deduplication assumptions: Compression<sup>1</sup>, full backup change rate, and incremental backup change rate. The change rate can be estimated or can be calculated, as we explain in the following method sections.
- ▶ All the data gathered in the data protection survey (see “Data protection survey” on page 59 for more details).

We describe the formula gathering the main factors to simplify and better understand the meaning of each factor. As described in 3.3.6, “Factoring ratio considerations” on page 77, the factoring ratio is the nominal capacity divided by the physical capacity and so we are going to explicitly discuss these two factors.

$$\text{NominalCapacity} = \text{FullCapacityVersions} + \text{IncrementalCapacityVersions}$$

Where:

<b>NominalCapacity</b>	This parameter represents the overall capacity stored in the repository during the retention period and is composed of all the full and incremental versions stored.
<b>FullCapacityVersions</b>	This parameter represents the overall full backup capacity (expressed in GB) stored during the retention period. In the following formula, you can see how the FullCapacityVersions depends on the FullCapacity, FullRetention, and FullFrequency parameters.
<b>FullCapacity</b>	This parameter represents the capacity (expressed in GB) stored during full backup.
<b>FullRetention</b>	This parameter represents the retention period (expressed in days) for the full backup jobs (for example, you might decide to retain your full jobs for 30 days).

<sup>1</sup> The compression that ProtecTIER uses is called Delta Compression. It is a customized version of an open source standard compression algorithm. It behaves like LZH, but is not LZH.

**FullFrequency** This parameter indicates how often you perform the full jobs during the retention period (for example, four versions in 30 days, that is, one full job a week, so this parameter must be set to a value of 7).

**Note:** The number of versions is obtained by dividing FullRetention by FullFrequency.

In the following formula, you can see the relationship between these parameters.

$$FullCapacityVersions = FullCapacity \times \left( \frac{FullRetention}{FullFrequency} \right)$$

Where:

**IncrementalCapacityVersions** This parameter represents the overall incremental backup capacity (expressed in GB) stored during the retention period. In the formula below, you can see how the IncrementalCapacityVersions depends on the IncrementalCapacity, IncrementalFrequency, IncrementalRetention, FullRetention, and FullFrequency parameters.

**IncrementalCapacity** This parameter represents the capacity (expressed in GB) stored during incremental backup.

**IncrementalRetention** This parameter represents the retention period (expressed in days) for the incremental backup jobs.

**IncrementalFrequency** This parameter indicates how often you perform the incrementals during the retention period (this parameter must be set to the value 1 if you perform an incremental every day).

**FullRetention** This parameter represents the retention period (expressed in days) for the full backup jobs (for example, you might decide to retain your full jobs for 30 days).

**FullFrequency** This parameter indicates how often you perform the full jobs during the retention period (for example, four versions in 30 days, that is, one full job a week, so this parameter must be set to a value of 7).

**Note:** In the formula below, you can see that you have to remove the number of full versions because during full backups, incremental backups are not performed.

For the physical capacity, we have the following formula.

$$IncrementalCapacityVersions = \frac{IncrementalCapacity}{IncrementalRetention - FullRetention} \times \frac{FullRetention}{IncrementalFrequency \times FullFrequency}$$

Where:

**PhysicalCapacity** This parameter represents the physical capacity (expressed in GB) effectively required in the repository to satisfy the nominal capacity of your environment.

**FullPhysicalCapacity** This parameter indicates the full physical capacity (expressed in GB) effectively required in the repository. In the formula below, note that a first full backup must be entirely stored because no data is in the repository.

Therefore, it is not possible to make an initial delta comparison.

**IncrementalPhysicalCapacity** This parameter indicates the incremental physical capacity (expressed in GB) effectively required in the repository.

**CompressionRate** This parameter describes the compression rate obtainable in the ProtecTIER through its Delta Compression. Note that it is possible to reduce the initial backup of unique new data as well.

In the formula shown below, you can calculate the FullPhysicalCapacity parameter.

$$\text{FullPhysicalCapacity} = \text{FullCapacity} + (\text{FullCapacityVersions} - \text{FullCapacity}) \times \text{FullChangeRate}$$

FullChangeRate indicates the estimated change rate between full backups in your current environment. Again, note that a first full backup must be entirely stored because no data is present on the repository, and so it is not possible to make an initial delta comparison.

The following formula shows how to calculate the incremental physical capacity.

$$\text{IncrementalPhysicalCapacity} = \text{IncrementalCapacityVersions} \times \text{IncrementalChangeRate}$$

IncrementalChangeRate indicates the estimated change rate between incremental backups in your current environment. Note that a first full backup must be entirely stored, because no data is present on the repository, and so it is not possible to make an initial delta comparison.

Finally, the factoring ratio is shown in the following formula.

$$\text{FactoringRatio} = \frac{\text{NominalCapacity}}{\text{PhysicalCapacity}}$$

This formula is quite complex, but it might give you an idea of the impact of the estimated data change rate on the estimated factoring ratio. Increasing the data change rate leads to a decreased factoring ratio. Also note how the compression rate is inversely proportional to the physical capacity.

Another relationship involves the nominal capacity, the retention period, and the backup frequency. Increasing the retention period or decreasing the backup frequency leads to an increasing factoring ratio.

### 3.3.7 Storage sizing

In this section, we discuss sizing the storage subsystem in terms of the number of physical disks, technology, LUNs, array, and RAID protection.

There are three types of the ProtecTIER system: the TS7650G gateway model, the TS7650 appliance model, and the TS7610 small medium business model. The TS7650 and TS7610 models include the storage subsystem, and the TS7650G does not include the storage subsystem. Therefore, the customer has to order the storage subsystem separately. This section describes the storage sizing for the different models.

## Disk configurations for TS7650G

The TS7650G can attach to a variety of disk subsystems from different market ranges. The key to a well performing and stable system is the selection of the disk subsystem to match the customer needs and requirements.

### Selecting drives and a disk subsystem for TS7650G

The speed, the type of drives used, and the type of storage impacts the performance of your system. Typically, the faster the drive, the higher the performance. This increase in performance comes at a cost. The faster drives typically have a higher cost than the lower performance drives. In addition, the modern virtualized disk storage outperforms older technology. As shown in Table 3-9, the key factors to select the needed disk subsystem are shown. The table is based on cross calculations for a single disk drive in the system listed. For the midrange segment, you can see that Fibre Channel disks outperform the SATA technology. An additional new option with IBM Storwize V7000 broadens the bandwidth for the ProtecTIER Repository.

Table 3-9 Listing of different storage technologies

Factor	Fibre Channel 10 KRPM	Fibre Channel 15 KRPM	SATA	V7000	IBM XIV®
Spin speed	10 K	15 K	7.2 K	N/N, virtual	N/N, virtual
Command queuing	16 Max	16 Max	1 Max	Yes	256 per vol. mapped
Single disk I/O rate (number of 512 byte IOPS) <sup>a</sup>	280	340	88	~150	~12,000
Read bandwidth (MBps)	69	76	60	~200	N/A
Write bandwidth (MBps)	68	71	30	~100	N/A

a. Note that IOPS and bandwidth figures are from disk manufacturer tests in ideal lab conditions. In practice, you will see lower numbers for Fibre Channel and SATA technology, but the ratio between SATA and FC disks still applies. The V7000 and XIV have roughly the same average according to the current available test results and can differ in practice.

The speed of the drive is the number of revolutions per minute (RPM). A 15 K drive rotates 15,000 times per minute. With the faster speeds comes the ability to have greater throughput.

Seek time is how long it takes for the drive head to move to the correct sectors on the drive to either read or write data. It is measured in thousandths of a second (milliseconds or ms). The faster the seek time, the quicker the data can be read from or written to the drive. The average seek time reduces when the speed of the spindle increases. Typically, a 7.2 K drive will have an average seek time of around 9 ms, a 10 K drive will have an average seek time of around 5.5 ms, and a 15 K drive will have an average seek time of around 3.5 ms. In fact, the seek time drops dramatically when a physical RAID array is holding more than one file system volume. Configuring more than one logical RAID arrays on a single physical RAID group dramatically impacts the system performance and must not be configured on storage products that do not virtualize the disk drives. The use of smaller disks to increase the amount of disks in such systems should be considered. Using a disk subsystem that virtualizes the RAID groups across many physical disk drives in a smaller configuration could be a beneficial solution instead.

Command queuing allows for multiple commands to be outstanding to the disk drive at the same time. The drives have a queue where outstanding commands can be dynamically rescheduled or re-ordered, along with the necessary tracking mechanisms for outstanding and completed portions of workload. The SATA disks do not have command queuing and the Fibre Channel disks currently have a command queue depth of 16. For the IBM Storwize V7000, discuss with an IBM storage specialist for this product what will be the best configuration.

For the virtualized storage product IBM Storwize V7000, the blocks of a virtual disk volume configured are distributed across the physical disk drives used. This method improves the performance. In this configuration option for ProtecTIER, the amount of V7000 modules and spindles are key factors in outstanding performance.

The blocks of a virtual disk volume in IBM System Storage XIV are distributed across different physical disks and modules. The key factor of performance is the amount of modules and the amount of used Fibre Channel connections.

### **Disk size for TS7610**

TS7610 is the IBM System Storage ProtecTIER Entry Edition. This system can be configured for two types of capacity:

- ▶ A Small TS7610 Appliance repository binary capacity of 4 TB
- ▶ A Medium TS7610 Appliance repository binary capacity of 5.4 TB

It is possible to change from 4 TB to 5.4 TB by the applying feature code 9314. See the document *Increasing Capacity on the 3959 SM1 from 4.0 TB to 5.4 TB (Feature Code 9314)* for more details. This machine has a 1 TB SATA disk drive for data storage.

### ***Planning the storage structure***

It is important to configure a storage system in accordance with the needs of the user. An important question and primary concern for most users or storage administrators is how to configure the storage subsystem to achieve good performance. Keep in mind that the I/O operations for all system activities of an IBM TS7600 with ProtecTIER are add up. In the worst case, consider a many to many hub that performs local backups and restores in parallel by performing additional sending and receiving data for native replication from three group members. This situation requires more I/O bandwidth from the back-end storage subsystem, as through only local backups and restores are running on this many to many hub. The storage subsystems capable of high IOPS, such as those with Fibre Channel drives or IBM StoreWize V7000, help deliver better TS7600 performance.



The number of physical drives within a disk storage subsystem contributes to higher performance. The theoretical impact of a balanced sizing of the TS7650 repository disk structure in regards of performance is based on the metadata file system, as can be seen in Figure 3-14, where the different RAID group configurations for metadata are shown as an example.




Figure 3-14 Example metadata RAID group performance: 450 GB FC 15 K RPM disk subsystem for 40 TB repository

The values are based on Fibre Channel disks with 15,000 RPM spindle speed and a disk size of 450 GB. The output is based on the IBM internal planner tool V2.5 for a 40 TB TS7650G repository. To be able to achieve a performance of 1100 MB per second, and considering that every I/O for user data must be referenced to the metadata, the RAID group selected configuration must perform well. Therefore, two RAID groups of 6+6 disk drives or three RAID groups of 4+4 disk drives can be selected. To be sure that the system achieves this performance and that the selected metadata RAID groups do not limit the overall performance of the TS7650, the next higher possible RAID group configuration can be selected. This configuration depends on the available amount of spindles in the system and your budget. If the file system for metadata is filled in a production environment, this file system can be expanded or you can simply add an additional RAID group. The used capacity for the metadata disks is not included in the total capacity view of the ProtecTIER manager GUI; it only displays the capacity of the available user data.

The second factor that can limit performance is the size and the number of the defined RAID groups for the user data file systems. To continue our example, refer to Figure 3-15, which shows the difference between the theoretical performance of two different RAID groups. A RAID group 7+1 will perform better than a RAID group 5+1 built on a 15 K RPM Fibre Channel disk subsystem with 450 GB single disk capacity. The more disks in a RAID group, the higher the performance will be on this system. To scale performance up, simply add additional RAID arrays to the system. As a performance of 1100 MBps is required in this example, the user data file systems must provide a similar performance. Therefore, the 7+1 RAID configuration can be selected. Having 26 7+1 RAID groups provides the required performance.




Figure 3-15 Example user data RAID group performance: 450 GB FC 15 K RPM disk subsystem for a 40 TB repository

The 26 7+1 RAID group is the disk RAID group configuration for this given example. It is important to consult with your IBM ProtecTIER sales specialist so that he can provide you with the latest available planning tools and information. You could discover that the configured amount of file systems do not meet your expectations regarding performance. Adding additional RAID groups to scale up the performance might not result in better performance in the future. The most important thing is to have the initial system designed and the hardware configured correctly.

The amount of disk storage controller cache memory and its fast write efficiency are also factors in overall performance. There is no simple answer and no best guideline for storage performance optimization that is valid in every environment and for every particular situation. You can find a preliminary (and less detailed) performance discussion in this section.

Also in this section, we review other aspects of the system configuration that can help optimize the storage capacity and resilience of the system. In particular, we review and discuss the RAID levels, array size, and array configuration.

The performance of the TS7650 and TS7610 appliances with ProtecTIER is limited by the amount of disks available in these products.

### RAID levels

In this section, we go through the different RAID levels and explain why we choose a particular level in a particular situation, and then you can draw your own conclusions.

### **RAID 0**

RAID 0 is also known as *data striping*. It is well suited for program libraries requiring rapid loading of large tables or, more generally, applications requiring fast access to read-only data or fast writing. RAID 0 is only designed to increase performance. There is no redundancy, so any disk failures require reloading from backups. Select RAID 0 for applications that would benefit from the increased performance capabilities of this RAID level. Never use this level for critical applications that require high availability.

### **RAID 1**

RAID 1 is also known as *disk mirroring*. It is most suited to applications that require high data availability, good read response times, and where cost is a secondary issue. The response time for writes can be somewhat slower than for a single disk, depending on the write policy. The writes can either be executed in parallel for speed or serially for safety. Select RAID 1 for applications with a high percentage of read operations and where cost is not a major concern. Because the data is mirrored, the capacity of the logical drive when assigned RAID 1 is 50% of the array capacity.

Here are some recommendations when using RAID 1:

- ▶ Use RAID 1 for the disks that contain your operating system. It is a good choice because the operating system can usually fit on one disk.
- ▶ Use RAID 1 for transaction logs. Typically, the database server transaction log can fit on one disk drive. In addition, the transaction log performs mostly sequential writes. Only rollback operations cause reads from the transaction logs. Therefore, we can achieve a high rate of performance by isolating the transaction log on its own RAID 1 array. Use write caching on RAID 1 arrays. Because a RAID 1 write will not complete until both writes have been done (two disks), performance of writes can be improved through the use of a write cache. When using a write cache, be sure that it is battery-backed up.

**Note:** RAID 1 is implemented only as RAID 10 (described in “RAID 10” on page 88) on DS4000 products.

## RAID 5

RAID 5 (Figure 3-16) stripes data and parity across all drives in the array. RAID 5 offers both data protection and increased throughput. When you assign RAID 5 to an array, the capacity of the array is reduced by the capacity of one drive (for data-parity storage). RAID 5 gives you higher capacity than RAID 1, but RAID 1 offers better performance. RAID 5 is best used in environments requiring high availability and fewer writes than reads. RAID 5 is good for multi-user environments, such as database or file system storage, where the typical I/O size is small and there is a high proportion of read activity. Applications with a low read percentage (write-intensive) do not perform as well on RAID 5 logical drives because of the way that a controller writes data and redundancy data to the drives in a RAID 5 array. If there is a low percentage of read activity relative to write activity, consider changing the RAID level of an array for faster performance. Use write caching on RAID 5 arrays, because RAID 5 writes will not be completed until at least two reads and two writes have occurred. The response time of writes will be improved through the use of write cache (be sure that it is battery-backed up).

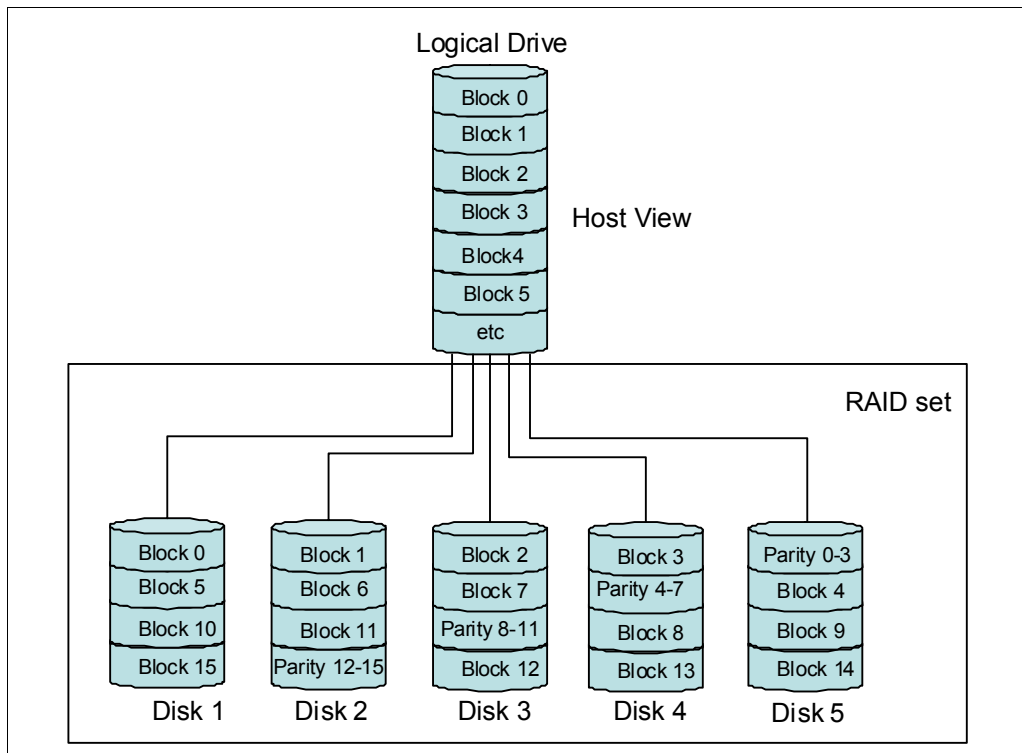


Figure 3-16 RAID 5

RAID 5 arrays with caching can give as good a performance as any other RAID level, and with some workloads, the striping effect gives better performance than RAID 1.

## RAID 6

RAID 6 (Figure 3-17) provides a striped set with dual distributed parity and fault tolerance from two drive failures. The array continues to operate with up to two failed drives. This makes larger RAID groups more practical, especially for high availability systems. This becomes increasingly important because large capacity drives lengthen the time needed to recover from the failure of a single drive. Single parity RAID levels are vulnerable to data loss until the failed drive is rebuilt. The larger the drive, the longer the rebuild will take. Dual parity gives time to rebuild the array without the data being at risk if one drive, but no more, fails before the rebuild is complete. RAID 6 can be used in the same workloads in which RAID 5 excels.

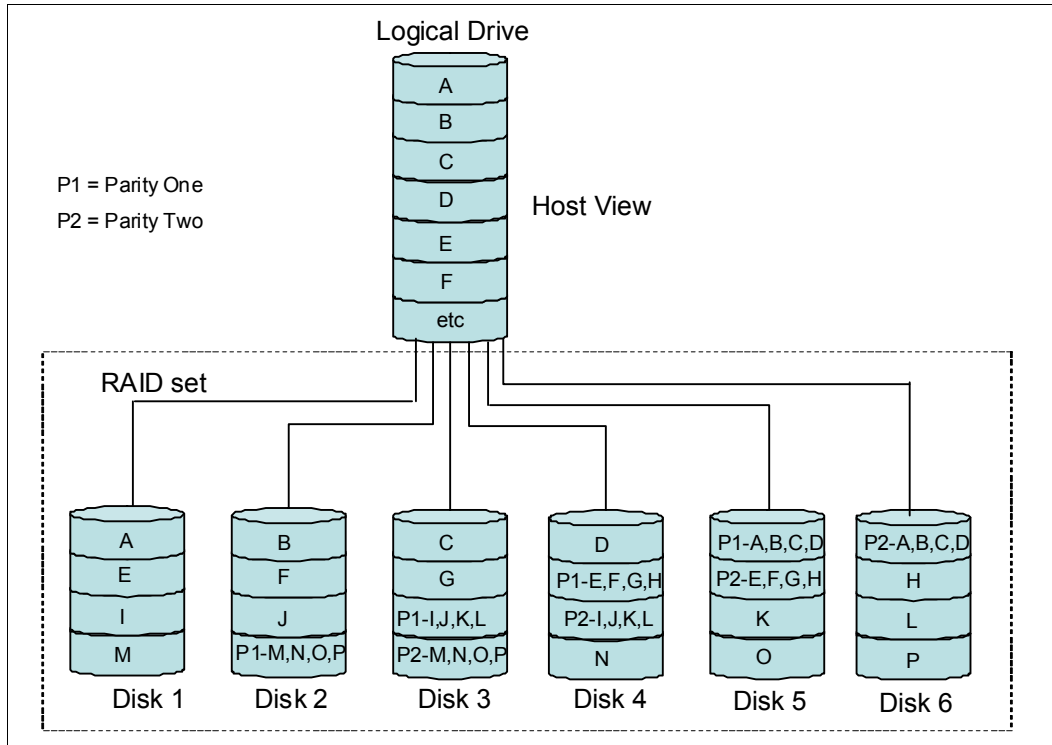


Figure 3-17 RAID 6

## RAID 10

RAID 10 (Figure 3-18), also known as RAID 1+0, implements block interleave data striping and mirroring. In RAID 10, data is striped across multiple disk drives, and then those drives are mirrored to another set of drives. The performance of RAID 10 is approximately the same as RAID 0 for sequential I/Os. RAID 10 provides an enhanced feature for disk mirroring that stripes data and copies the data across all the drives of the array. The first stripe is the data stripe. The second stripe is the mirror (copy) of the first data stripe, but it is shifted over one drive. Because the data is mirrored, the capacity of the logical drive is 50% of the physical capacity of the hard disk drives in the array.

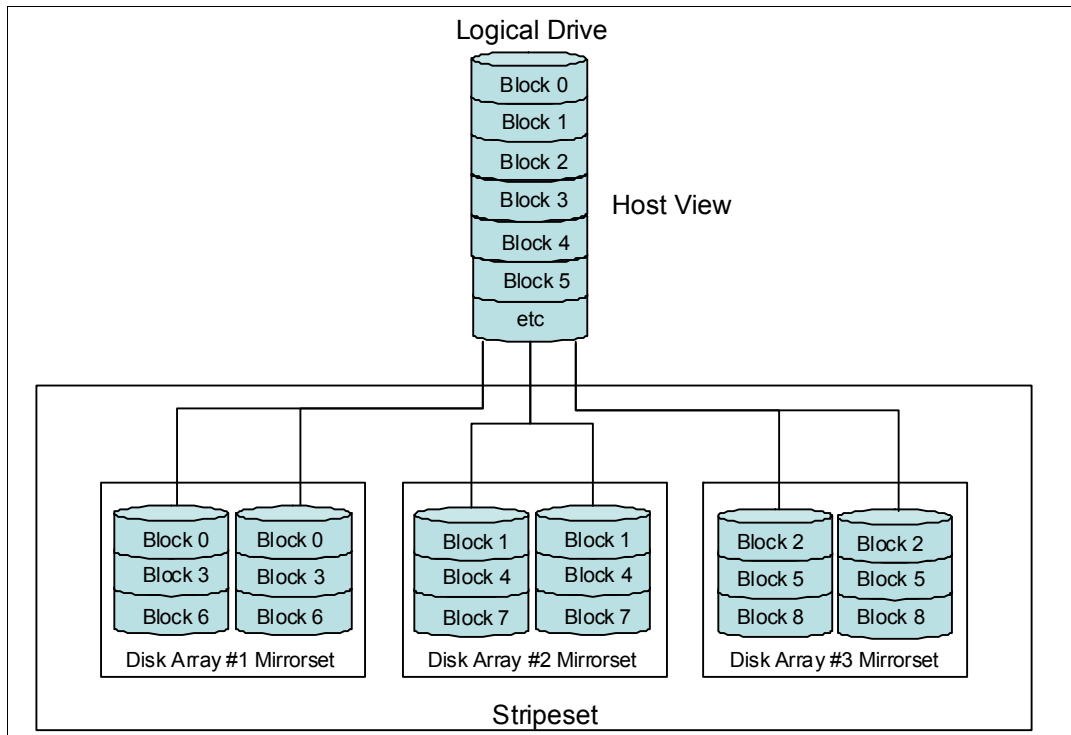


Figure 3-18 RAID 10

The recommendations for using RAID 10 are:

- ▶ Use RAID 10 whenever the array experiences more than 10% writes. RAID 5 does not perform as well as RAID 10 with a large number of writes.
- ▶ Use RAID 10 when performance is critical. Use write caching on RAID 10. Because a RAID 10 write will not be completed until both writes have been done, write performance can be improved through the use of a write cache (be sure that it is battery-backed up).
- ▶ When comparing RAID 10 to RAID 5:
  - RAID 10 writes a single block through two writes. RAID 5 requires two reads (read original data and parity) and two writes. Random writes are significantly faster on RAID 10.
  - RAID 10 rebuilds take less time than RAID 5 rebuilds. If a real disk fails, RAID 10 rebuilds it by copying all the data on the mirrored disk to a spare. RAID 5 rebuilds a failed disk by merging the contents of the surviving disks in an array and writing the result to a spare. RAID 10 is the best fault-tolerant solution in terms of protection and performance, but it comes at a cost.

## 3.4 Replication considerations

One of the benefits of deduplication is that it allows you to consider remote electronic transfer of data for disaster recovery purposes. The reduced amount of data to send to the remote site can make WAN data transfer of data more feasible in terms of bandwidth and costs. If you are going to replicate all data to the remote IBM System Storage TS7600 with ProtecTIER, then all of the topics in the previous sections apply. This is the case if the system initially started operation and replication of backed up cartridges starts the first time. The data of all the replicated cartridges on the single receiving hub does refer to the replicated cartridges to maintain the content of those for replication. On the second replication of the same cartridges, only a subset of data is transferred to the remote hub. It is necessary that the reduced amount of data to store at the remote site will factor in your sizing.

The IBM TS7600 with ProtecTIER Version 3.1 allows different scenarios for replication:

- ▶ The one directional many to one network replication of up to 12 spokes to replicate to a single hub while the spokes and hub can perform local backup while replication.
- ▶ The bidirectional four way many to many replication of up to four hubs where data simultaneously can be replicated to and from one to three hubs.
- ▶ The OST of an implicit hub mesh of up to 12 repositories.
- ▶ One replication manager to be running on a replication grid. Usually it is available on the nodes with the ProtecTIER software and can be activated or deactivated. The replication manager should run on the replication hub.
- ▶ The ProtecTIER replication grid manager can manage up to 24 repositories. If the amount of repositories to be managed is larger than 24, then it is possible to install the replication grid manager on a separate server other than TS7650, which requires an RPQ. With this option, 64 grids can be managed that can hold a total number of 256 repositories at each grid.

The planning process of ProtecTIER systems with replication deployed requires more input and consideration beyond the individual capacity and performance planning needed for a system that will only be used as a local VTL. Before introducing the replication capabilities, the ProtecTIER system was merely a target for a backup application data stream. Now when offering the replication capabilities of the local repository to a secondary DR site, the ProtecTIER system becomes a piece of the much larger picture of the organization's disaster recovery paradigm. The methodology behind a DR plan today calls for an intricate combination of policies, procedures, operations, and IT systems. The DR plans employed by IT organizations across different locations, regions, and types of business are as varied as are the businesses that they support. However, there is one fundamental building block upon which all DR plans rest: data accessibility at a secondary DR site. To achieve preparedness for a potential disaster, IT organizations can use the ProtecTIER replication capabilities to ensure that the mission-critical data that the business needs to continue operations is safely housed at a remote site, intact, with enterprise-class data integrity.

### 3.4.1 Supported replication configurations

The following list shows all of the supported replication configurations.

- ▶ New DD4 VTL installation, clustered or stand-alone, with replication
- ▶ New DD4 VTL installation, clustered or stand-alone, without replication
- ▶ New DD4 OpenStorage installation, clustered or stand-alone, with replication
- ▶ New DD4 OpenStorage installation, clustered or stand-alone, without replication

- ▶ Existing TS7650G installation with DD3 servers, adding new DD4 VTL servers for replication
- ▶ Existing TS7650G installation with DD4 VTL servers, adding new DD4 VTL servers for replication
- ▶ Existing TS7650G installation with DD4 OpenStorage servers, adding new DD4 OpenStorage servers for replication

### 3.4.2 Virtual shelf

The ProtecTIER replication feature introduces the concept of a virtual shelf. As a general rule, replication always occurs from the source shelf to a destination shelf. As with physical tape shelves, there is a limit to the number of tapes that can be put on a virtual shelf. The limit is the result of subtracting the total number of tapes in all libraries from the total number of tapes supported in a single repository. In other words, cartridges that can use for shelf will be reduced if you create and allocate new tape cartridges to a virtual library. Therefore, the shelf cartridges that can be used for replication will be reduced while the library cartridges are increase. For example, assuming a repository with three libraries, each contains 50,000 cartridges, the total number of cartridges in a single repository are 512,000. In this case, the number of cartridges that can be put on the virtual shelf is  $512,000 - 150,000 = 362,000$  cartridges.

**Note:** An OST machine does not have a shelf and there are no tape cartridges and drives. OST only has a repository for replication. ProtecTIER Version 3.1 supports replication on OST machines, but does not support replication between OST and VTL machines.

Take care when planning a repository to ensure that its shelf and the sum of tapes in all of its libraries do not exceed this max limit per repository (currently 512,000 for the TS7650G (Gateway) and 128,000 for the TS7650 Appliance).

The shelf can be used by the user without enabling the replication capability. In this case, every eject or export operation *sends* the cartridges to the virtual shelf. Again, the total number of cartridges cannot exceed the maximum number defined during the installation. For the calculation of this limit, all cartridges are taken into account, regardless of where they reside (in a library or on the shelf). This calculation uses the following formula:

Shelf cartridges + all library cartridges from all libraries = maximal amount of cartridges for entire repository

For example, assuming a repository with three libraries, if each library contains up to 10,000 cartridges and the user wants to be able to put 30,000 cartridges on the shelf, then they must plan the entire repository to have  $30,000 + 30,000 = 60,000$  cartridges.

This calculation should be used to plan the repository and the volume of backups and replication.

While data transfer is running, the remote site's cartridge cannot be relocated from the shelf. If the remote cartridge was manually relocated from the shelf, it might cause replication errors if you attempt to replicate to the cartridge.



### 3.4.3 Replication parallelism scheme

A single cartridge will have up to 32 concurrent data streams. Up to 128 cartridges can replicate simultaneously. The cartridge data transferred across the dedicated replication network is causing heavy network traffic at the initial first replication, so only changes made to the replicated cartridge will cross the network, which improves replication performance.

### 3.4.4 Initial synchronization considerations

The initial synchronization can require a long time to complete if the replication of the first cartridges is not immediately started after local backups are finished. It can be sped up if local and remote repositories are synchronized back-to-back in case the whole content of a repository must be mirrored to a remote site. Generally, the remote repository should receive the data through the replication function. There is no option to clone data from local to tape, then to restore, at the remote site, the replication policies over time while using the amount of bandwidth that was planned to support the steady state, which takes effect after the deduplication factor *occurs*. Therefore, you must consider the following options for the initial synchronization:

► Gradual Management of policies over time

This is the preferred method, whether the user is deploying a new system or adding replication to an existing system. In this method, the user adds new replication policies over time while manually ensuring that the total daily volume of replicated data remains within the bandwidth limit. Given that every new replication policy will send the full nominal volume each day, the user can calculate the total bandwidth consumption and stay within the limit. The following steps provide an overview of this process:

- a. Following the completion of a backup, manually select tapes that are used in a backup policy intended for replication and execute the replication policy.
- b. Following the replication completion of the previous replication policies, perform the following tasks within the backup application:
  - iii. Delete scratch tapes.
  - iv. Create new barcodes.
  - v. Point backup policies to the newly defined tape pool.
  - vi. Repeat steps i–iii until the secondary system is fully primed.

► Priming the secondary repository at a common locality with the primary system

The option of priming the secondary system at the primary site first and then moving it to its DR location has limited practical value and should not be considered as a first choice. If this approach is taken, the user must manage the synchronization process once again when the systems are placed in their final location. Given that this procedure might introduce weeks of catch-up, the above recommended approaches should be adopted as the best practice methods to successfully deploy a secondary/DR system.

However, if you are synchronizing a full, partial, or even a newly started repository, either you should have enough network bandwidth allotted to allow for the primary and secondary systems to synchronize within the available time frame, or the target system could be installed first at the primary location, and the full duration of the synch process must be scheduled locally while the system is connected locally to allow the maximum throughput between the servers.

No deduplication occurs for the initial synchronization. The new remote repository does not contain any data. The nominal data transfer equals the physical transfer, and uses reduced performance for the first replication.

A single node ProtecTIER system has 2 x 1 GB dedicated replication Ethernet ports:

- ▶ It can deliver a maximum throughput of 200 MBps in physical data-transfer terms.
- ▶ It delivers up to 600 MBps replication throughput of nominal data (amount of data before deduplication).

A dual-node clustered system has 4 x 1 GB dedicated replication Ethernet ports.

- ▶ It can deliver a maximum throughput of 400 MBps in physical data-transfer terms.
- ▶ It delivers up to 850 MBps replication throughput of nominal data (amount of data before deduplication).

**Note:** Your actual throughput could be much less depending on your WAN bandwidth capabilities and current utilization. Therefore, it is important to immediately use replication beginning from day one right after the first backup.

### 3.4.5 Bandwidth sizing and requirements

The ProtecTIER HyperFactor deduplication engine, which eliminates redundant data that must traverse the network from site to site, is a key enabler of cost-efficient replication for the DR architecture. Combining the deduplication and replication features of ProtecTIER provides dramatic improvements to disaster recovery planning and operations for an organization of any size. It enables ProtecTIER to extend the high level of data protection provided by replication to a wider range of data and applications in the data center and beyond.

Table 3-10 demonstrates the potential savings for ProtecTIER Replication users over Generation 1 VTL with no deduplication function, focusing only on the bandwidth cost.

*Table 3-10 Potential savings*

<b>Traditional VTL (no deduplication)</b>	<b>Value</b>
Daily backup (TB)	5
Data transferred (TB)	5
Replication window (hrs)	16
Mbps (megabits) required	728.18
Required bandwidth	OC12 +
Cost of bandwidth per month	\$30,000.00
<b>ProtecTIER</b>	
Daily backup (TB)	5
Data transferred (TB)	0.5
Replication window (hrs)	16
Mbps (megabits) required	72.8
Required bandwidth	OC3

Traditional VTL (no deduplication)	Value
Cost of bandwidth per month	\$11,000
BW savings per month	\$19,000
Total 3-year bandwidth savings	\$684,000

The effect of ProtecTIER deduplication capability on the required bandwidth for replication is similar to the effect that the daily deduplication ratio has on the daily backup workload. For example, if the daily change rate is 10%, which means only 10% of the data changes from one day to the next, then not only does the system backup adjust the changed, unique data, but more significantly, it only requires one-tenth of the replication network bandwidth that would otherwise be needed without data deduplication.

Table 3-11 shows the average costs by bandwidth type. Bandwidth costs vary dramatically by geography and distance, so the figures are approximate as an average for the U.S. Using this table, the potential cost savings of ProtecTIER are easy to see. For example, if ProtecTIER allows the user to deploy an OC12 rather than an OC48, the annual cost savings will be close to \$700,000 (the difference of \$960,000 and \$279,000, as seen in the last column).

Table 3-11 Bandwidth, capacities, and costs

Bandwidth type	Mbps	MBps	Lease per month (long haul)	Total MB in 24 hours	Total GB in 24 hours	Cost/year
T1/DS1	1.54	0.2	\$2,200	16,632	17	\$26,400
T3/DS3	44.7	5.6	\$5,000	482,760	483	\$60,000
OC3	155.5	19.4	\$11,000	1,679,400	1,679	\$132,000
OC12	622	77.8	\$23,325	6,717,600	6,718	\$279,900
OC48	2488	311.0	\$80,000	26,870,400	26,870	\$960,000
OC192	8853	1106.6	\$254,524	95,612,400	95,612	\$3,054,288
OC768	39813	4976.6	\$1,144,624	429,980,400	429,980	\$13,735,488

### 3.4.6 Network bandwidth sizing tips

ProtecTIER replication only replicates the new or unique elements of the deduplicated data. Data that was deduplicated on the primary will not be physically sent to the remote site (saving bandwidth), but the remote site still must read the entire amount of data to ensure 100% data integrity. With ProtecTIER V3.1, the new feature of a hub of replicating data to and receiving replication data from members of a four way many to many replication group of four members affects the replication network bandwidth, especially if the customer does not have the possibility to provide dedicated networks for replication purposes and must share networks with other network traffic.

A simple example demonstrates the impact of deduplication on the replication operation. Assuming that there are two cartridges at the primary site, cartridge A and cartridge B, that contain the exact same 1 GB of data:

- ▶ Replicating cartridge A transfers 1 GB of physical (which equals nominal) data, as at this point in time the data is new to the remote repository.
- ▶ Replicating cartridge B transfers 0 GB of physical data to the remote site, as all of the data already exists at the remote repository. In effect, 1 GB of nominal data is represented and indexed as cartridge B following its replication action.

There are two types of replication data-transfer throughput barriers:

- ▶ Physical data-transfer barrier, which stems from the fact that each ProtecTIER node has two 1 GbE ports:
  - A single-node system supports up to 200 MBps physical data transfer.
  - A two-node clustered system supports up to 400 MBps physical data transfer (as it has 4 x 1 GB replication Ethernet ports).
- ▶ Nominal data barrier: Nominal data is the original amount of backed-up data before applying ProtecTIER's deduplication factor, which stems from the maximum processing capability of a given ProtecTIER system:
  - A single-node system supports up to 600 MBps of nominal data replication.
  - A two-node clustered system supports up to 850 MBps of nominal data replication.

**Note:** The actual maximum replication throughput for any given scenario will be dictated by the barrier type (physical or nominal) that the ProtecTIER system will reach first.

Network failures and resource allocation problems leads to retries of the replication process for a specific cartridge for a period of seven days.

The formula shown in Figure 3-19 should be used to calculate the replication data transfer for the system replicating data to a hub. The formula calculates the actual changed data that must be sent across the network (daily backup \* change rate), and adds overhead capacity that is equal to 0.5% of the daily backup workload.

$$\text{Daily network workload out} = (\text{daily backup} * \text{change rate}) + (0.5\% * \text{daily backup})$$

Figure 3-19 Calculate daily network load created by system replicating data to a hub

As an example, for 6 TB daily backup with a change rate of 10%, the formula would be:

$$(6000 \text{ GB} * 10\%) + (0.5\% * 6000 \text{ GB}) = 630 \text{ GB}$$

Thus, in this scenario, 630 GB of physical data is replicated to the second site, rather than the 6 TB that would otherwise be transferred without deduplication.

The following formulas should be used to calculate the replication bandwidth required:

- ▶ Divide the daily network workload by the available hours in a day for replication. For example:

$$\text{Time frame for replication window (preferred mode of operation) of 10 hour} - 630 \text{ GB} / 10 \text{ hour} = 63 \text{ GB/hour required replication bandwidth}$$

- ▶ Continuous replication operation (24 hour period, concurrent with backup operation, rarely the recommended mode of operation):

630 GB / 24 hours = 26 GB/hour required replication bandwidth

**Note:** We recommend adding 10% of the required bandwidth for “headroom” due to network outages or slowdown periods.

### 3.4.7 Bandwidth validation tool

The ProtecTIER Replication Network Performance Validation Utility (pt\_net\_perf\_util) is a tool that tests and verifies the user’s replication network performance before replication is deployed on the IBM System Storage TS7600 with ProtecTIER starting with the new ProtecTIER R2.3 software installation or upgrade. This tool should be used to ensure that the replication network is capable of delivering the expected performance. The first element to consider is the bandwidth availability across the network, but when assessing the network, keep two other major components in mind, that is, latency and packet loss, which determine the network quality. The latency in any user’s WAN is dependant upon many factors along the network span and might vary, but should never exceed 200 ms. If it does, it might significantly decrease the system replication throughput. Packet loss across the network should be 0%. Any other value implies a major network problem that must be addressed before replication is deployed.

The pt\_net\_perf\_util network testing utility is included as part of the ProtecTIER software package and the installer needs the physical ProtecTIER server nodes at both sites to run this utility concurrently. At this point in the installation process, however, it is not yet necessary to build a repository or configure the ProtecTIER back-end disk.

#### How to use the utility

The pt\_net\_perf\_util utility’s objective is to test maximum replication performance between two future ProtecTIER repositories by emulating the network usage patterns of ProtecTIER’s replication component. This utility does not predict replication performance, but it might discover performance bottlenecks.

The pt\_net\_perf\_util utility and the iperf and nuttcp tools that it uses are installed as part of the ProtecTIER software installation. To test the replication performance, use one of the following tools:

- ▶ iperf 2.0.4, which is found in /usr/local/bin/iperf.
- ▶ nuttcp 6.1.2, which is found in /usr/local/bin/nuttcp-6.1.2.

This utility has two modes of operation, client and server. The server must be started before the client. Before running the utility, shut down all other programs on both the client and the server ProtecTIER systems. The client is the ProtecTIER system that transmits the test data and the server is the ProtecTIER system that receives the data (also known as the target server). Based on the data sent by the client and received by the server, the script outputs key network parameter values that indicate certain attributes of the network. The goal of these tests is to benchmark the throughput of the network. The most important benchmark is the direction in which replication will actually take place, that is, the target should be tested as the server because the flow of data will be to that server from the client. However, it is also important to also test the reverse direction to measure the bandwidth performance during disaster recovery failback. Network bandwidth is not always the same in both directions.

In the following procedure, the goal is to test network performance between two machines on a WAN, server1 and server2. Each test runs for five minutes. Because there are five tests, the process takes a total of 25 minutes.

Complete the following steps:

1. Start the server mode of the utility on server1 by entering the following commands on the command line (refer to Example 3-1 and Example 3-2 on page 97):

```
cd /opt/dtc/app/sbin
./pt_net_perf_util -s
```

**Note:** The above commands use the iperf tool. To use the nuttcp tool instead, add -n to the command. Enter one of the following series of commands to use nuttcp:

```
cd /opt/dtc/app/sbin
./pt_net_perf_util -sn

or

cd /opt/dtc/app/sbin
./pt_net_perf_util -s -n
```

**Note:** You have to stop the VTFD service. If you did not stop the VTFD service, the \*\*\* VTFD service is not down, please stop it and then run the utility message is displayed. Refer to Example 3-1

*Example 3-1 Server1 output of performance utility*

---

```
root@iceland ~]# service vtfd stop
Shutting down vtfd: [ OK ]
```

Please wait

```
.....
.....
[root@iceland ~]# cd /opt/dtc/app/sbin
[root@iceland sbin]# ./pt_net_perf_util -s
-----
Server listening on TCP port 5001
TCP window size: 85.3 KByte (default)
-----
```

---

2. Start the client mode of the utility on server2 by entering the following command on the command line:

```
./pt_net_perf_util -c server1 -t 300
```

**Note:** This step uses the iperf external utility. To use nuttcp instead, add -n to the command. Enter the following command to use nuttcp:

```
./pt_net_perf_util -c server1 -t 300 -n
```

3. The utility automatically performs the tests in sequence. The client output for server2 is shown Example 3-2.

**Note:** In the sample output shown in Example 3-2, the test ran for only 5 seconds instead of 300.

*Example 3-2 Server2 output*

---

```
[root@barcelona ~]# cd /opt/dtc/app/sbin
[root@barcelona sbin]# ./pt_net_perf_util -c iceland -t 300
*** Latency
PING iceland (9.11.201.112) 56(84) bytes of data.
--- iceland ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4000ms
rtt min/avg/max/mdev = 0.159/0.253/0.334/0.065 ms

*** Throughput - Default TCP
[ 3] 0.0- 5.0 sec 56.7 MBytes 94.8 Mbits/sec

*** Throughput - 1 TCP stream(s), 1024KB send buffer
[ 3] 0.0- 5.0 sec 57.0 MBytes 94.9 Mbits/sec

*** Throughput - 16 TCP stream(s), 1024KB send buffer
[SUM] 0.0- 5.9 sec 66.0 MBytes 94.4 Mbits/sec

*** Throughput - 127 TCP stream(s), 1024KB send buffer
[SUM] 0.0-11.3 sec 127 MBytes 94.1 Mbits/sec

Number of TCP segments sent: 231357
Number of TCP retransmissions detected: 1 (0.000%)
```

---

See Example 3-3 for information about interpreting the results of the tests.

*Example 3-3 Server1 out put*

---

```
[ 4] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42115
[ ID] Interval      Transfer      Bandwidth
[ 4] 0.0- 5.0 sec 56.7 MBytes 94.1 Mbits/sec
[ 5] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42116
[ ID] Interval      Transfer      Bandwidth
[ 5] 0.0- 5.1 sec 57.0 MBytes 94.1 Mbits/sec
[ 4] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42117
[ 5] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42118
[ 6] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42119
[ 7] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42120
[ 8] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42121
[ 9] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42122
[10] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42123
[11] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42124
[12] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42125
[13] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42126
[14] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42127
[15] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42128
[16] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42129
[17] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42130
```

```

[ 18] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42131
[ 19] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42132
[ ID] Interval      Transfer      Bandwidth
[ 10] 0.0- 5.2 sec  3.00 MBytes  4.87 Mbits/sec
[ ID] Interval      Transfer      Bandwidth
[ 15] 0.0- 5.3 sec  3.00 MBytes  4.77 Mbits/sec
[ ID] Interval      Transfer      Bandwidth
[ 17] 0.0- 5.3 sec  4.00 MBytes  6.29 Mbits/sec
[ ID] Interval      Transfer      Bandwidth
[ 13] 0.0- 5.5 sec  4.00 MBytes  6.15 Mbits/sec
[ ID] Interval      Transfer      Bandwidth
[  5] 0.0- 5.5 sec  4.00 MBytes  6.13 Mbits/sec
[ ID] Interval      Transfer      Bandwidth
[  7] 0.0- 5.5 sec  4.00 MBytes  6.08 Mbits/sec
[ ID] Interval      Transfer      Bandwidth
[  6] 0.0- 5.5 sec  4.00 MBytes  6.06 Mbits/sec
[ ID] Interval      Transfer      Bandwidth
[  4] 0.0- 5.6 sec  8.00 MBytes  12.0 Mbits/sec
[ ID] Interval      Transfer      Bandwidth
[ 12] 0.0- 5.7 sec  5.00 MBytes  7.32 Mbits/sec
[ ID] Interval      Transfer      Bandwidth
[  8] 0.0- 5.7 sec  4.00 MBytes  5.85 Mbits/sec
[ ID] Interval      Transfer      Bandwidth
[  9] 0.0- 5.8 sec  4.00 MBytes  5.79 Mbits/sec
[ ID] Interval      Transfer      Bandwidth
[ 11] 0.0- 5.8 sec  3.00 MBytes  4.30 Mbits/sec
[ ID] Interval      Transfer      Bandwidth
[ 16] 0.0- 5.9 sec  3.00 MBytes  4.28 Mbits/sec
[ ID] Interval      Transfer      Bandwidth
[ 14] 0.0- 5.9 sec  4.00 MBytes  5.70 Mbits/sec
[ ID] Interval      Transfer      Bandwidth
[ 19] 0.0- 5.9 sec  5.00 MBytes  7.13 Mbits/sec
[ ID] Interval      Transfer      Bandwidth
[ 18] 0.0- 5.9 sec  4.00 MBytes  5.70 Mbits/sec
[SUM] 0.0- 5.9 sec  66.0 MBytes  94.0 Mbits/sec
[ 20] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42133
[  4] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42134
[  5] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42135
[  6] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42136
[  7] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42137
[  8] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42138
[  9] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42139
[ 10] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42140
[ 11] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42141
[ 12] local 9.11.201.112 port 5001 connected with 9.11.201.120 port 42142

```

---- output omitted -----

### Interpreting the results

The utility performs five foreground tests (tests 1 - 5) and one background test (test 6). The example outputs are from the client side, with the script using iperf (not nuttcp) in tests 2 - 5. Each of the first five tests ran for 30 seconds (-t 300), while the last test monitored TCP performance during that time.



### **Test 1: Latency**

This test checks the nominal network link latency and packet loss. Here is an example result:

```
*** 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
```

Where:

- ▶ The average round-trip-time (rtt) was 78.4 ms and there was 0% packet loss.
- ▶ The latency in WAN topologies might vary, but should never exceed 200 ms. Contact your network administrator if latency reports more than 200 ms, as it might significantly decrease replication throughput.
- ▶ Higher latency values will cause a major deterioration in replication throughput.
- ▶ Packet loss should be 0%. Any other value implicates a major network problem.

### **Test 2: Throughput: Default settings**

This test checks the maximal TCP throughput using a single data stream with the default TCP settings. Here is an example result:

```
*** Throughput - Default TCP
[ 3] 0.0-120.1 sec 2.41 GBytes 173 Mbits/sec
```

The test ran for 120.1 seconds, transferred 2.41 GB, with an average throughput of 173 Mbps.

**Note:** 1 MB is equal to 1,048,576 bytes. 1 Mbps is equal to 1,000,000 bits per second.

### **Test 3: Throughput: Single stream with 1 MB send buffer**

This test checks maximal TCP throughput using a single data stream with a 1 MB send buffer. Here is an example result:

```
*** Throughput - 1 TCP stream(s), 1MB send buffer
[ 3] 0.0-120.0 sec 2.51 GBytes 180 Mbits/sec
```

The test ran for 120.0 seconds and transferred at 2.51 GBps, with an average throughput of 180 Mbps. There might be an improvement if you use high-latency links.

### **Test 4: Throughput: 16 streams with 1 MB send buffer**

Here is an example result:

```
*** Throughput - 16 TCP stream(s), 1MB send buffer
[SUM] 0.0-121.4 sec 5.91 GBytes 418 Mbits/sec
```

Where:

- ▶ The test ran for 121.4 seconds and transferred 5.91 GB, with an average throughput of 418 Mbps.
- ▶ The extra streams yielded higher utilization of the connection.
- ▶ The Mb per second reported in this test is the maximum replication performance that your system will achieve if your backup environment is using up to two to three cartridges in parallel.

### **Test 5: Throughput: 127 streams with 1 MB send buffer**

Here is an example result:

```
*** Throughput - 127 TCP stream(s), 1MB send buffer  
[SUM] 0.0-126.1 sec 8.08 GBytes 550 Mbits/sec
```

Where:

- ▶ The test ran for 126.1 seconds and transferred 8.08 GB, with an average throughput of 550 Mbps.
- ▶ TCP takes a while to reach its maximal throughput. Longer testing times (300 seconds or more) produce more accurate results.
- ▶ The throughput value given by this test is the potential physical replication throughput for this system. It is directly affected by the available bandwidth, latency, packet loss, and retransmission rate.
- ▶ The Mb 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.

### **Test 6: TCP retransmissions versus total TCP segments sent**

Here is an example result:

```
Number of TCP segments sent: 1619061  
Number of TCP retransmissions detected: 201038 (12%)
```

Where:

- ▶ A total of 1619061 TCP segments were sent during the five tests, out of which 201038 were lost and retransmitted.
- ▶ The retransmission rate imposes a direct penalty on the throughput, as the retransmission of these packets takes up bandwidth. The retransmission can be caused by the underlying network (for example, packet dropping by an overflowed router) or by the TCP layer itself (for example, retransmission due to packet reordering).
- ▶ Segment loss can be caused by each of the network layers.
- ▶ TCP retransmission larger than 2% might cause performance degradation and unstable network connectivity. Contact your network administrator to resolve this issue and reduce it to approximately 0%.

**Tip:** We recommend running these tests again to test the reverse throughput in the network. To run the tests in reverse, change server1 to the client and server2 to the server and repeat the procedures.

## **3.4.8 Repository sizing for replication with performance**

There are some considerations for sizing replication, which are performance, local/remote repository sizing, and network bandwidth sizing. We have to consider performance for repository sizing.

Performance and capacity sizing are essential for calculating the data amount to be backed up and the amount of data to be replicated.

When conducting the system performance sizing for a ProtecTIER deployment, there are several key factors to consider between the primary sites and the DR site. As discussed at length in Chapter 2, “IBM System Storage ProtecTIER architecture” on page 11, there are two available modes of operation for the replication activity: scheduled replication (with a predefined time window) and continuous replication, which runs concurrently with the backup operation. In most use cases, the scheduled replication approach should be explored first, as it enables the user to accurately plan for performance and it better ensures that SLAs are met.

The chosen approach has major consequences on the expected or required performance rates from the system and should be discussed thoroughly with the user. The performance requirements from the system are based on calculating the amount of data to be backed up and ingested during the backup window time frame, and the amount of data to be replicated during the replication window. As mentioned, the preferred and thus recommended strategy is to have separate windows for backup and replication with no overlap. In this scenario, the system performs the backup action at its peak available performance, and then during the replication window, it will deliver the same max performance for the replication operation.

Here are the key factors to be considered when sizing the system for all spokes and the hub for required performance:

- ▶ Assume that backup and replication operations require the same amount of system resources (CPU, I/Os, and so on).
- ▶ Sizing performance should take into account the planned mode of operation for the particular deployment:
  - Backup and replication activities are running during separate time windows; there is a planned backup window and then a planned replication window with no overlap.
  - Continuous operation of replication, in this case backup and replication activity, may run concurrently on the same repository.
- ▶ Under any mode of operation, a system designed for X MBps will deliver this maximum throughput:
  - Backup/restore alone will run at X MBps
  - Replication alone in a time frame mode will run at X MBpS. However, because of the rate-control function, it will be [X MBpS – 5%] while in continuous mode.
  - When continuous replication is deployed, concurrent backup/restore and replication will yield the MAX of X MBps for all activities combined.

**Note:** While running in a continuous replication mode, the system will balance priorities of backup and replication in real time based on a special dynamic algorithm (rate-control) that in most cases prioritizes the backup operation higher than the replication activity.

## Storage requirements

Replication metadata storage requires only up to 1.5% of the local repository physical size. The metadata resides on the repository associated with the ProtecTIER Replication Manager. It is not required for the TS7650 and TS7610 to size additional metadata.

## 3.5 Planning for OpenStorage

With OpenStorage, ProtecTIER can be integrated with NetBackup to provide the means for backup-to-disk without having to emulate traditional tape libraries. Using a plug-in that is installed on an OST-enabled media server, ProtecTIER can implement a communication protocol that supports data transfer and control between the backup server and the ProtecTIER server. Therefore, to support the plug-in, ProtecTIER implements a storage server emulation (see Figure 3-22 on page 115).

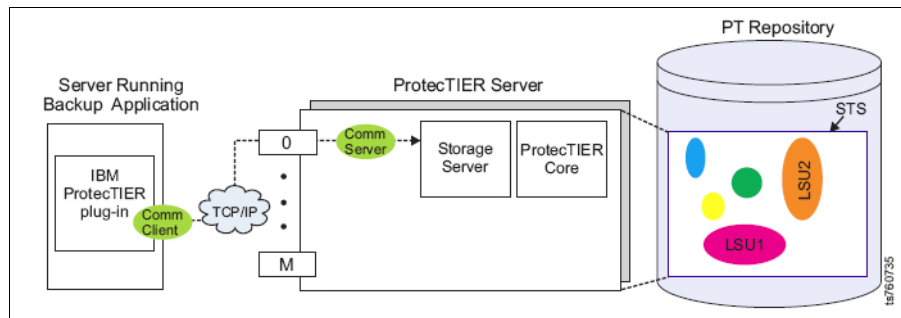


Figure 3-20 OST configuration map

OST has the following benefits:

- ▶ Treat disk as disk
  - Avoids the limitation of tape emulation.
  - Provides enhanced functionality only possible with disk.

- ▶ Tight integration

NetBackup is aware of all backup image copies and manages the creation, movement, and deletion of all backup images.

In regards to the size of needed repository, apply the same calculations used for the standard VTL. The difference here is the incoming channels for the attached hosts are on the VTL Fibre Channel adapters and are used for the TS7650 OST model 10 GB Ethernet connections. A single 10 GB Ethernet connection provides a bandwidth of 1.000 MBps. Depending of what TS7600 model is used, the interfaces can scale up to 6.000 MBps on a single node cluster. ProtecTIER is virtualizing the 10 GB OST ports to place those clusters into VLANs. Each port of the OST interfaces requires an IP address, each in a different subnet. More than one OST interface physical port can be grouped in one virtual port by using bonding. Then the network load is automatically balanced by the round robin method, which provides the best possible redundancy and bandwidth.

Symantek OST plug-ins are available for Windows, Solaris, and LINUX operating systems.

See Chapter 9, “IBM System Storage ProtecTIER with Symantec OpenStorage” on page 433 for more details.

### 3.5.1 Replication for OST

OST does not implement a virtual tape library or a virtual shelf. OST is the disk storage for NetBackup software. Therefore, If you try to configure OST on ProtecTIER system, you do not need to consider shelf or tape cartridges.

OST and VTL each other have different topologies, as shown in Figure 3-21:

- ▶ VTL: Pair (V2.3), hub and up to 12 spokes (ProtectTIER Version 2.4)
- ▶ OST: Implicit hub mesh of up to 12 (ProtectTIER Version 2.5)
- ▶ VTL: Four way many to many bidirectional replication of one to three hubs (ProtectTIER Version 3.1)

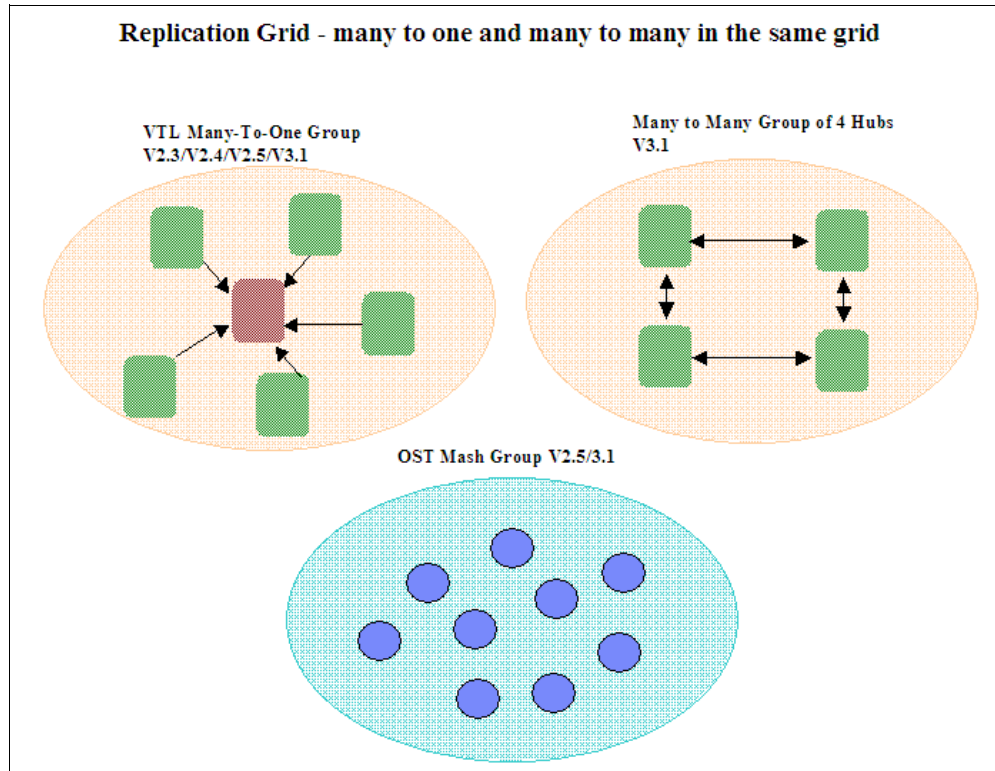


Figure 3-21 Replication topology VTL and OST

Each of the shown topologies can exist in a replication grid separately.

For the default IP addresses of the different TS7600 models, refer to Appendix A, "Installation and implementation checklists" on page 687.

The cabling and ports needed for OST are shown in Table 3-12.

Table 3-12 Required cabling and standard

System	Function	Cabling standard	Dedicated separate subnet	Ports used on TS7650	Amount of cables needed	Amount of cables if clustered	Responsibility
TS7610 SMB VTL	Customer local LAN	1 GB Ethernet	Yes, network 1	eth0	1	N/N	Customer
	Replication network	1 GB Ethernet	2 separate networks	eth1, eth2	2	N/N	Customer
	RAS/ TS3000	100/1000	Yes, 172.31.1.xx	eth3	1	N/N	IBM/Customer depending where TS3000 is located
	Host-Connection	Fibre Channel 8 GB	SAN	Port 0 to 4	4	N/N	Customer
TS7610 SMB OST	Customer local LAN	1 GB Ethernet	Yes, network 1	eth0	1	N/N	Customer
	Replication network	1 GB Ethernet	2 separate networks	eth1, eth4	2	N/N	Customer
	OST Network	10 GB Ethernet	2 separate networks	eth2, eth3	2	N/N	Customer
	RAS/ TS3000	100/1000	Yes, 172.31.1.xx	eth 5	1	N/N	IBM/Customer depending where TS3000 is located
TS7650/ TS7650 G VTL	Customer local LAN	1 GB Ethernet	Yes, network 1	eth0	1	2	Customer
	Cluster Network	1 GB Ethernet	Yes, 10.0.0.5x	eth1, eth4	2	4	IBM
	Replication network	1 GB Ethernet	2 separate networks	eth2, eth5	2	4	Customer
	RAS/ TS3000	100/1000	Yes, 172.31.1.xx	eth3, RSA-II	2	4	IBM/Customer depending where TS3000 is located
	Host-Connection	Fibre Channel 8 GB	SAN, separate zone from back-end ports	Port 0 to 4	4	8	Customer
	Backend Storage	Fibre Channel 8 GB	SAN separate zone from front-end ports and dedicated for TS7650	Port 5 to 8	4	8	Customer only for TS7650G

System	Function	Cabling standard	Dedicated separate subnet	Ports used on TS7650	Amount of cables needed	Amount of cables if clustered	Responsibility
TS7650/ TS7650 G OST	Customer local LAN	1 GB Ethernet	Yes, network 1	eth0	1	2	Customer
	Cluster Network	1 GB Ethernet	yes, 10.0.0.5x	eth1, eth12	2	4	IBM
	Replication network	1 GB Ethernet	2 separate networks	eth7,eth11	2	4	Customer
	RAS/ TS3000	100/1000	Yes, 172.31.1.xx	eth3	1	2	IBM/Customer depending where TS3000 is located
	Host-Connection	10 GB Ethernet	6 separate networks	eth4,eth5, eth6,eth8, eth9,eth10	6	12	Customer
	Backend Storage	Fibre Channel 8 GB	SAN separate zone	Port 5...8	4	8	Customer only for TS7650G

### 3.5.2 Deployment planning guidelines example

The following items needs to be determined:

- ▶ How much data a spoke can back up daily if all spokes do the same volume of work.
- ▶ How much data can a hub back up daily.
- ▶ Approach/configurations:
  - Quantify a maximum configuration of hub and spokes to be deployed: 12 spokes and one hub (two-node cluster).
  - Assume equal backup workload at all spokes.
  - Assume all backup data at each spoke (entire repository) is replicated to the hub.
  - Assume the hub system is also performing local backups (at the DR site).
  - Assume adequate bandwidth between all spokes and hub.
- ▶ Assumptions:
  - There are 8 hour backup windows (hub and spoke).
  - There is a 16 hour replication window.
  - All windows are aligned, meaning the 8 hour backup window is at the same actual time on all 13 ProtecTIER systems (hub and spokes).
  - There is adequate bandwidth between all spokes and hub.
  - There is a 10:1 deduplication ratio throughout the system.
  - The data change rate at the spokes does not saturate the physical reception capabilities at the hub.

- ▶ Maximum workloads assumed:
  - Hub backup:
    - There is an 8 hour backup window.
    - 3.6 TB/hr (1,000 MBps).
    - There is a 28.8 TB *nominal* daily backup.
  - Hub incoming replication:
    - There is a 16 hour replication window.
    - There is 3 TB/hr (850 MBps) replication performance.
    - 48 TB of nominal data is replicated from the spokes.
- ▶ Spoke backup:
  - There is an 8 hour backup window.
  - 48 TB for all 12 spokes = 4 TB daily backup data per spoke.
  - 4 TB / 8 Hrs = 500 GB/hr or 138 MBps sustained for 8 hours
  - A spoke could potentially back up 28.8 TB of nominal data, but can only replicate 4 TB due to configuration constraints.

### Primary (source) repository sizing

Use the ProtecTIER Planner tool to size the repository to enable the required performance. Keep in mind the maximum throughputs of which the configurations are capable. The specifications below are based on a realistic customer workload, assuming there are properly configured back-end disk arrays for the repository:

- ▶ A single node server is capable of 500 MBps of I/O activity (backup or replication).
- ▶ A two-node cluster is capable of:
  - 1000 MBps of I/O activity when only the backup is running.
  - 850 MBps when only replication activity is running.
  - 920 MBps when running concurrently.

The following two examples demonstrate the calculation of required performance from a ProtecTIER system under the two different scenarios: scheduled mode and continuous mode of replication operation:

- ▶ Example A (scheduled replication mode of operation)
  - There is backup activity running for 10 hours a day at a 500 MBps ingest rate.
  - The replication activity is running in a separate time slot of 12 hours at 500 MBps
  - The repository should support a sustained rate of 500 MBps.
- ▶ Example B (replication activity runs in a continuous mode of operation)
  - There is backup activity running 24 x 7 at an ingest rate of 400MBps.
  - Replication runs concurrently and in parallel to the backup activity at 400MBps.
  - The repository should support a sustained rate of 800 MBps.
- ▶ Bandwidth considerations per spoke. For example:
  - A spoke backs up 4 TB per night of nominal data.
  - Approximately 400 GB of new/unique data must be replicated to the hub (assuming a 10:1 deduplication ratio).



- Assuming a 16 hour replication window:
  - 400 GB in 16 hours requires that 7 MBps physical bandwidth be sustained for that time period.
  - This spoke will need about 38% of one OC3 link's capacity.

Capacity planning for the spoke is exactly the same as in previous releases, and must consider the following items:

- ▶ Nominal data
- ▶ Data change rates
- ▶ Retention times
- ▶ Spare capacity

In this example:

- ▶ Each spoke ingests 4 TB of nominal data daily.
- ▶ Each spoke needs 400 GB of physical space daily for local backups (4 TB at a 10:1 deduplication ratio).
- ▶ The total daily space for 27 incrementals is 11.2 TB physical (400 GB \* 27 - 10.8 TB), plus 2 TB for the first full backup (4 TB at 2:1 compression):  $10.8 + 2 = 12.8$  TB.

**Note:** The total physical repository that is needed for each spoke is  $12.8 + 10\%$  spare = 14.08 TB.

## Secondary (target) repository sizing

Consider the following items:

- ▶ The capacity design must accommodate the receipt of replication from all of its spokes, receive local backups (if applicable), and have some spare capacity. The formula for sizing this capacity is as follows:
 

Total hub Capacity = Physical capacity for the replication of all spokes + local backup capacity + spare capacity
- ▶ Critical success factors:
  - You must assess the change rate for each spoke, and then derive the physical data transferred per day.
  - You must assess the retention of both local and replicated data at the hub.
  - The total hub capacity must account for all local backups and all replicated data, plus some spare capacity and room for growth.
  - You may leave more than 10% spare capacity, depending upon the number of spikes and their data types.

In this example:

- ▶ Capacity for replication:
  - There is 24 TB for the first “full” backup for all 12 spokes [ $4 * 12 / 2$  (2:1 compression) = 24]
  - Each spoke will send 400 GB new data daily (4 TB of nominal data at a 10:1 deduplication ratio).
  - 4.8 TB of new data is received daily at the hub from all the spokes (400 GB \* 12 spokes).

- 4.8 TB \* 27 incrementals = 129.6 TB.
- All together: 24 TB + 129.6 TB = 153.6 TB
- ▶ Capacity for hub local backup:
  - There is 14.5 TB for the first “full” backup (29 TB of nominal dat at a 2:1 compression).
  - There is 2.9 TB of space daily for local backups (29 TB of nominal at a 10:1 ratio).
  - 2.9 TB \* 27 incrementals = TB
  - All together: 14.5 TB + 78.3 TB = 92.8 TB

**Attention:** The total space required for the hub in this example is 153.6 TB Replication + 92.8 TB local backup + 10% spare capacity = 271 TB.

For the hub system repository planning, the max ProtecTIER performance assumptions are:

- ▶ A single node server is capable of 500 MBps of incoming replication activity or local backup.
- ▶ Two-node cluster is capable of:
  - 850 MBps of incoming replication activity alone
  - 920 MBps when incoming replication and local backup run concurrently
  - 1000 MBps when performing only local backup

### 3.5.3 Replication policy

Replication policies are a defined set of cartridges from a local repository that need to be replicated to a DR site repository, and establish the rules for when and how that replication will take place.

Policies can be run either manually or automatically (that is, continuously). Whenever replication events are received, policies are continuously run. The most common types of triggers for automatic replication are:

- ▶ Backup
- ▶ Eject
- ▶ Unload cartridge

Manually run policies create replication jobs for all the valid cartridges in their list, whether or not they need to be replicated. Running a policy leads to lining up replication jobs in their respective priority queues, where they wait for resources and the replication time frame to start replicating.

To run a policy, select **Replication** → **Policy** → **Execute policy**. Replication policy is the only means to transfer data from a source repository to a destination repository. A policy is defined on a set of cartridges. Policies define a few attributes for replication:

- ▶ Destination repository
- ▶ Cartridges for replication

Policies have three optional types of priorities:

- ▶ HIGH
- ▶ MEDIUM
- ▶ LOW

The default is LOW for every policy. Replication policies may be enabled or disabled by the user. Policies are repository wide. There can be 254 policies on a single repository. Policies can also be set up by time frame so as to avoid contention with backups.

When writing to a cartridge that is part of a policy, the ProtecTIER system checks whether it must be replicated. The priority of the replication assigned to the cartridge is also considered. The cartridge is then created at the remote shelf. Data segments are transferred from local to a remote cartridge.

### **Policy execution**

Policies can be executed either manually or continuously. Manual execution creates replication jobs for all the valid cartridges.

Policies can be run either manually or automatically (for example, during the time frames set).

In automatic mode, whenever replication events are received, policies run continuously, whether cartridges are being written to by the backup application, ejected from a library (if visibility switching is activated), or unloaded from a drive. These policies start the actual replication activity during the time frame set by the user when created.

Manually run policies create replication jobs for all the valid cartridges included in their list, whether or not they need to be replicated. With both modes of operation, running a policy leads to lining up replication jobs in their respective priority queues where they wait for resources and the replication time frame to start replicating.

Policies should be created and executed in line with the performance capacity of the ProtecTIER system, and with the bandwidth capacity available between the two sites kept in mind. We discuss this concept at length Chapter 5, “IBM System Storage TS7600 with ProtecTIER initial setup” on page 171. In line with this principle, it is typically not a best practice to define a single policy with thousands or even hundreds of cartridges and execute that policy in manual mode, as this action may create a large number of events and replication triggers that can potentially overload the system. Instead, users should use a manual cartridge.

## **3.6 Choosing a replication mode of operation: Time frame versus continuous**

This section provides an operational description and a best practice for a replication mode of operation.

### **3.6.1 Operation**

ProtecTIER offers two modes of operation for the replication activity:

- ▶ Time frame (scheduled replication) with a predefined time window
- ▶ Continuous replication, which runs concurrently with the backup operation

The mode of operation is configured at the source system, and all replication policies defined operate in one of these modes. In almost all cases, scheduled replication is the recommended approach, as it enables the user to accurately plan for performance and better ensure that SLAs are met.

There are several variables, such as payload, available bandwidth, network utilization, time windows, and required SLAs, that might influence the final decision about which mode to choose. Given this mix and understanding the importance to the user of any one of these relative to others, the choice of approaches becomes clearer.

We discuss the two choices in the following sections.

### **Dedicated replication time frame window**

There is *no backup precedence* in ProtecTIER Manager. The window initiates the procedure 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 the matching backup catalog/DB for every 24 hour period. This is the best practice mode.

In dedicated mode, the user defines a time slot during which replication operations will execute throughout the day. After this daily period is exhausted, all replication activities are halted at a consistent state and queued until the following day's window is resumed.

The purpose of dedicated mode is to provide separation between backup/restore and replication operations. Because both backup/restore and replication jobs access the same back-end disk repository, contention between the two elongates both jobs in an unpredictable manner, which might impact the backup SLA and overall RTO. Thus, it is a fundamental requirement that replication tasks will not have to compete with general backup and restore operations over system resources within the dedicated time frame. This approach allows the backups to finish without replication impacting performance/backup window times.

**Note:** The virtual tape library remains fully available to the backup application throughout the dedicated window. No operations are actually blocked. In the event that a backup or restore operation must be performed during this time frame, there will be nothing preventing it.

Because backup/restore operations are not impeded in any way, it is important to be cautious if utilizing the system during the replication window. Even if the replication is not running at nearly the maximum throughput that the system was configured for, there is risk of resource contention if backup/restore jobs are executed that use physical disk spindles used by replication. Therefore, we recommend that no backup/restore operations be carried out in conjunction with the dedicated replication window except in emergency circumstances.

In summary, the dedicated replication time frame window does the following actions:

- ▶ Defines the start and end of the replication time.
- ▶ Begins replication activity as soon as the time window begins. It will be prioritized the same as backup/restore and compete on the same resources.
- ▶ Stops replication activity as soon as the time window ends.
  - Each cartridge in transit stops at a consistent point after the window ends.
  - Will not replicate at all outside of the time window scope.

**Note:** That there is no calendar schedule in ProtecTIER V2.3. The time frame that is established is available for all days of the week and applies to all replication policies.

## Continuous replication

Continuous/simultaneous precedence to back up in the ProtecTIER Manager is rarely used, but is available for unique scenarios. This mode enables concurrent operations and has no specific replication window. This mode might be considered when a user has consistently lower bandwidth and when the operation calls for few backup windows that are spread throughout the day. Care must be taken when choosing this option, as there will be an impact on backup performance. Read performance of the disk subsystem will be shared between the different processes needed to deduplicate the backup stream and read the data to be replicated. This mode is not the recommended one, as it stresses the system resources and impairs performance and bandwidth utilization. It will only apply to certain user environments that meet specific criteria, such as when a user has consistent low network bandwidth available and the operation policy calls for few backup windows spread throughout a day.

In this mode of operation, no specific window is defined and replication activities may run at any time continuously throughout the day. However, there is a core difference in how system resources are managed in this mode compared to a dedicated window.

With a dedicated window, there is no weighing of which internal service request will be executed at a higher priority. Replication and backup/restore resource requests are equally weighted and processed in a first-in, first-out fashion. Although the overall system throughput (backup/restore plus replication) can reach the maximum configured rate, because there is no precedence given to backup/restore jobs, the backup duration might vary from day to day, potentially impacting the overall SLA. For example, if a system is capable of delivering 500 MBps and backups are currently running at 75 MBps, replication will not be throttled, and could use roughly up to 425 MBps, which is the maximum capability of the system. However, when backup throughput increases to 200 MBps, replication will be throttled down to lower performance consumption until the backup throughput decreases. In this case, the minimum throughput that is expected from replication would be roughly 75 MBps (15% of 500 MBps), and the maximum (with no backup activity) will be about 425 MBps (85% of 500 MBps).

The rate-control function is designed so that when both replication and backup operations are running concurrently, it will respond quickly to changes in the backup activity/pattern, so the overall system resources will be allocated and used in the most efficient way between both activities.

Conversely, in continuous replication mode, there is a dynamic resource-governing mechanism (rate control) that gives precedence to backup/restore requests and throttles down replication traffic whenever backup/restore activity increases above what the system calls *idle* until the backup/restore workload drops below that idle threshold. During these time frames (called *busy*), when backup/restore activity picks up, the replication activity receives limited resources equal to roughly 15% of the system's aggregate performance capability. Both the threshold and the minimum allocated resource percentage are preconfigured settings, based on the specific ProtecTIER node expected performance during deployment. The rate control is a preconfigured dynamic function, based on real-time machine I/O activity statistics and the actual server capabilities. There is also a safe-ground mechanism built into this function that reserves about 15% (preconfigured) of the machine resources for backup/restore activity at times when only replication activity is running, which means that the maximum performance of the replication activity during backup idle state is roughly 85% of the available machine resources. This is done to prevent race conditions (over machine resources) when backup/restore activity starts, and allows for the system to dynamically control the allocation of machine resources between the activities.

In summary, continuous replication does the following functions:

- ▶ Data will automatically start replicating to a remote repository as soon as it is written to a cartridge.
- ▶ Replication will run faster if the system is considered idle (no backup/restore activity).
- ▶ Replication will be prioritized lower compared to backup/restore if running concurrently. Both will compete on the same CPU and disk resources.
- ▶ Typically, it requires a significantly larger system to enable concurrent operations.

### 3.6.2 Best practices for choosing the mode of operation

As discussed, deploying ProtecTIER Replication with a dedicated window will almost always be the best option for the user. The sum of backup/restore and replication throughput, measured nominally, cannot exceed the overall throughput that the source system is capable of delivering. Thus, when operating in continuous mode where replication tasks compete for resources with backup/restore operations, it can be difficult to ensure that all replication jobs complete within the mandated RTO time frame. Thus, as a best practice, use the dedicated replication time frame window mode of operation. A primary benefit of the dedicated replication time frame window is the ability to strictly classify when ProtecTIER uses the network infrastructure, allowing the user to accurately isolate various applications' usage of the network. This mode of operation is aligned with the current backup and DR activity, in which users typically manage a specific backup window and schedule cloning or vaulting jobs that follow the backup window.

### 3.6.3 Remote repository

By default, the remote repository serves as a replication target only (without local backup activity). The same tools and considerations for the primary system repository should be used, with the assumption that the required replication performance is equal to the required backup performance. The use of nominal data being replicated makes this planning stage simpler and almost similar to the primary system, except that at the DR site, the ProtecTIER nodes process the replication data stream while the primary deals with backup data stream ingest.

For the remote system repository planning, the maximum ProtecTIER performance assumptions are that the:

- ▶ Single-node server is capable of 600 MBps of incoming replication activity
- ▶ Two-node cluster is capable of 850 MBps of incoming replication activity

For sizing, evaluate the amount of data to be replicated. The following factors should be considered when sizing both the primary and the secondary DR site repositories:

- ▶ Users may replicate their entire repository to their DR site or may select sets of cartridges to be replicated to the secondary site.
- ▶ Evaluate the amount of data to be replicated:
  - If all data in the primary repository is planned to be replicated, the remote repository should be sized at least to the same size as the local repository.
  - If only part of the data is replicated, the remote repository capacity can be reduced accordingly.

- ▶ If the user's DR plan calls for using the DR site as the primary site for backup activity:
  - Evaluate the amount of data to be locally backed up and its retention period at the remote DR site, during a disaster while the primary site is unavailable.
  - Add this to the repository size needed for replication alone to get the final capacity size needed.

Most likely, you will find that your data falls into several categories and requires several different policies, for example:

- ▶ Non-critical data would not be replicated.
- ▶ Data not easily reconstructed or that has a impact on business operations might have a weekly or daily replication schedule.
- ▶ Data that would have a severe impact on business is considered critical and is replicated much more frequently.

### 3.7 Tips for using the visibility switch control feature

Cartridges can only be present and available at one location at a time. Visibility control is the means to initiate a move (as opposed to a copy) of a tape/cartridge from one location to another. ProtecTIER Replication software can emulate a move using the VTL import/export slots (Table 3-13) the same way that a physical tape library does.

*Table 3-13 Detailed information about import/export slots in a VTL library*

Column	Definition
Imp/Exp No.	The number of the import/export slot.
Address	The import/export slot's address number.
Barcode	If the import/export slot contains a cartridge, this column displays the cartridge's barcode.
Capacity	If the import/export slot contains a cartridge, this column displays the cartridge's estimated data capacity in megabytes.
Data Size	If the import/export slot contains a cartridge, this column displays, in megabytes, the amount of nominal data currently stored on the cartridge.

This method of moving cartridges from one location to another is useful in one domain backup environments. For example, backup applications such as Symantec NetBackup (NBU), Legato, and Backup, Recovery, and Media Services (BRMS) can be deployed in a single domain environment (check the specific backup application's literature for more details). In these environments, there is one catalog across multiple sites. If a user uses the VTL export/import slots to move tapes from one site to the other, the catalog will always be updated and aware of the location of all cartridges.

If you are planning on using the VTL export/import slots, you might want to increase the number of such available slots to the maximum of 1022 (the default when creating a library is 8).

The ProtecTIER Manager Imports/Exports tab displays detailed information about the import/export slots in the selected library.

## ProtectTIER VTL cartridge handling concepts

This section outlines the manner in which ProtectTIER manages cartridges, library slots, and import/export slots of the virtual library. When it comes to using the ProtectTIER cartridge Visibility Switch Control feature, the user should consider the following items:

- ▶ Cartridge slots inside the ProtectTIER VTL

A cartridge imported into a VTL library will be placed in an empty slot. The operation of moving a cartridge from an import/export slot into a specific library slot fails if there are no free slots in the library at that time. Therefore, the remote DR site libraries should be created with enough free slots to host *replicated cartridges* before users attempt to move them into a library following a DR scenario when the DR site becomes active.

- ▶ Import/export (I/E) slots

These VTL slots are used to inject or eject cartridges to or from a specific virtual library inside the repository. The default number of import/export slots configured during the *create library* operation is eight slots. Consider increasing this number to the maximum number of import/export slots supported by the backup application (up to 1022 per ProtectTIER virtual library), as this increases the capability of the backup application process to import or export cartridges to or from a ProtectTIER VTL library.

- ▶ Visibility control switch

Although this ProtectTIER feature can be used within any backup application environment, this mode is more suitable for use when the backup application master server controls both the primary and the remote libraries, that is, a single-domain backup environment, where multiple sites share the same catalog or database. In this scenario, the backup application master server knows the content of the ejected cartridges at the primary site and the injected cartridges at the secondary (DR) site, because the shared catalog or database already contains entries for them. Furthermore, the catalog is always updated with the locations of all cartridges at both sites.

- ▶ When to export cartridges

Cartridges should be exported (ejected from the library) immediately after the backup is completed to prevent the backup application from using it for subsequent backups. Use the same strategy for exporting (ejecting) cartridges that is being practiced when handling physical tapes.

- ▶ When a cartridge is available at the remote site, using the Visibility Switch Control feature

After the exported/ejected cartridges from the primary site are fully synchronized (have finished their replication) at the remote site, they will immediately be moved to the remote library import/export slots. Cartridges that are not in sync between the primary and remoter repositories must transfer (replicate) their data in full, and only then are they moved to the remote library import/export slots. The user can check the remote library for cartridges that completed this stage and were moved into the library import/export slots. The user can monitor the replication tasks view at the primary location's ProtectTIER Manager to learn about the status of these cartridges that are still transferring data to the secondary site. Alternatively, the user can create a cartridge report at the remote site to learn about which cartridges already finished replication and were moved into the library import/export slots.



## 3.8 Planning for OST

With OpenStorage, ProtecTIER can be integrated with NetBackup to provide backup to disk without having to emulate traditional tape libraries. Using a plug-in that is installed on an OST-enabled media server, ProtecTIER can implement a communication protocol that supports data transfer and control between the backup server and the ProtecTIER server. Therefore, to support the plug-in, ProtecTIER implements a storage server emulation (see Figure 3-22).




Figure 3-22 OST configuration map

OST has follow benefits.

- ▶ Treat disk as disk
  - Avoids the limitations of tape emulation.
  - Provides enhanced functionality only possible with disk.
- ▶ Tight integration
  - NetBackup is aware of all backup image copies and manages the creation, movement and deletion of all backup images.

See Chapter 9, “IBM System Storage ProtecTIER with Symantec OpenStorage” on page 433 for more details.

## 3.9 Planning for cartridges

Prior to discussing the various elements of the TS7650G or TS7650 and TS7610 that monitor and manage capacity, we should have a background discussion.

IBM System Storage TS7650, TS7610, and TS7650G with ProtecTIER are designed to mimic the behavior (and management) of a traditional tape library as closely as possible. They are designed to be used intuitively by the backup administrator, who is typically trained and experienced in managing tape capacity. IBM System Storage TS7600 with ProtecTIER introduces certain challenges not associated with traditional tape libraries:

- ▶ Nominal capacity cannot be directly observed or calculated.
- ▶ Nominal capacity can fluctuate over time.

IBM System Storage TS7650, TS7610, and TS7650G with ProtecTIER have internal mechanisms to manage the nominal capacity of the system and communicate with the backup administrator in the language of tape.

### 3.9.1 Capacity management in the traditional tape library paradigm

To explain how the IBM System Storage TS7600 with ProtecTIER adheres to tape management principles, here we review key points of managing tape capacity.

The overarching objective of the backup administrator is to make sure that there is enough capacity for the foreseeable future. In the tape backup world, backup administrators pay close attention to the number of tapes available for the day's operations. Through the backup application console, the number of available cartridges is visible. The backup application cannot directly see the capacity of a cartridge, but the administrator knows the capacity based on cartridge type and compression ratios. The administrator must answer the following questions each day: How many tapes are available, and what is the capacity of each one? By calculating the total available capacity, the administrator knows whether the daily payload will fit on the available tapes.

Over time, traditional tape libraries reach an equilibrium state. Tapes are recycled, which means that they are put back into the pool of cartridges available for new backups. In equilibrium, the number of cartridges that are returned to the available pool roughly equals the number required for the given day's payload.

What does this mean in terms of impacting the backup operator? It means that capacity shortages are usually easy to predict. Typically, if the number of new tapes used exceeds the number of tapes being returned to the pool, capacity shortages will happen.

One other key point to note is that in the physical tape world, there are early warning (EW) signals provided to the backup application by the tape drive when a tape cartridge is nearing its end. This signal allows the backup application to change to a fresh cartridge efficiently. This EW signal is relevant to understanding IBM System Storage TS7600 with ProtecTIER capacity management.

### 3.9.2 ProtecTIER systems versus traditional tape libraries

In many ways, the IBM System Storage TS7600 with ProtecTIER virtual tape solution is similar to traditional tape libraries. It is designed for easy use by backup operators who do not have to learn a new paradigm from an operational perspective.

Once installed, IBM System Storage TS7600 with ProtecTIER behaves and is managed like a standard library. It has a nominal capacity that is available for use by the backup application. This capacity is represented by a certain number of cartridges, each with a given capacity. Just as with real tape libraries, with the IBM System Storage TS7600 with ProtecTIER, the backup administrator uses the number of tapes available and the capacity of each as the key indicators. Just as with real libraries, IBM System Storage TS7600 with ProtecTIER reaches an equilibrium point in which cartridges are returned to the scratch pool at roughly the same rate at which they are consumed by the new day's backup totals. However, while there are many capacity management similarities between IBM System Storage TS7650, TS7610, and TS7650G with ProtecTIER and traditional tape libraries from the backup administrator's point of view, there are also differences.

Cartridge capacities within the IBM System Storage TS7600 with ProtecTIER fluctuate. As a result, the number of tapes used per day fluctuates, even when the payload stays constant. Space reclamation is also different. New data sets have a bigger impact on capacity, and capacity expansion requires additional steps that must be learned as part of the operational process.

### 3.9.3 How the TS7600 with ProtecTIER manages changes in nominal capacity

The initial factoring ratio is always an estimate and the actual data change rate of data that enters the IBM System Storage TS7600 with ProtecTIER fluctuates each day and over time. The result is that the *nominal capacity* of the overall system fluctuates. IBM System Storage TS7600 with ProtecTIER manages these changes in a nominal capacity with an internal *learning algorithm*. The learning algorithm enables changes in nominal cartridge capacities that reflect changes to the system-wide nominal capacity.

**Note:** As the size of the cartridges configured in the IBM TS7600 with ProtecTIER depends on the actual available physical space in the repository, the actual factoring ratio, and changes in the data structure that is stored continuously, the shown capacity is not a fixed value. The formula shown below gives guidance about how many cartridges can be configured for a repository:

$$\text{Cartridge size} = \text{PhysicalSizeRepository} \times \text{FactoringRatio} / \text{NumberOfCartridges}$$

The purpose of this algorithm is to help ensure that all capacity is fully used. It also provides an intuitive way for a backup administrator to manage fluctuating capacity. The results of the learning algorithm are visible to the backup administrator through the usage of cartridges on a daily basis. If capacities decline, the EW of an individual tape arrives sooner, which in turn requires the backup application to request a new tape. The overall effect on a daily basis is that more tapes are consumed. Inversely, if capacities increase, the EW of a cartridge arrives later, and less tapes are used overall. Just as with a traditional tape library, the administrator is able to track the tape usage statistics to manage the system capacity. Thus, the paradigm of traditional tape library management is closely mimicked.

### 3.9.4 Managing capacity fluctuations

As mentioned, cartridge capacity changes to reflect the system-wide shift in nominal capacity. In Table 3-14, the factoring ratio at *equilibrium* is greater than the factoring ratio that was used when the IBM System Storage TS7600 with ProtecTIER was first installed. There are 1000 tape cartridges, but because the factoring ratio has stabilized at a higher number (12:1 versus 10:1), the nominal capacity has increased from 100 TB to 120 TB. To accommodate the change, the capacity per cartridge has increased from 100 GB to 120 GB. The IBM System Storage TS7600 with ProtecTIER handles this situation by using the learning algorithm mentioned earlier. The backup application still manages 1000 cartridges in its catalog, and because it will only *change* cartridges when an *end of cartridge* signal is sent by the IBM System Storage TS7600 with ProtecTIER, the increase in cartridge capacity is transparent to the backup application.

Table 3-14 Effect of learning algorithm with higher than expected factoring ratio

Day	Physical capacity	Number of cartridges	Factoring ratio	Nominal capacity	Capacity per cartridge
Day 1	10 TB	1000	10:1	100 TB	100 GB
Day 30	10 TB	1000	12:1	120 TB	120 GB

In Table 3-15, the factoring ratio at equilibrium is less than the factoring ratio that was used when the TS7600 with ProtecTIER system was first installed. As you can see in Table 3-15, there are 1000 tape cartridges, but because the factoring ratio has stabilized to a lower value (8:1 versus 10:1), the nominal capacity has decreased from 100 TB to 80 TB. To accommodate the change, the capacity per cartridge has decreased from 100 GB to 80 GB.

*Table 3-15 Effect of learning algorithm with a lower than expected factoring ratio*

Day	Physical capacity	Number of cartridges	Factoring ratio	Nominal capacity	Capacity per cartridge
Day 1	10 TB	1000	10:1	100 TB	100 GB
Day 30	10 TB	1000	8:1	80 TB	80 GB

As the cartridge size changes, the EW signal arrives sooner or later than originally. In the example shown in Table 3-14 on page 117, the EW for each cartridge occurs 20 GB later on day 30 than on day 1, allowing more data to fit on a given cartridge. In the example shown in Table 3-15, the EW for each cartridge occurs 20 GB earlier on day 30 than on day 1, allowing less data to fit on a given cartridge. As a result of the learning algorithm, more or fewer tapes will be consumed during a given day's workload.

**Note:** Backup administrators for ProtecTIER must keep track of the number of cartridges, as this number is used as a key indicator of capacity fluctuations.

### 3.9.5 Capacity management implications for the TS7650, TS7610, or TS7650G

Capacity management begins from the moment that the TS7650G starts to accept backups. Different things happen at different phases of the TS7650G implementation.

### 3.9.6 Capacity management implications: Initialization phase

During the first weeks of an IBM System Storage TS7600 with ProtecTIER implementation, the daily fluctuations in capacity are more pronounced. This is normal system behavior as IBM System Storage TS7600 with ProtecTIER learns the true data change rate. Do not be alarmed during this phase when cartridge capacity oscillates, sometimes significantly. After the system runs a full backup cycle (for all data sets that it will manage), the capacity changes should stabilize.

### 3.9.7 Management of IBM System Storage TS7600 with ProtecTIER

From an ongoing management perspective, the backup administrator must be aware of the following situation when running an IBM System Storage TS7600 with ProtecTIER virtual tape solution within the corporate backup environment.

As cartridge size changes, the EW signal arrives sooner or later than originally. In the example in Table 3-14 on page 117, the EW for each cartridge occurs 20 GB later on day 30 than on day 1, allowing more data to fit on a given cartridge. As a result, more or fewer tapes will be consumed during a given day's workload. Therefore, backup administrators for IBM System Storage TS7650, TS7610, and TS7650G with ProtecTIER must keep track of the number of cartridges used as a key indicator to indicate capacity fluctuations.

### 3.9.8 Capacity management implications: Adding new data sets to an existing IBM System Storage TS7600 with ProtecTIER

Often during the initial phase, or sometime thereafter, you will decide to send more data to a TS7650, TS7610, or TS7650G. Although this is a common occurrence, this situation creates new implications of which the backup administrator must be aware.

A new data stream will have a high change rate (given that all of the data is new to the IBM System Storage TS7600 with ProtecTIER). This causes an increase in the system-wide change rate and a decrease in the nominal capacity, because the factoring ratio is going to decrease. As the new data set runs through a full cycle, the nominal capacity might or might not return to what it was previously, depending on the data change rate of the new data set. Given the variability that is inherent in this situation, you must be aware of the phenomenon and understand the impact. The best way to add new data streams is to first sample the data to project the likely impact. In some cases, this action might create a need for more physical disk that might or might not have been built in to the original TS7650, TS7610, or TS7650G design.

### 3.9.9 Space reclamation and steady state

As mentioned previously, with real tape systems, cartridges expire and are returned to the available pool of cartridges to be used for new backups. This process is easy to manage and understand: When a cartridge is returned to the available pool, its full capacity is available for new backups. See 2.6.2, “Steady state” on page 34 for more details.

From the backup application point of view, the process of tape recycling is exactly the same in IBM System Storage TS7600 with ProtecTIER: Cartridges expire and are returned to the available pool for new backups. However, the underlying process of space reclamation in the IBM System Storage TS7600 with ProtecTIER is unique.

As soon as a cartridge begins to receive new data at the beginning of its media, IBM System Storage TS7650, TS7610, or TS7650G with ProtecTIER knows that it has been recycled by the backup application. Any data elements that the recycled cartridge alone references (that is, elements that are not used by other cartridges) become available for space reclamation, so the physical space is then returned to the repository as ready for new data. In equilibrium, the rate at which old data expires is approximately equal to the rate at which new unique data is written to the IBM System Storage TS7600 with ProtecTIER repository. The implication of this process is that the actual physical capacity that is returned to the available pool when a cartridge is recycled is not readily observable. It is usually a fraction of the nominal cartridge size, on the order of 5 - 10%.

### 3.9.10 Summary of TS7600 with ProtecTIER capacity management

In summary, the benefit of disk savings enabled by the IBM System Storage TS7600 with ProtecTIER also introduces new capacity management challenges of which the administrator must be aware. IBM System Storage TS7600 with ProtecTIER has an internal learning algorithm that allows nominal capacity to adjust to changing data change rates. This is done while maintaining a traditional tape management paradigm. System capacity changes are reflected in cartridge capacities. This allows the TS7650, TS7610, or TS7650G to maintain the *language of tape* to the backup administrator as follows:

- ▶ If nominal capacity increases, fewer tapes are consumed per day.
- ▶ If nominal capacity decreases, more tapes are consumed per day.





# Hardware planning for IBM System Storage ProtecTIER

In this chapter, we describe which options and features can be configured with the IBM System Storage TS7650G ProtecTIER Deduplication Gateway (3958-DD4), the IBM System Storage TS7650 ProtecTIER Deduplication Appliance (3958-AP1), and the IBM System Storage TS7610 ProtecTIER Deduplication SMB Appliance (3959-SM1).

We provide information about which configuration options are available and useful. We also discuss the installation tasks required and the responsibilities assigned to the IBM System Service Representative (SSR) and to the customer.

This chapter discusses the following topics:

- ▶ Hardware and software components
- ▶ Configuration options
- ▶ Installation tasks

## 4.1 General overview of the TS7610, TS7650, and TS7650G

There are three models that run ProtecTIER for open environments:

- ▶ IBM System Storage TS7610 ProtecTIER Deduplication SMB Appliance (3959-SM1)
- ▶ IBM System Storage TS7650 ProtecTIER Deduplication Appliance (3958-AP1)
- ▶ IBM System Storage TS7650G ProtecTIER Deduplication Gateway (3958-DD4)

The IBM System Storage TS7610 ProtecTIER Deduplication SMB Appliance (3959-SM1) is a rack unit (2U) bundled appliance. It comes with the disk storage and the ProtecTIER software in a 2U device. The rail kit it is 3U. Two levels of disk capacity can be ordered:

- ▶ 4.0 TB
- ▶ 5.4 TB

For the 3959-SM1, the repository and one virtual library is already configured, as well as the file systems over the disk storage. The customer is responsible for hardware installation and software configuration. After your host is attached and zoned to the 3958-SM1 server, you can set up your applications on your host.

The IBM System Storage TS7650 ProtecTIER Deduplication Appliance (3958-AP1) comes with disk storage in the same frame, and the ordered configuration will be pre-cabled in manufacturing. The IBM System Storage TS7650 ProtecTIER Deduplication Appliance (3958-AP1) can be ordered as a single node or as a two-node cluster configuration.

The disk controller for the 3958-AP1 server is an IBM System Storage DS5020 Express and the disk expansion module is an IBM System Storage EXP810 Storage Expansion Unit. The DS5020 storage is configured for optimal performance at manufacturing. Three levels of disk capacity can be ordered:

- ▶ 7 TB
- ▶ 18 TB
- ▶ 36 TB

For the 3958-AP1 server, the repository is already configured, as well as the file systems over the disk storage. After the 3958-AP1 server is installed by a IBM SSR and the ProtecTIER Manager is installed on a workstation, you can start to work on the 3958-AP1 server. After your host is attached and zoned to the 3958-AP1 server, you can set up your application on your host.

The IBM System Storage TS7650G ProtecTIER Deduplication Gateway (3958-DD4) can be ordered with IBM System Storage systems. By combining the advantages of IBM disk and tape storage subsystems, a high reliability solution is delivered.

The IBM System Storage TS7650G ProtecTIER Deduplication Gateway (3958-DD4) can also be ordered without back-end disk storage. The 3958-DD4 server also supports non-IBM disk storage.

**Note:** For a list of disk subsystems that are supported by the TS7650G, refer to the interoperability matrix, found at:

[ftp://service.boulder.ibm.com/storage/tape/ts7650\\_support\\_matrix.pdf](ftp://service.boulder.ibm.com/storage/tape/ts7650_support_matrix.pdf)

Refer to the TS7650G in the System Storage Interoperation Center (SSIC) for a list of supported environments:

<http://www-03.ibm.com/systems/support/storage/config/ssic/index.jsp>



The functionality of the 3959-SM1, 3958-AP1, and the 3958-DD4 servers are identical, but there are differences in performance and maximum capacity. Table 4-1 shows the specifications for a single server.

Table 4-1 Differences for the 3959-SM1, 3958-AP1, and the 3958-DD4

Component	3959-SM1	3958-AP1	3958-DD4
Number of processors	4	16	32
Amount of memory	24 GB	32 GB	64 GB
Number of virtual libraries	Up to 4	Up to 12	Up to 16
Number of virtual tape drives	Up to 64	Up to 256	Up to 256
Number of virtual cartridges	Up to 8,192	Up to 128,000	Up to 512,000
Amount of supported disk capacity	Up to 5.4 TB	Up to 36 TB	Up to 1 PB
IBM Path Failover technology	Yes, <sup>1</sup> with IBM TS3500 virtual tape library definition and multipath driver	Yes, <sup>1</sup> with IBM TS3500 virtual tape library definition and multipath driver	Yes, <sup>1</sup> with IBM TS3500 virtual tape library definition and multipath driver
Two-node cluster configuration	No	Yes	Yes
IP-based replication configuration	Yes	Yes	Yes
Disaster recovery failover	Yes	Yes	Yes
Flexible disk-based storage options	No	No	Yes
Sustained inline throughput <sup>2</sup>	Up to 80 MBps	Up to 500 MBps	Up to 1400 MBps one node
Data reduction	Up to 25:1 or more	Up to 25:1 or more	Up to 25:1 or more
Preinstalled disk storage	Yes	Yes	No
Server comes in a rack	No	Yes	No

1. Not for the Quantum P3000 virtual library.

2. Depending on the back-end storage attached and workload.

You might find references to the 3958-DD1 server in older documentation. The 3958-DD1 server is the first generation of the IBM System Storage TS7650G ProtecTIER Deduplication Gateway. The 3958-DD1 server was withdrawn from marketing in 2009 and the replacement is the 3958-DD3 server. The newest server is 3958-DD4 server, which was announced at the end of October of 2010 and replaces the 3958-DD3 server. The difference is the hardware on which the ProtecTIER Enterprise Edition V2.5 Base Software is installed. The 3958-DD3 server is based on the IBM eServer™ xSeries® x3850 M2/x3950 M2, and the CPU used with this server is a 4X hex cores (24 cores) CPU (the 3958-DD1 server uses a 4X quad core (16 cores) CPU; the 3958-DD4 server is based on IBM eServer xSeries X5, and the CPU is 4X octo-cores (32 cores).

**Note:** The following sections are written for the 3958-DD4 with ProtecTIER V2.5. Many sections are valid for both the 3958-DD4 and the 3958-DD3 servers or explained for 3958-DD3 server as well. Read all the material to better understand the TS7650 and TS7650G. For example, when you order a 3958-AP1 server in a two-node cluster configuration, it comes preconfigured from manufacturing.

## 4.2 Hardware and software components for 3959-SM1, 3958-AP1, and 3958-DD4

This section describes in detail the components of the ProtecTier systems

### 4.2.1 The 3958-SM1 server features

The IBM System Storage TS7610 ProtecTIER Deduplication SMB Appliance solution, composed of IBM System Storage ProtecTIER Entry Edition V2.5 software and the IBM System Storage TS7610 Deduplication SMB Appliance (3959-SM1) hardware, is designed to address the disk-based data protection needs of enterprise data centers.

The IBM System Storage ProtecTIER Entry Edition V2.51 software and Red Hat Enterprise Linux 5.4 Advanced Platform 64-bit software are installed by the IBM SSR on the 3959-SM1.

The TS7610 ProtecTIER Deduplication SMB Appliance offers:

- ▶ Inline data deduplication powered by HyperFactor technology.
- ▶ Supports replication. A TS7610 system may participate in topology groups with TS7650 Appliances, TS7650G Gateways, and other TS7610 SMB Appliance systems.
- ▶ Fibre Channel ports for host and server connectivity.
- ▶ Inline data deduplication performance of up to 80 MBps support.
- ▶ Two storage choices and options, available with a 4.0 TiB and 5.4 TiB repository.
- ▶ Virtual tape emulation of up to four virtual tape libraries, up to 64 virtual tape drives on each node, and up to 8192 virtual cartridges.
- ▶ Emulation of the IBM TS3500 tape library with IBM LTO Ultrium 3 tape drives.
- ▶ Emulation of the Quantum P3000 tape library with DLT7000 tape drives and IBM LTO Ultrium 2 tape drives.
- ▶ Emulation of the DTC VTF 0100 tape library with DLT7000 tape drives and IBM LTO Ultrium 2 tape drives.
- ▶ Emulation of the IBM V-TS3500 tape library with IBM LTO Ultrium 3 tape drives.
- ▶ IBM Path Failover technology.
- ▶ LUN masking.
- ▶ Optional replication support to allow virtual cartridges to be copied to a target (remote) ProtecTIER based on defined policies.
- ▶ The OpenStorage (OST) feature implements the functionality to directly connect the ProtecTIER server to hosts running Symantec NetBackup.

## 4.2.2 3959-SM1 server characteristics

The 3959-SM1 is based on a 2U rack-drawer IBM server. It supports the necessary performance required for midrange-level data protection and hardware redundancy.

The 3959-SM1 server is shipped with the following characteristics:

- ▶ One Intel Xeon 2.33 GHz 24M 4-core 6.4 Gb 130 W Turbo processors
- ▶ 12 x 1 TB SATA disk drives
- ▶ Six 4 GB dual in-line memory modules (DIMMs)
- ▶ Two internal 300 GB 2.5-inch 10 K RPM SAS Drives (RAID 1 mirror)
- ▶ An internal ServeRAID M5015 SAS/SATA controller with battery
- ▶ One dual port 1 Gbps 10/100/1000Base-TX PCIe Ethernet adapter
- ▶ SAS expander printed circuit board (PCB)
- ▶ Dual port Emulex Fibre Channel HBA, available with 4 Gb and 8 Gb speed
- ▶ External USB DVD Drive
- ▶ RAID battery backup unit (BBU)
- ▶ Two redundant power supplies
- ▶ 10 cooling fans

Figure 4-1 shows the front and back view of a 3959-SM1.



Figure 4-1 Front and back view of the 3959-SM1 Appliance Express model

## 4.2.3 3958-AP1 server features

The IBM System Storage TS7650G ProtecTIER Deduplication Appliance solution, composed of IBM System Storage ProtecTIER Appliance Edition V2.5 software and the IBM System Storage TS7650 Deduplication Appliance (3958-AP1) hardware, is designed to address the disk-based data protection needs of enterprise data centers.

The TS7650 ProtecTIER Deduplication Appliance offers:

- ▶ Inline data deduplication powered by HyperFactor technology.
- ▶ A powerful multicore virtualization and deduplication engine.
- ▶ Clustering support for a higher availability (available only for the 36 TB Appliance).
- ▶ Fibre Channel ports for host and server connectivity.
- ▶ Inline data deduplication performance of up to 500 MBps support.
- ▶ Virtual tape emulation of up to 12 virtual tape libraries, up to 256 virtual tape drives on each node, and up to 128,000 virtual cartridges.
- ▶ Emulation of the IBM TS3500 tape library with IBM LTO Ultrium 3 tape drives.
- ▶ Emulation of the Quantum P3000 tape library with DLT7000 tape drives and IBM LTO Ultrium 2 tape drives.
- ▶ Emulation of the DTC VTF 0100 tape library with DLT7000 tape drives and IBM LTO Ultrium 2 tape drives.
- ▶ Emulation of the IBM V-TS3500 tape library with IBM LTO Ultrium 3 tape drives.
- ▶ DS5020 cache subsystem with 7 TB, 18 TB, or 36 TB in single node configuration.
- ▶ DS5020 cache subsystem with 36 TB in dual node cluster configuration.
- ▶ IBM Path Failover technology.
- ▶ LUN masking.
- ▶ Optional replication support to allow virtual cartridges to be copied to a target (remote) ProtecTIER two-node cluster based on defined policies.
- ▶ The OpenStorage (OST) feature implements the functionality to directly connect the ProtecTIER server to hosts running Symantec NetBackup.

#### 4.2.4 3958-AP1 server characteristics

The 3958-AP1 server is based on a 4U rack-drawer IBM System x3850 X5 type 7145 plus DS5020 storage subsystem. It supports the necessary performance required for enterprise-level data protection and hardware redundancy.

The IBM System Storage ProtecTIER Appliance Edition V2.5 software and Red Hat Enterprise Linux 5.4 Advanced Platform 64-bit software are installed by an IBM SSR on the 3958-AP1 server.

The 3958-AP1 server is shipped with the following characteristics:

- ▶ Two Intel Xeon X7560 2.27 GHz 24M 8-core 6.4 Gb 130 W Turbo processors (16 cores in total, 32 threads with hyper-threading)
- ▶ Four memory cards, with 2 x 4 GB DDR3 DIMMs per memory card, total amount of memory 32 GB
- ▶ Two internal 300 GB 2.5-inch 10 K RPM SAS Drives (RAID 1 mirror)
- ▶ Internal ServeRAID M5015 SAS/SATA controller with battery
- ▶ Two dual-port Emulex LPe12002 8 Gbps Expansion Card (front end ports) for VTL configuration
- ▶ Two dual-port Qlogic QLE2562 8 Gbps Expansion Card (back-end ports)
- ▶ in VTL configuration, one Intel Pro/1000 PT Quad Port Ethernet Server Adapter

- ▶ in OST configuration, two Intel Pro/1000 PT Quad Port Ethernet Server Adapter or two QLE2562 dual port 10 Gb CNAs
- ▶ Two 1975 W standard hot swap Redundant Power supplies, Hot Swap Redundant fans, Hot Swap HDDs, and Hot Swap memory
- ▶ SATA CD/DVD-ROM
- ▶ Integrated Management Module (IMM) as the successor of Remote Supervising Adapter II (RSAIL)

You can see the front view of the 3958-AP1 server in Figure 4-2 on page 129.

You can see the rear view of the 3958-AP1 server, including the cards and port assignment for VTL and OST, in Figure 4-3 on page 130 and Figure 4-4 on page 130.

**Note:** For detailed information about Integrated Management Module (IMM), also visit the following address:

<http://www-947.ibm.com/systems/support/reflib/imm/index.html>

The *User's Guide for Integrated Management Module* can be found at:

[ftp://ftp.software.ibm.com/systems/support/system\\_x\\_pdf/imm\\_users\\_guide\\_60y1465.pdf](ftp://ftp.software.ibm.com/systems/support/system_x_pdf/imm_users_guide_60y1465.pdf)

## 4.2.5 3958-DD4 server features

The IBM System Storage TS7650G ProtecTIER Deduplication Gateway solution, composed of IBM System Storage ProtecTIER Enterprise Edition V2.5 software and the IBM System Storage TS7650G Deduplication Gateway (3958-DD4) hardware, is designed to address the disk-based data protection needs of enterprise data centers.

The TS7650G ProtecTIER Deduplication Gateway offers:

- ▶ Inline data deduplication powered by HyperFactor technology.
- ▶ A powerful multicore virtualization and deduplication engine.
- ▶ Clustering support for a higher availability.
- ▶ Fibre Channel ports for host and server connectivity.
- ▶ Inline data deduplication performance of up to 2000 MBps (two-node cluster) support.
- ▶ Flexible storage choices and options.
- ▶ Virtual tape emulation of up to 16 virtual tape libraries, up to 256 virtual tape drives on each node, and up to 512,000 virtual cartridges.
- ▶ Emulation of the IBM TS3500 Tape Library with IBM LTO Ultrium 3 tape drives.
- ▶ Emulation of the Quantum P3000 tape library with DLT7000 tape drives and IBM LTO Ultrium 2 tape drives.
- ▶ Emulation of the DTC VTF 0100 tape library with DLT7000 tape drives and IBM LTO Ultrium 2 tape drives.
- ▶ Emulation of the IBM V-TS3500 Tape Library with IBM LTO Ultrium 3 tape drives.
- ▶ IBM Path Failover technology.
- ▶ LUN masking.
- ▶ Optional replication support to allow virtual cartridges to be copied to a target (remote) ProtecTIER two-node cluster based on defined policies.

- ▶ The OpenStorage (OST) feature implements the functionality to directly connect the ProtecTIER server to hosts running Symantec NetBackup.

A wide variety of disk-based storage can be attached to this Gateway server. Check the supported devices at the following address:

[ftp://service.boulder.ibm.com/storage/tape/TS7650\\_support\\_matrix.pdf](ftp://service.boulder.ibm.com/storage/tape/TS7650_support_matrix.pdf)

The IBM System Storage ProtecTIER Enterprise Edition V2.5 software and Red Hat Enterprise Linux 5.4 Advanced Platform 64-bit software are installed by the IBM SSR on the 3958-DD4 server.

The 3958-DD4 server is available in a single node configuration and in a two-node cluster configuration. The single node configuration consists of one IBM System Storage TS7650G ProtecTIER Deduplication Gateway, which is based on an IBM System x3850 M2 Type 7145-PBR server. In this book, we refer to this server as IBM machine type and model 3958-DD4 for the IBM System Storage TS7650G ProtecTIER Deduplication Gateway and model 3958-AP1 for the IBM System Storage TS7650 ProtecTIER Deduplication Appliance.

The two-node cluster configuration includes the following hardware components:

- ▶ Two IBM 3958-DD4 Gateway servers
- ▶ Two network Ethernet switches
- ▶ One remote Network Power Switch

To allow installation, service, and maintenance of the 3958-DD4 server, a console must be provided by the customer that can be directly attached to the 3958-DD4 server.

The 3958-DD4 server requires the purchase of the following additional hardware components (one or more frames, one or more disk arrays, and expansions) to be fully functional:

- ▶ The disk array is the term used in this document to refer to a disk storage subsystem.
- ▶ The expansion refers to a disk expansion attached to the disk array.
- ▶ A frame is a 19-inch rack supplied by the customer and used to house the Gateway servers and the TS3000 System Console. A second frame can be used to house disk arrays and expansions.

**Note:** Disk arrays, expansions, and frames are required but are not included with the 3958-DD4 shipment.

The supported disk arrays have the following characteristics:

- ▶ Support for the 3958-DD4 server operating system with the correct update level.
- ▶ Dual active-active controller for compatibility with the Linux Multipath software included in the 3958-DD4 server operating system to allow path fail over.
- ▶ Fibre Channel, SAS, or SATA disk systems.
- ▶ Support for the Back End Fibre Channel Host Bus Adapter (HBA) brand, model, and firmware level installed on the Gateway server. The back-end HBAs are used to direct or SAN attach the 3958-DD4 server to the disk array. In case of SAN attachment, the disk array must also support the fabric switches used.
- ▶ The array should not perform its own compression by default. ProtecTIER does not require additional compression to be effective. ProtecTIER performs compression, by default, after the deduplication process.

We describe in detail the hardware components and the corresponding feature codes in the following sections.

#### 4.2.6 3958-DD4 server characteristics

The 3958-DD4 server is based on a 4U rack-drawer IBM System x3850 X5 Type 7145-PBR. It is an improvement to the 3958-DD3 server in performance, for example, due to memory and processor upgrades and a faster bus speed.

The 3958-DD4 server is shipped with the following characteristics:

- ▶ Four Intel Xeon X7650 2.27 GHz 24M 8 core 6.4 Gb 130 W Turbo processors (32 cores in total, 64 threads with hyper-threading)
- ▶ Eight memory cards, with 2 x 4 GB DDR3 DIMMs per memory card, total amount of memory 64 GB
- ▶ Two internal 300 GB 2.5-inch 10 K RPM SAS Drives (RAID 1 mirror)
- ▶ Internal ServeRAID M5015 SAS/SATA controller with battery
- ▶ Two dual-port Emulex LPe12002 8 Gbps Expansion Card (front-end ports) for VTL configuration
- ▶ Two dual-port Qlogic QLE2562 8 Gbps Expansion Card (back-end ports)
- ▶ In the VTL configuration, one Intel Pro/1000 PT Quad Port Ethernet Server Adapter
- ▶ In the OST configuration, two Intel Pro/1000 PT Quad Port Ethernet Server Adapter or two QLE2562 dual port 10 Gb CNAs
- ▶ Two 1975 W standard hot swap Redundant Power supply, Hot Swap Redundant fans, Hot Swap HDDs, and Hot Swap memory
- ▶ SATA CD/DVD-ROM
- ▶ Integrated Management Module (IMM) as successor of Remote Supervising Adapter II (RSAIL)

Figure 4-2 shows the front view of the IBM System x3850X5 rack mounted server



Figure 4-2 Front view of a IBM System x3850 X5 (AP1 and DD4 server)

Figure 4-3 shows the rear view of a AP1 and DD4 server with VTL port assignment.

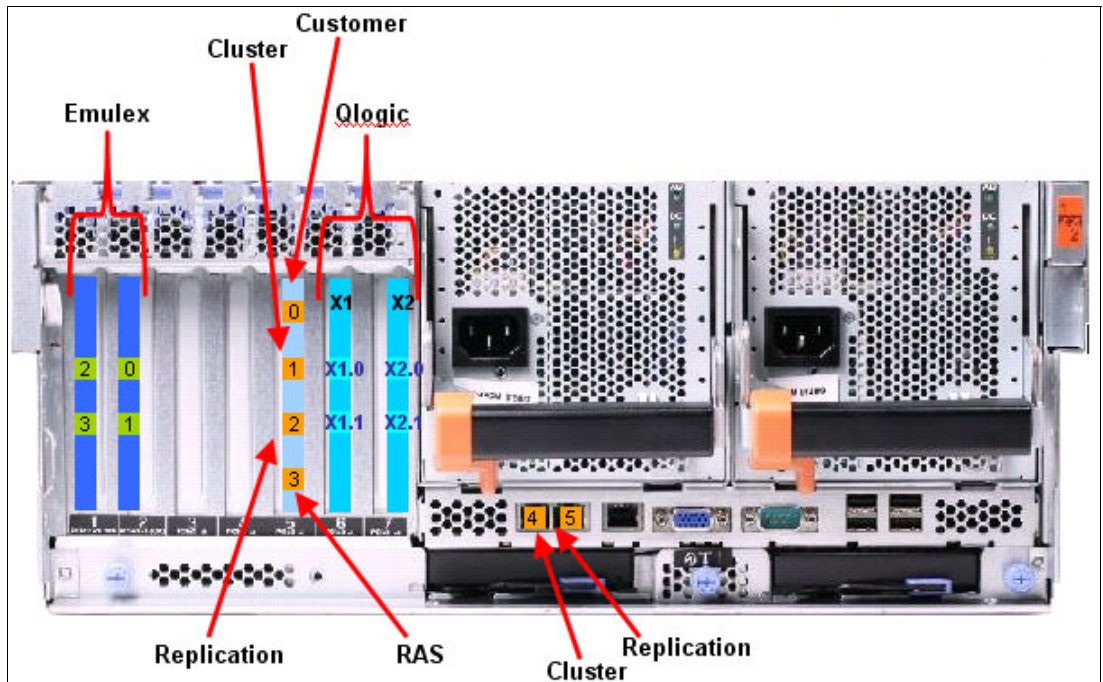


Figure 4-3 Rear view of AP1 and DD4 server with VTL port assignment

Figure 4-4 shows the rear view of a AP1 and DD4 server with OST port assignment.




Figure 4-4 Rear view of a AP1 and DD4 server with OST port assignment



## 4.2.7 TS3000 System Console

A TS3000 should be included in the order with the TS7610, TS7650, or TS7650G. One TS3000 console can be used to monitor and manage several ProtecTIER servers. If a TS3000 console was not included in the order, we recommend that one be ordered. Existing TS3000 models 306 or lower are not supported ProtecTIER servers.

If you need a TS3000 System Console, it can be ordered with all TS7600 family products, and then they are shipped together. The TS3000 System Console is a one-unit (1U) System x server that allows an IBM System Service Representative (SSR) to perform maintenance and, if enabled by the customer, the TSCC can remotely monitor the installation and automatically call home with any hardware errors.

**Note:** If you want open a service call at IBM for your ProtecTIER system, always use the model name (3959-SM1, 3958-AP1, or 3958-DD4) and the serial number shown on the front of the server.

## 4.3 3959-SM1, 3958-AP1, and 3958-DD4 feature codes

This section describes the 3959-SM1, 3958-AP1, and 3958-DD4 components and lists the required and optional feature codes for each component.

### 4.3.1 Features codes for 3959-SM1

This section lists the feature codes to use when you order the required and optional features for the IBM 3959 Model SM1:

► FC 3448 - Dual Port Ethernet Card

This feature provides a dual port 1 Gbps 10/100/1000Base-TX PCIe Ethernet adapter. This adapter has an RJ-45 connector for attaching CAT6 cables. This adapter conforms to the IEEE 802.3ab 1000Base-T standard. It supports distances of up to 100 meters using four pairs of CAT6 balanced copper cabling. The feature is chargeable and plant installed only.

Prerequisite feature: FC 9023.

► FC 3458 - 8 Gb FC Dual Port HBA - Host

This 8 Gb FC dual-port PCIe HBA provides the connections to the host servers or switches. The feature is chargeable and plant installed only.

Prerequisite feature: FC 9022.

► FC 9022 - Virtual Tape Library (VTL)

This feature designates that this ProtecTIER system will be configured for use as a Virtual Tape Library management system. The feature is chargeable and plant installed only.

Prerequisite features: FC 3458 or FC 3459 is required. Mutually exclusive with FC 9023 and FC 3448.

► FC 9023 - OpenStorage (OST)

This feature designates that this ProtecTIER system will be configured for use with OpenStorage. The feature is chargeable and plant installed only.

Prerequisite features: FC 3448 is required. Mutually exclusive with FC 9022 and FC 3459.

- ▶ FC 9313 - SW Capacity 4.4 TB

This feature instructs IBM manufacturing to load ProtecTIER Entry Edition software with a usable repository of 4.4 TB (4.0 TiB) on the TS7610 ProtecTIER Deduplication Appliance Express. A separate order for the software is required. The feature is chargeable and plant installed only.

Prerequisite features: One of FC 9313 or FC 9314 is required.

- ▶ FC 9314 - SW Capacity 5.9 TB

This feature indicates that ProtecTIER Entry Edition software with a usable repository of 5.9 TB (5.4 TiB) is running on this TS7610 ProtecTIER Deduplication Appliance Express. If ordered for Plant Installation, IBM manufacturing will load ProtecTIER Entry Edition software with a usable repository of 5.9 TB (5.4 TiB) on the TS7610 ProtecTIER Deduplication Appliance Express. An order for Field installation indicates that this ProtecTIER Deduplication Appliance Express will be upgraded from ProtecTIER Entry Edition software with a usable repository of 4.4 TB (4.0 TiB). A separate order for the software is required in all cases. The feature is chargeable and plant or field installed.

Prerequisite features: One of FC 9313 or FC 9314 is required.

### 4.3.2 Features codes for 3958-AP1 server

This section lists the feature codes to use when you order the required and optional features for the IBM 3958 Model AP1 server:

- ▶ FC 3170 - TS7650 Appliance Server

This is the feature code for the ProtecTIER appliance server itself.

- ▶ FC 2714 - Console Expansion

This feature provides an attachment cable for the connection of a unit using the TS3000 System Console (TSSC), and an Ethernet hub for expanding the number of units that can be attached to the TSSC. Up to 14 additional connections are provided by this feature for connection of FC 2715 or another FC 2714. The feature is chargeable and can be plant or field installed.

- ▶ FC 2715 - Console Attachment

This feature provides a cable to attach to the Ethernet hub provided by an existing TS3000 System Console, IBM TotalStorage Master Console for Service (feature FC 2718), or Console Expansion (FC 2714). A maximum of 40 of feature FC 2715 may be included in a single console facility. The feature is chargeable and can be plant or field installed.

- ▶ FC 2722 - TS3000 System Console

This feature provides the enhanced rack-mountable TS3000 Service Console, an Ethernet switch, and a cable and connectors for connection of one machine to an IBM-supplied modem to enable remote enhanced service. This feature is an enhanced replacement of the IBM TS3000 System Console for Service (FC 2721). It includes a console upgrade previously provided as FC 2719 (memory upgrade to 2 GB total RAM) and a second Ethernet card for the Service Console to allow redundant connections into the service network. The feature is chargeable and can be plant or field installed.

Prerequisite features: One of FC 2714, FC 2715 or FC 2722 is required.

Corequisite features: FC 2733, FC 5510, and FC 5512.

- ▶ FC 2733 - TS3000 Internal Modem

This feature provides an internal modem for installation. The feature is chargeable and plant installed only.

Corequisite features: FC 2732 is required in countries where the modem provided in FC 2733 is homologated.

▶ FC 3170 - 16 x 2.27 GHz Processor Cores

This feature provides a System x3850 X5 processor.

Prerequisite features: If two of FC 3120 or FC 3170 are installed, two FC 3706 and six FC 3707 are required.

FC 3437 is required when two FC 3170 or one FC 3120 and one FC 3170 are installed.

▶ FC3437 - Cluster Connection Kit

This feature provides a network power control switch, 2 Ethernet switches, associated cables, and mounting hardware to create a clustered configuration between two TS7650 Servers. The network power switch is used to automatically control the power to a failed node. The feature is chargeable and can be plant or field installed.

The Cluster Connection Kit can cluster two 3958-AP1 servers.

▶ FC 3449 - 4 GB FC Dual Port HBA - Disk

This Qlogic 4 GB FC dual-port PCIe HBA provides the connections to the application disk storage subsystem. The feature code is chargeable and can be plant or field installed.

▶ FC 3456 - 1 Gb Quad Port Ethernet Adapter

This feature provides a quad port 1 Gbps 10/100/1000Base-TX PCIe Ethernet adapter. This adapter has an RJ-45 connector for attaching CAT6 cables. This adapter conforms to the IEEE 802.3ab 1000Base-T standard. It supports distances of up to 100 meters using four pairs of CAT6 balanced copper cabling. The feature is chargeable and can be ordered as plant or field installed.

Prerequisite feature: At least one FC 9023 is required. When FC 9023 is installed, two FC 3456 are required for each FC 3170.

▶ FC 3457 - 10 Gb Dual Port Ethernet Adapter

This feature provides a dual port 10 Gbps Ethernet longwave adapter. This adapter has an SC Duplex connector for attaching 9 micron single mode fibre cables. It is a standard longwave (1,310 nm) adapter conforming to the IEEE 802.3ae standards. It supports distances up to 10 km. The feature is chargeable and can be ordered as plant or field installed.

Prerequisite feature: At least one FC 9023 is required. When FC 9023 is installed, two FC 3456 or FC 3457 are required for each FC 3170.

▶ FC 3458 - 8 Gb FC Dual Port HBA

This 8 Gb FC dual-port PCIe HBA provides the connections to the host servers or switches. The feature is chargeable and can be plant or field installed.

Prerequisite feature: At least one FC 9022 is required. When FC 9022 is installed, two FC 3458 are required for each FC 3170.

▶ FC 3707 - 4.8 TB FC Disk Expansion

A DS4000 expansion drawer with 16 x 450 GB HDDs. The feature is chargeable and can be plant or field installed.

Prerequisite features: If two FC 3706 are installed, you must have six FC 3707. If two FC 3120 are installed, you must have two FC 3706.

▶ FC 3708 - 4.8 TB FC Disk Controller

Disk control unit with 16 x 450 GB HDDs. The feature is chargeable and can be plant or field installed.

Prerequisite features: At least one FC 3706 or FC 3708 is required. If two FC 3706 plus FC 3708 are installed, you must have six FC 3707. If two FC 3120 plus FC 3170 are installed, you must have two FC 3706 plus FC 3708.

▶ FC 5510 - Plant Install KVM Switch

This feature provides the KVM switch, cables, and mounting hardware to be plant installed in the frame or delivered for installation in a rack. The feature is chargeable and can be plant or field installed.

Corequisit features: FC 2722 and FC 5512.

▶ FC 5512 - KVM Display, Keyboard, Mouse

This feature provides the KVM display, keyboard, mouse and mounting hardware to be plant installed in the frame or delivered for installation in a rack. The feature is chargeable and can be plant or field installed.

Prerequisite feature: FC 2722.

Corequisit feature: FC 5510.

▶ FC 5513 - KVM Adjacent Frame Connection

This feature provides a cable to connect one or two appliances or gateways to the KVM switch in an adjacent unit with feature FC 5510 and FC 5512 installed. This feature allows the keyboard, mouse, and video display installed with FC 5512 to be used by the ProtectTIER system. A maximum of four FC 5513 may be attached to a single FC 5510. The feature is chargeable and can be plant or field installed.

Prerequisite features: FC 2714 or FC 2715 is required. FC 5510 and FC 5512 must exist in an adjacent frame or rack.

▶ FC 5514 - Second Server in Adjacent Frame

This feature designates the second appliance or gateway that will use the FC 5513 cable to share the keyboard, mouse, and video display from an adjacent frame with FC 5512 installed. The feature is chargeable and can be plant or field installed.

Prerequisite features: FC 5514 is required and only allowed when FC 5513 is installed and FC 3120 plus FC 3170 equals two.

▶ FC 6025 - 25 Meter LC/LC Fibre Channel Cable

This feature provides a 25 meter (82 ft.) 50.0/125 micrometer short wavelength multimode fiber-optic cable with LC duplex connectors on both ends. The feature is chargeable and can be plant or field installed.

**Note:** This cable is for attaching a TS7600 server (that has an LC duplex connector) to switches, disk subsystems, or hosts with LC duplex Fibre Channel connectors.

Prerequisite features: One of FC 9700 or at least four of FC 6025 is required.

▶ FC 9022 - ProtectTIER Virtual Tape Library (VTL)

This feature designates that this ProtectTIER system will be configured for use as a virtual tape library management system. The feature is chargeable and can be plant or field installed.

Prerequisite features: Two FC 3459 are required for each FC 3120 installed. Two FC 3458 are required for each FC 3170 installed. Mutually exclusive with FC 3456 and FC 9023.

- ▶ FC 9023 - OpenStorage (OST)

This feature designates that this ProtecTIER system will be configured for use with OpenStorage. The feature is chargeable and can be plant installed only.

Prerequisite features: Two FC 3456 or FC 3457 are required for each FC 3170 installed. Mutually exclusive with FC 3120, FC 3458, and FC 9022.

- ▶ FC 9308 - ProtecTIER Appliance Software Version 2.5

This feature specifies that ProtecTIER Appliance Edition V2.5 software will be run on the server. A separate order for the software is required. The feature is chargeable and can be plant installed only.

- ▶ FC 9700 - No Factory Cables

This feature should be specified if you do not want the factory to ship any Fibre Channel cable features with the new machine. These Fibre Channel cables are for attachment to switches or host servers. The feature is available at no cost and can be ordered plant installed only.

Prerequisite features: One of FC 9700 or at least four of FC 6025 is required.

### 4.3.3 Feature codes for 3958-DD4 server

This section lists the feature codes to use when you order the required and optional features for the IBM 3958-DD4 Gateway Server:

- ▶ FC 2714 - Console Expansion

This feature provides an attachment cable for connection of a unit using the TS3000 System Console (TSSC), and an Ethernet hub for expanding the number of units that can be attached to the TSSC. Up to 14 additional connections are provided by this feature for connection of FC 2715 or another FC 2714. The feature is chargeable and can be plant or field installed.

- ▶ FC 2715 - Console Attachment

This feature provides a cable to attach to the Ethernet hub provided by an existing TS3000 System Console, IBM TotalStorage Master Console for Service (feature FC 2718), or Console Expansion (FC 2714). A maximum of 40 of feature FC 2715 may be included in a single console facility. The feature is chargeable and can be plant or field installed.

- ▶ FC 2722 - TS3000 System Console

This feature provides the enhanced rack-mountable TS3000 Service Console, an Ethernet switch, and a cable and connectors for connection of one machine to an IBM-supplied modem to enable remote enhanced service. This feature is an enhanced replacement of the IBM TS3000 System Console for Service (FC 2721). Includes a console upgrade previously provided as FC 2719 (memory upgrade to 2 GB total RAM) and a second Ethernet card for the Service Console to allow redundant connections into the service network. The feature is chargeable and can be plant or field installed.

Prerequisite features: One of FC 2714, FC 2715, or FC 2722 is required.

Corequisit features: FC 2733, FC 5510, and FC 5512.

- ▶ FC 2733 - TS3000 Internal Modem

This feature provides an internal modem for installation. The feature is chargeable and plant installed only.

Corequisite features: FC2732 is required in countries where the modem provided in FC 2733 is homologated.

► FC 3437 - Cluster Connection Kit

This feature provides a network power control switch, two Ethernet switches, and the associated cables and mounting hardware to create a clustered configuration between two TS7650 servers. The network power switch is used to automatically control the power of a failing node. This feature is chargeable and only available as plant installed. Only one Cluster Connection Kit is required for each pair of clustered nodes. Both cluster servers must be installed in the same rack. The feature provides the required power cords and Ethernet cables to attach the two-node servers with the network power switch and the Ethernet switches.

► FC 3456 - 1 Gb Quad Port Ethernet Adapter

This feature provides a quad port 10/100/1000Base-TX PCIe Ethernet adapter. This adapter has 4 port RJ-45 connectors for attaching CAT6 cables. This adapter conforms to IEEE 802.3ab 1000Base-T standard. It supports distances of up to 100 meters using four pairs of CAT6 balanced copper cables. It is a chargeable feature with a maximum of two installed in one node. Plant installed only.

Prerequisite features: Two FC3456 are needed.

► FC 3457 - 10 Gb Dual Port Ethernet Adapter

This feature provides a dual port 10 Gbps Ethernet longwave adapter. This adapter has an SC Duplex connector for attaching 9 micron single mode fibre cables. It is a standard longwave (1,310 nm) adapter conforming to the IEEE 802.3ae standards. It supports distances up to 10 km. The feature is chargeable and can be ordered as plant or field installed.

Prerequisite feature: At least one FC 9023 is required. When FC 9023 is installed, two FC 3456 or FC 3457 are required for each FC 3170.

► FC 3458 - 8 Gb FC Dual Port HBA

This feature provides a 8 Gb FC dual port PCIe HBA for connectivity to host servers or switches. This is a chargeable feature and plant installed only.

Prerequisite features: Two of FC3458 are needed.

► FC 5510 - KVM Switch

This feature provides the KVM switch, cables, and mounting hardware to be plant installed in the frame or delivered for field installation in a rack. This feature is chargeable.

Corequisite features: FC2722 and FC5512.

► FC 5512 - KVM Display, Keyboard, and Mouse

This feature provides the KVM display, keyboard, and mouse, including mounting hardware to be plant installed in a frame or delivered for field installation in a rack. The feature is chargeable.

Prerequisite feature: FC2722

Corequisite feature: FC5510

► FC 5513 - KVM Adjacent Frame Connection

This feature provides a cable to connect one or two appliances or gateways to the KVM switch in an adjacent unit with feature FC5510 and FC5512 installed. A maximum of four FC5513 may be attached to a single FC5510. This feature is chargeable. Could be plant or field installed.

Prerequisite features: FC2714 or FC2715, FC5510, and FC5512 must exist in an adjacent frame or rack.

► FC 5514 - Second Server in an Adjacent Frame

This feature designates the second appliance or gateway that will use the FC5513 cable to share the keyboard, mouse, and video display from an adjacent frame with the FC5512 installed. This feature is chargeable. Could be plant or field installed.

Prerequisite features: FC2714 or FC2715 and FC5513 on another ProtecTIER server in the same frame or rack.

► FC6025 - 25 meter LC/LC Fibre Channel Cable

This feature provides a 25 meter (82 feet) 50.0/125 micrometer short wavelength multimode fibre optic cable with LC duplex connectors on both ends. This cable is for attaching a 3958-DD4 server (that has a LC duplex connector) to switches, disk subsystems, or hosts with a LC duplex Fibre Channel connector. The feature is chargeable. A maximum number of four can be installed; in an environment with cluster support, a maximum of eight could be installed. Plant or field installation is possible.

Prerequisite: FC9700, and at least four FC6025 are required.

► FC9022 - ProtecTIER Virtual Tape Library (VTL)

This feature indicates that the 3958-DD4 server will be installed as a virtual tape library management system. The feature is chargeable. Plant installed only.

Prerequisite features: Two FC3458. Mutually exclusive with FC3456, FC3457, and FC9023.

► FC9023 - Open Storage (OST)

This feature indicates that the 3958-DD4 server will be configured for use with open storage only. The feature is chargeable and plant installed only.

Prerequisite features: Two FC3458. Mutually exclusive with FC3456, FC3457, and FC9022.

► FC9306 - ProtecTIER Preload

This feature instructs IBM manufacturing to load the ProtecTIER Enterprise Edition software on the TS7650G server. A separate order for the software is required. The feature is available at no cost and plant installed only.

► FC9340 - Rack Mount

This feature code indicates that the TS7650G server will be mounted in a customer provided industry standard 19-inch rack. The feature is available at no cost and plant installed only.

► FC9700 - No Factory Cables

This feature should be specified if you do not want the factory to ship any fibre cables with the new machine. This includes the fibre cables for the attachment of switches and host servers. The feature is available at no cost and plant installed only.

Prerequisite features: One of FC9700 or at least four of FC6025 are required.

## 4.4 IBM System Storage TS7600 with ProtecTIER software

The TS7610, TS7650, and TS7650G use the same ProtecTIER software. This section provides an overview of that software.

### 4.4.1 5639-XXB ProtecTIER Enterprise Edition (EE) V2.5 Base Software

The 5639-XXB ProtecTIER Enterprise Edition V2.5 Base Software is the software that provides the base functionality for the ProtecTIER Enterprise Edition and the TS7650G models. The software supports any capacity of FC or SATA disk drives in its disk storage pool. It is ordered on a *per server* basis.

ProtecTIER EE V2.5 supports:

- ▶ Inline data deduplication powered by HyperFactor technology.
- ▶ Up to 16 virtual tape libraries per cluster.
- ▶ Up to 256 virtual tape drives per cluster.
- ▶ Up to 512,000 virtual cartridges per cluster with a maximum capacity of XYZ.
- ▶ Emulation of IBM Ultrium 2 or Ultrium 3 tape drives.
- ▶ Emulation of Quantum DLT 7000 drives.
- ▶ Emulation of IBM TS3500 and IBM V-TS3500 tape libraries.
- ▶ Emulation of Quantum P3000 tape library.
- ▶ Emulation of DTC VTF 0100 virtual tape libraries.
- ▶ IBM path fail over technology.
- ▶ ProtecTIER appliance Edition V2.5 is available with 7 TB, 18 TB, and 36 TB in a single node configuration and also with 36 TB in a dual-node cluster configuration.
- ▶ Optional clustering support for higher availability.
- ▶ Optional IP replication support to allow virtual cartridges to be copied to a target ProtecTIER cluster based on defined policies.

#### Capacity support

This is the base functionality for ProtecTIER Enterprise Edition (ProtecTIER EE). It supports any capacity of either FC or SATA disk drives in its storage pool. It is ordered on a per terabyte basis for usable capacity of the whole configuration. The tier levels are shown in Table 4-2.

Table 4-2 Tier capacities

Tier capacities in terabytes
1 - 12 TB
13 - 32 TB
33 - 64 TB
65 - 100 TB
101 - 250 TB
250+ TB



## Clustering

This feature allows the server running this software to share a common disk storage pool with a second server running ProtecTIER EE. The second server also must have the cluster feature. This feature is ordered on a per terabyte basis. The tier levels are the same as for capacity support and shown in Table 4-2 on page 138.

## Replication

This feature allows the server running the ProtecTIER V2.5 EE to replicate data to a second system running the ProtecTIER software with the replication feature. This feature is also ordered on a *per terabytes* basis with the same tier capacities shown in Table 4-2 on page 138.

**Note:** Contact your responsible IBM sales representative to calculate the licensing for your already installed or future disk storage.

### 4.4.2 5639-XXP ProtecTIER Appliance Edition (AE) V2.5 Software

The IBM ProtecTIER 5639-XXP ProtecTIER Appliance Edition V2.5 software provides the basic functionality for the TS7650 Appliance (3958-AP1) model. The IBM ProtecTIER AP offers 7 TB, 18 TB, and 36 TB of disk space pre-installed in the TS7650-AP1 frame in a single node configuration. The ProtecTIER AP in dual node cluster configuration is available with 36 TB.

ProtecTIER AE V2.5 supports:

- ▶ Inline data deduplication powered by HyperFactor technology.
- ▶ Up to 12 virtual tape libraries per cluster.
- ▶ Up to 256 virtual tape drives per cluster.
- ▶ Up to 128,000 virtual cartridges per cluster.
- ▶ Emulation of IBM Ultrium 2 or Ultrium 3 tape drives.
- ▶ Emulation of Quantum DLT 7000 drives.
- ▶ Emulation of IBM TS3500 and IBM V-TS3500 tape libraries.
- ▶ Emulation of Quantum P3000 tape library.
- ▶ Emulation of DTC VTF 0100 virtual tape libraries.
- ▶ IBM path failover technology.
- ▶ ProtecTIER appliance Edition V2.5 is available with 7 TB, 18 TB, or 36 TB in single node configuration and also with 36 TB in a dual-node cluster configuration.
- ▶ Optional clustering support for higher availability.
- ▶ Optional IP replication support to allow virtual cartridges to be copied to a target ProtecTIER cluster based on defined policies.

### 4.4.3 ProtecTIER Manager V2.5 console software

The 3958-DD4, the 3958-SMB, and 3958-AP1 servers must have a console workstation for installation and maintenance. This workstation runs the ProtecTIER Manager application. This application is a graphical user interface (GUI) used to install, configure, and monitor ProtecTIER. Install the ProtecTIER Manager on a workstation running either Windows or Linux. The ProtecTIER Manager workstation must be connected to the 3958-DD3/DD4, 3958-SMB and 3958-AP1 servers through your network (see Chapter 5, “IBM System Storage TS7600 with ProtecTIER initial setup” on page 171).

- ▶ The console must be capable of operating one of the following operating systems:
  - Windows 2000
  - Windows 2003
  - Windows XP
  - Windows 7
  - Linux Red Hat 32/64-bit (Red Hat Enterprise 4 or 5)
- ▶ Hardware requirements:
  - x86 (Pentium or higher) microprocessor
  - At least 1.2 GB MB of available disk space
  - At least 256 MB of RAM
  - Keyboard, mouse, and CD-ROM drive

In addition, consider configuring the monitor for ProtecTIER Manager software with the following settings:

- ▶ Resolution of 1024 x 768 pixels or higher
- ▶ 24-bit color or higher

**Note:** The console must have access to the IP address of the ProtecTIER service nodes (port 3501 has to be open on the firewall).

If you are planning to run ProtecTIER Manager on a Linux system, configure your graphics card and X Window System. This is done either manually or by using the Xconfigurator utility. For instructions, refer to the appropriate Linux documentation.

## 4.5 Feature codes for Red Hat Linux

This section provides a description of the Red Hat Enterprise Linux Edition and its features installed on the ProtecTIER server.

### 4.5.1 ProtecTIER Enterprise Edition V2.5

The ProtecTIER Enterprise Edition V2.5 is installed on a Red Hat Enterprise Linux 5.4 Advanced Platform (RHELAP) x86\_64, either on a 3958-DD4, 3958-SM1, or 3958-AP1 server with your initial order of the server.

The customer has to purchase the Red Hat Enterprise Linux 5.4 Advanced Platform. One Red Hat license is required for each ProtecTIER gateway server. The appropriate number of licenses must be purchased from IBM along with a one-year or three-year Red Hat support subscription that matches the maintenance period purchased for the TS7650G. Use PID 5639-RHL with feature 0017 (1-year support) or feature 0019 (3-year support) to purchase the appropriate number of Red Hat licenses and support.

Red Hat 5.4 Advanced Platform is shipped on a recovery disk that can be used for restoring the operating system on the ProtecTIER gateway server in case of a hardware failure.

## 4.6 3958-DD4 server configuration options

The 3958-DD4 server consists of mandatory and optional components. We cover only the 3958-DD4 Gateway model connection in this section. The 3958-AP1 and 3959-SM1 servers come from the plant in a predefined configuration and no connections changes are needed. These components can be configured to meet customer requirements in terms of scalability and performance. Our intent is to describe the recommended configurations with all the components that must be installed and deployed to provide a fully functional system. As mentioned before, a tape frame, disk array, and expansions can be supplied by other vendors, but it is essential that the disk array is supported and that the rack satisfies specific customer requirements in terms of adequate space, and adequate air flow, layout, and so on.

**Note:** In this book, we describe the fully functional system using an IBM System Storage DS5020 as a disk array, but you can use any supported disk array. Disk arrays supported by the 3958-DD4 server are listed in the document found at the following address:

[ftp://service.boulder.ibm.com/storage/tape/TS7650\\_support\\_matrix.pdf](ftp://service.boulder.ibm.com/storage/tape/TS7650_support_matrix.pdf)

In the following sections, we describe the two possible configuration options for the 3958-DD4 server (they are also valid for the 3958-DD3 server):

- ▶ The single node configuration
- ▶ The two-node cluster configuration

For easier installation and maintenance in both cases, the components included in the purchase of the 3958-DD4 server and the TS3000 System Console should occupy one frame (the server frame), while the disk array components occupy a second frame (the cache frame).

### 4.6.1 Single node configuration

The TS7650G single node configuration can be built, depending on the customer needs, in terms of capacity and performance, as noted in Chapter 3, “Planning for deduplication and replication” on page 53.

The recommended rack layout in the case of a single node configuration is shown in Figure 4-5. As you can see, two frames are used: one for the 3958-DD4 server and the other for the storage server. The server frame is used to host the 3958-DD4 server, the TS3000 System Console, the TS3000 System Console switch, and two Ethernet switches. A cache frame might be used to host specific supported disk arrays and expansions. For example, the DS5020 might be mounted in a customer-supplied frame, while an IBM System Storage DS8300 is shipped with its own dedicated enclosure rack. Therefore, it does not need a customer-supplied cache frame.

106	36	Empty (1u)	36	106	106	36	Empty (1u)	36	106
103	35	Empty (1u)	35	103	103	35	Empty (1u)	35	103
100	34	Empty (1u)	34	100	100	34	Empty (1u)	34	100
97	33	Empty (1u)	33	97	97	33	Empty (1u)	33	97
94	32	Empty (1u)	32	94	94	32	Empty (1u)	32	94
91	31	Empty (1u)	31	91	91	31	Empty (1u)	31	91
88	30	Empty (1u)	30	88	88	30	Empty (1u)	30	88
85	29	Empty (1u)	29	85	85	29	Empty (1u)	29	85
82	28	Empty (1u)	28	82	82	28	Empty (1u)	28	82
79	27	Empty (1u)	27	79	79	27	Empty (1u)	27	79
76	26	Empty (1u)	26	76	76	26	Empty (1u)	26	76
73	25	Empty (1u)	25	73	73	25	Empty (1u)	25	73
70	24	Empty (1u)	24	70	70	24	Empty (1u)	24	70
67	23	Empty (1u)	23	67	67	23	Empty (1u)	23	67
64	22	Empty (1u)	22	64	64	22	Empty (1u)	22	64
61	21	Empty (1u)	21	61	61	21	Empty (1u)	21	61
58	20	Empty (1u)	20	58	58	20	Empty (1u)	20	58
55	19	Empty (1u)	19	55	55	19	Empty (1u)	19	55
52	18	TSSC (1u)	18	52	52	18	Empty (1u)	18	52
49	17	KVM Tray + TSSC sw (1u)	17	49	49	17	Empty (1u)	17	49
46	16	Empty (1u)	16	46	46	16	Empty (1u)	16	46
43	15	Empty (1u)	15	43	43	15	Empty (1u)	15	43
40	14	Empty (1u)	14	40	40	14	Empty (1u)	14	40
37	13	Empty (1u)	13	37	37	13	Empty (1u)	13	37
34	12	Empty (1u)	12	34	34	12	Empty (1u)	12	34
31	11	Empty (1u)	11	31	31	11	Disk Expansion (3u)	11	31
28	10	Empty (1u)	10	28	28	10	Empty (1u)	10	28
25	9	Empty (1u)	9	25	25	9	Empty (1u)	9	25
22	8	Empty (1u)	8	22	22	8	Disk Expansion (3u)	8	22
19	7	ProtectTIER Server (4u)	7	19	19	7	Empty (1u)	7	19
16	6		6	16	16	6	Empty (1u)	6	16
13	5		5	13	13	5	Disk Expansion (3u)	5	13
10	4	Empty (1u)	4	10	10	4	Empty (1u)	4	10
7	3	Power Distribution Unit (PDU)	3	7	7	3	Empty (1u)	3	7
4	2	Power Distribution Unit (PDU)	2	4	4	2	Disk Controller (3u)	2	4
1	1	Empty (1u)	1	1	1	1	Empty (1u)	1	1

Figure 4-5 TS7650G frame layout for single node and adjacent storage frame

The server frame includes all the hardware components necessary to build the single node configuration, but also allows further expansion of the single node configuration to the two-node cluster configuration. Figure 4-5 shows the recommended position of each hardware component in the server frame and storage frame. Some of the rack units are left empty for a potential upgrade to the two-node cluster configuration and storage capacity growth.

The single node configuration example is composed of:

- ▶ Two 36 U frames
- ▶ One 3958-DD4 server
- ▶ One TS3000 System Console and an accompanying Keyboard/Video/Monitor (KVM) kit
- ▶ Disk array, in our example, a dual controller IBM System Storage DS5020
- ▶ Three expansion units IBM TotalStorage DS5020 EXP520 Expansion Enclosures frame).

This solution can be used in environments where the configuration results from an adequate capacity and performance planning assessment are met by a single node configuration. In this section, we discuss the hardware planning in depth, pointing out the consequent sizing considerations, guidelines, and best practices.

The 3958-DD4 server contains Four Intel Xeon X7650 2.27 Ghz eight-core processors and 64 GB of RAM. Each processor core can manage two simultaneous compressed jobs, allowing 32 concurrent backup application jobs. As a guideline, HyperFactor has a Memory Resident Index, like a table of contents, that can map the contents of a 1 PB repository in only 4 GB of RAM, eliminating the need to store indexes on disk or on large amounts of RAM.

### DS5020 disk array controller and EXP520 expansion configurations

In this section, we describe the features of the IBM System Storage DS5020 in a single node configuration as an example. You can also install all the other listed cache storage systems.

**Note:** To see the entire list of supported disk controllers and disk expansion that can be connected to TS7650G, access the TS7650/TS7650G/TS7610 ProtecTIER Deduplication interoperability matrix at the following address:

[ftp://service.boulder.ibm.com/storage/tape/TS7650\\_support\\_matrix.pdf](ftp://service.boulder.ibm.com/storage/tape/TS7650_support_matrix.pdf)

This disk array has controller and power redundancy to avoid a single point of failure (SPOF). Also, this disk array can host up to 16 disk drives on FC or SATA technology and with different sizes. The disk drive characteristics in terms of technology, size, and speed must be chosen after a capacity and performance planning assessment. To see the description of the assessment phases, refer to Chapter 3, “Planning for deduplication and replication” on page 53.

When attaching expansions, drive loops are configured as redundant pairs using one port from each controller, which helps ensure data access in the event of a path/loop or controller failure. This controller can be expanded up to six expansions and in this full configuration the disk array is cabled using all two-drive channel pairs, assuming that there are six total expansion enclosures evenly spread out across the drive channel pairs (three each).

Figure 4-6 shows the rear view of the controller disk array that shows the P1 and P2 ports for FC connections to expansion drawers and the H1 and H2 ports for FC connections to the 3958-DD4 server.




Figure 4-6 Rear view of the DS5020 controller

Figure 4-7 shows the rear view of the expansion drawer that shows the ports available for connections. Only the Ports 1A and 1B are used to connect to the cache controller and the expansion units. Port 2A and 2B are not used.




Figure 4-7 Rear view of the DS5000 Expansion unit

Figure 4-8 shows the FC connectivity layout of the connections between a 3958-DD4 server and a DS5020 server. In this example, the 3958-DD4 server is configured for VTL, but the back-end storage attachment would be the same in an OST configuration. Table 4-3 shows the detailed port connections.




Figure 4-8 Single node FC cabling

Table 4-3 Fibre Channel connectivity to back-end storage for stand-alone VTL and OST

Label	From	To
1	DD4A S6 1	Disk Array, Disk Controller A-1 H1
2	DD4A S7 1	Disk Array, Disk Controller B-1 H1
5	CNA-1 P2	Disk Array, Disk Expansion A1-1 1B

Label	From	To
6	EXPA1-1 1A	Disk Array, Disk Expansion A1-2 1B
7	CNA-1 P1	Disk Array, Disk Expansion A1-3 1B
9	CNB-1 P1	Disk Array, Disk Expansion A1-2 1B
10	EXPB1-1 1B	Disk Array, Disk Expansion B1-2 1A
11	CNB-1 P2	Disk Array, Disk Expansion B1-3 1B

**Note:** In Table 4-4 on page 145, we show other detailed port connections that will be used only when we have two DS5020 controllers connected to a single node 3958-DD4 server. This topic is further discussed in 4.6.2, “Two-node cluster configuration” on page 147.

Table 4-4 Fibre Channel connectivity to back-end storage for stand-alone VTL and OST

Label	From	To
3	DD4A S6 2	Disk Array, Disk Controller A-2 H2
4	DD4A S7 2	Disk Array, Disk Controller B-2 H2
12	CNA-2 P2	Disk Array, Disk Expansion A2-1 1B
13	EXPA2-1 1A	Disk Array, Disk Expansion A -2 1B
14	CNA-2 P1	Disk Array, Disk Expansion A2-3 1B
15	CNB-2 P1	Disk Array, Disk Expansion B2-2 1B
16	EXPB2-1 1B	Disk Array, Disk Expansion B2-2 1
17	CNB-2 P2	Disk Array, Disk Expansion B2-3 1B

## TS3000 System Console configuration

In this section, we show the recommended Ethernet connections from the TS7650G server to the TS3000 System Console. The TS3000 System Console is a one-unit IBM System x server with a one-unit KVM, and a network switch provided for the network connection to the 3958-DD4 server and the disk array. In Figure 4-9 on page 146, you can see how to connect the TSSC and 3958-DD4 server to a KVM switch.

The 3958-DD4 server uses the following port to connect to the TS3000 System Console through the TS3000 System Console Ethernet Switch:

- In the 3958-DD4 server, the Remote Supervisor Adapter (RSA) is replaced by the Integrated Management Module (IMM). The functionality of the IMM is the same as in the RSAII adapter.

**Note:** For detailed information about the IMM, refer to the following address:

<http://www-947.ibm.com/systems/support/reflib/imm/index.html>

- In a 3958-DD4 server, either with an OST or VTL configuration, you have two links between the TS3000 Ethernet switch and the ProtecTIER server: One from TS3000 Ethernet switch port 4 to the Ethernet card in slot 5 on port 4 in the 3958-DD4 server, and one from the Ethernet switch port 5 to the built-in Ethernet port of the 3958-DD4 server.

The DS5020 disk array does not have a dedicated connection to the TS3000. You should, however, consider installing the Storage Manager GUI on the TS3000 and connect to the DS5020 through your LAN.

The TS3000 System Console uses one port to connect to the TS3000 System Console Switch to communicate with the disk array and the Gateway server. Another port is used by the TS3000 System Console for the customer connection.

**Note:** Figure 4-9 on page 146 shows the Ethernet and KVM switch connectivity for a single cluster 3958-DD4 server. In this example, the 3958-DD4 server is configured for VTL, but the connectivity to KVM and TS3000 would be the same in an OST configured server. Refer to Table 4-5 on page 146 for a detailed connectivity explanation.




Figure 4-9 TSSC, KVM, and customer Ethernet connections for single node

Table 4-5 TSSC and KVM connections for stand-alone 3958-DD4 server

Label	From	On Device	To	On Device/Location
9	Port 5	TSSC Ethernet switch	Port 1	TSSC
10	Port 2	TSSC	Customer's local area network	Customer specified device
11	Port 4	TSSC Ethernet switch	Port 4, slot 5	Server A
13	Port 6	TSSC Ethernet switch	Port B	Server A
21	Video port	KVM Switch	Video port	SC KVM



Label	From	On Device	To	On Device/Location
22	Port 3	KVM Switch	Video port - Port U3	SC Server
23	Port 1	KVM Switch	Video port - Port E1	Server A
24	Port U1	KVM Switch	Video Port	SC KVM

## 4.6.2 Two-node cluster configuration

The two-node cluster configuration can be built, depending on the customer needs, in terms of capacity, performance, and high availability, as discussed in Chapter 3, “Planning for deduplication and replication” on page 53.

Figure 4-10 shows our rack layout for a two-node cluster configuration. As you can see, we used two frames: one for the 3958-DD4 servers and the other for the storage servers. The server frame is used to host both 3958-DD4 servers, the TS3000 System Console, the TS3000 System Console switch, and two Ethernet switches. A cache frame might be used to host specific supported disk arrays and expansions. For example, both DS5020 might be mounted in a customer-supplied cache frame, while an IBM System Storage DS8300 is shipped with its own dedicated enclosure rack. Therefore, it does not need a customer-supplied cache frame.

106	36	Empty (1u)	36	106	106	36	Empty (1u)	36	106
103	35	Empty (1u)	35	103	103	35	Empty (1u)	35	103
100	34	Empty (1u)	34	100	100	34	Empty (1u)	34	100
97	33	Empty (1u)	33	97	97	33	Disk Expansion (3u)	33	97
94	32	Empty (1u)	32	94	94	32	Empty (1u)	32	94
91	31	Empty (1u)	31	91	91	31	Empty (1u)	31	91
88	30	Empty (1u)	30	88	88	30	Disk Expansion (3u)	30	88
85	29	Empty (1u)	29	85	85	29	Empty (1u)	29	85
82	28	Empty (1u)	28	82	82	28	Empty (1u)	28	82
79	27	Empty (1u)	27	79	79	27	Disk Expansion (3u)	27	79
76	26	Empty (1u)	26	76	76	26	Empty (1u)	26	76
73	25	Empty (1u)	25	73	73	25	Empty (1u)	25	73
70	24	Empty (1u)	24	70	70	24	Disk Controller (3u)	24	70
67	23	Empty (1u)	23	67	67	23	Empty (1u)	23	67
64	22	Empty (1u)	22	64	64	22	Empty (1u)	22	64
61	21	Empty (1u)	21	61	61	21	Empty (1u)	21	61
58	20	Empty (1u)	20	58	58	20	Empty (1u)	20	58
55	19	Empty (1u)	19	55	55	19	Empty (1u)	19	55
52	18	TSSC (1u)	18	52	52	18	Empty (1u)	18	52
49	17	KVM Tray + TSSC sw (1u)	17	49	49	17	Empty (1u)	17	49
46	16	Empty (1u)	16	46	46	16	Empty (1u)	16	46
43	15	1 GB Network Switch (1u)	15	43	43	15	Empty (1u)	15	43
40	14	1 GB Network Switch (1u)	14	40	40	14	Empty (1u)	14	40
37	13		13	37	37	13	Empty (1u)	13	37
34	12	ProtecTIER Server (4u)	12	34	34	12	Empty (1u)	12	34
31	11		11	31	31	11	Disk Expansion (3u)	11	31
28	10		10	28	28	10	Empty (1u)	10	28
25	9	Empty (1u)	9	25	25	9	Empty (1u)	9	25
22	8		8	22	22	8	Disk Expansion (3u)	8	22
19	7	ProtecTIER Server (4u)	7	19	19	7	Empty (1u)	7	19
16	6		6	16	16	6	Empty (1u)	6	16
13	5		5	13	13	5	Disk Expansion (3u)	5	13
10	4	WTI Power Switch (1u)	4	10	10	4	Empty (1u)	4	10
7	3	Power Distribution Unit (PDU)	3	7	7	3	Empty (1u)	3	7
4	2	Power Distribution Unit (PDU)	2	4	4	2	Disk Controller (3u)	2	4
1	1	Empty (1u)	1	1	1	1	Empty (1u)	1	1

Figure 4-10 TS7650G frame layout for two-node cluster and adjacent storage frame

The server frame includes all the hardware components necessary to build the two-node cluster configuration. A additional cache frame might be used to host specific supported disk arrays and expansions. Figure 4-10 on page 147 shows the recommended position of each hardware component in the server frame and storage frame. Some of the storage rack units are left empty for a potential storage capacity growth.

The two-node cluster configuration example is composed of:

- ▶ Two 36 U frames.
- ▶ Two 3958-DD4 servers. A combination of 3958-DD1 and 3958-DD3 servers is also supported.

**Note:** A Cluster Connection Kit feature (FC3437) is required in a two-node cluster configuration, which includes two Ethernet switches and one remote network power switch.

- ▶ One TS3000 System Console and an accompanying Keyboard/Video/Monitor (KVM) Kit.
- ▶ WTI Power switch to control the power of both nodes.
- ▶ Two Network Ethernet Switches for two-node cluster connections.
- ▶ Disk arrays, in our example, two dual controller IBM System Storage DS5020.
- ▶ Six expansions IBM TotalStorage DS5020 EXP520 Expansions. Each DS5020 is attached to three EXP520 expansions.

This two-node cluster configuration may be used in enterprise environments where higher performance is required. Again, consider a careful assessment for performance and capacity planning. Refer to Chapter 3, “Planning for deduplication and replication” on page 53 for a discussion of performance and capacity planning.

In the two-node cluster configuration, the two-node cluster infrastructure provides the functions for two 3958-DD4 nodes (servers) to work together as a storage two-node cluster. The storage two-node cluster provides a consistent file system image across the two-nodes in the cluster, allowing them to simultaneously read from and write to a single shared file system.

Using a two-node cluster system provides increased performance with sufficient disk resources. The two servers can share the backup load and increase ProtecTIER's performance.

Figure 4-11 illustrates the details of the ProtecTIER two-node cluster setup.

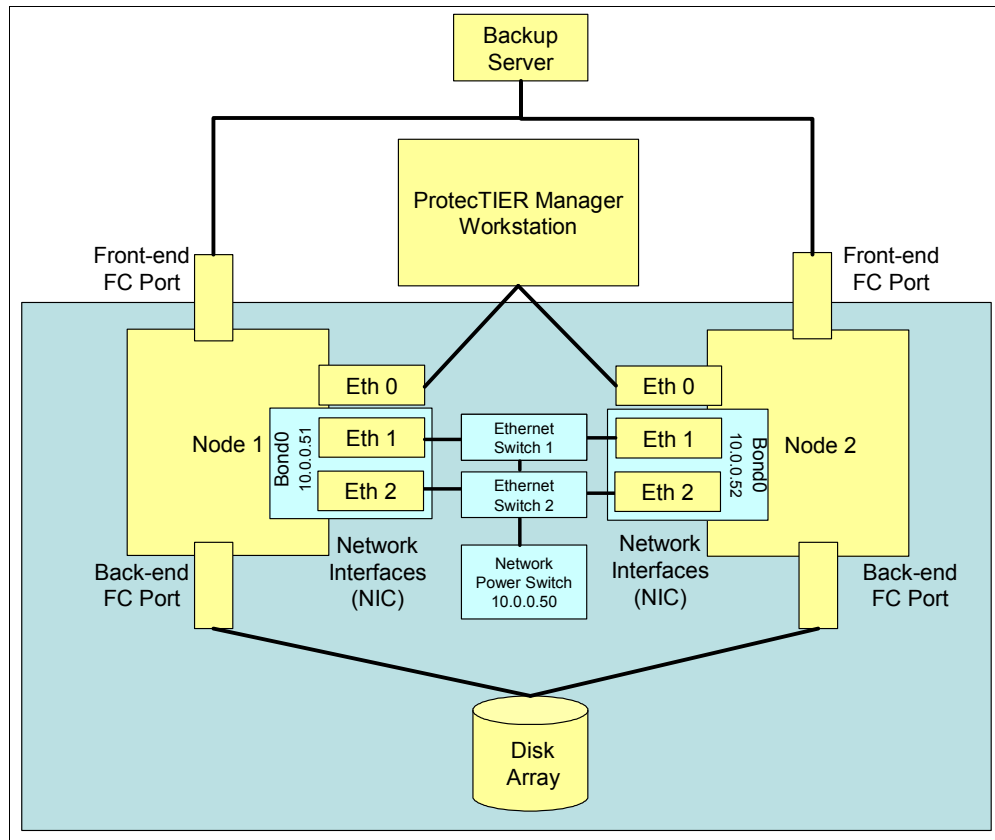


Figure 4-11 Two-node cluster diagram

Each node connects through front-end Fibre Channel ports to the backup server and through back-end Fibre Channel ports to disk arrays. The network interface cards enable the nodes to connect to each other and to the ProtecTIER Manager workstation.

The Eth0 network interface connects to the ProtecTIER Manager workstation, and Eth1 and Eth2 NICs are used in the two-node cluster internal network. Eth1 and Eth2 are usually bonded to a virtual master device called Bond0. By default, Eth2 is the active port and Eth1 is the standby port associated with Bond0.

### DS5020 disk array configuration and EXP520 expansion configuration

In this section, we describe the features of the IBM System Storage DS5020 in a two-node cluster configuration as an example. You can also install all the other listed cache storage systems.

**Note:** To see the entire list of supported disk controllers and disk expansions that can be connected to TS7650G, refer to the TS7650/TS7650G/TS7610 ProtecTIER Deduplication interoperability matrix at the following address:

[ftp://service.boulder.ibm.com/storage/tape/TS7650\\_support\\_matrix.pdf](ftp://service.boulder.ibm.com/storage/tape/TS7650_support_matrix.pdf)

This disk array has controller and power redundancy to avoid a single point of failure (SPOF). Each disk array can host up to 16 disk drives on FC or SATA technology and at different sizes. The disk drive characteristics in terms of technology, size, and speed must be chosen after a capacity and performance planning assessment, as described in Chapter 3, “Planning for deduplication and replication” on page 53.

When attaching expansions, drive loops are configured as redundant pairs using one port from each controller, which helps ensure data access in the event of a path/loop or controller failure. This controller can be expanded up to six expansions and in this full configuration the disk array is cabled using all two drive channel pairs, assuming that there are six total expansion enclosures evenly spread out across the drive channel pairs (three each).

Figure 4-12 shows the rear view of the controller disk array that shows the P1 and P2 ports for FC connections to expansion drawers and the H1 and H2 ports for FC connections to the 3958-DD4 server.




Figure 4-12 Rear view of DS5020 controller

Figure 4-13 shows the rear view of the expansion drawer that shows the ports available for connections. Only the Ports 1A and 1B are used to connect to the cache controller and the expansion units. Port 2A and 2B are not used.




Figure 4-13 Rear view of the DS5000 Expansion unit

Figure 4-14 shows the FC connectivity layout of the connections between both 3958-DD4 servers and both DS5020 controllers. In this example, both 3958-DD4 servers are configured for VTL, but the back-end storage attachment would be the same in OST configuration. In Table 4-6 on page 152, you can find the detailed port connections.

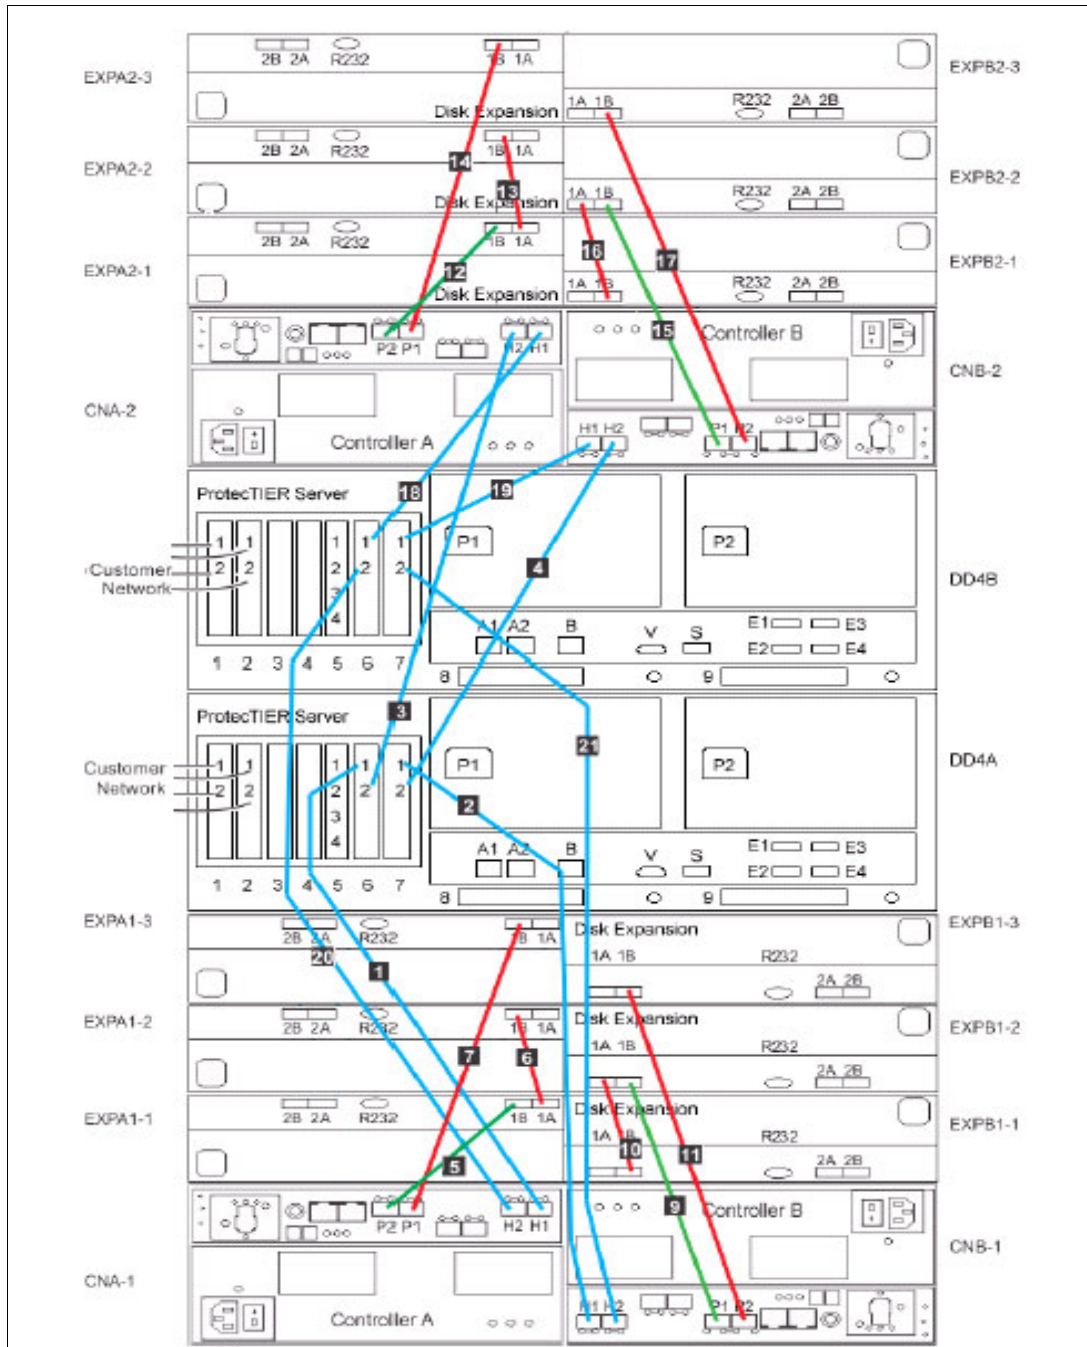


Figure 4-14 Two-node cluster Fibre Channel cabling

In each 3958-DD4 server back-end storage repository attachment, there are two dual-port Qlogic QLE2562 8 Gbps Expansion Cards installed, in VTL configuration as well as in OST configuration. The back-end storage cabling is an example of a DS5020 storage unit with two controllers and three expansions on each controller attached to a two-node cluster 3958-DD4 server configured for VTL, as shown in Figure 4-14 on page 151. The detailed description of the connected ports is shown in Table 4-6

**Note:** Consider mounting the two disk arrays and their expansions in a dedicated frame (cache frame) for easier installation and maintenance. We do not show the frames (server and cache) in order to provide an easy view of the connections between the components.

Table 4-6 Fibre Channel connectivity to back-end storage for stand-alone VTL and OST

Label	From	To
1	DD4A S6 1	Disk Array, Disk Controller A-1 H1
2	DD4A S7 1	Disk Array, Disk Controller B-1 H1
3	DD4A S6 2	Disk Array, Disk Controller A-2 H2
4	DD4A S7 2	Disk Array, Disk Controller B-2 H2
5	CNA-1 P2	Disk Array, Disk Expansion A1-1 1B
6	EXPA1-1 1A	Disk Array, Disk Expansion A1-2 1B
7	CNA-1 P1	Disk Array, Disk Expansion A1-3 1B
9	CNB-1 P1	Disk Array, Disk Expansion A1-2 1B
10	EXPB1-1 1B	Disk Array, Disk Expansion B1-2 1A
11	CNB-1 P2	Disk Array, Disk Expansion B1-3 1B
12	CNA-2 P2	Disk Array, Disk Expansion A2-1 1B
13	EXPA2-1 1A	Disk Array, Disk Expansion A -2 1B
14	CNA-2 P1	Disk Array, Disk Expansion A2-3 1B
15	CNB-2 P1	Disk Array, Disk Expansion B2-2 1B
16	EXPB2-1 1B	Disk Array, Disk Expansion B2-2 1
17	CNB-2 P2	Disk Array, Disk Expansion B2-3 1B

## TS3000 System Console configuration

In this section, we show the Ethernet connections between a TS7650g two-node cluster and the TS3000 System Console. The TS3000 is a one unit System x server with a one unit KVM and a network switch provided for the network connection to the 3958-DD4 server and disk array. In Figure 4-15, you can see how to connect the TSSC and both 3958-DD4 servers to the KVM switch.




Figure 4-15 TTSSC, KVM, and customer Ethernet connections in the two-node cluster

The 3958-DD4 server uses the following ports to connect to the TS3000 System Console through the TS3000 System Console Ethernet Switch:

- ▶ In the 3958-DD4 server, the Remote Supervisor Adapter (RSA) is replaced by the Integrated Management Module (IMM). The functionality of the IMM is the same as the RSAll adapter.

**Note:** For more information, refer to the following address:

<http://www-947.ibm.com/systems/support/reflib/imm/index.html>

- ▶ In a 3958-DD4 server, either with OST or VTL configuration, you have two links between the TS3000 Ethernet switch and the ProtecTIER server: one from the TS3000 Ethernet switch port 4 to an Ethernet card in slot 5 on port 4 in the 3958-DD4 server, and one from the Ethernet switch port 5 to the built-in Ethernet port of the 3958-DD4 server.

The two DS5020 disk arrays do not have a dedicated connection to the TS3000. However, consider installing the Storage Manager GUI on the TS3000 and connect to the DS5020 serve through your LAN.

The TS3000 System Console uses one port to connect to the TS3000 System Console switch and communicate with the cache controller and the Gateway server. Another port is used by TS3000 System Console for your connection.

**Note:** In Figure 4-15 on page 153, you can see the TS3000 and KVM switch connections in a 3958-DD4 server configured for VTL. These connections are exactly the same for a 3958-DD4 server with an OST configuration. In Table 4-7 on page 154 you can find the TSSC and KVM connections for a clustered 3958-DD4 server.

Table 4-7 TSSC and KVM connections for the clustered 3958-DD4 server

Callout	From	On Device	To	On Device/Location
9	Port 5	TSSC Ethernet switch	Port 1	TSSC
10	Port 2	TSSC	Customer's local area network	Customer specified device
11	Port 4	TSSC Ethernet switch	Port 4, slot 5	Server A
12	Port 3	TSSC Ethernet switch	Port 4, slot 5	Server B
13	Port 6	TSSC Ethernet switch	Port B	Server A
14	Port 1	TSSC Ethernet switch	Port B	Server B
21	Video port	KVM Switch	Video port	SC KVM
22	Port 3	KVM Switch	Video port - Port U3	SC Server
23	Port 1	KVM Switch	Video port - Port E1	Server A
24	Port U1	KVM Switch	Video Port	SC KVM
25	Port 2	KVM Switch	Video port - Port E1	Server B

### 4.6.3 Host attachment ports configuration for VTL and OST

In this section, we show the 3958-DD4 port layout for host attachment connections. These ports are the same for both a single node configuration or two-node cluster configuration.



For host attachment of a 3958-DD4 VTL, there are two dual-port Emulex LPe12002 8 Gbps Expansion Cards installed in slots 1 and 2, as shown in Figure 4-16.




Figure 4-16 Host attachment for a 3958-DD4 VTL configuration

For host attachment of a 3958-DD4 OST configuration, there are two quad port 1 Gbps 10/100/1000Base-TX PCIe Ethernet adapters installed in slots 3 and 4, as shown in Figure 4-17. You also have the option to have two dual port 10 Gbps PCIe Ethernet adapters installed in the same slots instead.




Figure 4-17 Host attachment for a clustered 3958-DD4 OST configuration

## 4.7 Usage considerations

The 3958-DD4, the 3958-AP1, and the 3959-SM1 solutions have some features that we discuss in the following sections. We describe how to use them in a customer environment and explain what you should consider, as well as the pitfalls. This section discusses these features of the ProtecTIER and the best practices for planning the deployment to guarantee performance.

### 4.7.1 Virtual tape libraries and drives

The following virtual tape libraries are emulated with ProtecTIER:

- ▶ DTC VTF virtual tape libraries
- ▶ ATL P3000 virtual tape libraries
- ▶ IBM TS3500 virtual tape libraries
- ▶ IBM V-TS3500 virtual tape libraries

The type of tape drive model depends on the library type chosen.

### 4.7.2 Fibre Channel ports and host assignment considerations

It is important to configure properly the connection between the FC ports of the backup server and the front-end ports of the TS7600 family in a VTL configuration. We recommend that the backup server have at least two HBAs to obtain redundancy paths to ProtecTIER. All FC front-end ports must be used to obtain best performance and redundancy. This connection can be done in loop mode (direct attach) or in SAN Fabric mode. To get high redundancy, you should use a loop configuration or two independent fabrics. You may also use only one fabric, but in this case the single fabric becomes a single point of failure (SPOF).

In Figure 4-18, you can see a possible architecture where the backup server is direct attached to both a virtual tape library and a physical tape library. Note that in this sample the virtual tape library is composed of:

- ▶ The TS7650G ProtecTIER server, for emulating physical tape and for deduplicating backup streams.
- ▶ The disk array on which the data repository is stored (user data and metadata).

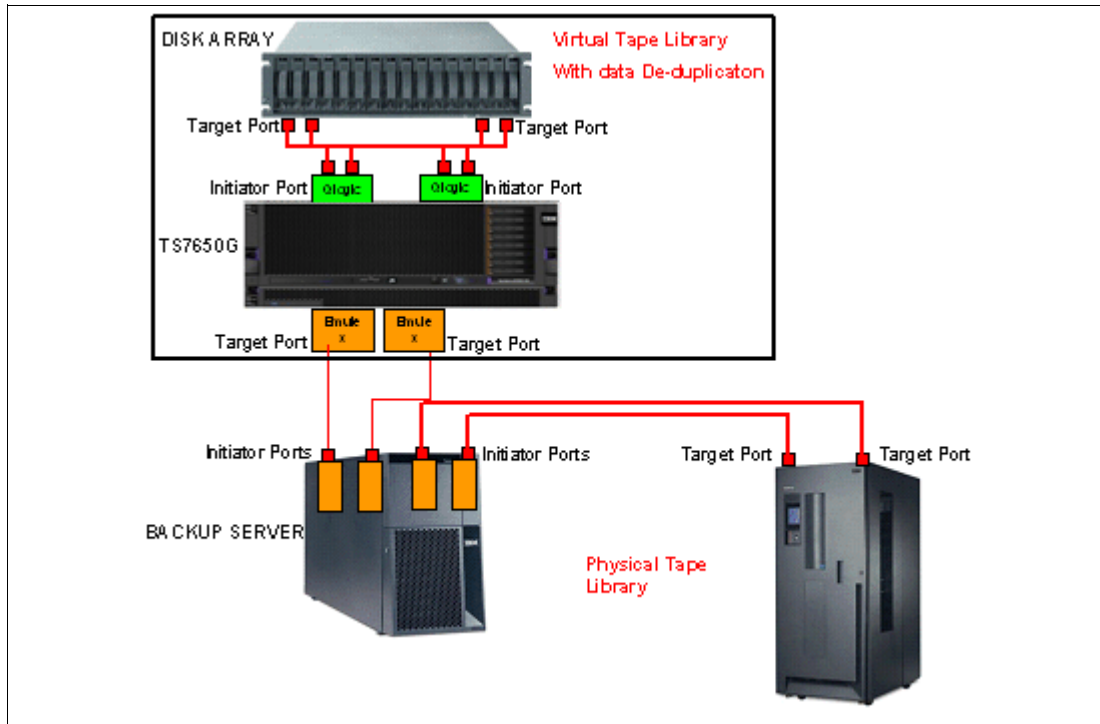


Figure 4-18 Architecture composed of a backup server, a virtual tape library, and a physical tape library

**Note:** The backup server should have dedicated HBA ports for the 3958-DD4, 3958-AP1, and 3959-SM1 virtual tape libraries. These ports could be shared with a real tape library, but it is a best practice that the physical tape library not be in the same zone as the ProtecTIER front-end ports.

### 4.7.3 Firewall environments: Ports assignments in ProtecTIER Replication Manager

Ethernet ports on all ProtecTIER family involved in the native replication operation at both the primary and secondary sites should be configured as detailed in Chapter 2, “Setting up ProtecTIER”, of *IBM System Storage TS7600 with ProtecTIER User’s Guide*, GC53-1156. In addition, in user environments where firewalls are used, it is important to open the following TCP ports for IP replication to function:

- ▶ The ProtecTIER Replication Manager uses ports 6202, 3501, and 3503.
- ▶ The replication operation between any two repositories requires the following TCP ports: 6520, 6530, 6540, 6550, 3501, and 3503.

For more information about replication, refer to Chapter 11, “Native replication and disaster recovery” on page 575.

**Note:** ProtecTIER Replication does not use any UDP ports.

## 4.8 Installation planning

In this section, we present important details to help you plan for the installation and basic setup of the TS7600 family. Planning is primarily a customer responsibility.

### 4.8.1 Installation worksheets

We provide implementation worksheets in Appendix A, “Installation and implementation checklists” on page 687 that can be used by you to log the required information to implement a TS7600 family. Most of the worksheets must be completed prior to the installation of TS7600 family, as the values inserted are needed to complete the installation.

### 4.8.2 Supported backup server operating environments

The TS7600 family can be used with a wide range of operating environments. For the most current list of supported products or for more information about support, refer to the System Storage Interoperation Center (SSIC) at the following address:

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

Currently, the TS7600 family supports the following operating systems at the minimum levels indicated:

- ▶ IBM AIX 5L™ V5.1, V5.2, V5.3, and V6.1
- ▶ Sun Solaris 8, 9, and 10
- ▶ Microsoft Windows 2003 and 2008
- ▶ HP-UX 11v1, 1v2, and 11v3
- ▶ Red Hat ES4 and ES5
- ▶ SUSE 10
- ▶ IBM i V5R4 and V6R1
- ▶ Linux on System z/SUSE 10

### 4.8.3 Planning the ProtecTIER installation

In this section, we describe the prerequisites necessary to install ProtecTIER, focusing on:

- ▶ Hardware prerequisites
- ▶ Software prerequisites

#### **Prerequisites**

Before installing the ProtecTIER servers, verify that the following prerequisites are met.

#### ***infrastructure prerequisites***

- ▶ All cabling and physical connections between hardware components must be done prior to installation. This includes the local LAN connections with an assigned IP address.
- ▶ A workstation is connected to the customer network.
- ▶ RAID groups have been created (for a Gateway).
- ▶ The Network Power Switch is configured (for a dual node cluster).

- ▶ Rack space is available.
- ▶ There are power connections.

### ***Operating system prerequisites for 3958-DD4 and 3958-AP1 servers***

IBM System Storage ProtecTIER Enterprise Edition V2.5 Software and Red Hat Enterprise Linux 5.4 Advanced Platform (RHELAP) x86\_64 are loaded on the 3958-DD4 server with your initial order of the Gateway. Ordering both the Red Hat PID (use PID 5639-RHL with feature 0017 1-year support or feature 0019 3-year support) and the ProtecTIER PID (ProtecTIER Enterprise Edition V2.5 PID 5639-XXB for 3958-DD3 or ProtecTIER Appliance Edition V2.5 Software PID 5639-XXP for 3958-AP1) is required when the Gateway hardware is ordered.

### ***Connectivity prerequisites***

The supported Fibre Channel switches are listed at the System Storage Interoperation Center (SSIC) at the following address:

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

Because the ProtecTIER system is accessed through the ProtecTIER Manager installed on another computer, network connections must be configured to allow access.

### ***File system prerequisites***

ProtecTIER stores data in Red Hat Global File Systems (GFS). The number and size of the file systems depends on the size and parameters of your configured repository (disk array LUNs). For information about the configured capacity of subsystem and sizing of all LUNs, call your IBM representative.

IBM Tivoli Storage Manager and other compatible software offerings can provide storage and tape management software for the TS7600 family of products. The supporting software and applications must be obtained separately from IBM, IBM Business Partners, or ISVs. For a current list of compatible software, call your IBM representative or visit the following address:

<http://www-03.ibm.com/systems/storage/tape/library.html#compatibility>

### ***ProtecTIER Manager workstation prerequisites***

You must use the ProtecTIER Manager software to configure, manage, and monitor the operation of the TS7600 family of products. You are responsible for obtaining the workstation where the ProtecTIER Manager software is installed. You can install the software from the supplied ProtecTIER Manager application CD or download the ProtecTIER Manager from the IBM website at the following address:

<http://www-03.ibm.com/systems/storage/tape/>

The following are hardware and operating environment requirements for the ProtecTIER Manager workstation:

- ▶ Hardware requirements:
  - x86 (Pentium or higher) microprocessor
  - 256 MB of memory
  - 1.2 GB of disk space access to the ProtecTIER service node's IP address (port 3501 and 3502 have to be open on the firewall)
  - Keyboard, mouse, and CD-ROM drive
  - Screen resolution of 1024x768 or higher
  - 24-bit color or higher

- ▶ Operating environments supported:
  - Windows 32-/64-bit (2003, XP or 7)
  - Linux Red Hat 32-/64-bit (Enterprise 4 or later)

#### 4.8.4 Installation tasks

Installation involves the IBM Systems Service Representative (SSR), the IBM ProtecTIER Specialist, and customer personnel.

##### **IBM System Services installation responsibilities**

The IBM System Services Representative (SSR) and the IBM ProtecTIER Specialist or Lab Services specialist, will install and configure the TS7650G or TS7650 server and complete the following tasks.

##### ***IBM System Service Representative (SSR) tasks***

- ▶ Installs the TS7650G or TS7650 hardware components purchased with the solution into the server and disk storage frames.
- ▶ Labels and connects power, Ethernet, and Fibre Channel cables, including OST front-end cables, as necessary.
- ▶ Connects the TS7650G or TS7650 server to the customer's local area network and replication network, if applicable.
- ▶ Starts the system.
- ▶ Verifies the accuracy of hardware installation and cabling. Performs a visual check of fault indicator LEDs.
- ▶ Configures the TSSC for use with the TS7650G or TS7650 server.
- ▶ Configures the RAS package on the servers.
- ▶ Tests Call Home on the servers.

##### ***IBM ProtecTIER specialist and Lab Services specialist tasks***

- ▶ Oversee project management for the installation and integration of the engagement.
- ▶ Oversee change management and process control for the installation.
- ▶ Coordinate and schedule IBM resources for customer installations, and act as a focal point of contact for coordination of installation services.
- ▶ Schedule and facilitate planning and solution assurance conference calls.
- ▶ Create and document the installation service process.
- ▶ Configure ProtecTIER on the stand-alone server or Server A in a cluster, *including OST if applicable*, and create the file system.
- ▶ Install ProtecTIER Replication Manager on one of the ProtecTIER servers being used for replication, if applicable.
- ▶ Install ProtecTIER Manager on the TSSC or ProtecTIER Manager workstation, register each server as a new ProtecTIER node, and create the repository.
- ▶ Add the second server (Server B) to the ProtecTIER cluster, if applicable.
- ▶ Verify cluster operation, if applicable.
- ▶ Perform RAS verification tasks.
- ▶ Release the system to the customer. Advise the customer that it is their responsibility to create and configure the replication grid, if applicable.
- ▶ Document and report installation results.

**Note:** You should add a step to your implementation plan to verify that the SSR completed all the steps.

## Customer installation responsibilities

Customers are responsible for preparing the installation site prior to the installation of the TS7650G solution. All physical planning for the TS7650G is a customer responsibility. In summary, the customer is responsible for:

- ▶ Completing the planning, preparation, and installation tasks described in *IBM System Storage TS7600 with ProtecTIER Introduction and Planning Guide for the TS7650G (3958 DD4)*, GC53–1152.
- ▶ Meeting the pre-installation requirements outlined in *IBM System Storage TS7600 with ProtecTIER Introduction and Planning Guide for the TS7650G (3958 DD4)*, GC53–1152.
- ▶ Completing the worksheets provided in *IBM System Storage TS7600 with ProtecTIER Introduction and Planning Guide for the TS7650G (3958 DD4)*, GC53–1152. For convenience, blank copies of the worksheets are included in the appendixes of this document.
- ▶ Purchasing, installing, and configuring (if necessary) all the hardware components not included in the purchase of the Gateway.
- ▶ Confirming that an existing TSSC that is being used with the TS7650G serve has FC 2719. Call your IBM representative to find this FC. Unless an existing TSSC resides in the server frame and is used in conjunction with a KVM switch, the customer must provide a USB keyboard and graphics-capable monitor for server configuration.
- ▶ Ensuring that a separate USB keyboard and graphics-capable monitor are available for use during installation.
- ▶ Providing the site preparation, that is, cabling and wiring for connections to the host, cooling and heating, telephone service, safety, and security.
- ▶ Providing network connections (cables and switches).
- ▶ Providing an SNMP catcher, host clients, and email services for RSA/IMM alerts.
- ▶ Providing all the necessary IP addresses for the installation of the TS7600, which includes IP addresses for the TS7600, including the IP addresses for replicating and one IP address for the TS3000 System Console.
- ▶ For TS7650G installations where customer-supplied frames, disk arrays, and expansions are used, these items must be fully installed and operational before the installation of the 3958-DD4 can begin.
- ▶ Providing a ProtecTIER Manager workstation and the appropriate adapters and cables. Fibre Channel cables are required to attach the TS7650 to various server adapters if the console was not shipped.
- ▶ Providing client machines and the required Fibre Channel and Ethernet switches.
- ▶ Downloading or obtaining from IBM and installing designated Machine Code (microcode, basic input/output system code (called BIOS), utility programs, device drivers, and diagnostics delivered with an IBM machine) and other software updates in a timely manner from an IBM intranet website or from other electronic media, and following the instructions provided.

**Note:** Customers may request that IBM install machine code changes, but might be charged for that service.

### ***Hardware, cabling, and infrastructure responsibilities***

In general, the customer is responsible for providing the appropriate infrastructure resources, such as frames and power distribution units, and the cables and switches required to support TS7600 server connections to the ProtecTIER Manager console, cache controllers, and cache modules when applicable. The customer must provide the following hardware and cabling resources prior to the installation of the 3958-DD4 server:

- ▶ One server frame with two power distribution units to provide redundancy.
- ▶ One or two cache frames, as required, based on the amount of rack space needed for the cache components. Each cache frame must have two power distribution units to provide redundancy.
- ▶ CAT5e or CAT6 Ethernet cables for connecting the TS7600 family to the ProtecTIER Manager workstation, cache controllers, and TS3000 System Console. Note that CAT5 cables cannot be used because they do not provide sufficient data transfer rates for the TS7600 family.
- ▶ Fibre Channel cables for connecting the cache controllers to the cache modules and to SAN switches.
- ▶ A ProtecTIER Manager workstation, as described in “ProtecTIER Manager workstation prerequisites” on page 159.

**Note:** All customer supplied hardware components, including frames, switches, client host platforms, and media, must be installed and verified as operational by the customer prior to the installation of the ProtecTIER servers.

## **4.8.5 Host attachment considerations**

The TS7600 family is supported by a variety of Fibre Channel (FC) switches and Fibre Channel directors that support Fibre Channel Arbitrated Loop (FC-AL). The support is dependent on the server, operating system, and host bus adapter that are being used.

For a current list of supported products or more information about the support and prerequisites, go to the following address:

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

For the supported host bus adapter (HBA) firmware levels and SAN fabric components, go to the following address:

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

## **4.8.6 SAN configuration**

One of the benefits of ProtecTIER systems is that it takes advantage of an existing SAN environment. Media servers and one or two (single node or two-node cluster configurations) disk arrays can connect to one or two 3958-DD4 servers:

- ▶ Through loop, for example, direct-attached to the ProtecTIER systems
- ▶ Through a SAN fabric

Consider using patch cords and loop technology between the ProtecTIER system and the disk array to avoid expensive fabric switches and to simplify the configuration. When the environment needs a SAN configuration, consider using at least two FC switches to provide redundancy against single points of failure. The multipath software installed on the TS7600 family is the native Linux multiple devices administrative interface (MDADM).



## Loop environment

It is possible to use simple patch cords to connect the front-end (FE) and back-end (BE) ports of the 3958-DD4 to disk arrays and the TS7600 family to media servers, respectively.

Remember that the 3958-DD4 Gateway server is shipped with two Emulex dual-port LPe12002 8 Gbps expansion cards for the front-end connections (VTL) and two Qlogic dual-port QLE2562 8 Gbps expansion cards for the back-end connections (VTL and OST). The two Emulex cards for front-end connections will be removed in an OST configured system and the new Ethernet cards (FC 3456 or FC 3457) are added. For more information on this topic, refer to 4.3.2, “Features codes for 3958-AP1 server” on page 132 or 4.3.3, “Feature codes for 3958-DD4 server” on page 135.

This method does not need a SAN configuration or zoning, and therefore is simpler. Because each 3958-DD4 or 3958-AP1 server is shipped with four front -end FC ports, it is possible to attach:

- ▶ Up to four media servers without redundancy
- ▶ Up to two media servers with redundancy (multipath redundancy)

For more information about the supported host bus adapter media servers, go to the following address:

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

In Figure 4-19, you see the configuration without multipath redundancy, where each of the four media servers has only one FC port available to connect to the 3958-DD4 or 3958-AP1 server with VTL configuration. This configuration is less desirable because of the lack of redundancy.

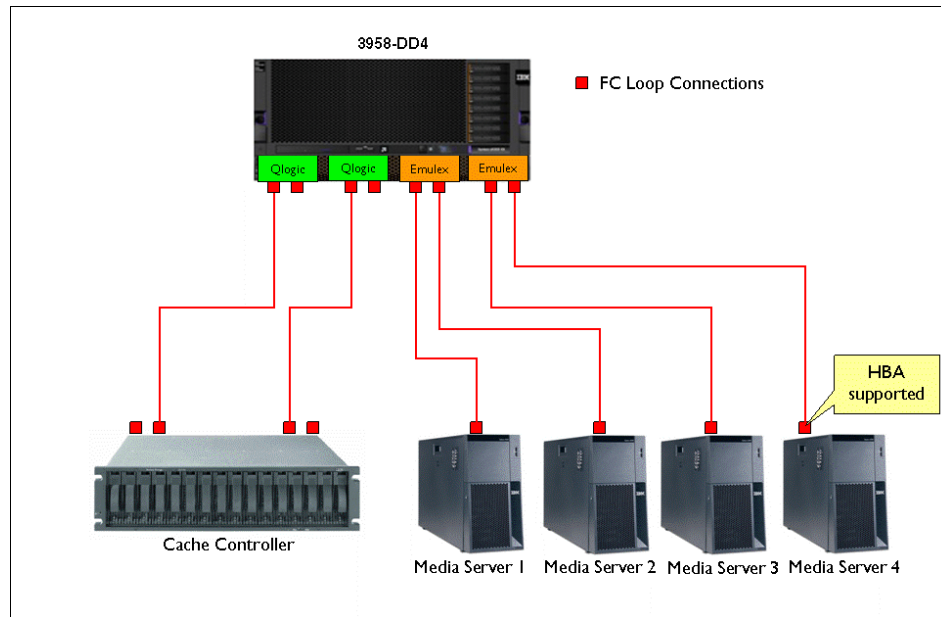


Figure 4-19 Non-redundant loop connections

In Figure 4-20, you see the configuration with multipath redundancy, where each of the two media servers has two FC ports available to connect to the 3958-DD4 or 3958-AP1 server with VTL configuration. Consider using this configuration and, on each media server, consider using two single-port FC expansion cards instead of one dual-port FC expansion card to protect the server from a failure of the expansion card.

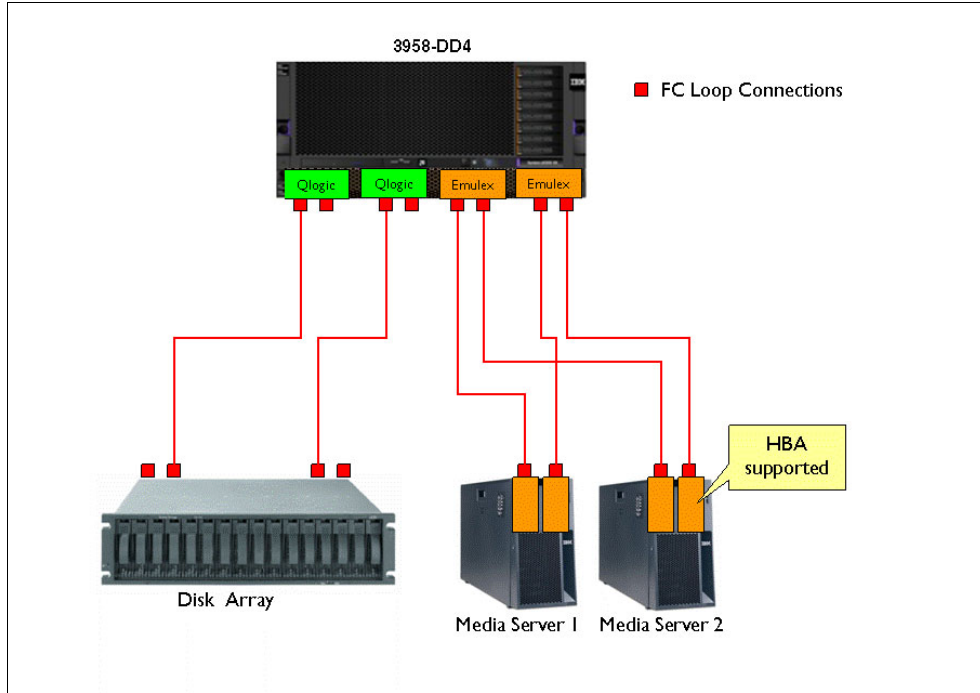


Figure 4-20 Redundant loop connections

In Figure 4-21, you see the configuration with multipath redundancy where each of the four media servers has two FC ports available to connect to the two-node cluster with VTL configuration. Consider using this configuration. On each media server, consider mounting two single-port FC expansion cards instead of one dual-port FC expansion card to protect the server from a failure of the expansion card.




Figure 4-21 Loop connections in a two-node clustered configuration

## SAN environment

In this section, we describe how to deploy ProtecTIER in an existing or new SAN infrastructure. The advantage of this configuration is that the ProtecTIER system ports can be managed by SAN tools, such as Brocade tools, and more media servers can be attached. The maximum number of media servers that can be direct attached to the ProtecTIER are four (single node configuration) or eight (two-node cluster configuration), but these scenarios are not recommended because of the lack of redundancy.

In addition, you should consider the topic of oversubscription when there are many backup servers. When designing a SAN, it is important to consider the possible traffic patterns to determine the possibility of oversubscription, which might result in degraded performance. Oversubscription to a storage device may be overcome by adding another adapter to the storage array and connecting into the fabric.

The 3958-DD4 and 3958-AP1 FC ports are either initiators or targets, depending on whether they are back-end or front-end ports. Therefore, standard SAN configuration rules apply (number of fabric switches between nodes, uplink saturation rules, and so on). The example in Figure 4-22 shows the required zoning for a fabric configuration. All backup servers must be in separate zones with their 3958-DD4 and 3958-AP1 front-end ports. Disk array ports must also be zoned to isolate them from other devices.

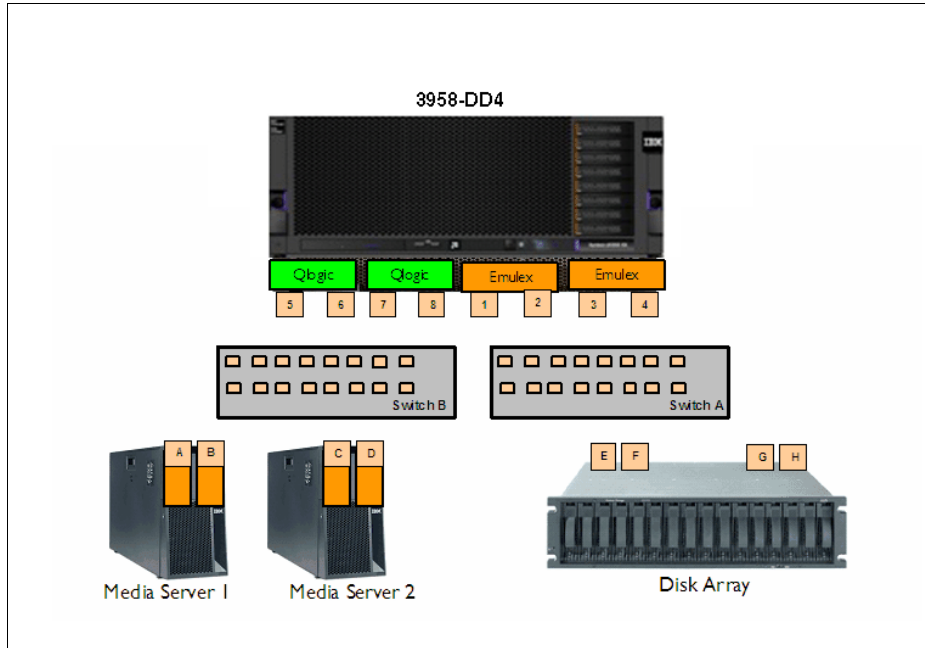


Figure 4-22 Zoning configuration

**Note:** ProtecTIER system front-end ports and back-end ports must never be in the same zone or on the same unzoned fabric. Do not put storage LUNs in the same zone as ProtecTIER system front-end ports.

Ensure that the following checklist of items is addressed:

- ▶ Approved IBM switch make/model
- ▶ Approved IBM firmware level (typically, the latest firmware level from the vendor)
- ▶ Free ports to connect ProtecTIER system front-end and back-end ports
- ▶ AC power to switch (if new)
- ▶ Ethernet port available on network
- ▶ Ethernet cables run to switch/fabric
- ▶ Zoning identified
- ▶ Zoning configured
- ▶ Fibre cables run to ProtecTIER system back-end ports
- ▶ Fibre cables run to ProtecTIER system front-end ports
- ▶ Fibre cables run to disk array ports
- ▶ Fibre cables run to media server ports

To achieve high redundancy, you should use two independent fabrics, as shown in Table 4-8 and Figure 4-22 on page 166. However, load balancing and failover also work with one SAN fabric.

Table 4-8 Zoning table

Zone Media Server	FE port (Emulex)	Media Server port	BE port (Qlogic)	Disk array port
Zone Media Server 1-1A	1	A	-	-
Zone Media Server 1-3A	3	A	-	-
Zone Media Server 1-1B	1	B	-	-
Zone Media Server 1-3B	3	B	-	-
Zone Media Server 2-2C	2	C	-	-
Zone Media Server 2-4C	4	C	-	-
Zone Media Server 2-2D	2	D	-	-
Zone Media Server 2-4D	4	D	-	-
Zone Storage	-	-	E-F-G-H	5-6-7-8

You should consider performing the following actions to properly configure your SAN environment:

- ▶ 3958-DD4 or 3958-AP1 front end-ports should not be in a zone with any other 3958-DD4 or 3958-AP1 front-end ports.
- ▶ If possible, dedicated HBA ports, in the media server, should be zoned to single 3958-DD4 or 3958-AP1 front-end ports.
- ▶ Ensure that Inter-Switch Links (ISL) between switches, connected to 3958-DD4 or 3958-AP1 ports and backup servers or disk arrays, are not oversubscribed.
- ▶ Use 8 Gbps for the 3958-DD4 or 3958-AP1 back-end and front-end connections.
- ▶ Do not put storage LUNs in the same zone as 3958-DD4 or 3958-AP1 ProtecTIER front-end ports.
- ▶ For best performance and reliability, use separate HBA ports on the media server to connect to 3958-DD4 or 3958-AP1 servers only. If this is not possible, establish zones where ProtecTIER system ports only see media server HBA ports and not other tape devices.
- ▶ Where possible, distribute LUNs evenly between ProtecTIER system back-end ports.

Native Red Hat load balancing and failover software (MDADM) is the multipath software included in the Red Hat Enterprise Linux 5 Advanced Platform 64-bit license. There is nothing additional to purchase. Native Red Hat failover works with any disk array that supports active-active path failover.

For enabling back-end failover / load-balancing products, you must perform the following actions:

- ▶ This will require no change in the 3958-DD4 and 3958-AP1, as it relies on the SCSI naming conventions for each of the tools.
- ▶ After the failover software is installed, the file systems should be created and mounted as /mnt/fs1 to /mnt/fsx (x = total number of metadata + user data file systems).
- ▶ All file systems should be added to the fstab and mounted upon Linux bootup or the 3958-DD4 or 3958-AP1 server will not use them to create a repository.

Finally, it is also possible to use ProtectTIER servers in a mixed environment, where some of the ports are connected to a fabric and some are connected through loop (patch cords). ProtectTIER Manager is used to configure each port as a fabric or loop.

**Note:** The 3958-DD4 server currently supports only back-end Fibre Channel connectivity to the disk array, even if the disk array is allowed to be accessed through iSCSI technology.

# **Implementing and administering the IBM System Storage TS7650G and TS7650 servers**

In this part, we describe the setup and host implementation of the TS7650G and TS7650 servers. We also provide you with operational guidance and all the information required to monitor the servers.







# IBM System Storage TS7600 with ProtecTIER initial setup

In this chapter, we provide information about how to install ProtecTIER Manager and how to set up the ProtecTIER system so that it is ready to use with backup applications. We cover:

- ▶ Enabling SNMP support
- ▶ Getting started
- ▶ Adding nodes in ProtecTIER Manager
- ▶ Creating repositories
- ▶ Setting up the virtual libraries and cartridges
- ▶ Setting up replication and replication policies

At this point, the IBM System Service Representative (SSR) and the IBM ProtecTIER Specialist already have:

- ▶ Installed the TS7650 or TS7650G servers, TS3000, KVM kit, Ethernet switch, and network power switch
- ▶ Checked the disk arrays configuration (provided by the customer)
- ▶ Connected the network cables (provided by the customer) to the network switches (provided by the customer) in the customer configuration
- ▶ Connected the Fibre Channel cables
- ▶ Applied cable labels
- ▶ Set IP addresses according to customer-provided network assignments
- ▶ Created the backup set for the TS7650 or TS7650G servers
- ▶ Verified that the IBM System Storage TS7600 with ProtecTIER hardware is functioning properly
- ▶ Set up the TS3000, which includes connecting the customer-supplied analog phone line and network cable
- ▶ Installed any clustering hardware and software, if applicable
- ▶ Verified fencing functionality, if applicable

For more details, refer to *IBM System Storage TS7600 with ProtecTIER Introduction and Planning Guide*, GC53-1152.

After these tasks are complete, the following steps must be performed to fully set up the ProtecTIER system:

- ▶ Enable ProtecTIER SNMP support.
- ▶ Install ProtecTIER Manager software.
- ▶ Add one or more nodes.
- ▶ Create a repository (TS7650G only).
- ▶ Create a virtual library (robot, drives, slots, and cartridges).
- ▶ Set up ProtecTIER Replication Manager.

## 5.1 Enabling ProtecTIER SNMP support

ProtecTIER responds to SNMP, implementing MIB-2 and generating the appropriate traps. The server responds to SNMP discovery and queries and to the standard MIB requests.

### 5.1.1 Defining the IP address

When the server is installed at the user site, the IP address of the SNMP management station, that is, the IBM TotalStorage Service Console (TSSC), must be made available to the ProtecTIER servers to enable SNMP support. To do this task, perform the following steps:

1. Edit the configuration file. The `snmpd` file is found in the `/etc/snmp` directory.
2. From the command line, use the `vi` editor:

```
vi /etc/snmp/snmpd.conf <Enter>
```

The following output is displayed:

```
#####  
#  
# snmpd.conf  
#  
# - created by the snmpconf configuration program  
#  
#####  
# SECTION: System Information Setup  
#  
# This section defines some of the information reported in  
# the "system" mib group in the mibII tree.  
  
# syslocation: The [typically physical] location of the system.  
# Note that setting this value here means that when trying to  
# perform an snmp SET operation to the sysLocation.0 variable will make  
# the agent return the "notWritable" error code. IE, including  
# this token in the snmpd.conf file will disable write access to  
# the variable.  
# arguments: location_string  
  
#syslocation Unknown (edit /etc/snmp/snmpd.conf)  
  
# syscontact: The contact information for the administrator  
# Note that setting this value here means that when trying to
```

```

# perform an snmp SET operation to the sysContact.0 variable will make
# the agent return the "notWritable" error code. IE, including
# this token in the snmpd.conf file will disable write access to
# the variable.
# arguments: contact_string

#syscontact Root <root@localhost> (configure /etc/snmp/snmp.local.conf)

#####
# SECTION: Access Control Setup
#
# This section defines who is allowed to talk to your running
# snmp agent.

# rwcommunity: a SNMPv1/SNMPv2c read-write access community name
# arguments: community [default|hostname|network/bits] [oid]

rwcommunity ProtecTIER
trapsink localhost

```

3. Scroll down the list to the bottom of the file and locate this line:  
trapsink localhost
4. Use the arrow key on the keyboard to place the cursor under the letter *l* in the word localhost.
5. Press the Delete key until localhost is removed.
6. Press the “a” key (this is for *add* or *insert* mode).
7. Enter the IP address of the TS3000 SNMP management station (Service Console).
8. After the IP address has been entered, press the Esc key, type **wq!** (write-quit), and press Enter. This saves the file.

## 5.1.2 IBM MIB definition file

IBM supplies an MIB definition file called DILIGENT-MIB.txt that can be found in the /usr/share/snmp/mibs directory of the TS7650G or TS7650 AP1 server. This file describes each of the supported traps and should be imported into the SNMP application used at the customer site.

For the TS7610 (3959 SM1 server), the MIB file names are DILIGENT-MIB.txt and IBM-TS-RAS-MIB.txt. The files are located in the /usr/share/snmp/mibs directory.

## 5.1.3 SNMP compatibility

ProtecTIER implements SNMP as follows:

- ▶ MIB implementation

In the MIB-2 System Group, the following fields are implemented:

- sysDescr
- sysObjectID
- sysUpTime
- sysContact

- sysName
- sysLocation
- sysServices

All other parts of the MIB-2 respond in such a way that management tools understand that they are not implemented.

► Traps

The traps generated are:

- coldStart
- warmStart
- authenticationFailure
- operatorInterventionRequired (proprietary)
- recoverableErrorNotification (proprietary)

If authentication of a user fails more than five times in a row, an SNMP authenticationFailure trap is generated.

► Startup consistency checks

ProtectTIER checks its persistent data on startup, using a Power On Self Test (POST) procedure to verify that the application can run. Initialization files are checked for consistency. Resource allocation (specifically, memory) is checked.

Errors encountered might be recoverable or unrecoverable.

- Recoverable errors: A recoverable error is an error from which the ProtecTIER system can recover without losing the user's data. Recoverable errors are logged in the ProtecTIER logs and generate an SNMP warning notification trap.
- Unrecoverable errors: An unrecoverable error is an error that prevents the platform from booting correctly, an unrecoverable consistency check error during ProtecTIER server startup, or any other error that could cause or has caused loss of user data. The server is left in an offline state (booted, responding to TCP/IP SNMP enquiries, and responding to console, Telnet, and modem logins). Unrecoverable errors are logged in the ProtecTIER logs and generate an SNMP error trap.

► Restarting on error

If the ProtecTIER server detects an error at run time, it recovers by rebooting and restarting the ProtecTIER process. If multiple restarts are detected within a short time period, the ProtecTIER server declares an unrecoverable error.

► Alerts

The server generates SNMP traps to higher level management frameworks. In particular, whenever the system enters the online state, it generates a coldStart trap if the platform has rebooted, and a warmStart trap if the system only returned to online state from the offline state. A recoverable error generates a recoverableErrorNotification trap. An unrecoverable error generates an operatorInterventionRequired trap.

## 5.2 Installing ProtecTIER Manager

ProtectTIER Manager is an application that enables you to monitor the status of nodes and two-nodes cluster in your ProtecTIER system, along with the accompanying repositories and services. ProtecTIER Manager is used to initially configure your ProtecTIER system, and can be used to change the configuration of the system. You must install the ProtecTIER Manager application on one or more workstations.

**Note:** The ability to have ProtecTIER Manager (or any additional software) installed directly on the ProtecTIER systems is not supported.

If you are installing ProtecTIER Manager on a workstation on which an older version of ProtecTIER Manager is already installed, uninstall the older version first.

## 5.2.1 Prerequisites

Before you start with the installation of the ProtecTIER Manager on your workstation, make sure that the following prerequisites are met:

- ▶ Hardware requirements
  - x86 (Pentium or higher) microprocessor.
  - 256 MB memory.
  - 100 MB of disk space.
  - Access to the ProtecTIER service node's IP address (port 3501 is open on the firewall).
  - Keyboard, mouse, and CD-ROM drive.
  - Screen resolution of 1024 x 768 or higher. (This is the minimum resolution supported, however, 1280x1024 is recommended.)
  - 24 bit color or higher.
- ▶ Operating system requirements
  - Supported operating environments are:
    - Windows 32-/64-bit (2003/2000/XP)
    - Linux Red Hat 32-/64-bit (Red Hat Enterprise 3 or higher)
  - ProtecTIER Manager requires a properly installed version of the Java runtime environment.
  - If you are planning to run ProtecTIER Manager on a Linux system, configure your graphics card and X Window System. This is done either manually or by using the Xconfigurator utility. For instructions, refer to the appropriate Linux documentation.

For the latest information, refer to *IBM System Storage TS7600 with ProtecTIER Introduction and Planning Guide*, GC53-1152.

## 5.2.2 Installing on Windows

To install on a Windows system, complete the following steps:

1. Insert the IBM System Storage ProtecTIER Manager Enterprise Edition V2.5 DVD into the DVD-ROM drive of the designated ProtecTIER Manager workstation.
2. If the autorun process does not launch automatically, select **Start** → **Run**, type D: (where D: is your CD-ROM drive), and press Enter. From the files listed on the CD, select PT\_Manager\_V2.5.0.0, then select the Windows or Windows 64 version that matches your system OS type. The ProtecTIER Manager wizard will start and the Introduction window opens.

Click **Next**. The Software License Agreement window opens.

**Note:** You can print the License Agreement by clicking **Print**. If you want to read the non-IBM terms of the license agreement, click **Read Non-IBM Terms** and a window opens with the corresponding text.

3. Select **I accept both the IBM and the non IBM-terms** and click **Next**.
4. The License Agreement window opens. Select **I accept the terms of the License Agreement** and click **Next**.

The Choose Install Folder window opens (Figure 5-1).

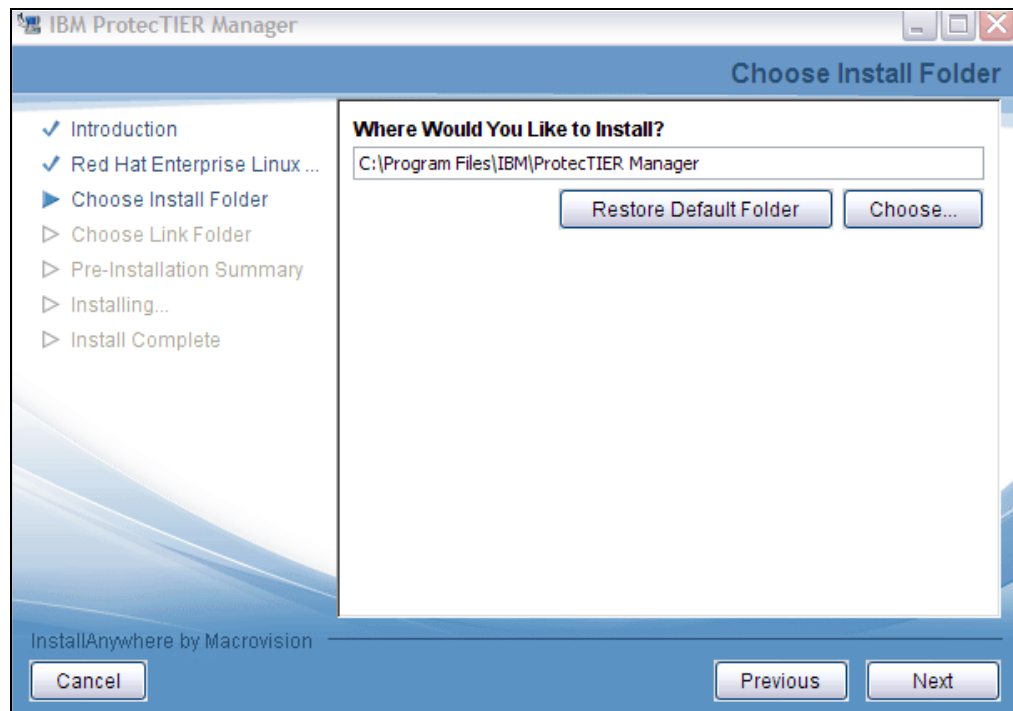


Figure 5-1 Choose Install Folder window

5. Enter the path where you want to install ProtecTIER Manager or click **Choose** to browse for a location.

**Note:** Click **Restore Default Folder** to revert to the default path.

Click **Next**. The Choose Shortcut Folder window opens (Figure 5-2).

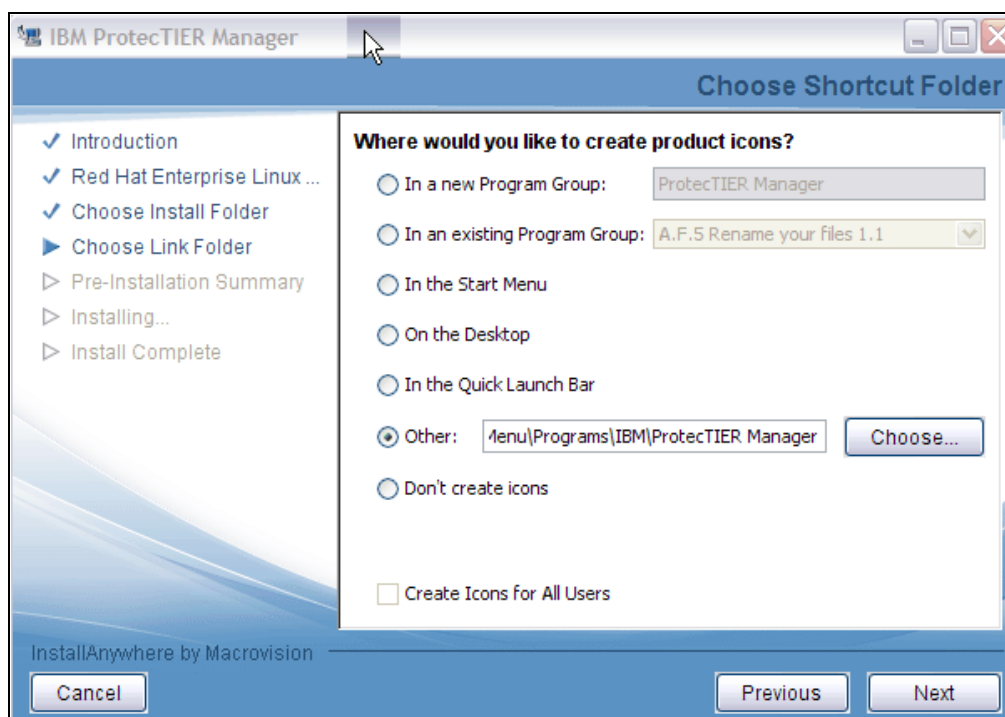


Figure 5-2 Choose Shortcut Folder window

6. Select one of the following locations for the ProtecTIER Manager shortcut:

- |                                     |   |
|-------------------------------------|---|
| <b>In a new Program Group</b>       | Creates a new program group in the Program list of the Start menu.  |
| <b>In an existing Program Group</b> | Adds the shortcut to an existing program group in the Program list of the Start menu.                         |
| <b>In the Start Menu</b>            | Creates shortcuts directly in the Start menu.   |
| <b>On the Desktop</b>               | Creates shortcuts on the desktop.   |
| <b>In the Quick Launch Bar</b>      | Creates shortcuts in the Quick Launch Bar.  |
| <b>Other</b>                        | Enables you to enter a path location for the shortcut or to browse for a location by clicking <b>Choose</b> . |
| <b>Don't create icons</b>           | No shortcuts are created.   |

You can select **Create Icons for All Users** to create a shortcut in the defined location for all user accounts on the workstation. In our example, we used the default **Other**.

Click **Next**. The Pre-Installation Summary window opens (Figure 5-3).

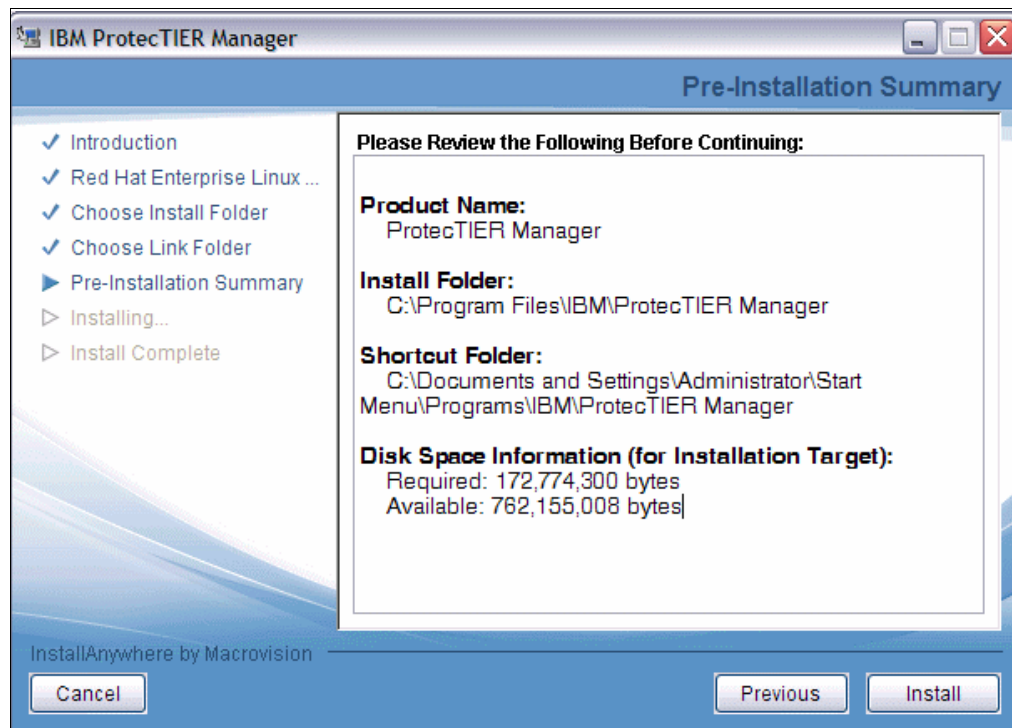


Figure 5-3 Pre-Installation Summary window

7. Click **Install**. The Installing ProtecTIER Manager window opens and ProtecTIER Manager is being installed on your computer (Figure 5-4).

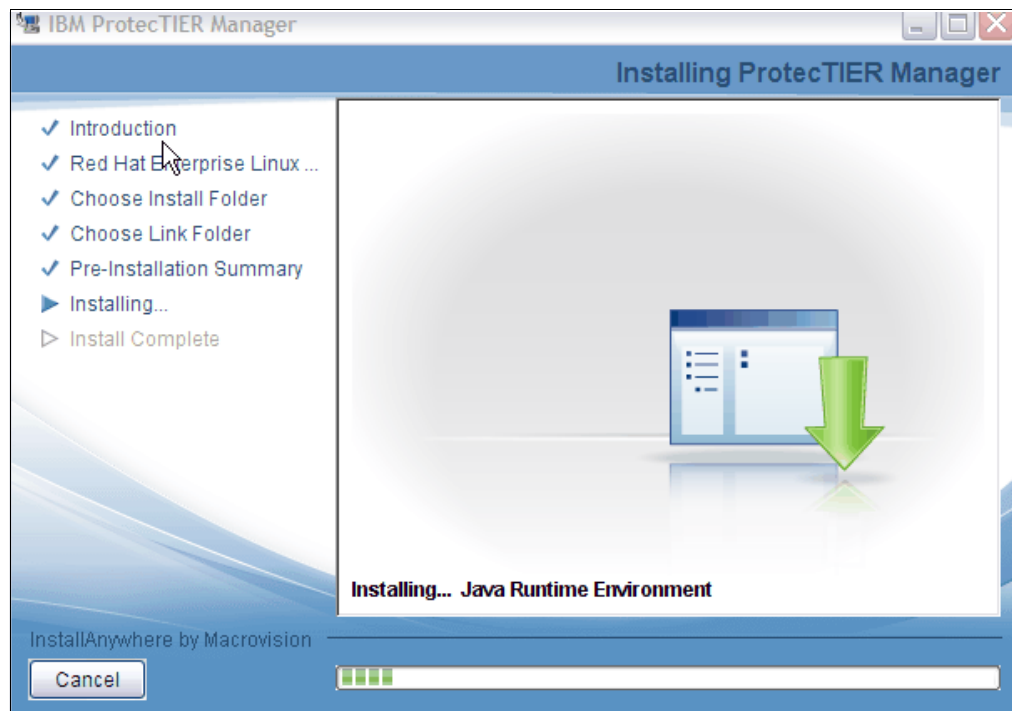


Figure 5-4 Installing ProtecTIER Manager window



When the installation is complete and ProtecTIER Manager has been successfully installed, the Install complete window opens (Figure 5-5).

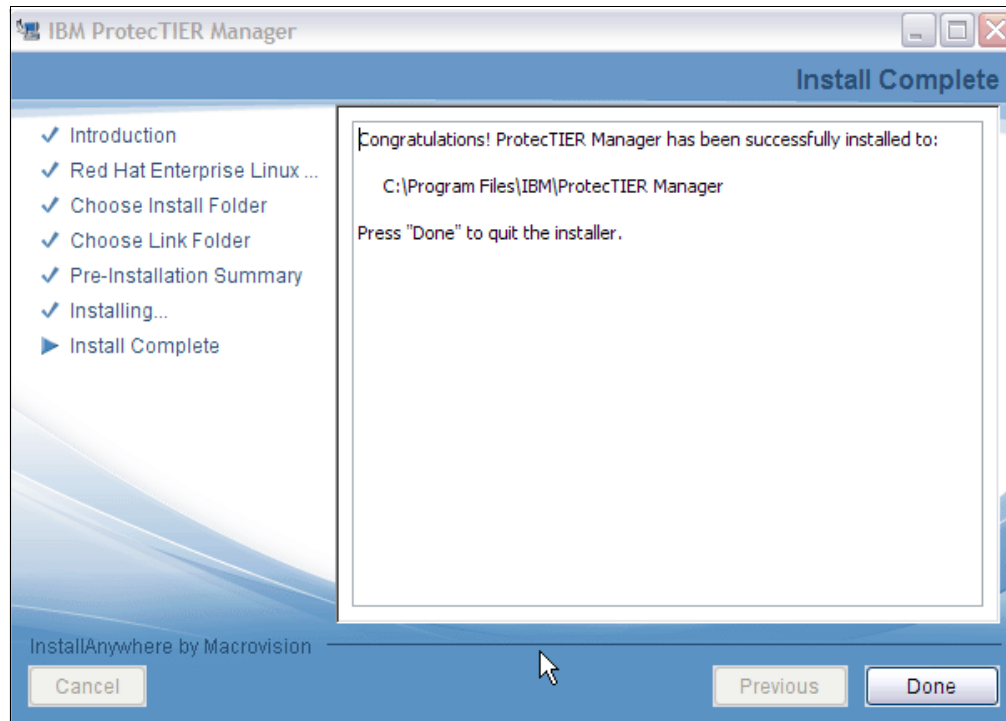


Figure 5-5 Install Complete window

8. Click **Done**. The ProtecTIER Manager wizard closes.

### 5.2.3 Installing ProtecTIER Manager on Linux

The following procedure assumes that the workstation on which ProtecTIER Manager is being installed has a Linux graphical user interface (GUI). A GUI is required for ProtecTIER Manager operation on Linux.

To install the ProtecTIER Manager on Linux, complete these steps:

1. Insert the IBM System Storage ProtecTIER Manager Enterprise Edition V2.5 CD into the CD-ROM drive of the designated ProtecTIER Manager workstation.

2. Run the ProtecTIER installer. From the Linux Desktop, double-click the CD-ROM icon (Figure 5-6).



Figure 5-6 Linux desktop

After double-clicking, three folders will be shown (Figure 5-7).

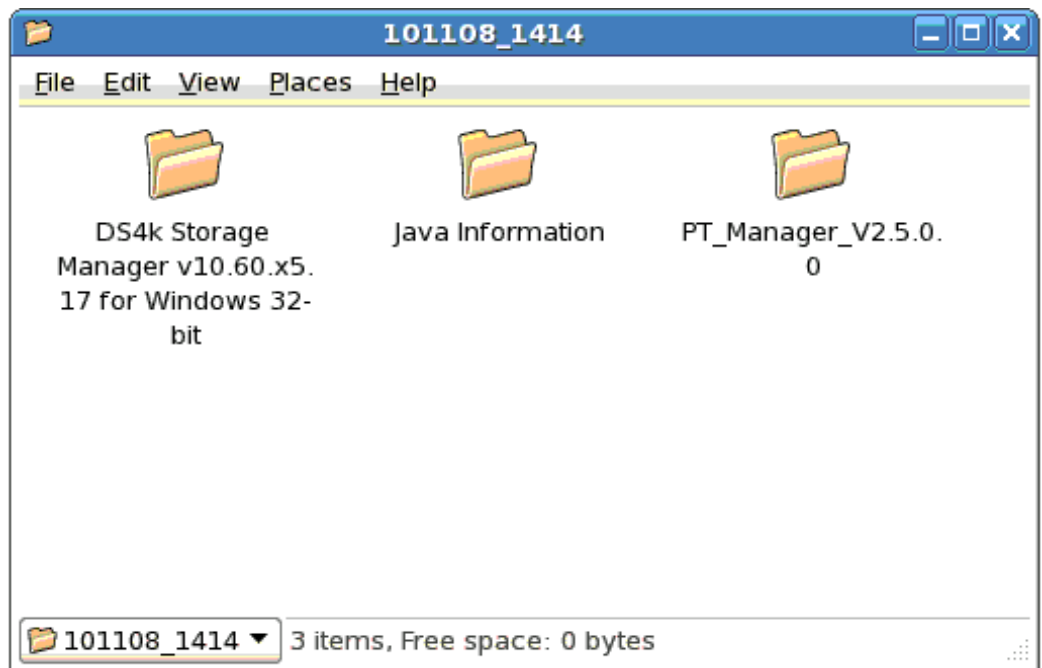


Figure 5-7 Three available folders

3. Double-click **PT\_Manager\_V2.5.0.0**. Open the folder that is suitable for your Linux server (Figure 5-8).

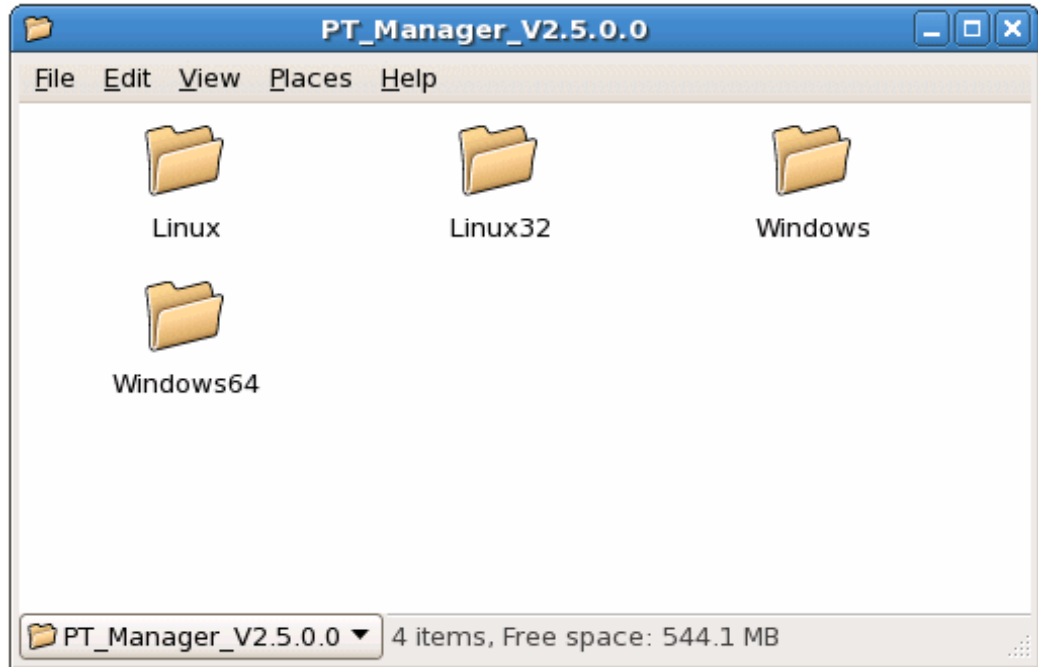


Figure 5-8 Four available folders for the installed operating system

4. From the installation folder, select the `InstallLinuxXX.bin` file (where `XX` is 32 or 64, depending on the folder that you are in) and drag the file onto your desktop. You can also drag the file to any other location. For this example, we copied it to the desktop. See Figure 5-9.

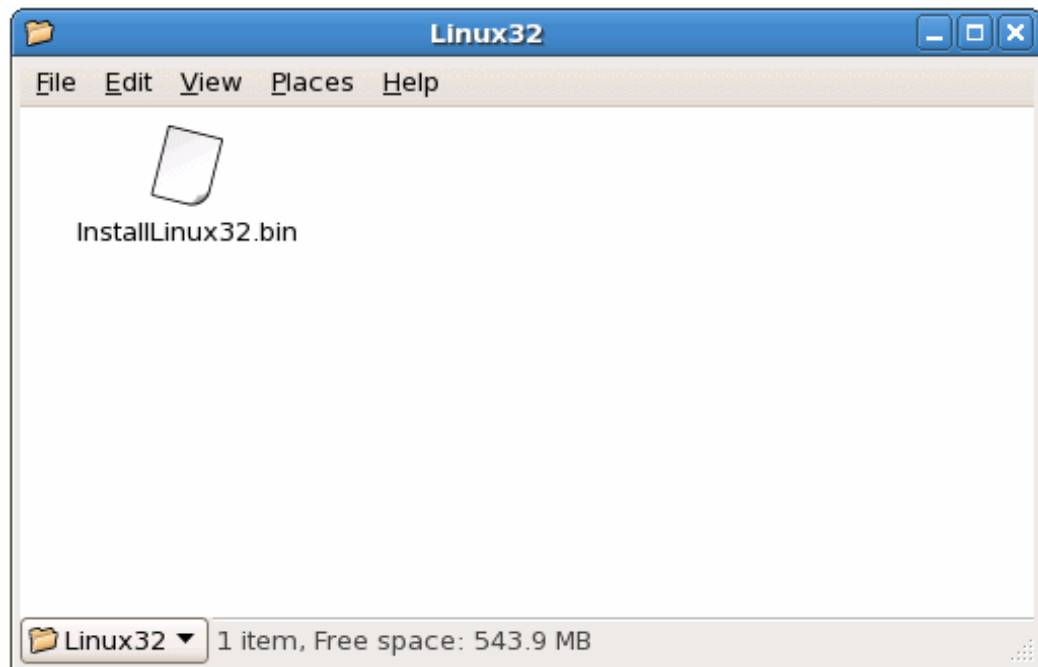


Figure 5-9 Linux content folder

5. Close any open windows.
6. Right-click any open area of the desktop, and from the menu that displays, click **Open Terminal**. The terminal windows opens.
7. At the terminal command prompt, change to the Desktop directory by running the following command:

```
cd Desktop
```

Be aware that the commands in Linux are case-sensitive. Type Desktop using a capital D.

8. From the Desktop directory, run the ProtecTIER Manager installer. Enter the `./InstallLinuxXX.bin` command (where *XX* is 64 or 32) and press Enter.

The IBM ProtecTIER Manager Wizard Introduction window opens. When you see the message Permission Denied, enter the `chmod +x InstallLinuxXX.bin` command. This changes the permissions of the file to be executable.

Figure 5-10 shows the Introduction window of the ProtecTIER installation.

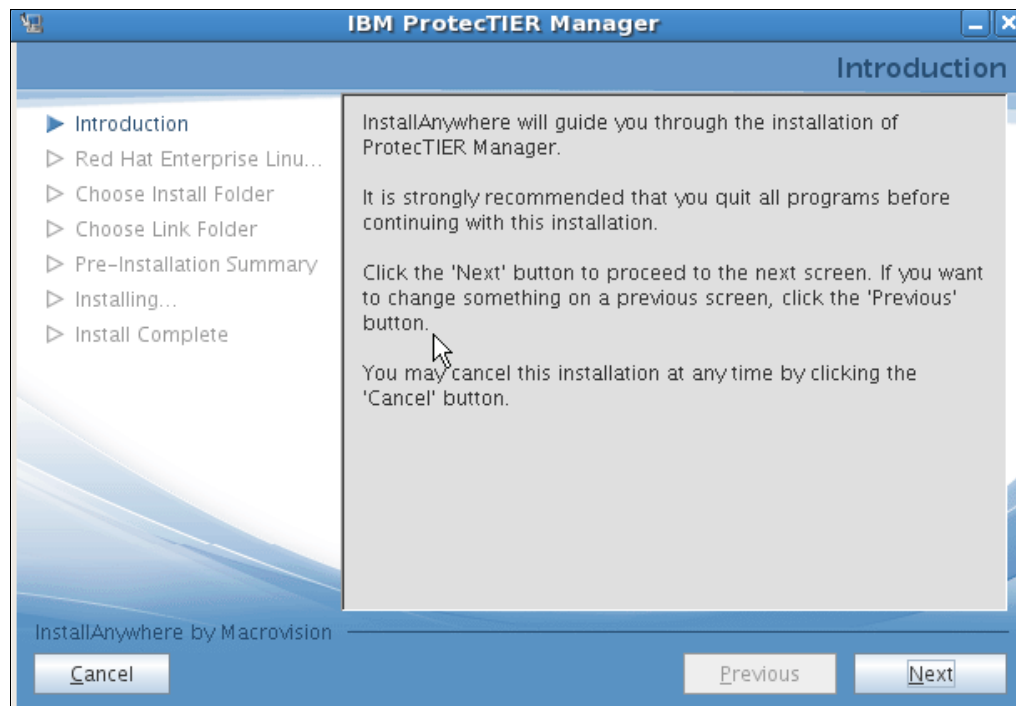


Figure 5-10 Introduction window installing PT Manager on Linux

Click **Next**. The Software License Agreement window opens (Figure 5-11).



Figure 5-11 Software License Agreement window

**Note:** You can print the license agreement by clicking **Print**. If you want to read the non-IBM terms of the license agreement, click **Read Non IBM Terms** and a window opens with the corresponding text.

9. Select **I accept both the IBM and non-IBM terms** and click **Next**. The Red Hat Linux Enterprise Linux Agreement window opens (Figure 5-12).

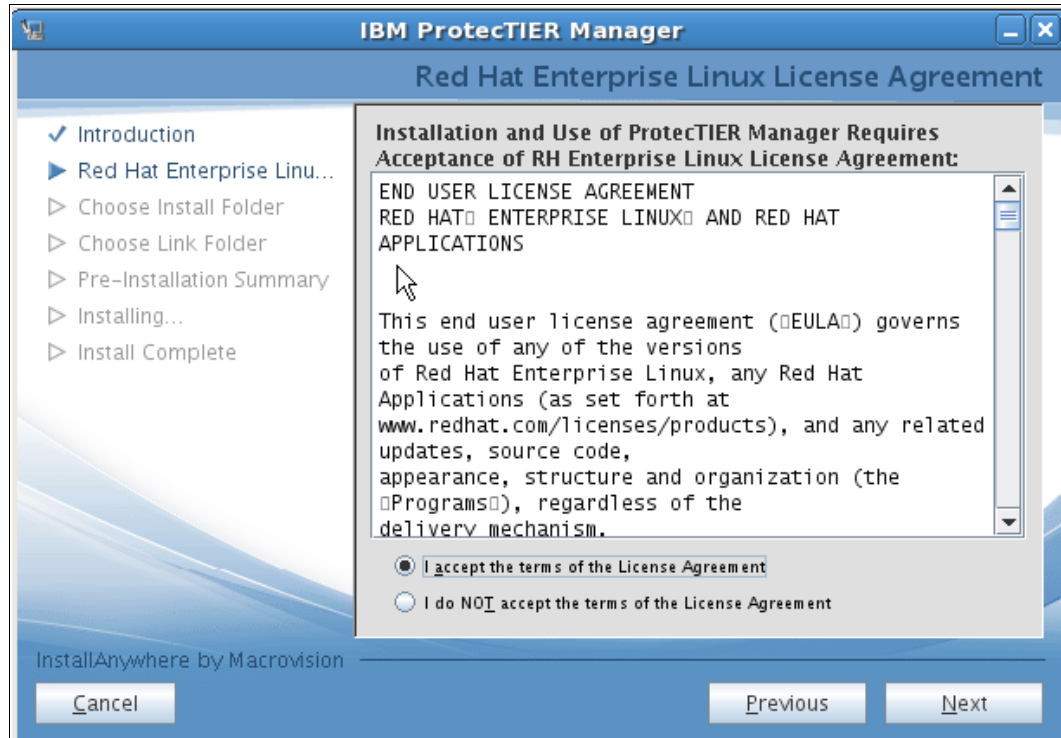


Figure 5-12 Red Hat Enterprise Linux License Agreement

10. Read the license agreement, select **I accept the terms of the License Agreement**, and click **Next**. The Choose Install Folder window opens (Figure 5-13).

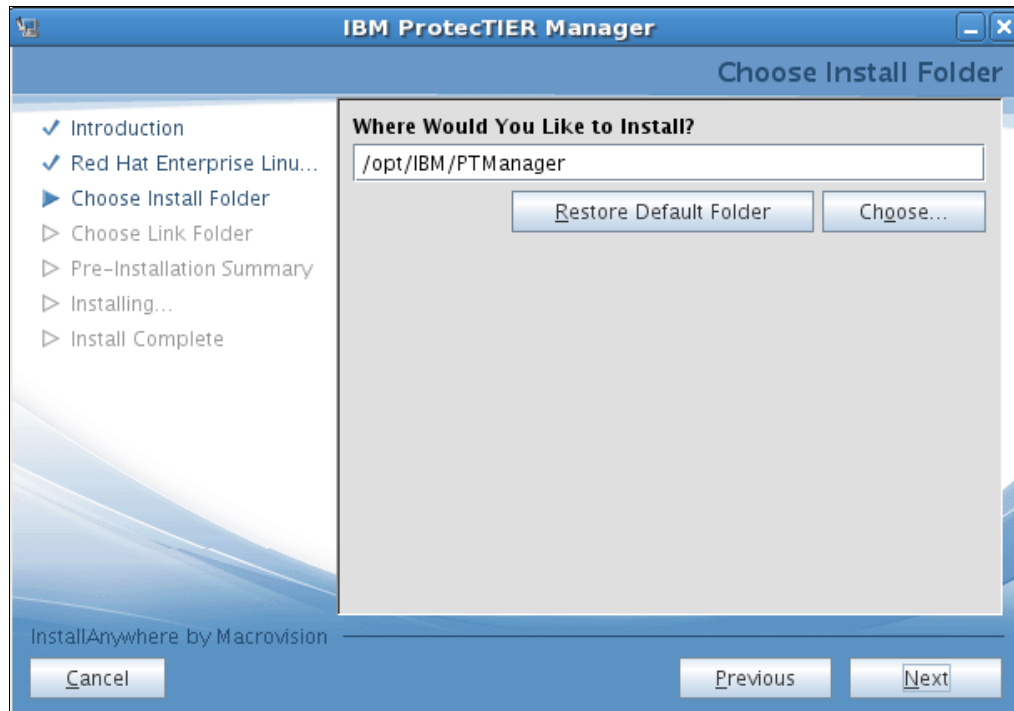


Figure 5-13 Choose Install Folder

11. Enter the path where you want to install ProtecTIER Manager or click **Choose** to browse for a location.

**Note:** Click **Restore Default Folder** to revert to the default path.

Click **Next**. The Choose Link Folder window opens (Figure 5-14).

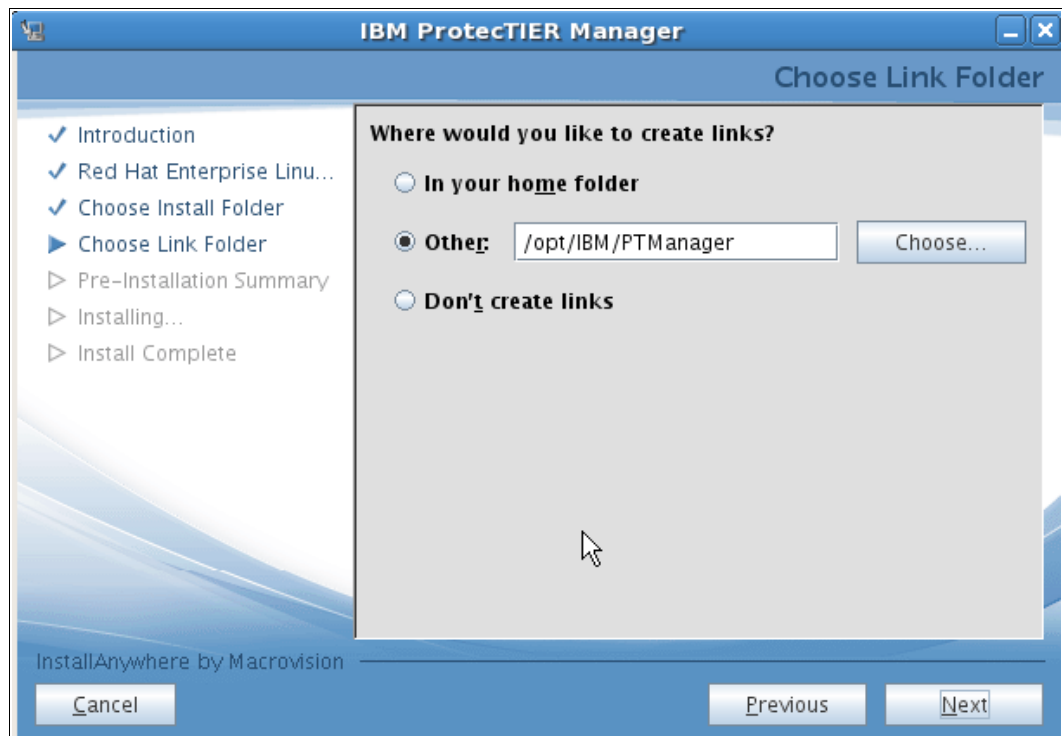


Figure 5-14 Choose Link Folder

12. Select one of the following locations for the ProtectTIER Manager shortcut:

**In your home folder**

Creates the links in your home folder.

**Other**

Enables you to enter a path location for the link or to browse for a location by clicking **Choose**.

**Don't create links**

No links are created.



Click **Next**. The Pre-Installation Summary window opens (Figure 5-15).



Figure 5-15 Pre-Installation Summary window

13. Click **Install**. The Installing ProtectTIER Manager window opens and ProtectTIER Manager will be installed on your computer (Figure 5-16).



Figure 5-16 Installing Protect Manager

When the installation is complete and ProtecTIER Manager has been successfully installed, the Install Complete window opens (Figure 5-17).



Figure 5-17 Install Complete

14. Click **Done**. The ProtecTIER Manager wizard closes and the upgrade process is complete.

The ProtecTIER Manager is installed in /opt/IBM/PTManager and the files shown in Figure 5-18 are installed in this folder.

```
File Edit View Terminal Tabs Help
[root@localhost:/opt/IBM/PTManager ] # ls -l
total 17096
-rwxrwxr-x 1 root root 347095 Oct 25 01:11 alloy.jar
drwxrwxr-x 2 root root 4096 Nov 16 04:06 bin
-rwxrwxr-x 1 root root 165191 Oct 25 01:11 binding-2.0.0.jar
-rwxrwxr-x 1 root root 137438 Oct 25 01:11 branding.jar
-rwxrwxr-x 1 root root 1000 Oct 25 01:11 cli_stat_profile_backup_stats_report.xml
-rwxrwxr-x 1 root root 934 Oct 25 01:11 cli_stat_profile_replication_stats_report.xml
-rwxrwxr-x 1 root root 86268 Oct 25 01:11 commons-cli2-1.0b.jar
-rwxrwxr-x 1 root root 821300 Oct 25 01:11 ibm-commons.jar
lrwxrwxrwx 1 root root 37 Nov 16 04:06 IBM_ProtectTIER_Manager -> /opt/IBM/PTManager/P
rotectTIER_Manager
-rwxrwxr-x 1 root root 3297074 Oct 25 01:11 ibm-pt-manager.jar
-rwxrwxr-x 1 root root 175909 Oct 25 01:11 ibm-shared.jar
-rwxrwxr-x 1 root root 307510 Oct 25 01:11 jcommon-1.0.5.jar
-rwxrwxr-x 1 root root 1137683 Oct 25 01:11 jfreechart-1.0.2.jar
-rwxrwxr-x 1 root root 1904171 Oct 25 01:11 jmf-2.1.1e.jar
drwxr-xr-x 4 525 113 4096 Jun 2 2009 jre
-rwxrwxr-x 1 root root 3195464 Oct 25 01:11 jviews-diagrammer.jar
-rwxrwxr-x 1 root root 5280013 Oct 25 01:11 jviews-framework-lib.jar
-rw-rw-r-- 1 root root 41405 Nov 16 04:06 lax.jar
drwxrwxr-x 2 root root 4096 Nov 16 04:06 license
drwxrwxr-x 3 root root 4096 Nov 16 04:06 media
-rwxrwxr-x 1 root root 4713 Oct 25 01:11 packer-2.0.jar
-rwxr-xr-x 1 root root 48037 Nov 16 04:06 ProtecTIER_Manager
-rwxrwxr-x 1 root root 3630 Nov 16 04:06 ProtecTIER_Manager.lax
-rwxr-xr-x 1 root root 48037 Nov 16 04:06 ptcli
-rwxrwxr-x 1 root root 3598 Nov 16 04:06 ptcli.lax
-rwxrwxr-x 1 root root 252914 Oct 25 01:11 trilead-ssh2-build213.jar
lrwxrwxrwx 1 root root 76 Nov 16 04:06 Uninstall__IBM_ProtectTIER_Manager -> /opt/IBM/
PTManager/Uninstall_ProtectTIER_Manager/Uninstall_ProtectTIER_Manager
drwxrwxr-x 2 root root 4096 Nov 16 04:07 Uninstall_ProtectTIER_Manager
-rwxrwxr-x 1 root root 114517 Oct 25 01:11 xmlrpc-1.2-b1.jar
[root@localhost:/opt/IBM/PTManager ] #
```

Figure 5-18 Linux directory view

15. To start the ProtecTIER Manager, run the `./ProtectTIER_Manager` command and the GUI of the ProtecTIER Manager starts. You can start to configure your ProtecTIER system.
16. You can also create a shortcut on your desktop. Right-click your desktop and select **Create Launcher** and a window opens (Figure 5-19).

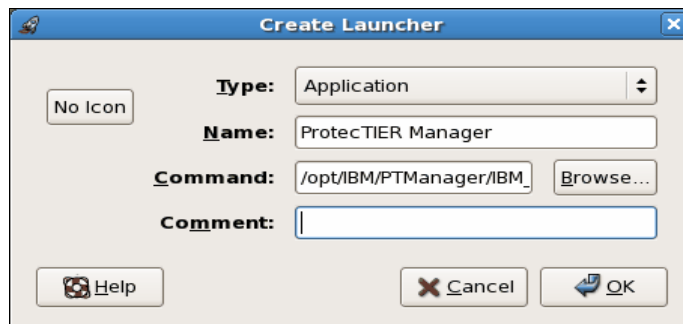


Figure 5-19 Create launcher

17. Select **Application** for the type, in the Name field type a name of your choice, and in the Command field enter the path to the executable ProtecTIER file:

/opt/IBM/PTManager/ProtectTIER\_Manager

## 5.3 Getting started

After ProtecTIER is installed and the ProtecTIER Manager GUI workstation is in place, complete the ProtecTIER system setup.

Perform the following tasks in this order to bring your ProtecTIER system online:

- ▶ Add a node to ProtecTIER Manager. For more information, see 5.3.1, “Adding nodes to ProtecTIER Manager” on page 190
- ▶ Plan a repository for the ProtecTIER (TS7650G model only). For more information see Chapter 3, “Planning for deduplication and replication” on page 53
- ▶ Create a set of file systems for the repository (TS7650G model only). For more information see 5.5.1, “Creating file systems” on page 209.
- ▶ Create the repository (TS7650G model only). For more information see 5.5.3, “Creating the repository” on page 216.

Creating the repository on the node results in a functional one-node ProtecTIER or a two-node cluster system. If you are using a one-node system, create a library on the one node. For more information, see 5.6, “Setting up the virtual library and cartridges” on page 230.

### 5.3.1 Adding nodes to ProtecTIER Manager

Adding nodes registers the node’s IP address and port number with the instance of ProtecTIER Manager at your workstation. Similarly, removing nodes removes the node’s registration from ProtecTIER Manager at that workstation

Figure 5-20 shows the ProtecTIER window after starting ProtecTIER Manager without any nodes added.

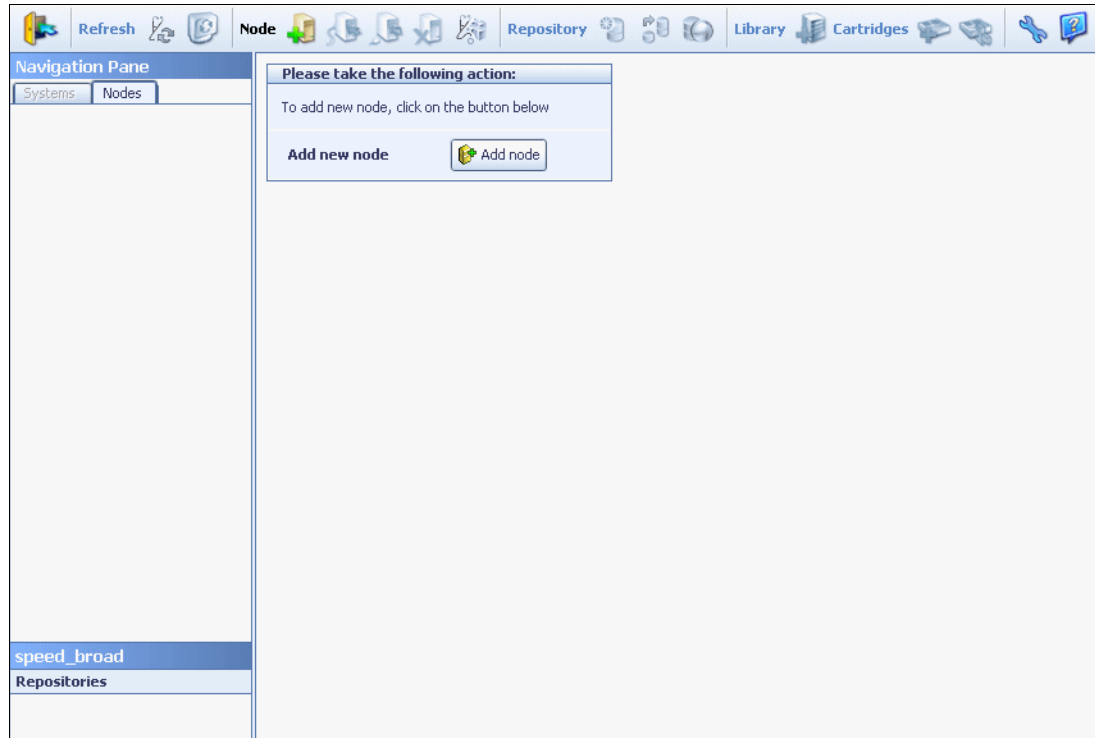


Figure 5-20 ProtecTIER Manager default window

Complete these steps:

1. Click **Add node**. The Add node window opens (Figure 5-21), prompting you for the IP address and port number of the node that you want to add.

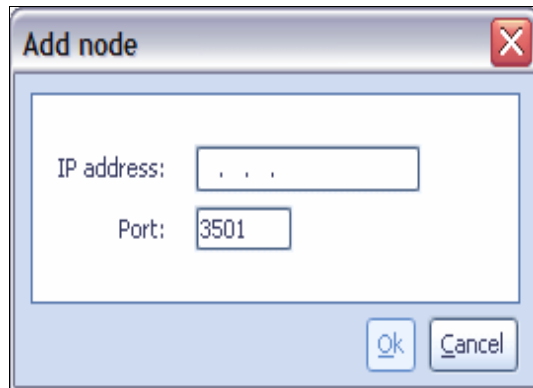


Figure 5-21 ProtecTIER Add node window

2. Enter the IP address of the node and click **OK**. The node is displayed in the Nodes pane, and the Login button is displayed in the View pane.

**Note:** Do not change the port number of the node unless directed by IBM support.

3. Click **Login**. You are prompted for your user name and password.

4. Enter your user name and password and click **OK**. ProtecTIER Manager displays the information for that node. If the node has a repository, the node's cluster is displayed in the PT Systems tab of the Navigation pane.

**Note:** If that two-node cluster already contains a second node, the second node is also displayed in the Nodes pane.

For more information about logging in and out, refer to 5.3.2, “Logging in and out” on page 192.

For more information about managing nodes in ProtecTIER Manager, refer to 10.1, “Managing nodes in ProtecTIER Manager” on page 474.

### 5.3.2 Logging in and out

Log in to each ProtecTIER system that you want to manage using ProtecTIER Manager:

1. On the navigation pane, select a system that you want to log in to or select a node on the nodes pane of the system in which you want to log in (Figure 5-22).

**Note:** Even if you selected a single node to log in to, you are logged in to the clustered system after a successful authentication check.

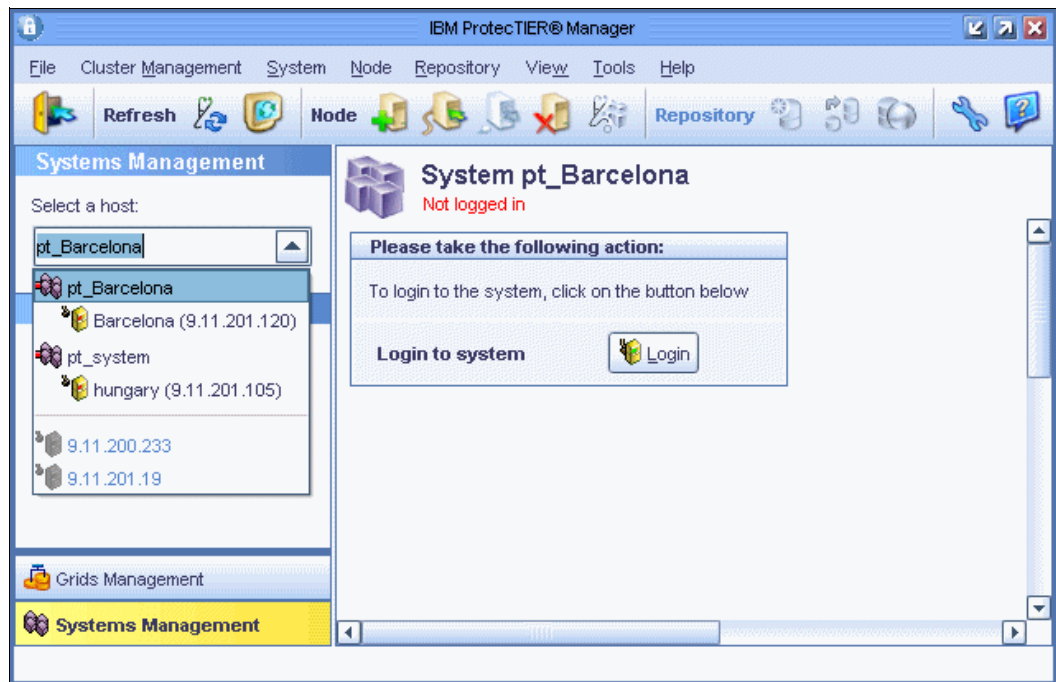


Figure 5-22 ProtecTIER Manager Start window

2. Click **Login**. The Login window opens (Figure 5-23).

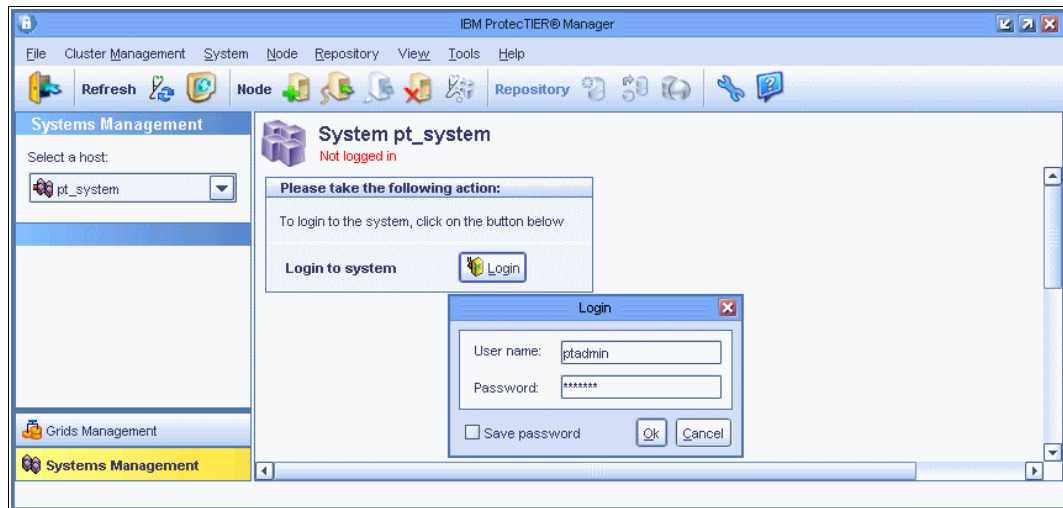


Figure 5-23 ProtecTIER Manager Login window

3. Enter your user name and password.
4. Click **OK**. The Login window closes and you are logged into ProtecTIER Manager.

ProtecTIER Manager has default user accounts corresponding to three user permission levels:

- ▶ Administrator
- ▶ Operator
- ▶ Monitor

For more information, see Chapter 12, “Monitoring and reporting of the IBM System Storage TS7600 with ProtecTIER” on page 619. Table 5-1 lists the default user name and password for each of these accounts.

Table 5-1 Default user names and passwords

Permission level	Default user name	Default password
Administrator	ptadmin	ptadmin
Operator	ptoper	ptoper
Monitor	ptuser	ptuser

**Note:** Consider changing or replacing these default user accounts. For more information, see 12.1, “ProtecTIER Manager user management” on page 620.

Only one administrator can be logged into a ProtecTIER system at a time. Consider logging out at the end of each session by clicking **Logout**.

If you log in with administrator level permissions while another administrator is already logged in, a message box opens with the message shown in Figure 5-24.

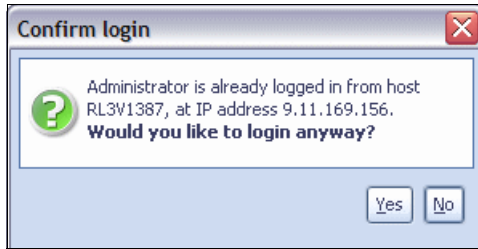


Figure 5-24 ProtecTIER Manager Confirm login window

Click **Yes** to force the other administrator to log out.

To discover which node is currently being used by ProtecTIER Manager to monitor the two-node cluster (using the GUI proxy), in the navigation pane select the **Systems** tab and click the desired system. The ProtecTIER systems overview window opens with a section providing information about the cluster members (Figure 5-25).

Cluster members						Options ▾
IP address	DNS	GUI proxy	Status	Management service	Applications	
10.0.200.233	italy	✓	● Ok	● Online	● VT	
10.0.200.19	naples		● Ok	● Online	● VT	

Figure 5-25 ProtecTIER Manager Cluster members window

**Note:** If the responsible server for the GUI proxy has been stopped, the second server takes over the responsibility automatically and retains it even if the first server is started again.

### 5.3.3 TS7610 start message

TS7610 SM1 is a small and medium model that a customer can install without assistance from an IBM SSR or Business Partner. When it is shipped to the customer, it will come configured with the repository and VTL (for VTL model only). The customer will be able to start using the ProtecTIER system with minimal setup.

When you first login to the ProtecTIER Manager, you will see a message display, as shown in Figure 5-26. Click **OK** to close the message.

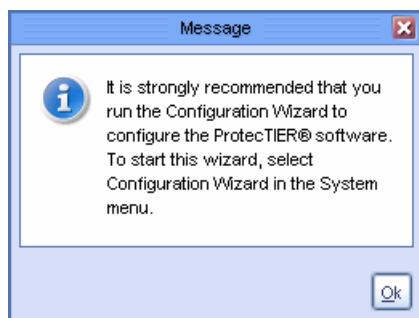


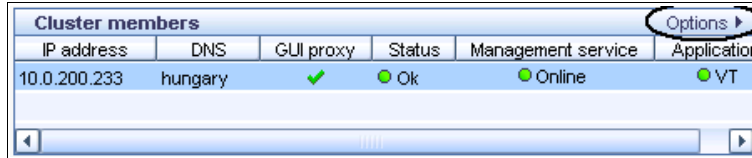
Figure 5-26 Message on TS7610 first login



Refer to 5.4, “Initial setup for the TS7610 model” on page 199” about how to use the Configuration Wizard to set up the TS7610.

### 5.3.4 Saving and printing data

You can save or print the data in the ProtecTIER Manager. Click **Options** at the top right hand corner of any window where it is available (Figure 5-27). The Save and Print options appear. Select **Save** or **Print** to save the information or print it using the standard saving or printing procedures for your operating system. For example, it will save as a Comma Separated Value (CSV) file on a Windows workstation.



IP address	DNS	GUI proxy	Status	Management service	Application
10.0.200.233	hungary	✓	● Ok	● Online	● VT

Figure 5-27 ProtecTIER Manager save or print data

### 5.3.5 Refreshing ProtecTIER Manager

Many windows and panes of ProtecTIER Manager automatically refresh to reflect the most accurate information. Some, however, must be refreshed manually. In addition, because you can set up multiple ProtecTIER Manager workstations on the same system, changes can be made on another workstation that are not automatically reflected on your workstation. Periodically refresh the Navigation pane and the View pane of ProtecTIER Manager (Figure 5-28).

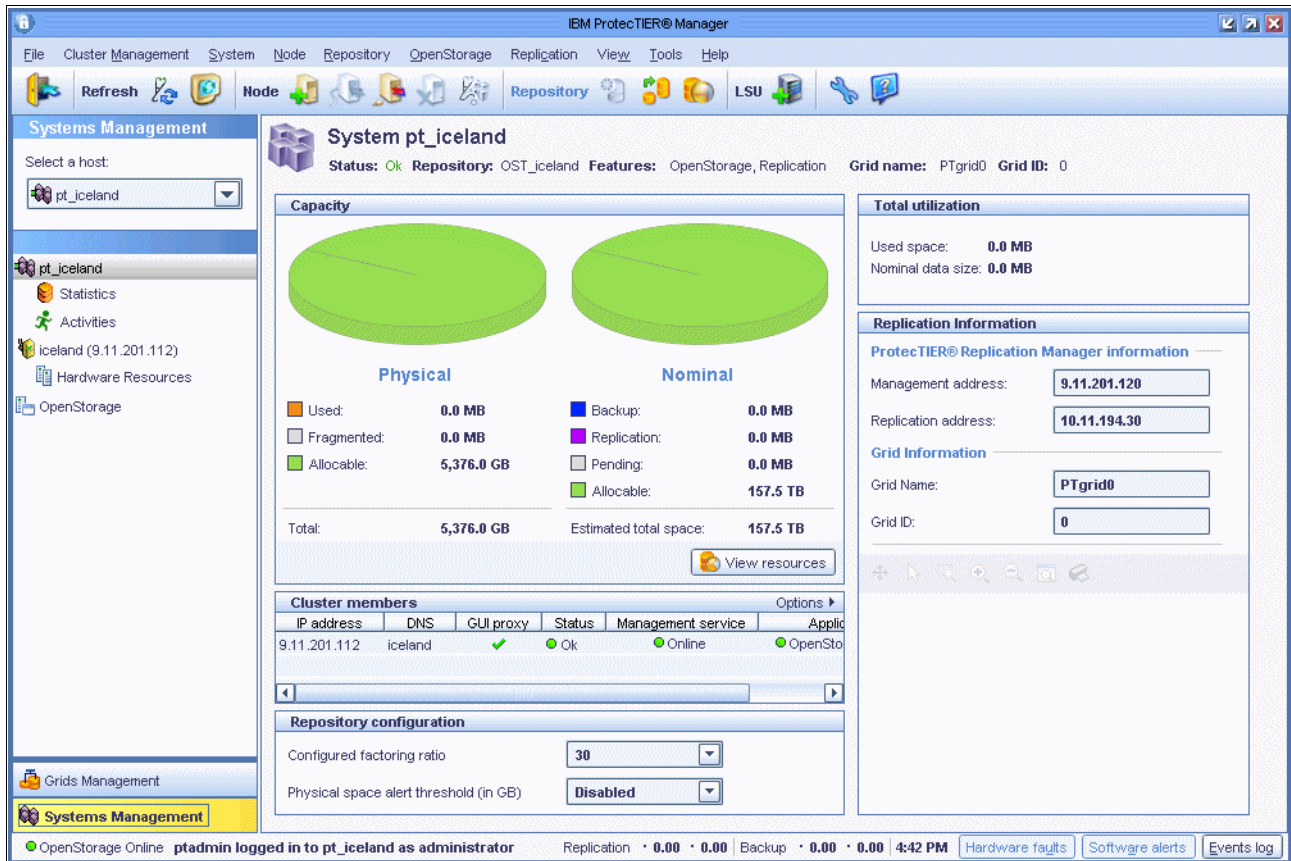


Figure 5-28 ProtecTIER Manager Navigation and View pane

Complete these steps:

1. Click **Refresh navigation pane** to refresh the Navigation pane (Figure 5-29).

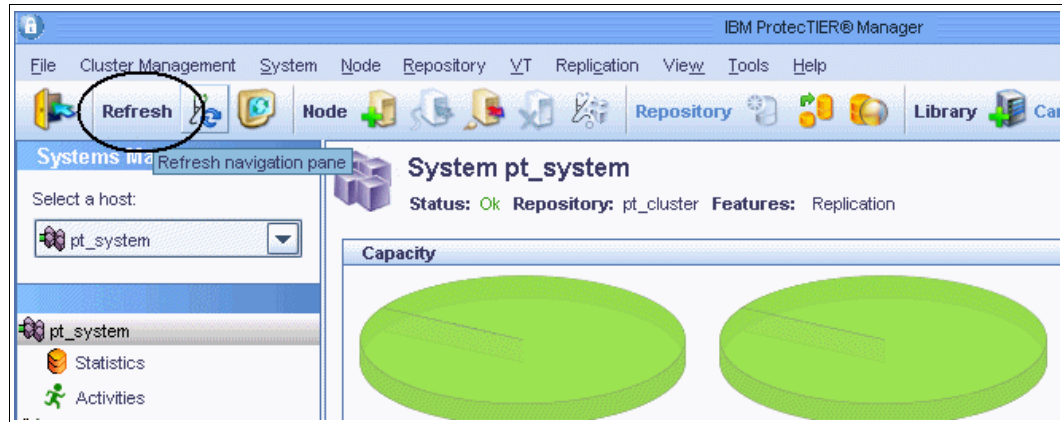


Figure 5-29 Refresh navigation pane

2. Click **Refresh current view** to refresh the View pane (Figure 5-30).

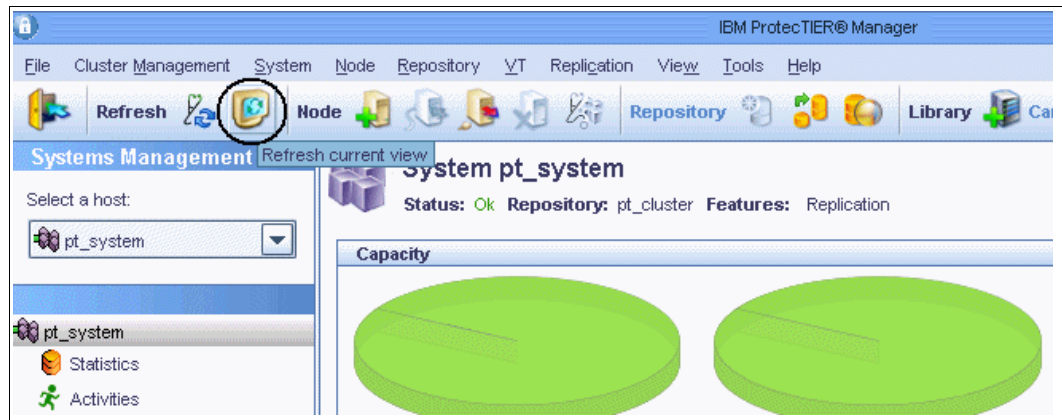


Figure 5-30 Refresh current view

### 5.3.6 Renaming the system

You can rename the ProtecTIER system according to your own naming convention by completing the following steps:

1. In the **Systems Management** pane, select the host you want to rename (Figure 5-31).




Figure 5-31 Rename system

2. The Rename system window opens (Figure 5-32). Enter the new System name according to your naming convention for easy identification later. Click **Ok**.

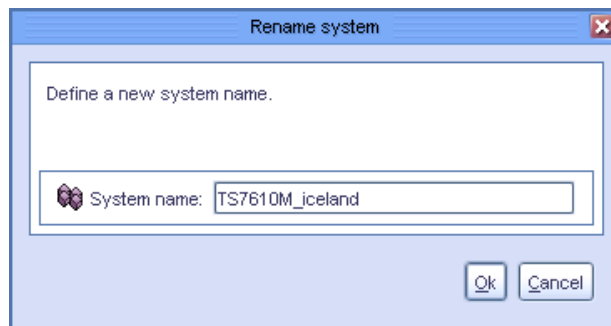


Figure 5-32 Rename system window

- The new system name will be immediately be reflected in the navigation pane (Figure 5-33). In this example, this system is a TS7610 Medium model with the host name of iceland, so we use the name TS7610M\_iceland to identify it.

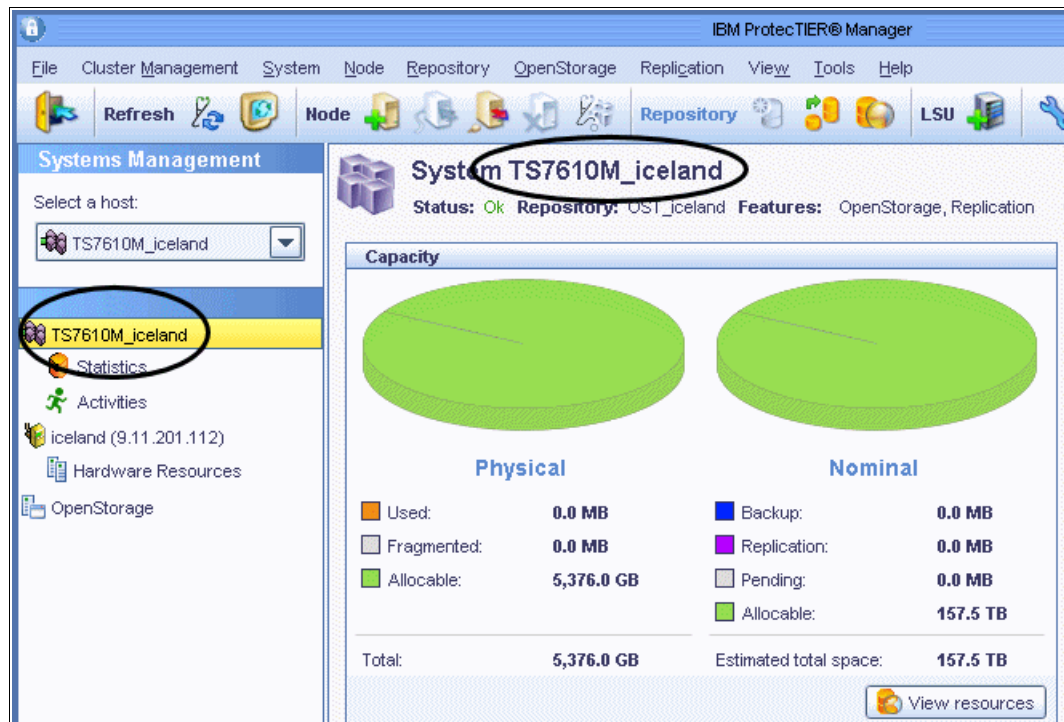


Figure 5-33 New system name

Now that you are getting familiar with the user interface of the ProtecTIER Manager, we will proceed to configure the system. For TS7610 models, go to 5.4, “Initial setup for the TS7610 model” on page 199. For TS7650 or TS7650G, go to 5.5, “Creating a ProtecTIER repository for TS7650G” on page 208.

## 5.4 Initial setup for the TS7610 model

In this section, you use the ProtecTIER Manager Configuration Wizard to enable and set up email or SNMP trap hardware fault notification, and enable IBM Call Home.

Typically, hardware faults are reported through the PT Manager GUI. However, certain conditions will prevent the GUI from communicating with the server, rendering the GUI unable to provide hardware fault information for the failed component. By enabling email or SNMP trap notifications, you have greater assurance that hardware faults will not go undetected.

You should complete the wizard the first time you run it. However, if you are unable to provide some of the information requested, simply leave the fields blank. The message shown in Figure 5-26 on page 194 displays each time you log into ProtecTIER Manager, as a reminder to provide the remaining details. You can complete the wizard (or change information previously provided) at any time, by performing the steps shown below. This configuration wizard is the same for both a TS7610 VTL system and a OST system.

## 5.4.1 TS7610 configuration wizard setup

If you have not set up the ProtecTIER Manager, refer to 5.2, “Installing ProtecTIER Manager” on page 174 to learn how to install it, and 5.3, “Getting started” on page 190 to learn how to start using it.

Complete the following the steps to set up the TS7610 using the configuration wizard:

1. Log in to the ProtecTIER Manager GUI. Make sure that the Systems Management pane is selected (which will be highlighted in yellow) at the bottom left corner of the GUI, as shown in Figure 5-34.
2. Go to the **System** menu and select **Configuration wizard**, as shown in Figure 5-34.

**Note:** You should configure the SNMP traps and email notification to ensure that you receive the alerts in time when there is any error in the TS7610.

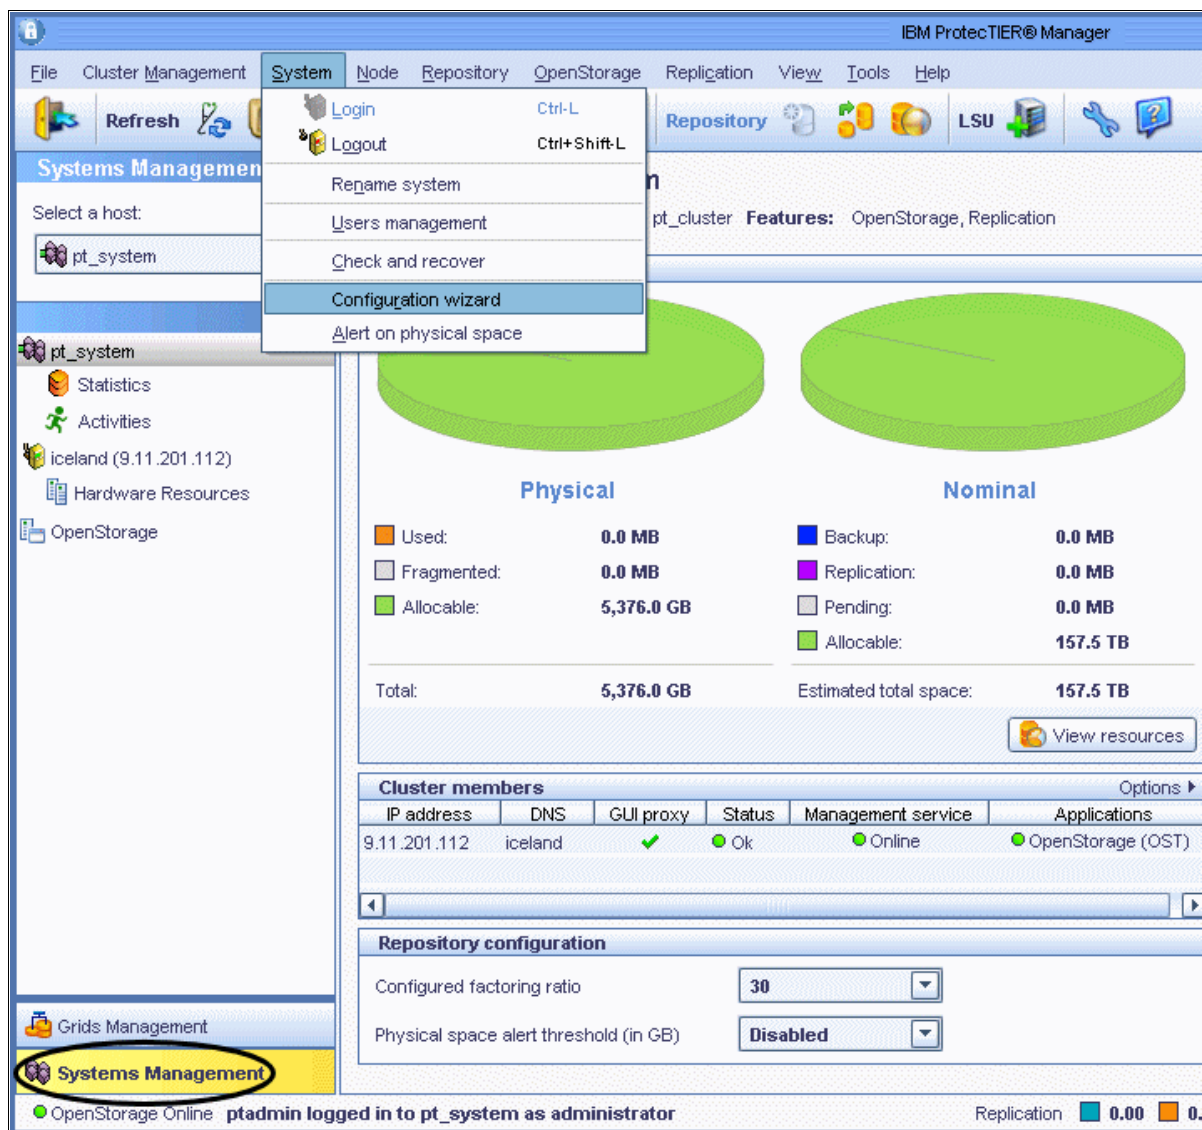


Figure 5-34 Getting to the Configuration Wizard

3. The Welcome window opens. Click **Next**. See Figure 5-35.

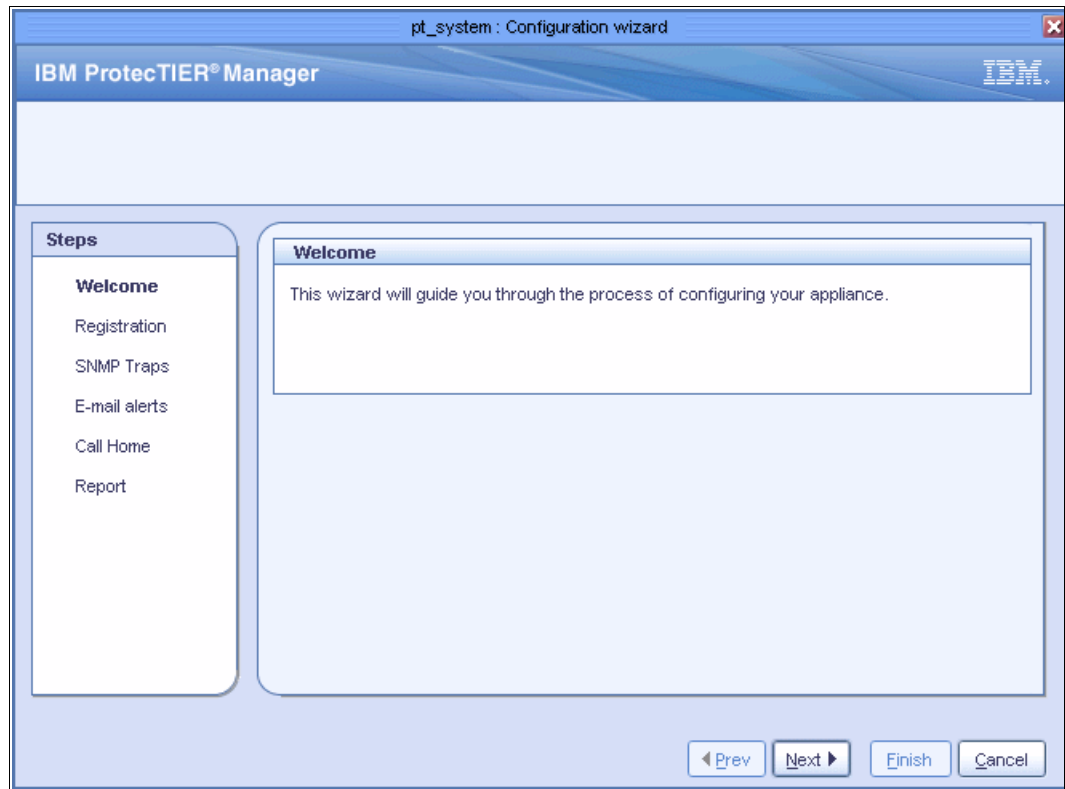


Figure 5-35 TS7610 Configuration Wizard Welcome window

- The Registration information window opens (Figure 5-36). In this window, enter the Company and System Administration information and click **Next**.

**Note:** If you plan to use Call Home, you must enter your country code and IBM customer number where indicated. Call Home will not function properly without this information.

The screenshot shows the 'Registration information' window in the IBM ProtecTIER Manager Configuration wizard. The window title is 'pt\_system : Configuration wizard'. The main title is 'IBM ProtecTIER Manager'. On the left, there is a 'Steps' sidebar with 'Welcome' checked and 'Registration' selected. The main area is divided into two sections: 'Company information' and 'System administration information'. The 'Company information' section includes fields for Name (ABC), Street, City (TUCSON), State (ARIZONA), Postal code (123456), and Customer number (1234567). The 'System administration information' section includes fields for Name (Dave), Phone (1234567), and E-mail (abcd@ibm.com). At the bottom right, there are buttons for 'Prev', 'Next', 'Finish', and 'Cancel'.

Figure 5-36 TS7610 Configuration wizard registration information window

- The SNMP Traps window appears, with the **Enable SNMP traps** check box cleared and the input fields disabled.
  - If you do not want to receive SNMP notifications, leave the **Enable SNMP Traps** check box clear, click **Next**, and then go on to step 6 on page 205.
  - If you want to receive SNMP notifications, complete the following steps:



- i. Select the **Enable SNMP Traps** check box. The SNMP trap input fields become active, as shown in Figure 5-37.

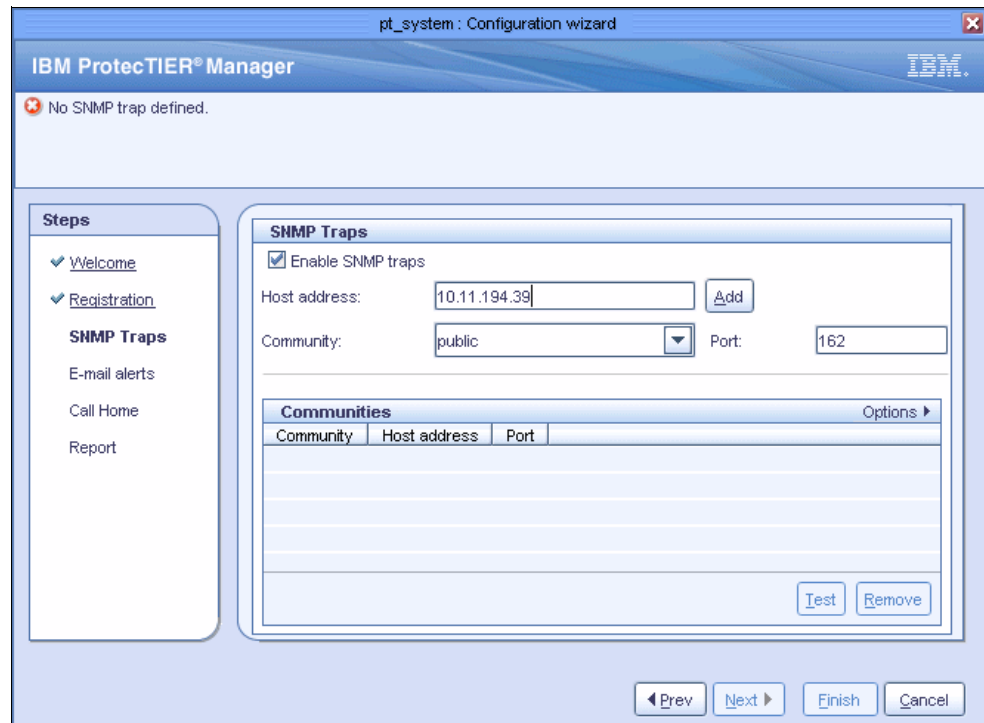


Figure 5-37 Configuration wizard: SNMP Traps window

- ii. In the **Host address field**, enter the SNMP server's host name or IP address.
- iii. In the **Community** drop-down list, select **public** or **private**. The default port used is 162.

- iv. Click **Add**. The wizard will test the SNMP connection with the server and add a line under the communities section (Figure 5-38).

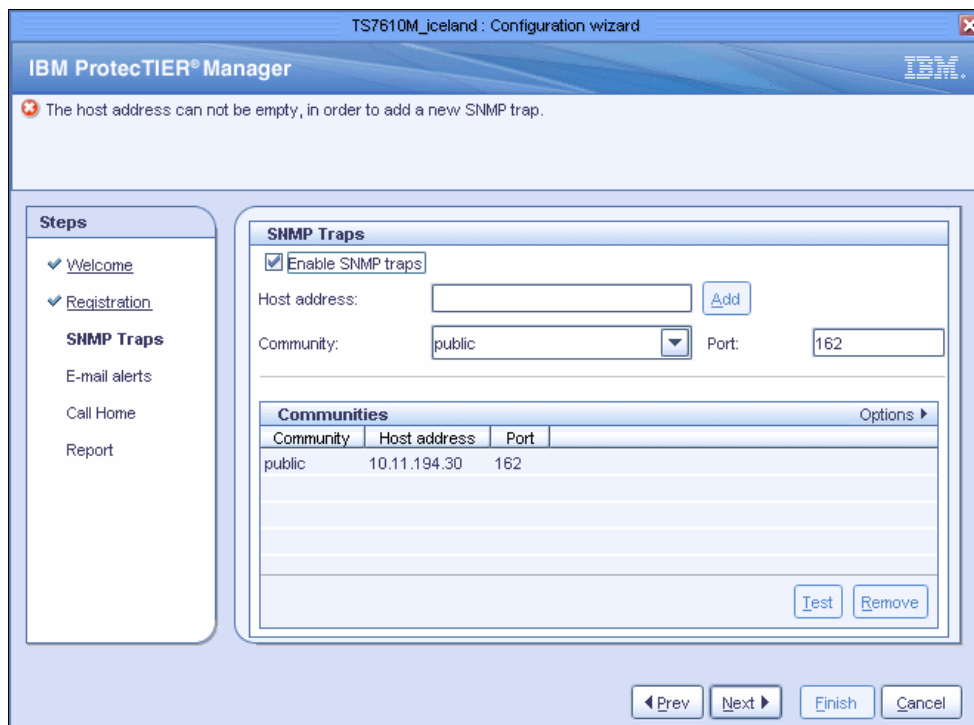


Figure 5-38 Configuration wizard: SNMP Traps window

**Note:** In a Public Community, access to the SNMP trap agent is read-only. This allows the underlying network management system (NMS) to retrieve data from the associated network elements. Public Community is the default.

In a Private Community, access to the SNMP trap agent is read-write, which allows the NMS to change the values set on the associated network elements.

(Optional) Even after the configuration has completed, if you want to manually send a test SNMP trap to the specified SNMP host server for troubleshooting purpose, you can always come back to this configuration wizard SNMP Traps windows and click **Test**. The system will send a test SNMP trap report to the SNMP server. If the test is successful, a window opens with the message “SNMP trap test succeeded”, as shown in Figure 5-39.




Figure 5-39 SNMP trap test succeeded

Click **Next**.

**Note:** To receive and display the SNMP trap reports, you must have SNMP trap receiver software installed on the SNMP trap server. For instructions about installing, configuring, and using your SNMP trap receiver software, refer to the manufacturer's documentation that came with the product.

To enable the trap receiver to translate and display the SNMP trap data, you need to specify the file name and location of the management information base (MIB) files for the SNMP trap receiver. On the 3959 SM1 server, the MIB file names are DILIGENT-MIB.TXT and IBM-TS-RAS-MIB.txt. The files are located in the /usr/share/snmp/mibs/ directory on the server.

6. After the information for the SNMP traps has been configured, the Configure outgoing E-mail server window opens, with the Enable E-mail alerts check box cleared and the input fields disabled.

In the Configure outgoing E-mail server window, you can specify to have notification of hardware faults sent to one or more email addresses:

If you do not want to receive email notifications, leave the **Enable E-mail Alerts** check box clear, click **Next**, skip this step, and go to step 7 on page 206.

If you want to receive email notifications, complete the following steps:

- a. Select the **Enable E-mail Alerts** check box. The E-mail alerts input fields become active, as shown in Figure 5-40.

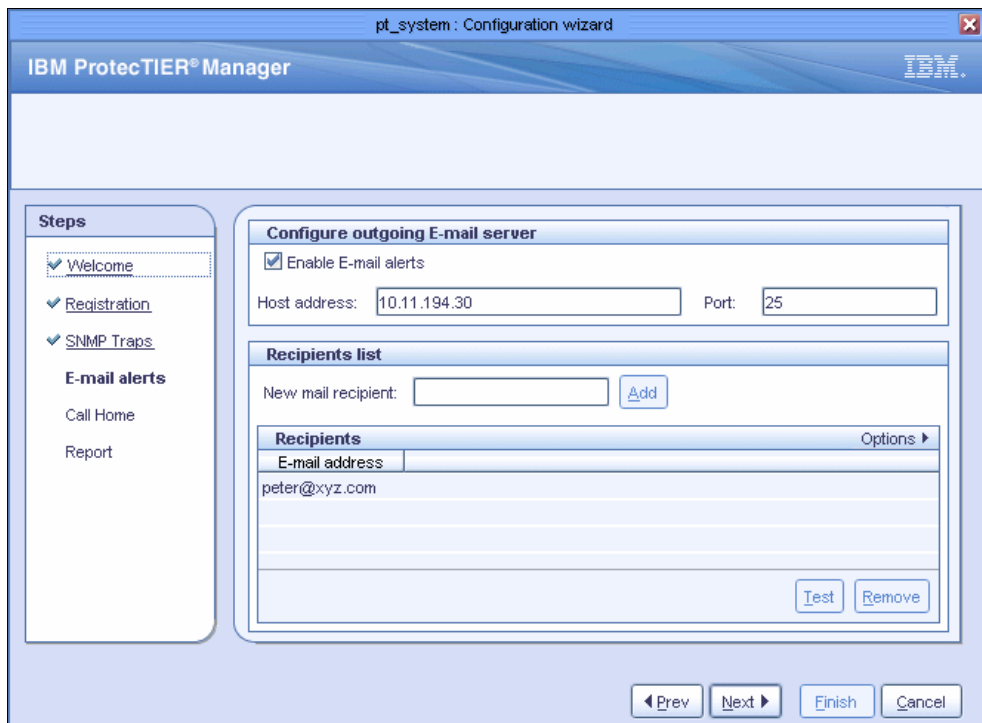


Figure 5-40 Configure outgoing E-mail server window

- b. Enter the Host address for the Outgoing E-mail Server.
- c. Leave the Port set to the default value of 25.
- d. Enter an alert recipient's email address in the New mail recipient field. Use only lowercase characters when entering email addresses. Click **Add**.

- e. Repeat the above step for each additional email alert recipient you want to specify.
  - f. To manually send a test message to a recipient, select the recipient from the list and click **Test**. The system sends a test message to the selected recipient's email address.
  - g. Click **Next**.
7. (Optional) The Call Home window opens, with the Enable Call Home check box cleared. If your environment allows the Call Home feature, you should enable this feature.

If you do not want to use Call Home, leave the Call Home check box clear, click **Next**, and go to step 8 on page 208.

If you want to enable the Call Home feature, select the Enable Call Home check box as shown in Figure 5-41. You must also make sure that you entered your Customer Number and Country Code on the Configuration Wizard's Registration window, as shown in Figure 5-36 on page 202. Click **Test** to verify that Call Home is configured correctly.

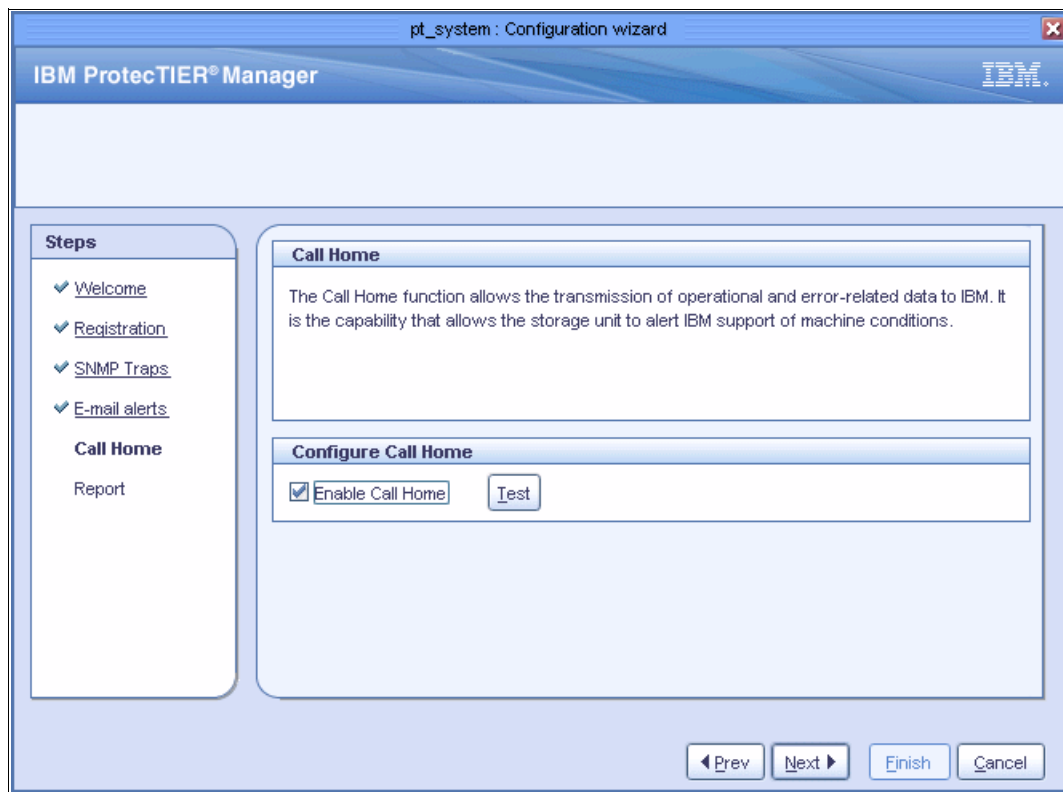


Figure 5-41 Call Home window

If the Call Home Test Succeeded window opens (as shown in Figure 5-42), the Call Home configuration is correct. Click **OK** to close the window. Click **Next**.

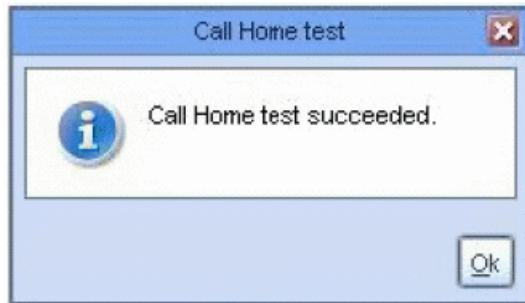


Figure 5-42 Call Home test result

If any of the configuration is not correct, you will see an error message similar to Figure 5-43. You will need to check your configuration and environment if it allows the Call home feature. Contact IBM for assistance.

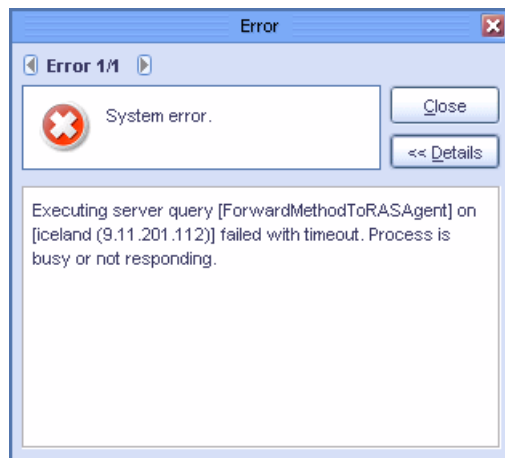


Figure 5-43 Error message when Call Home fails

8. The Summary report window opens, as shown in Figure 5-44.

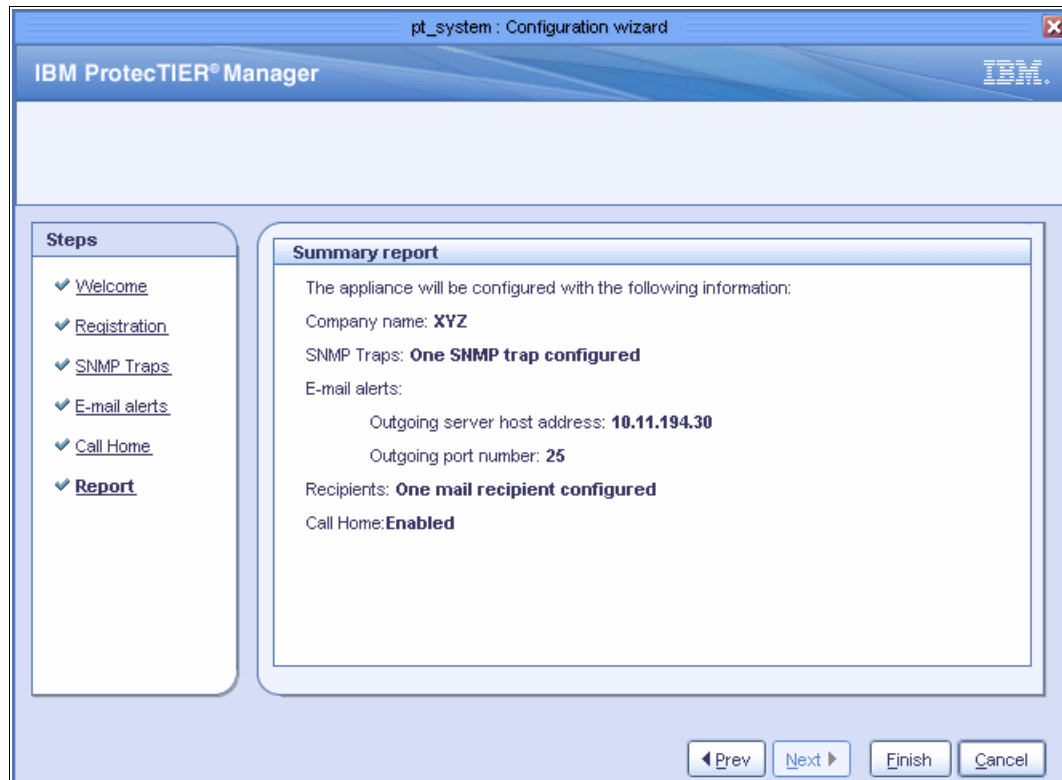


Figure 5-44 Summary report for the Configuration wizard

9. Review the information in the Summary Report. If you need to make changes to the information displayed, click **Prev** to return to the applicable window(s). Make the necessary changes and then click **Next** until you return to the Summary Report.
10. Click **Finish**. The wizard will configure the system based on the information you provided. When configuration finishes, the processing window closes.
11. If you are using the system for VTL, go to 5.6, "Setting up the virtual library and cartridges" on page 230 if you need more tape libraries. If you are using the system for OST, go to 5.7, "Working with OpenStorage using ProtecTIER" on page 240.

## 5.5 Creating a ProtecTIER repository for TS7650G

This step is needed only for TS7650G, because the storage is not included with the product. For TS7610 and TS7650 Appliances, the file systems and the repository are already created for you.

The process of creating a repository requires a careful planning of the repository, which enables you to create the necessary file systems and the repository itself for high performance. For details, refer to 3.3.5, "Local repository sizing" on page 75, and 3.3.7, "Storage sizing" on page 80.

## 5.5.1 Creating file systems

File systems are automatically created with the `fsCreate` tool. This tool scans for LUNs presented by the disk array and creates the Red Hat Global File System (GFS) file systems for each of these LUNs. More specifically, the tool performs the following commands associated with creating GFS file systems:

- ▶ Creates partitions for each LUN.
- ▶ Creates a physical volume for each partition.
- ▶ Creates a volume group for each partition.
- ▶ Creates a logical volume for each partition.
- ▶ Creates a GFS file system for each logical volume.
- ▶ Creates a mount point for each GFS file system.
- ▶ Adds a mount point for each GFS file system in `/etc/fstab`.
- ▶ Mounts all GFS file systems.

**Note:** If there are two-nodes, make sure to shut down the `vtfd` service on node B and power off node B (if it exists) before creating file systems to avoid any possibility of corrupting metadata on the repository's management file system.

This tool is used for the creation of file systems during the first installation. It is also used to create the new GFS file systems for new LUNs presented to the ProtecTIER system during a capacity upgrade.

**Note:** This tool generates a file system for each LUN presented from the disk array. In addition, this tool only supports creating a single file system that makes up the entire size of the LUN for each of the LUNs presented by the disk array. For example, if the disk array presents five LUNs of 1 TB each, the tool will create five 1 TB file systems.

There is no option to select certain LUNs to be used for ProtecTIER file system; all the LUNs presented on the server and unused will be formatted. Remove any LUNs that you do not want to be used by the ProtecTIER so that the Linux OS does not see it before you run the `fsCreate` command.

To create the file systems, complete the following steps:

1. Open a secured shell (SSH) session to the node and log into a ProtecTIER node using the root user ID and admin password.
2. Check that the LUNs are seen by the Red Hat Linux OS (Example 5-1). Run the following command:

```
/sbin/multipath -ll
```

*Example 5-1 Checking LUN visibility*

```
[root@barcelona ~]# /sbin/multipath -ll
mpath2 (360080e500017c00c000018f34cc4b3ba) dm-2 IBM,1814          FASTT
[size=1.6T][features=1 queue_if_no_path][hwandler=1 rdac][rw]
  \_ round-robin 0 [prio=100][active]
     \_ 5:0:0:8 sdp 8:240 [active][ready]
  \_ round-robin 0 [prio=0][enabled]
     \_ 3:0:0:8 sdd 8:48 [active][ghost]
mpath1 (360080e500017c00c000018f04cc4b397) dm-1 IBM,1814          FASTT
[size=2.0T][features=1 queue_if_no_path][hwandler=1 rdac][rw]
  \_ round-robin 0 [prio=100][active]
     \_ 5:0:0:7 sdo 8:224 [active][ready]
```

```
\_ round-robin 0 [prio=0][enabled]
\_ 3:0:0:7 sdc 8:32 [active][ghost]
```

---

**Note:** If the disk system has presented 10 LUNs to the ProtecTIER system, you should see 10 mpath devices, for example, mpath0 to mpath9. It is alright that the mpath number may not be in sequence when displayed using the `multipath` command, as shown in the Example 5-1; this sequence does not affect the GFS file system creation. You can also see that there are two paths to each device here.

3. Run the following command:

```
cd /opt/dtc/app/sbin
```

4. To display all the unused devices, run the following command:

```
./fsCreate -u
```

Example 5-2 shows the output of this command.

*Example 5-2 Displaying unused devices*

---

```
[root@nodeA ~]# /opt/dtc/app/sbin/fsCreate -u
[INFO] [09/02/17-16:35:24] New devices are:
[INFO] [09/02/17-16:35:24] /dev/mapper/mpath7
[INFO] [09/02/17-16:35:24] /dev/mapper/mpath6
[INFO] [09/02/17-16:35:24] /dev/mapper/mpath5
[INFO] [09/02/17-16:35:24] /dev/mapper/mpath11
[INFO] [09/02/17-16:35:24] /dev/mapper/mpath4
[INFO] [09/02/17-16:35:24] /dev/mapper/mpath10
[INFO] [09/02/17-16:35:24] /dev/mapper/mpath3
```

---

**Note:** Make sure that all the devices that are unused are the ones that will be used to create the ProtecTIER repository shown in Example 5-2. Remove any LUNs that you do not want ProtecTIER to format.

5. Check if any file systems are part of a repository by running the following command:

```
./fsCreate -r
```

Example 5-3 shows the output of this command.

*Example 5-3 No file systems belonging to any repository*

---

```
[root@Tuscany sbin]# ./fsCreate -r
[INFO] [10/11/22-14:36:36] No repository file systems found
```

---

**Note:** For a new TS7650G, you should not see any file system listed as part of a repository, as shown in Example 5-3. If you see any file systems listed as part of repository, as shown in Example 5-4, there are two options you can take.

- ▶ Option 1: Remove the entire repository. Make sure you do not have any data that you need that is stored on this repository.
- ▶ Option 2: Add capacity to the existing repository if you have data on your repository.

*Example 5-4 File systems that are part of an existing repository*

---

```
[root@barcelona sbin]# ./fsCreate -r
[INFO] [10/11/19-10:24:51] Repository file systems are:
```



```
[INFO] [10/11/19-10:24:51] /dev/mapper/vg0-lv_vg0 /mnt/vg0-lv_vg0
[INFO] [10/11/19-10:24:51] /dev/mapper/vg1-lv_vg1 /mnt/vg1-lv_vg1
[INFO] [10/11/19-10:24:51] /dev/mapper/vg2-lv_vg2 /mnt/vg2-lv_vg2
[INFO] [10/11/19-10:24:51] /dev/mapper/vg3-lv_vg3 /mnt/vg3-lv_vg3
[INFO] [10/11/19-10:24:51] /dev/mapper/vg4-lv_vg4 /mnt/vg4-lv_vg4
[INFO] [10/11/19-10:24:51] /dev/mapper/vg5-lv_vg5 /mnt/vg5-lv_vg5
```

---

- If you are creating file systems for the first time and there is no file systems that is part of any repository, go to step 7.
  - If you want to delete all the data and recreate the repository, there are existing file system you need to remove; go to 5.5.5, “Deleting existing repository and file systems” on page 227.
  - If you want to add capacity using new available multipath devices to an existing repository, go to step 10 on page 214.
6. If you are working in a clustered configuration, shut down the vtfd service on Server B during the file system creation to avoid any possibility of corrupting metadata on the repository's management file system and then power off the server. Run the following commands:

```
service vtfd shutdown
poweroff
```

7. After making sure that all the LUNs are available, to create the repository for the first time, run the following command:

```
./fsCreate -n
```

The **fsCreate** command removes any existing data on the disk array as a result of creating the file systems (Example 5-5). Enter **data loss** to continue.

*Example 5-5 Removing existing data*

---

```
[root@Tuscany sbin]# ./fsCreate -n
[WARNING] [10/11/22-14:43:03] This tool will create filesystems for the repository and
will remove any existing data on the repository.
[WARNING] [10/11/22-14:43:03] To continue type "data loss"
data loss
[INFO] [10/11/22-14:43:13] Creating partition for mpath1
[INFO] [10/11/22-14:43:14] Partition for /dev/mapper/mpath1 created successfully
[INFO] [10/11/22-14:43:14] Assigning partition size for mpath1
[INFO] [10/11/22-14:43:15] Partition size of 1074MB for /dev/mapper/mpath1 assigned
successfully
[INFO] [10/11/22-14:43:15] Creating physical volume for /dev/mapper/mpath1p1
[INFO] [10/11/22-14:43:16] Physical volume for /dev/mapper/mpath1p1 created successfully
[INFO] [10/11/22-14:43:16] Creating volume group vg0 for /dev/mapper/mpath1p1
[INFO] [10/11/22-14:43:17] Volume group vg0 for /dev/mapper/mpath1p1 created
successfully
[INFO] [10/11/22-14:43:17] Creating logical volume lv_vg0 with volume group vg0
[INFO] [10/11/22-14:43:19] Logical volume lv_vg0 with volume group vg0 created
successfully
[INFO] [10/11/22-14:43:19] Creating filesystem gfs_lv_vg0
[INFO] [10/11/22-14:43:22] Filesystem gfs_lv_vg0 created successfully
[INFO] [10/11/22-14:43:22] Creating partition for mpath0
[INFO] [10/11/22-14:43:23] Partition for /dev/mapper/mpath0 created successfully
[INFO] [10/11/22-14:43:23] Assigning partition size for mpath0
[INFO] [10/11/22-14:43:25] Partition size of 1197GB for /dev/mapper/mpath0 assigned
successfully
[INFO] [10/11/22-14:43:25] Creating physical volume for /dev/mapper/mpath0p1
[INFO] [10/11/22-14:43:26] Physical volume for /dev/mapper/mpath0p1 created successfully
```

```

[INFO] [10/11/22-14:43:26] Creating volume group vg1 for /dev/mapper/mpath0p1
[INFO] [10/11/22-14:43:27] Volume group vg1 for /dev/mapper/mpath0p1 created
successfully
[INFO] [10/11/22-14:43:27] Creating logical volume lv_vg1 with volume group vg1
[INFO] [10/11/22-14:43:28] Logical volume lv_vg1 with volume group vg1 created
successfully
[INFO] [10/11/22-14:43:28] Creating filesystem gfs_lv_vg1
[INFO] [10/11/22-14:43:33] Filesystem gfs_lv_vg1 created successfully
[INFO] [10/11/22-14:43:33] Creating partition for mpath9
[INFO] [10/11/22-14:43:34] Partition for /dev/mapper/mpath9 created successfully
[INFO] [10/11/22-14:43:34] Assigning partition size for mpath9
[INFO] [10/11/22-14:43:35] Partition size of 591GB for /dev/mapper/mpath9 assigned
successfully
[INFO] [10/11/22-14:43:35] Creating physical volume for /dev/mapper/mpath9p1
[INFO] [10/11/22-14:43:36] Physical volume for /dev/mapper/mpath9p1 created successfully
[INFO] [10/11/22-14:43:36] Creating volume group vg2 for /dev/mapper/mpath9p1
[INFO] [10/11/22-14:43:37] Volume group vg2 for /dev/mapper/mpath9p1 created
successfully
[INFO] [10/11/22-14:43:37] Creating logical volume lv_vg2 with volume group vg2
[INFO] [10/11/22-14:43:39] Logical volume lv_vg2 with volume group vg2 created
successfully
[INFO] [10/11/22-14:43:39] Creating filesystem gfs_lv_vg2
[INFO] [10/11/22-14:43:43] Filesystem gfs_lv_vg2 created successfully
[INFO] [10/11/22-14:43:43] Creating partition for mpath8
[INFO] [10/11/22-14:43:44] Partition for /dev/mapper/mpath8 created successfully
[INFO] [10/11/22-14:43:44] Assigning partition size for mpath8
[INFO] [10/11/22-14:43:45] Partition size of 591GB for /dev/mapper/mpath8 assigned
successfully
[INFO] [10/11/22-14:43:45] Creating physical volume for /dev/mapper/mpath8p1
[INFO] [10/11/22-14:43:46] Physical volume for /dev/mapper/mpath8p1 created successfully
[INFO] [10/11/22-14:43:46] Creating volume group vg3 for /dev/mapper/mpath8p1
[INFO] [10/11/22-14:43:48] Volume group vg3 for /dev/mapper/mpath8p1 created
successfully
[INFO] [10/11/22-14:43:48] Creating logical volume lv_vg3 with volume group vg3
[INFO] [10/11/22-14:43:49] Logical volume lv_vg3 with volume group vg3 created
successfully
[INFO] [10/11/22-14:43:49] Creating filesystem gfs_lv_vg3
[INFO] [10/11/22-14:43:53] Filesystem gfs_lv_vg3 created successfully
[INFO] [10/11/22-14:43:53] Creating partition for mpath7
[INFO] [10/11/22-14:43:54] Partition for /dev/mapper/mpath7 created successfully
[INFO] [10/11/22-14:43:54] Assigning partition size for mpath7
[INFO] [10/11/22-14:43:56] Partition size of 591GB for /dev/mapper/mpath7 assigned
successfully
[INFO] [10/11/22-14:43:56] Creating physical volume for /dev/mapper/mpath7p1
[INFO] [10/11/22-14:43:57] Physical volume for /dev/mapper/mpath7p1 created successfully
[INFO] [10/11/22-14:43:57] Creating volume group vg4 for /dev/mapper/mpath7p1
[INFO] [10/11/22-14:43:59] Volume group vg4 for /dev/mapper/mpath7p1 created
successfully
[INFO] [10/11/22-14:43:59] Creating logical volume lv_vg4 with volume group vg4
[INFO] [10/11/22-14:44:00] Logical volume lv_vg4 with volume group vg4 created
successfully
[INFO] [10/11/22-14:44:00] Creating filesystem gfs_lv_vg4
[INFO] [10/11/22-14:44:06] Filesystem gfs_lv_vg4 created successfully
[INFO] [10/11/22-14:44:06] Creating partition for mpath6
[INFO] [10/11/22-14:44:08] Partition for /dev/mapper/mpath6 created successfully
[INFO] [10/11/22-14:44:08] Assigning partition size for mpath6
[INFO] [10/11/22-14:44:09] Partition size of 591GB for /dev/mapper/mpath6 assigned
successfully
[INFO] [10/11/22-14:44:09] Creating physical volume for /dev/mapper/mpath6p1
[INFO] [10/11/22-14:44:10] Physical volume for /dev/mapper/mpath6p1 created successfully

```

```

[INFO] [10/11/22-14:44:10] Creating volume group vg5 for /dev/mapper/mpath6p1
[INFO] [10/11/22-14:44:12] Volume group vg5 for /dev/mapper/mpath6p1 created
successfully
[INFO] [10/11/22-14:44:12] Creating logical volume lv_vg5 with volume group vg5
[INFO] [10/11/22-14:44:14] Logical volume lv_vg5 with volume group vg5 created
successfully
[INFO] [10/11/22-14:44:14] Creating filesystem gfs_lv_vg5
[INFO] [10/11/22-14:44:18] Filesystem gfs_lv_vg5 created successfully
[INFO] [10/11/22-14:44:18] Creating partition for mpath10
[INFO] [10/11/22-14:44:19] Partition for /dev/mapper/mpath10 created successfully
[INFO] [10/11/22-14:44:19] Assigning partition size for mpath10
[INFO] [10/11/22-14:44:20] Partition size of 591GB for /dev/mapper/mpath10 assigned
successfully
[INFO] [10/11/22-14:44:20] Creating physical volume for /dev/mapper/mpath10p1
[INFO] [10/11/22-14:44:21] Physical volume for /dev/mapper/mpath10p1 created
successfully
[INFO] [10/11/22-14:44:21] Creating volume group vg6 for /dev/mapper/mpath10p1
[INFO] [10/11/22-14:44:23] Volume group vg6 for /dev/mapper/mpath10p1 created
successfully
[INFO] [10/11/22-14:44:23] Creating logical volume lv_vg6 with volume group vg6
[INFO] [10/11/22-14:44:24] Logical volume lv_vg6 with volume group vg6 created
successfully
[INFO] [10/11/22-14:44:24] Creating filesystem gfs_lv_vg6
[INFO] [10/11/22-14:44:30] Filesystem gfs_lv_vg6 created successfully
[INFO] [10/11/22-14:44:30] Adding mount point /mnt/vg0-lv_vg0 to /etc/fstab
[INFO] [10/11/22-14:44:30] Mount point /mnt/vg0-lv_vg0 successfully added to /etc/fstab
[INFO] [10/11/22-14:44:30] Adding mount point /mnt/vg1-lv_vg1 to /etc/fstab
[INFO] [10/11/22-14:44:30] Mount point /mnt/vg1-lv_vg1 successfully added to /etc/fstab
[INFO] [10/11/22-14:44:30] Adding mount point /mnt/vg2-lv_vg2 to /etc/fstab
[INFO] [10/11/22-14:44:30] Mount point /mnt/vg2-lv_vg2 successfully added to /etc/fstab
[INFO] [10/11/22-14:44:30] Adding mount point /mnt/vg3-lv_vg3 to /etc/fstab
[INFO] [10/11/22-14:44:30] Mount point /mnt/vg3-lv_vg3 successfully added to /etc/fstab
[INFO] [10/11/22-14:44:30] Adding mount point /mnt/vg4-lv_vg4 to /etc/fstab
[INFO] [10/11/22-14:44:30] Mount point /mnt/vg4-lv_vg4 successfully added to /etc/fstab
[INFO] [10/11/22-14:44:30] Adding mount point /mnt/vg5-lv_vg5 to /etc/fstab
[INFO] [10/11/22-14:44:30] Mount point /mnt/vg5-lv_vg5 successfully added to /etc/fstab
[INFO] [10/11/22-14:44:30] Creating mount point /mnt/vg6-lv_vg6
[INFO] [10/11/22-14:44:30] Mount point /mnt/vg6-lv_vg6 created successfully
[INFO] [10/11/22-14:44:30] Adding mount point /mnt/vg6-lv_vg6 to /etc/fstab
[INFO] [10/11/22-14:44:30] Mount point /mnt/vg6-lv_vg6 successfully added to /etc/fstab
[INFO] [10/11/22-14:44:35] Mounting filesystems
[INFO] [10/11/22-14:44:37] Filesystems successfully mounted

```

8. Check that the file systems are actually mounted on the system, as shown in Example 5-6. Run the following command:

```
df -k
```

*Example 5-6 Checking the mount state*

```

[root@Tuscany sbin]# df -k
Filesystem            1K-blocks      Used Available Use% Mounted on
/dev/sda2              19840924    10259376   8557408   55% /
/dev/sda3              101865496     875264   95732160    1% /pt_work
/dev/sda1              147764       21027    119108   16% /boot
tmpfs                  32995404         0   32995404    0% /dev/shm
/dev/mapper/vg0-lv_vg0
                        651088         20    651068    1% /mnt/vg0-lv_vg0
/dev/mapper/vg1-lv_vg1
                        1168223680     444 1168223236    1% /mnt/vg1-lv_vg1

```

/dev/mapper/vg2-lv_vg2	576266624	228	576266396	1%	/mnt/vg2-lv_vg2
/dev/mapper/vg3-lv_vg3	576266624	228	576266396	1%	/mnt/vg3-lv_vg3
/dev/mapper/vg4-lv_vg4	576266624	228	576266396	1%	/mnt/vg4-lv_vg4
/dev/mapper/vg5-lv_vg5	576266624	228	576266396	1%	/mnt/vg5-lv_vg5
/dev/mapper/vg6-lv_vg6	576266624	228	576266396	1%	/mnt/vg6-lv_vg6

9. To check that the file systems will be mounted automatically after reach reboot, check that the file systems are listed in the `/etc/fstab` file. Each LUN will be formatted into a logical volume. There are seven LUNs used for this system, as shown in Example 5-7, which have been formatted and mounted as `vg0-lv_vg0` to `vg6-lv_vg6`. You can run the following command to check on your system:

```
cat /etc/fstab
```

*Example 5-7 Checking automatic reboot*

```
[root@Tuscany sbin]# cat /etc/fstab
LABEL=/ / ext3 defaults 1 1
LABEL=/pt_work /pt_work ext3 defaults 1 2
LABEL=/boot /boot ext3 defaults 1 2
tmpfs /dev/shm tmpfs defaults 0 0
devpts /dev/pts devpts gid=5,mode=620 0 0
sysfs /sys sysfs defaults 0 0
proc /proc proc defaults 0 0
LABEL=SWAP-sda5 swap swap defaults 0 0
/dev/mapper/vg0-lv_vg0 /mnt/vg0-lv_vg0 gfs defaults,noatime,nodiratime,noquota 0 0
/dev/mapper/vg1-lv_vg1 /mnt/vg1-lv_vg1 gfs defaults,noatime,nodiratime,noquota 0 0
/dev/mapper/vg2-lv_vg2 /mnt/vg2-lv_vg2 gfs defaults,noatime,nodiratime,noquota 0 0
/dev/mapper/vg3-lv_vg3 /mnt/vg3-lv_vg3 gfs defaults,noatime,nodiratime,noquota 0 0
/dev/mapper/vg4-lv_vg4 /mnt/vg4-lv_vg4 gfs defaults,noatime,nodiratime,noquota 0 0
/dev/mapper/vg5-lv_vg5 /mnt/vg5-lv_vg5 gfs defaults,noatime,nodiratime,noquota 0 0
/dev/mapper/vg6-lv_vg6 /mnt/vg6-lv_vg6 gfs defaults,noatime,nodiratime,noquota 0 0
```

10. To create file systems on new available multipath devices, we need to register them to `/etc/fstab` and mount them. Run the following command on the ProtecTIER system to list the newly available multipath devices:

```
./fsCreate -u
```

Newly available multipath devices are shown in Example 5-8.

*Example 5-8 Listing new devices*

```
[root@Tuscany sbin]# ./fsCreate -u
[INFO] [10/11/22-15:04:26] New devices are:
[INFO] [10/11/22-15:04:26] /dev/mapper/mpath11
```

Example 5-9 show the command on the ProtecTIER system to create the new file system on the new multipath devices:

```
./fsCreate -e
```

*Example 5-9 Creating new file system*

```
[root@Tuscany sbin]# ./fsCreate -e
[INFO] [10/11/22-15:07:24] Creating partition for mpath11
```

```

[INFO] [10/11/22-15:07:26] Partition for /dev/mapper/mpath11 created successfully
[INFO] [10/11/22-15:07:26] Assigning partition size for mpath11
[INFO] [10/11/22-15:07:27] Partition size of 591GB for /dev/mapper/mpath11 assigned
successfully
[INFO] [10/11/22-15:07:27] Creating physical volume for /dev/mapper/mpath11p1
[INFO] [10/11/22-15:07:28] Physical volume for /dev/mapper/mpath11p1 created
successfully
[INFO] [10/11/22-15:07:28] Creating volume group vg7 for /dev/mapper/mpath11p1
[INFO] [10/11/22-15:07:29] Volume group vg7 for /dev/mapper/mpath11p1 created
successfully
[INFO] [10/11/22-15:07:29] Creating logical volume lv_vg7 with volume group vg7
[INFO] [10/11/22-15:07:30] Logical volume lv_vg7 with volume group vg7 created
successfully
[INFO] [10/11/22-15:07:30] Creating filesystem gfs_lv_vg7
[INFO] [10/11/22-15:07:34] Filesystem gfs_lv_vg7 created successfully
[INFO] [10/11/22-15:07:34] Creating mount point /mnt/vg7-lv_vg7
[INFO] [10/11/22-15:07:34] Mount point /mnt/vg7-lv_vg7 created successfully
[INFO] [10/11/22-15:07:34] Adding mount point /mnt/vg7-lv_vg7 to /etc/fstab
[INFO] [10/11/22-15:07:34] Mount point /mnt/vg7-lv_vg7 successfully added to /etc/fstab
[INFO] [10/11/22-15:07:34] Mounting filesystems
[INFO] [10/11/22-15:07:34] Filesystems successfully mounted

```

**Note:** The fsCreate script removes any existing data on the disk array as a result of creating the file systems.

## 5.5.2 fsCreate parameters

There are a few options that you can use with the **fsCreate** command. Table 5-2 lists the syntax of these options and the description of what each one does.

Table 5-2 *fsCreate* parameters

Parameter	Description	Return code (rc=0)	Return code (rc=1)
-n	Creates GFS file systems for all mpath devices during first-time installation.	All file systems built.	Error encountered building file systems or detected physical volumes created from mpath devices.
-e	Creates GFS file systems for new mpath devices during capacity upgrade.	All new file systems built.	Error encountered building file systems or detected physical volumes created from mpath devices.
-t	Creates mount points and register GFS file systems to /etc/fstab.	All mount points created and registered to /etc/fstab.	Error encountered creating and registering mount points to /etc/fstab or detected physical volumes created from mpath devices.
-r	Displays all repository GFS file systems.	Call completed without error.	Failed to determine status of an existing file system or detected physical volumes created from mpath devices.

Parameter	Description	Return code (rc=0)	Return code (rc=1)
-u	Displays unused devices.	Call completed without error.	Detected physical volumes created from mpath devices.
-g	Displays all non-repository GFS file systems.	Call completed without error.	Failed to determine status of an existing file system or detected physical volumes created from mpath devices.

### 5.5.3 Creating the repository

After the necessary file systems have been created, use the information generated during the repository planning process to create the repository.

#### Repository planning in ProtecTIER Manager

The ProtecTIER Manager has a tool that helps with the planning of the repository. It is named *Create repository planning* wizard. This method is available only if the first node is configured after the file systems are mounted. This is the point in time where the storage must be available.

In addition to guidance from IBM Support, run the Create repository planning wizard to determine the optimum repository size and metadata file system arrangement for your repository.

Complete the following steps:

1. Select **Repository** → **Create repository planning** (Figure 5-45). The Create repository planning wizard opens.

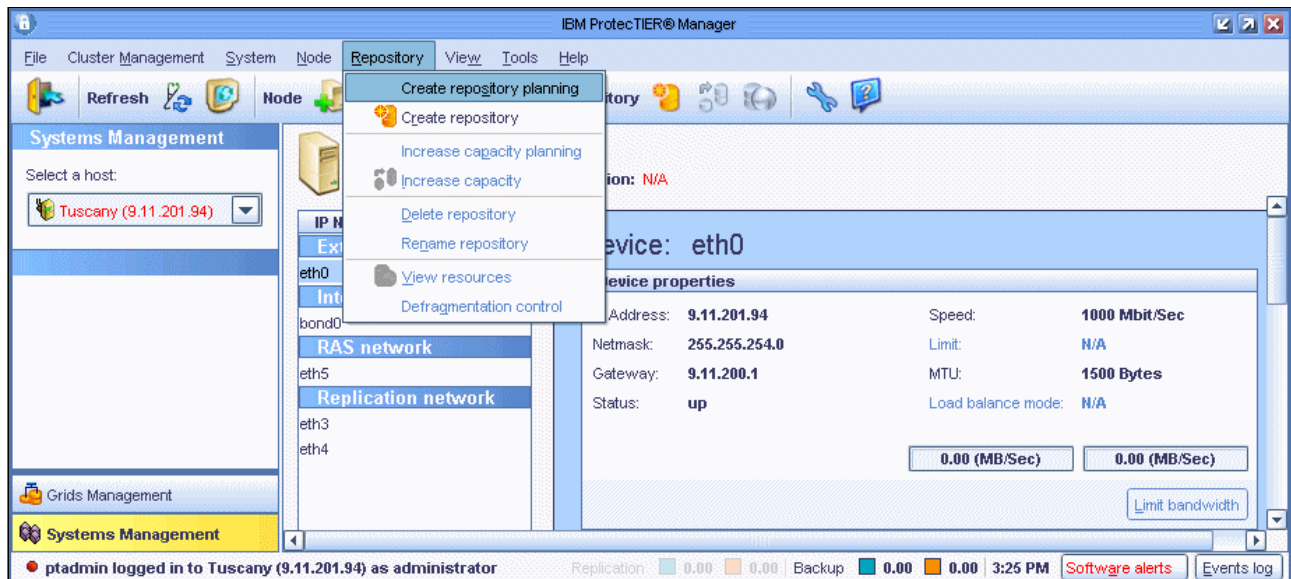


Figure 5-45 Create repository planning

2. The Plan repository creation window opens (Figure 5-46). Select the backup interface. It can be either Virtual Tape Library or OpenStorage (OST); it cannot be both at the same time.

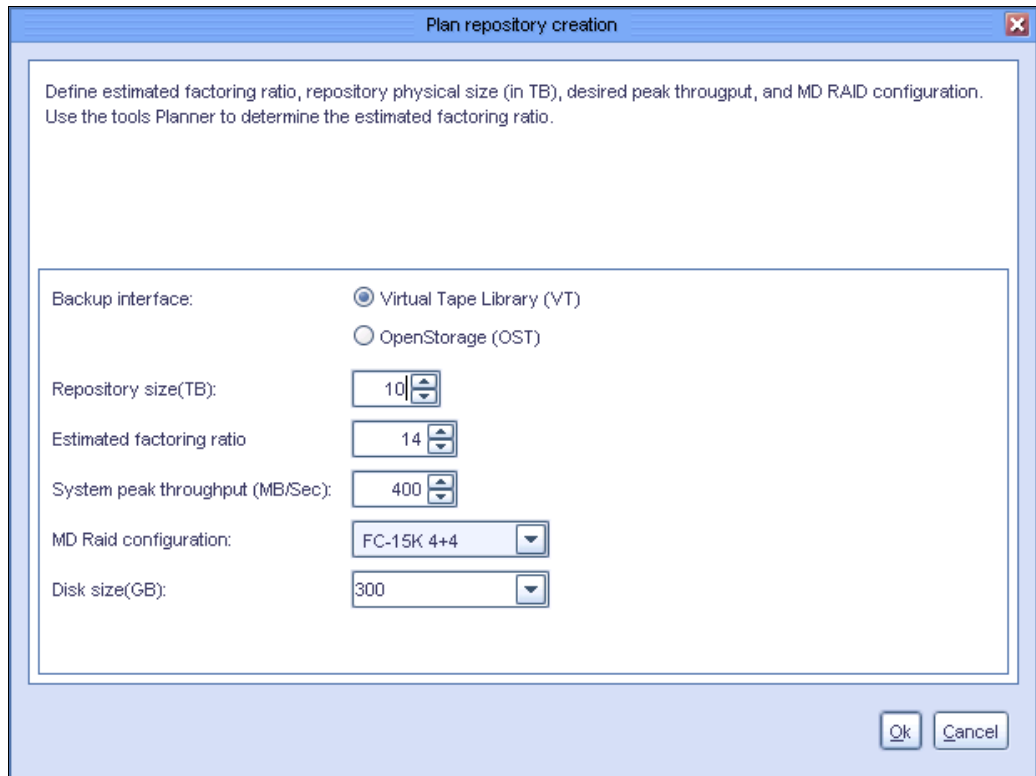


Figure 5-46 Create repository planning wizard

3. In the Repository size field, enter the size value of the repository that you want to create.

**Note:** The maximum possible repository physical size is 1 PB.

4. In the Estimated factoring ratio field, enter the value estimated for your environment based on your data change rate, backup policies, and retention period.
5. In the System peak throughput field, specify the rate of system peak throughput that your metadata file systems can support.
6. In the MD RAID configuration field, select the RAID configuration of the logical volumes on which the repository metadata file systems are to be created. For example, select FC-15K 4+4 for a configuration of RAID 10 4+4 with Fibre Channel 15 K RPM disks.
7. In the Disk size field, enter the size of the disks that you use in your storage array.

- Click **OK**. A processing window opens, as shown in Figure 5-47). The Repository metadata storage requirements window opens with a list of file system arrangement options that are suitable for your needs (Figure 5-48).

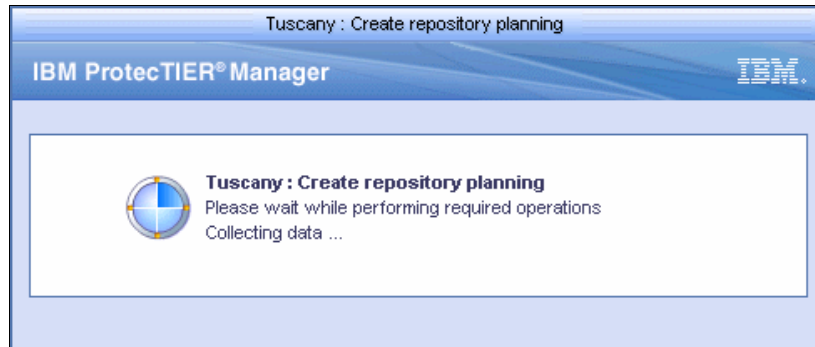


Figure 5-47 Create repository planning

**Repository meta data storage requirements**

The table shows metadata LUNs and file system requirements based on the provided repository parameters (such as repository size and factoring ratio), to the required disk LUN size column.

⚠ LUN's space may not be 100% utilized due to performance considerations.  
 ⚠ Using the recommended disk LUN size will enable easier MD growth in the future.

Repository requirements				
Size (TB)	File systems #	Memory (GB)	Recommended disk LUN size	Required MD size
14	3	16.0	~1,024.0 MB	1,024.0 MB
			~1,117.6 GB	311.5 GB
			~1,117.6 GB	339.7 GB
15	3	16.0	~1,024.0 MB	1,024.0 MB
			~1,117.6 GB	328.6 GB
			~1,117.6 GB	361.0 GB
16	3	16.0	~1,024.0 MB	1,024.0 MB
			~1,117.6 GB	345.8 GB
			~1,117.6 GB	382.4 GB
17	3	16.0	~1,024.0 MB	1,024.0 MB
			~1,117.6 GB	365.0 GB
			~1,117.6 GB	405.7 GB
18	3	16.0	~1,024.0 MB	1,024.0 MB
			~1,117.6 GB	382.1 GB
			~1,117.6 GB	427.0 GB
19	3	16.0	~1,024.0 MB	1,024.0 MB
			~1,117.6 GB	403.3 GB
			~1,117.6 GB	448.4 GB

**Summary report**

Repository plan is configured with the following properties:

- Repository size: **10 TB**
- Estimated factoring ratio: **14**
- Peak throughput(MB/Sec): **400**
- Raid configuration: **FC-15K 4+4**
- Raid member disk size: **300GB**
- MD Raid configuration: **FC-15K 4+4**
- Backup interface: **VTL**

Figure 5-48 Repository metadata storage requirements



9. Click **Options** to print or save as a .csv file the information in the Repository metadata storage requirements window using the standard procedures for your operating system.  
Click **OK**. The Repository metadata storage requirements window closes.
10. Using the information provided in the Repository metadata storage requirements window, choose the metadata file system arrangement that is most appropriate for your needs. For example, your projected repository size is 15 TB, but you notice that the file system arrangement for 16 TB more closely matches the logical disk arrangement of your disk array. In that case, choose 16 TB instead of the original planned value.

**Note:** For existing customers who are using the 3958-DD1 or 3958-DD3 model, there is a new option to increase the memory of the ProtecTIER system from 32 GB to 64 GB. This upgrade is only necessary if you want to have a repository size greater than 512 TB with replication enabled and using more than 128 virtual drives per node.

In a clustered configuration, the upgrade procedure should allow for one node to remain online providing customer data access while the peer node in the cluster undergoes the upgrade procedure

## Creating a repository

A repository can only be created on a single node cluster.

**Note:** Creating a repository is a prerequisite for adding a second node to a cluster. The repository creation *must* be done from node A.

To create a repository, complete the following steps:

1. In the Nodes pane, select the node on which to create the repository.
2. Click **Create new repository**. The Create repository wizard Welcome window opens (Figure 5-49).

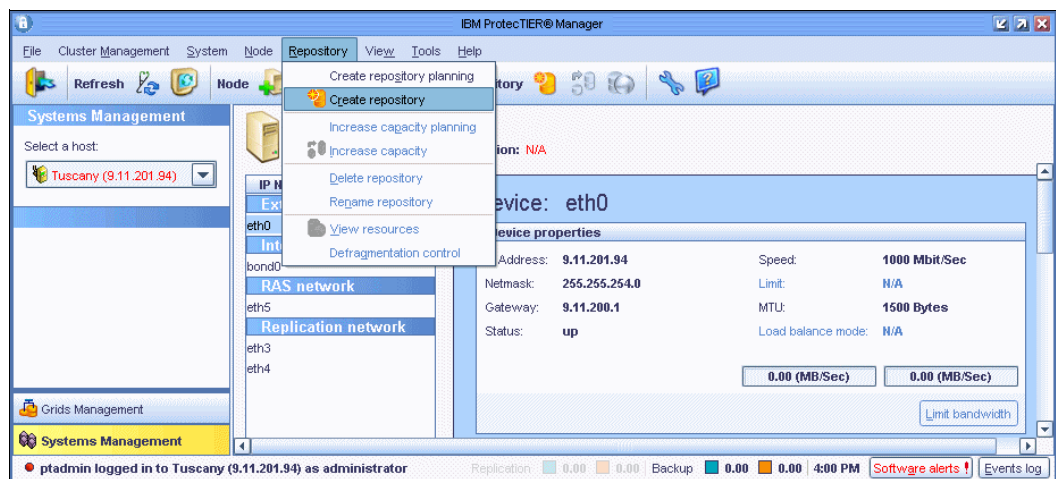


Figure 5-49 Create repository welcome window

Click **Next**. The Repository Name window opens (Figure 5-50).

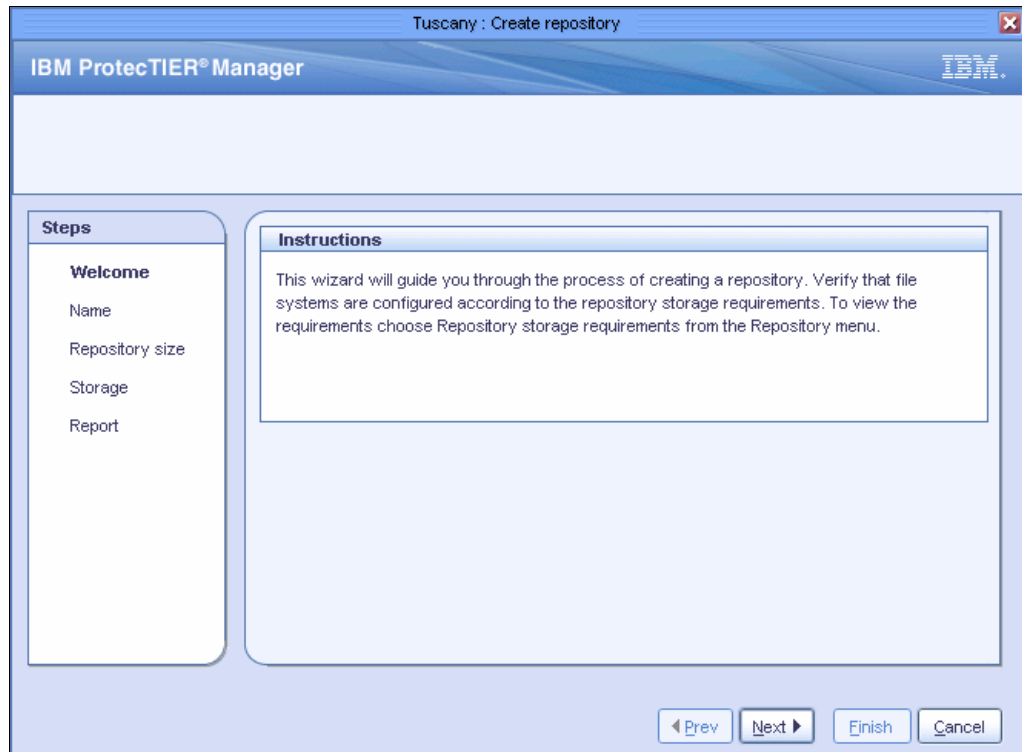


Figure 5-50 Repository name window

3. In the System name field, enter the name of the system on which the repository will be created. In the Repository name field, enter the name of the repository that you are creating.

Click **Next** and the Repository window opens (Figure 5-51).

The screenshot shows the 'IBM ProtectTIER® Manager' window titled 'Tuscany : Create repository'. At the top, there are two information messages: 'System name field limited to 16 characters.' and 'Repository name field limited to 32 characters.'. On the left, a 'Steps' sidebar lists 'Welcome' (checked), 'Name', 'Repository size', 'Storage', and 'Report'. The main area is titled 'Name' and contains two text input fields: 'System name:' with the value 'Tuscany' and 'Repository name:' with the value 'Tuscany\_repository'. At the bottom right, there are four buttons: 'Prev', 'Next', 'Finish', and 'Cancel'.

Figure 5-51 Repository size window

4. In the Repository size field (Figure 5-52), enter the repository size in terabytes that you determined using the Create repository planning wizard.
5. In the Estimated factoring ratio field, enter the estimated factoring ratio value that was determined with the assistance of your IBM System Services Representative (SSR), Lab Services (LBS), and Field Technical Service and Support (FTSS) personnel.
6. In the System peak throughput field, specify the rate of system peak throughput that your metadata file systems can support.
7. In the Metadata RAID configuration field, select the RAID configuration of the logical volumes on which the repository metadata file systems are to be created. For example, select FC-15K 4+4 for a configuration of RAID 10 4+4 with Fibre Channel 15 K RPM disks. This configuration depends on the way that your storage arrays are configured. Other choices are Serial Attached SCSI (SAS) storage arrays and Serial Advanced Technology Attachment (SATA) storage arrays.
8. In the Disk size field, enter the size of the disks that are in your storage array. Click **Next**.

**Note:** The TS7650 Appliance comes with FC-15 K 450 GB 4 Gbps Fibre Channel disk drive modules (DDMs) configured on an IBM System Storage DS4700.

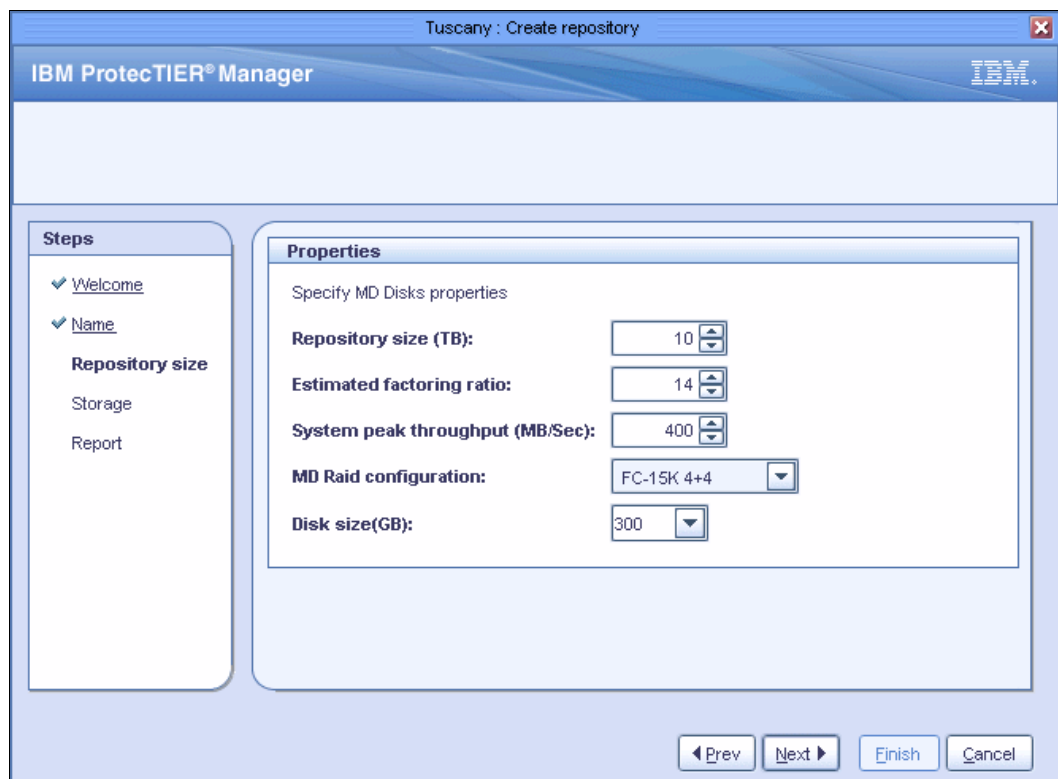


Figure 5-52 Repository size properties window

The Metadata window opens (Figure 5-53). The Allocated metadata size field shows the amount of disk space allocated for metadata, and the Allocated user data size field shows the amount of disk space allocated for user data, based on the estimated factoring ratio and the set of existing file systems (Figure 5-52 on page 222).

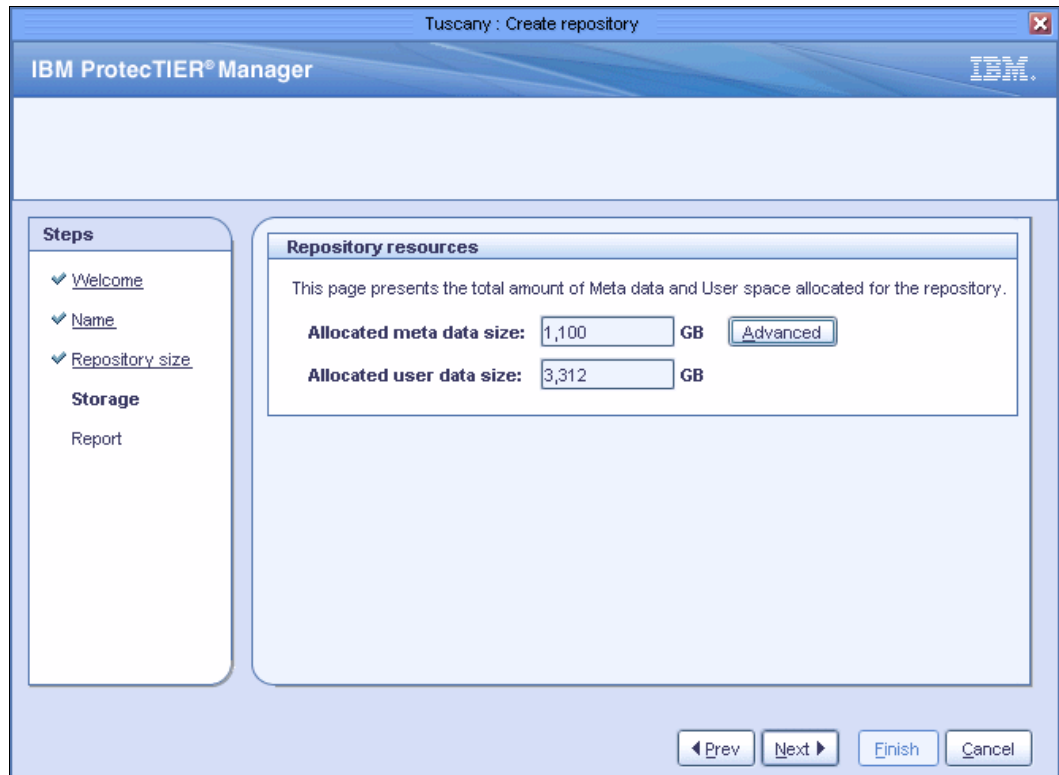


Figure 5-53 Metadata and user data allocation

- The Repository resources window also automatically open (by using the Advanced button) for you to see the actual LUNs assignment (Figure 5-54).

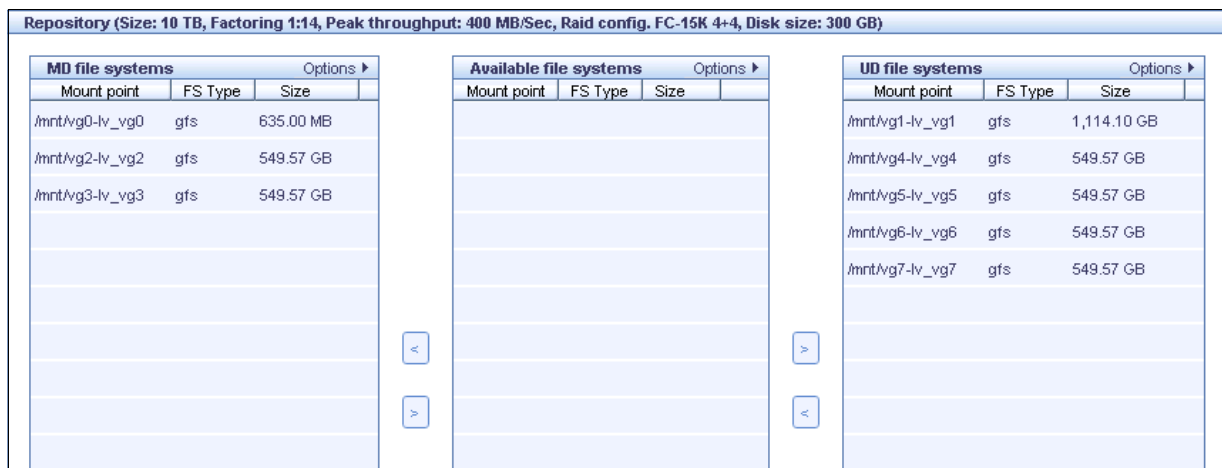


Figure 5-54 Repository resources window

- Verify that the correct file systems are selected for metadata and user data, based on the metadata file system sizes indicated by the repository planning process.

**Note:** By default, the ProtecTIER system generally selects the smallest available file systems for use as metadata file systems. The remaining file systems are available for user data. You cannot assign more storage space for user data than the repository size defined in the Repository Size window.

If the file systems selected by ProtecTIER for metadata and user data do not match the file systems that you created for those purposes, change the assignment by selecting file systems from the Available file systems list and click the arrow buttons to move file systems to and from the MD file systems and UD file systems lists.

11. Click **OK**. The Metadata resources window closes.

**Note:** Optionally, reopen the Metadata resources window by clicking **Advanced** in the Metadata window.

Click **Next**. The Summary Report window opens (Figure 5-55).

12. Click **Finish**.

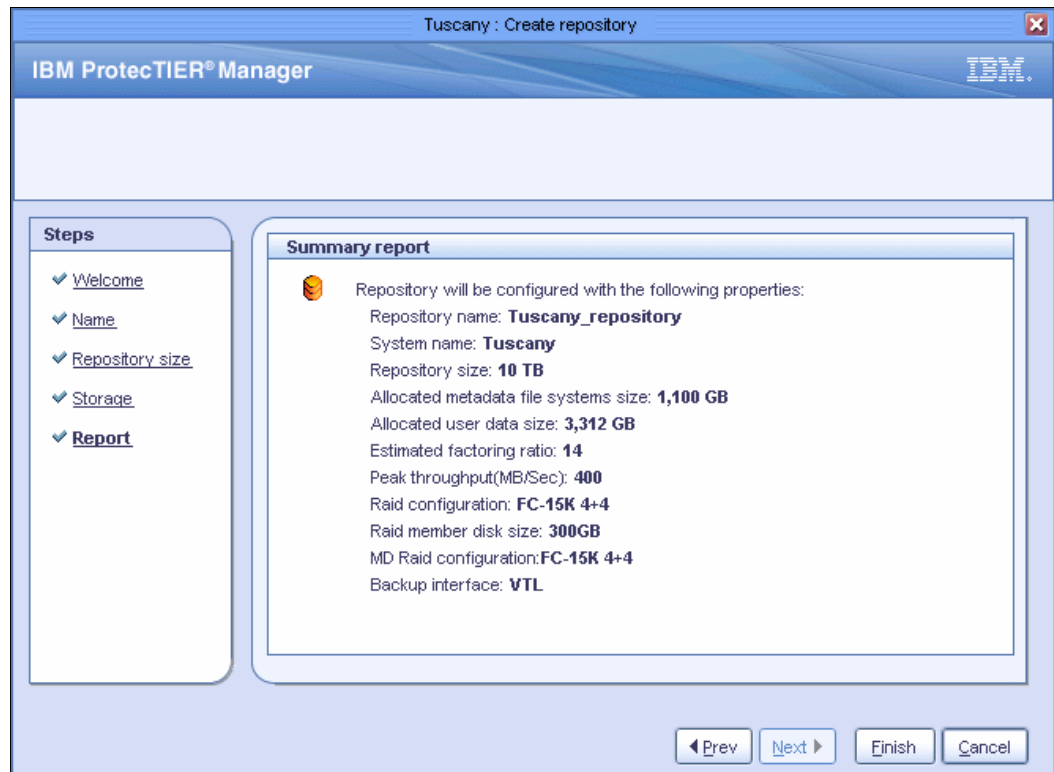


Figure 5-55 Summary Report for create repository

13. The Create repository wizard closes and a confirmation window opens (Figure 5-56).
14. Click **Yes**. The ProtecTIER system temporarily goes offline to create the repository. This operation might take a while.

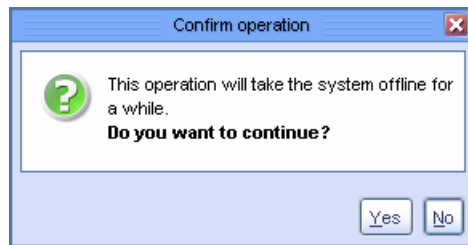


Figure 5-56 Confirm operation window

15. The Create repository window stay open until the repository is created (Figure 5-57).

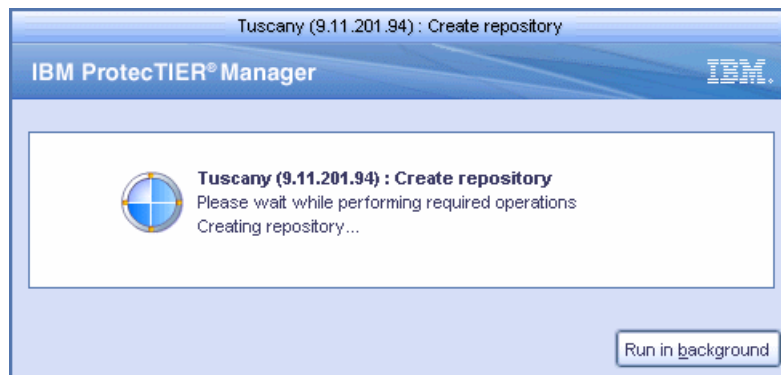


Figure 5-57 Creating repository

#### 5.5.4 Renaming the repository

By default, the repository will be created with the name `pt_cluster`. You change the name of the repository to something more unique so that it will be easier to identify when you are doing VTL or OST replication.

To rename the repository, complete the following steps:

1. In the Systems Management pane, select the host on which to rename the repository. Click the system name (above Statistics). From the Repository menu, select **Rename repository** (Figure 5-58).

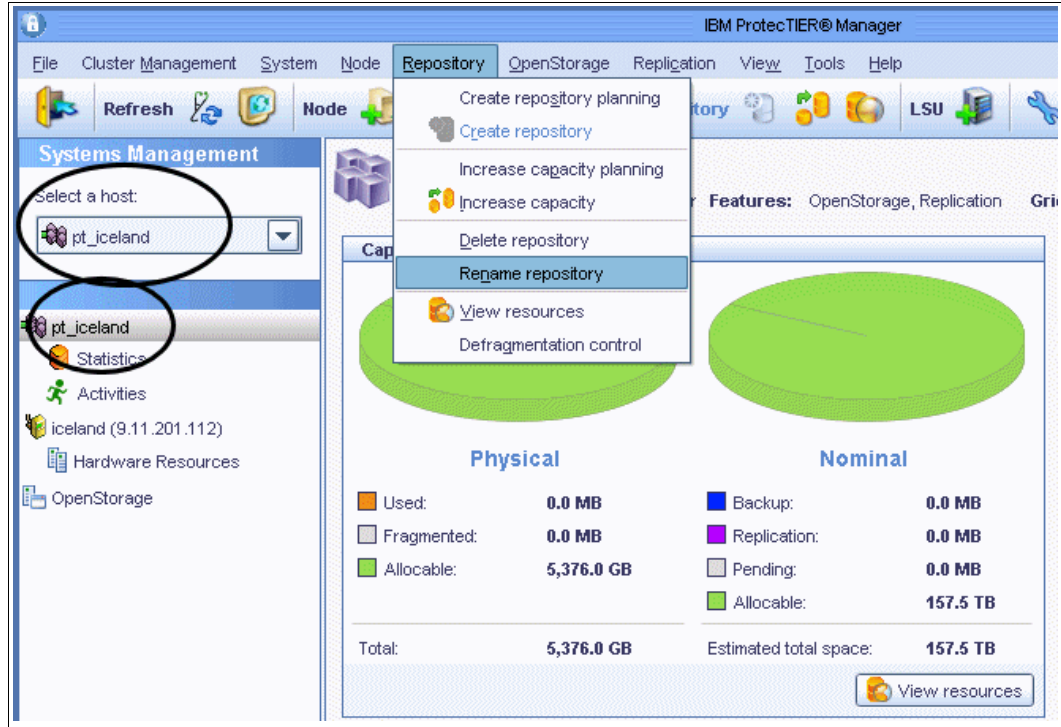


Figure 5-58 Select Rename repository

**Note:** You should change the name of the repository now, so that when you add the repository to a replication grid, you can identify each individual repository better.

2. The Rename repository window opens (Figure 5-59). Enter the new name and click **OK**.



Figure 5-59 Rename repository window



3. The new repository name will be updated immediately (Figure 5-60).

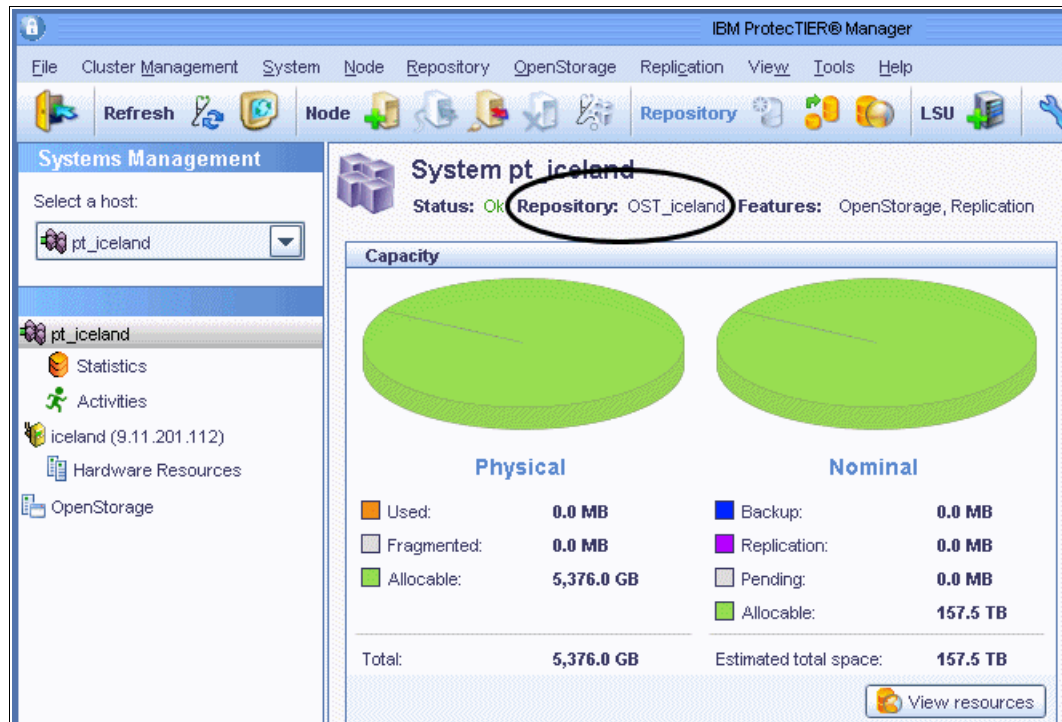


Figure 5-60 New repository name

For more information about managing repositories, refer to 10.2, “Managing repositories” on page 476.

### 5.5.5 Deleting existing repository and file systems

The section describes the steps to delete a repository.

In a two-node cluster, where the repository was created on one of the nodes, you must remove the second node before you delete the repository.

In general, it is not necessary to delete a repository. However, you might have to delete a repository in certain troubleshooting scenarios or when you want to change the file system during the initial setup.

**Note:** Deleting the repository results in the loss of all the data contained in the repository.

Complete the following steps to delete a repository:

1. Select the ProtecTIER system, as shown in Figure 5-61.

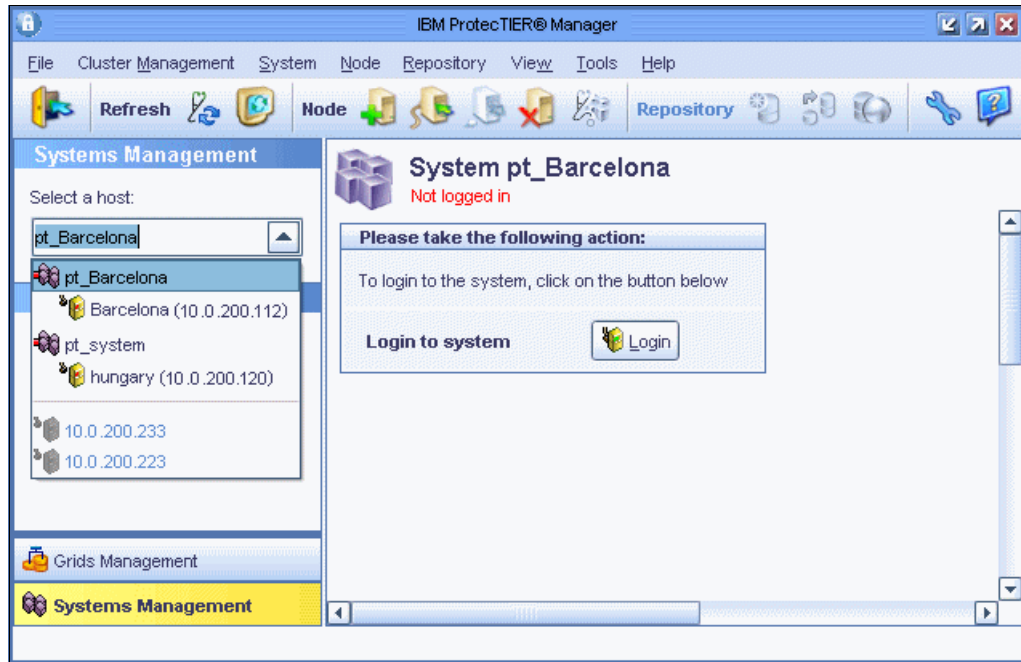


Figure 5-61 Select the ProtecTIER system

2. Select **Repository** → **Delete repository** (Figure 5-62).

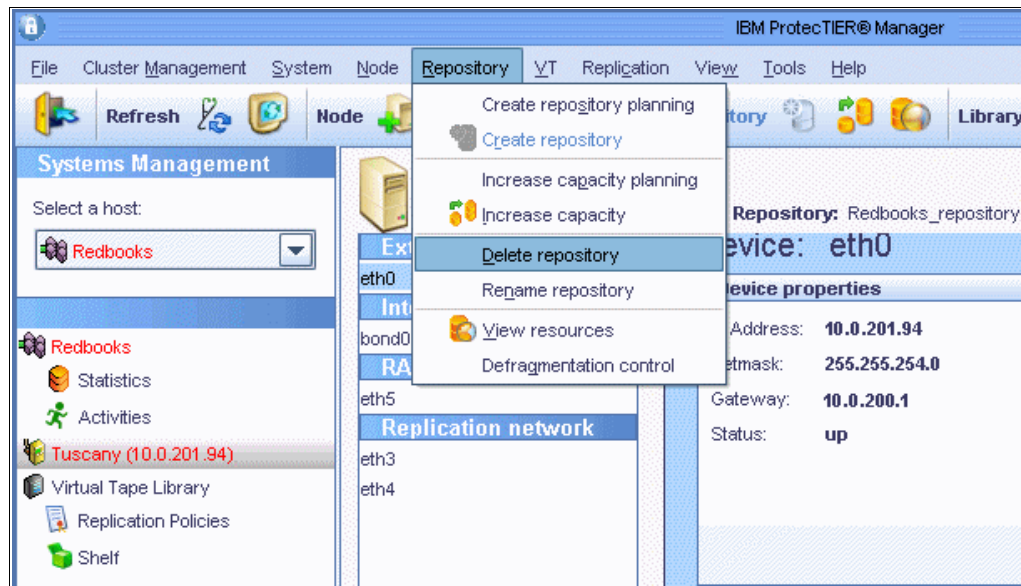


Figure 5-62 Delete ProtecTIER repository

3. A confirmation window opens. Click **Yes** (Figure 5-63).

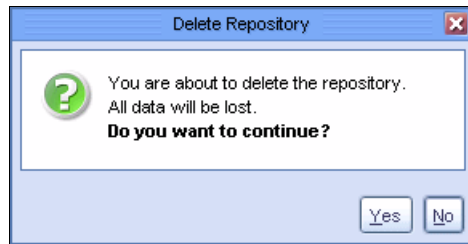


Figure 5-63 Delete the repository and continue

4. The Data Loss Confirmation window opens, as shown in Figure 5-64. In the field, enter "data loss" and click **OK**.

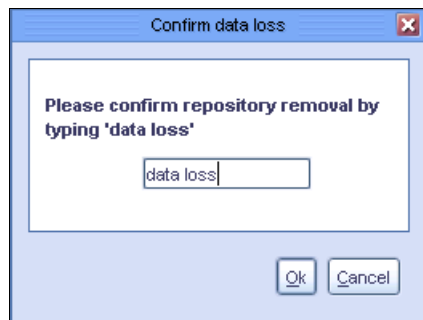


Figure 5-64 ProtecTIER data loss confirmation

5. The Confirm Operation window opens (Figure 5-65). Click **Yes**.

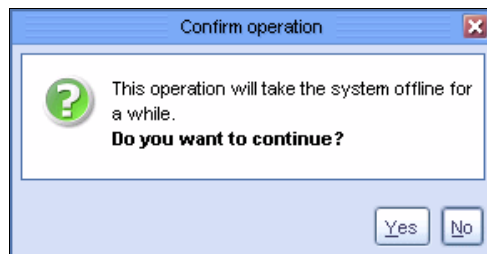


Figure 5-65 ProtecTIER data loss Confirm operation window

6. The ProtecTIER system temporarily goes offline to delete the repository (Figure 5-66).

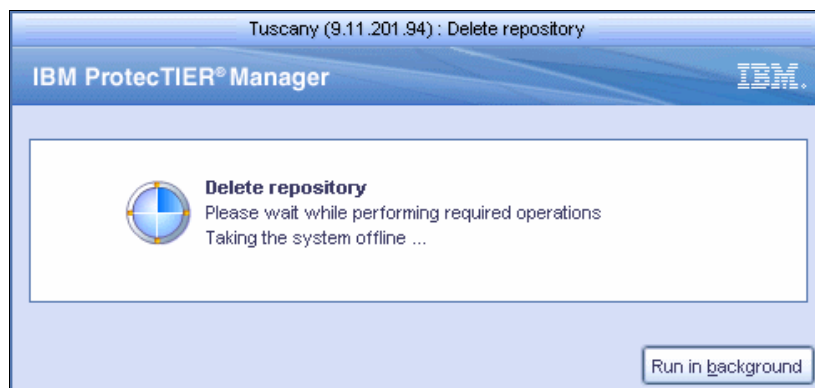


Figure 5-66 ProtecTIER Delete repository: System temporarily goes offline

## 5.6 Setting up the virtual library and cartridges

The ProtecTIER system enables you to create virtual tape libraries on which the backup application stores your data. These libraries and their components are part of the virtual tape service.

### 5.6.1 Creating libraries

A library can be created on a ProtecTIER system of either a one node cluster or a two-node cluster.

**Note:** Use the **Scan** button of the Port attributes pane to verify that the ports of the ProtecTIER system to which the virtual devices of the library are to be assigned are connected to the correct host. For more information, see Chapter 12, “Monitoring and reporting of the IBM System Storage TS7600 with ProtecTIER” on page 619.

Complete the following steps:

1. Log in to the system on which you want to add a library.
2. In the Systems pane, select a two-node cluster on which to add a library.
3. From the menu bar, select **VT** → **VT Library** → **Create new library**. The Create new library wizard Welcome window opens (Figure 5-67).

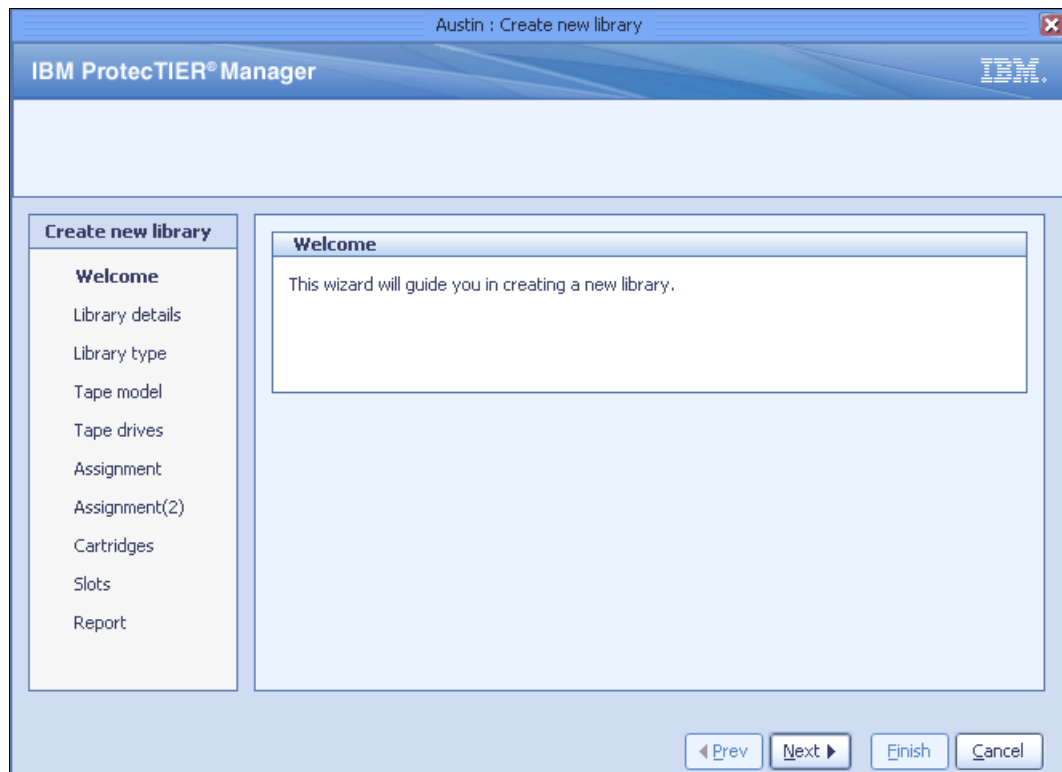


Figure 5-67 Create new library Welcome window

Click **Next**. The Library details window opens (Figure 5-68).

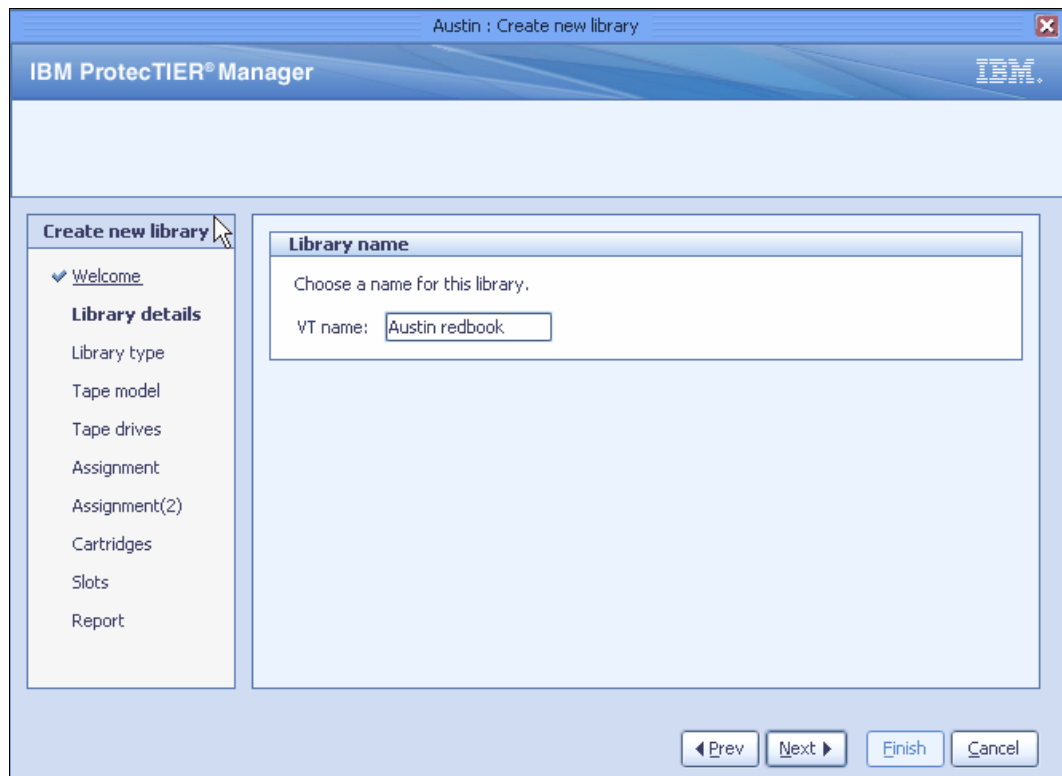


Figure 5-68 Create new library details window

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

5. Click **Next**. In the Library details pane, specify the type of library that you want to use for your application (Figure 5-69).

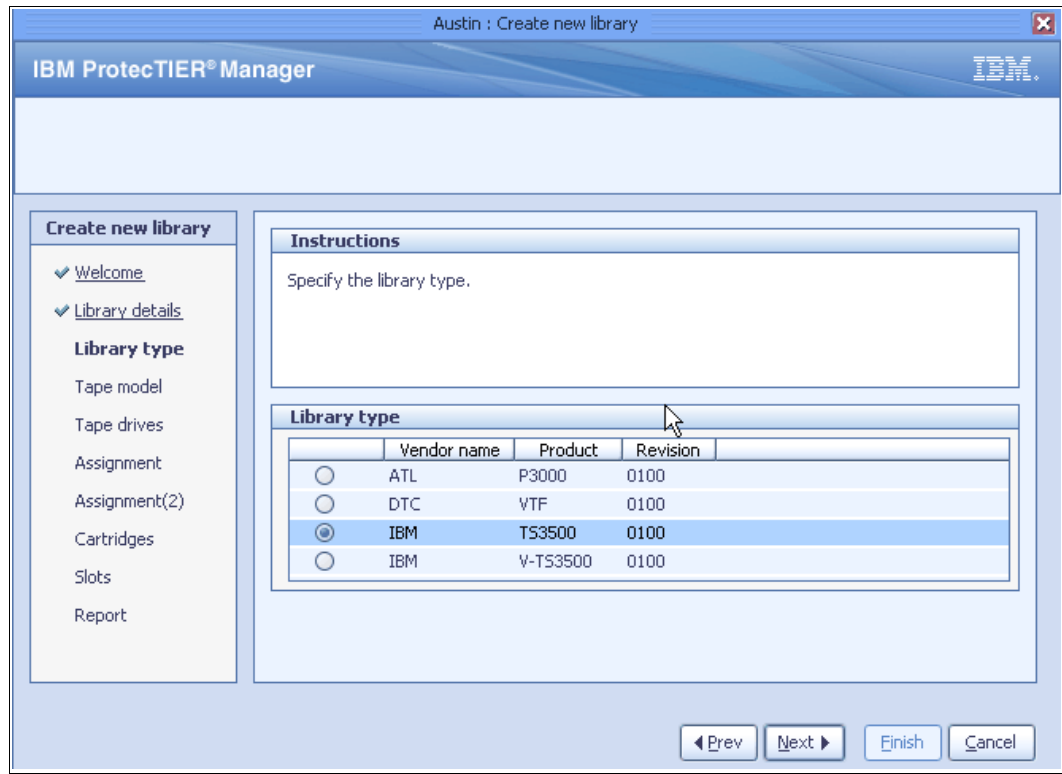


Figure 5-69 Create library details window

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

**Note:** 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 use the V-TS3500.
- ▶ If you are using Tivoli Storage Manager (TSM) software, you should use the TS3500.

6. Click **Next**. The Tape Model window opens. Select the tape drive model that you to use for your virtual library (Figure 5-70).

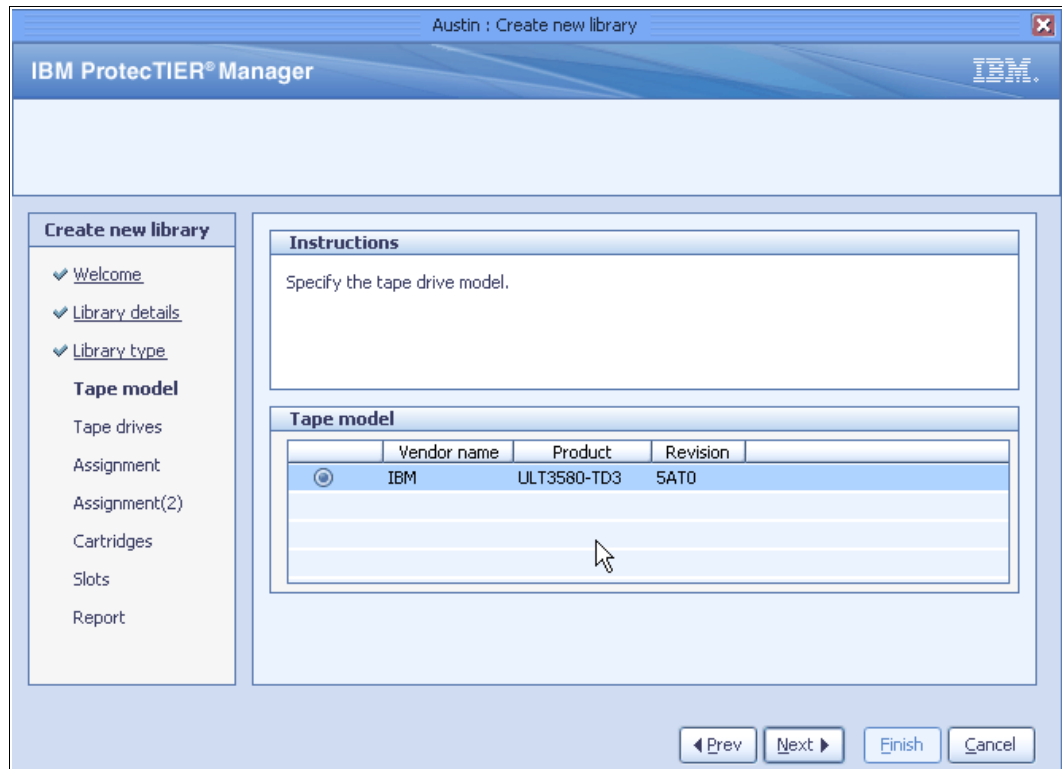


Figure 5-70 Create new library Tape Model window

- Click **Next**. The Tape drives window opens. In the Number of tape drives field for each node, enter the number of tape drives to assign to the node. To maximize load balancing, distribute tape drives across the nodes in a two-node cluster based on the relative power of the nodes (Figure 5-71).

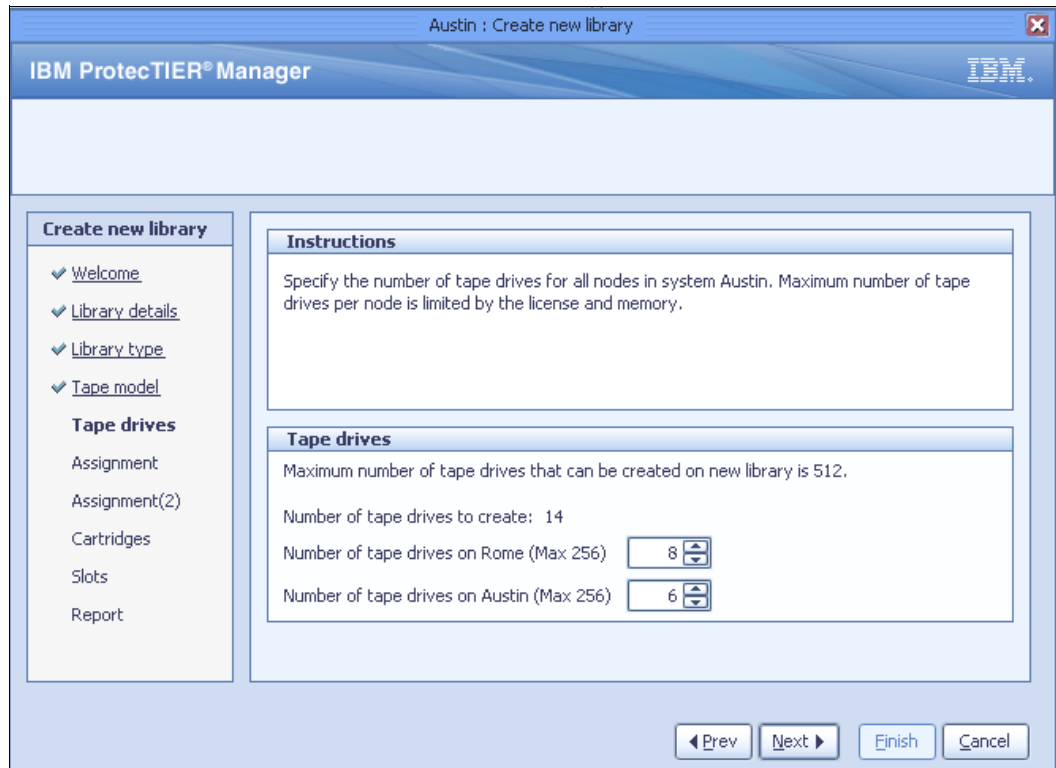


Figure 5-71 Create new library Tape Model window

In this example, we are creating eight drives on node Rome and six drives on node Austin.

**Note:** Check with your backup application administrator to ensure the number of drives and cartridges supported by your application. The value of the maximum number of tape drives per ProtecTIER node is 256, as shown in Table 5-3.

Table 5-3 Maximum number of libraries and drives allowed

	TS7610 Appliance	TS7650 Appliance	TS7650G
Maximum number of libraries	4	12	16
Maximum number of drives	256	256 (dual-node 512)	256 (dual-node 512)



Click **Next**. The Assignment window opens (Figure 5-72).

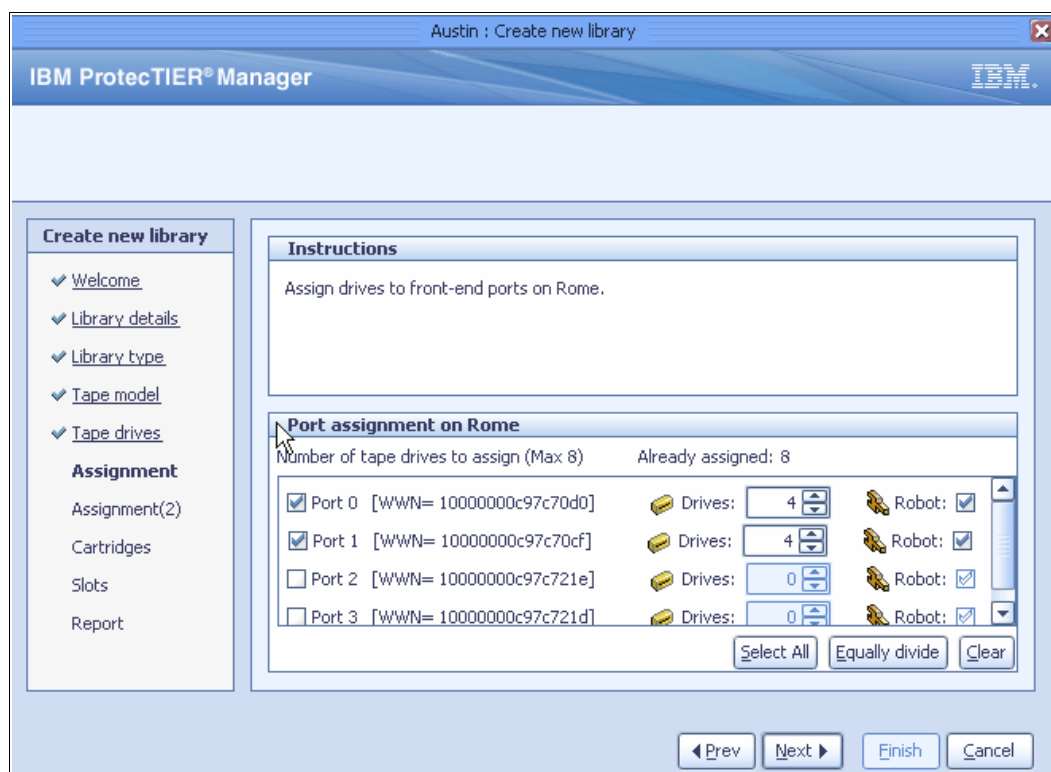


Figure 5-72 Create new library assignment window

Select or deselect the check boxes next to each port to define which of the node's ports are assigned virtual devices. In our example, we set up an IBM System Storage TS3500 server, and by default all the robots are selected and enabled. We deselected the robots for ports 2 and port 3. If you have chosen a library model other than an IBM model, the robots are not checked and only one 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.

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

8. Click **Next**. If a second node exists in your cluster, the Assignment (2) window opens. Follow the same steps as you did already for your first cluster.

**Note:** The backup application can only access 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 using the Scan button in the port attributes pane.

Click **Next**. The Cartridges window opens (Figure 5-73).

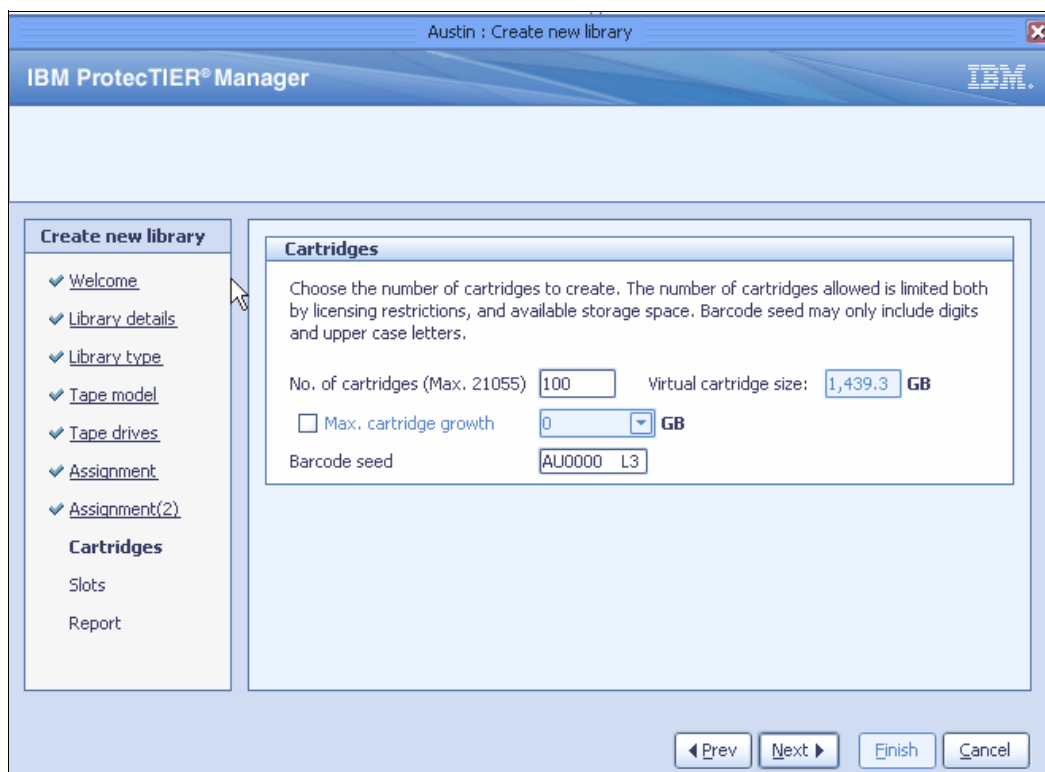


Figure 5-73 Create new library Cartridges window

9. In the No. of cartridges field, enter the number of cartridges that you want to have in the library.

In the Barcode seed field, enter a value for the barcode seed. The barcode seed is the barcode that is assigned to the first cartridge created. Every cartridge added after the first cartridge is assigned a barcode following the initial barcode seed.

**Note:** The barcode seed must contain only numbers and capital letters and be only six characters in length (for example, DS0006).

The Virtual cartridge size field automatically displays the maximum possible size for virtual cartridges for your system, based on the number of cartridges entered, the total amount of available storage space in your repository, and the current HyperFactor ratio.

The value of the maximum number of cartridges possible on a system depends on the amount of storage space available on your system.

Optionally, select the **Max. cartridge growth** check box. When selected, you can limit the maximum amount of nominal data that a cartridge can contain.

Click **Next**. The Slots window opens (Figure 5-74).

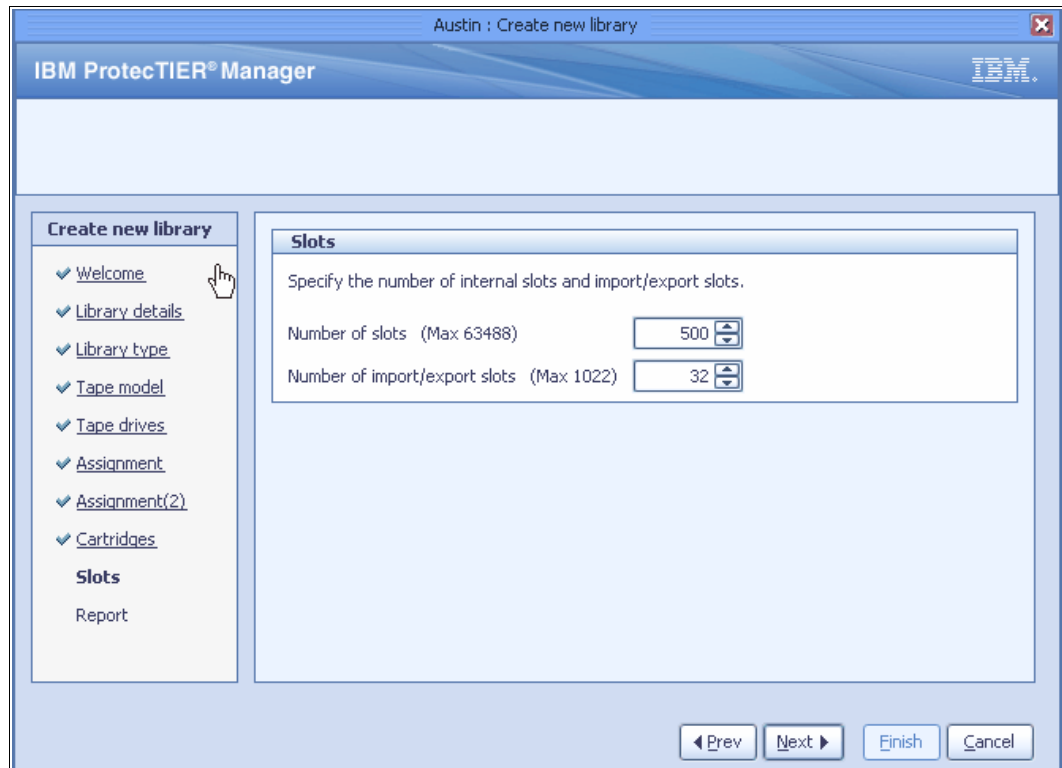


Figure 5-74 Creating new library Slots window

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

**Notes:** The number of cartridge slots must be equal or more than the number of cartridges that you are creating. You should create additional slots now, if you expect the number of cartridges to increase later.

Libraries attached to IBM i can have a maximum of 4096 positions where media can reside. (That is, the total of the # drives, # convenience IO slots + number of media slots + 1 (for the picker) must be 4096 or less.)

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 in the entire system is 1022.

10. Click **Next**.

11. The Create new library wizard closes and a Summary report window opens (Figure 5-75). Click **Finish**.

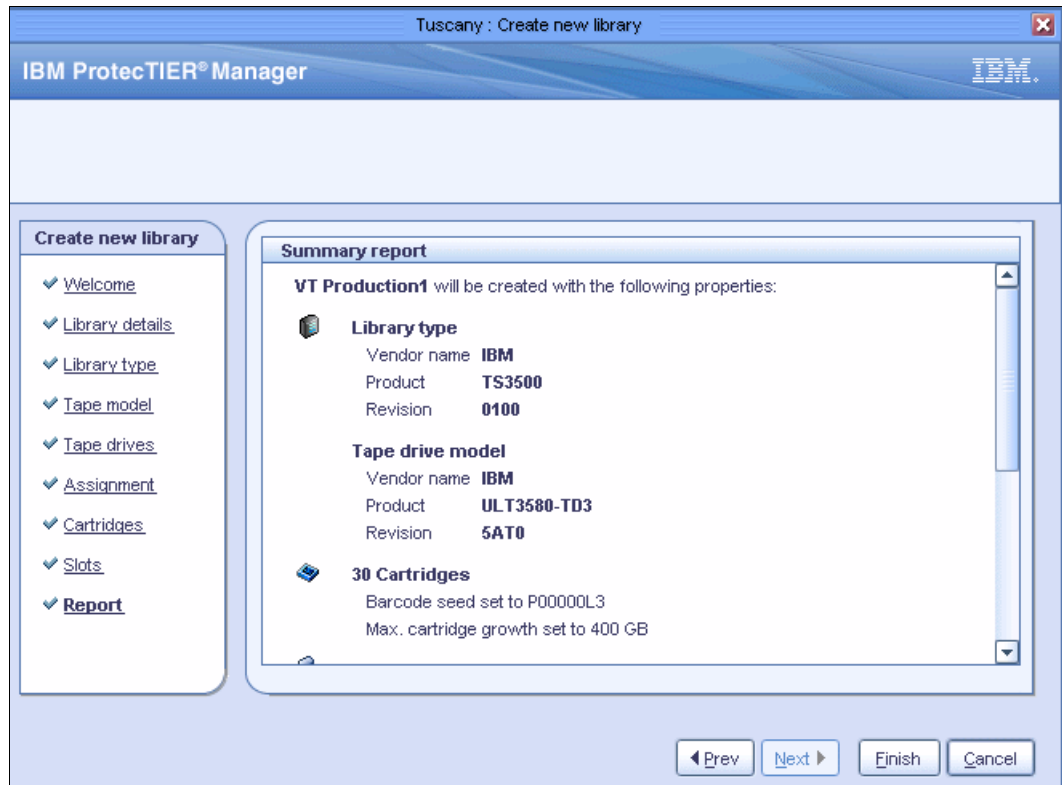


Figure 5-75 Summary report

12. The Confirm operation window opens (Figure 5-76). Click **Yes**. The ProtecTIER system temporarily goes offline to create the library.




Figure 5-76 Create new library Confirm operation window

13. The newly created library is displayed in the left pane and the right pane shows the details of the virtual library you have just created, as shown in Figure 5-77.

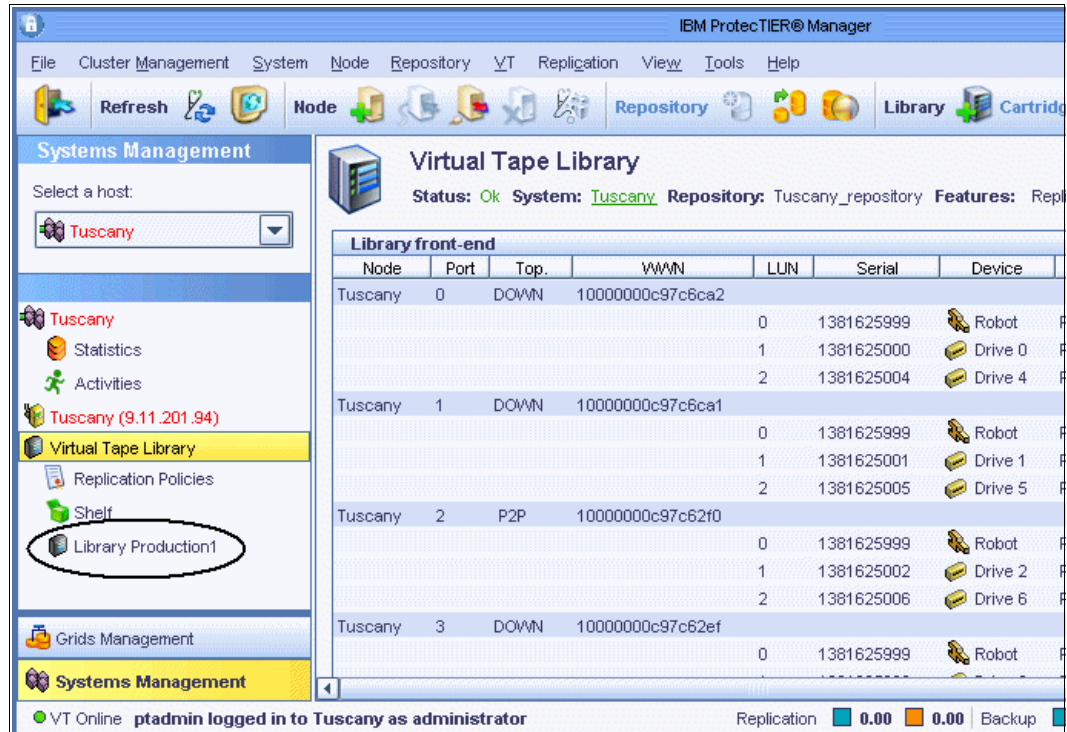


Figure 5-77 New virtual tape library created

See Figure 5-78 to view the libraries configuration in the right pane.

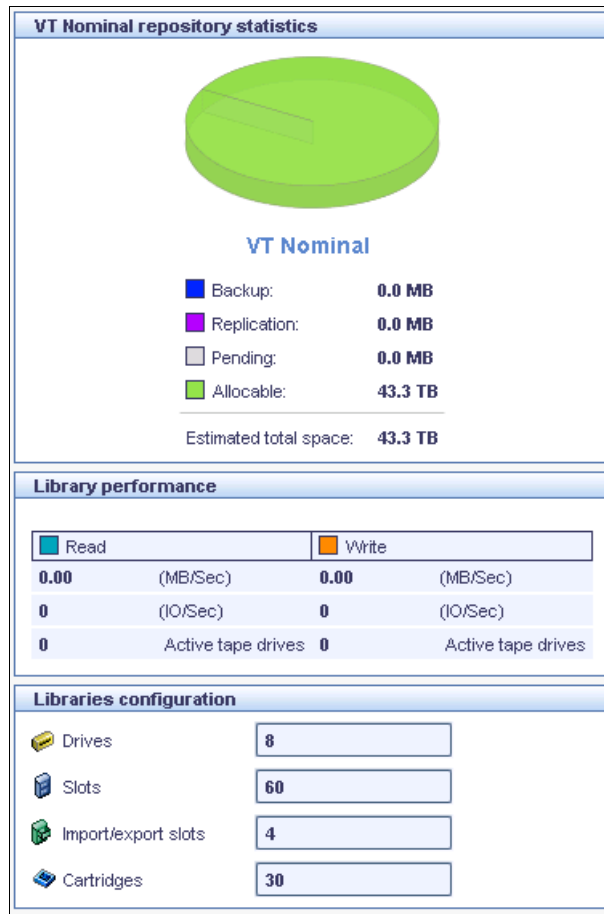


Figure 5-78 Libraries configuration on the right pane

For more information about managing virtual libraries and cartridges, refer to 10.3, “Managing virtual libraries and cartridges” on page 487.

## 5.7 Working with OpenStorage using ProtecTIER

This section describes the IBM ProtecTIER storage solution using OpenStorage (OST).

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

### 5.7.1 The OpenStorage operating environment

There are two major components that make up the OpenStorage operating environment and communicate through a TCP IP network:

- ▶ The storage server
- ▶ The plug-in

The storage server is an entity that runs on the ProtecTIER servers and uses the major internal functionality of the ProtecTIER platform (such as DHF, clustering, and replication).

The plug-in is a shared library (that is, a stateless software component) that resides on the NetBackup machine and is dynamically linked to the NetBackup application for data transfer to the ProtecTIER storage server emulation.

## 5.7.2 Installing the ProtecTIER storage appliance

Installing the ProtecTIER storage appliance and the appropriate software is generally done either in manufacturing, or is the responsibility of a trained ProtecTIER specialist.

Completing the ProtecTIER system setup tasks for new installations is a customer responsibility. The ProtecTIER storage appliance must meet the prerequisites to install and run the OST plug-in effectively.

After the storage appliance prerequisites have been met, you can move on to 5.7.3, “Configuring ProtecTIER to work with the OpenStorage environment” on page 241.

## 5.7.3 Configuring ProtecTIER to work with the OpenStorage environment

The following section describes how to configure the ProtecTIER system for use with OpenStorage using ProtecTIER Manager.

Before you begin configuring the ProtecTIER system, make sure that a repository has been created with OpenStorage as the backup interface.

**Note:** You can verify this configuration during repository creation by reviewing the Summary report of the Create repository wizard. The Backup interface field should appear as OST.

After a repository has been created, use the ProtecTIER Manager to create, configure, and monitor the storage server.

The storage server (STS) is a high-level component defined by the OST API. In simple terms, it is a “container” of logical storage units (LSUs) and provides the means to access the logical storage units and their contents. Currently, only one STS can be defined for each OST storage appliance. For more information, refer to Chapter 4, “Hardware planning for IBM System Storage ProtecTIER” on page 121 and Chapter 9, “IBM System Storage ProtecTIER with Symantec OpenStorage” on page 433.

## 5.7.4 Configuring a storage server

This section describes how to add and configure a storage server.

Complete the following steps to add the storage server (STS) and then define the logon credentials. NetBackup uses these credentials so that the media server can log in to the storage server for storage access. (See the relevant Symantec NetBackup OpenStorage documentation for more information about OpenStorage server credentials.) Adding a storage server is a one-time operation.

1. Log in to ProtecTIER Manager (Figure 5-79). You should see that it is an OpenStorage system.

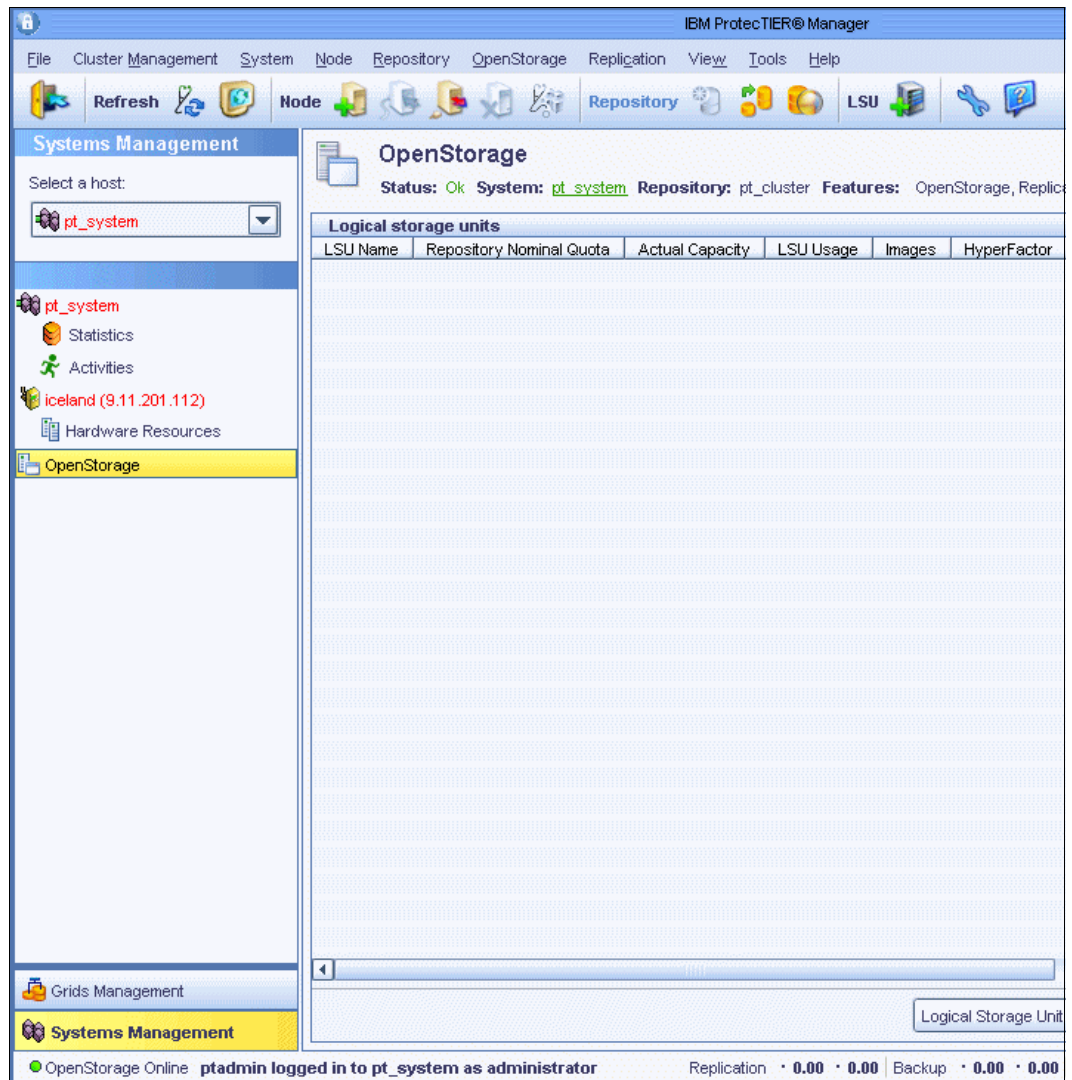


Figure 5-79 Log in to the OpenStorage system



- From the OpenStorage menu, select **Storage server** → **Add Storage server** (Figure 5-80). The STS configuration window opens.

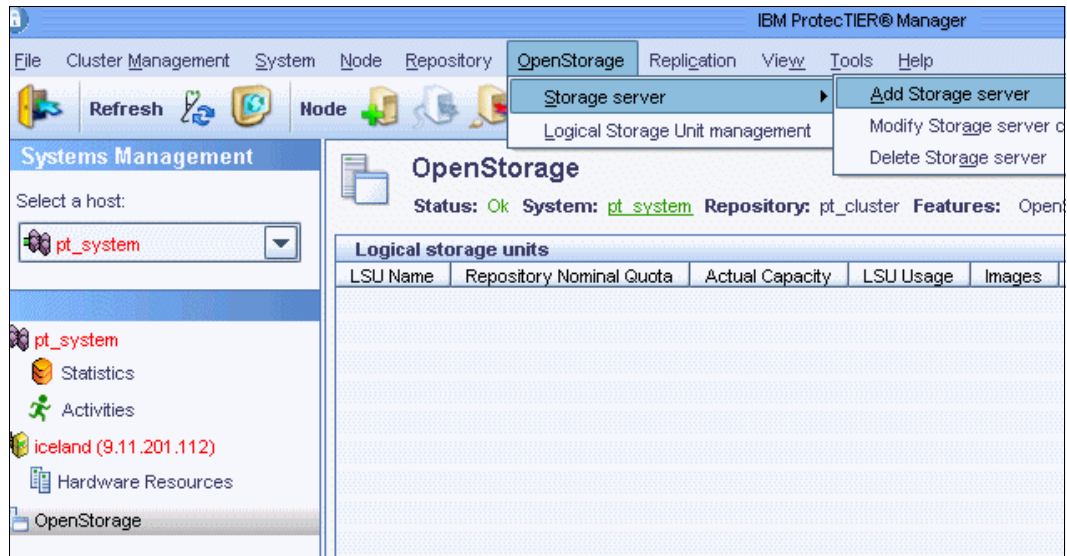


Figure 5-80 Add Storage Server

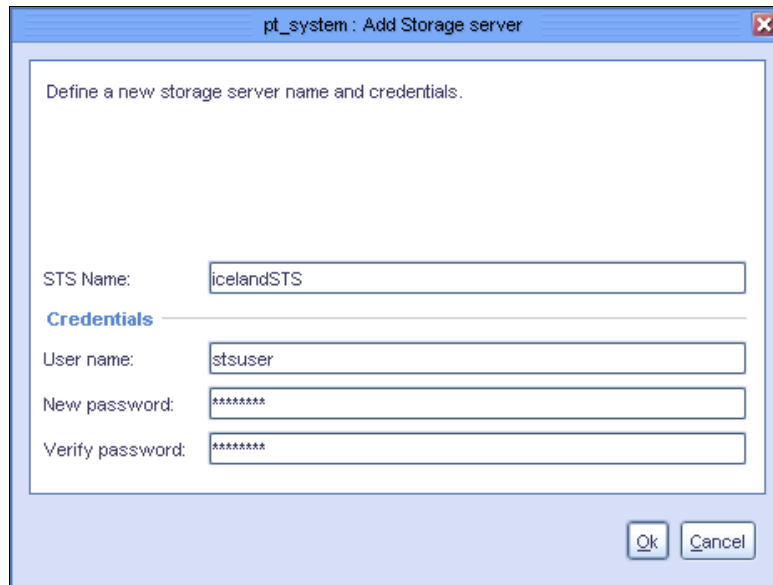
**Note:** Refer to Table 5-4 for the field specifications when defining the logon credentials.

Table 5-4 Logon credentials field specifics

Field name	Field length	Field type (alphanumeric or special characters)
STS Name	4 - 16 characters	alphanumeric, period (.), or underscore (_).
User name	4 - 16 characters	ASCII.
New Password	4 - 16 characters	All alphanumeric and special characters. The use of special characters may be unrecognized by NBU. Therefore, if special characters are used (for example, space, &, -, and others), use quotation marks (" ") before and after the string.

3. Enter a unique name for the new storage server in the STS Name field (Figure 5-81).

**Note:** The STS unique name cannot be modified after the STS has been created.



pt\_system : Add Storage server

Define a new storage server name and credentials.

STS Name: icelandSTS

**Credentials**

User name: stuser

New password: \*\*\*\*\*

Verify password: \*\*\*\*\*

Ok Cancel

Figure 5-81 Add Storage server window

4. Enter a user name in the User name field. Enter a password in the New Password field and re-enter the password in the Verify password field. Click **OK** to complete the task.
5. The STS is added and configured on the ProtecTIER system and can be monitored through ProtecTIER Manager. If you would like to modify the STS user name or password, go to 5.7.5, “Modifying the storage server credentials” on page 244 or continue to 5.7.7, “Configuring a logical storage unit” on page 246.

### 5.7.5 Modifying the storage server credentials

This section describes how to modify the credentials of the selected storage server using ProtecTIER Manager. The STS unique name cannot be modified after the STS has been created.

Complete the following steps:

1. Log in to ProtecTIER Manager.

- From the OpenStorage menu, select **Storage server** → **Modify Storage server credentials** (Figure 5-82).

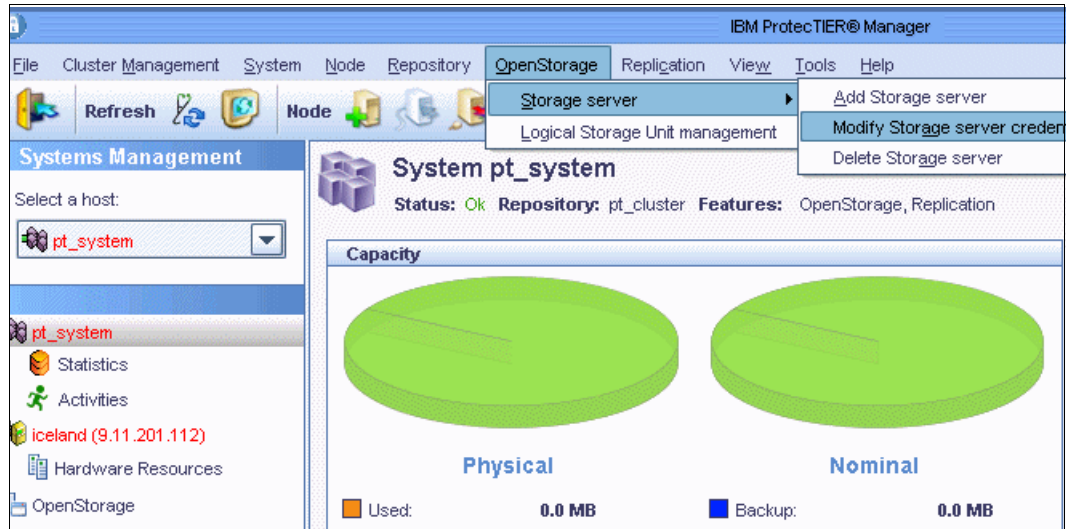


Figure 5-82 Select Modify Storage server credentials

- The STS modify credentials window opens (Figure 5-83). If you are modifying the user name, enter a new user name in the User name field. If you are modifying the password, enter a new password in the New password field and re-enter the password in the Verify password field.

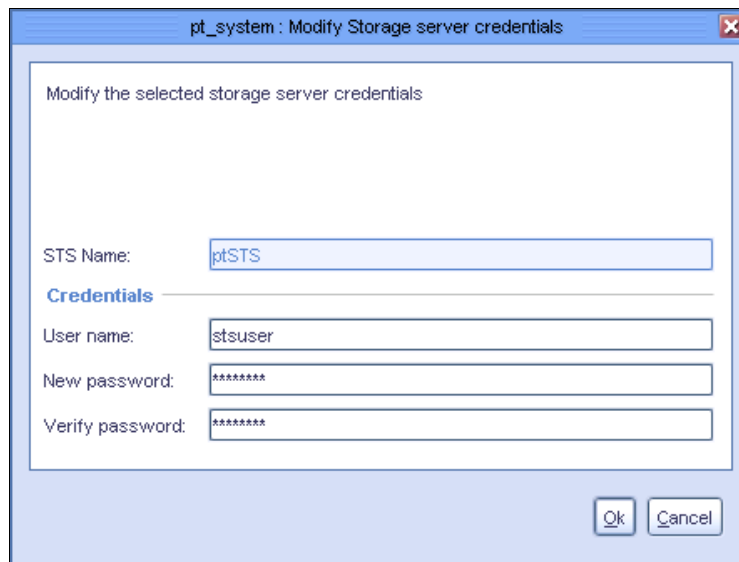


Figure 5-83 Modify Storage server credentials

- Click **OK** to complete the task.

### 5.7.6 Deleting a storage server

This task describes how to delete a selected storage server. An STS can be deleted at any time, but this action might potentially cause data loss.

**Note:** Before deleting an OpenStorage component from ProtecTIER, you must first delete the component from NetBackup. If the OpenStorage component is not deleted first from NetBackup, the delete action will be blocked and you will be unable to delete the component from NetBackup at a later time. Refer to the relevant Symantec NetBackup OpenStorage documentation for more information.

Complete the following steps:

1. Log in to the ProtecTIER Manager.
2. From the OpenStorage menu, select **Storage server** → **Delete Storage server** (Figure 5-84).

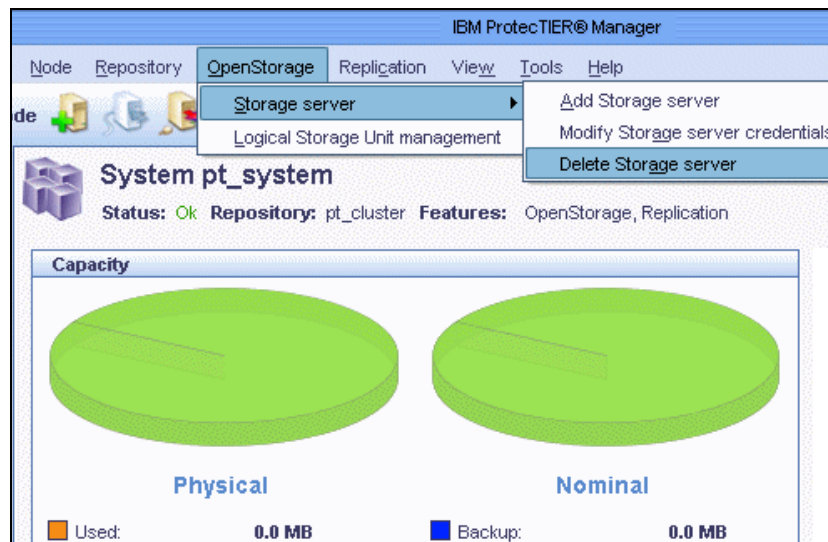


Figure 5-84 Delete Storage server window

3. Enter "data loss" to confirm the operation (Figure 5-85).

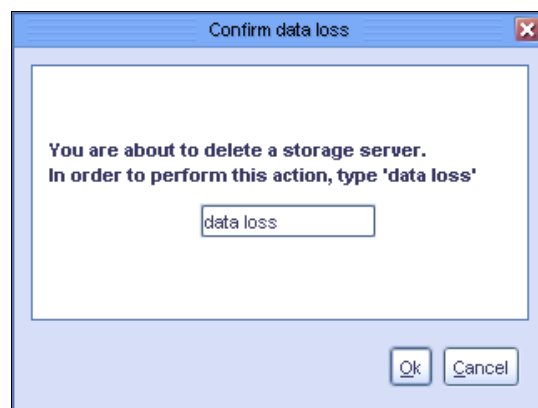


Figure 5-85 Confirm data loss window

## 5.7.7 Configuring a logical storage unit

The following section describes how to configure a logical storage unit (LSU) on an STS.

Configuring LSUs on a storage server divides the appliance into one or more logical units of space. An LSU, like an STS, is also defined by the OST API and is a “container” of storage and images (which consume storage). Up to 256 LSUs can be defined per STS and are identified by a name string that is unique within the storage server.

In ProtecTIER for OST, an LSU's storage properties are defined in nominal terms. This means that the LSU's storage capacity is defined as a nominal percentage of the repository's overall nominal capacity, taking into consideration the configured percentages of the other LSUs that all share the repository's physical storage.

**Note:** The overall capacity of all LSU's together can be less than or equal to the repository's nominal capacity.

### 5.7.8 Adding an logical storage unit

Complete the following steps to add an LSU to an STS.

1. From ProtecTIER Manager, select **OpenStorage** → **Logical Storage Unit management** (Figure 5-86). The Logical Storage Unit management window opens:

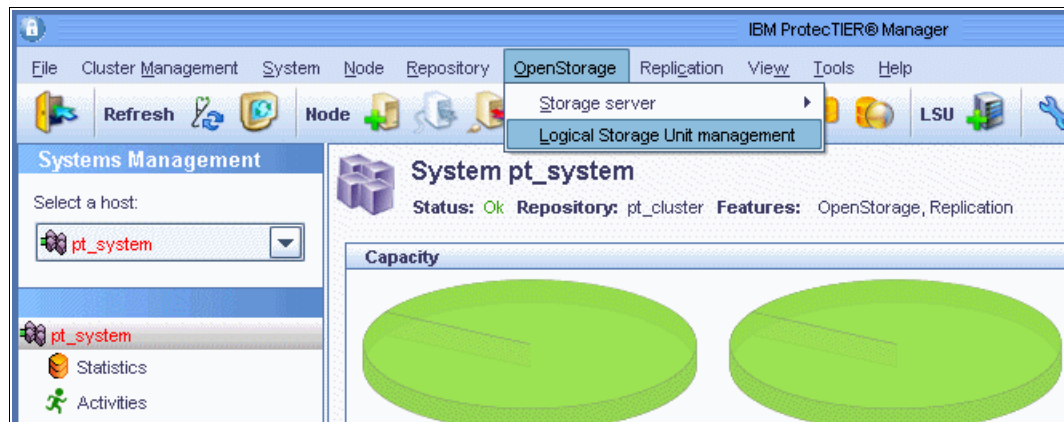


Figure 5-86 Add LSU to STS

2. Click **Add** to configure a new LSU on the storage server (Figure 5-87).

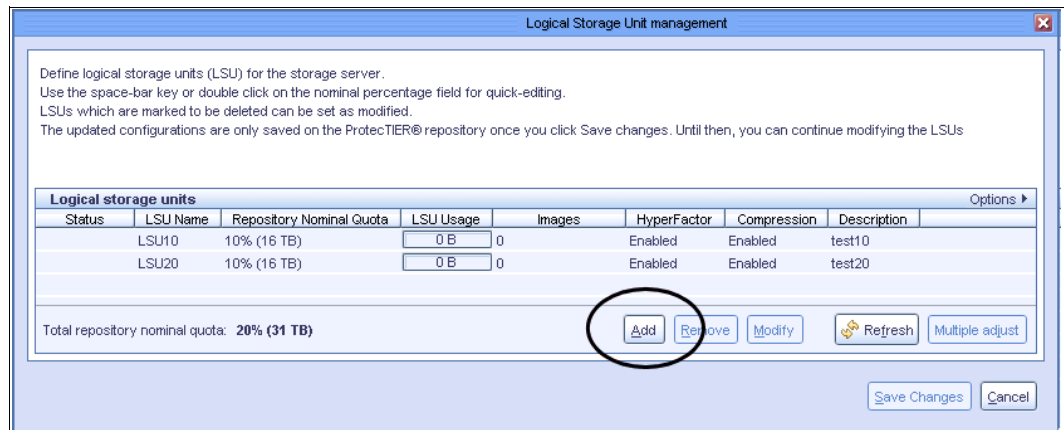


Figure 5-87 Logical Storage Unit management

The Adding an LSU window opens (Figure 5-88).

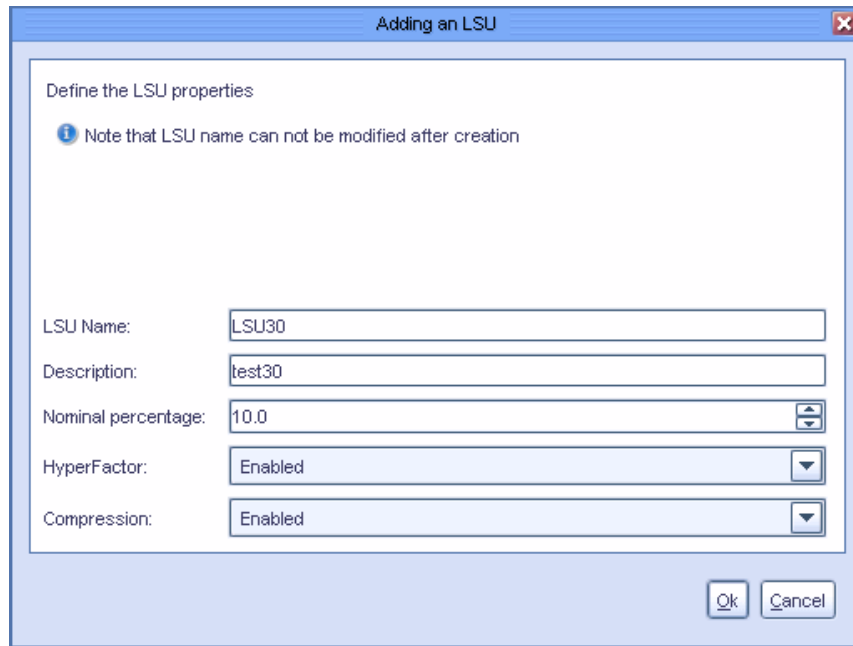


Figure 5-88 Adding an LSU window

3. Enter a description of the LSU in the Description field. Enter the size of the LSU as the nominal percentage (up to two precision digits, for example, 5.55) of the repository in the Nominal percentage field.

**Note:** Refer to Table 5-5 on page 248 for the field specifics when defining an LSU.

Table 5-5 LSU field specifics

Field name	Field length	Field type (alphanumeric and special characters)
LSU name	4 - 16 characters	alphanumeric, period (.), and underscore (_)
Description	0- 1023 characters	All alphanumeric and special characters

4. Select the data factoring mode from the HyperFactor drop-down menu.
5. Select either **On** or **Off** in the Compression field to enable or disable data compression. Click **OK**.

6. Click **Save** to save the LSU properties (Figure 5-89).

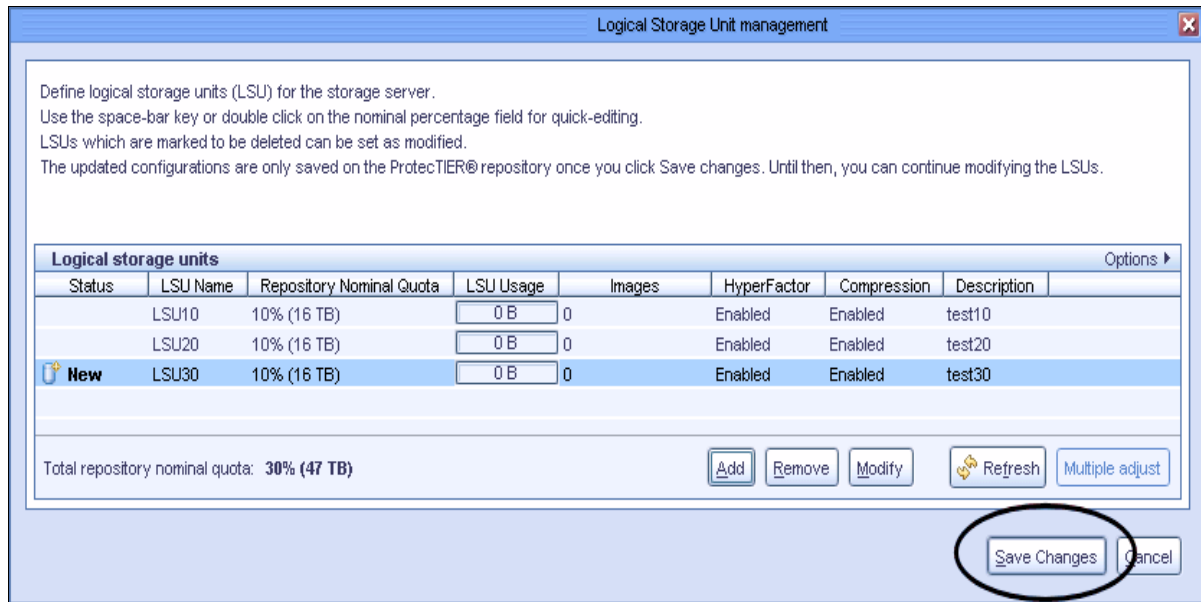


Figure 5-89 Save the changes on the Logical Storage Unit management window

### 5.7.9 Modifying an logical storage unit configuration

Complete the following steps to modify an LSU on an STS:

1. From ProtecTIER Manager, select **OpenStorage** → **Logical Storage Unit management**. The Logical Storage Unit management window opens (Figure 5-90).

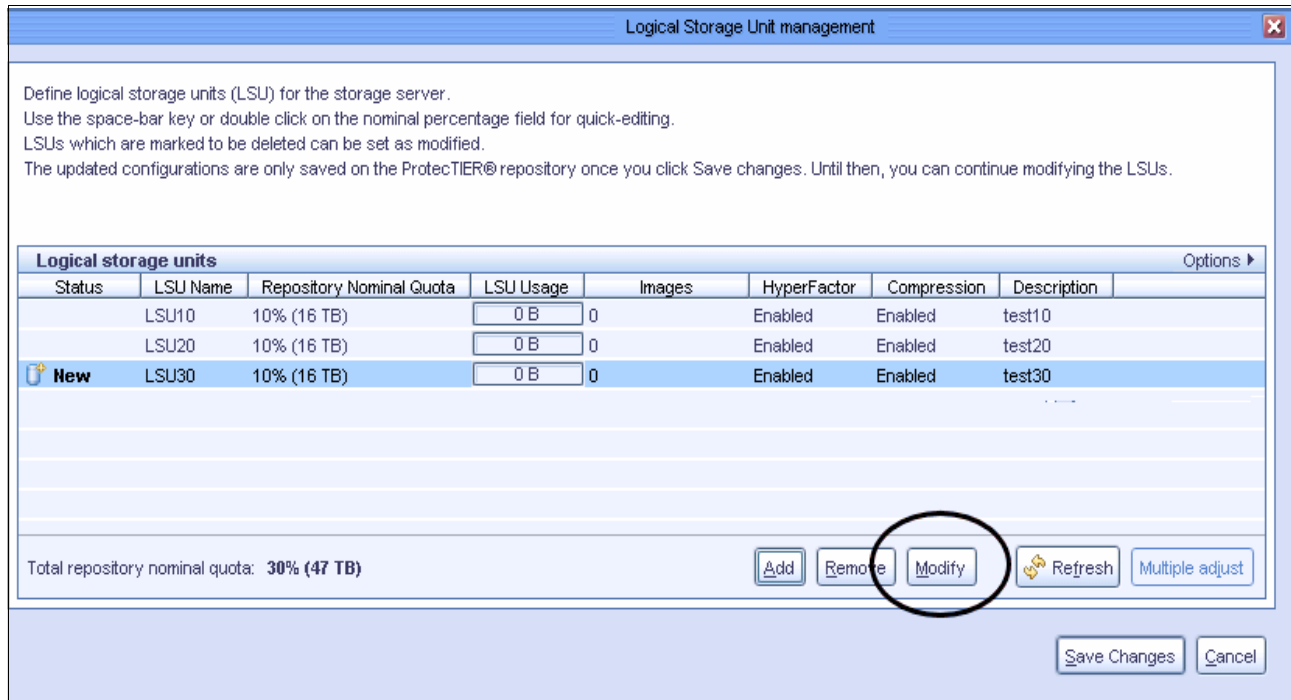


Figure 5-90 Logical Storage Unit management

2. Highlight one of the LSUs to be modified by clicking once on the LSU name. Click **Modify** to change the LSU properties previously defined on the storage server. The Modifying an LSU window opens (Figure 5-91).

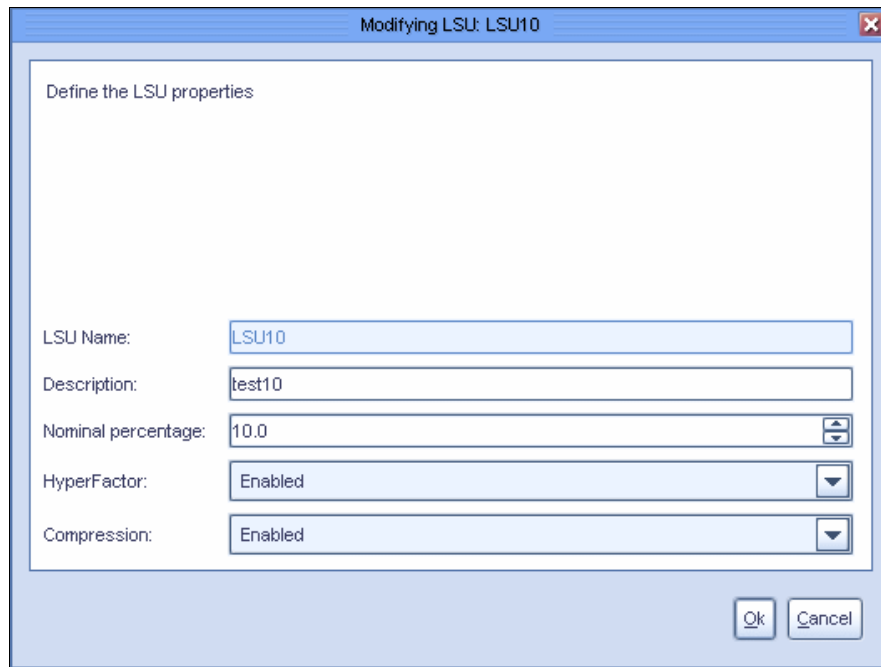


Figure 5-91 Modify LSU window

3. Modify either one or all of the following items:
  - Enter a new description (up to 1023 characters) to describe the LSU in the Description field.
  - Change the size of the LSU as the nominal percentage (up to two precision digits, for example, 5.55) of the repository in the Nominal percentage field.
  - Modify the factoring mode option from the HyperFactor drop-down menu:
    - Enabled to enable hyperfactoring
    - Disabled to disable hyperfactoring
    - Baseline for no deduplication, but an index is built
  - Modify the data compression option by selecting either **On** or **Off** in the Compression field to enable or disable data compression.
4. Click **OK** to save the LSU properties

**Note:** After an LSU has been created and saved, the LSU name can no longer be modified.

### 5.7.10 Managing the logical storage unit configuration

Use ProtectTIER Manager to manage the LSU configuration. The Logical Storage Unit management window displays the LSU configuration properties. Use the spacebar or double-click the **Repository Nominal Quota** or **LSU usage** fields for “quick editing” of the LSU storage capacity.



Complete the following steps:

1. Select **Logical Storage Unit management** from the OpenStorage menu. The Logical Storage Unit management window opens and (Figure 5-92) displays the following fields:

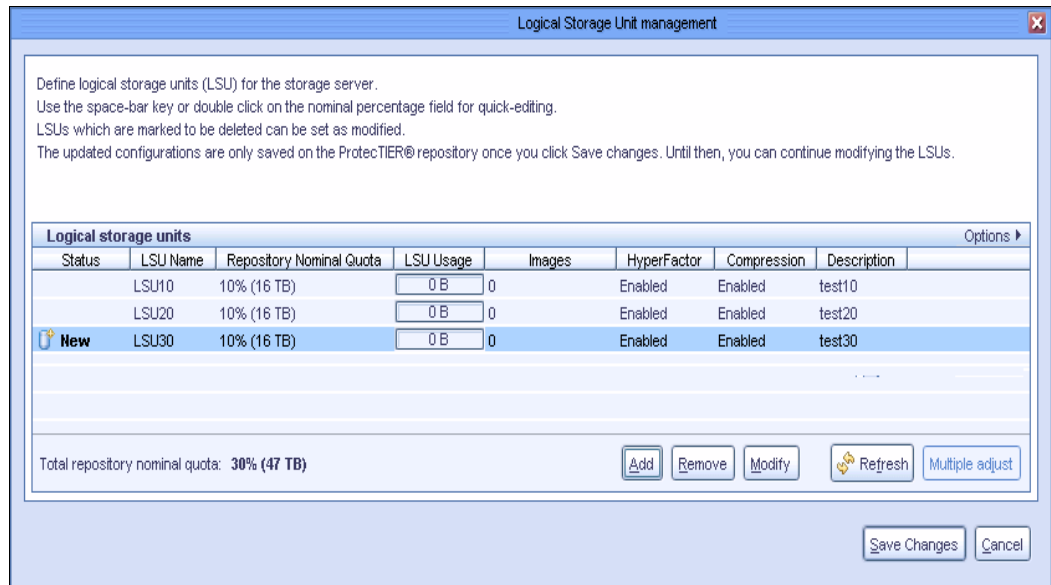


Figure 5-92 Logical Storage Unit management window

- Status: Displays whether the LSU is new, or if it has been modified or deleted.
  - LSU name: The unique name defined for the LSU.
  - Repository nominal quota: The LSU's nominal percentage of the repository's overall nominal capacity.
  - LSU usage: The actual amount (in MB) used from the LSU's nominal percentage.
  - Images: Number of images contained on an LSU.
  - Hyperfactor: The data factoring mode status.
  - Compression: Data compression is either enabled (on) or disabled (off).
  - Description: Describes the LSU.
2. Click **Refresh** to update the LSU usage column.
  3. Click **Multiple adjust** to equally divide the repository nominal quota between all of the selected LSUs (that is, Available percentage / Number of selected LSUs), or set all of the selected LSUs with a selected percentage value.
  4. Click **Save Changes** to save the updated configurations, or continue modifying the LSUs.

Now you have completed the initial setup of OST on ProtecTIER. To continue on the OST implementation with Symantec NetBackup, go to 9.3, “The ProtecTIER OpenStorage plug-in” on page 446.

## 5.8 Setting up native replication

This section provides the information about setting up native replication on ProtecTIER.

Native replication lets you replicate data objects between ProtecTIER repositories. In order for a logical set of repositories to replicate from one to another, you must create a replication grid. The replication grid is remotely created and managed by the Replication Manager.

Since its release date in September 2009, ProtecTIER V2.3 code provides one-to-one replication capabilities. Version 2.4 (released May 2010) provides many-to-one replication capabilities, with up to 12 primary sites (spokes) and one central DR location (hub). Whether deploying replication in a one-to-one configuration or a many-to-one configuration, you should use the latest Version 2.5 code to take advantage of many new replication features and enhancements.

Replication Manager is installed together with the ProtecTIER code.

This section describes how to set up your native replication using the ProtecTIER Replication Manager for VTL and OST.

The ProtecTIER Replication Manager controls the following items:

1. Managing the repositories in the replication grid to which you can replicated. Two repositories are in a replication relationship:
  - A source repository
  - A target repository
2. Maintaining the IP addresses of all repositories.
3. Updating repositories leaving and joining the grid.
4. High-level monitoring and statistics of traffic in the replication grids. ProtecTIER Manager connects to the ProtecTIER Replication Manager using the IP address of the ProtecTIER Replication Manager.

Setting up your ProtecTIER Replication Manager is done in three steps:

1. Activating the Replication Manager
2. Adding the nodes to the grid
3. Setting up a replication policy (application group)

**Note:** ProtecTIER Replication Manager is a software component that may be installed on a dedicated host through a special request for IBM support.

### 5.8.1 Replication throughput control

The incoming data stream rate control is not enforced at the remote site, so pacing of the replication process is up to the primary source site. The synchronization of available replication time windows between both sites is up to the user.

The user can configure the ProtecTIER system to limit the replication throughput regardless of backup ingest activity. Because only the primary site ProtecTIER can accept local backups, this feature is effective only at the primary source site. This means that the user cannot limit the throughput at the secondary site, which is the replication destination, but can affect it indirectly by limiting the throughput of the source system.

## 5.8.2 ProtecTIER Replication Manager

The ProtecTIER Replication Manager is a software module that manages the replication grid's configuration in the user's organization. The ProtecTIER Replication Manager module will reside on one of the active ProtecTIER servers in the grid. ProtecTIER Manager connects to the ProtecTIER Replication Manager using the IP address of the ProtecTIER Replication Manager node.

The ProtecTIER Replication Manager module manages the repositories in the replication grid by:

- ▶ Maintaining the IP addresses of all repositories
- ▶ Updating repositories leaving and joining the grid
- ▶ High-level monitoring and statistics of traffic in the replication grids

Only a single ProtecTIER Replication Manager instance within an organization is supported. However, multiple instances of ProtecTIER Replication Manager for multiple stand-alone networks (their repositories can never be connected to repositories in other networks) may be created in a single instance per stand-alone network fashion. For the ProtecTIER Replication Manager module to be able to manage a grid with multiple repositories, it must have access and communicate with them all on its network.

After a single ProtecTIER Replication Manager instance or a number of independent ProtecTIER Replication Managers are set up, a repository cannot be moved from one ProtecTIER Replication Manager instance to another. After a repository has joined a ProtecTIER Replication Manager, it will always belong to it even if it is removed (or deleted) from the grid.

Any ProtecTIER Replication Manager software module running on a ProtecTIER node is limited to a single grid with up to 16 active repositories (eight replication pairs) on that replication grid. For managing more than one grid, the ProtecTIER Replication Manager software must be installed on a separate, dedicated server, through an RPQ approval process.

Figure 5-93 shows a typical replication network deployment topology for a replication pair. In this example, the replication pair consists of a two-node cluster in each location. The diagram shows how to use the two available network ports that each node has and connect them to two separate subnets for redundancy.

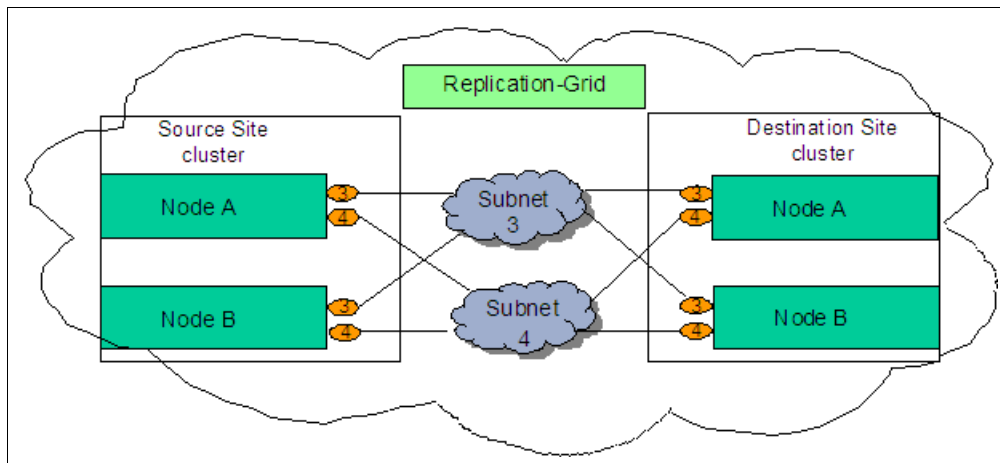


Figure 5-93 Replication grid topology example: Replication pair with two-node cluster systems

### 5.8.3 Adding replication to an existing production system

This section discusses the options for upgrading an existing TS7650 or TS7650G server to a system that will use the replication feature to send data to a remote, secondary (DR) TS7650 or TS7650G ProtecTIER system. The specific user scenario will dictate the correct approach to upgrading the primary system and deploying the target system.

An important detail to keep in mind here during the planning process is that all of the nominal data associated with the initial cartridges must physically be replicated (there will be no deduplication benefit regarding the bandwidth during these initial replication jobs). Also, the full contents of each cartridge associated with the replication policy must be replicated, even if only a fraction of the cartridge was used in the latest backup job. For example, cartridge AB1234 is used during the first backup window following the deployment of replication, and 10 GB is written on it. However, it also holds 190 GB from previous backups. All 200 GB must be replicated during the first replication window for that cartridge. For this reason, a new user of ProtecTIER replication should start with new cartridges for all the jobs that are associated with the newly created replication policies so that no previously backed-up data will need to be physically replicated as new data is appended to an existing cartridge.

A deployment of a second ProtecTIER server at a secondary (DR) site has a significant impact on the planning cycle. The first replication jobs will consume much more bandwidth than required after deduplication takes effect. So when preparing for deployment, the field engineer must help the user's team plan and prepare the infrastructure and resources that are temporarily needed to support this configuration. The plan must take into account the amount of physical data to be replicated, the amount of dedicated bandwidth, and the extra time needed for these first several replication runs. The user may need to allow the first replication job to complete before the next backup activity starts.

#### Before deploying the secondary (DR) system

The first step prior to deploying the target IBM System Storage TS7600 with ProtecTIER system is to confirm the bandwidth needed (or a conservative estimate of it) with the user. Work with the user's network team to obtain IP addresses, port connections, and additional network cabling, if needed. Verify that the network connections are clean and provide adequate speed. Use the Network Validation Utility available with the IBM System Storage TS7600 with ProtecTIER system to verify the amount of bandwidth and the quality of the available user's network. See *IBM System Storage TS7600 with ProtecTIER Users Guide*, GC53-1156 for more information about this utility.

The second step is to determine the approach to repository synchronizing between the local and the new secondary site systems that the user will take. As mentioned above, the key to a successful deployment here is a gradual replication load approach. The most efficient way to deploy replication with an existing operational IBM System Storage TS7600 with ProtecTIER system requires a load-throttle that introduces workload into the replication operation (policies) while staying under the available network bandwidth ceiling assigned between the two sites, until a steady state is reached.

This would be the only feasible and practical method to bring the secondary system to the point that it will be fully synchronized with the primary system while using the available network bandwidth that was planned for the steady state after the deduplication factor takes effect. For example, if the first backup set to be replicated is 6 TB (full nominal) and the replication window is 10 hours, then the amount of bandwidth that the user requires to allot will be  $6 \text{ TB}/10 \text{ hours} = 600 \text{ GB/hour}$  for the replication to be completed in the 10 hour window.

Later on, after the target and the primary systems are completely in sync, the deduplication factor takes effect, and assuming a 10% data change rate for that same job, the network bandwidth required will be:

$$[(6000 \text{ GB} * 10\%) + (0.5\% * 6000 \text{ GB})] / 10 \text{ hours} = 63 \text{ GB/hour}$$

This means that for the initial period of time the user must allot enough network bandwidth to account for the full nominal size of the data to be replicated. Thus, creating replication policies with a gradual increase approach should be done to stay within the available network bandwidth boundaries and within the replication window time frame.

## Upgrading the existing system

After the sync approach has been chosen and planned for, your IBM representative will perform the upgrade. Follow these step-by-step instructions, which describe how to upgrade an existing ProtecTIER system to Version 2.5 or later and then begin replicating to a target system:

1. Install the target IBM System Storage TS7600 with ProtecTIER system, either locally for the initial synchronization or at the target location (depending on the chosen approach and available network bandwidth).
2. Upgrade the primary ProtecTIER servers to Version 2.5 or later.
3. Add an additional GigEx2 Card, if the existing nodes are 3958-DD1.
4. The installation wizard confirm that the primary system has enough metadata space (all appliances are already configured with it). If not, it will add it automatically. After this task has been done, it should be possible to complete all other activity while the primary system is still in use.
5. Install ProtecTIER Replication Manager on one of the IBM System Storage TS7600 with ProtecTIER systems, either the primary or the remote one. If an existing ProtecTIER Replication Manager is installed in the user's environment already from a previous deployment, it may be used and there is no need to install a new one. It should be installed at the remote site.
6. Create a new replication grid and add the primary IBM System Storage TS7600 system and the target IBM System Storage TS7600 systems to this grid. If there is already a grid for another system in the environment, the existing grid should be used rather than creating a new grid.
7. *Pair* the primary IBM System Storage TS7600 with ProtecTIER system with the remote IBM System Storage TS7600 with ProtecTIER system.
8. Create the replication policy. This step defines which of the tapes that were previously created must be replicated. In most cases, the user will probably begin replicating new tapes, versus all the tapes that were previously written.

In the initial installation of the primary system, planning for the use of separate tape pools (or virtual libraries) for replication will be beneficial when it comes time to upgrade the system to support the replication feature (for example, one tape pool for tapes that do not need to be replicated and another tape pool for those that do need to be replicated). This setup makes the creation of the replication policies easier.

If the *visibility control switch* function will be used, the user should create scripts to automate the ejection of virtual tapes identified as those to be replicated. This must be done for both local and remote systems.

Consult Chapter 2, "Setting up ProtecTIER", in *IBM System Storage ProtecTIER User Guide for Enterprise Edition and Appliance Edition*, GC53-1156 for more details.

**Note:** During the upgrade process, before the replication operation commences, and if at all feasible (that is, both the primary and the remote system are installed in their locations prior to starting the sync process), run the ProtecTIER Replication Network Performance Validation Utility, which is covered in *IBM System Storage TS7600 with ProtecTIER Users Guide*, GC53-1156.

## 5.8.4 Planning for a new installation

This section outlines the process of installing new ProtecTIER Gateways or Appliances at two locations (source and target) and configuring replication between them. Prior to the installation, capacity, performance, and bandwidth planning must be complete. Refer to 3.5.1, “Replication for OST” on page 102 for more details.

The most efficient way to deploy two new IBM System Storage TS7600 with ProtecTIER systems using replication between them requires a load-throttle that introduces workload into the replication operation (policies) while staying under the available network bandwidth ceiling assigned by the user between the two sites, until a steady state is reached. This setup is the only feasible and practical method to bring the secondary system to the point that it will be fully synchronized with the primary system while using the available network bandwidth that was planned for the steady state after the deduplication factor takes effect.

The key here, as mentioned in the previous section, is to introduce workload by backup policies such that nominal data stays under the bandwidth ceiling available between the sites until the secondary system is fully primed.

## 5.8.5 Activating the ProtecTIER Replication Manager

The ProtecTIER (PT) Replication Manager is installed together with the ProtecTIER code in the TS7610, TS7650, or TS7650G system. When the replication feature is required, you just need to activate the ProtecTIER Replication Manager using the ProtecTIER Configuration Menu.

To activate the ProtecTIER Manager, complete the following steps:

1. Make sure that the ProtecTIER server is powered on. You can either use the console or SSH to the Linux login prompt. Login as ptconfig and press Enter.
2. At the Password prompt, enter ptconfig and press Enter. The ProtecTIER Configuration Menu displays, as shown in Example 5-10.

*Example 5-10 ProtecTIER Configuration menu*

---

```
login as: ptconfig
ptconfig@10.0.201.112's password:
Access denied
ptconfig@10.0.201.112's password:
Last login: Tue Nov 30 22:34:04 2010 from 10.0.158.138
ProtecTIER Configuration Menu:
=====
```

1. Update Time, Date, Timezone & Timeserver(s)
2. Update System Name
3. Update Customer Network
4. Enable/Disable Call Home
5. Activate/Deactivate Replication Manager

- 6. Update Replication Network
- 7. Configure Static routes
- 8. Configure Application Interfaces
- 9. Restore Network Configuration

q. Quit

Please choose an option:

---

3. From the menu, select option 5 for Activate/Deactivate Replication Manager. At the prompt Please choose an option, enter 5, and press Enter.
4. If the Replication Manager is currently not activated, a message displays, as shown in Example 5-11. When prompted by the Replication Manager is currently deactivated, do you wish to activate it? (yes|no) message, enter Yes. Read the warning, and at the next prompt, Are you sure you want to activate the ProtecTIER Replication Manager on this node? (yes|no), enter Yes.

*Example 5-11 Activating Replication Manager*

---

```
Please choose an option:5
Replication Manager is currently deactivated, do you wish to activate it?
(yes|no) yes
Warning: You should not activate Replication Manager on more than one server in
a grid, as doing so may cause conflicts within the grid. It is recommended that
you designate the Target server (hub) as the Replication Manager.
Are you sure you want to activate the ProtecTIER Replication Manager on this
node? (yes|no) yes
Gathering information [ Done ]
activatePTReplicationManager ended successfully
Press the ENTER key to continue...
```

---

5. Press Enter and you will go back to the main Configuration Menu. You have now successfully activated the ProtecTIER Replication Manager.

**Note:** You can also activate the Replication Manager using the command-line interface (CLI). Log in as root to the ProtecTIER node using SSH using the default password admin, and run the following command:

```
/opt/dtc/install/ptconfig -activatePTRepMan
```

Enter yes when prompted if you want to activate the ProtecTIER Replication Manager and yes when prompted if you want to reinstall it on this node.

## 5.8.6 Deactivating the ProtecTIER Replication Manager

The ProtecTIER (PT) Configuration Menu also has an option of deactivating the ProtecTIER Replication Manager. You can deactivate the Replication Manager using the PT Configuration Menu or using the CLI.

To deactivate the ProtecTIER Manager using the Configuration Menu, complete the following steps:

1. Make sure that the ProtecTIER server is powered on. You can either use the console or SSH to the Linux login prompt. Log in as ptconfig and press Enter.

2. At the Password prompt, enter `ptconfig` and press Enter. The ProtecTIER Configuration Menu displays, as shown in Example 5-12.

*Example 5-12 ProtecTIER Configuration Menu*

---

```
login as: ptconfig
ptconfig@10.0.201.112's password:
Access denied
ptconfig@10.0.201.112's password:
Last login: Tue Nov 30 22:34:04 2010 from 10.0.158.138
ProtecTIER Configuration Menu:
=====
```

1. Update Time, Date, Timezone & Timeserver(s)
2. Update System Name
3. Update Customer Network
4. Enable/Disable Call Home
5. Activate/Deactivate Replication Manager
6. Update Replication Network
7. Configure Static routes
8. Configure Application Interfaces
9. Restore Network Configuration

```
q. Quit
Please choose an option:
```

---

3. From the menu, select option 5 for Activate/Deactivate Replication Manager. At the prompt Please choose an option, enter 5 and press Enter.
4. When the Replication Manager is currently active, you receive the prompt Do you wish to deactivate it; enter yes. Read the warning message. If you want to deactivate the ProtecTIER Replication Manager, at the prompt Are you sure you want to deactivate the ProtecTIER Replication Manager on this node?, enter yes, as shown in Example 5-13.

*Example 5-13 Deactivating PT Replication Manager*

---

```
Please choose an option:5
Replication Manager is currently active, do you wish to deactivate it? (yes|no)
yes
Warning: Before deactivating the Replication Manager, please note that
re-activating the Replication Manager later on will cause an inconsistency with
the grid data files. This can be solved using the Replication Manager Restore
wizard, using the ProtecTIER Manager.
Please follow the user's manual for further information.
Are you sure you want to deactivate the ProtecTIER Replication Manager on this
node? (yes|no) yes
deactivatePTReplicationManager ended successfully
Press the ENTER key to continue...
```

---

5. You have now successfully deactivated the PT Replication Manager on this node. Press Enter to go back to the main Configuration Menu.



**Note:** You can also deactivate the Replication Manager using the CLI. Log in using root to the ProtecTIER node using SSH using the default password admin, and run the following command:

```
/opt/dtc/install/ptconfig -deactivatePTRepMan
```

Enter yes when prompted if you want to deactivate the ProtecTIER Replication Manager on this node. When successful, the deactivatePTReplicationManager ended successfully message displays.

## 5.8.7 Adding the repository to the grid

The procedure in this section uses ProtecTIER Replication Manager to create a replication pair. Complete the following steps:

1. To start the ProtecTIER Replication Manager, click the Grid Management icon (the left icon shown in Figure 5-94) to open the ProtecTIER Replication Manager.

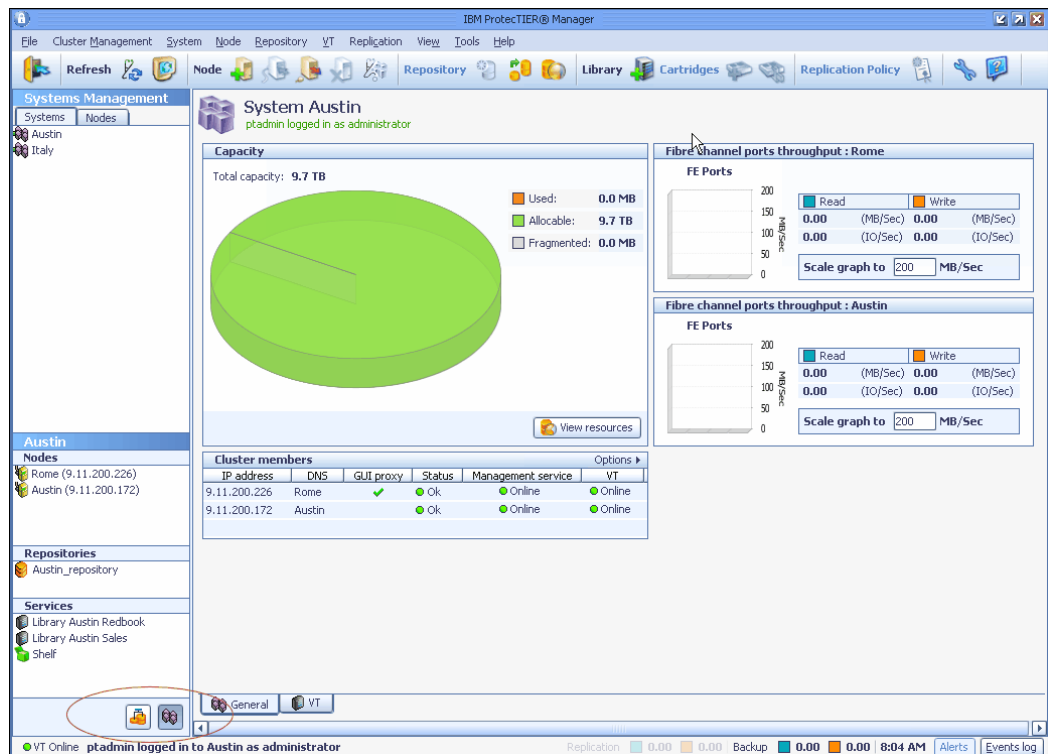


Figure 5-94 ProtecTIER Replication Manager: Select ProtecTIER Replication Manager icon

- The ProtecTIER Replication Manager application window opens. Select **ProtecTIER Replication Manage** → **Add ProtecTIER Replication Manager** (Figure 5-95) to create a new grid.

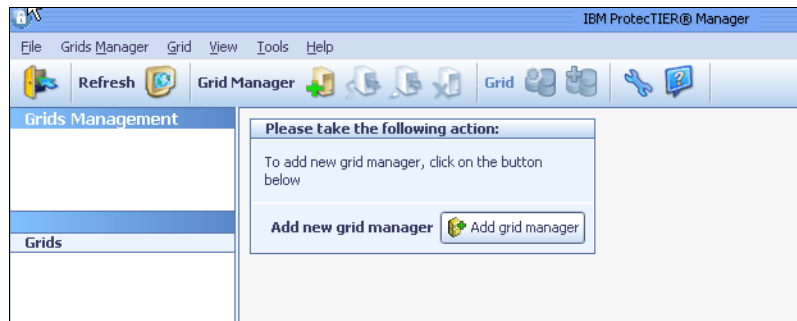


Figure 5-95 PT Replication Manager: Add ProtecTIER Replication Manager window

- Enter the IP address of the node that you want to add. This is the external IP address of the node (Eth0). In our example, this is IP address 9.11.200.172 (Figure 5-96).

**Note:** A maximum of 256 repositories can exist in a grid. A repository cannot be a grid member of more than one grid.




Figure 5-96 ProtecTIER Replication Manager entering the IP of the first node

- The login window is displayed. Log in with your user ID. Another available user ID is gmadmin with password gmadmin (Figure 5-97). Another available user is gmuser with password gmuser.



Figure 5-97 ProtecTIER Replication Manager login window

5. After you are logged in, click **Create new grid** (Figure 5-98).

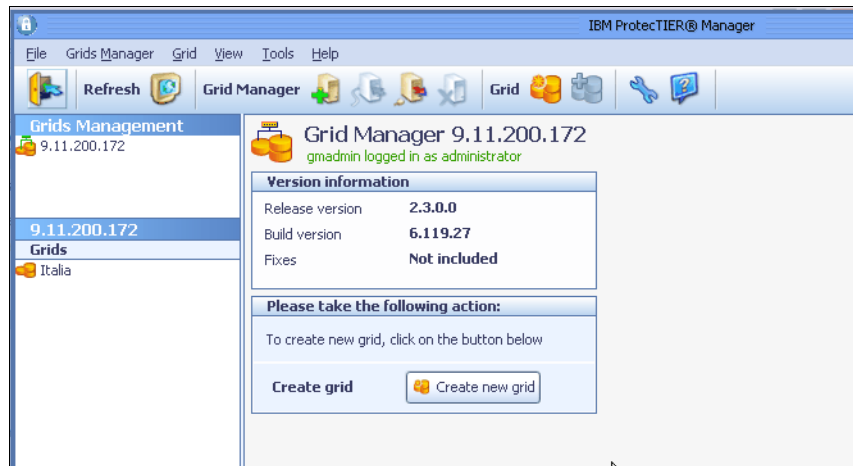


Figure 5-98 ProtecTIER Replication Manager Create new grid window

6. On the next window, you must provide a name for your grid that you are setting. We call our grid Redbooks Grid (Figure 5-99).

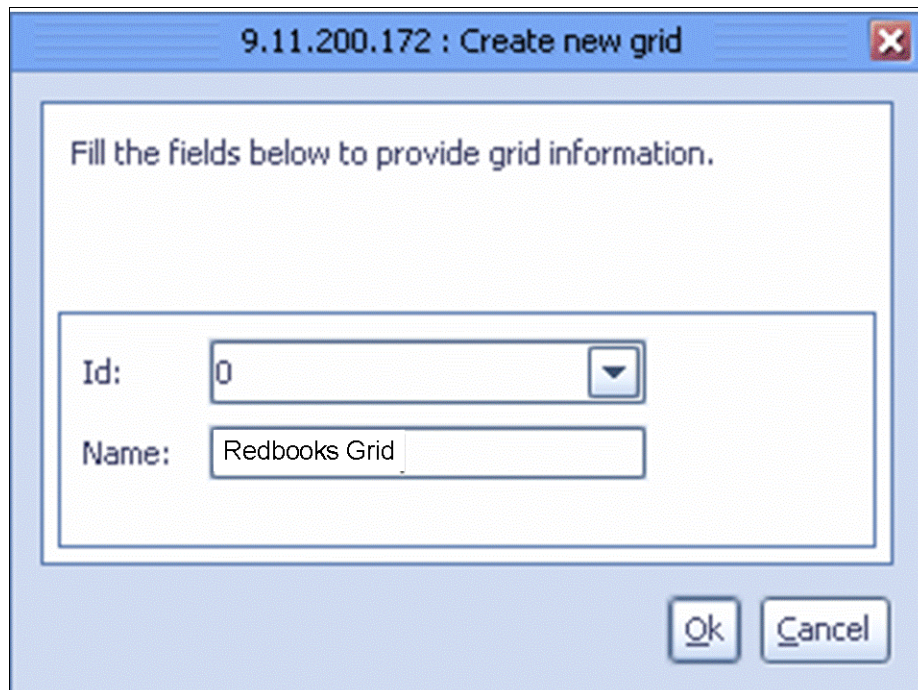


Figure 5-99 ProtecTIER Replication Manager naming window

ProtecTIER does not allow you to choose an ID that has already been used.

**Note:** Grid IDs and repository IDs are numbers that are never recycled. Even if a repository *leaves* the grid and *re-enters* the grid, it will not receive the same ID. Thus, actions like leaving a grid are expected to be rare and should not be planned as part of the normal flow of work.

7. Click **OK**. The repository information window opens. Provide the IP address of the node that you want to add to the grid and the repository connection parameters (Figure 5-100). Leave the port numbers at their default values and use the ptadmin user account for the Login information fields. This step only must be done once per repository. If dealing with a two-node cluster, pick only one.

**Important:** For the Replication IP address field, you must enter one of the replication IP addresses (eth3 or eth4) of any source repository node, *not* a customer IP address (eth0).

9.11.200.172 : Add repository to grid

Fill the fields below to provide repository information. This information must include the administrator user name and password, as well as, the repository connectivity parameters.

**Repository connection information**

Replication IP address :

Port:

Ping port:

**Login information**

User name:

Password:

Ok Cancel

Figure 5-100 ProtecTIER Replication Manager repository information

An add to grid window opens (Figure 5-101).

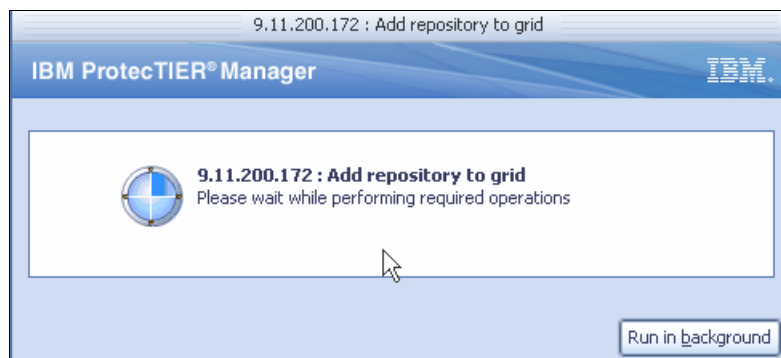


Figure 5-101 ProtecTIER Replication Manager add to repository window

The add repository to grid action will take your node offline, and after it is added successfully, your added repository appears in the ProtecTIER Replication Manager window (Figure 5-102).

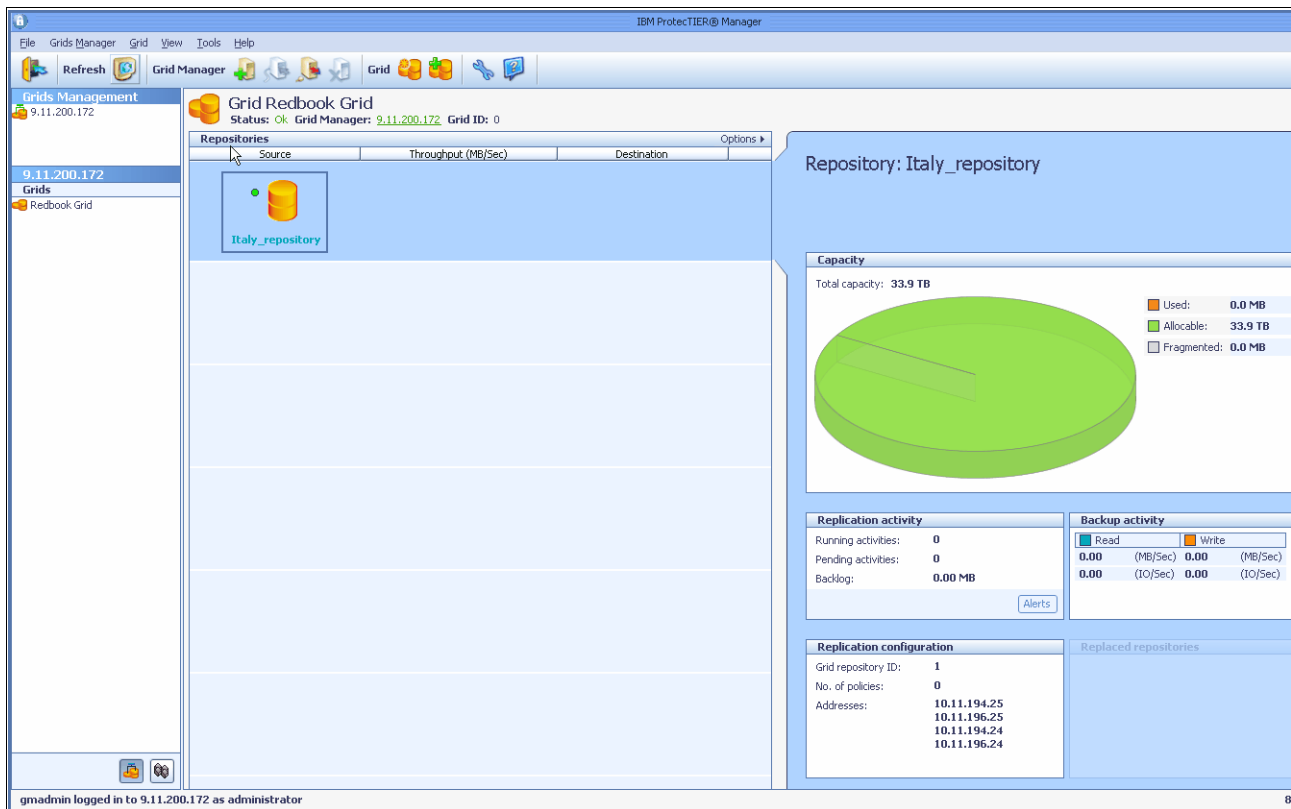


Figure 5-102 ProtecTIER Replication Manager after a successful addition of the repository

- After a successful addition of the source repository, all steps must be repeated for the target repository. After you add the source repository, you can set up the replication policies. Figure 5-103 shows both repositories in the grid.

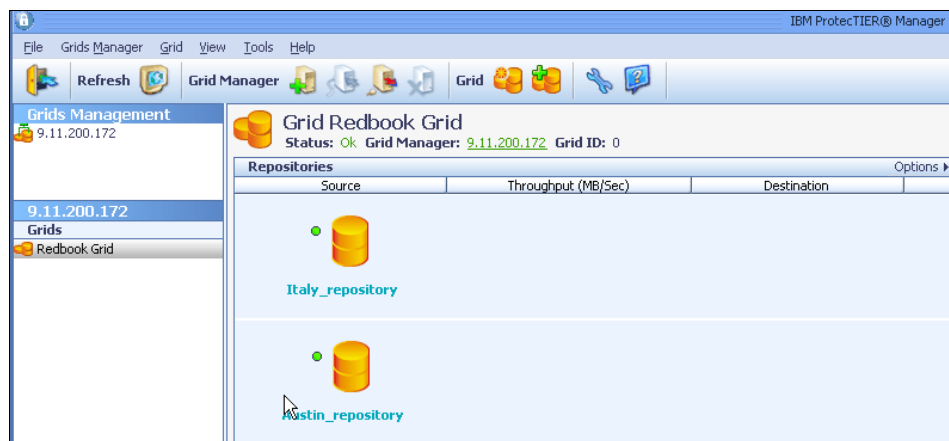


Figure 5-103 ProtecTIER Replication Manager view with both repositories

## 5.8.8 Setting up the replication policies

Before creating a replication policy, you must create a replication pair. When defining a replication pair, you also define the direction of the replication from the source repository to the destination repository. Typically, the source is the primary repository and the destination is the disaster recovery (DR) site.

A replication policy defines a set of cartridges from a local repository that must be replicated to a remote repository and establishes the rules for when and how that replication will take place. When an event matching a replication policy rule occurs, such as data on a cartridge changing or the user ejecting a cartridge into an I/E slot, a trigger is created for replication activity to take place.

Replication policies are defined through the Systems Management view of ProtecTIER Manager. A policy can only be created on a repository that is the source in a grid's pair and the policy only applies to the repository on which it is defined.

When creating a replication policy, the following parameters are defined:

- ▶ User object (that is, cartridges)
- ▶ Time frame in which the policy will take place (the replication window)
- ▶ Replication destination, with visibility change (specific library) or without visibility change (virtual shelf) at the remote location
- ▶ The priority: Each policy can be set with three priority levels:
  - Low
  - Medium
  - High

Users can set one policy with low priority and another with medium or high. The system default is low.

By default, the destination of the replication is the virtual shelf at the remote repository. If the visibility control switch is enabled, the destination is defined as a specified target library. This means that if you eject a cartridge that belongs to a policy with the enabled visibility switch, the cartridge is *moved* to the shelf of the local repository, and at the destination repository, once replicated, the cartridge is placed in the import/export slots of the specified destination library.

Create a single policy or a small number of them for any single library for ease of management. It might be easier to manage one policy per library rather than many policies to reduce maintenance burden and to prevent human errors.

One shelf is configured for each repository in the grid and when the cartridges are in the shelf and cannot be seen by the backup application. The cartridges are *outside* the library. ProtecTIER GUI has access to cartridges in the shelf in a special view.

When the cartridges are in the shelf, the following actions can be done:

- ▶ Relocated to a library's import slot with the ProtecTIER GUI or by visibility switching.
- ▶ Be replicated to another repository.
- ▶ The cartridges can be deleted.

## Creating a replication policy

When creating a replication policy, the following parameters are defined:

- ▶ User object type (for example, cartridges)
- ▶ User object (for example, barcodes)
- ▶ Replication destination (for example, visibility change)

If the destination is defined as a target library, visibility switching is enabled. This means that if you eject a cartridge that belongs to the respective policy, the cartridge is relocated to the shelf of the local repository, and on the destination repository the cartridge is placed in the import/export slots of the destination library.

Visibility control is the means by which you can determine *where* cartridges actually exist. From a backup application standpoint, a specific cartridge or barcode can *exist* in only one location at a given time. After they are exported by the backup application, cartridges can be placed on a virtual *shelf*, which is a container for cartridges, that is visible through ProtecTIER Manager, providing more flexibility in managing cartridges and where they are kept (similar to keeping physical tapes on an actual shelf outside of the tape library).

**Note:** Visibility of a cartridge is automatically transferred to the shelf when an eject operation from ProtecTIER or from the backup application is done on it.

The following steps provide an example of creating a replication policy:

1. In the Grid Manager, select **Grid Manager** → **Create Replication Pair** (Figure 5-104).

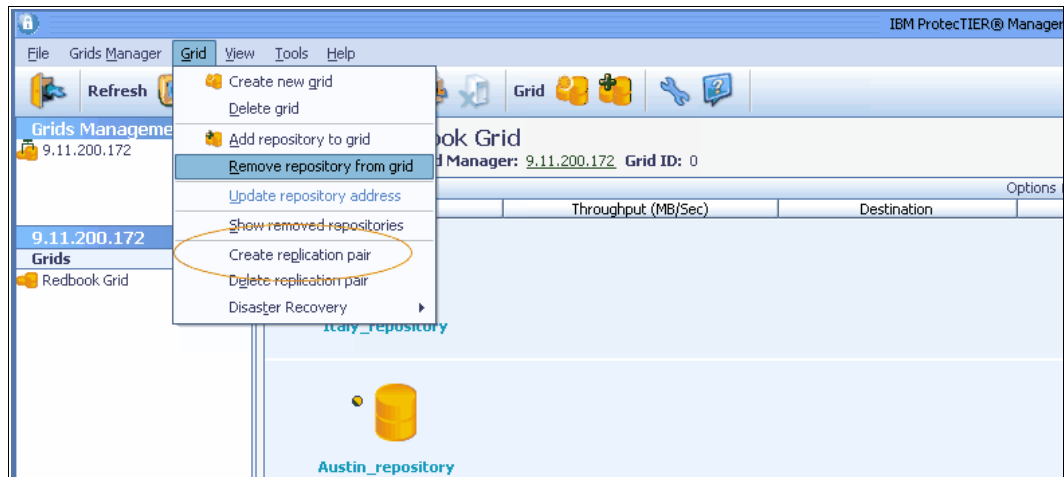


Figure 5-104 Creating a replication pair

2. In the next window, set up the source repository and destination repository. In our example, you can see that the source and the destination repository must be different. A message will be displayed if the source repository and the destination repositories are the same (Figure 5-105).

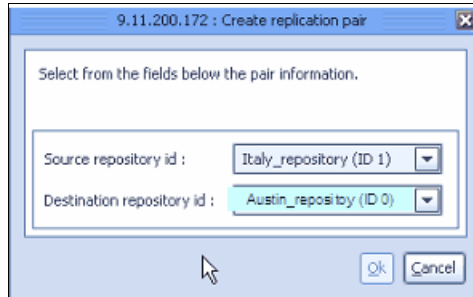


Figure 5-105 Source and destination repository

The create replication runs for a short time, and after the replication pair is created successfully, the source and destination repository are displayed (Figure 5-106).

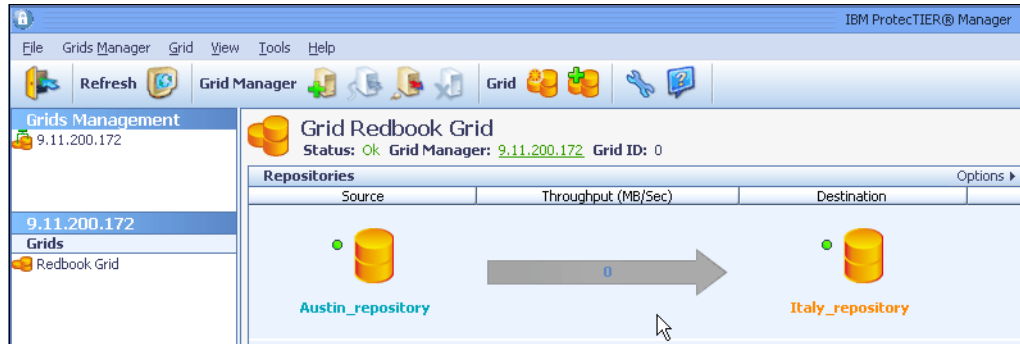


Figure 5-106 Source and destination repository view

The next step is to set up the replication policy on the *source* repository. We now describe how to set up a replication policy.

Creating a replication policy is done from the System Management View Pane ProtecTIER web GUI by completing the following steps:

1. When you are still in the Grid Manager View pane of the web GUI, change to the System Management View pane.



2. Select **Replication** → **Policy** → **Create policy**. The Welcome window of the Create policy wizard opens (Figure 5-107).

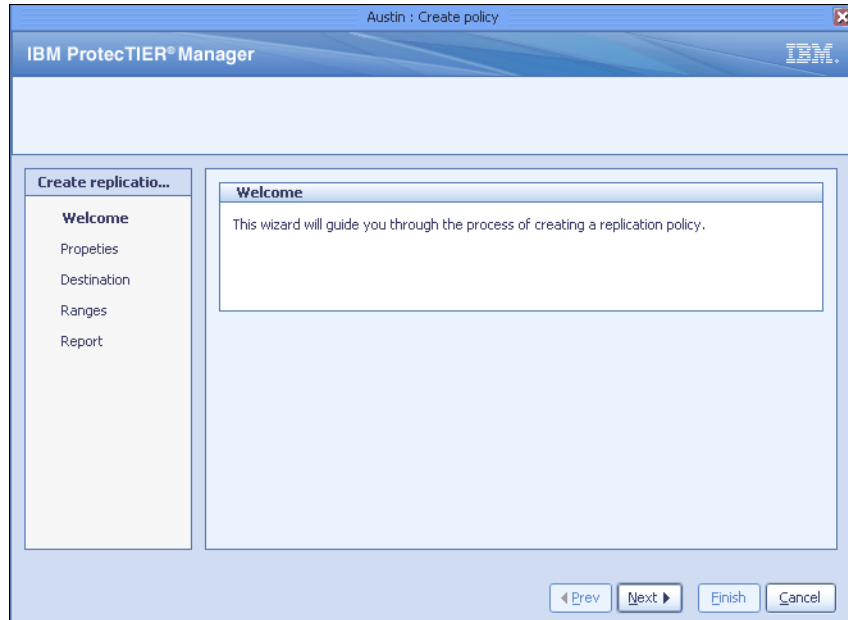


Figure 5-107 Create Replication Policy welcome window

Click **Next**. The properties window opens (Figure 5-108).

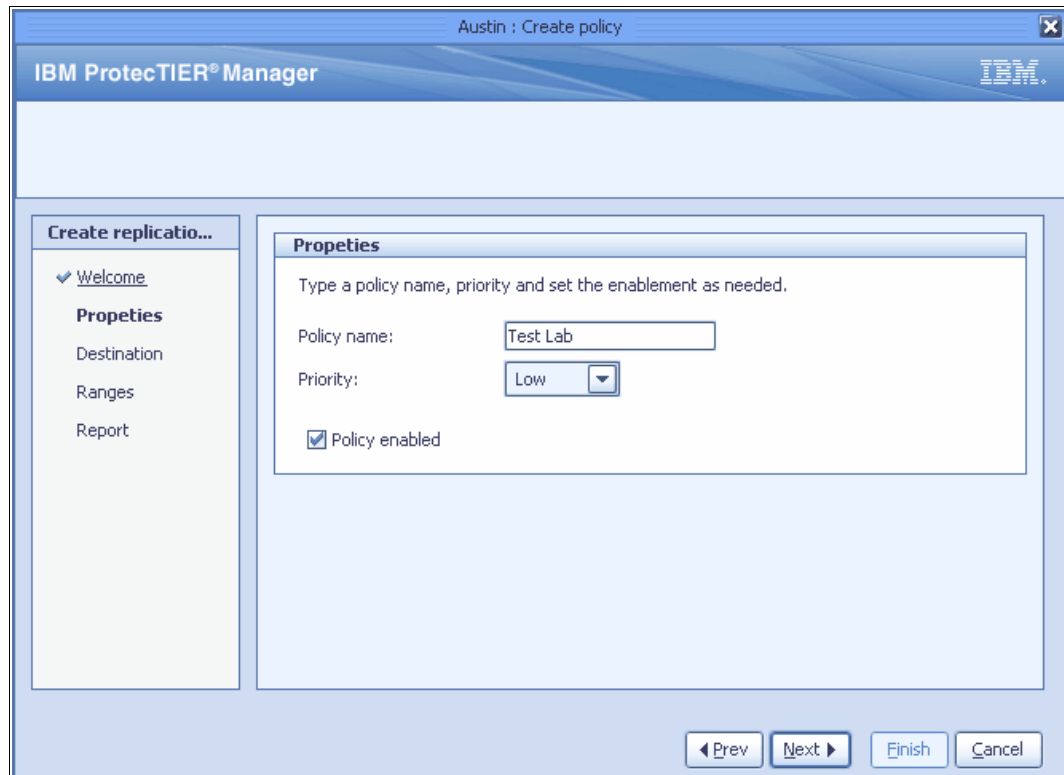


Figure 5-108 Create Replication Policy Properties window

- a. Type a unique policy name in the Policy name field. In our example, we used *Test Lab* policy). You cannot have the same policy name twice. If you do define the same policy name, an error message is displayed.
- b. Policies have three options of priority:
  - High
  - Medium
  - Low

Define the policy's priority according to the importance or urgency of the data that must be transferred. For example, a policy with a high priority is transferred first, then a policy with medium priority, followed by low priority. The default is Low for every policy.

- c. Selecting **Policy enabled** automatically runs the policy within the time frame defined. If Policy enabled is not selected, no activities take place.

Click **Next**. The Replication Destination window opens (Figure 5-109).

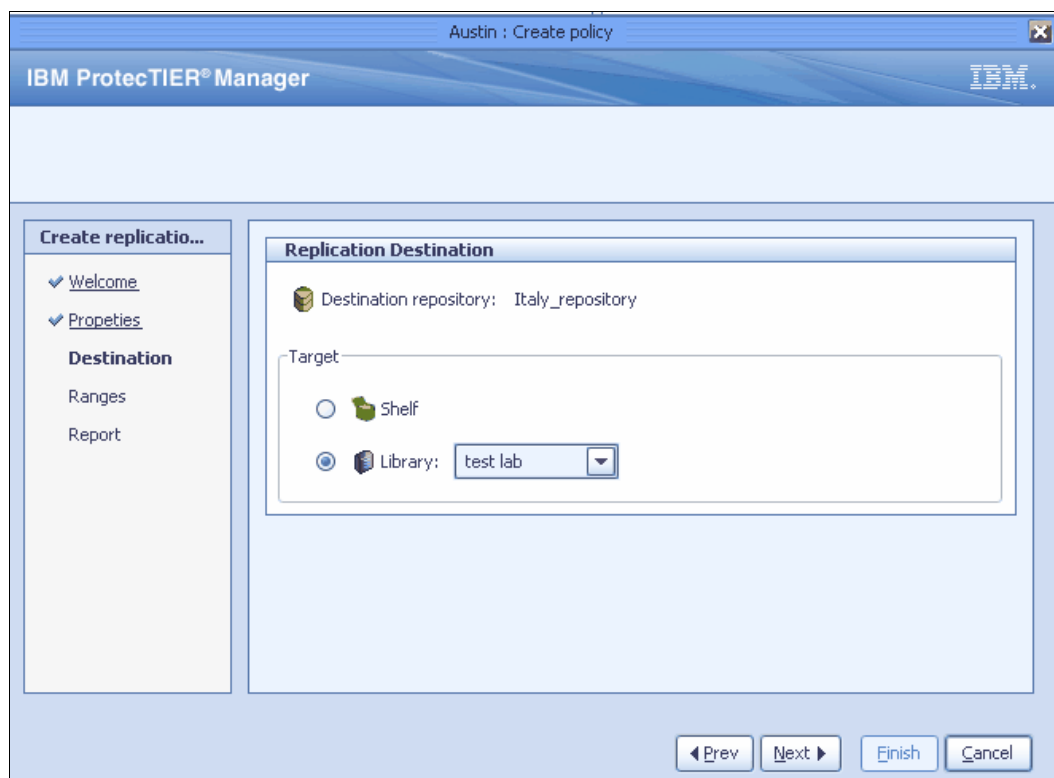


Figure 5-109 Create policy: Replication Destination window

The Destination repository option controls the visibility of replicated cartridges at the destination repository. Destination replica cartridges can be *invisible* (if you choose the shelf) or *visible* (if you choose a library). The target is the destination of the cartridge. The target is either the shelf or a library:

- If you choose the shelf, the visibility switching feature is not activated.
- If the target is a library, the visibility switching feature is activated. When the cartridge is ejected, it is first relocated to the shelf on the local repository. On the remote repository, the cartridge is relocated to the import/export slot of a library so that the backup application on the remote site has visibility to it.

In this example, the test lab library was used as a target.

Click **Next**. The Barcode ranges window opens (Figure 5-110).

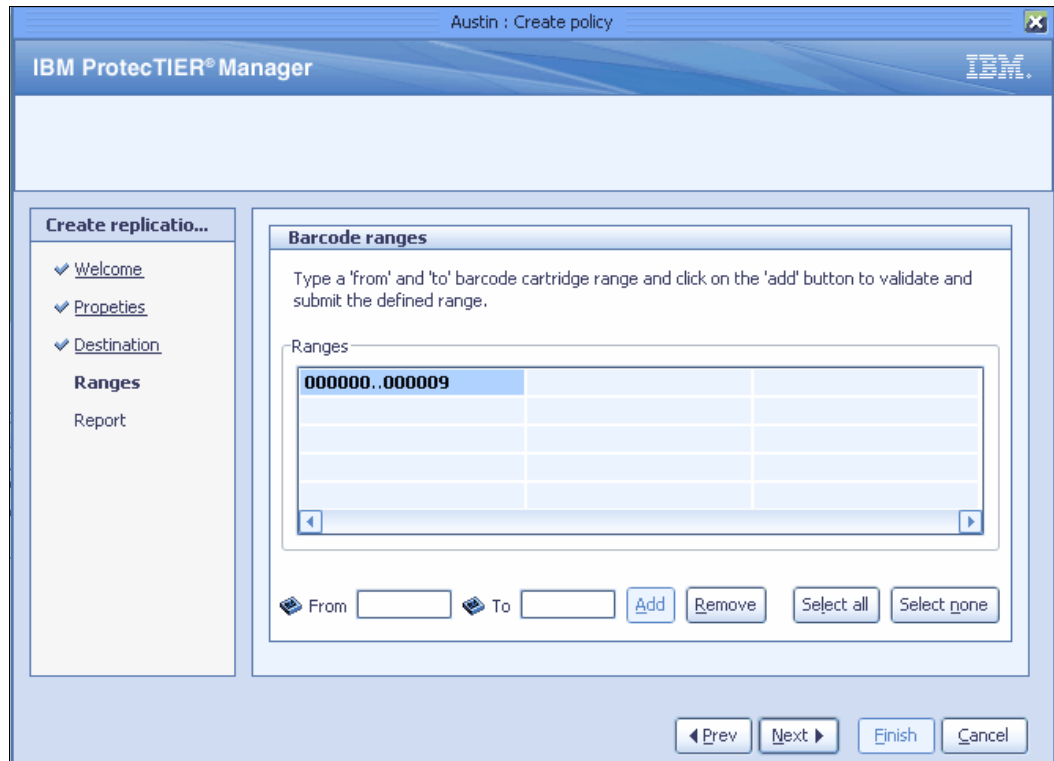


Figure 5-110 Create policy: Barcode ranges window

3. The policy objects and the cartridges are defined as barcode ranges. There can be up to 256 ranges in a single policy.

Enter the From and To barcodes for a range of cartridges to be replicated. Click **Add** to view the range in the Ranges table. If a barcode number or barcode range appears more than once, an error message is displayed with the conflict. To delete the barcode ranges from the table, select **Select all** → **Remove**, or click **Select none** to deselect.

Click **Next**. The Summary report window opens with the policy name and the number of ranges that were defined (Figure 5-111).

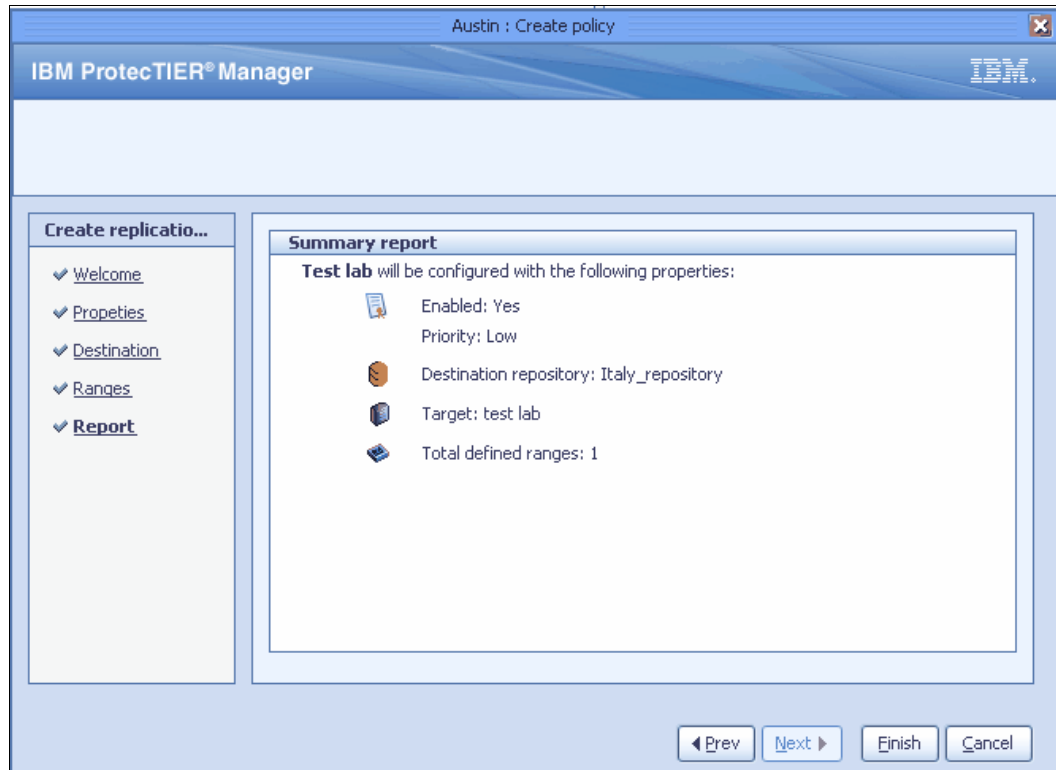


Figure 5-111 Create policy: Summary report window

In the System Management View Pane, you can see your newly created replication policy (Figure 5-112).

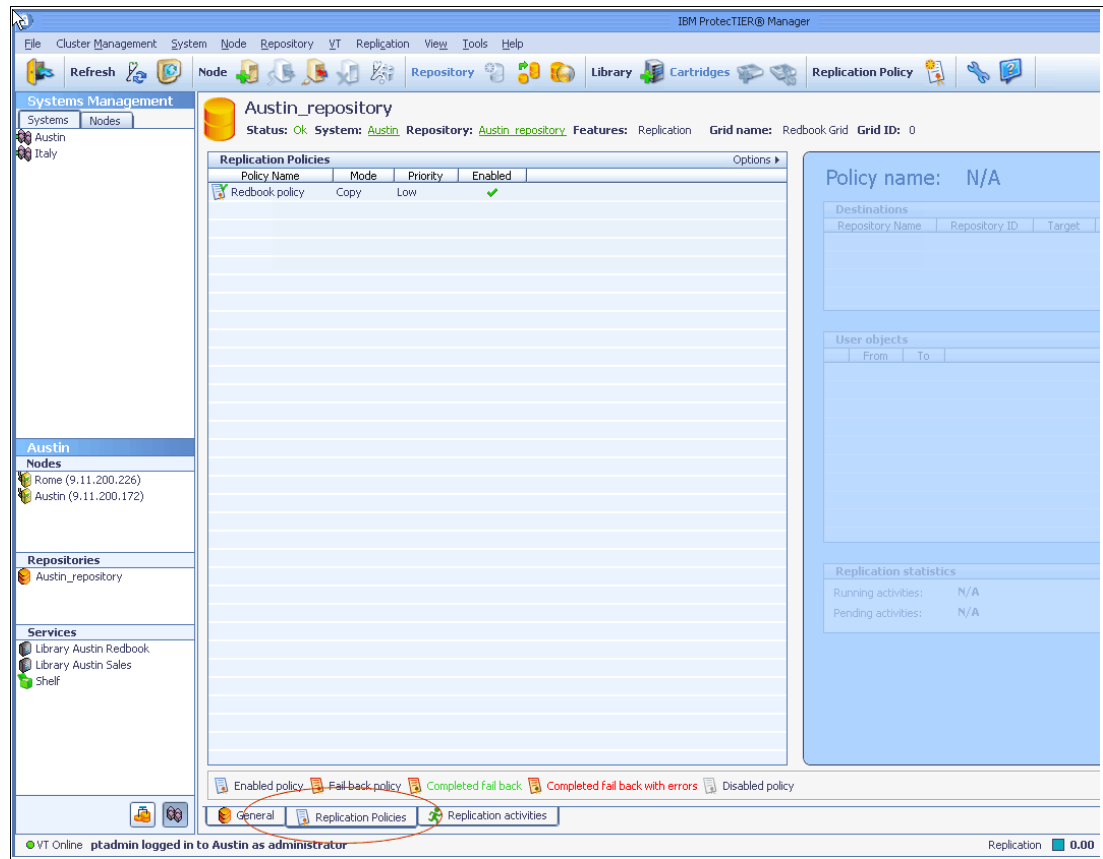


Figure 5-112 Creating Replication Policy view

### Global replication time frame

The objective of configuring the replication time frame is to control when replication activities occur. This is a system-wide option, meaning that it affects all policies in the system and is configured only at the source repository. There is no replication window definition for the receiving (target) system.

When selecting the replication mode, the user is given two options:

- ▶ No Backup Precedence: This is the Dedicated Window mode. With this option, the user can schedule a time window per day during which replication activities will occur. This is the preferred mode of operation that will work for almost all customer use cases.
- ▶ Precedence to Backup: This is the continuous mode of operation in which replication occurs concurrent to the backup operation. This mode is available as an option for select cases, but should be rarely used, as discussed above.

Now set up your replication policy. The next and final step is creating a replication time frame by completing the following steps:

1. In the System Management View Pane, select **Select Replication** → **Set replication timeframe**. The Set replication time frame window opens (Figure 5-113).

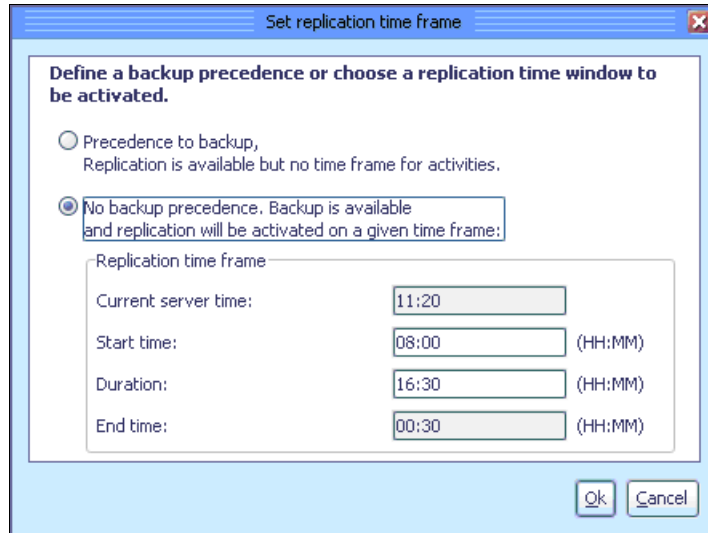


Figure 5-113 Set replication time frame

2. Select an option. You can either set a precedence for backup where the replication policy will take precedence in the order of other policies to be replicated, or you can set a given time frame by setting the start time, duration, and end time. Click **OK**.

Replication setup is now complete. Review Chapter 11, “Native replication and disaster recovery” on page 575 to familiarize yourself with disaster recovery situations.

## 5.8.9 Enabling and disabling a policy

A successfully created policy is enabled by default, which means that all incoming replication events will apply their rules onto the policy's definition. The user may choose to disable a policy at any time. When a policy is disabled, all incoming replication events will *ignore* the policy from the moment that it is disabled, which does not affect the currently running and pending activities. Any disabled policy can be enabled from within ProtecTIER Manager at any time as well. This does not affect current running activities.

## 5.8.10 Running a policy

Policies can be run either manually or automatically (that is, during the time frame set).

In automatic mode, whenever replication events are received, policies run continuously, whether cartridges are being written to by the backup application, ejected from a library (if visibility switching is activated), or unloaded from a drive. These policies start the actual replication activity during the time frame set by the user when created.

Manually run policies create replication jobs for all the valid cartridges included in their list, whether or not they must be replicated.

With both modes of operation, running a policy leads to lining up replication jobs in their respective priority queues where they wait for resources and the replication time frame to start replicating.

Policies should be created and executed in line with the performance capacity of the ProtecTIER system, and the bandwidth capacity available between the two sites kept in mind. This concept is covered at length further in this document. In line with this principle, we typically do not recommend defining a single policy with thousands or even hundreds of cartridges and executing it in the manual mode, as this might create a large number of events and replication triggers that can potentially overload the system. Instead, users should use manual cartridge replication with multiple-selection of the needed cartridges.

The following example demonstrates the difference between the two approaches. A repository has 500,000 cartridges with a single replication policy defined that includes all cartridges. In this case, with manual policy execution, the user will replicate all cartridges, including the empty cartridges' metadata. However, if manual cartridge replication with multiple-selection is used, specific cartridges will be selected by the user and only those cartridges will be replicated.

## 5.9 Setup replication on OpenStorage systems

This section provides the information needed to manage native replication configuration and monitoring of ProtecTIER for OpenStorage.

Native replication lets you replicate data objects between ProtecTIER repositories. In order for a logical set of repositories to replicate from one to another, you must create a replication grid. The replication grid is remotely created and managed by the Replication Manager.

The ProtecTIER Replication Manager is a server that remotely manages the replication grids within an organization. The ProtecTIER Manager connects to the ProtecTIER Replication Manager using the IP address of the ProtecTIER Replication Manager server. The ProtecTIER Replication Manager can be installed on a dedicated host (which requires an RPQ), or on a ProtecTIER node. If the ProtecTIER Replication Manager is installed on a ProtecTIER node, it can manage up to one grid with 24 repositories. If ProtecTIER Replication Manager is installed on a dedicated server, it can manage up to 64 grids with 256 repositories in each grid.

Each ProtecTIER Replication Manager has a unique identity. A repository, after it has joined a replication manager, cannot join a replication grid managed by a different replication manager, even if it has left the grid. This configuration prevents data collision.

In an OpenStorage environment, replication policies and activities are managed by NetBackup. ProtecTIER Replication Manager is used to manage the replication grid, create application groups, and define the replication connections within the replication grid. Each repository, or grid member, can replicate to each other and replication is bidirectional. A maximum of 12 repositories (that is, members) can replicate to each other at any given time, with up to 256 repositories existing in the entire grid.

**Note:** OpenStorage duplication and ProtecTIER replication are closely related. In this book, we use the same term, replication, for VTL replication and as well as when we refer to the OpenStorage duplication. If you are not using the replication feature, you may skip this section.







## Host implementation for virtual tape libraries

In this chapter, we discuss LUN Masking, configuring host systems for attachment to ProtecTIER systems, and what is required to achieve optimum performance from ProtecTIER systems.

## 6.1 Connecting hosts to ProtecTIER systems

You can connect your backup server to the SAN and then create a zone to include both the host HBA WWN and the WWN of ProtecTIER front-end ports. But if you intend to take advantage of path redundancy or you have a two-node clustered system, then the connections should be well designed, not only for redundancy purposes, but also for performance reasons. For more details about redundancy connections topology, refer to Chapter 4, “Hardware planning for IBM System Storage ProtecTIER” on page 121.

### 6.1.1 What bandwidth you need for ProtecTIER

Many SANs are designed for disk I/O, which means that they are optimized for high input/output per second (IOPS) and are not necessarily for the high bandwidth that is needed for tape and backup I/O. During backup sessions, high volumes of data are generally transferred from disk to tape or virtual tape system. High-speed transfer is also a necessity to minimize backup time. Current tape or virtual tape drives, such as the IBM ProtecTIER system with up to 500 MBps per node and up to 1000 MBps on a two-node clustered system, can easily saturate a single Fibre Channel (FC) link if several virtual drives operate concurrently. Using a SAN Fabric configuration, you can potentially attach and access many tape drives through one or more HBAs. But how will this affect performance if several tape virtual drives run concurrently? The theoretical maximum data transfer rate for one Fibre Channel connection in a 2 Gb SAN is 200 MBps. In reality, we typically see an effective data transfer rate of about 160 MBps. In a 4 Gb SAN, we see an effective data transfer rate of just around 300 MBps. So if you have only one 4 Gb SAN HBA in your backup media server, you cannot maximize the power of your ProtecTIER system.

As well as the number of HBAs, the overall SAN design must be capable of supporting the volume of data transferred during the backup process (Figure 6-1 on page 277). The SAN might be able to easily sustain the overall load and normal data traffic, but still have an insufficient bandwidth when backups take place. This is especially true in complex fabrics with multiple switches, where the inter-switch links (ISLs) can become saturated. You can increase your bandwidth if you install additional HBAs and increase the number of ISLs.

Figure 6-1 shows a typical SAN layout with ISL between the servers and tapes.




Figure 6-1 SAN bandwidth

## 6.1.2 Multiple paths to tape drives

For some 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 definitions on your server. To solve this problem:

- ▶ You can enable *zoning*, which allows one HBA to see only certain tape drives.
- ▶ You can use the *alternate path function* within the IBM Tape Device Driver.
- ▶ You can use *persistent binding* for the HBAs to see only certain WWNs.
- ▶ You can ignore the extra copy of each device.

### Relationship of virtual tape resources and device name in OS

Regardless of the solution that you choose, you should be aware of the relationship between the virtual tape resource and the device name in the OS. Unlike a physical tape library, you cannot distinguish the tape devices by WWN because we can get more than one device from one FC port on the virtual tape library, as they might have the same WWN. To identify the tape devices, you must use the serial number. You can get the serial number by running the `lscfg` command in AIX.

Example 6-1 shows three virtual tape drives in our environment. rmt0, rmt1, and rmt8 have the same WWN, 10000000C979518D. They are all assigned to the ProtecTIER FE port 10000000C979518D. But rmt0 and rmt8 have the same serial number, 4604034000, so rmt0 and rmt8 are devices generated for the same virtual tape drive by the device driver program in the operating system.

*Example 6-1 lscfg example*

---

```
highway> lscfg -vpl rmt0
  rmt0          U787A.001.DNZ08E8-P1-C4-T2-W10000000C979518D-L1000000000000
IBM 3580 Ultrium Tape Drive (FCP)
```

```
Manufacturer.....IBM
Machine Type and Model.....ULT3580-TD3
Serial Number.....4604034000
Device Specific.(FW).....5AT0
```

PLATFORM SPECIFIC

```
Name: tape
Node: tape
Device Type: byte
```

```
highway> lscfg -vpl rmt1
  rmt1          U787A.001.DNZ08E8-P1-C4-T2-W10000000C979518D-L2000000000000
IBM 3580 Ultrium Tape Drive (FCP)
```

```
Manufacturer.....IBM
Machine Type and Model.....ULT3580-TD3
Serial Number.....4604034001
Device Specific.(FW).....5AT0
```

PLATFORM SPECIFIC

```
Name: tape
Node: tape
Device Type: byte
```

```
highway> lscfg -vpl rmt8
  rmt8          U787A.001.DNZ08E8-P1-C6-T2-W10000000C979518D-L1000000000000
IBM 3580 Ultrium Tape Drive (FCP)
```

```
Manufacturer.....IBM
Machine Type and Model.....ULT3580-TD3
Serial Number.....4604034000
Device Specific.(FW).....5AT0
```

---

You should make a table of virtual tape devices, as shown in Table 6-1.

Table 6-1 Tape library worksheet

Device in OS	Type	VTL system	VTL node	VTL port	VTL port WWN	Serial number	Element number

**Note:** There are differences from the real tape library that you are simulating. The virtual drives created in ProtecTIER might have the same WWN.

IBM i does not support multipath for tape and should be zoned so that multipath conditions do not exist.

### 6.1.3 Tape and disk on the same HBA

Tape and disk I/O are, by nature, different. Virtual tape devices transport data in the same way that physical tape devices do. So when we talk about transporting data, we refer to the virtual tape devices as tape devices, although they use disk space to store data. While tape drives use large blocks and *data streaming* to maximize performance, disk drives use smaller block sizes appropriate for random I/O. Therefore, mixing disk and tape on a single SCSI bus is not recommended and is rarely used. Tape drives use large blocks and data streaming that can tie up the SCSI bus for long periods of time, meaning that disk drives get less than their fair share of the SCSI bus. Conversely, tape drives might have difficulty getting access to the bus because disk drives can respond faster. The best practice is to keep disk and tape on separate SCSI buses.

With Fibre Channel, there is, in principle, less contention. I/O can be multiplexed, and small blocks can be interleaved in the middle of large blocks. There is no shared bus to keep a slow device from sending whenever it is ready, so it is possible with Fibre Channel to have both disk and tape sending or receiving with little interference. In a switched fabric, data is not directly sent to the device, but rather to the Fibre Channel switch.

As long as there is no other disk I/O activity during the backup operation, it might work without problems. But if other applications that also access the disk are running during the backup, performance will be impacted. Because disk access uses a smaller block size and is quicker and more frequent, the disk I/O will occupy the FC link and cause the tape to wait for access to the link, causing backhitches and thus affecting performance. Or worse, the disk I/O will keep the command queue filled so that no tape commands can go through.

For a backup server, where a high I/O load exists, and for LAN-free clients, where during the backup some other disk I/O occurs, you should use multiple HBAs and separate the disk and tape I/O (Figure 6-2). This is especially true for backup servers using a disk storage pool (such as with IBM Tivoli Storage Manager) where the backup data is staged to disk and then migrated to tape. If the backup and migration run simultaneously, it will cause a high I/O load. Keep in mind that your environment will evolve over time, and usually the data traffic will increase. Although it might have been justified (generally from a cost perspective) in a simple installation with low traffic to share the HBA, you must reevaluate the situation on a regular basis. In most cases, you will reach a point where it becomes necessary to install additional HBAs.




Figure 6-2 HBA shared between disk and tape

Most SANs are designed for disk I/O (high IOPS) rather than high bandwidth (such as tape), which means that there might not be enough ISLs for high bandwidth. As discussed previously, you need at least one ISL for every two to three tape drives. If you have more than one ISL between two switches, you should also consider enabling trunking. This is because there is no load balancing on the basis of real I/O over several ISLs if trunking is not enabled.

Many switches (for example, IBM 2109 and IBM 2005) have DLS (dynamic load balancing on the basis of the routing table) disabled, which means that load balancing is done during startup of the switch only. Periodically check whether all ISLs are working. After installing an additional ISL, run the dynamic load balancing at least once by setting `dlsset`. When done, remember to disable it again as recommended by IBM.

Another potential issue is the device driver level. IBM attempts to be on the same supported FC HBA device driver levels for tape and disk, but it is possible that because of a temporary issue, a given tape device will require a different device driver level than disk or vice versa.

For all of these reasons, wherever possible, a single HBA should *not* be shared for concurrent disk and tape operation. IBM supports mixing disk and tape on an HBA, and the IBM Support Center will accept problems reported on these configurations. However, if the problem determination process reveals that the cause is the mixing of tape and disk traffic, the customer might be told that the only fix is to separate the traffic.

### 6.1.4 SAN zoning

The ProtecTIER Gateway and appliance have specific recommendations about how SAN zones are created:

- ▶ Use zones based on World Wide Port Name (WWPN).
- ▶ Use two-member zones, that is, one initiator port and one target port per zone.
- ▶ For each backup server, create a separate zone for each HBA that will access ProtecTIER virtual resources.

Before creating WWN zones at a SAN switch, you must get the WWPN of each port for both your ProtecTIER and your host computer.

For ProtecTIER, you need the WWPN of each front-end port. You can use the ProtecTIER Manager GUI (Figure 6-3) to accomplish this task.

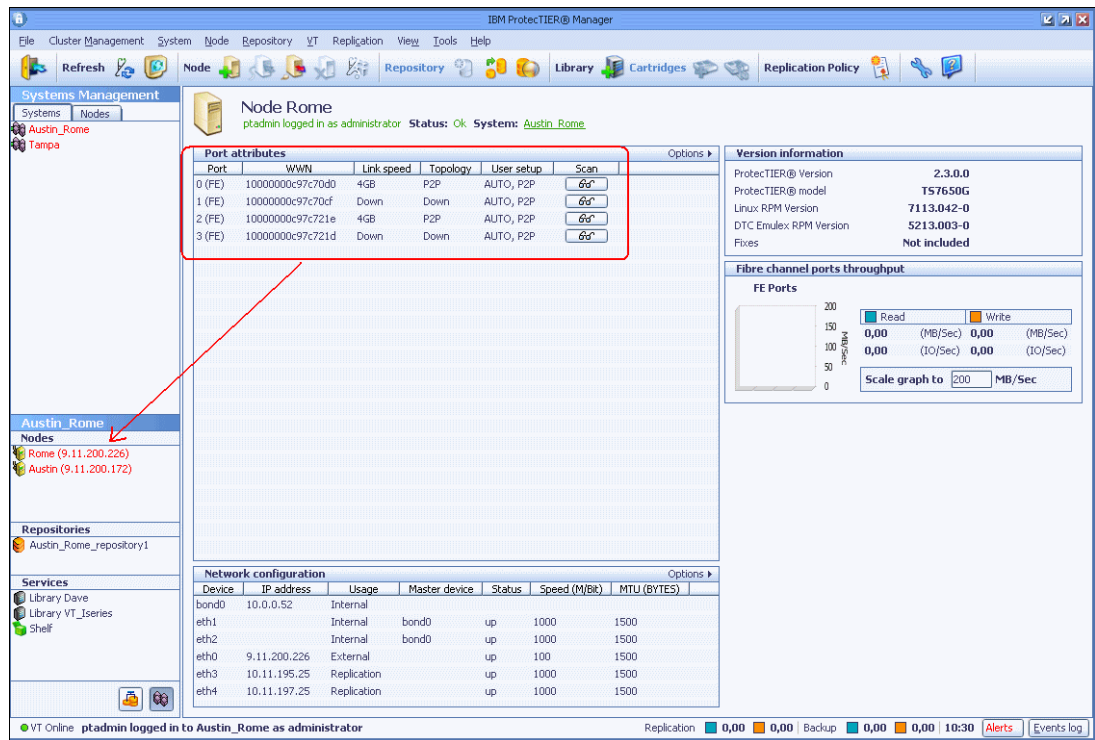


Figure 6-3 ProtecTIER front-end port WWPN

You also need to plug the FC cables into the correct port in the rear of ProtecTIER. Refer to Figure 6-4 and Table 6-2. You can obtain the physical locations for each port listed in the ProtecTIER Manager GUI.




Figure 6-4 ProtecTIER front -end ports

Table 6-2 Front-end port location

Location	Slot 1	Slot2
Upper	port0(FE)	port2(FE)
Lower	port1(FE)	port3(FE)

For your host computer, you can issue a command to get the WWPN of each HBA. In our case, we used AIX (Example 6-2).

*Example 6-2 lscfg command*

```

highway> lscfg -vpl fcs0
fcs0                U787A.001.DNZ08E8-P1-C6-T1  FC Adapter

Part Number.....03N5029
EC Level.....A
Serial Number.....1F6050C37D
Manufacturer.....001F
Customer Card ID Number.....5759
FRU Number..... 03N5029
Device Specific.(ZM).....3
Network Address.....10000000C952184E
ROS Level and ID.....02C82132
Device Specific.(Z0).....1036406D
Device Specific.(Z1).....00000000
Device Specific.(Z2).....00000000
Device Specific.(Z3).....03000909
Device Specific.(Z4).....FFC01155
Device Specific.(Z5).....02C82132
Device Specific.(Z6).....06C12132
Device Specific.(Z7).....07C12132
Device Specific.(Z8).....20000000C952184E
Device Specific.(Z9).....BS2.10X2

```



```
Device Specific.(ZA).....B1F2.10X2
Device Specific.(ZB).....B2F2.10X2
Device Specific.(ZC).....00000000
Hardware Location Code.....U787A.001.DNZ08E8-P1-C6-T1
```

#### PLATFORM SPECIFIC

```
Name: fibre-channel
Model: LP11000
Node: fibre-channel@1
Device Type: fcp
Physical Location: U787A.001.DNZ08E8-P1-C6-T1
```

---

### 6.1.5 LUN Masking for VTL systems

LUN Masking is used to control the visibility of a device (for example, tape drives or robots) by allowing specific devices to be seen only by a selected group of host initiators (backup servers).

LUN Masking allows you to assign specific devices to a specific host running backup application modules. It enables multiple initiators (backup servers) to share the same target FC port without having conflicts on the devices being emulated.

The LUN Masking setup can be monitored and modified at all times during system operation. LUN Masking in ProtecTIER influences the visibility of the devices by the hosts systems.

**Note:** Every modification in the LUN Masking in ProtecTIER may affect the host configuration and require re-scanning by the host and by the backup application.

#### Enabling or disabling LUN Masking

The LUN Masking settings are disabled by default. It can be enabled or disabled at anytime. LUN Masking groups can be created, edited, or deleted in either mode, but the relevant changes will only be applied to the tape emulation devices when LUN Masking is enabled.

When LUN Masking is disabled, all the devices are seen by all hosts connecting to that Fibre Channel (FC) port. When LUN Masking is enabled, devices are seen only by the hosts that are defined with them in the same LUN Masking group.

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

1. Log in to the PT Manager GUI. From the VT drop-down menu, select **LUN Masking** and select **Enable/Disable LUN Masking**, as shown in Figure 6-5.

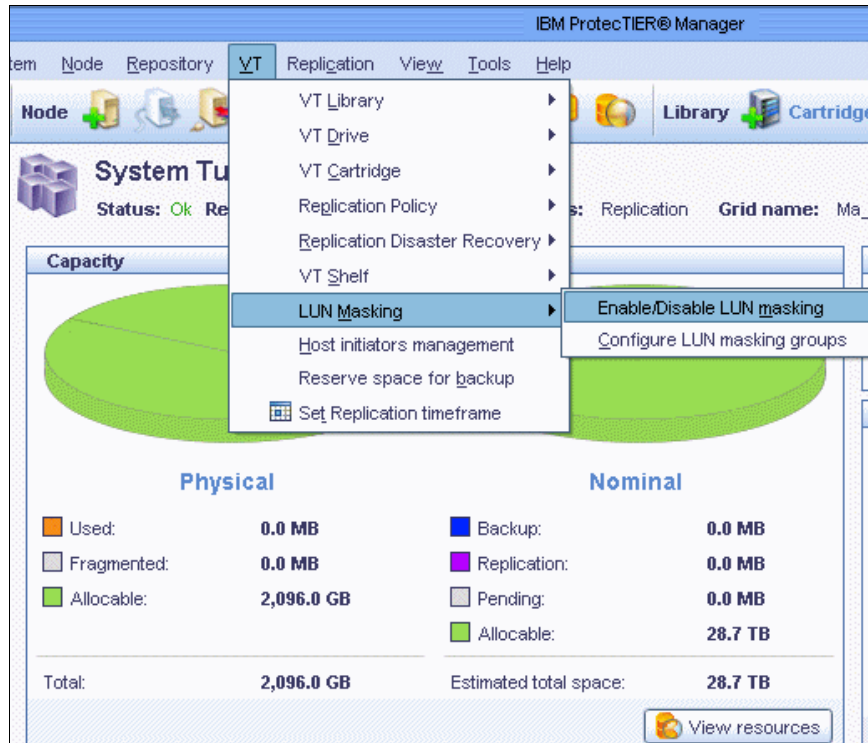


Figure 6-5 Select Enable/Disable LUN Masking

2. The Enable/Disable LUN Masking window opens. Select **Enable LUN Masking** to turn on LUN Masking, or **Disable LUN Masking** to turn off LUN Masking, as shown in Figure 6-6. Click **OK**.

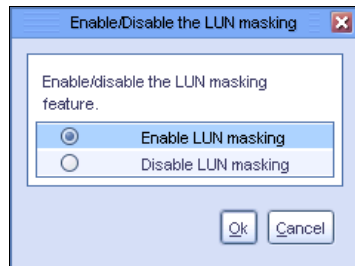


Figure 6-6 Enable/Disable LUN Masking

After you have enabled or disabled the LUN Masking option, rescan the devices from the host systems. You should perform this action so that the host will have the most up-to-date information in the list of visible devices and their associated LUN numbers.

### Adding LUN Masking groups

LUN Masking groups define the connection between the host initiators and the VT library devices (for example, robots or tape drives). Devices that belong to a certain group can be seen by the host initiators of the group. A device can belong to multiple LUN Masking groups, but a host initiator can only belong to one LUN Masking group.

**Note:** A maximum of 512 groups can be configured per PT system and a maximum of 512 drives can be configured per group. (A maximum of 1024 host initiators can be defined on a system.) Each group needs to contain at least one host initiator and one device (that is, drive or robot). Robots can be added as required.

To configure the LUN Masking groups, complete the following steps:

1. From the PT Manager GUI, click the VT drop-down menu, select **LUN Masking**, and select **Configure LUN Masking groups**, as shown in Figure 6-7.

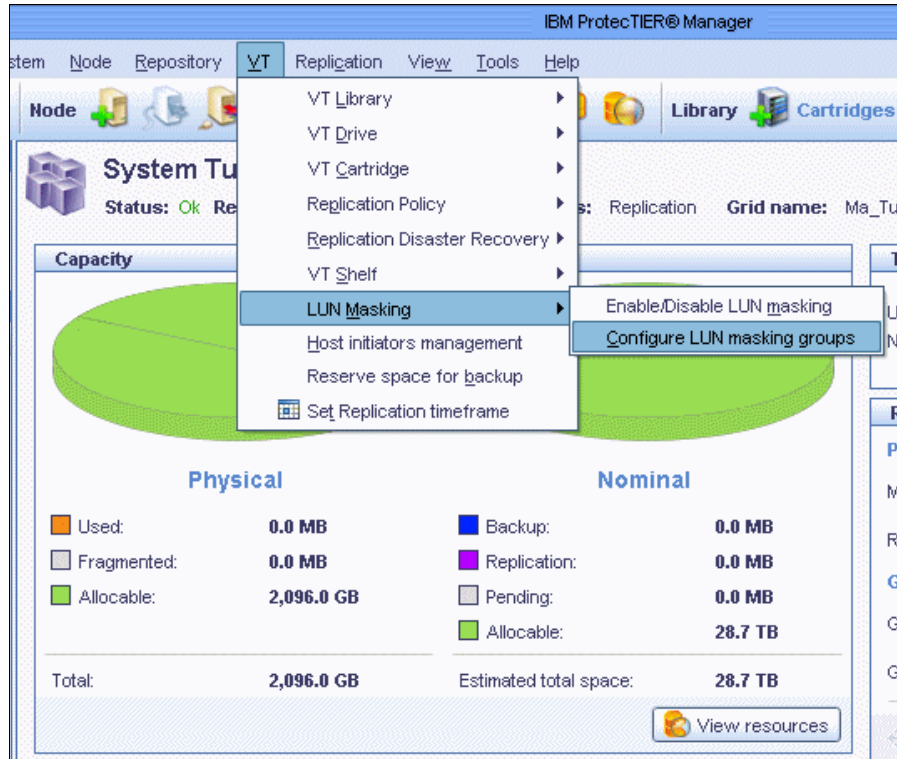


Figure 6-7 Select Configure LUN Masking groups



3. Enter a new name for the group in the Group name field, as shown in Figure 6-9.




Figure 6-9 LUN Masking





6. If the zoning has been done correctly, the available Host Adapters WWPN appears after it has rescanned the ports, as shown in Figure 6-12. Alternatively, you can add the Host Adapters WWPN manually by clicking **Add** and entering the WWPN manually.

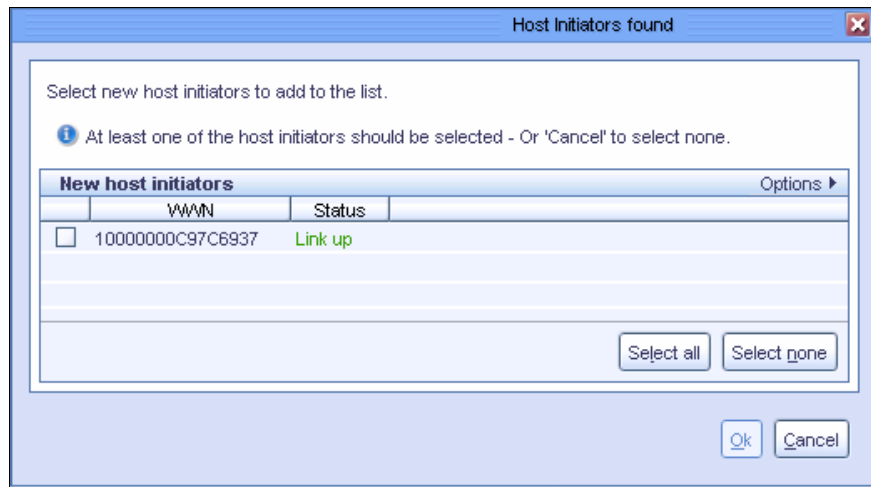


Figure 6-12 Host Initiators found window

7. Select the WWPN that you want to add to the list. Click **OK**, as shown in Figure 6-13. When the Host Initiators found window closes, we return to the Host Initiator Management window.

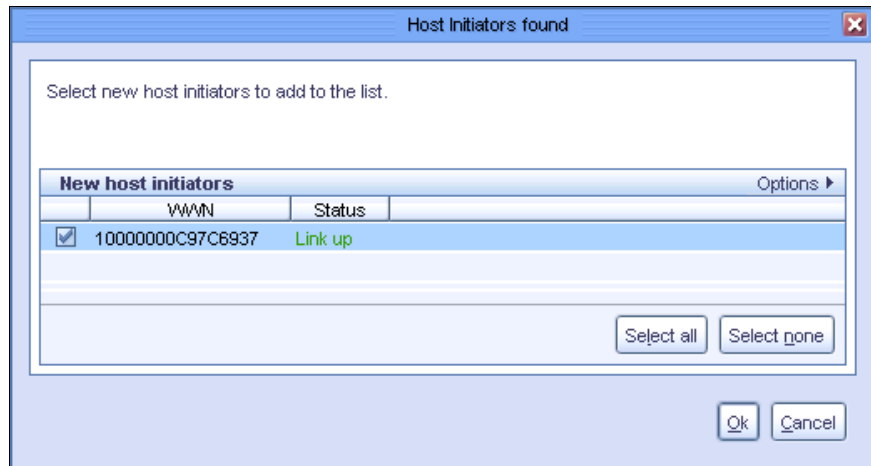


Figure 6-13 Host Initiators found window



8. The list of Host Initiators is displayed, as shown in Figure 6-14. Select the Host Initiators you want to add to the LUN Masking group from the list displayed and click **OK**.

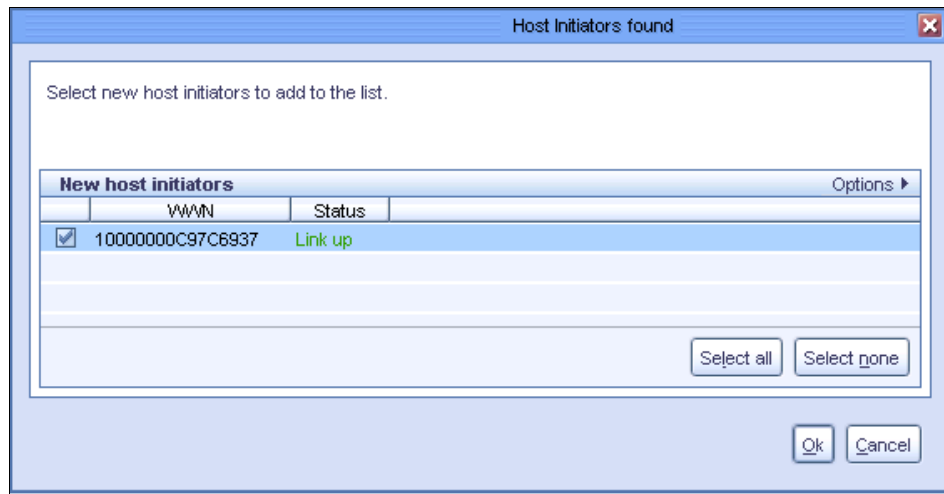


Figure 6-14 Host Initiators found window

9. After adding the Host Initiators, go to the Library Mappings pane and click **Add** (as shown in Figure 6-15) to select the tape library devices you want to be visible to the host initiators in the group.

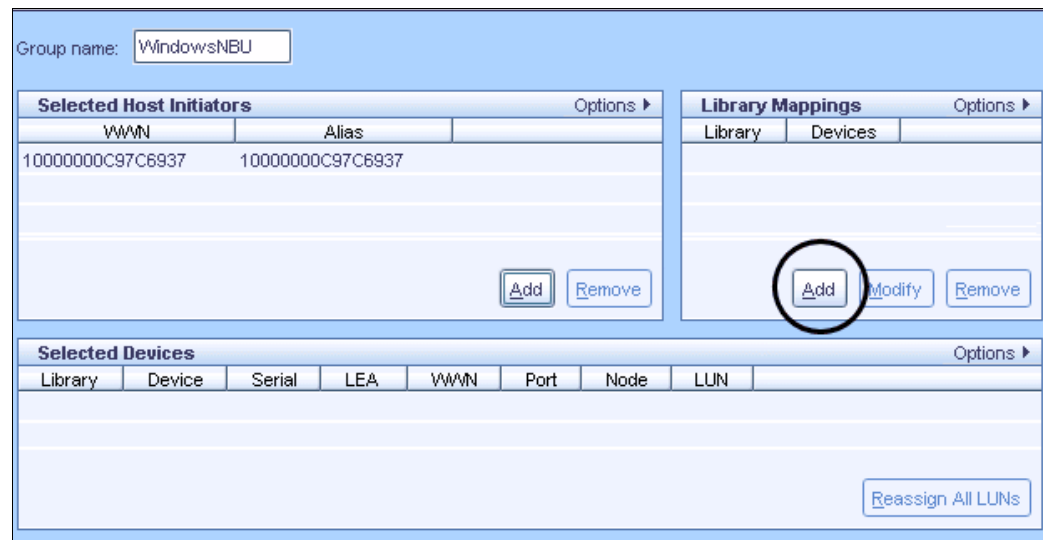


Figure 6-15 LUN Masking window

10. The Library devices selection window opens, as shown in Figure 6-16. A list of libraries in the system is displayed in the Library drop-down list. Select the relevant libraries from the list. Select the tape robot and drives that you want to add to this group and click **OK**.

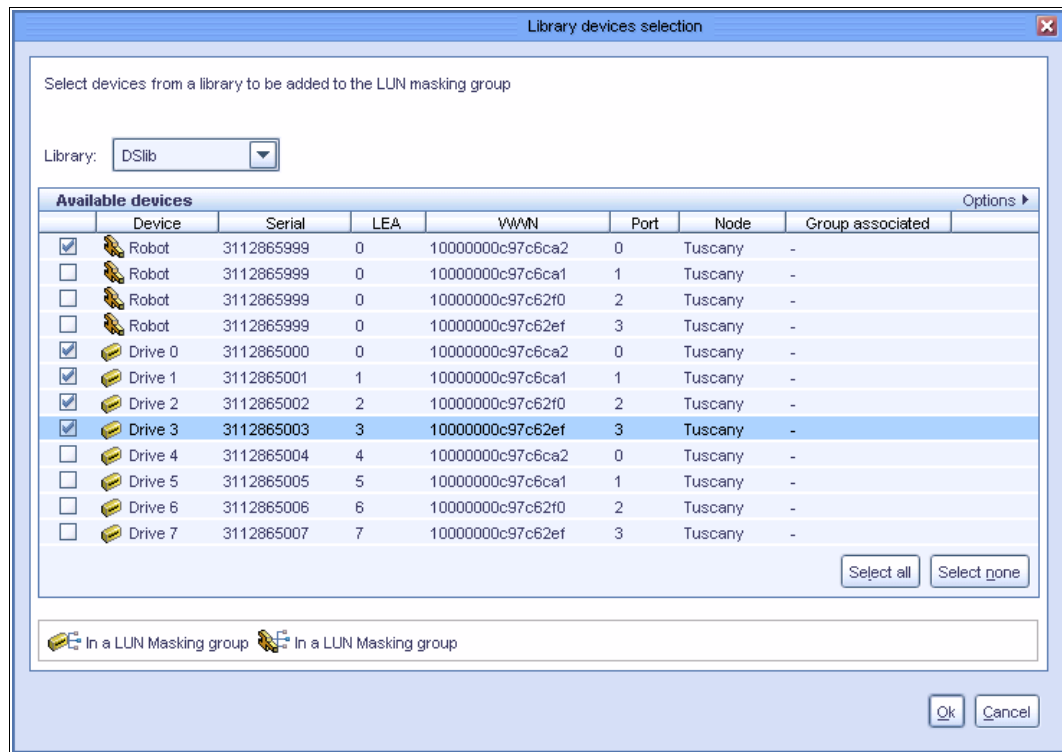


Figure 6-16 Library devices selection window

- The LUN Masking window should show the tape devices selected for the group, as shown in Figure 6-17. Confirm that the Host Initiators and tape devices are correct and click **Save Changes**.

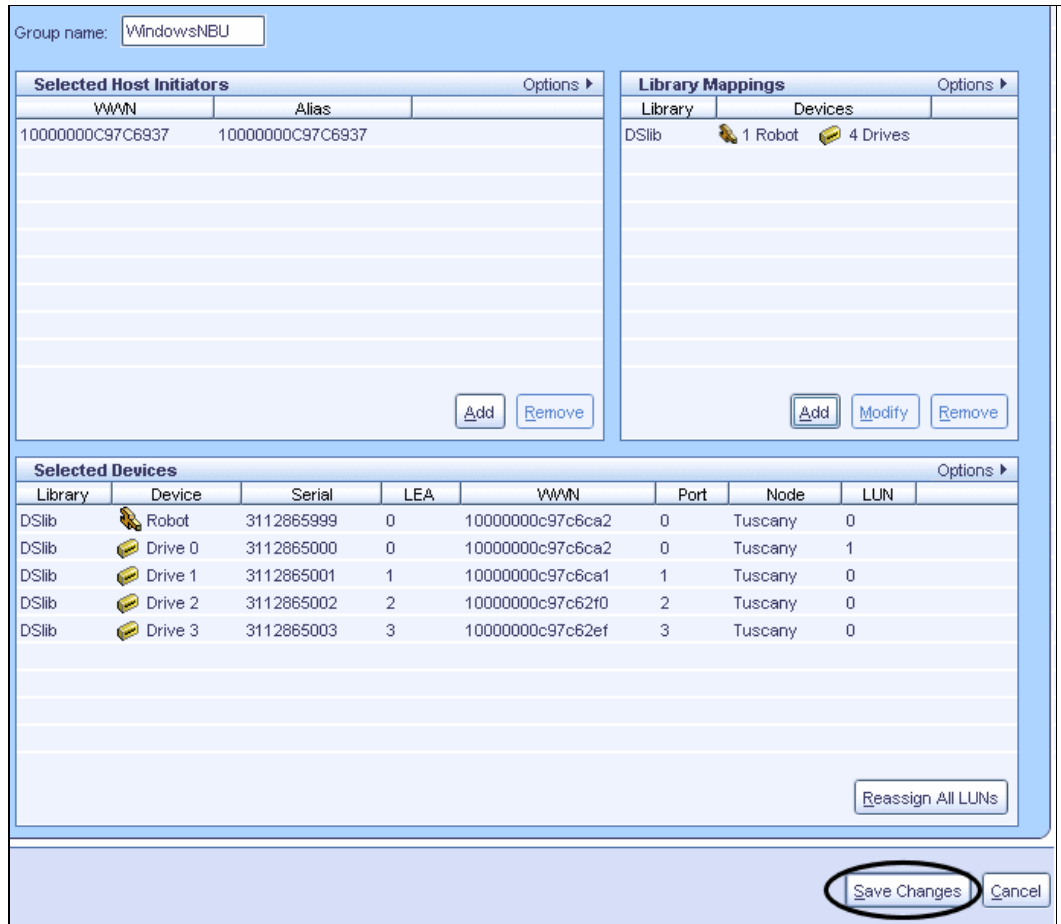


Figure 6-17 LUN Masking window

### Modifying a LUN Masking group

From the Selected Devices pane, you can modify the libraries and devices defined in the LUN Masking group. After you have finished, click **Save changes** to save your modifications and close the LUN Masking window.

After modifying a LUN Masking group, undesired holes may occur within the LUN numbering sequence. For example, removing a device from an existing group will cause holes in the LUN numbering if this device does not have the highest LUN number and, therefore, the backup application may have trouble scanning the devices.

If your backup application does not have trouble scanning the devices, you should re-enumerate the LUN by completing the following steps:

1. Reassign the LUNs by clicking **Reassign All LUNs** at the bottom of the Selected Devices pane, as shown in Figure 6-18.
2. The Reassign LUNs window opens and informs you that you are about to renumber all the LUN values of the available devices in the group and the connected hosts will be rescanned.
3. Click **Yes** to renumber. The LUN values will be renumbered and all the devices in the Selected Devices pane will be assigned new LUN numbers, sequentially, starting with "0" (zero).

**Note:** After modifying a LUN Masking group, undesired holes may occur within the LUN numbering sequence and the backup application may have trouble scanning the device. You should re-enumerate the LUNs.




Figure 6-18 Reassign All LUNs in LUN Masking



## Library-related operations

Some library configuration operations may affect the LUN Masking groups. Most of these changes will be blocked by the system. The only exception to this is the Delete library operation and reducing the number of drives through the Change library dimensions operations where the removal of the devices might directly affect the relevant LUN Masking group.

### *Deleting a library*

If devices from a deleted library are members of LUN Masking groups, a warning is issued to the log file and to the user (through the GUI). If the user chooses to proceed with the operation, the relevant devices are removed from the relevant LUN Masking groups, and the configuration file data is updated accordingly.

Removing a library can cause holes in the LUN Masking numbering and the backup application may have trouble scanning the devices. In this case, you must reassign the LUN. Refer to “Modifying a LUN Masking group” on page 293 for information about how to reassign LUN numbers.

**Note:** When you modify the dimensions of a library or reassign the library devices, if a device already belongs to a LUN Masking group, the following rules apply:

- ▶ The device cannot be deleted or reassigned to another Fibre Channel port. To reassign the device to a different port, you must remove the device from all of its related groupings.
- ▶ The number of drives in the library cannot be reduced.

## 6.1.6 Deduplication rate consideration

Backup and replication may run concurrently on the local repository. The best practice is to use separate time windows for the backup and the replication operations. ProtecTIER balances priorities of backup and replication activity based on an algorithm that, in most cases, prioritizes backup over replication. If you are upgrading to ProtecTIER V2.5 or later and plan to do replication, you must resize your environment to ensure the additional, and possibly competing, workload does not exceed the capabilities of the ProtecTIER system.

## 6.1.7 Path failover

Path failover may be used to help enhance availability.

### **Control path failover**

Control path failover is designed to provide automatic control path failover to a preconfigured redundant control path in the event of a loss of a host adapter or control path drive, without aborting the current job in progress. This function might also be referred to as data path failover (DPF). Support is provided under various operating systems, such as AIX, Linux, Solaris, HP-UX, and Windows for Fibre Channel attachments when the IBM device driver is used.

You can enable and disable the DPF in ProtecTIER Manager by each virtual tape library. Select the virtual tape library that you want to change and then select **VT** → **VT Library** → **Set control path failover mode**. The Set control path failover mode window opens. You can enable or disable this function (Figure 6-20).



Figure 6-20 Set control path failover mode

## Data path failover

Data path failover (DPF) supports virtual Ultrium Tape Drives in the TS3500 virtual tape library using the IBM device driver for AIX, Linux, Windows, Solaris, and HP-UX. Data path failover is designed to provide a failover mechanism in the IBM device driver to enable configuration of multiple redundant paths in a SAN environment. In the event of a path or component failure, the failover mechanism is designed to automatically provide error recovery to retry the current operation using an alternate, preconfigured path without aborting the current job in progress. This provides flexibility in SAN configuration, availability, and management.

**Note:** ProtecTIER path failover functions are not supported with Quantum P3000 virtual tape library definitions.

## 6.2 Installing and configuring the device driver in OS

In this section, we describe how to install the device driver in AIX and Windows. For other OS types, the information can be found in the following manuals:

- ▶ *IBM Tape Device Drivers: Installation and User's Guide*, GC27-2130
- ▶ *Implementing IBM Tape in UNIX Systems*, SG24-6502

### 6.2.1 Getting device drivers

The IBM device driver must be installed on all operating systems that will use ProtecTIER virtual drives and libraries. The latest device driver code can be downloaded from the IBM website at the following address:

<http://www-01.ibm.com/support/docview.wss?rs=577&uid=ssg1S4000784>

Or you can go to <http://www.ibm.com> and select **Support & Downloads** → **Download** → **Fixes, updates & drivers**, choose the product, click **Go** on the Support for TS7650G or TS7650 with ProtecTIER page, then select **Device drivers**.

There are additional pages for tape device driver downloads:

- ▶ AIX
- ▶ HP-UX

- ▶ Linux
- ▶ Solaris
- ▶ Microsoft Windows

Go to the page for the OS that you are using. From here you can download the following documentation:

- ▶ Fixlist
- ▶ README
- ▶ *IBM Tape Device Drivers Installation and User's Guide, GC27-2130*
- ▶ *IBM Tape Device Drivers Programming Reference, GA32-0566*

Then you can go to the device driver ftp site by choosing **IBM Tape Device Drivers FTP site** from the list of directories displayed and select your specific device driver.

## 6.2.2 Installing IBM tape device drivers for AIX

Complete the steps described in 6.2.1, "Getting device drivers" on page 297 to go to the download site for the device drivers. In our environment, we chose the AIX directory. You will then be presented with a list of files similar to Example 6-3.

*Example 6-3 Files*

---

```
Atape.10.7.3.0.bin
Atape.11.6.0.0.bin
Atape.8.4.9.0.bin
Atape.9.7.5.0.bin
Atape.fixlist
Atape.README
atldd.5.5.1.0.bin
atldd.6.7.7.0.bin
atldd.fixlist
atldd.Readme
README
```

---

Table 6-3 provides the definitions for each of the device driver files.

*Table 6-3 File definitions*

Device driver file	Description
Atape.n.n.n.n.bin	Atape device driver for IBM TotalStorage and Ultrium products
Atape.fixlist	Fixlist for Atape
Atape.README	Readme information concerning Atape driver
atldd.n.n.n.n.bin	3494 Library Daemon
atldd.fixlist	Fixlist for 3494 Library Daemon
atldd.Readme	Readme information concerning atldd driver



You must download the correct version of `Atape.n.n.n.n.bin`. Refer to Table 6-4 to decide which one you need.

Table 6-4 *Atape levels*

Atape level	AIX version
11.x.x.x	AIX Version 5.2 and later
10.x.x.x	AIX Version 5.x with Data Encryption
9.x.x.x	AIX Version 5.1 and later
8.x.x.x	AIX Versions 4.3, 5.1, and 5.2
6.x.x.x and 7.x.x.x	AIX Versions 4.3 and 5.1

Issue the `oslevel` command in AIX to get the OS version, as shown in Example 6-4.

Example 6-4 *Obtaining the OS version*

---

```
dili5> oslevel
5.3.0.0
dili5> oslevel -r
5300-09
```

---

Download the driver in binary to your AIX host. The following instructions assume that you named the file `/tmp/Atape.x.x.x.x.bin`.

### Atape driver installation using the command-line interface

Assuming that you have downloaded the driver to your local system, run the following command:

```
installp -acXd /directory/filename Atape.driver
```

In our environment, we downloaded the `Atape.11.6.0.0.bin` file to the `/temp` directory. The command string is:

```
installp -acXd /temp/Atape.11.6.0.0.bin Atape.driver
```

This command installs and commits the Atape driver in your system. Example 6-5 shows the example `installp` command output.

Example 6-5 *installp command output*

---

```
dili5> installp -acXd /temp/Atape.11.6.0.0.bin Atape.driver
+-----+
                        Pre-installation Verification...
+-----+
Verifying selections...done
Verifying requisites...done
Results...

SUCSESSES
-----
  Filesets listed in this section passed pre-installation verification
  and will be installed.

  Selected Filesets
  -----
```

Atape.driver 11.6.0.0 # IBM AIX Enhanced Tape and Me...

<< End of Success Section >>

+-----+  
BUILDDATE Verification ...

+-----+  
Verifying build dates...done

FILESET STATISTICS

-----

1 Selected to be installed, of which:  
1 Passed pre-installation verification

----

1 Total to be installed

0503-409 installp: bosboot verification starting...

installp: bosboot verification completed.

+-----+  
Installing Software...

+-----+

installp: APPLYING software for:  
Atape.driver 11.6.0.0

. . . . << Copyright notice for Atape >> . . . . .

IBM AIX Enhanced Tape and Medium Changer Device Driver

(C) COPYRIGHT International Business Machines Corp. 1993 2005  
All Rights Reserved  
Licensed Materials - Property of IBM

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

. . . . << End of copyright notice for Atape >>. . . .

Checking for existing Atape devices...

Installing AIX Version 5.3 Atape.driver...

Adding device prototype...

Adding odm and smit entries...

Adding catalogs...

Adding trace template...

Adding error template...

6 entries added.

0 entries deleted.

0 entries updated.

Adding utility programs...

Finished processing all filesets. (Total time: 10 secs).

0503-409 installp: bosboot verification starting...

installp: bosboot verification completed.

0503-408 installp: bosboot process starting...

bosboot: Boot image is 38789 512 byte blocks.  
0503-292 This update will not fully take effect until after a  
system reboot.

\* \* \* A T T E N T I O N \* \* \*  
System boot image has been updated. You should reboot the  
system as soon as possible to properly integrate the changes  
and to avoid disruption of current functionality.

installp: bosboot process completed.

+-----+  
Summaries:  
+-----+

#### Installation Summary

Name	Level	Part	Event	Result
Atape.driver	11.6.0.0	USR	APPLY	SUCCESS

### Installing using SMIT

Start the System Management Interface Tool (SMIT) by entering `smit` at the command line. Choose the following options: Software Installation and Maintenance Install and Update Software Install and Update from all Available Software (includes devices and printers).

We show the ASCII SMIT interface. However, if you have an X Window System display, you will see the GUI version. Both have identical functions. Your screen will look similar to Example 6-6. Enter the directory where you downloaded the driver (or `/dev/cd0` if on a CD).

#### Example 6-6 Atape installation using SMIT

Install and Update from ALL Available Software

Type or select a value for the entry field.  
Press Enter AFTER making all desired changes.

[Entry Fields]  
\* INPUT device / directory for software [ /temp ] +

Press Enter. A list displays. Select Atape and press Enter again. Another screen opens with a line similar to:

@ 11.6.0.0 IBM AIX Enhanced Tape and Medium Changer Device Driver

Select this line by pressing F7 and Enter. Complete the installation options on the next screen and change the ACCEPT new license agreements field to yes, as illustrated in Example 6-7.

#### Example 6-7 Accepting new license agreements

Install and Update from ALL Available Software

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

[Entry Fields]

```

* INPUT device / directory for software
.
* SOFTWARE to install
[@ 11.6.0.0 IBM AIX E> +
  PREVIEW only? (install operation will NOT occur)          no
+
  COMMIT software updates?
yes                  +
  SAVE replaced files?
no                   +
  AUTOMATICALLY install requisite software?
yes                  +
  EXTEND file systems if space needed?
yes                  +
  OVERWRITE same or newer versions?                          no
+
  VERIFY install and check file sizes?
no                   +
  DETAILED output?
no                   +
  Process multiple volumes?
yes                  +
  ACCEPT new license agreements?
yes                  +
  Preview new LICENSE agreements?
no                   +

F1=Help             F2=Refresh         F3=Cancel          F4=List
F5=Reset            F6=Command         F7=Edit            F8=Image
F9=Shell            F10=Exit           Enter=Do

```

---

After this selection is complete, the installation process begins. You see the same messages as in Example 6-5 on page 299.

You are now ready to configure the devices.

### Configuring tape and medium changer devices

After the driver software is installed and the tape drives are connected to the server, the device can be configured and made available for use. You cannot access the devices until this step is completed.

Configure a tape device by using either of the following procedures:

- ▶ Enter the following command without parameters:
 

```

      cfgmgr
      
```

 This command configures all devices automatically, and will make available any new tape or medium changer devices.
- ▶ Power off your system and reboot to configure all devices automatically during the startup, and make available any new tape or medium changer devices in the system.

## Control path failover support for tape libraries

The Atape device driver path failover support configures multiple physical control paths to the same logical library within the device driver and provides automatic failover to an alternate control path when a permanent error occurs on one path. This is transparent to the running application.

### ***Configuring and unconfiguring path failover support***

Path failover support is not enabled automatically when the device driver is installed. It must be configured initially on each logical device after installation. When path failover support is enabled for a logical device, it remains set until the device is deleted or unconfigured. The alternate path failover setting is retained even if the system is rebooted.

To enable or disable the support on a single logical device, use the SMIT menu to Change/Show Characteristics of a Tape Drive, select the logical device to change, such as smc0, smc1, and so on, then select Yes or No for Enable Path Failover Support. The support can also be enabled or disabled by using the **chdev** command. In Example 6-8, we turn the CPF on and then turn it off.

#### *Example 6-8 Path failover support*

---

```
highway> lsdev -Cc tape | grep smc
smc0 Available 06-09-02      IBM 3584 Library Medium Changer (FCP)
smc1 Available 0B-09-02      IBM 3584 Library Medium Changer (FCP)
highway> chdev -l smc0 -aalt_pathing=yes
smc0 changed
highway> chdev -l smc1 -aalt_pathing=yes
smc1 changed
highway> 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)
highway> chdev -l smc0 -aalt_pathing=no
smc0 changed
highway> chdev -l smc1 -aalt_pathing=no
smc1 changed
highway> lsdev -Cc tape | grep smc
smc0 Available 06-09-02      IBM 3584 Library Medium Changer (FCP)
smc1 Available 0B-09-02      IBM 3584 Library Medium Changer (FCP)
highway>
```

---

### ***Primary and alternate paths***

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

When a second logical device is configured with path failover support enabled for the same physical device, it configures as an alternate 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. A third logical device is also configured as an alternate path with either -A or -ALT appended, and so on. The device driver supports up to 16 physical paths for a single device.

The labeling of a logical device as either a primary or an alternate path is for information only, to:

- ▶ Be able to identify the actual number of physical devices configured on the system and a specific logical device associated with them. There is only one logical device labeled as the primary path for each physical device. However, there can be many (multiple) logical devices labeled as an alternate path for the same devices.
- ▶ Provide information about which logical devices configured on the system have path failover support enabled.

### ***Querying primary and alternate path configurations***

You can display the primary and alternate path configuration for all devices with the `lsdev` command. There can be two or more logical devices configured for a single physical device, but the first device configured is labeled the primary device. All other logical devices configured after the first device are labeled as alternate devices. To see this configuration, run the `lsdev -Cc tape` command and look at the location field in the data. Run the `lsdev -Cc tape` command (Example 6-9).

#### *Example 6-9 lsdev output*

---

```
highway> lsdev -Cc tape
rmt0 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt1 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt2 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt3 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt4 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt5 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt6 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt7 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt8 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt9 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt10 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt11 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt12 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt13 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt14 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt15 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
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 and unconfiguring primary and alternate devices***

Logical devices configured as alternate paths can be unconfigured and reconfigured at any time after the initial configuration is run. Unconfiguring an alternate path device removes that device from the primary device path list, removes the `-A` or `-ALT` appended to the location field, and changes the device to the defined state. The primary and any other alternate devices are still available.

Likewise, configuring a new alternate path device or reconfiguring an existing one in the defined state adds that device to the primary device path list, appends `-A` or `-ALT` to the location field, and makes the device available.

Logical devices that are configured as primary paths can also be unconfigured and reconfigured at any time after initial configuration is run. However, the operation is different for a primary device. When a primary device is unconfigured, the following events occur:

- ▶ All alternate devices are unconfigured, as described previously.
- ▶ The primary device is unconfigured.
- ▶ The -P or -PRI appended to the location field is removed.
- ▶ The device is changed to the defined state.
- ▶ All alternate devices that were unconfigured are reconfigured. The first device that is reconfigured becomes the new primary device. All remaining alternate devices are reconfigured as alternate paths.

These methods provide the ability to unconfigure and reconfigure physical devices on the system when device connections or addressing changes are made.

### **Data path failover and load balancing support for tape drives**

The AIX operating system only supports a static configuration of devices, which also applies to the path failover and failover support. When devices are initially configured at a specific SCSI ID and physical connection (drive port, host bus adapter, and switch number/port, if applicable) and in the available state, changing the physical device address/connection without either rebooting or unconfiguring and reconfiguring the devices has unpredictable results and is not supported.

#### ***Installing the data path failover license key***

Use the following command line scripts to query, add, or delete license keys for this feature before enabling the path failover feature as described below. The key is a 16-digit hexadecimal value, for example, 1234567890abcdef.

All key values A–F should be entered in lowercase letters (a–f).

- ▶ Query installed keys: `dpf_keys`
- ▶ Install a license key: `dpf_keys -a key`
- ▶ Delete a license key: `dpf_keys -d key`

#### ***Configuring and unconfiguring path failover support***

Path failover support is not enabled automatically when the device driver is installed. It must be configured initially on each logical device after installation. When path failover support is enabled for a logical device, it remains set until the device is deleted or the support is unconfigured. The path failover setting is retained even if the system is rebooted.

Path failover support can be enabled on all configured devices at one time, or it can be enabled or disabled selectively by logical device. It might be desirable at times to configure some, but not all, logical paths to a device with the support enabled.

To enable the support globally on all currently configured devices, run the following command:

```
/usr/lpp/Atape/instAtape -a
```

This command unconfigures all devices that have path failover set to no and reconfigures all devices, setting path failover to yes (Example 6-10).

#### ***Example 6-10 Enabling support***

---

```
highway> lsdev -Cc tape  
rmt0 Available 06-09-02 IBM 3580 Ultrium Tape Drive (FCP)
```

```

rmt1 Available 06-09-02      IBM 3580 Ultrium Tape Drive (FCP)
rmt2 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt3 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt4 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt5 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt6 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt7 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt8 Available 0B-09-02      IBM 3580 Ultrium Tape Drive (FCP)
rmt9 Available 0B-09-02      IBM 3580 Ultrium Tape Drive (FCP)
rmt10 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt11 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt12 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt13 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt14 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt15 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
smc0 Available 06-09-02      IBM 3584 Library Medium Changer (FCP)
smc1 Available 0B-09-02      IBM 3584 Library Medium Changer (FCP)
highway> /usr/lpp/Atape/instAtape -a
Setting alternate pathing support on smc0...
smc0 changed
Setting alternate pathing support on rmt0...
rmt0 changed
Setting alternate pathing support on rmt1...
rmt1 changed
Setting alternate pathing support on smc1...
smc1 changed
Setting alternate pathing support on rmt8...
rmt8 changed
Setting alternate pathing support on rmt9...
rmt9 changed
highway> lsdev -Cc tape
rmt0 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt1 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt2 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt3 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt4 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt5 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt6 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt7 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt8 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt9 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt10 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt11 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt12 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt13 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt14 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt15 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
smc0 Available 06-09-02-PRI IBM 3584 Library Medium Changer (FCP)
smc1 Available 0B-09-02-ALT IBM 3584 Library Medium Changer (FCP)

```

---



To enable or disable the support on a single logical device, use the SMIT menu to Change/Show Characteristics of a Tape Drive, then select Yes or No for Enable Path Failover Support. The support can also be enabled or disabled using the `chdev` command (Example 6-11):

```
chdev -l rmt0 -aalt_pathing=yes
chdev -l rmt0 -aalt_pathing=no
```

*Example 6-11 Enabling single drive support*

---

```
highway> chdev -l rmt0 -aalt_pathing=no
rmt0 changed
highway> chdev -l rmt8 -aalt_pathing=no
rmt8 changed
highway> lsdev -Cc tape
rmt0 Available 06-09-02 IBM 3580 Ultrium Tape Drive (FCP)
rmt1 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt2 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt3 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt4 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt5 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt6 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt7 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt8 Available 0B-09-02 IBM 3580 Ultrium Tape Drive (FCP)
rmt9 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt10 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt11 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt12 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt13 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt14 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt15 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
smc0 Available 06-09-02-PRI IBM 3584 Library Medium Changer (FCP)
smc1 Available 0B-09-02-ALT IBM 3584 Library Medium Changer (FCP)
```

---

**Primary and alternate paths**

When the device driver configures a logical device with path failover support enabled, the first device configured always becomes the primary path and PRI is appended to the location field of the device. When a second logical device is configured with path failover support enabled for the same physical device, it configures as an alternate path and ALT is appended to the location field. A third logical device is configured as the next alternate path with ALT appended, and so on. The device driver supports up to 16 physical paths for a single device. For example, suppose that in our system rmt1 and rmt9 are a pair. If rmt1 is configured first, then rmt9, the `lsdev -Cc tape` command output is similar to the following:

```
rmt1 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt9 Available 0B-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
```

If rmt1 is configured first, then rmt9, the command output is similar to the following:

```
rmt1 Available 06-09-02-ALT IBM 3580 Ultrium Tape Drive (FCP)
rmt9 Available 0B-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
```

The labeling of a logical device as either a primary or alternate path is for information only, in order to:

- ▶ Be able to identify the actual number of physical devices configured on the system and a specific logical device associated with them. There is only one logical device labeled the primary path for each physical device. However, there might be many (multiple) logical devices labeled as an alternate path for the same devices.
- ▶ Provide information about which logical devices configured on the system have path failover support enabled.

### ***Querying primary and alternate path configuration***

You can display the primary and alternate path configuration for all devices with the `lsdev` command. There might be two or more logical devices configured for a single physical device, but the first device configured is labeled the primary device. All other logical devices configured after the first device are labeled as alternate devices.

To see this configuration, run the `lsdev -Cc tape` command and look at the location field in the data. By running `lsdev -Cc tape | grep PRI` (Example 6-12), you can easily determine how many physical devices on the IBM RS/6000® or IBM System p® server are configured with path failover support.

#### *Example 6-12 Querying devices*

---

```
highway> lsdev -Cc tape | grep PRI
rmt1 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt2 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt3 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt4 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt5 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt6 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
rmt7 Available 06-09-02-PRI IBM 3580 Ultrium Tape Drive (FCP)
smc0 Available 06-09-02-PRI IBM 3584 Library Medium Changer (FCP)
```

---

### ***Configuring and unconfiguring primary and alternate devices***

Logical devices configured as alternate paths can be unconfigured and reconfigured at any time after the initial configuration is run. Unconfiguring an alternate path device removes that device from the primary device path list, removes the ALT appended to the location field, and changes the device to the defined state. The primary and any other alternate devices are still available.

Likewise, configuring a new alternate path device or reconfiguring an existing one in the defined state adds that device to the primary device path list, appends ALT to the location field, and makes the device available.

Logical devices that are configured as primary paths can also be unconfigured and reconfigured at any time after the initial configuration is run. However, the operation is different for a primary device. When a primary device is unconfigured, the following events occur:

1. All alternate devices are unconfigured as described previously.
2. The primary device is unconfigured.
3. The PRI appended to the location field is removed.
4. The device is changed to the defined state.

5. All alternate devices that were unconfigured are reconfigured. The first device that is reconfigured becomes the new primary device. All remaining alternate devices are reconfigured as alternate paths.

These methods provide the ability to unconfigure and reconfigure physical devices on the system when device connections or addressing changes are made.

### 6.2.3 Installing IBM tape device drivers for Windows 2003 and 2008

The Windows tape and medium changer device driver is designed specifically to take advantage of the features provided by the IBM tape drives and medium changer devices. The goal is to give applications access to the functions required for basic tape operations (such as backup and restore) and medium changer operations (such as mount and unmount the cartridges), as well as to the advanced functions needed by full tape management systems. Whenever possible, the driver is designed to take advantage of the device features transparent to the application.

The software described in this chapter covers the Windows device driver and the interface between the application and the tape device.

#### Downloading the drivers

Complete the steps described in 6.2.1, “Getting device drivers” on page 297 to go to the download site of the device drivers. In our environment, we choose the Windows directory. Chose the Windows200X depending on the Windows version you are using. Then open the Latest directory, and find the files:

- ▶ IBMTape.x86\_6xxx.zip: Used for Windows 200x running on x86
- ▶ IBMTape.i64\_6xxx.zip: Used for Windows 200x running on IA64
- ▶ IBMTape.x64\_6xxx.zip: Used for Windows 200x running on AMD64 and EM64T

Download one of these to your host computer, according to your hardware platform.

## Extracting the drivers

The IBM tape device driver will be downloaded as a compressed file. Unzip the package. We used the C:\IBMdrv\IBMTape.x86\_6201.zip directory, as shown in Figure 6-21.

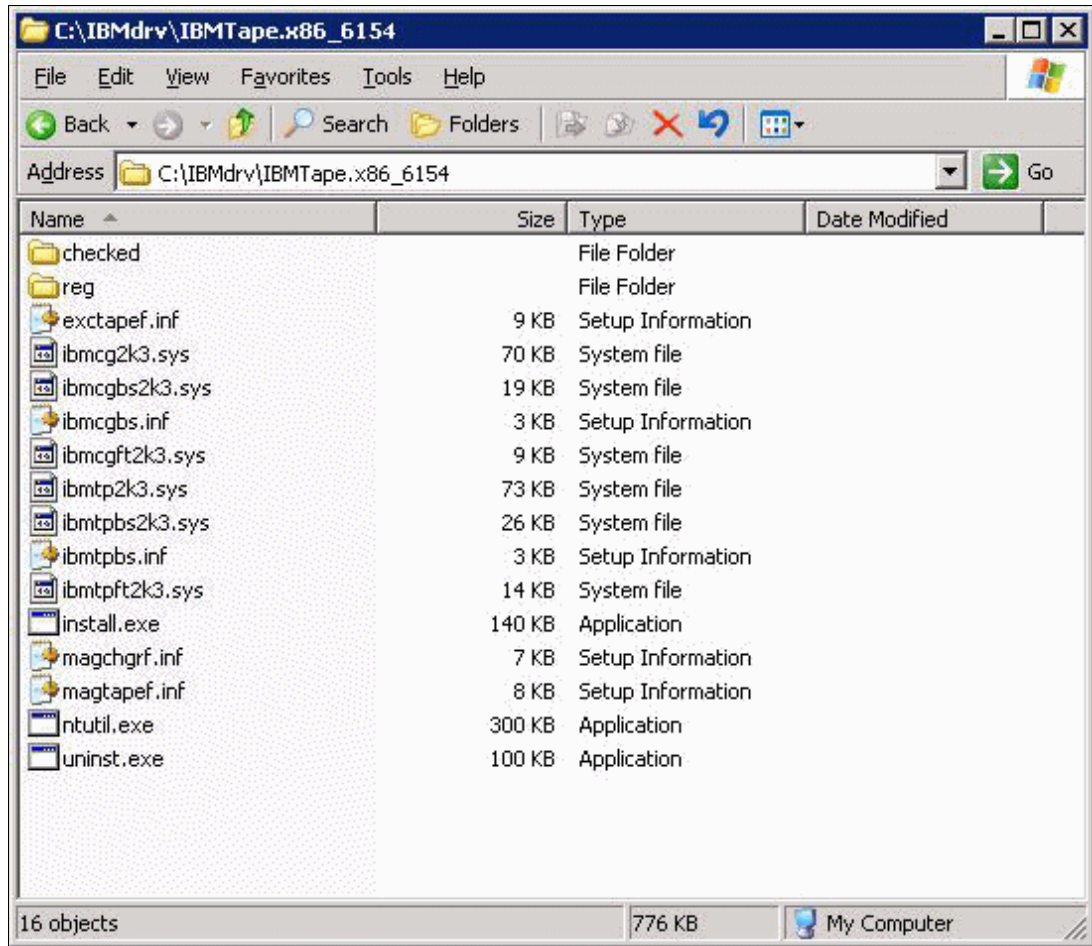


Figure 6-21 Extracting the drivers

## Detecting the added devices

If the tape library is physically attached, Windows 200x attempts to automatically install the drivers upon boot. The Driver Installation Wizard appears and requests details regarding the location of the drivers and information files. The IBM tape library may be hot-plugged into the SCSI or Fibre adapter. In this case, you might need to perform a hardware scan from Device Manager, as shown in Figure 6-22, to detect the added devices and initiate the Device Driver Install Wizard.

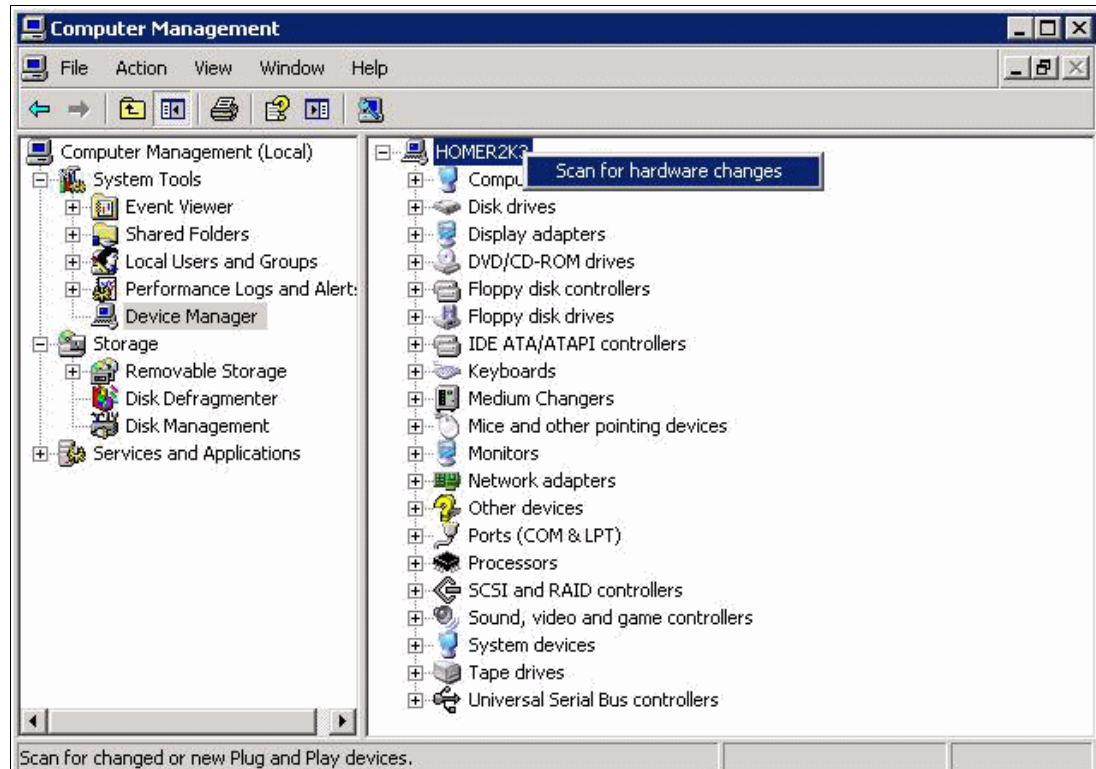


Figure 6-22 Detecting added devices

## Uninstalling the previous version

If an older version of IBM tape device driver is already installed on the system, uninstall it first.

## Installing the medium changer and tape device driver

In order to reduce complexity, the `install.exe` file is included in the driver package to perform the installation steps automatically. Simply double-click the `install.exe` icon from within the driver install directory. The virtual and actual hardware will be installed correctly on your system. For more details, refer to the *IBM Tape Device Drivers: Installation and User's Guide*, GC27-2130.

**Note:** If the Windows Found New Hardware Wizard begins during installation, cancel the wizard. `install.exe` will perform the necessary steps.

install.exe prompts you to select between a driver customized for IBM Tivoli Storage Manager and a generic driver that should be used if you intend to run other applications, such as Microsoft Removable Storage Manager or Microsoft Backup, as shown in Example 6-13.

**Note:** In most cases, select exctapef.inf.

*Example 6-13 Selecting the driver*

---

Running installation application, version 6.x.x.x.

Select the driver file to be used for tape drives.

- 1) exctapef.inf (default)
- 2) magtapef.inf

This file installs a driver that has been customized for IBM Tivoli Storage Manager (TSM). If you are running other applications, such as Microsoft Removable Storage Manager or Microsoft Backup, then you should select magtapef.inf. For more information, consult the README file on the CD or FTP site to determine which driver to use.

Select number, or press Enter for default: 1  
exctapef.inf has been selected

Preparing system for driver installation...  
Uninstalling existing medium changers...  
Existing medium changers successfully uninstalled.  
Uninstalling existing tape drives...  
Existing tape drives successfully uninstalled.  
Installing/updating the changer bus enumerator...  
Successfully installed/updated the changer bus enumerator.  
Installing/updating the tape bus enumerator...  
Successfully installed/updated the tape bus enumerator.  
Installing driver for medium changers...  
Driver for medium changers installed.  
Installing driver for tape drives...  
Driver for tape drives installed.

Program successful.

---

You will see a Program Status window with the message that the installation was successful. Click **OK**. You should be able to verify that the device was installed correctly. Repeat this procedure for every device that you install. Remember that if you are installing a tape library you must install drivers for both the tape drives and the medium changer.

**Note:** To reduce installation complexity, the file install.exe is included in the driver package to perform the installation steps automatically. Beginning in Version 6.1.6.7 of the device driver, the non-exclusive option for use with Tivoli Storage Manager is implemented through an installer option. Double-click install.exe or run install.exe without any option if the user runs IBM Tivoli Storage Manager.

Run install.exe with the -n option from a command line if you want to run with RSM (for example, C:\>install -n).

If you are installing a driver that has not been certified by the Microsoft Windows Hardware Quality Laboratories (WHQL), you will be presented with a warning window. If you want to continue installing the driver, select **Continue Anyway**.

## **Control path failover support for tape libraries**

To take advantage of control path failover (CPF) support in Windows, the appropriate feature code must be installed.

### ***Configuring and unconfiguring control path failover support***

Control path failover support is enabled automatically when the device driver is installed. It may be disabled or re-enabled for the entire set of attached medium changers by modifying the registry. To modify the registry, complete the following steps:

1. Open the reg folder of the driver package.
2. Double-click `DisableCPF.reg` or `EnableCPF.reg`.
3. Reboot the system. This is necessary for any registry modification to take effect.

### ***Querying primary and alternate path configuration***

To check whether the control path failover has been enabled in the device driver and to display the primary and alternate paths, use the tape diagnostic and utility tool.

**Note:** Display the primary and alternate path configuration for any device using tape diagnostic and utility functions. Refer to *IBM Tape Device Drivers: Installation and User's Guide*, GC27-2130.

### ***Checking disablement of the control path failover setting***

If you have disabled the control path failover in the device driver's setting by double-clicking the `DisableCPF.reg` file and rebooting your system, you may go into the registry by issuing the Windows **regedit** command to confirm that CPF has been disabled. Look for a line similar to the following if your system is Windows Server 2003 or Windows Server 2008:

```
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ibmccg2kx]
"FailoverDisabled"=dword:00000001
```

This indicates that CPF has been disabled in the driver. This setting only takes effect after your system is rebooted.

## **Data path failover support for tape drives**

To take advantage of DPF support in Windows, the appropriate feature code must be installed.

### ***Configuring and unconfiguring data path failover support***

Data path failover support is enabled automatically when the device driver is installed. It may be disabled or re-enabled for the entire set of attached drives or medium changers by modifying the registry:

1. Open the reg folder of the driver package.
2. Double-click `DisableDPF.reg` or `EnableDPF.reg`.
3. Reboot the system. This is necessary for any registry modification to take effect.

### ***Querying primary and alternate path configuration***

To check whether the data path failover has been enabled in the device driver and display the primary and alternate paths, you may use the tape diagnostic and utility tool.

**Note:** Display the primary and alternate path configuration for any device using tape diagnostic and utility functions. Refer to *IBM Tape Device Drivers: Installation and User's Guide*, GC27-2130.

### **Checking disablement of data path failover setting**

If you have disabled the data path failover in the device driver's setting by double-clicking the `DisabledDPF.reg` file and rebooting your system, you may go into the registry by issuing the Windows **regedit** command to confirm that DPF has been disabled. Look for a line like the following if your system is Windows Server 2003 or Windows Server 2008:

```
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ibmtp2kx]
"FailoverDisabled"=dword:00000001
```

This indicates that DPF has been disabled in the driver. This setting only takes effect after your system is rebooted.

## **6.2.4 Installing the IBM lin\_tape driver for Red Hat Linux**

In this section, we describe how to install the IBM Linux Tape and Medium Changer Device Driver (`lin_tape`) on a Red Hat Linux platform and connect it to the ProtecTIER system. For tape diagnostic and utility functions, we provide a Tape Utility Program (`IBMtapeutil`). For more details, refer to the following information:

- ▶ *IBM Tape Device Drivers: Installation and User's Guide*, GC27-2130
- ▶ *Implementing IBM Tape in Linux and Windows*, SG24-6268

### **Interoperability**

Check the IBM Interoperability Matrix to ensure that the version of backup server and the operating system that you are running on are supported for ProtecTIER. Also check the HBA and FW level from the host platform to ensure that your end-to-end environment is supported.

**Note:** Not all backup applications require the IBM Tape Device Driver installation. For some vendors, the SCSI pass-through or native OS driver is used. Check the vendor requirements and ISV.



The most current information about supported hardware and software configurations for lin\_tape is available in the README files. The most current information is found in the README files on the IBM Device Drivers FTP website (Figure 6-23), which is located at the following address:

<http://www-933.ibm.com/support/fixcentral/>

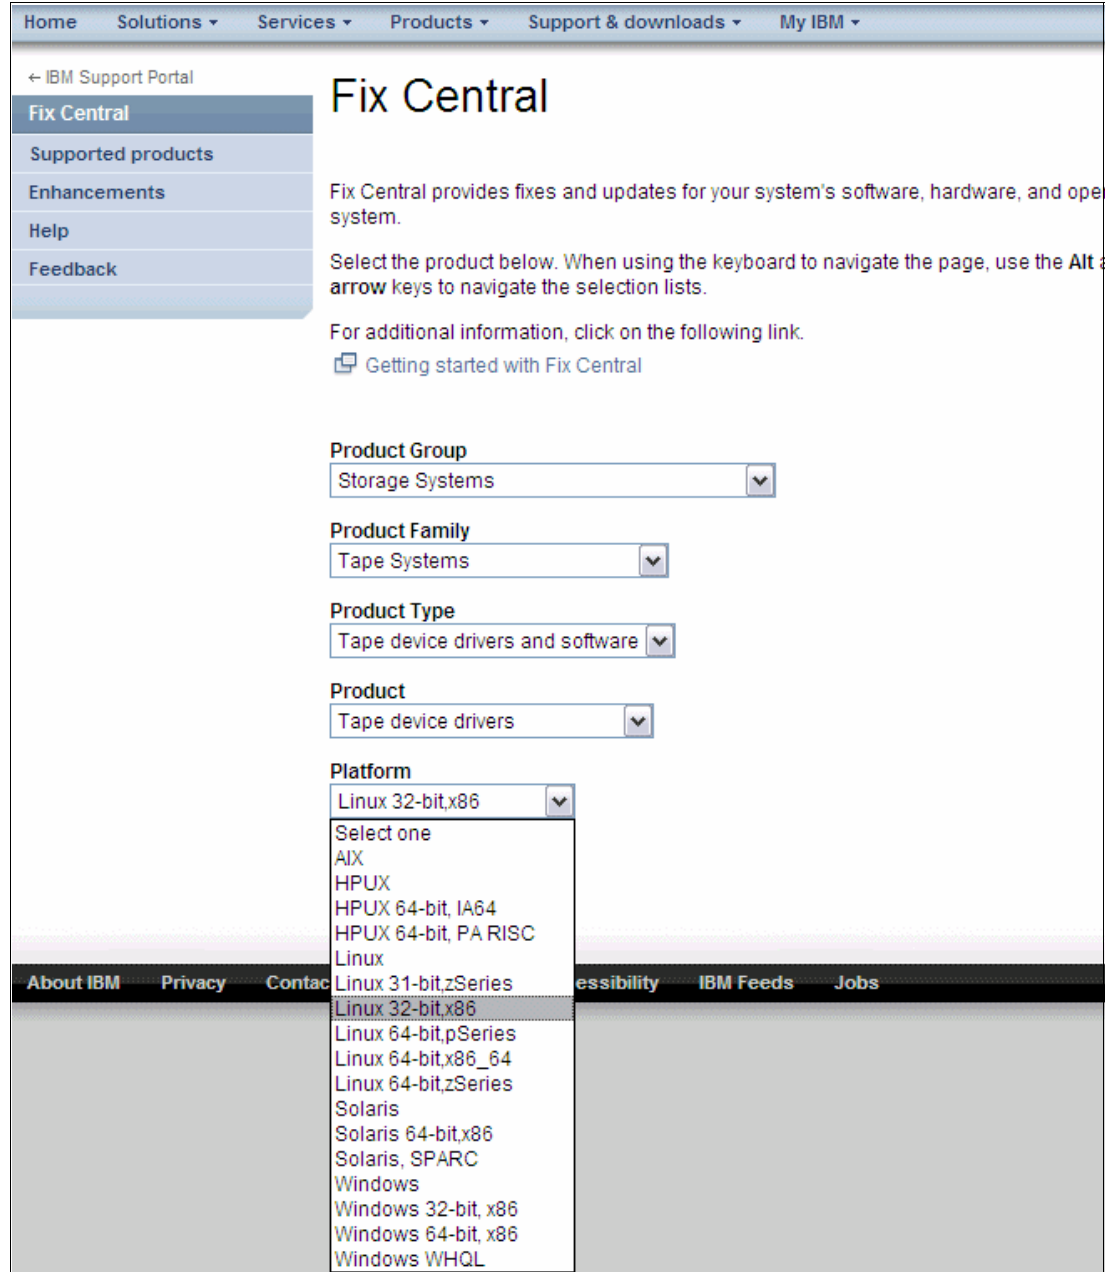


Figure 6-23 Fix Central website

Refer to the SSIC and ISV websites for release information:

- ▶ System Storage Interoperation Center (SSIC) website, found at:

[http://www-03.ibm.com/systems/support/storage/config/ssic/displayessearchwithoutjs.wss?start\\_over=yes](http://www-03.ibm.com/systems/support/storage/config/ssic/displayessearchwithoutjs.wss?start_over=yes)

- ▶ Independent Software Vendors (ISV) website, found at:

<http://www-03.ibm.com/systems/storage/tape/library.html#interoperabilityv>

## The Linux Device Driver (lin\_tape)

The lin\_tape and medium changer device driver is designed specifically to take advantage of the features provided by the IBM tape drives and medium changer devices. The goal is to give applications access to the functions required for basic tape operations (such as backup and restore) and medium changer operations (such as mounting and unmounting the cartridges), as well as to the advanced functions needed by full tape management systems. Whenever possible, the driver is designed to take advantage of the device features transparent to the application.

The software described in this section covers the Linux Device Driver (lin\_tape device driver) and the interface between the application and the tape device.

Figure 6-24 illustrates a typical data flow process.




Figure 6-24 Data flow for Linux Device Driver (lin\_tape)

## Installation and configuration instructions

The lin\_tape device driver for Linux is provided in a source *rpm* package. The utility tools for lin\_tape are supplied in binary rpm packages. They can be downloaded from the following FTP site:

<http://www-933.ibm.com/support/fixcentral/>

The following sections describe installation, configuration, uninstalling, and verification procedures for lin\_tape and its utility tools. Refer to Linux documentation for **tar** command information and any Linux distribution supporting rpm for rpm command information. You must have root authority to proceed with the installation of the driver. See the README file at the following address:

<http://www-933.ibm.com/support/fixcentral/>

This file contains the latest driver information and supersedes the information in this chapter.

The lin\_tape package consists of the device driver and a number of associated files. Also, the IBMtapeutil package consists of IBMtapeutil, IBMtapeconfig, and the source files for IBMtapeutil (from the rpm package). All the components are listed in the *IBM Tape Device Drivers: Installation and User's Guide*, GC27-2130.

### ***Explanation for lin\_tape, lin\_taped, and IBMtapeutil***

The lin-tape, lin-taped, and IBMtapeutil utilities are explained here:

- ▶ **lin\_tape**: the device driver for managing IBM libraries and IBM drives.
- ▶ **lin\_taped**: The **lin\_taped** daemon program can automatically create or delete special files under the `/dev` directory that correspond to the attached or detached tape devices.
- ▶ **IBMtapeutil**: IBMtapeutil is a tape utility program that exercises or tests the functions of the Linux device driver, IBMtape. It performs tape and medium changer operations. These tools are not available separately, but are found in `itdt`.
- ▶ **IBMtapeconfig**: IBMtapeconfig is a script that creates and destroys IBMtape device special files according to the information logged in the `/proc/scsi/IBMtape` and `/proc/scsi/IBMchanger` files.

### **Installation procedure**

Complete the following steps

1. Download the appropriate level of the source RPM package to a directory of your choice on the Linux kernel on which you want to install it.
2. Run **rpmbuild --rebuild <filename>**, where *<filename>* is the name of the RPM file. This creates a binary RPM package for your kernel from the source RPM package. For example:

```
>rpmbuild --rebuild lin_tape-1.x.x.x.0-1.src.rpm
```

3. The output from the build is output to your screen. Near the end of the output, there is a line that indicates the file name and location of your binary RPM package. For example, a line similar to the following is output to your screen:

```
Wrote: /usr/src/redhat/RPMS/i386/lin_tape-1.x.x.x.0-1.i386.rpm
```

4. To install the **lin\_tape** driver from the binary package, run **>rpm -ivh <filename>**. For example:

```
>rpm -ivh /usr/src/redhat/RPMS/i386/lin_tape-1.x.x.x.0-1.i386.rpm
```

5. To install the **lin\_taped** daemon, download it to your Linux file system and run **rpm -ivh** on the daemon RPM file. For example:

```
>rpm -ivh /usr/src/redhat/RPMS/i386/lin_taped-1.x.x.x.0-rhel5.i386.rpm
```

### **Steps to install IBMtapeutil manually**

This procedure is only used when IBMtapeutil does not install automatically when you run **rpm -ivh**. On most platforms, it will install automatically when you run **rpm -ivh lin\_tape**.

Complete the following steps:

1. Download the IBMtapeutil package from:

<http://www-933.ibm.com/support/fixcentral/>

For example, download `IBMtapeutil.1.5.1.rhel4.x86_64.tar.bin`.

2. Install IBMtapeutil by running the following command:

```
> rpm -iv IBMtapeutil.1.5.1.sles9.x86_64.tar.bin
```

Refer to the IBMtapeutil README at the following website for more details:

<http://www-933.ibm.com/support/fixcentral/>

## Update procedure

Before using this procedure to update your device driver to a different level, run the following command to obtain your existing `lin_tape` device driver version if there is at least one tape device attached to your system:

```
> IBMtapeutil -f /dev/IBMtape0 qryversion
```

This command uses tape diagnostic and utility functions. If your current `lin_tape` device driver was installed from a rpm package previously, you may uninstall the driver first, then install the newer version. For example:

```
>rpm -e lin_tape
>rmbuild -rebuild lin_tape.x.x.x.i386.rpm>rpm -i lin_tape.x.x.x.i386.rpm
```

**Note:** All tape devices that use the `lin_tape` device driver must be closed and cannot be in use when `lin_tape` is uninstalled.

## Querying the installed package

The query is supported for the `lin_tape` device driver rpm package only. The installed rpm package can be queried by running the commands given below to display information associated with the package (Example 6-14).

To display information about `lin_tape`, run the following command:

```
>rpm -qi lin_tape
```

### Example 6-14 Query the installed package

---

```
# rpm -qi lin_tape
Name       : lin_tape                Relocations: (not relocatable)
Version    : 1.22.0                 Vendor: IBM
Release    : 1                     Build Date: Mon 19 Jan 2009 01:52:16
PM EST
Install Date: Mon 19 Jan 2009 01:59:35 PM EST
Build Host: xxxxxxx.storage.tucson.ibm.com
Group      : System Environment/Kernel   Source RPM: lin_tape-1.22.0-1.src.rpm
Size       : 1956921                  License: GPL
Signature  : (none)
Packager   : IBM Tape SCSI Device Driver Development
Summary    : IBM Tape SCSI Device Driver for Linux
Description:
The IBM Tape Device Driver, lin_tape, product is a device driver
that provides attachment for the IBM TotalStorage and System
Storage tape devices to Linux compatible platforms.
```

---

To display the file list, run the following command, as shown in Example 6-15:

```
>rpm -ql lin_tape
```

### Example 6-15 Display the file list

---

```
# rpm -ql lin_tape
/etc/init.d/lin_tape
/etc/udev/rules.d/98-lin_tape.rules
/lib/modules/2.6.18-92.el5/kernel/drivers/scsi/lin_tape.ko
/sbin/udev.get_lin_tape_id.sh
/usr/share/doc/lin_tape-1.22.0
/usr/share/doc/lin_tape-1.22.0/COPYING
```

---

```
/usr/share/doc/lin_tape-1.22.0/COPYING.LIB
/usr/share/doc/lin_tape-1.22.0/lin_tape_359X.ReadMe
/usr/share/doc/lin_tape-1.22.0/lin_tape_Ultrium.ReadMe
```

---

To display the states of files in the package (for example, normal, not installed, or replaced), run the following command, as shown in Example 6-16:

```
>rpm -qs lin_tape
```

*Example 6-16 Display the states of files*

---

```
rpm -qs lin_tape
normal      /etc/init.d/lin_tape
normal      /etc/udev/rules.d/98-lin_tape.rules
normal      /lib/modules/2.6.18-92.el5/kernel/drivers/scsi/lin_tape.ko
normal      /sbin/udev.get_lin_tape_id.sh
normal      /usr/share/doc/lin_tape-1.22.0
normal      /usr/share/doc/lin_tape-1.22.0/COPYING
normal      /usr/share/doc/lin_tape-1.22.0/COPYING.LIB
normal      /usr/share/doc/lin_tape-1.22.0/lin_tape_359X.ReadMe
normal      /usr/share/doc/lin_tape-1.22.0/lin_tape_Ultrium.ReadMe
```

---

## Verifying the installation and updating

Run the following command to verify the `lin_tape` device driver version if there is at least one tape device attached to the system (Example 6-17):

```
IBMtapeutil -f /dev/IBMtape0 qryversion
```

*Example 6-17 Verify the lin\_tape device driver version*

---

```
IBMtapeutil -f /dev/IBMtape0 qryversion
Issuing query driver version...
IBMtape driver version: 1.22.0
```

---

**Note:** This command uses the tape diagnostic and utility functions.

## Uninstallation procedure

**Note:** All tape devices that use the `lin_tape` driver must be closed and cannot be in use when `lin_tape` is uninstalled or the uninstallation fails.

Run the following command:

```
>rpm -e lin_tape ---to remove
```

Further information can be obtained from the *IBM Tape Device Drivers: Installation and User's Guide*, GC27-2130.

## Discovering the virtual ProtecTIER devices

The IBMtaped daemon is running and IBMtaped will automatically create the special files under the /dev directory for you. The kernel module automatically creates the special files. In our case, the IBMtaped started automatically. Otherwise, you must run **IBMtapeconfig** to manage the creation of special files for the attached devices, as shown in Example 6-18.

### *Example 6-18 Manual creation of special files*

---

```
#!/IBMtapeconfig
Creating IBMtape special files
major number: 251
Attached devices: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23
24 25 26 27 28 29 30 31
mknod -m 0666 /dev/IBMtape0 c 251 0
mknod -m 0666 /dev/IBMtape0n c 251 1024
mknod -m 0666 /dev/IBMtape1 c 251 1
mknod -m 0666 /dev/IBMtape1n c 251 1025
mknod -m 0666 /dev/IBMtape2 c 251 2
mknod -m 0666 /dev/IBMtape2n c 251 1026
mknod -m 0666 /dev/IBMtape3 c 251 3
mknod -m 0666 /dev/IBMtape3n c 251 1027
mknod -m 0666 /dev/IBMtape4 c 251 4
mknod -m 0666 /dev/IBMtape4n c 251 1028
mknod -m 0666 /dev/IBMtape5 c 251 5
mknod -m 0666 /dev/IBMtape5n c 251 1029
mknod -m 0666 /dev/IBMtape6 c 251 6
mknod -m 0666 /dev/IBMtape6n c 251 1030
mknod -m 0666 /dev/IBMtape7 c 251 7
mknod -m 0666 /dev/IBMtape7n c 251 1031
mknod -m 0666 /dev/IBMtape8 c 251 8
mknod -m 0666 /dev/IBMtape8n c 251 1032
mknod -m 0666 /dev/IBMtape9 c 251 9
mknod -m 0666 /dev/IBMtape9n c 251 1033
mknod -m 0666 /dev/IBMtape10 c 251 10
mknod -m 0666 /dev/IBMtape10n c 251 1034
mknod -m 0666 /dev/IBMtape11 c 251 11
mknod -m 0666 /dev/IBMtape11n c 251 1035
mknod -m 0666 /dev/IBMtape12 c 251 12
mknod -m 0666 /dev/IBMtape12n c 251 1036
mknod -m 0666 /dev/IBMtape13 c 251 13
mknod -m 0666 /dev/IBMtape13n c 251 1037
mknod -m 0666 /dev/IBMtape14 c 251 14
mknod -m 0666 /dev/IBMtape14n c 251 1038
mknod -m 0666 /dev/IBMtape15 c 251 15
mknod -m 0666 /dev/IBMtape15n c 251 1039
mknod -m 0666 /dev/IBMtape16 c 251 16
mknod -m 0666 /dev/IBMtape16n c 251 1040
mknod -m 0666 /dev/IBMtape17 c 251 17
mknod -m 0666 /dev/IBMtape17n c 251 1041
mknod -m 0666 /dev/IBMtape18 c 251 18
mknod -m 0666 /dev/IBMtape18n c 251 1042
mknod -m 0666 /dev/IBMtape19 c 251 19
mknod -m 0666 /dev/IBMtape19n c 251 1043
mknod -m 0666 /dev/IBMtape20 c 251 20
mknod -m 0666 /dev/IBMtape20n c 251 1044
mknod -m 0666 /dev/IBMtape21 c 251 21
```

```
mknod -m 0666 /dev/IBMtape21n c 251 1045
mknod -m 0666 /dev/IBMtape22 c 251 22
mknod -m 0666 /dev/IBMtape22n c 251 1046
mknod -m 0666 /dev/IBMtape23 c 251 23
mknod -m 0666 /dev/IBMtape23n c 251 1047
mknod -m 0666 /dev/IBMtape24 c 251 24
mknod -m 0666 /dev/IBMtape24n c 251 1048
mknod -m 0666 /dev/IBMtape25 c 251 25
mknod -m 0666 /dev/IBMtape25n c 251 1049
mknod -m 0666 /dev/IBMtape26 c 251 26
mknod -m 0666 /dev/IBMtape26n c 251 1050
mknod -m 0666 /dev/IBMtape27 c 251 27
mknod -m 0666 /dev/IBMtape27n c 251 1051
mknod -m 0666 /dev/IBMtape28 c 251 28
mknod -m 0666 /dev/IBMtape28n c 251 1052
mknod -m 0666 /dev/IBMtape29 c 251 29
mknod -m 0666 /dev/IBMtape29n c 251 1053
mknod -m 0666 /dev/IBMtape30 c 251 30
mknod -m 0666 /dev/IBMtape30n c 251 1054
mknod -m 0666 /dev/IBMtape31 c 251 31
mknod -m 0666 /dev/IBMtape31n c 251 1055
```

Creating IBMchanger special files

```
major number: 251
```

```
Attached devices: 0 1
```

```
mknod -m 0666 /dev/IBMchanger0 c 251 2048
```

```
mknod -m 0666 /dev/IBMchanger1 c 251 2049
```

---

Figure 6-25 shows a part of our created virtual library named Decanter\_1. It has one medium changer (Robot) and 16 drives configured for both nodes in that cluster (2 x 16 drives, 2 x changers).

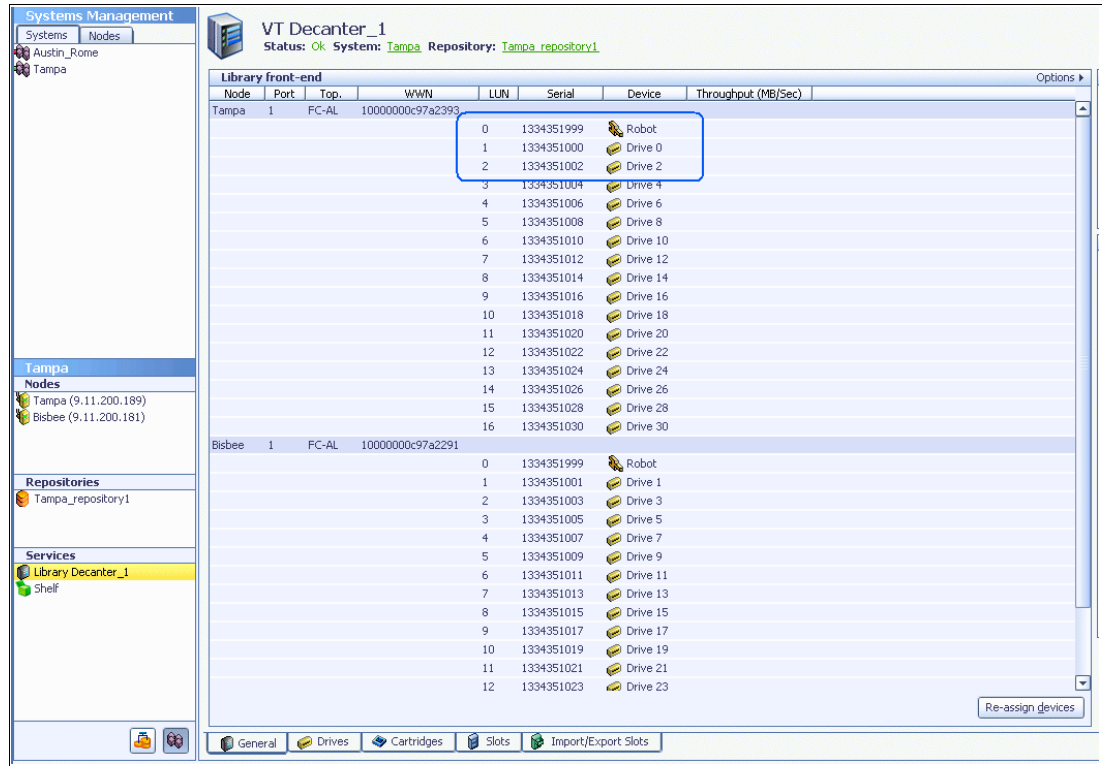


Figure 6-25 Changer and drives on node1

The following commands allow you to view the virtual ProtecTIER devices (tape drives and changer) created by the IBMtape device driver based on the ProtecTIER configuration:

- ▶ **cat /proc/scsi/IBMtape**
- ▶ **cat /proc/scsi/IBMchanger**
- ▶ **cat /proc/scsi/IBM\***

The output that is created by executing the previous commands show the ProtecTIER devices by model and by SN (Example 6-19).

*Example 6-19 Viewing devices created by IBMtaped*

```
# cat /proc/scsi/IBMtape
lin_tape version: 1.22.0
lin_tape major number: 251
Attached Tape Devices:
Number  model          SN              HBA              FO Path
0       ULT3580-TD3   1334351001     lpfc             NA
1       ULT3580-TD3   1334351003     lpfc             NA
2       ULT3580-TD3   1334351005     lpfc             NA
3       ULT3580-TD3   1334351007     lpfc             NA
4       ULT3580-TD3   1334351009     lpfc             NA
5       ULT3580-TD3   1334351011     lpfc             NA
6       ULT3580-TD3   1334351013     lpfc             NA
7       ULT3580-TD3   1334351015     lpfc             NA
8       ULT3580-TD3   1334351017     lpfc             NA
```



```

9      ULT3580-TD3 1334351019      lpfc      NA
10     ULT3580-TD3 1334351021      lpfc      NA
11     ULT3580-TD3 1334351023      lpfc      NA
12     ULT3580-TD3 1334351025      lpfc      NA
13     ULT3580-TD3 1334351027      lpfc      NA
14     ULT3580-TD3 1334351029      lpfc      NA
15     ULT3580-TD3 1334351031      lpfc      NA
16     ULT3580-TD3 1334351000      lpfc      NA
17     ULT3580-TD3 1334351002      lpfc      NA
18     ULT3580-TD3 1334351004      lpfc      NA
19     ULT3580-TD3 1334351006      lpfc      NA
20     ULT3580-TD3 1334351008      lpfc      NA
21     ULT3580-TD3 1334351010      lpfc      NA
22     ULT3580-TD3 1334351012      lpfc      NA
23     ULT3580-TD3 1334351014      lpfc      NA
24     ULT3580-TD3 1334351016      lpfc      NA
25     ULT3580-TD3 1334351018      lpfc      NA
26     ULT3580-TD3 1334351020      lpfc      NA
27     ULT3580-TD3 1334351022      lpfc      NA
28     ULT3580-TD3 1334351024      lpfc      NA
29     ULT3580-TD3 1334351026      lpfc      NA
30     ULT3580-TD3 1334351028      lpfc      NA
31     ULT3580-TD3 1334351030      lpfc      NA

```

```
# cat /proc/scsi/IBMchanger
```

```
lin_tape version: 1.22.0
```

```
lin_tape major number: 251
```

```
Attached Tape Devices:
```

Number	model	SN	HBA	FO Path
0	03584L32	0013343519990402	lpfc	NA
1	03584L32	0013343519990402	lpfc	NA

## Control path failover support for tape libraries

The Linux `lin_tape` device driver control path failover support configures multiple physical control paths to the same logical library within the device driver and provides automatic failover to an alternate control path when a permanent error occurs on one path. This action is transparent to the running application.

## Configuring and unconfiguring path failover support

Control path failover support (Example 6-20) is not enabled automatically when the device driver is installed. The Linux `lin_tape` device driver provides a driver parameter `alternate_pathing` for you to enable the library control path failover.

*Example 6-20 /etc/modprobe.conf with enabled control path failover*

```

# cat /etc/modprobe.conf
alias eth0 e1000e
alias eth1 e1000e
alias eth2 bnx2
alias eth3 bnx2
alias scsi_hostadapter megaraid_sas
alias scsi_hostadapter1 ata_piix
alias scsi_hostadapter2 qla2xxx
alias scsi_hostadapter3 lpfc
alias scsi_hostadapter4 usb-storage

```

```
options scsi_mod max_luns=256
alias scsi_hostadapter5 lpfc
alias scsi_hostadapter6 lpfc
alias scsi_hostadapter7 lpfc
options lin_tape alternate_pathing=1
```

---

To enable the failover support in the `lin_tape` device driver software, complete the following steps after installing the `lin_tape` rpm package:

1. Run **`lin_taped stop`**.  
Stops the `lin_taped` daemon.
2. Run **`rmmod lin_tape`**.  
Unloads the `lin_tape` driver from the memory.
3. Add the following line to your `/etc/modules.conf` file for 2.4 kernels or to the `/etc/modprobe.conf.local` file for 2.6 kernels:  

```
options lin_tape alternate_pathing=1
```
4. Run **`depmod`**.
5. Run **`modprobe lin_tape`**.  
Reloads the `lin_tape` driver into memory.
6. Run **`lin_taped`**.  
Restarts the `lin_taped` daemon.

The kernel can be checked by running **`uname -a`** (Example 6-21).

*Example 6-21 Kernel check*

---

```
uname -a
Linux atlantic.storage.tucson.ibm.com 2.6.18-92.el5 #1 SMP Tue Apr 29 13:16:15 EDT
2008 x86_64 x86_64 x86_64 GNU/Linux
```

---

The Red Hat version can be checked by running **`cat /etc/redhat-release`** (Example 6-22).

*Example 6-22 Version check*

---

```
cat /etc/redhat-release
Red Hat Enterprise Linux Server release 5.2 (Tikanga)
```

---

You can check whether the `lin_tape` driver has recognized multiple control paths for your library by reading the `/proc/scsi/IBMchanger` file:

```
cat /proc/scsi/IBMchanger
```

If your library lists Primary or Alternate under FO Path, you have successfully enabled the control path failover feature for your library. If NA is listed under FO Path, then the control path failover is not enabled.

After control path failover support is enabled, it remains set until the `lin_tape` driver is reloaded with the `alternate_pathing` driver parameter set to OFF. The path failover setting is retained even if the system is rebooted. If you want to turn off the control path failover feature in the `lin_tape` device driver, complete the following steps:

1. Run **`lin_taped stop`**.
2. Run **`rmmod lin_tape`**.

3. Delete the following line in your `/etc/modules.conf` file:

```
options lin_tape alternate_pathing=1
```

4. Run `depmod`.

5. Run `modprobe lin_tape`.

6. Run `lin_taped`.

## Primary and alternate paths

When `lin_tape` is loaded into kernel memory, the first logical medium changer device that `lin_tape` sees in the system is the primary path for that medium changer. The other logical medium changers that `lin_tape` attached for the same medium changer are configured as alternate paths. The device driver supports up to 16 physical paths for a single device. The primary and alternate path information can be obtained by running the following command:

```
cat /proc/scsi/IBMchanger
```

Example 6-23 is an example of a `/proc/scsi/IBMchanger` file.

*Example 6-23* `cat /proc/scsi/IBMchanger`

---

```
cat /proc/scsi/IBMchanger
lin_tape version: 1.22.0
lin_tape major number: 251
Attached Tape Devices:
Number  model      SN                HBA                FO Path
0       03584L32   0013343519990402 lpfc                Primary
1       03584L32   0013343519990402 lpfc                Alternate
```

---

The labeling of a logical device as either a primary or an alternate path is for information only, in order to:

- ▶ Be able to identify the actual number of physical devices configured on the system and a specific logical device associated with them. There is only one logical device labeled as the primary path for each physical device. However, there can be multiple logical devices labeled as an alternate path for the same devices.
- ▶ Provide information about which logical devices are configured on the system that have path failover support enabled.

## Querying primary and alternate path configuration

You can display the primary and alternate path configuration for all devices by reading the `/proc/scsi/IBMchanger` file, as explained in “Primary and alternate paths” on page 325.

## Disabling and enabling primary and alternate paths

After you load the `lin_tape` device driver with the `alternate_pathing` parameter set to ON, by default, all the available paths for a physical device are enabled. If it is necessary to disable a path and not perform path failover (for example, due to maintenance), run commands to disable and then later enable the primary and alternate paths. The commands to enable and disable primary and alternate paths are tape diagnostic and utility functions.

**Note:** See *IBM Tape Device Drivers: Installation and User's Guide*, GC27-2130, for information about the IBM Tape Diagnostic Tool (ITDT) and the System - Tape Utility Program (IBMtapeutil).

## Data path failover and load balancing support for tape drives

Data path failover support is not enabled automatically when the device driver is installed. The Linux `lin_tape` device driver provides a driver parameter `alternate_pathing` for you to enable the data path failover. To enable the failover support in the `lin_tape` device driver software, you must complete the following steps after installing the `lin_tape rpm` package:

1. Run `lin_taped stop`.

Stops the `lin_taped` daemon.

2. Run `rmmod lin_tape`.

Unloads the `lin_tape` driver from the memory.

For IBM LTO tape drives, the library must have the path failover feature code. The data path failover license keys are needed to enable the failover if you are running LTO2 drives or if you are running LTO3 drives with old levels of drive code. DPF keys do not need to be added if you are running the latest drive code on LTO3 or LTO4 drives. Add the following line in your `/etc/modules.conf` file for 2.4 kernels or the `/etc/modprobe.conf.local` file for 2.6 kernels:

```
options lin_tape alternate_pathing=1 dpf_keys="abcdefghijklmnop"
```

“`abckdefghijklmnop`” is an example of a data path failover feature key. If you have multiple libraries and multiple data path failover feature keys, input your keys as follows:

```
dpf_keys="key1;key2;..."
```

Save the file, then run the following commands:

1. `depmod`

2. `modprobe lin_tape`

Reloads the `lin_tape` driver into memory.

3. `lin_taped`

Restarts the `lin_taped` daemon.

You can check whether the `lin_tape` driver has recognized multiple paths for your tape drive by reading the `/proc/scsi/IBMtape` file:

```
>cat /proc/scsi/IBMtape
```

If your tape drive lists Primary or Alternate under FO Path, you have successfully enabled the data path failover feature for your tape drive. If NA is listed under FO Path, then the data path failover is not enabled. After the path fail over support is enabled, it remains set until the `lin_tape` driver is reloaded with the *alternate pathing* driver parameter set to OFF. The path failover setting is retained even if the system is rebooted. If you want to turn off the data path failover feature in the `lin_tape` device driver, complete the following steps:

1. Run `lin_taped stop`.

2. Run `rmmod lin_tape`.

3. Delete the following line in your `/etc/modules.conf` file:

```
options lin_tape alternate_pathing=1
```

4. Run `depmod`.

5. Run `modprobe lin_tape`.

6. Run `lin_taped`.

## Primary and alternate paths

When the `lin_tape` device driver is loaded into kernel memory with path failover support enabled, the first logic device that `lin_tape` sees always becomes the primary path. The other logical devices that `lin_tape` sees are configured as the alternate paths. The device driver supports up to 16 physical paths for a single device. The primary and alternate path information can be obtained by running the following command:

```
>cat /proc/scsi/IBMtape
```

Example 6-24 is an example of the output of `/proc/scsi/IBMtape`.

*Example 6-24 Example of /proc/scsi/IBMtape*

---

```
cat /proc/scsi/IBMtape
lin_tape version: 1.22.0
lin_tape major number: 251
Attached Tape Devices:
Number  model      SN          HBA          FO Path
0       ULT3580-TD3 1691612001  lpf         Primary
1       ULT3580-TD3 1691612000  lpf         Primary
```

---

The labeling of a logical device as either a primary or an alternate path is for informational purposes only, to:

- ▶ Be able to identify the actual number of physical devices configured on the system and a specific logical device associated with them. There is only one logical device that is labeled as the primary path for each physical device. However, there might be many (multiple) logical devices labeled as an alternate path for the same devices.
- ▶ Provide information about which logical devices configured on the system have path failover support enabled.

## Querying primary and alternate path configuration

You can display the primary and alternate path configuration for all devices by reading the `/proc/scsi/IBMtape` file, as explained in “Primary and alternate paths” on page 325.

**Note:** You can display the primary and alternate path configuration for any device using tape diagnostic and utility functions that come with the operation system or device driver installation. Refer to the IBM Tape Diagnostic Tool (ITDT) or the Linux System - Tape Utility Program (IBMtapeutil) at the following address:

<http://www-933.ibm.com/support/fixcentral/>

## Disabling and enabling primary and alternate paths

If it is necessary to disable a path and not perform path failover (for example, due to maintenance), run commands to disable and then later enable the primary and alternate paths. The commands to enable and disable primary and alternate paths are tape diagnostic and utility functions.

## 6.2.5 Open source device driver: `lin_tape`

The `lin_tape` device driver is the new device driver for the Linux 2.6 kernels to replace the `IBMtape` driver. In most respects, it behaves the same as the `IBMtape` device driver. This section covers significant differences between the `IBMtape` driver and the `lin_tape` driver.

## IBMtape and lin\_tape comparison

The table shown in Figure 6-26 compares the names for various components of the IBMtape and the lin\_tape device drivers.

Component	IBMtape	Lin_tape
Driver name	IBMtape	lin_tape
Module name	IBMtape.ko	lin_tape.ko
Special files	/dev/IBMtape0 /dev/IBMchanger0, etc.	No change
proc entry	/proc/scsi/IBMtape /proc/scsi/IBMchanger	No change
Daemon name	IBMtaped	lin_taped
Daemon configuration file	/etc/IBMtaped.conf	/etc/lin_taped.conf
Daemon trace files	/var/log/IBMtape.trace /var/log/IBMtape.errorlog	/var/log/lin_tape.trace /var/log/lin_tape.errorlog

Figure 6-26 Comparison of IBMtape and lin\_tape names

The IBMtapeutil utility is the same for both the IBMtape driver and the lin\_tape driver.

**Note:** This task is performed using tape diagnostic and utility functions.

## Installation

Installation of the lin\_tape driver is the same as for the IBMtape driver, except that IBMtape should be replaced with lin\_tape in all of the installation instructions. Refer to the “Installation and configuration instructions” on page 316, for details. The lin\_tape driver cannot be installed if the IBMtape driver is already installed. If the IBMtape driver is installed, first uninstall the IBMtape driver, then install the lin\_tape driver. With RHEL4 and SLES10, driver removal also requires a reboot of the server, because the IBMtape driver module is permanent in these distributions.

## Driver parameters and special device files

The driver parameters have not changed for the lin\_tape driver. However, it is important to note that the module parameters, such as alternate pathing and dpf\_keys, must now be applied to the lin\_tape module instead of the IBMtape module. For example, in the /etc/modprobe.conf (for Red Hat) or /etc/modprobe.conf.local (for SUSE) file, add the following line for the LTO library’s path failover:

```
options lin_tape alternate_pathing=1 dpf_keys="abcdefghijklmnop"
```

“abckdefghijklmnop” is an example of a data path failover feature key.

The special device files for the lin\_tape driver are the same as for the IBMtape driver. Refer to “Special Files for the Tape Device” and “Special Files for the Medium Changer Device” sections in the *IBM Tape Device Drives Installation and User’s Guide*, GC27-21230.

## Path failover support

Path failover support in `lin_tape` is the same, except that with the `lin_tape` driver, failover support is provided through the `lin_taped` daemon. If the `lin_taped` daemon is not running, control path failover or data path failover is not attempted. The `lin_taped` daemon is started automatically when the `lin_tape` driver is loaded. To check whether the `lin_taped` daemon is running, run the following command:

```
lin_taped status
```

Example 6-25 shows the output of this command.

*Example 6-25 lin\_taped status*

---

```
# lin_taped status
lin_taped is running, pid 15476.
```

---

This command indicates whether the `lin_taped` daemon is running. If the `/proc/scsi/IBMtape` and `/proc/scsi/IBMchanger` files indicate not applicable (NA) for Failover (FO) Path, this indicates that failover support for that device is not enabled. If all other settings are correct, but FO Path is incorrectly indicating NA, confirm that the `lin_taped` daemon is running.

For details about the path failover support, refer to “Control path failover support for tape libraries” on page 303 and “Data path failover and load balancing support for tape drives” on page 305.

## lin\_taped daemon

The `lin_taped` daemon uses the same command-line arguments as the `IBMtaped` daemon. The `lin_taped` configuration file is the same as the `IBMtaped` configuration file, but has been renamed to `lin_taped.conf`. Refer to *IBM Tape Device Drives: Installation and User's Guide*, GC27-2130 for detailed information.

## How to get the WWPN and WWNN of a HBA

Run `cd` to change the directory to this location:

```
cd /sys/class/scsi_host
```

Run `ls` to list the host adapter numbers. In our case, we had two Dual Port Emulex HBAs and two Dual QLogic HBA (total of eight ports), as shown in Example 6-26.

*Example 6-26 List host adapters*

---

```
# ls -l
[00mtotal 0
drwxr-xr-x 2 root root 0 Sep 27 10:42 [00;34mhost0[00m
drwxr-xr-x 2 root root 0 Sep 26 23:26 [00;34mhost1[00m
drwxr-xr-x 2 root root 0 Sep 26 23:26 [00;34mhost2[00m
drwxr-xr-x 2 root root 0 Sep 26 23:26 [00;34mhost3[00m
drwxr-xr-x 2 root root 0 Sep 26 23:27 [00;34mhost4[00m
drwxr-xr-x 2 root root 0 Sep 27 10:46 [00;34mhost5[00m
drwxr-xr-x 2 root root 0 Sep 26 23:27 [00;34mhost6[00m
drwxr-xr-x 2 root root 0 Sep 26 23:28 [00;34mhost7[00m
drwxr-xr-x 2 root root 0 Sep 26 23:28 [00;34mhost8[00m
```

---

Run `cd` to enter each directory and run `ls -l` to list the contents. Any entries that list `lpfc` are Emulex HBAs. In our case, these are `host5`–`host8`. `lpfc` is the Emulex Driver on Red Hat. `lpfc` is also the abbreviation for Light Pulse Fibre Channel.

For host5 (Emulex HBA), run `cd /sys/class/fc_host/host5/` and run:

- ▶ `cat port_name` for the WWPN
- ▶ `cat node_name` for the WWNN

Example 6-27 shows the output of these commands.

*Example 6-27 Checking the WWNN and WWPN*

---

```
# cat node_name
0x20000000c979617a
# cat port_name
0x10000000c979617a
```

---

Host1 - host4 are used for the QLogic HBAs in our case, and for them, you can run the following command:

```
cd /sys/class/scsi_host/hostx/device/fc_host:hostx (replace the x with the host number)
```

Then you can check the WWNN and WWPN in the same way as for the Emulex HBA by running the following commands (See Example 6-28 for the output of these commands):

- ▶ `cat port_name` for the WWPN
- ▶ `cat node_name` for the WWNN

*Example 6-28 Checking the WWNN and WWPN*

---

```
cat node_name
0x2000001b3211a4d9
# cat node_nameport_name
0x2100001b3211a4d9
```

---

## 6.3 Setting up ProtecTIER on IBM i

In this section, we describe setting up ProtecTIER on IBM i.

### 6.3.1 Prerequisites and test environment

For IBM i, you must have a ProtecTIER minimum code level of 2.2.3.0 or later.

TS7650G and TS7650 systems can be connected through IBM i Fibre Channel adapters. See the System Storage Interoperation Center (SSIC) for supported adapters at the following address:

[http://www-03.ibm.com/systems/support/storage/config/ssic/displaysssearchwithoutjs.wss?start\\_over=yes](http://www-03.ibm.com/systems/support/storage/config/ssic/displaysssearchwithoutjs.wss?start_over=yes)

The following IBM i operating system levels support connection of ProtecTIER:

- ▶ V5R4
- ▶ V5R4M5
- ▶ V6R1



We use the following host config for our test environment:

- ▶ IBM i partition in POWER5™
- ▶ IBM i V6R1
- ▶ Two 4 Gb IOP-based adapters for tape connection, feature 5761 and Custom Card Identification Number (CCIN) 280D
- ▶ Backup software Backup, Recovery, and Media Services

### 6.3.2 Interoperability

Check the IBM Interoperability Matrix to ensure that the version of the backup server and the operating system that you are running on are supported for ProtecTIER.

Refer to the SSIC and ISV websites for release information:

- ▶ SSIC website:

[http://www-03.ibm.com/systems/support/storage/config/ssic/displayessearchwithoutjs.wss?start\\_over=yes](http://www-03.ibm.com/systems/support/storage/config/ssic/displayessearchwithoutjs.wss?start_over=yes)

- ▶ ISV website:

<http://www-03.ibm.com/systems/storage/tape/library.html#interoperability>

### 6.3.3 Connecting the ProtecTIER FC ports to IBM i Fibre Channel adapters

The following Fibre Channel protocols are supported by a front-end port in ProtecTIER:

- ▶ Point to point (direct connection)
- ▶ Arbitrated loop

When connecting the ProtecTIER through SAN switches, the arbitrated loop protocol is translated to Switched Fabric protocol in the switch. In our test, we connect the ProtecTIER front-end ports to IBM i I/O adapters (IOAs) directly. The Fibre Channel port properties can be changed in ProtecTIER Manager by selecting **Nodes** → **Port Attributes**, as shown in Figure 6-28 on page 333. The topology point-to-point should work in most cases and changes automatically if a switch is connected. IBM i supports one path to a tape device, whether a real tape device or virtual tape drive, as shown in Figure 6-27.

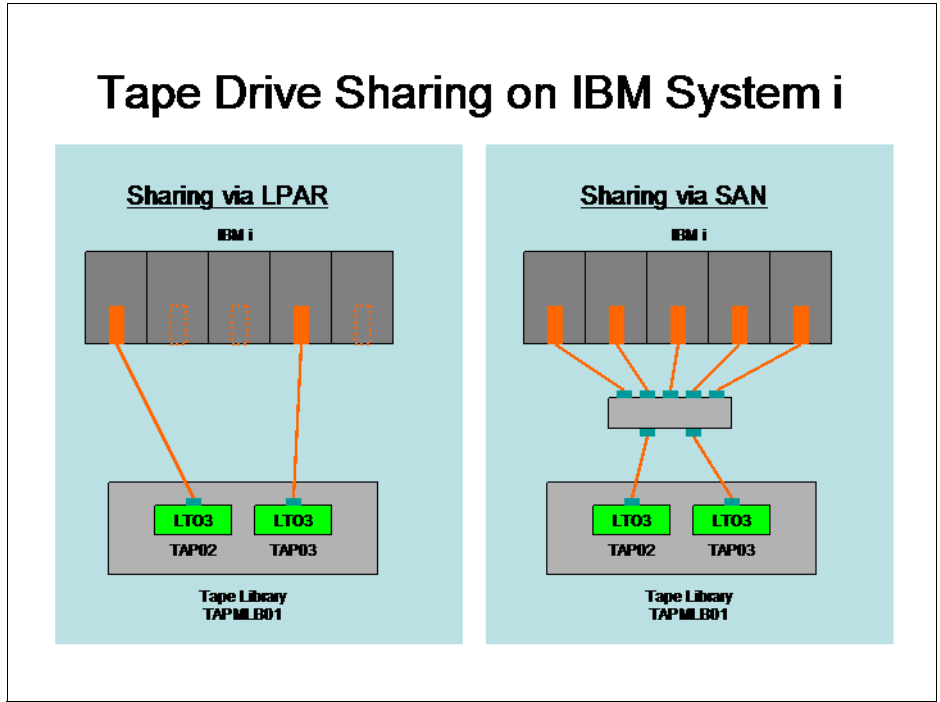


Figure 6-27 Tape drive sharing on IBM i

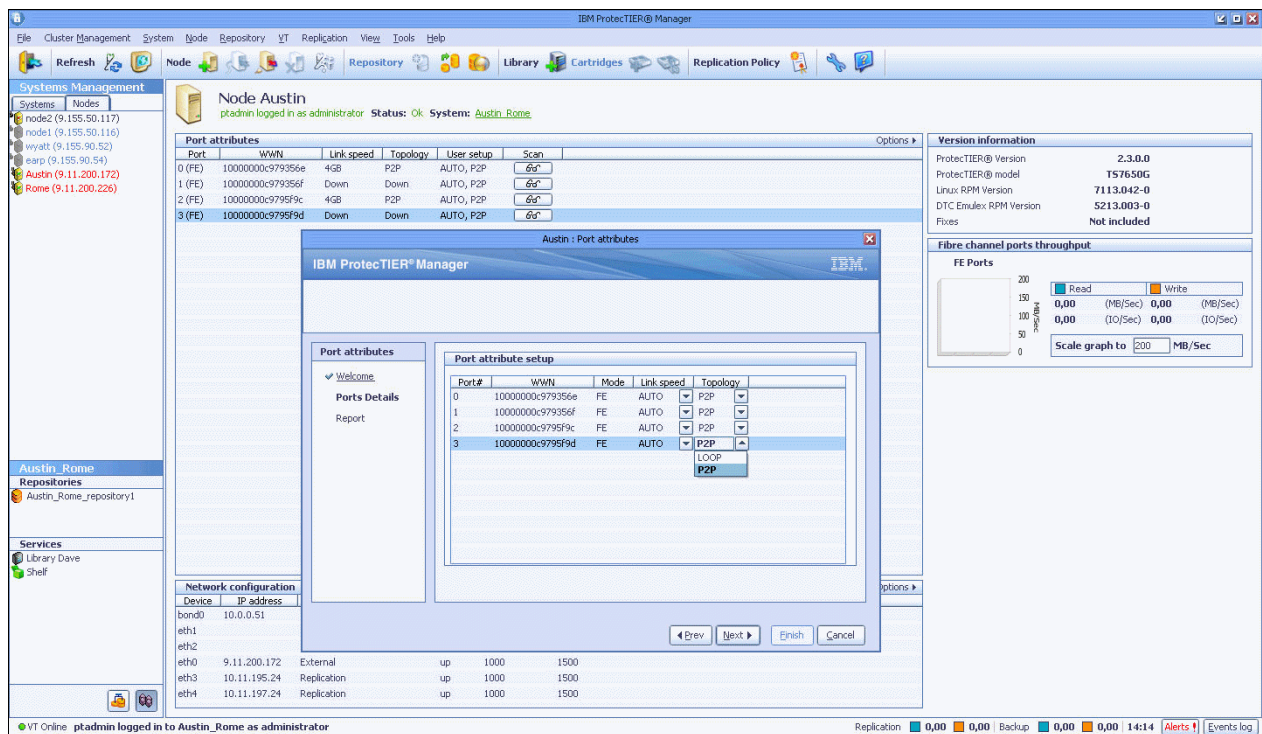


Figure 6-28 Port Attribute menu

Configurations in IBM i must have only one path from an IBM i partition to any tape device, because there currently is no multipath support. Therefore, the switches should be zoned so that only one IOA per IBM i partition sees a port in a ProtecTIER system (we do not allow multiple IOAs to one front-end port), but multiple ports in a ProtecTIER system can see the same IOA.

**Note:** It is acceptable for several different IBM i partitions to all see the same ProtecTIER Fibre port, just not two tape adapters in a single partition.

Even if a tape library with control or data path failover capability is connected to an IBM i partition, the same tape drive cannot be connected through multiple IBM i Fibre Channel adapters to achieve failover and load balancing. The reason for this is that IBM i licensed internal code (LIC) contains the driver functions for supported adapters, and LIC does not support more than one path to a tape drive, so a control path failover or data path failover are not possible.

## 6.3.4 Creating a VTL for IBM i

We create a Virtual Tape Library for IBM i by completing the following steps:

1. Log in to the ProtecTIER GUI. On the toolbar, click **Library** to start the Create New Library wizard, as shown in Figure 6-29.

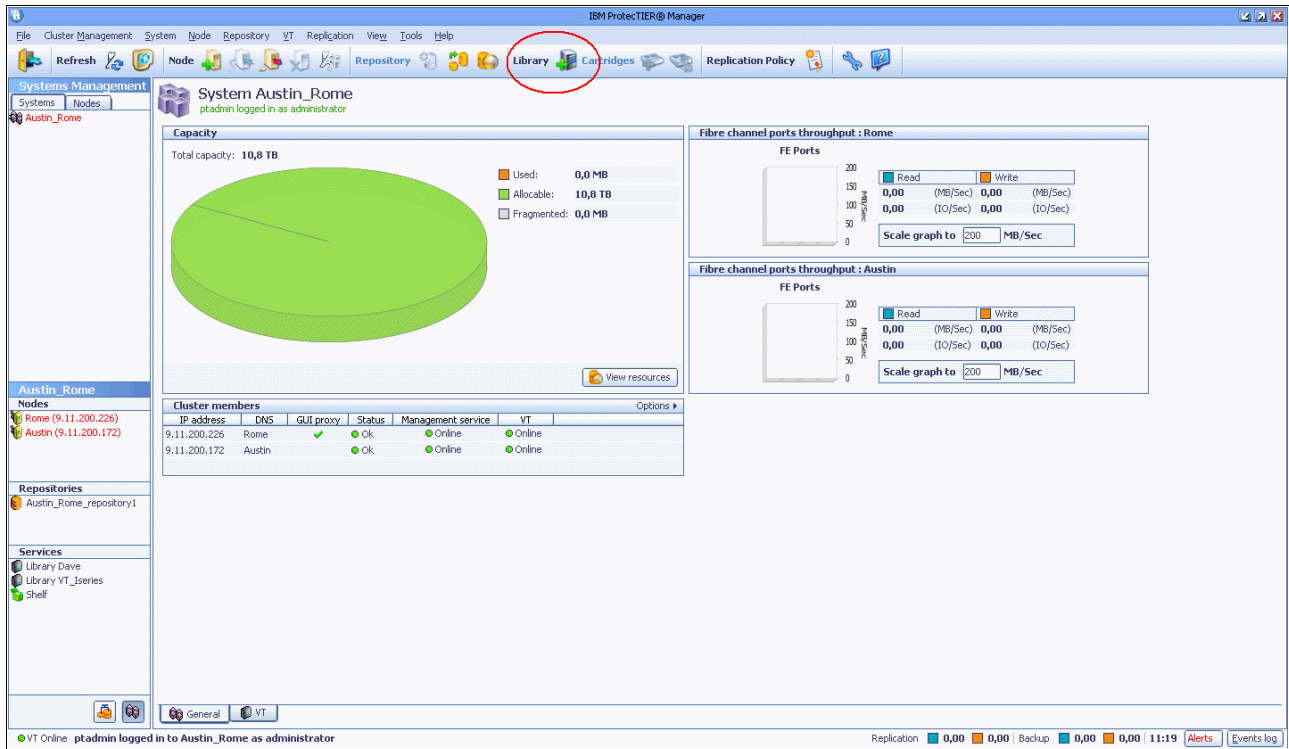


Figure 6-29 Add/create library

2. Click **Next** on the Welcome window. On the Library name window, insert the name of the new VTL (Figure 6-30).

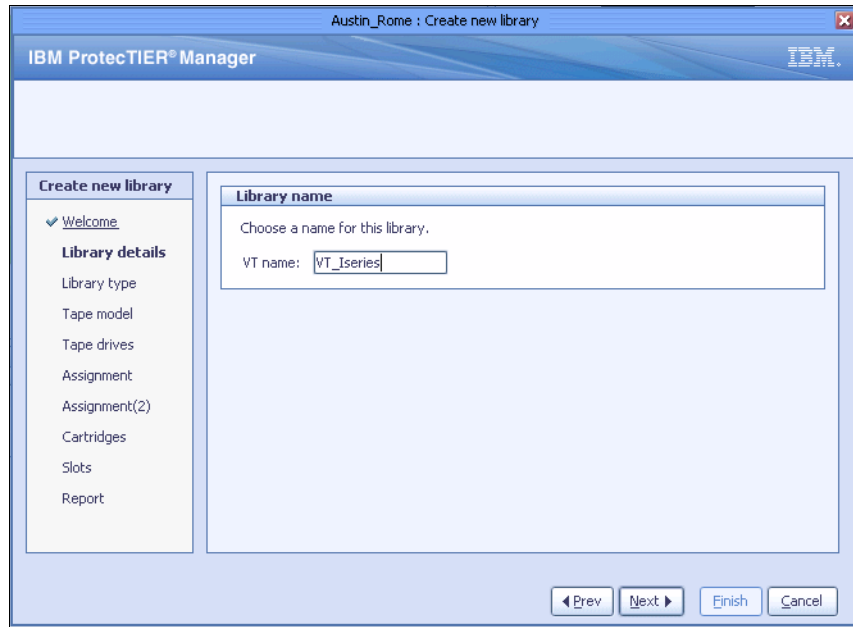


Figure 6-30 Library details menu

3. On next window, specify the type of virtual tape library being created. We specify TS3500, as shown in Figure 6-31.

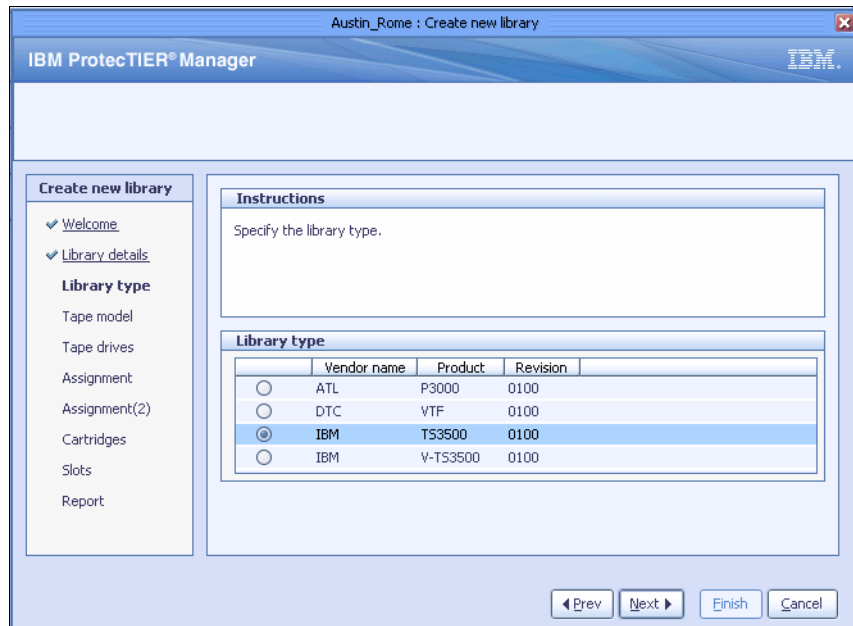


Figure 6-31 Library type menu

- After confirming the default tape drive model, specify the number of virtual tape drives within the VTL on each node of the ProtecTIER two-node cluster. We define one tape drive on node1 and one on node2 (Figure 6-32).

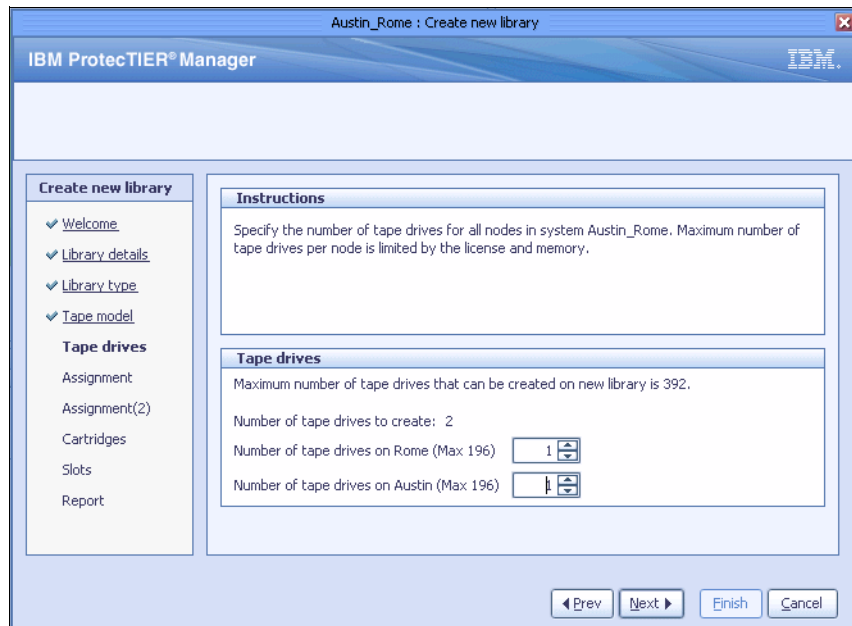


Figure 6-32 Tape drives menu

- We have the ability to assign the created tape drives to ports in each node and to define or change control paths. Because we defined only one drive per node, ProtecTIER Manager, by default, assigns it to the first port on the node. The checked Robot box means that the control path will be enabled for that drive. Clear the other ports, as shown Figure 6-33.

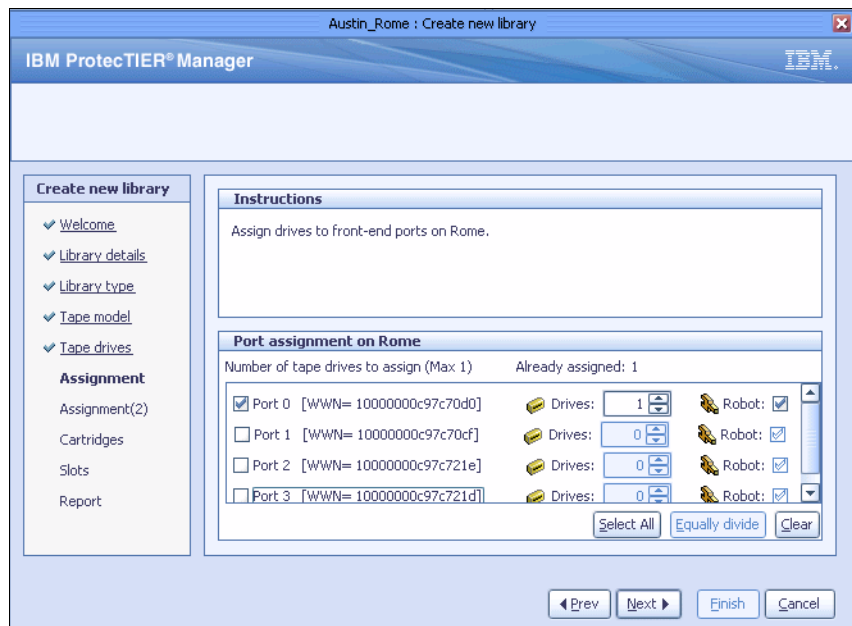


Figure 6-33 Port Assignment menu

6. By assigning the drive to a ProtecTIER port connected to the IBM i adapter, IBM i sees that the drive and the tape library providing the system value autoconfig is set to yes. We now define the number of virtual cartridges and their volume serial numbers or barcodes. We define 16 cartridges with volsers I00000–I00015, as shown in Figure 6-34.

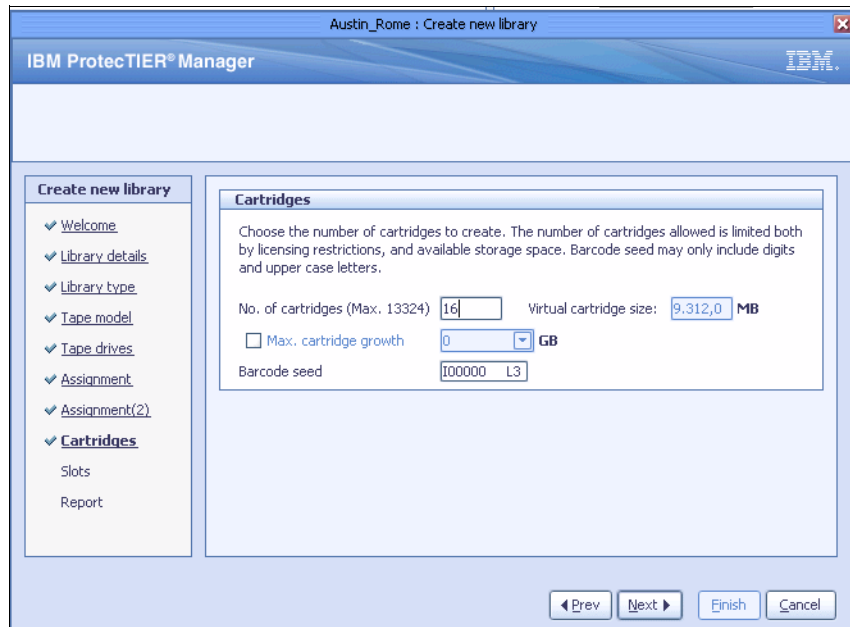


Figure 6-34 Cartridge define/create menu

7. We specify an equal number of slots and cartridges. We do not need any import or export slots, so we leave those numbers as 0 (Figure 6-35).

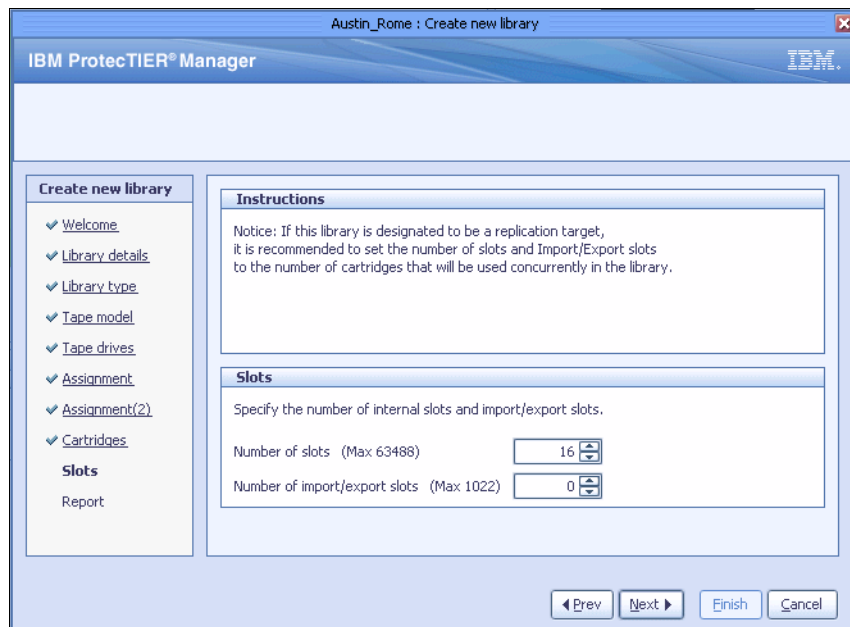


Figure 6-35 Slot configuration menu

8. Before the VTL is actually created, check its summary. If you click the **Previous** button, you can still change any feature. After you agree with the defined library, click **Finish**. The summary of our VTL is shown in Figure 6-36.

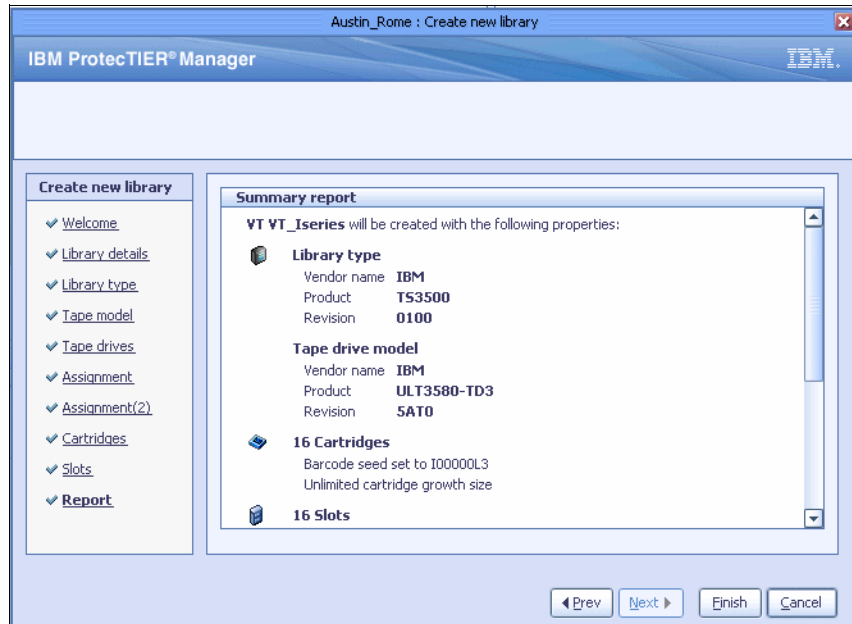


Figure 6-36 Summary report

9. After confirming the creation of the new VTL, ProtecTIER Manager informs you of the library being created in the background, as shown in Figure 6-37.

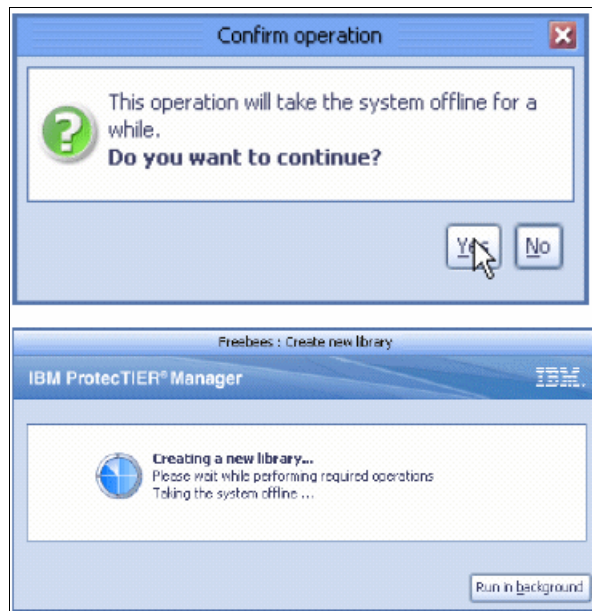


Figure 6-37 Creating a new library message



The new VTL can be seen in the ProtecTIER GUI. Our example is shown in Figure 6-38.

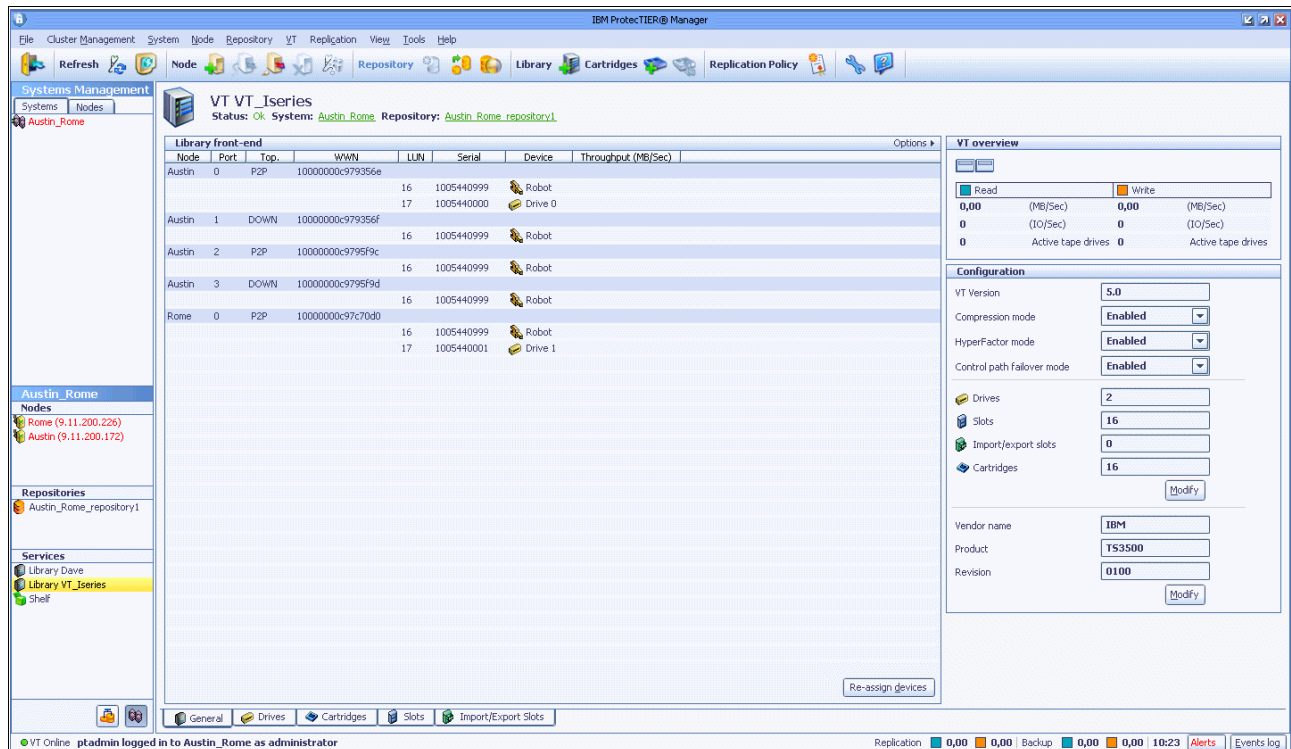


Figure 6-38 New VTL

As can be seen in Figure 6-38, we create one drive on each ProtecTIER node, and each drive is connected through a separate front-end port (FE port 0) in a ProtecTIER system.

**Note:** For host operating systems other than IBM i, consider LUN Masking in the HBA device on the host operating systems.

### 6.3.5 Managing the virtual devices in IBM i

After the front-end ports in the ProtecTIER system are connected to IBM i adapters and the virtual tape library is created and assigned to these ports, the tape library is recognized by the IBM i partition.

In System Service Tools (SST) menu, complete the following steps:

1. Select Option 1 = Start a service tool.
2. Select Option 7 = Hardware service manager.
3. Select Option 2 = Logical hardware resources (buses, IOPs, controllers, and so on).
4. Select Option 1 = System bus resources.
5. Select Option 9 = Show resources associated with the IOP to which the FC adapters is attached.

The tape library (robot) and the tape drive under that FC adapter are shown as operational (Figure 6-39).

Logical Hardware Resources Associated with IOP				
Type options, press Enter.				
2=Change detail 4=Remove 5=Display detail 6=I/O debug				
7=Verify 8=Associated packaging resource(s)				
		Resource		
Opt	Description	Type-Model	Status	Name
	Combined Function IOP	2844-001	Operational	CMB08
	Storage IOA	280D-001	Operational	DC11
	Tape Library	3854-032	Operational	TAPMLB15
	Tape Unit	3580-003	Operational	TAP29
F3=Exit F5=Refresh F6=Print F8=Include non-reporting resources				
F9=Failed resources F10=Non-reporting resources				
F11=Display serial/part numbers F12=Cancel				

Figure 6-39 Logical HW Resources Associated with IOP

By specifying Option 5 = Display detail at the tape library, you can see its serial number, as shown in Figure 6-40.

Library Hardware Resource Detail	
Description	Tape Library
Type-model	3854-032
Status	Operational
Serial number	<b>00-9990402</b>
Part number	
Resource name	TAPMLB15
PCI bus	
System bus	13
System board	0
System card	32
Library	
I/O adapter	2
I/O bus	
Library label	1
Controller	
Device	
F3=Exit F5=Refresh F6=Print	
F9=Change detail F12=Cancel	

Figure 6-40 Library Hardware Resource Detail

Provided that automatic configuration is enabled in IBM i, device descriptions of the drives and robots are automatically created as soon as the devices are connected. Therefore, the descriptions of both tape library (TAPMLBxxx) and tape drive (TAPxxx) can be seen in IBM i. Initially, both tape libraries in IBM i have the status of available. Also, initially the status of both tape drives is unavailable. Each tape library has one tape drive assigned. See Figure 6-41.

```

Work with Device Descriptions                               System: I5PFE

Position to . . . . . Starting characters

Type options, press Enter.
 2=Change 3=Copy 4=Delete 5=Display 6=Print 7=Rename
 8=Work with status 9=Retrieve source

Opt Device  Type  Text
TAPMLB14  3854  CREATED BY AUTO-CONFIGURATION
TAPMLB15  3854  CREATED BY AUTO-CONFIGURATION
TAP29     3580  CREATED BY AUTO-CONFIGURATION
TAP30     3580  CREATED BY AUTO-CONFIGURATION

Bottom

Parameters or command
==>
F3=Exit F4=Prompt F5=Refresh F6=Create F9=Retrieve F12=Cancel
F14=Work with status

```

Figure 6-41 Work with Device Descriptions

### Tape pooling

To see both tape drives under the tape library TAPMLB15, we first make both tape libraries unavailable by using Option 8 = Work with status at each library. Then we use Option 2 = Make unavailable. Next, we make the tape library TAPMLB15 available, as shown in Figure 6-42.

```

I5PFE                               System:

Work with Devices

Type options below, then press Enter.
 1=Make available 2=Make unavailable 5=Display details
 7=Display message 8=Work with controller and line 9=Rename
13=Change description

Opt Device  Type  Status
2 TAPMLB15  3854  Available to use

Bottom

F1=Help F3=Exit F5=Refresh F9=Command line F11=Display descriptions
F12=Cancel F17=Top F18=Bottom F21=Select assistance level
TAPMLB15 made available.

```

Figure 6-42 Work with Devices

With these actions, both tape drives are seen in the tape library TAPMLB15 (tape pooling).





# Backup and restore applications

In this chapter, we provide setup recommendations for some of the major backup applications, considerations for operations, and how to prepare for a disaster.

## 7.1 Considerations for all backup servers

Many backup servers have features and settings that can be used to optimize performance when writing data to physical cartridges. Because ProtecTIER provides a *virtual tape library* with *virtual drives* and cartridges to the *backup server*, some of the settings that are optimized for *real tape* are no longer required, and might even have a detrimental effect on the ProtecTIER *factoring ratio* and performance.

The following considerations are fairly generic and are common to all backup servers. Check the current settings of your backup server and apply the settings that can be implemented.

Some of the backup server specific sections of this chapter have more detailed considerations than these general topics. When this is the case, the specific consideration should take precedence.

### 7.1.1 General considerations

As a general rule, the preferred method of operation is to imitate the procedure 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 make sure that there is enough bandwidth and time allotted so that there will be 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 into ProtecTIER.
3. The system should be set up so that replication will then start and be finished before the next backup cycle starts.
4. At this point, the user should have a complete and easily recoverable set of their latest daily backup, including the backup application catalog image.
5. In case of a disaster, the user can revert back to that 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).

These are general considerations that are not specific to any backup server. More specific considerations are given in the following sections.

#### Interoperability

Check the IBM Interoperability Matrix to ensure that the version of the backup server and the operating system that you are running on are supported for ProtecTIER. You can view the matrix at the following address:

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

## Software compatibility

Make sure that your backup server version, platform, and operating system version are on the supported hardware and software list for ProtecTIER. You can view the list at the following address:

<http://www-03.ibm.com/systems/storage/tape/library.html#compatibility>

## Software currency

Ensure that the backup server has the latest patches or maintenance level to improve the overall factoring performance.

Ensure that the operating system of the platform that the backup server is running on has the latest patches or maintenance level to improve the overall HyperFactor performance.

## Tape library zoning

The backup server should have dedicated HBA ports for the ProtecTIER virtual tape library. These ports can be shared with a physical tape library. However, the physical tape library must not be in the same SAN zone as the virtual tape library.

## Compression

Standard compression will effectively scramble the data sent to ProtecTIER, making pattern matching difficult; real-time compression in databases and IBM Storwize does not scramble this data. As can be expected, this situation has an effect on data-matching rates, even if the same data is sent each time. ProtecTIER compresses the data that it sends to the back-end physical disk, after it has been received by the virtual tape drives and deduplicated.

If a data stream is not already compressed, our experience suggests that it is most efficient to let ProtecTIER compress data after deduplication. Compressed data can have fewer duplicate blocks than uncompressed data, so the effectiveness of deduplication can be diminished. Workloads and results vary, so we encourage experimentation.

Compressed data can reach backup systems and ProtecTIER in a number of ways. Consider the following use cases where the data stream going to ProtecTIER remains compressed:

- ▶ File backups may contain files from remote systems that enable compression in their backup client software to conserve network bandwidth. File system backups may also contain compressed file formats, such as GIF images, which remain compressed until opened by a GIF reader.
- ▶ Block-level database backups may contain compressed database objects. Database compression features deliver about 5:1 data reduction for Oracle, DB2, and Informix, and a bit less for databases with free compression.
- ▶ NDMP image backups may contain data compressed by Storwize or another process. NDMP uses a back channel to create backups, so Storwize or other rehydration processes are bypassed.

Some of the actions you can take regarding data being compressed before reaching ProtecTIER are:

- ▶ ProtecTIER does not require additional compression to be effective. ProtecTIER performs compression, by default, after the deduplication process. Do not allocate disk for ProtecTIER VTLs that compresses data by default.
- ▶ ProtecTIER can manage multiple VTLs, each with its own configuration. For compressed data streams, create a new ProtecTIER VTL with compression turned off. Compressing data a second time can cause data expansion, so compressed data should be segregated in ProtecTIER whenever possible.

- ▶ File systems with small files (under 32 KB), whether or not they are compressed, should not be sent directly to ProtecTIER. The following options should be considered to prevent ProtecTIER from bypassing small files:
  - Large NAS systems with small files should use NDMP for image backups and then send those files to ProtecTIER.
  - File level backups should first back up to backup application Disk Storage Pools and then the Disk Storage Pool can be copied to ProtecTIER.
- ▶ If a data stream is not already compressed, our experience suggests it is most efficient to let ProtecTIER compress data after deduplication. Compressed data can have fewer duplicate blocks than uncompressed data, so the effectiveness of deduplication can be diminished. Workloads and results vary, so we encourage experimentation.
- ▶ Always encrypt last. Deduplicating encrypted data is ineffective. Compressing encrypted data can decrease security. Drive-level encryption has no performance impact, and it ensures that encryption occurs last.

## Encryption

Encryption makes each piece of data sent to ProtecTIER unique, including duplicate data. As can be expected, this has an effect on data matching rates and the factoring performance because even if the same data is sent each time, it will appear differently to the deduplication engine.

You should disable any encryption features for the ProtecTIER storage pool in the backup server.

## Multiplexing

Do not use the multiplexing feature of any backup application with the ProtecTIER storage pool. Although ProtecTIER works with these features, the benefits (disk savings) of the HyperFactor algorithm and compression are greatly reduced.

You should disable any multiplexing features in the backup server for the ProtecTIER storage pool.

## Tape block sizes

To optimize the backup server, set the block size for data sent to the (virtual) tape drives to be at least 256 KB.

## Operating system clock time

If possible, ensure that the system time on the backup servers and ProtecTIER systems are synchronized by the same source, such as an NTP server or other means. This makes any problem diagnosis activities easier to conduct, should they be required.

## Ports

If possible, provision the HBA ports by connecting to the IBM System Storage TS7650, TS7610, or TS7650G storage systems with no other devices (disk or real tape) connected or zoned to those ports.

Ensure that the HBAs in the backup server are spread across all the PCI buses.



## 7.1.2 Additional factors

Another factor that affects performance in a ProtecTIER environment is the type of data being targeted for backup. Some data is well suited to 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. You might want to consider some of the following options:

- ▶ Larger NAS systems should use NDMP for image backups and then be sent to ProtecTIER.
- ▶ File level backups should first back up to backup application Disk Storage Pools, and then the Disk Storage Pool can be copied to ProtecTIER.

We discuss known configuration changes suited to specific data types in 7.7, “Data types” on page 408.

## 7.1.3 Assessing cartridge status and synchronizing with the catalog

After the disaster recovery (DR) site backup application server is recovered, the user must review the status of the replicated cartridges to ensure that their replication is consistent with the backup catalog or database. If there are a small number of cartridges, the user can inquire through the ProtecTIER Manager cartridge view for the last synchronization time of the cartridges from the list obtained through the backup application included in the catalog image.

This section describes an available report that shows the status of all cartridges involved in replication if there are a large number of cartridges. The report works with any backup application and is generated from ProtecTIER Manager. We explain the process for assessing cartridge status and synchronizing the backup application catalog with the cartridges for disaster recovery. Before running a restore for disaster recovery, you must verify that the list of associated cartridges is completely replicated to the remote site, or an earlier full backup image must be used for recovery.

The easiest way to determine the time of the last full backup is by having a specific time each day when your replication backlog is zero (that is, there is no pending data to replicate and backups are not running). If this is not the case, 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 backup application catalog

There are several ways to obtain a copy of the catalog at the remote site:

- ▶ From a catalog backup on a virtual cartridge that will be replicated to the remote site
- ▶ From disk-based replication, or by other means

If the catalog is backed up to a virtual cartridge, through the cartridge view of the library in ProtecTIER Manager, query each of the cartridges used for catalog backup to find the most recent synchronization dates marked on the cartridges. Assuming that there are multiple backup copies, you must find the latest backup that finished 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 remote site:

Each cartridge has a last synchronization time that displays the last time that the cartridge’s data was fully replicated to the remote site. (The synchronization time is updated during the replication, not only when the replication for this cartridge is finished.)

The cartridge marked with the most recent last synchronization time stamp should be used to recover the backup application catalog.

The procedure for recovering the selected catalog backup depends on the backup application and should be documented in the backup application's official documentation.

## Recovering the data

Scan the backup application catalog and search for the full backup image that you want to recover:

- ▶ Get the start and end backup time of the full backup image.
- ▶ View the list of cartridges associated with this full backup.

After it is invoked, the report might take a few minutes to run, during which time ProtecTIER Manager will be unavailable. When the report is completed, the system produces a .csv file that lists each cartridge and several key metrics per cartridge.

The elements that are provided in this report include:

- ▶ Cartridge unique ID
- ▶ Barcode, nominal size
- ▶ Last update time
- ▶ Last synchronization point repository unique ID
- ▶ Last synchronization point source time
- ▶ Last synchronization point destination time

Using ProtecTIER Manager, generate the cartridges' .csv file report by selecting **Replication** → **Create** and downloading the statistics (Figure 7-1).

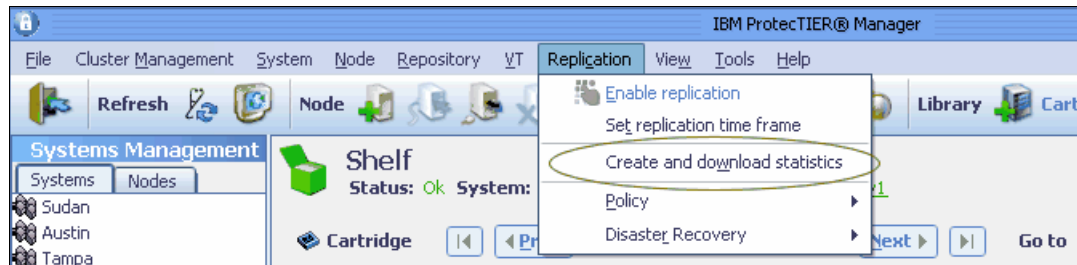


Figure 7-1 Create and download statistics

The .csv file is created in the /pt\_work directory, and a window opens that prompts you to download the file to your workstation (Figure 7-2). You can open the file with any spreadsheet application that you want to use (the default is Excel).

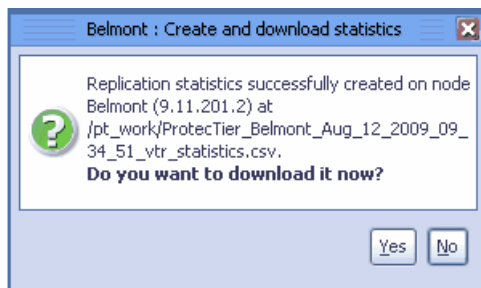


Figure 7-2 Create statics window

The report helps you determine the status of each cartridge at the remote site should a disaster strike while the replication operation is midstream and, in addition, if some of the replicated cartridges at the remote site only contain partial data.

The report, which includes statistics for all the cartridges at the remote site, indicates whether the replica cartridge data is consistent with the associated sync time stamp (that is, the current data on the tape represents what it was at the specified sync time).

To use this tool during a DR situation, complete the following steps:

1. Get the end time of the full backup image that the user is attempting to restore from the backup application catalog or database. Compare the last update time, which represents the last time that the replica was updated, with 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 a consistent point in time. Otherwise, the cartridge is incomplete or in transit. If the cartridge has a consistent point in time, ensure that this time stamp is larger than the full backup image end time, which indicates that the cartridge contains all the required data for this recovery operation. Otherwise, you must use a previous full backup image for recovery.
2. View the list of cartridges that are required for restoring the primary backup environment.
3. Using ProtecTIER Manager, generate the cartridges' .csv file report by selecting **Replication** → **Create** and downloading the statistics. The .csv file is created in the /pt\_work directory and a window opens that prompts you to download the file to your workstation. You can open the file with any spreadsheet application that you desire, such as IBM Lotus Symphony™, but the default is Excel (Figure 7-3).

	A	B	C	D	E	F	G
1	cart unique id	barcode	nominal size	last update time	last sync	last sync point source time	last sync point destination time
2	1000	0003E8	1,048,576.00	05/07/2009 02:31	21	05/07/2009 02:31	05/07/2009 02:31
3	1001	0003E9	2,097,152.00	05/07/2009 03:30	21	05/07/2009 03:31	05/07/2009 03:31
4	1002	0003EA	3,145,728.00	05/07/2009 04:31	21	05/07/2009 04:31	05/07/2009 04:31
5	1003	0003EB	4,194,304.00	05/07/2009 05:31	21	05/07/2009 05:31	05/07/2009 05:31
6	1004	0003EC	5,242,880.00	05/07/2009 06:31	21	05/07/2009 06:31	05/07/2009 06:31
7	1005	0003ED	1,048,576.00	05/07/2009 07:31	21	05/07/2009 07:31	05/07/2009 07:31
8	1006	0003EE	2,097,152.00	05/07/2009 08:31	21	05/07/2009 08:31	05/07/2009 08:31
9	1007	0003EF	3,145,728.00	05/07/2009 09:31	21	05/07/2009 09:31	05/07/2009 09:31

Figure 7-3 Cartridge status report

4. The report, which includes statistics for all the cartridges at the remote site, indicates whether the replica cartridge data is consistent with the associated synchronization time stamp (that is, the current data on the tape represents what it was at the specified synchronization time).
5. Compare the last update time, which represents the last time that the replica was updated, with the last synchronization point destination time. If the last update time is less than or equal to the destination synchronization time, the replica cartridge has a consistent point in time. Otherwise, the cartridge is incomplete or in transit.

If the cartridge has a consistent point in time, ensure that this time stamp is larger than the end time of the full image that you are attempting to restore. This indicates that the cartridge contains all the required data for this recovery operation. Otherwise, you must restore an earlier copy of the same full backup image.

**Important:** When processing the cartridge list to discover a complete set of DR tapes, you must keep track of 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 it might be set to the incorrect date and time and therefore be unreliable. Use the source date and time rather than the destination synchronization time when comparing cartridge states to the backup catalog/database. The destination synchronization time should only be used to determine which cartridges are whole.

In addition, 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 the DR administrator or operators. For example, if the backup server is two hours behind, a cartridge might have a synchronization time that precedes its backup complete time, that is, it will appear as a previous, old backup.

If there is uncertainty regarding the time differences, compare the nominal size of the cartridge to the catalog/DB value as an additional (not a substitute) layer of verification.

You might have a case where the cartridge sync point is after the backup start time, but before the end of the backup. This 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 they receive a sync point earlier than the backup end time.

As such, 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.

## 7.2 IBM Tivoli Storage Manager

IBM Tivoli Storage Manager is a storage software suite that addresses the challenges of complex storage management in distributed heterogeneous environments. It protects and manages a broad range of data, from workstations to the corporate server environment.

When performing standard backups, IBM Tivoli Storage Manager uses the incremental backup method. Progressive backup methodology (often referred to as incremental backup) saves time and storage space by backing up only new and modified files. The progressive backup feature uses the IBM Tivoli Storage Manager relational database to track data wherever it is stored, delivering direct one-step file restore. Progressive backup eliminates the need for traditional full-plus-incremental or full-plus-differential backup and restore procedures, commonly used by other storage management products.

There are some additional IBM Tivoli Storage Manager products designed specifically to interact with third-party software products, such as email (Microsoft Exchange) and databases (Oracle or MS-SQL). These are known as Tivoli Data Protection (TDP) modules.

Other considerations are:

- ▶ For improved performance in terms of factoring, use IBM Tivoli Storage Manager V5.4 or later.
- ▶ Client compression should be disabled.
- ▶ When using Windows-based IBM Tivoli Storage Manager servers, the IBM Tivoli Storage Manager driver for tape and libraries for Windows must be used. Native Windows drivers for the emulated P3000 libraries and DLT7000 drives are not supported.

Figure 7-4 illustrates a typical Tivoli Storage Manager environment using ProtecTIER. 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 backed up from the clients it is serving. Every action and backup set that Tivoli Storage Manager processes 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 that is contained on the storage devices.

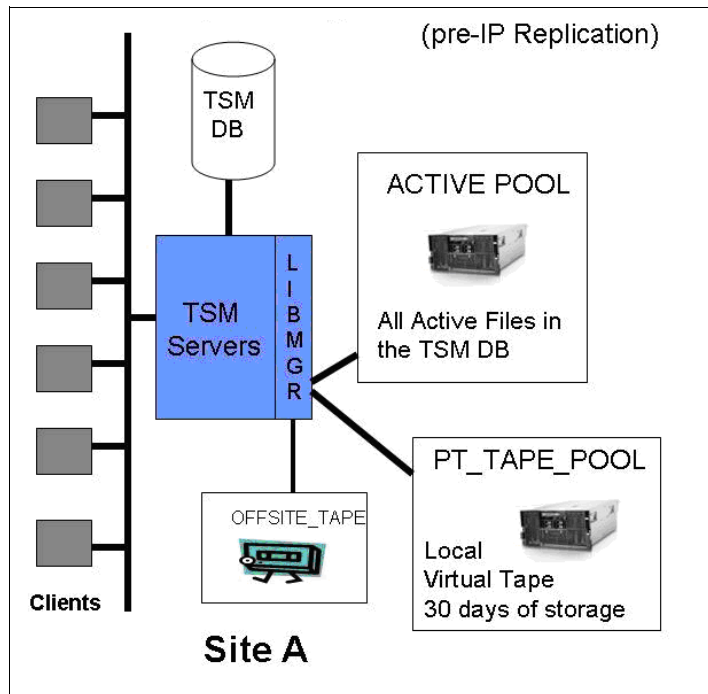


Figure 7-4 Typical Tivoli Storage Manager environment with ProtecTIER (pre-replication)

ProtecTIER 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 in Figure 7-4)

The user can also maintain another storage pool to create real physical tapes to take offsite (called OFFSITE\_TAPE in our example). The user has sized the ProtecTIER system to store all active and about 30 days of inactive client files on virtual tape. The customer also created an ACTIVE Tivoli Storage Manager pool, which is also hosted on the ProtecTIER system, which contains the most recent (active) file backed up from all client servers. The ACTIVE pool is from where client restores will come. The advantage of this architecture is that it has eliminated the use of physical tape in the data center and allows restores to occur much faster, as they are coming from the ProtecTIER disk-based virtual tape versus real tape.

## 7.2.1 ProtecTIER IP replication in the Tivoli Storage Manager environment

The ProtecTIER IP replication functionality provides a powerful tool that enables users to design a robust disaster recovery architecture. Thanks to the deduplication of data, users can now electronically vault backup data with much less bandwidth, thus changing the paradigm of how data is taken offsite for safe-keeping. ProtecTIER IP replication can eliminate the expensive and labor-extensive handling, transporting, and storing of real physical tapes for DR purposes.

Figure 7-5 illustrates how ProtecTIER's IP replication functionality can be used in a Tivoli Storage Manager environment. This user has chosen to use ProtecTIER to replicate all of the virtual tapes in the PT\_TAPE\_POOL offsite for DR purposes. The user also backs up the Tivoli Storage Manager database to virtual tapes. These database backup virtual tapes are also replicated to site B. In the event of a disaster, the user now has the ability to restore the Tivoli Storage Manager server at site B, which is then connected to a ProtecTIER virtual tape library that contains the Tivoli Storage Manager database on virtual tape and all of the client ACTIVE files on virtual tapes. Offsite physical tape is now only used for long-term archiving. The physical tapes are also contained in the Tivoli Storage Manager database that is protected on virtual tape at site B.

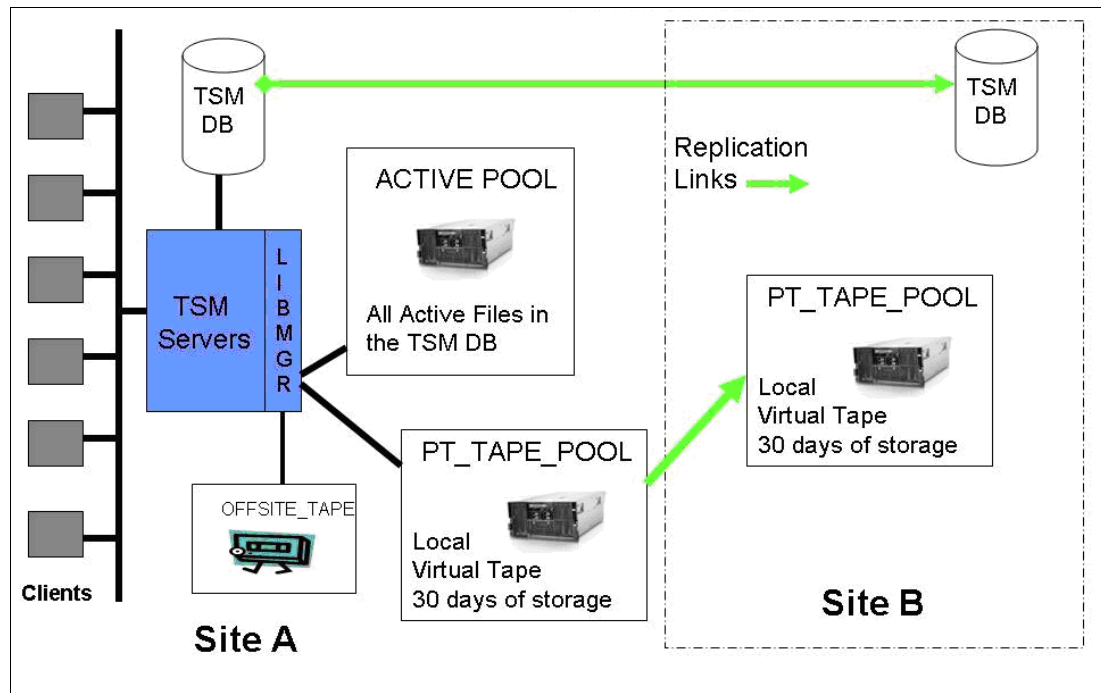


Figure 7-5 ProtecTIER replication and Tivoli Storage Manager

## 7.2.2 Implementing a virtual tape library

In the following sections, we walk through the steps required to define and connect a virtual tape library (VTL) in ProtecTIER to an IBM Tivoli Storage Manager server and make it usable.

For the purposes of this example, the Tivoli Storage Manager server is named FREEWAY and is running Version 5.5 with Level 2.0. The host server is called freeway.storage.tucson.ibm.com and is running AIX 5L Version 5.3.9.0.

## 7.2.3 The ProtecTIER virtual tape library definition

Using ProtecTIER Manager, a virtual tape library named TSM-Libr has already been created for use by IBM Tivoli Storage Manager (Figure 7-6). This library name is an internal name used only within ProtecTIER and is not visible anywhere outside of it. For details about how to create a virtual tape library using ProtecTIER Manager, refer to 5.6.1, “Creating libraries” on page 230.

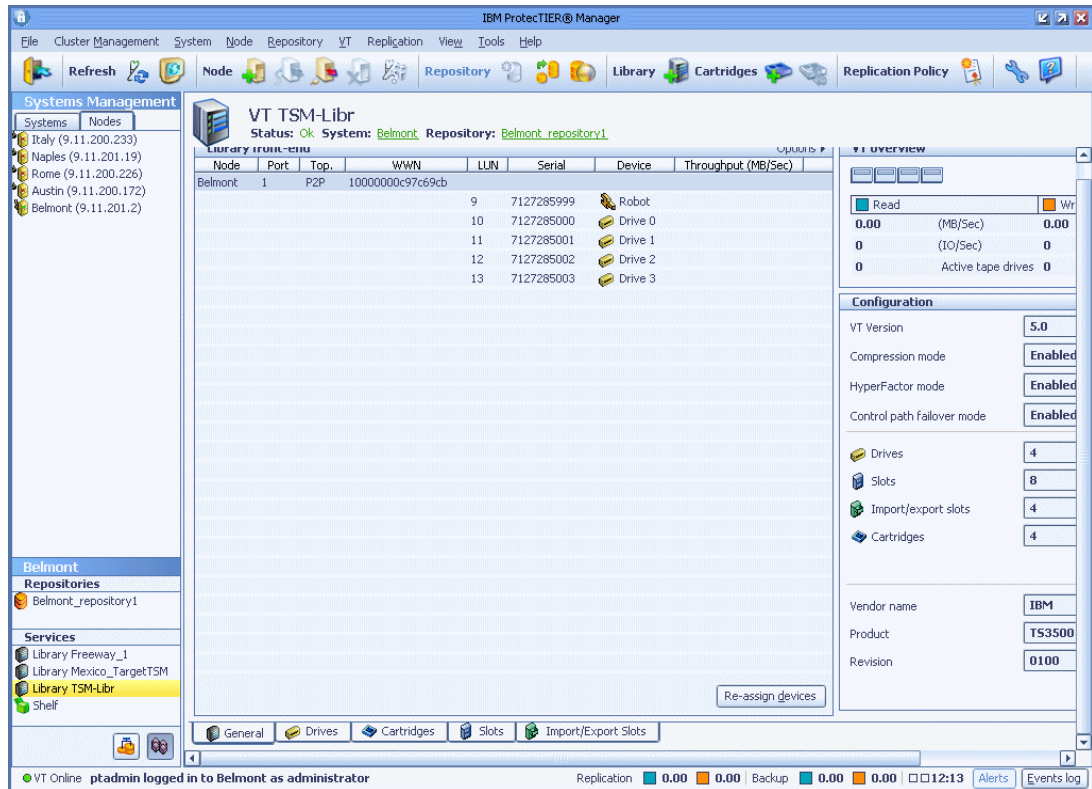


Figure 7-6 Summary of virtual tape library TSM-Libr (in ProtecTIER Manager)

The library was given four virtual LTO3 cartridge drives. The tape drives all use the same front-end port (this was due to the hardware resources that were chosen while writing this book, rather than for any performance or configuration reasons). The virtual robot is defined on the ProtecTIER node called Belmont.

Eight library slots were created along with four import/export (I/E) slots for completeness of the library emulation. Four virtual LTO2 cartridges were created also.

The four virtual LTO3 tape drive definitions are detailed in Figure 7-7.

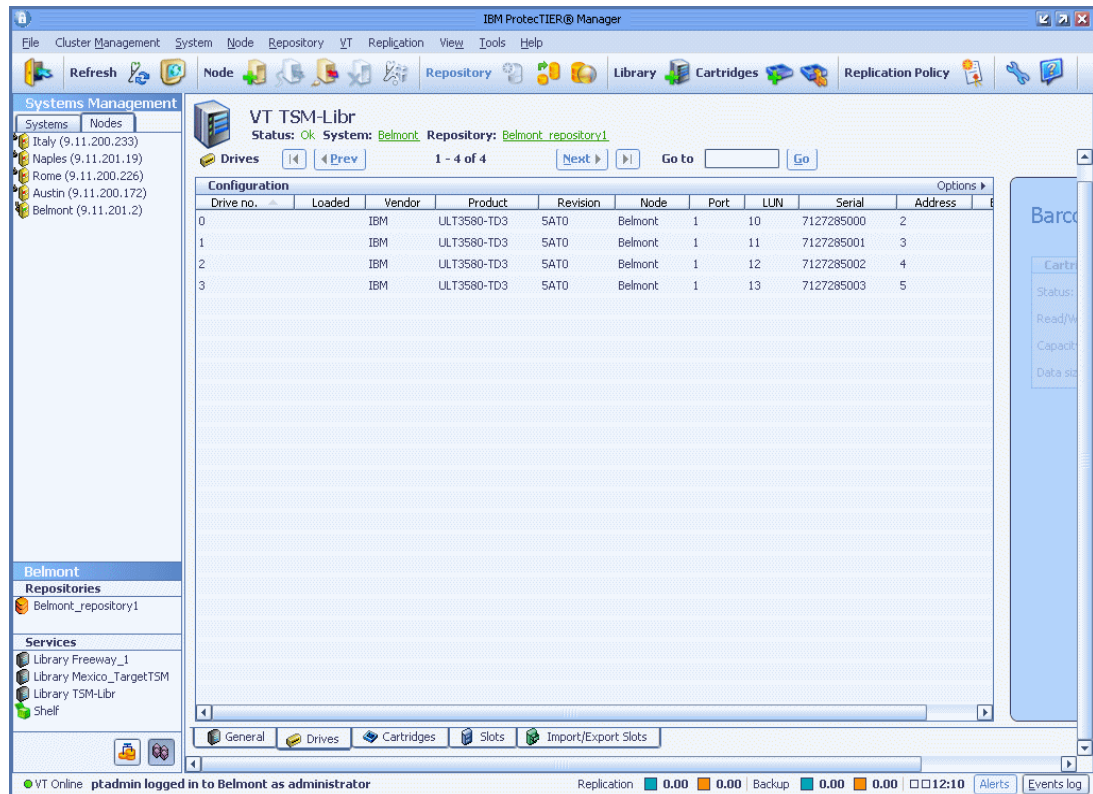


Figure 7-7 LTO3 drive definitions in virtual tape library TSM-Libr (in ProtecTIER Manager)

Note the logical unit number (LUN) assigned to each virtual tape drive. It is important to know that this is the ProtecTIER LUN number only and the host operating system will almost certainly assign each drive a different LUN than that which appears here, as the host will have more logical units or resources than just these tape drives.

During the definition of the library, ProtecTIER assigns serial numbers to each drive, seeded from a random number. These serial numbers can be important later when defining IBM Tivoli Storage Manager paths to link the host devices and the IBM Tivoli Storage Manager tape drive definitions. It can be helpful to know the drive serial number if you should experience any problems when defining the paths. You can use the serial number as a common reference TSM point when matching a drive with its host device file name.

An element number is the other address of each object in the tape library. Everything that can be located (such as robot;drive;slot;io slot) has its unique element number in the library. Tivoli Storage Manager (and any other tape library related software) uses this number as the address in the SCSI command to drive the robot to work. The element number can also be important later when defining the IBM Tivoli Storage Manager tape drives.



The four virtual LTO3 cartridges definitions are detailed in Figure 7-8.

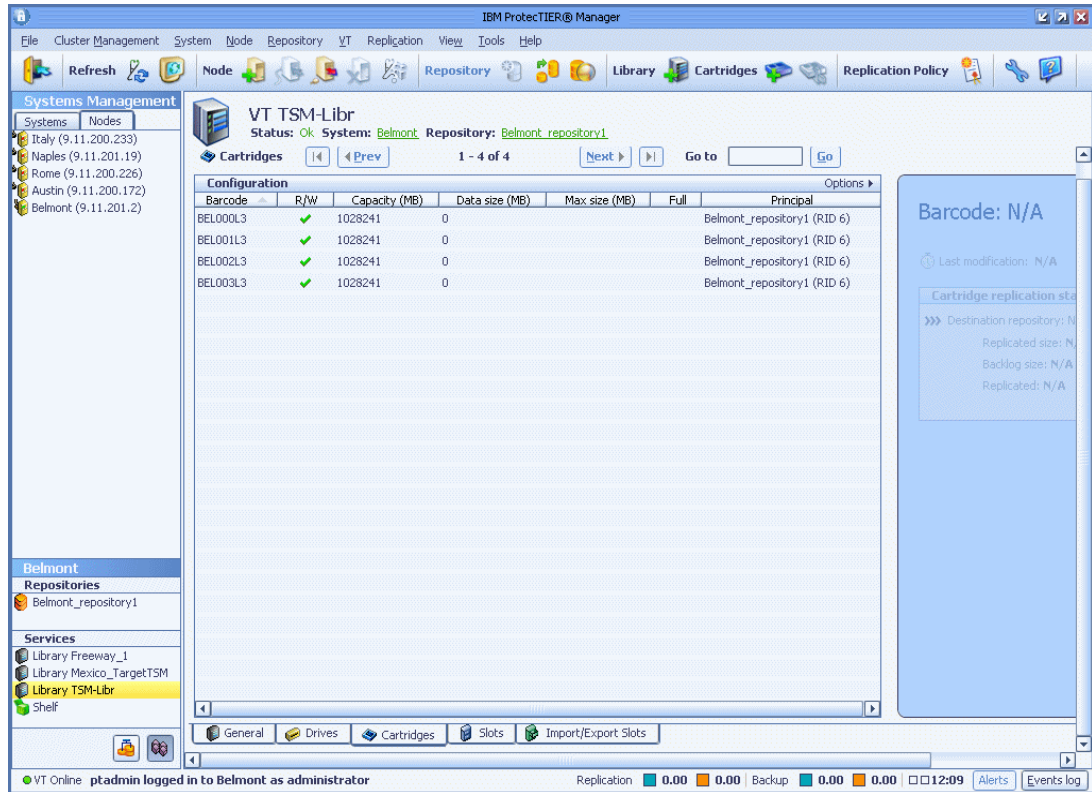


Figure 7-8 LTO3 cartridge definitions in virtual tape library TSM-Libr (in ProtecTIER Manager)

All the cartridges begin with a barcode prefix of BEL. They are all write enabled and currently hold no data. Their current estimated capacity is approximately 4 TB, but this capacity will change over time, as the cartridges have not been defined with a maximum capacity or size.

## 7.2.4 Defining the virtual tape library to AIX with IBM Tivoli Storage Manager

The steps needed to define a ProtecTIER virtual tape library and drives to IBM Tivoli Storage Manager are identical to those required for the corresponding physical tape library and drives.

To define a physical or virtual library in Tivoli Storage Manager, complete the following steps:

1. Install Device Drivers for tape, as demonstrated in 6.2, “Installing and configuring the device driver in OS” on page 297.
2. Run the AIX configuration manager command (**cfgmgr**) to assign device special files to each tape drive, as shown in Example 7-1.

*Example 7-1 Configuration and recognition of virtual tape drives in AIX*

```
freeway> cfgmgr

freeway> lsdev -Cc tape
rmt0 Available 04-08-02 IBM 3580 Ultrium Tape Drive (FCP)
rmt1 Available 04-08-02 IBM 3580 Ultrium Tape Drive (FCP)
rmt2 Available 08-08-02 IBM 3580 Ultrium Tape Drive (FCP)
rmt3 Available 08-08-02 IBM 3580 Ultrium Tape Drive (FCP)
rmt4 Available 08-08-02 IBM 3580 Ultrium Tape Drive (FCP)
```

rmt5	Available	0B-08-02	IBM 3580 Ultrium Tape Drive (FCP)
rmt6	Available	0B-08-02	IBM 3580 Ultrium Tape Drive (FCP)
rmt7	Available	0B-08-02	IBM 3580 Ultrium Tape Drive (FCP)
rmt8	Available	0B-08-02	IBM 3580 Ultrium Tape Drive (FCP)
rmt9	Available	0B-08-02	IBM 3580 Ultrium Tape Drive (FCP)
rmt10	Available	0B-08-02	IBM 3580 Ultrium Tape Drive (FCP)
rmt11	Available	0B-08-02	IBM 3580 Ultrium Tape Drive (FCP)
rmt12	Available	0B-08-02	IBM 3580 Ultrium Tape Drive (FCP)
rmt13	Available	0B-08-02	IBM 3580 Ultrium Tape Drive (FCP)
smc0	Available	04-08-02	IBM 3584 Library Medium Changer (FCP)
smc1	Available	0B-08-02-PRI	IBM 3584 Library Medium Changer (FCP)
smc2	Available	0B-08-02	IBM 3584 Library Medium Changer (FCP)

---

3. Obtain information about the tape devices.

In our environment, we have four tape drives and one robot created and assigned. We must ensure that each device of rmtx and smc search the new drives by issuing the **lscfg** command in AIX. See Example 7-2.

*Example 7-2 Obtain information about each device in AIX*

---

```
freeway> lscfg -vpl rmt10
rmt10          U787A.001.DPM04K9-P1-C6-T1-W1000000C97C69CB-LA000000000000
IBM 3580 Ultrium Tape Drive (FCP)
```

```
Manufacturer.....IBM
Machine Type and Model.....ULT3580-TD3
Serial Number.....7127285000
Device Specific.(FW).....5AT0
```

```
freeway> lscfg -vpl rmt11
rmt11          U787A.001.DPM04K9-P1-C6-T1-W1000000C97C69CB-LB000000000000
IBM 3580 Ultrium Tape Drive (FCP)
```

```
Manufacturer.....IBM
Machine Type and Model.....ULT3580-TD3
Serial Number.....7127285001
Device Specific.(FW).....5AT0
```

```
freeway> lscfg -vpl rmt12
rmt12          U787A.001.DPM04K9-P1-C6-T1-W1000000C97C69CB-LC000000000000
IBM 3580 Ultrium Tape Drive (FCP)
```

```
Manufacturer.....IBM
Machine Type and Model.....ULT3580-TD3
Serial Number.....7127285002
Device Specific.(FW).....5AT0
```

```
freeway> lscfg -vpl rmt13
rmt13          U787A.001.DPM04K9-P1-C6-T1-W1000000C97C69CB-LD000000000000
IBM 3580 Ultrium Tape Drive (FCP)
```

```
Manufacturer.....IBM
Machine Type and Model.....ULT3580-TD3
Serial Number.....7127285003
Device Specific.(FW).....5AT0
```

```

freeway> lscfg -vpl smc2
smc2          U787A.001.DPM04K9-P1-C6-T1-W10000000C97C69CB-L9000000000000
IBM 3584 Library Medium Changer (FCP)

Manufacturer.....IBM
Machine Type and Model.....03584L32
Serial Number.....0071272859990402
Device Specific.(FW).....0100

```

---

You also can obtain information about the element number by using the `tapeutil` tool in AIX, as shown in Example 7-3. *XXX Address* is the element number.

*Example 7-3 Use `tapeutil` to get more information about the tape library*

---

```

freeway> tapeutil -f /dev/smc2 invent
Reading element status...

Robot Address 0
  Robot State ..... Normal
  ASC/ASCQ ..... 0000
Media Present ..... No
  Source Element Address Valid ... No
Media Inverted ..... No
  Volume Tag .....

Import/Export Station Address 64514
  Import/Export State ..... Normal
  ASC/ASCQ ..... 0000
Media Present ..... No
  Import Enabled ..... Yes
  Export Enabled ..... Yes
  Robot Access Allowed ..... Yes
  Source Element Address Valid ... No
Media Inverted ..... No
  Volume Tag .....

Import/Export Station Address 64515
  Import/Export State ..... Normal
  ASC/ASCQ ..... 0000
Media Present ..... No
  Import Enabled ..... Yes
  Export Enabled ..... Yes
  Robot Access Allowed ..... Yes
  Source Element Address Valid ... No
Media Inverted ..... No
  Volume Tag .....

Import/Export Station Address 64516
  Import/Export State ..... Normal
  ASC/ASCQ ..... 0000
Media Present ..... No
  Import Enabled ..... Yes
  Export Enabled ..... Yes
  Robot Access Allowed ..... Yes
  Source Element Address Valid ... No
Media Inverted ..... No

```

```

Volume Tag .....

Import/Export Station Address 64517
  Import/Export State ..... Normal
  ASC/ASCQ ..... 0000
Media Present ..... No
  Import Enabled ..... Yes
  Export Enabled ..... Yes
  Robot Access Allowed ..... Yes
  Source Element Address Valid ... No
Media Inverted ..... No
  Volume Tag .....

Drive Address 2
  Drive State ..... Normal
  ASC/ASCQ ..... 0000
Media Present ..... No
  Robot Access Allowed ..... Yes
  Source Element Address Valid ... No
Media Inverted ..... No
  Same Bus as Medium Changer ..... Yes
  SCSI Bus Address Vaild ..... No
  Logical Unit Number Valid ..... No
  Volume Tag .....

Drive Address 3
  Drive State ..... Normal
  ASC/ASCQ ..... 0000
Media Present ..... No
  Robot Access Allowed ..... Yes
  Source Element Address Valid ... No
Media Inverted ..... No
  Same Bus as Medium Changer ..... Yes
  SCSI Bus Address Vaild ..... No
  Logical Unit Number Valid ..... No
  Volume Tag .....

Drive Address 4
  Drive State ..... Normal
  ASC/ASCQ ..... 0000
Media Present ..... No
  Robot Access Allowed ..... Yes
  Source Element Address Valid ... No
Media Inverted ..... No
  Same Bus as Medium Changer ..... Yes
  SCSI Bus Address Vaild ..... No
  Logical Unit Number Valid ..... No
  Volume Tag .....

Drive Address 5
  Drive State ..... Normal
  ASC/ASCQ ..... 0000
Media Present ..... No
  Robot Access Allowed ..... Yes
  Source Element Address Valid ... No

```

Media Inverted ..... No  
Same Bus as Medium Changer ..... Yes  
SCSI Bus Address Vaild ..... No  
Logical Unit Number Valid ..... No  
Volume Tag .....

Slot Address 1026

Slot State ..... Normal  
ASC/ASCQ ..... 0000  
Media Present ..... Yes  
Robot Access Allowed ..... Yes  
Source Element Address Valid ... No  
Media Inverted ..... No  
Volume Tag ..... BEL00L3

Slot Address 1027

Slot State ..... Normal  
ASC/ASCQ ..... 0000  
Media Present ..... Yes  
Robot Access Allowed ..... Yes  
Source Element Address Valid ... No  
Media Inverted ..... No  
Volume Tag ..... BEL001L3

Slot Address 1028

Slot State ..... Normal  
ASC/ASCQ ..... 0000  
Media Present ..... Yes  
Robot Access Allowed ..... Yes  
Source Element Address Valid ... No  
Media Inverted ..... No  
Volume Tag ..... BEL002L3

Slot Address 1029

Slot State ..... Normal  
ASC/ASCQ ..... 0000  
Media Present ..... Yes  
Robot Access Allowed ..... Yes  
Source Element Address Valid ... No  
Media Inverted ..... No  
Volume Tag ..... BEL003L3

Slot Address 1030

Slot State ..... Normal  
ASC/ASCQ ..... 0000  
Media Present ..... No  
Robot Access Allowed ..... Yes  
Source Element Address Valid ... No  
Media Inverted ..... No  
Volume Tag .....

Slot Address 1031

Slot State ..... Normal  
ASC/ASCQ ..... 0000  
Media Present ..... No

```

Robot Access Allowed ..... Yes
Source Element Address Valid ... No
Media Inverted ..... No
Volume Tag .....

Slot Address 1032
Slot State ..... Normal
ASC/ASCQ ..... 0000
Media Present ..... No
Robot Access Allowed ..... Yes
Source Element Address Valid ... No
Media Inverted ..... No
Volume Tag .....

Slot Address 1033
Slot State ..... Normal
ASC/ASCQ ..... 0000
Media Present ..... No
Robot Access Allowed ..... Yes
Source Element Address Valid ... No
Media Inverted ..... No
Volume Tag .....
freeway>

```

**Note:** For the virtual tape library created in ProtecTIER, the count of objects (such as drive;slot;io slot) and the element numbers are different from the real tape library that you are simulating.

Using this information, complete a tape library work sheet, as shown in Table 7-1. Do not forget to include the element numbers.

Table 7-1 Tape library worksheet

Device in OS	Type	VTL system	VTL node	VTL port	VTL port WWN	Serial number	Element number
smc2	3584	Belmont	Belmont	FE1	10000000C97C69CB	007127285999 0402	0
rmt10	LT3	Belmont	Belmont	FE1	10000000C97C69CB	7127285000	2
rmt11	LT3	Belmont	Belmont	FE1	10000000C97C69CB	7127285000	3
rmt12	LT3	Belmont	Belmont	FE1	10000000C97C69CB	7127285000	4
rmt13	LT3	Belmont	Belmont	FE1	10000000C97C69CB	7127285000	5

4. Define the library to IBM Tivoli Storage Manager using these commands:

```

DEFINE LIBRARY ...
QUERY LIBRARY

```

Example 7-4 shows the output of these commands.

*Example 7-4 Define the library in Tivoli Storage Manager*

```

tsm: FREEWAY>define library ts7650lib libtype=scsi
ANR8400I Library TS7650LIB defined.

```

```

tsm: FREEWAY>query libr ts7650lib f=d

      Library Name: TS7650LIB
      Library Type: SCSI
      ACS Id:
      Private Category:
      Scratch Category:
      WORM Scratch Category:
      External Manager:
      Shared: No
      LanFree:
      ObeyMountRetention:
      Primary Library Manager:
      WWN:
      Serial Number:
      AutoLabel: No
      Reset Drives: No
      Relabel scratch:Yes
      Last Update by (administrator): ADMIN
      Last Update Date/Time: 08/12/09 14:06:06

```

---

5. Define the path to the IBM Tivoli Storage Manager library using these commands:

```

DEFINE PATH ...
QUERY PATH

```

Example 7-5 shows the output of these commands. Note that there are many paths for other tape devices already in the Tivoli Storage Manager.

*Example 7-5 Define the path to the library in Tivoli Storage Manager*

```

tsm: FREEWAY>define path freeway ts7650lib srctype=server desttype=library
device=/dev/smc2
ANR1720I A path from FREEWAY to TS7650LIB has been defined.

```

```

tsm: FREEWAY>query path

```

Source Name	Source Type	Destination Name	Destination Type	On-Line
FREEWAY	SERVER	LTO_2	LIBRARY	Yes
FREEWAY	SERVER	DRIVE0	DRIVE	Yes
FREEWAY	SERVER	DRIVE1	DRIVE	Yes
FREEWAY	SERVER	LTO_7	LIBRARY	Yes
FREEWAY	SERVER	DRIVE0	DRIVE	Yes
FREEWAY	SERVER	DRIVE1	DRIVE	Yes
FREEWAY	SERVER	DRIVE2	DRIVE	Yes
FREEWAY	SERVER	DRIVE3	DRIVE	Yes
FREEWAY	SERVER	DRIVE4	DRIVE	Yes
FREEWAY	SERVER	DRIVE5	DRIVE	Yes
FREEWAY	SERVER	DRIVE6	DRIVE	Yes
FREEWAY	SERVER	TS7650LIB	LIBRARY	Yes

6. Define the drives to IBM Tivoli Storage Manager using these commands:

```

DEFINE DRIVE ...
QUERY DRIVE

```

Repeat the define command for each drive. Example 7-6 shows the output of these commands. Note that the last line is our definition.

*Example 7-6 Define the drives in Tivoli Storage Manager*

```
tsm: FREEWAY>define drive ts7650lib ts7650drv01
ANR8404I Drive TS7650DRV01 defined in library TS7650LIB.
```

```
tsm: FREEWAY>query drive
```

Library Name	Drive Name	Device Type	On-Line
LTO_2	DRIVE0	LTO	Yes
LTO_2	DRIVE1	LTO	Yes
LTO_7	DRIVE0	LTO	Yes
LTO_7	DRIVE1	LTO	Yes
LTO_7	DRIVE2	LTO	Yes
LTO_7	DRIVE3	LTO	Yes
LTO_7	DRIVE4	LTO	Yes
LTO_7	DRIVE5	LTO	Yes
LTO_7	DRIVE6	LTO	Yes
TS7650LIB	TS7650DRV01	UNKNOWN	Yes

7. Define the paths to the IBM Tivoli Storage Manager drives using these commands:

```
DEFINE PATH ...
QUERY PATH
```

Repeat the DEFINE command for each drive. Example 7-7 shows the output of these commands. Note that the last line is our definition.

*Example 7-7 Define the paths to the drives in Tivoli Storage Manager*

```
tsm: FREEWAY>define path freeway ts7650drv01 srctype=server desttype=drive
libr=ts7650lib device=/dev/rmt10
ANR1720I A path from FREEWAY to TS7650LIB TS7650DRV01 has been defined.
```

```
tsm: FREEWAY>query path
```

Source Name	Source Type	Destination Name	Destination Type	On-Line
FREEWAY	SERVER	LTO_2	LIBRARY	Yes
FREEWAY	SERVER	DRIVE0	DRIVE	Yes
FREEWAY	SERVER	DRIVE1	DRIVE	Yes
FREEWAY	SERVER	LTO_7	LIBRARY	Yes
FREEWAY	SERVER	DRIVE0	DRIVE	Yes
FREEWAY	SERVER	DRIVE1	DRIVE	Yes
FREEWAY	SERVER	DRIVE2	DRIVE	Yes
FREEWAY	SERVER	DRIVE3	DRIVE	Yes
FREEWAY	SERVER	DRIVE4	DRIVE	Yes
FREEWAY	SERVER	DRIVE5	DRIVE	Yes
FREEWAY	SERVER	DRIVE6	DRIVE	Yes
FREEWAY	SERVER	TS7650LIB	LIBRARY	Yes
FREEWAY	SERVER	TS7650DRV01	DRIVE	Yes



In this step, Tivoli Storage Manager retrieves the drive type and element number and changes the device type from Unknown to LTO, as shown in Example 7-8.

*Example 7-8 Query drive command*

```
tsm: FREEWAY>q drive
```

Library Name	Drive Name	Device Type	On-Line
LTO_2	DRIVE0	LTO	Yes
LTO_2	DRIVE1	LTO	Yes
LTO_7	DRIVE0	LTO	Yes
LTO_7	DRIVE1	LTO	Yes
LTO_7	DRIVE2	LTO	Yes
LTO_7	DRIVE3	LTO	Yes
LTO_7	DRIVE4	LTO	Yes
LTO_7	DRIVE5	LTO	Yes
LTO_7	DRIVE6	LTO	Yes
TS7650LIB	TS7650DRV01	LTO	Yes

```
tsm: FREEWAY>query drive TS7650LIB TS7650DRV01 f=d
```

```

Library Name: TS7650LIB
Drive Name: TS7650DRV01
Device Type: LTO
On-Line: Yes
Read Formats:
ULTRIUM3C,ULTRIUM3,ULTRIUM2C,ULTRIUM2,ULTRIUMC,ULTRIUM
Write Formats:
ULTRIUM3C,ULTRIUM3,ULTRIUM2C,ULTRIUM2
Element: 2
Drive State: UNKNOWN
Volume Name:
Allocated to:
WWN: 20000000C97C69CB
Serial Number: 7127285000
Last Update by (administrator): ADMIN
Last Update Date/Time: 08/12/09 14:25:06
Cleaning Frequency (Gigabytes/ASNEEDED/NONE): NONE

```

- Define an IBM Tivoli Storage Manager device class using these commands:

```
DEFINE DEVCLASS ...
QUERY DEVCLASS
```

Example 7-9 shows the output of these commands. Note that the last line is our definition.

*Example 7-9 Define the device class in Tivoli Storage Manager*

```
tsm: FREEWAY>define devclass ts7650dev devtype=lto format=drive libr=ts7650lib
ANR2203I Device class TS7650DEV defined.
```

```
tsm: FREEWAY>query devclass
```

Device Class Name	Device Access Strategy	Storage Device Pool Count	Device Type	Format	Est/Max Capacity (MB)	Mount Limit
-----	-----	-----	-----	-----	-----	-----

BIG	Sequential	1 LTO	DRIVE		DRIVES
DBFILE	Sequential	0 FILE	DRIVE	51,200.0	20
DISK	Random	3			
EXPORTFI- LE	Sequential	0 FILE	DRIVE	5,120,00 0.0	20
HIGHWAYD- EV	Sequential	0 SERVER		500.0	10
LTO_1	Sequential	0 LTO	DRIVE		DRIVES
LTO_2	Sequential	2 LTO	DRIVE		DRIVES
LTO_3	Sequential	0 LTO	DRIVE		DRIVES
LTO_4	Sequential	0 LTO	DRIVE		DRIVES
LTO_7	Sequential	2 LTO	DRIVE		DRIVES
TS7650DEV	Sequential	0 LTO	DRIVE		DRIVES

---

9. Define an IBM Tivoli Storage Manager storage pool that uses the device class defined in step 7 using these commands:

```
DEFINE STGPOOL ...
QUERY STGPOOL
```

Example 7-10 shows the output of these commands. Note that the last line is our definition.

*Example 7-10 Define the storage pool in Tivoli Storage Manager*

---

```
tsm: FREEWAY>define stgpool ts7650_lto ts7650dev maxscratch=20
ANR2200I Storage pool TS7650_LTO defined (device class TS7650DEV).
```

```
tsm: FREEWAY>query stgpool
```

Storage Pool Name	Device Class Name	Estimated Capacity	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Stora- ge Pool
ARCHIVEPOOL	DISK	2 G	0.0	0.0	90	70	LTO_2_POOL
BACKUPPOOL	DISK	10 G	0.0	0.0	20	0	LTO_2_POOL
BIGPOOL	BIG	0.0 M	0.0	0.0	90	70	
LTO_2_CPOOL	LTO_2	0.0 M	0.0				
LTO_2_POOL	LTO_2	9,131 G	5.9	80.0	90	70	
LTO_7_CPOOL	LTO_7	0.0 M	0.0				
LTO_7_POOL	LTO_7	99,839 G	3.2	4.0	90	70	
SPACEMGPOOL	DISK	0.0 M	0.0	0.0	90	70	
TS7650_LTO	TS7650DEV	0.0 M	0.0	0.0	90	70	

---

10. Label the cartridges in the library using these commands:

```
LABEL LIBVOLUME ...: Label the cartridges
Query REQ: To check if any prompt has been issued
Query PROC: To check if the label process finish or not
Query LIBVOL: To list the volumes we labeled.
```

Example 7-11 shows the output of these commands.

*Example 7-11 Labeling the virtual cartridges in Tivoli Storage Manager*

```
tsm: FREEWAY>label libvolume ts7650lib search=yes checkin=scratch
labelsource=barcode
ANS8003I Process number 1 started.
```

```
tsm: FREEWAY>query req
ANR8346I QUERY REQUEST: No requests are outstanding.
ANS8001I Return code 11.
```

```
tsm: FREEWAY>q proc
```

Process Number	Process Description	Status
1	LABEL LIBVOLUME	ANR8805I Labeling volumes in library TS7650LIB; 3 volume(s) labeled.

```
tsm: FREEWAY>query proc
ANR0944E QUERY PROCESS: No active processes found.
ANS8001I Return code 11.
```

```
tsm: FREEWAY>query libvol TS7650LIB
```

Library Name	Volume Name	Status	Owner	Last Use	Home	Device	Element Type
TS7650LIB	BEL000L3	Scratch				1,026	LTO
TS7650LIB	BEL001L3	Scratch				1,027	LTO
TS7650LIB	BEL002L3	Scratch				1,028	LTO
TS7650LIB	BEL003L3	Scratch				1,029	LTO

The virtual library is now defined to your IBM Tivoli Storage Manager server, and the virtual drives and cartridges are ready for use.

You now must use standard methods to alter your management class copy groups to change the destination value to point to the storage pool created for the virtual library.

**Note:** If you do not label the virtual cartridges before use, when Tivoli Storage Manager attempts to write data to them, the process will fail, and Tivoli Storage Manager will issue an error message saying that it could not read the internal label of the cartridge. If this error occurs, issue CHECKOUT LIBVOLUME commands to check out and reset the library status of all the cartridges (include the REMOVE=NO option so that they do not leave the virtual library) and label them again with the LABEL LIBVOLUME command.

If you forget to include the REMOVE=NO option in your CHECKOUT LIBVOLUME command, the library will place the virtual cartridges in the virtual import/export slots. You can view cartridges stored in these slots on the Import/Export tab of the Library window shown in Figure 10-45 on page 507. Using the menu options accessed by right-clicking the desired cartridge, you can relocate the cartridges back to standard virtual slots. After they are relocated to standard slot locations, use the LABEL LIBVOLUME command to label and check them in again. Alternatively, you can label them directly from the Import/Export slots by using the SEARCH=BULK option on the LABEL LIBVOLUME command.

**Note:** If IBM Tivoli Storage Manager has SANDISCOVERY ON when you are using the ProtecTIER Virtual Tape Library, it can cause problems with the tape drive path descriptions of a node if the node goes offline or the path to the port is broken. With this option on in this scenario, you must reconfigure all the tape paths or determine which device belongs to which path by serial number (which will take much longer). With a lot of virtual drives, this could be time consuming.

## 7.2.5 ProtecTIER and Tivoli Storage Manager attachment on Windows 2008

In this section, we show how to set up Tivoli Storage Manager with ProtecTIER in a Windows environment and make it usable for backup and restore. For the purposes of this example, the Tivoli Storage Manager server is named DECANter\_SERVER1 and is running Version 5.5 with Level 2.0. The host server is called Decanter and is running Windows 2008 Server Enterprise SP1.

Using the ProtecTIER Manager, we create a virtual tape library named VT Decanter\_1 for use by IBM Tivoli Storage Manager. This library name is an internal name used only within ProtecTIER and is not visible anywhere outside of it. For details about how to create a virtual tape library using ProtecTIER Manager, refer to 6.3.4, “Creating a VTL for IBM i” on page 334.

VT Decanter\_1 is on a two-node cluster. The VTL has 16 drives and one robot per node. Node Tampa devices are connected to Tampa FE port1 and node Bisbee devices are connected to Bisbee FE port1. For the Tivoli Storage Manager setup, we focus on robot, drive 0, and drive 2 on Tampa FE port1. Figure 7-9 shows the serial number and LUN for the robot and the drives on node Tampa (VT Decanter\_1). The library has 200 cartridges assigned. The barcode starts with TDK000L3 and goes to TDK199L3.

Library front-end								
Node	Port	Top.	WWN	LUN	Serial	Device	Throughput (MB/Sec)	Options
Tampa	1	FC-AL	10000000c97a2393			Robot		
				0	1334351999	Drive 0		
				1	1334351000	Drive 2		
				2	1334351002	Drive 4		
				3	1334351004	Drive 6		
				4	1334351006	Drive 8		
				5	1334351008	Drive 10		
				6	1334351010	Drive 12		
				7	1334351012	Drive 14		
				8	1334351014	Drive 16		
				9	1334351016	Drive 18		
				10	1334351018	Drive 20		
				11	1334351020	Drive 22		
				12	1334351022	Drive 24		
				13	1334351024	Drive 26		
				14	1334351026	Drive 28		
				15	1334351028	Drive 30		
16	1334351030							
Bisbee	1	FC-AL	10000000c97a2291			Robot		
				0	1334351999	Drive 1		
				1	1334351001	Drive 3		
				2	1334351003	Drive 5		
				3	1334351005			

Figure 7-9 Devices: Two-node cluster VT Decanter\_1

Figure 7-10 shows the drive definitions.

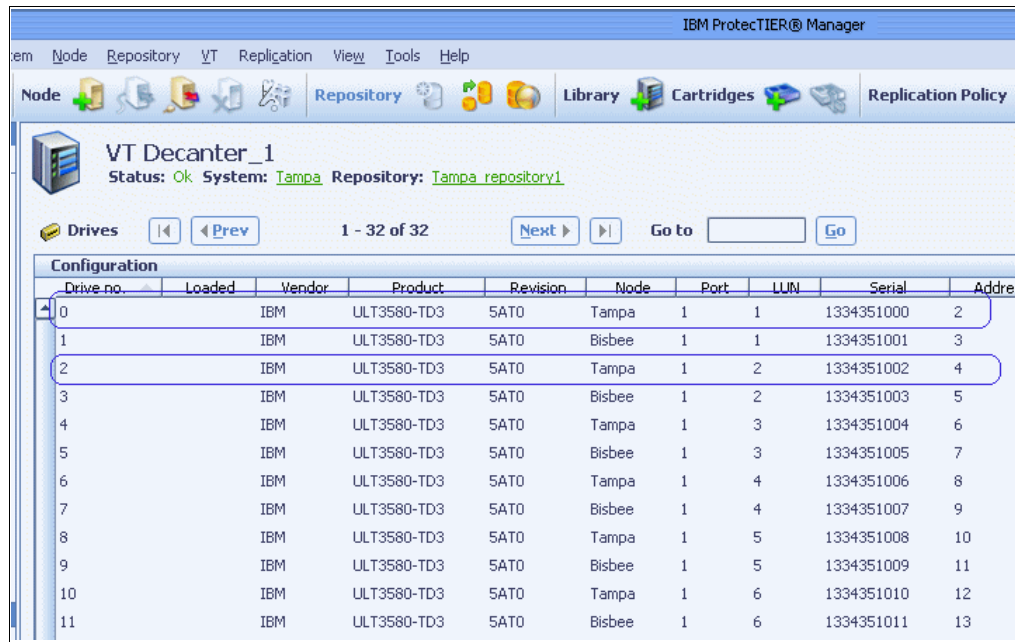


Figure 7-10 VT Decanter\_1 Drive definitions

Note the LUN assigned to each virtual drive. It is important to know that this is the ProtecTIER LUN number only, and the host operating system will almost certainly assign each drive a different LUN than that which appears here, as the host will have more logical units or resources than just these tape drives.

During the library definition, ProtecTIER assigns serial numbers to each drive, seeded from a random number. These can be important later when defining IBM Tivoli Storage Manager paths to link the host devices and the IBM Tivoli Storage Manager drive definitions. It can be helpful to know the drive serial number if you should experience any problems when defining the paths. You can use the serial number as a common reference point when matching a drive with its host device filename.

The element number is another address of each object in the tape library. Everything that can be located (such as robot;drive;slot;io slot) has its unique element number in the library. Tivoli Storage Manager (and any other tape library related software) uses this number as an address in the SCSI command to drive the robot. The element number can also be important later when defining the IBM Tivoli Storage Manager tape drives.

### Determining the IBM Tape Device Driver device names

First, we determine the IBM Tape Device Driver device names of the attached virtual devices on ProtecTIER using the `tsmdlst` utility and the `ntutil` utility by completing the following steps:

1. We open a command-line interface and change the directory to `C:\Program Files\Tivoli\TSM\console` and execute `tsmdslst`.

Example 7-12 shows the output of `tsmdslst`. In our example, for the Tivoli Storage Manager setup, we focus on Robot, Drive0, and Drive2, as previously mentioned.

Example 7-12 `tsmdslst` output

TSM Name	ID	LUN	Bus	Port	SSN	WWN	TSM Type
1b0.0.0.6	0	0	0	6	0013343519990402	10000000C97A2393	LIBRARY

```

mt0.1.0.6 0 1 0 6 1334351000 10000000C97A2393 LTO
mt0.10.0.6 0 10 0 6 1334351018 10000000C97A2393 LTO
mt0.11.0.6 0 11 0 6 1334351020 10000000C97A2393 LTO
mt0.12.0.6 0 12 0 6 1334351022 10000000C97A2393 LTO
mt0.13.0.6 0 13 0 6 1334351024 10000000C97A2393 LTO
mt0.14.0.6 0 14 0 6 1334351026 10000000C97A2393 LTO
mt0.15.0.6 0 15 0 6 1334351028 10000000C97A2393 LTO
mt0.16.0.6 0 16 0 6 1334351030 10000000C97A2393 LTO
mt0.2.0.6 0 2 0 6 1334351002 10000000C97A2393 LTO
mt0.3.0.6 0 3 0 6 1334351004 10000000C97A2393 LTO
mt0.4.0.6 0 4 0 6 1334351006 10000000C97A2393 LTO
mt0.5.0.6 0 5 0 6 1334351008 10000000C97A2393 LTO
mt0.6.0.6 0 6 0 6 1334351010 10000000C97A2393 LTO
mt0.7.0.6 0 7 0 6 1334351012 10000000C97A2393 LTO
mt0.8.0.6 0 8 0 6 1334351014 10000000C97A2393 LTO
mt0.9.0.6 0 9 0 6 1334351016 10000000C97A2393 LTO
lb0.0.0.8 0 0 0 8 0013343519990402 10000000C97A2291 LIBRARY
mt0.1.0.8 0 1 0 8 1334351001 10000000C97A2291 LTO
mt0.10.0.8 0 10 0 8 1334351019 10000000C97A2291 LTO
mt0.11.0.8 0 11 0 8 1334351021 10000000C97A2291 LTO
mt0.12.0.8 0 12 0 8 1334351023 10000000C97A2291 LTO
mt0.13.0.8 0 13 0 8 1334351025 10000000C97A2291 LTO
mt0.14.0.8 0 14 0 8 1334351027 10000000C97A2291 LTO
mt0.15.0.8 0 15 0 8 1334351029 10000000C97A2291 LTO
mt0.16.0.8 0 16 0 8 1334351031 10000000C97A2291 LTO
mt0.2.0.8 0 2 0 8 1334351003 10000000C97A2291 LTO
mt0.3.0.8 0 3 0 8 1334351005 10000000C97A2291 LTO
mt0.4.0.8 0 4 0 8 1334351007 10000000C97A2291 LTO
mt0.5.0.8 0 5 0 8 1334351009 10000000C97A2291 LTO
mt0.6.0.8 0 6 0 8 1334351011 10000000C97A2291 LTO
mt0.7.0.8 0 7 0 8 1334351013 10000000C97A2291 LTO
mt0.8.0.8 0 8 0 8 1334351015 10000000C97A2291 LTO
mt0.9.0.8 0 9 0 8 1334351017 10000000C97A2291 LTO

```

Completed in: 0 days, 0 hours, 0 minutes, 3 seconds.

C:\Program Files\Tivoli\TSM\console>

2. Start ntutil, the IBM Tape Device Driver Utility, to list all registered devices by running **ntutil**. Example 7-13 shows the main NTutil screen.

*Example 7-13 NTutil main menu*

```

Test tool version 6.2.0.2x64
Variable settings
===== LIBRARY MODE =====
tape-special-file-name: tape0, changer-special-file-name: changer0
gp->fd0=FFFFFFFFFFFFFFFF gp->fd1=FFFFFFFFFFFFFFFF block size=1024 block
count=1
hex block id = 0000000000000000
return_error_when_fail 1 exit_on_unexpected_result 0 trace_flag 0
manual test menu:
=====
1: set device special file          2: display symbols
3: set block size R/W (now !0 fixed) 4: library only mode (OFF)
5: set return error when fail       6: set/reset trace
7: set exit on unexpected result     8: Base Mode
=====
10: return library inventory        11: move medium

```

```

12: initialize element status          13: get changer parameter
14: exchange medium                   15: get changer bus info
16: get cartridge location
=====
20: open                               21: close
22: read                               23: write
24: read and display block            25: flush (buffer->media)
26: read block id                     27: erase
28: locate block id                   29: display block data
=====
30: write filemark(s)                 31: rewind
32: forward space filemark(s)         33: unload
34: reverse space filemark(s)         35: load
36: forward space record(s)           37: return error
38: reverse space record(s)           39: test unit ready
43: set media parms (block size)      44: set dev parms(compression)
46: get device information             47: restore data
48: get medium information             49: inquiry
50: poll registered devices
53: space EOD                         54: display message
59: get encryption state              60: encryption diagnostics
63: get tape bus info
=====
70: system command
=====
80: Force Dump                        81: Read Dump
82: Update MicroCode                  83: Log Sense
84: Get Last Sense                     85: Get Version
86: Associative/Persistent WProtect    87: Read/Write Test
88: List registered devices          89: Get MTDevice Info
=====
99: return to main menu

```

3. Display all registered devices using the menu function 88, as shown in Example 7-14.

*Example 7-14 Menu function 88 to list all registered devices*

```

enter selection: 88
Device found: Changer0 @"Scsi Port 6\Scsi Bus 0\Target Id 0\Logical Unit Id 0"
Device found: Tape0 @"Scsi Port 6\Scsi Bus 0\Target Id 0\Logical Unit Id 1"
Device found: Tape9 @"Scsi Port 6\Scsi Bus 0\Target Id 0\Logical Unit Id 10"
Device found: Tape10 @"Scsi Port 6\Scsi Bus 0\Target Id 0\Logical Unit Id 11"
Device found: Tape11 @"Scsi Port 6\Scsi Bus 0\Target Id 0\Logical Unit Id 12"
Device found: Tape12 @"Scsi Port 6\Scsi Bus 0\Target Id 0\Logical Unit Id 13"
Device found: Tape13 @"Scsi Port 6\Scsi Bus 0\Target Id 0\Logical Unit Id 14"
Device found: Tape14 @"Scsi Port 6\Scsi Bus 0\Target Id 0\Logical Unit Id 15"
Device found: Tape15 @"Scsi Port 6\Scsi Bus 0\Target Id 0\Logical Unit Id 16"
Device found: Tape1 @"Scsi Port 6\Scsi Bus 0\Target Id 0\Logical Unit Id 2"
Device found: Tape2 @"Scsi Port 6\Scsi Bus 0\Target Id 0\Logical Unit Id 3"
Device found: Tape3 @"Scsi Port 6\Scsi Bus 0\Target Id 0\Logical Unit Id 4"
Device found: Tape4 @"Scsi Port 6\Scsi Bus 0\Target Id 0\Logical Unit Id 5"
Device found: Tape5 @"Scsi Port 6\Scsi Bus 0\Target Id 0\Logical Unit Id 6"
Device found: Tape6 @"Scsi Port 6\Scsi Bus 0\Target Id 0\Logical Unit Id 7"
Device found: Tape7 @"Scsi Port 6\Scsi Bus 0\Target Id 0\Logical Unit Id 8"
Device found: Tape8 @"Scsi Port 6\Scsi Bus 0\Target Id 0\Logical Unit Id 9"
Device found: Changer0 @"Scsi Port 8\Scsi Bus 0\Target Id 0\Logical Unit Id 0"
Device found: Tape16 @"Scsi Port 8\Scsi Bus 0\Target Id 0\Logical Unit Id 1"
Device found: Tape25 @"Scsi Port 8\Scsi Bus 0\Target Id 0\Logical Unit Id 10"

```

```

Device found: Tape26 @"Scsi Port 8\Scsi Bus 0\Target Id 0\Logical Unit Id 11"
Device found: Tape27 @"Scsi Port 8\Scsi Bus 0\Target Id 0\Logical Unit Id 12"
Device found: Tape28 @"Scsi Port 8\Scsi Bus 0\Target Id 0\Logical Unit Id 13"
Device found: Tape29 @"Scsi Port 8\Scsi Bus 0\Target Id 0\Logical Unit Id 14"
Device found: Tape30 @"Scsi Port 8\Scsi Bus 0\Target Id 0\Logical Unit Id 15"
Device found: Tape31 @"Scsi Port 8\Scsi Bus 0\Target Id 0\Logical Unit Id 16"
Device found: Tape17 @"Scsi Port 8\Scsi Bus 0\Target Id 0\Logical Unit Id 2"
Device found: Tape18 @"Scsi Port 8\Scsi Bus 0\Target Id 0\Logical Unit Id 3"
Device found: Tape19 @"Scsi Port 8\Scsi Bus 0\Target Id 0\Logical Unit Id 4"
Device found: Tape20 @"Scsi Port 8\Scsi Bus 0\Target Id 0\Logical Unit Id 5"
Device found: Tape21 @"Scsi Port 8\Scsi Bus 0\Target Id 0\Logical Unit Id 6"
Device found: Tape22 @"Scsi Port 8\Scsi Bus 0\Target Id 0\Logical Unit Id 7"
Device found: Tape23 @"Scsi Port 8\Scsi Bus 0\Target Id 0\Logical Unit Id 8"
Device found: Tape24 @"Scsi Port 8\Scsi Bus 0\Target Id 0\Logical Unit Id 9"
Total elapsed time in seconds = 0.00
Return to continue:

```

4. To get more detailed information about the tape devices and the medium changer device, use the IBM Tape Device Driver Utility ntutil. This tool has many functions that can obtain detailed device information.
5. Complete the tape library worksheet, as shown in Table 7-2. Get the element number for each device.

Table 7-2 Library worksheet

Device in OS	Type	VTL system	VTL node	VTL port	VTL port WWN	Serial number	Element number
Changer0	3584	Tampa	Tampa	FE1	10000000C97A2393	0013343519990402	0
Tape0	LTO3	Tampa	Tampa	FE1	10000000C97A2393	1334351000	2
Tape1	LTO3	Tampa	Tampa	FE1	10000000C97A2393	1334351002	4

## Configuring the Tivoli Storage Manager

To configure Tivoli Storage Manager for backups to ProtecTIER, perform the following steps:

1. Define devclass PTclass library=PTlib devtype=lt.
2. Define library PTlib libtype=scsi shared=yes.
3. Query the defined library (Example 7-15).

*Example 7-15 Query the defined library*

```
tsm: DECANter_SERVER1>query libr PTLIB f=d
```

```

Library Name: PTLIB
Library Type: SCSI
ACS Id:
Private Category:
Scratch Category:
WORM Scratch Category:
External Manager:
RSM Media Type:
Shared: Yes

```



```

LanFree:
ObeyMountRetention:
Primary Library Manager:
    WWN:
        Serial Number: 0013343519990402
        AutoLabel: No
        Reset Drives: Yes
        Relabel Scratch: Yes
Last Update by (administrator): ADMIN
Last Update Date/Time: 09/29/2009 10:48:08

```

---

4. Define stgpool PTpool PTclass maxscratch=100.
5. Define path DECANter\_SERVER1 PTlib srctype=server desttype=library device=\\.\Changer0.
6. Define drive PTlib PTlibdr0.
7. Define drive PTlib PTlibdr1 (Example 7-16).

*Example 7-16 Query the defined drives*

---

```
tsm: DECANter_SERVER1>q drive
```

Library Name	Drive Name	Device Type	On-Line
PTLIB	PTLIBDR0	LTO	Yes
PTLIB	PTLIBDR1	LTO	Yes

```
tsm: DECANter_SERVER1>q drive PTLIB PTLIBDR0 f=d
```

```

Library Name: PTLIB
Drive Name: PTLIBDR0
Device Type: LTO
On-Line: Yes
Read Formats:
ULTRIUM3C,ULTRIUM3,ULTRIUM2C,ULTRIUM2,ULTRIUMC,ULTRI-
UM
Write Formats:
ULTRIUM3C,ULTRIUM3,ULTRIUM2C,ULTRIUM2
Element: 2
Drive State: LOADED
Volume Name: TDK000L3
Allocated to: DECANter_SERVER1
WWN: 20000000C97A2393
Serial Number: 1334351000
Last Update by (administrator): ADMIN
Last Update Date/Time: 09/29/2009 10:49:41
Cleaning Frequency (Gigabytes/ASNEEDED/NONE): NONE

```

---

8. Define path DECANter\_SERVER1 PTlibdr0 srctype=server desttype=drive library=PTlib device=\\.\Tape0.

9. Define path DECANter\_SERVER1 PTLibr1 srctype=server desttype=drive library=PTLib device=\\.\Tape1 (Example 7-17).

*Example 7-17 Query the paths*

tsm: DECANter\_SERVER1>q path

Source Name	Source Type	Destination Name	Destination Type	On-Line
DECANter_S- ERVER1	SERVER	PTLIB	LIBRARY	Yes
DECANter_S- ERVER1	SERVER	PTLIBDR0	DRIVE	Yes
DECANter_S- ERVER1	SERVER	PTLIBDR1	DRIVE	Yes

10. Label libvol PTLibr checkin=scratch search=yes overwrite=yes labelsource=barcode (Example 7-18).

*Example 7-18 Query libvol*

tsm: DECANter\_SERVER1>q libvol

Library Name	Volume Name	Status	Owner	Last Use	Home	Device Element	Type
PTLIB	TDK000L3	Private	DECANter_- SERVER1		Data	1,026	LTO
PTLIB	TDK001L3	Scratch				1,027	LTO
PTLIB	TDK002L3	Scratch				1,028	LTO
PTLIB	TDK003L3	Scratch				1,029	LTO
PTLIB	TDK004L3	Scratch				1,030	LTO
PTLIB	TDK005L3	Scratch				1,031	LTO
PTLIB	TDK006L3	Scratch				1,032	LTO
PTLIB	TDK007L3	Scratch				1,033	LTO
PTLIB	TDK008L3	Scratch				1,034	LTO
PTLIB	TDK009L3	Scratch				1,035	LTO

more... (<ENTER> to continue, 'C' to cancel)

11. Define domain PTdom.  
 12. Define policyset PTdom standard (Example 7-19).

*Example 7-19 Query policyset*

tsm: DECANter\_SERVER1>q policyset

Policy Domain Name	Policy Set Name	Default Mgmt Class Name	Description
PTDOM	ACTIVE	STANDARD	
PTDOM	STANDARD	STANDARD	
STANDARD	ACTIVE	STANDARD	Installed default policy set.
STANDARD	STANDARD	STANDARD	Installed default policy set.

13. Define mgmtclass PTdom standard standard (Example 7-20).

*Example 7-20 Query mgmtclass*

tsm: DECANter\_SERVER1>q mgmtclass

Policy Domain Name	Policy Set Name	Mgmt Class Name	Default Mgmt Class ?	Description
PTDOM	ACTIVE	STANDARD	Yes	
PTDOM	STANDARD	STANDARD	Yes	
STANDARD	ACTIVE	STANDARD	Yes	Installed default management class.
STANDARD	STANDARD	STANDARD	Yes	Installed default management class.

14. Define copygroup PTdom standard standard type=*backup* destination=PTpool.

15. Define copygroup PTdom standard standard type=*archive* destination=PTpool (Example 7-21).

*Example 7-21 Query copygroup*

tsm: DECANter\_SERVER1>q copygroup

Policy Domain Name	Policy Set Name	Mgmt Class Name	Copy Group Name	Versions Data Exists	Versions Data Deleted	Retain Extra Versions	Retain Only Version
PTDOM	ACTIVE	STANDARD	STANDARD	2	1	30	60
PTDOM	STANDARD	STANDARD	STANDARD	2	1	30	60
STANDARD	ACTIVE	STANDARD	STANDARD	2	1	30	60
STANDARD	STANDARD	STANDARD	STANDARD	2	1	30	60

16. Assign defmgmtclass PTdom standard standard..

17. Activate policysset PTdom standard (Example 7-22).

*Example 7-22 Query domain*

tsm: DECANter\_SERVER1>q domain

Session established with server DECANter\_SERVER1: Windows

Server Version 5, Release 5, Level 2.0

Server date/time: 09/30/2009 09:55:03 Last access: 09/30/2009 09:16:12

Policy Domain Name	Activated Policy Set	Activated Default Mgmt Class	Number of Registered Nodes	Description
PTDOM	STANDARD	STANDARD	1	
STANDARD	STANDARD	STANDARD	0	Installed default policy domain.

Next, you must assign the Tivoli Storage Manager Client to a Tivoli Storage Manager Policy Domain.

18. Determine your client node name. In our case, the client node name is DECANter. Run the following command:

query node

19. Change the domain of the node (Example 7-23) by running the following command:

```
update node client_node_name domain=PTdom
```

*Example 7-23 Query node*

```
tsm: DECANTER_SERVER1>q node
```

Node Name	Platform	Policy Domain Name	Days Since Last Access	Days Since Password Set	Locked?
DECANTER	WinNT	PTDOM	<1	<1	No

20. To configure the client node, you could use the Client Node Configuration Wizard GUI. It is part of the Tivoli Storage Manager Management console, and can be found in the Wizard section, as shown in Figure 7-11.

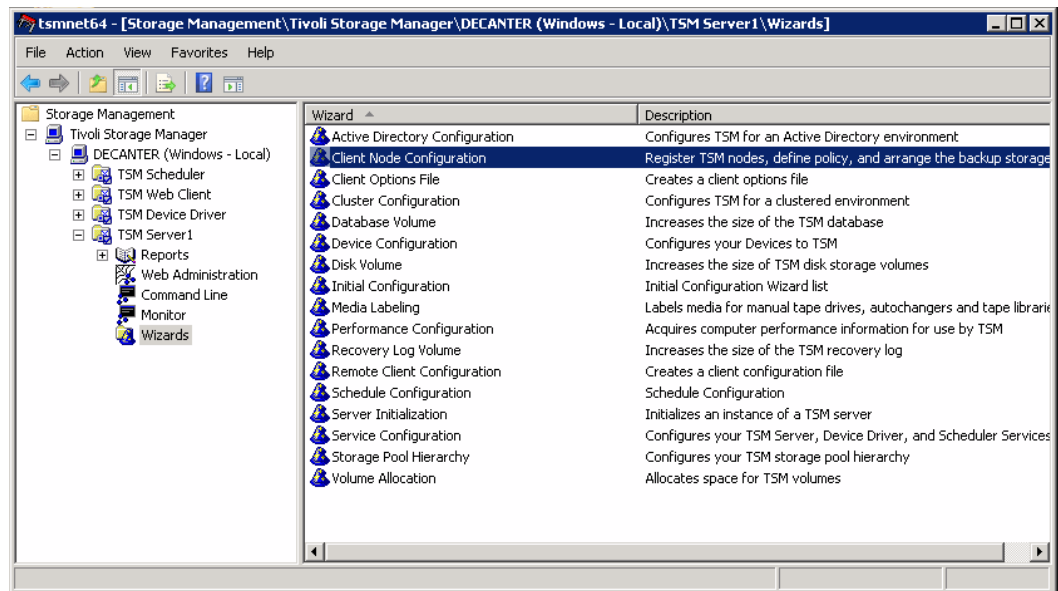


Figure 7-11 Tivoli Storage Manager Management console

The Client Node Configuration Wizard allows you to define Tivoli Storage Manager client nodes and policies, as shown in Figure 7-12.

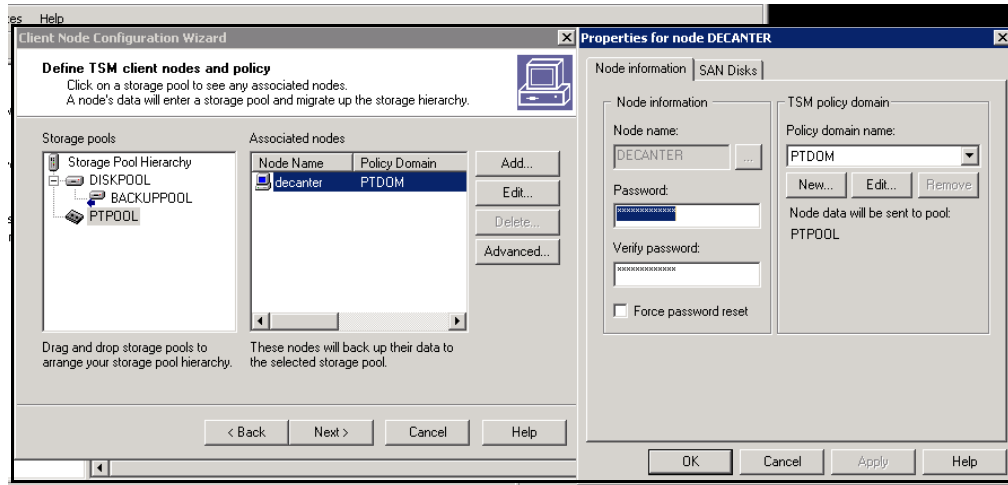


Figure 7-12 Client Node Configuration Wizard

Now you can perform a backup using the Tivoli Storage Manager Backup Archive (BA) Client, as shown in Figure 7-13.

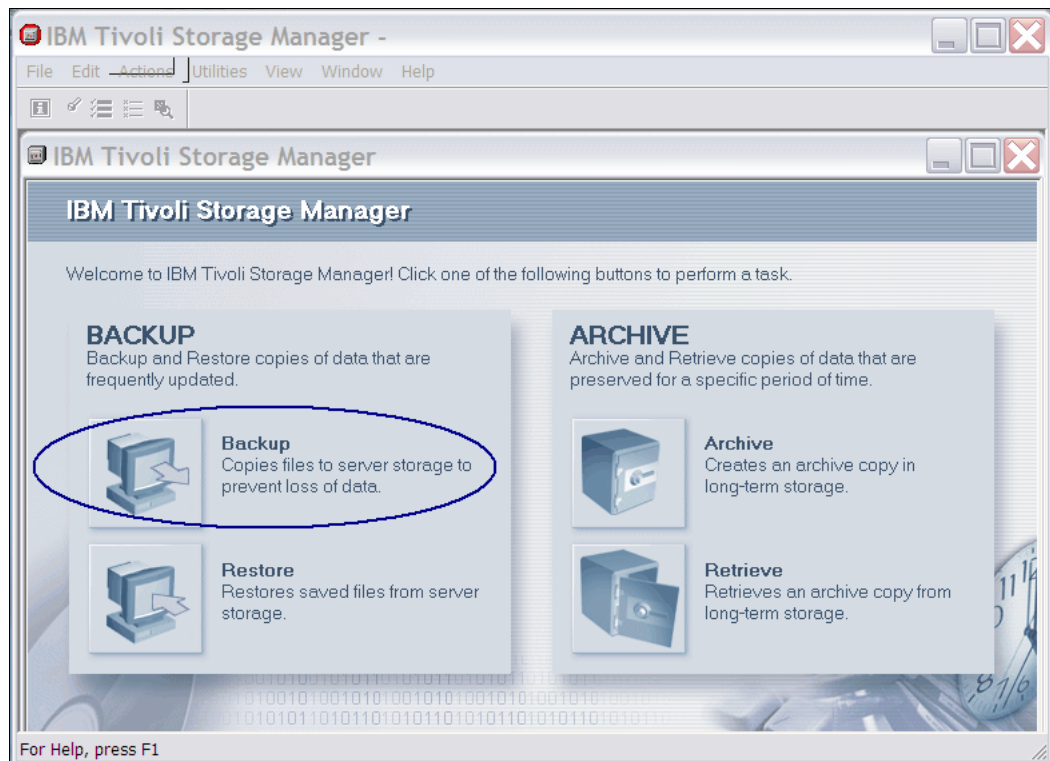


Figure 7-13 Tivoli Storage Manager BA client

21. Select the data for backup by checking the check box in front of the location. Figure 7-14 shows a total backup selection for DECANter local data.

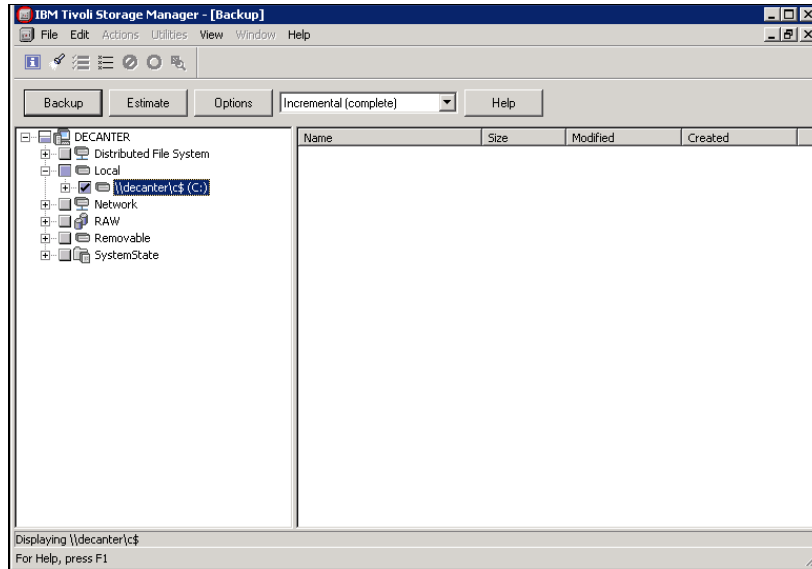


Figure 7-14 Select backup

After starting the backup, the window shown in Figure 7-15 opens. When the backup is finished, you will get a backup completed message (Figure 7-16) and a detailed status report (Figure 7-17 on page 377).

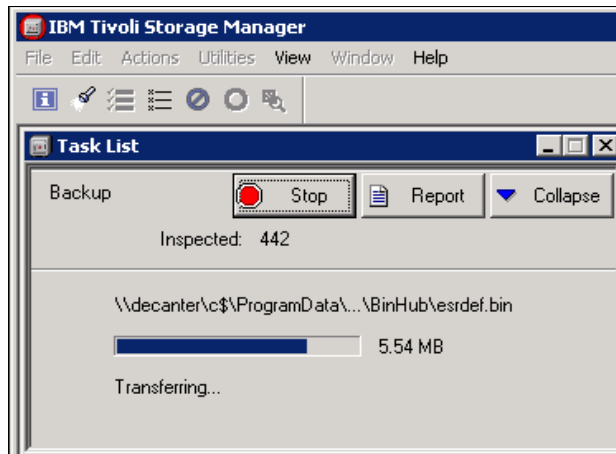


Figure 7-15 Backup data transfer

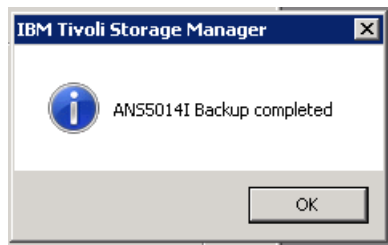


Figure 7-16 Backup completed

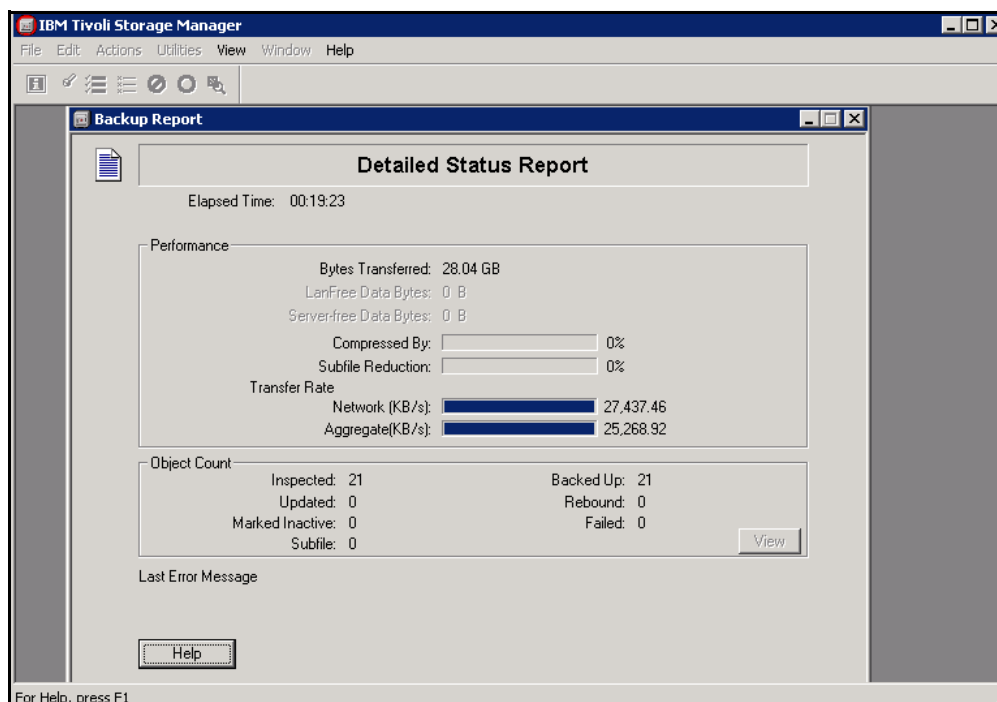


Figure 7-17 Detailed Status Report

22. Finally, we write our backup successfully to VT Decanter\_1 on cartridge TDK000L3 (Figure 7-18).

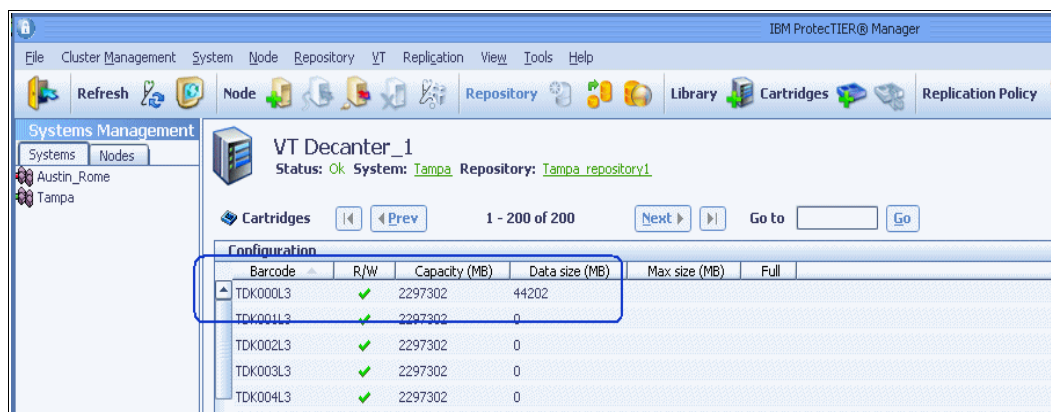


Figure 7-18 Data size on virtual data cartridge after backup

## 7.2.6 Tivoli Storage Manager database replication status

When designing a ProtecTIER IP 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, written to virtual tape in site A, to be completely replicated to and available in site B? The RPO for physical tape-based DR is typically 24 hours.

For example, in a generic user case, 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 is a 14-hour delay between the time the first evening's backup begins and when the data is safely offsite.

However, if a disaster occurs before the courier arrives (for example, a fire destroys the entire set of Tuesday's backup tapes early Wednesday morning), the customer will recover the applications from Monday's backup workload. Monday's workload is, by default, a day behind, providing a 24 hour RPO. With ProtecTIER IP Replication, it is possible to start getting the backups offsite almost immediately and replicate them at the same rate in which they were backed up with enough bandwidth. Because ProtecTIER is always working within the backup application paradigm, the RPO typically remains 24 hours. The improvements enabled by ProtecTIER are in recovery time (restoring data rapidly from disk) and reliability of the DR operation.

With Tivoli Storage Manager, nothing can be restored without a good copy of the Tivoli Storage Manager database. Therefore, the status of the virtual tapes holding the Tivoli Storage Manager database is of utmost importance.

Scenario one in Figure 7-19 illustrates a situation where the user strategy is to use ProtecTIER Replication in a scheduled time frame mode after the backup cycle is completed. This allows the data to start replicating to the offsite location immediately following the backup window's completion. In this first scenario, assuming the replication window ended at 6 a.m., the disaster strikes at 6:15 a.m., 15 minutes after the last night's replication cycle was completed. The status of all virtual cartridges is 100% completely replicated when the disaster event occurs.

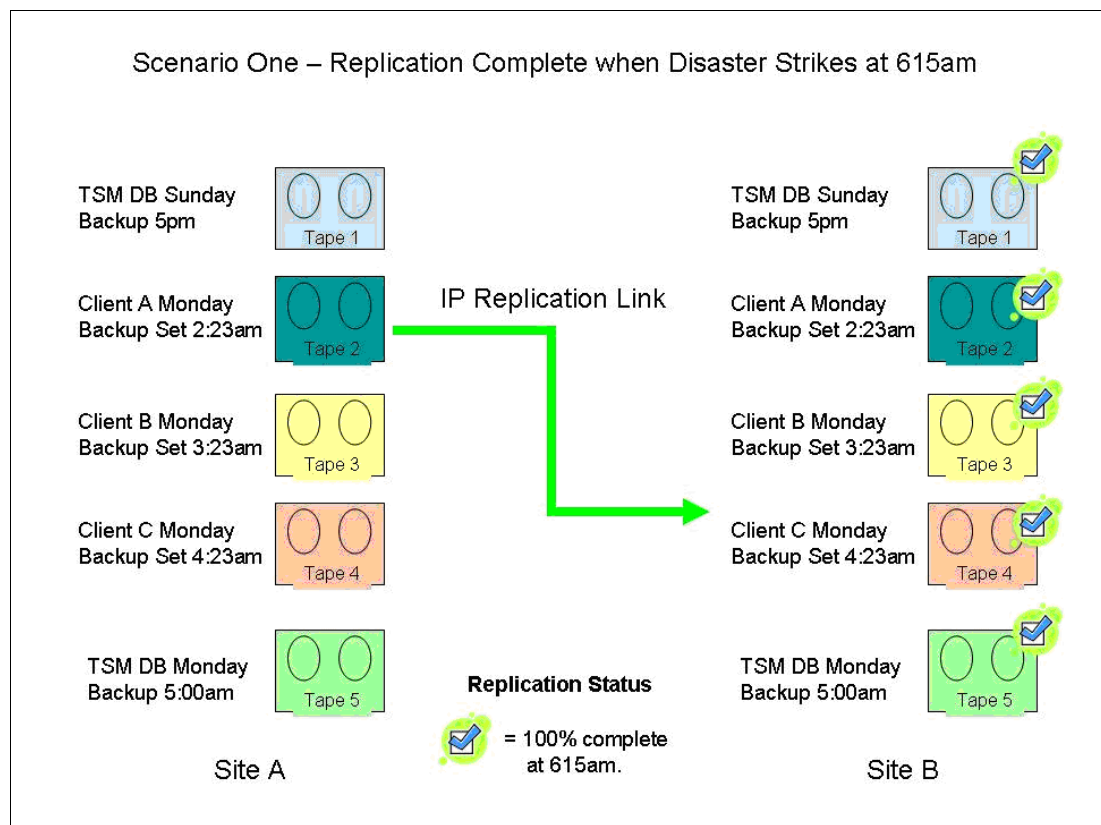


Figure 7-19 Scenario one: Replication complete

The DR restoration activities in scenario one are therefore straightforward. The user brings up the Tivoli Storage Manager server using Tivoli Storage Manager DB backup that occurred at 5 a.m. Monday on tape 5. The Tivoli Storage Manager DB has knowledge of all prior tapes, and restores can begin immediately from tapes 2 - 4, which hold the last client backups.



Scenario two, shown in Figure 7-20, illustrates the second possible disaster recovery situation. This disaster event occurs at 5:15 a.m., shortly (45 minutes) before the nightly replication cycle has completed. Thus, some of last night's backups have not yet been 100% replicated to site B. At 5:15 a.m., Monday morning, tapes 4 and 5 have not been completely replicated when the link goes down because of the disaster event. Attempts to restore the Tivoli Storage Manager database are unsuccessful, as the database was not replicated in its entirety yet. The user must therefore restore from the last previous Tivoli Storage Manager DB backup, which is on tape 1 from Sunday at 5 p.m. Because tapes 2 and 3 were created after Sunday at 5 p.m., they are not in the Tivoli Storage Manager database and cannot be used. Clients A, B, and C must therefore be restored from tapes that exist in the Tivoli Storage Manager DB as of Sunday at 5 p.m.

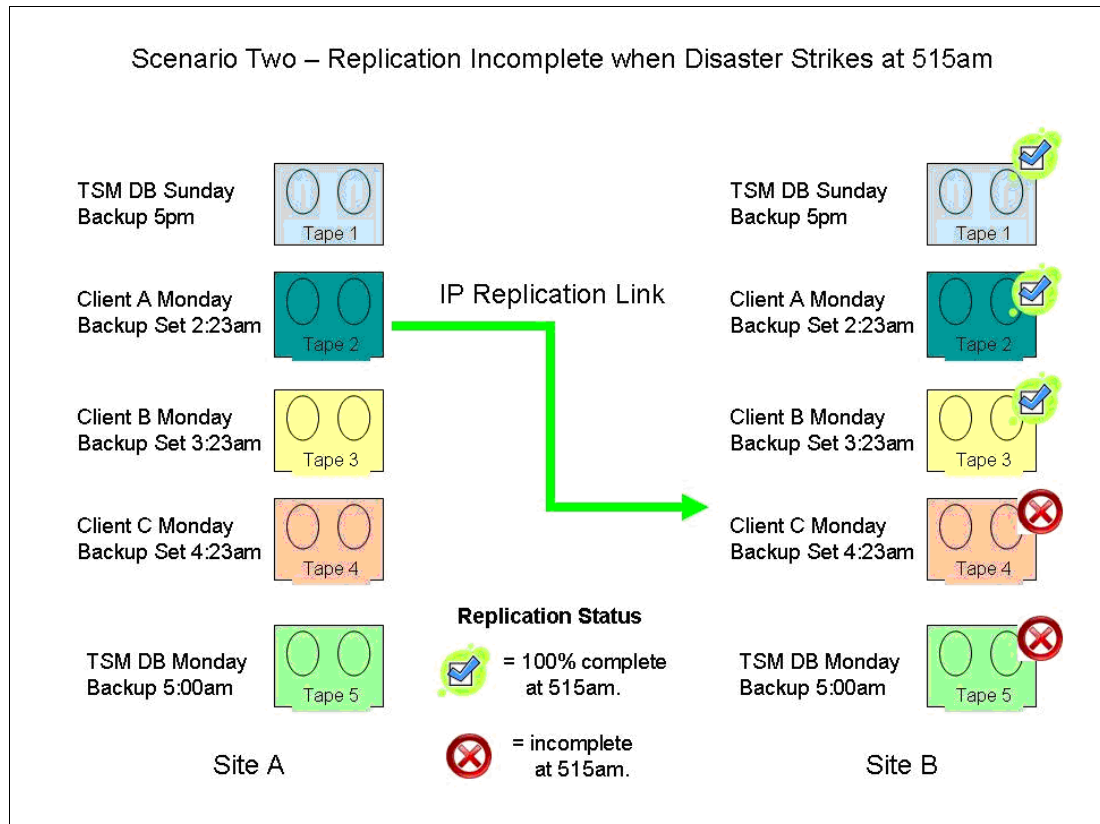


Figure 7-20 Scenario two: Replication incomplete

Scenario three, shown in Figure 7-21, illustrates another possibility where the most recent Tivoli Storage Manager database virtual tape is replicated, but not all of the associated tapes have completed replication when the disaster event occurs. As Figure 7-21 shows, the disaster strikes at 5:30 a.m. (30 minutes prior to the anticipated replication cycle completion). At this point, tapes 1, 2, 3, and 5 have replicated 100% and tape 4, due to the size of the backup dataset stored on it, has not completely finished replication when the disaster event occurs.

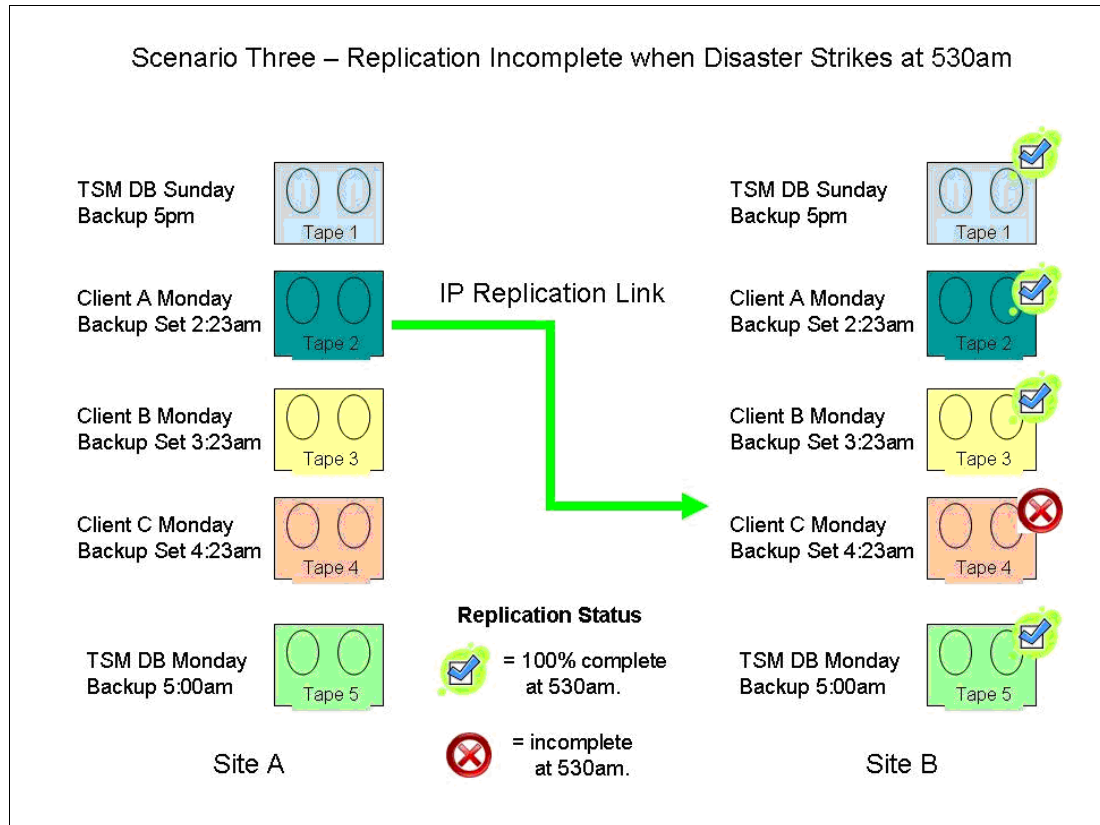


Figure 7-21 Scenario three: Replication incomplete, database complete

The Tivoli Storage Manager server is restored using the Tivoli Storage Manager database held on tape 5 (backed up at 5 a.m. Monday morning). However, because tape 4 had not completed replication when the disaster occurred, the tape must be audited and the Tivoli Storage Manager database *fixed* to represent exactly what is on the tape. To correct this error, use the following Tivoli Storage Manager command:

```
# audit volume 4 fix=yes
```

The audit/fix would need to be performed for every tape that had not been fully replicated when the disaster struck.

## 7.2.7 Reclamation considerations

Another important consideration when using ProtecTIER IP Replication in a Tivoli Storage Manager environment is the effect that reclamation will have. Reclamation is the Tivoli Storage Manager process that moves expired data off of tapes and moves any unexpired data to other tapes, thus freeing up space and returning empty tapes to the scratch pool. All reclamation tape movement is recorded in the Tivoli Storage Manager database and it must be replicated to reflect an accurate tape environment in the event of a disaster.

If the disaster strikes while reclamation is running and the database backup available for restore does not reflect the data on tape movements, an audit/fix of the volumes might be required.

The effect of reclamation can be delayed by using the REUSEDELAY parameter. When the user defines or updates a sequential access storage pool, the user can use the REUSEDELAY parameter. This parameter specifies the number of days that must elapse before a volume can be reused or returned to scratch status after all files have been expired, deleted, or moved from the volume. When the user delays reuse of such volumes and they no longer contain any files, they enter the pending state. Volumes remain in the pending state for as long as specified with the REUSEDELAY parameter for the storage pool to which the volume belongs.

Delaying reuse of volumes can be helpful under certain conditions for disaster recovery. When files are expired, deleted, or moved from a volume, they are not actually erased from the volumes. The database references to these files are removed. Thus, the file data might still exist on sequential volumes if the volumes are not immediately reused. This prevents a situation where reclamation has run after the last Tivoli Storage Manager database backup to virtual tape, and reclaimed tapes have been replicated. If the disaster occurs at a point after the reclamation ran and these tapes fully replicated, but this is not reflected in the Tivoli Storage Manager database, the user could get a mismatch.

A disaster might force the user to restore the database using a database backup that is not the most recent backup. In this case, some files might not be recoverable because the server cannot find them on current volumes. However, the files might exist on volumes that are in the pending state. The user 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 prior to file expiration.
2. Use a primary or copy storage pool volume that has not been rewritten and contains the expired file at the time of the database backup.

If the user backs up the primary storage pools, set the REUSEDELAY parameter for the primary storage pools to 0 to efficiently reuse primary scratch volumes. For their copy storage pools, the user should delay reuse of volumes for as long as they keep their oldest database backup.

Users can also disable the reclamation by changing the reclaim parameter of the **update stgpool** command:

```
update stgpool <stg pool name> reclaim=100
```

Consider disabling reclamation altogether if the user has deployed ProtecTIER with small virtual tape cartridge sizes (for example, 50 GB or less). This is because there is mathematically less to reclaim on a smaller volume than a larger one, and everything on a smaller volume is likely to expire together, requiring the tape to go immediately to the scratch pool.

## 7.2.8 Summary

In summary, consider the following actions:

- ▶ Ensure catalog/database backups are performed to virtual tape and replicated along with the daily workload each day. A database backup should be performed and replicated at the end of each backup cycle.
- ▶ A separate tape pool should be created for database backups.

- ▶ Consider adjusting Tivoli Storage Manager reclamation processing to ensure that actions are in sync with the replicated database. This includes setting the REUSEDELAY to provide a 2 day or 3 day delay in the reuse of tapes that have been reclaimed.
- ▶ Consider deploying additional network bandwidth to ensure that a replication backlog does not occur. Database or catalog synchronization becomes a greater issue if the backlog exceeds one backup cycle (typically 12 - 24 hours).

## 7.3 Determining what is available for restore at the DR site

This section suggests ways for the users to determine what catalog and data sets are complete, matched, and readily available to restore at the secondary/DR site.

### 7.3.1 Which database copy at the remote site is valid

Before running a restore for disaster recovery, you must verify that the list of associated cartridges is completely replicated to the remote 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 when you have a specific time each day where your replication backlog is zero (that is, there is no pending data to replicate).

If this is not the case, you can assess the cartridges by recovering the backup application catalog and scanning it to find the last full backup where the associated cartridges completed replication.

There are several ways to obtain a copy of the catalog at the remote site:

- ▶ From a catalog backup on a virtual cartridge that will be replicated to the remote site
- ▶ From disk-based replication, or by other means

If the catalog is backed up to a virtual cartridge, through the cartridge view of the library in ProtecTIER Manager, query each of the cartridges used for catalog backup to find the most recent sync dates marked on the cartridges. Assuming that there are multiple backup copies, you must find the latest backup that finished 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 remote site:

- ▶ Each cartridge has a last sync time that displays the last time that the cartridge's data was fully replicated to the remote site. (The sync time will be updated during the replication, not only when the replication for this cartridge is finished.)
- ▶ The cartridge marked with the most recent last sync time date should be used to recover the backup application catalog.

### 7.3.2 Determining which cartridges at the remote are valid for restore

After the DR site Tivoli Storage Manager server is recovered, you must review the status of the replicated cartridges to ensure their replication consistency with the Tivoli Storage Manager database. To achieve this goal, use the available ProtecTIER system Replicated Cartridges Status Report, as explained in 7.1.3, "Assessing cartridge status and synchronizing with the catalog" on page 347.

## Eject and inject commands from the backup application

To eject a cartridge from a library, run **checkout libvolume**. This command can be invoked through a web GUI or from the CLI. For example:

```
TSM:PUMA_SERVER1>CHECKOUT LIBVOL <name of library> REMOVE=BULK FORCE=yes  
CHECKLABEL=YES VOLLIST=<volume barcode>
```

To inject/import (insert) a cartridge into a library, run **add cartridge** from a web GUI or run **checkin libvolume** from a CLI. For example:

```
TSM:PUMA_SERVER1>CHECKIN libvol <name of library> search=BULK checklabel=barcode  
status=SCRATCH WAITTIME=60
```

Run **reply** with the ID process to finish the import operation:

```
TSM:PUMA_SERVER1>reply 2
```

### 7.3.3 Configuration changes

The following IBM Tivoli Storage Manager server and client options should be checked and, if necessary, changed to enable the optimum performance of ProtecTIER:

- ▶ Any IBM Tivoli Storage Manager client compression that will be stored in a ProtecTIER storage pool should have that compression disabled.
- ▶ Ensure that server option MOVEBATCHSIZE is set at 1000 (the default value).
- ▶ Ensure that server option MOVESIZETHRESHOLD is set at 2048 (the default value).
- ▶ When running IBM Tivoli Storage Manager on a Windows platform, the Tivoli Storage Manager driver for tape drives and libraries must be used. The native Windows drivers for the emulated P3000 and DLT7000 drives will not function.
- ▶ The only way for the VTL to realize that a volume has been deleted and its space can be reallocated is to write to the beginning of the newly returned scratch volume. The VTL will then see the volume as available. Tivoli Storage Manager can relabel volumes that have just been returned to scratch if the RELABELSCRATCH parameter is specified.

Given that ProtecTIER acts as a virtual tape library and a data deduplication device, the advantages associated with disk backup over tape backup apply here too. The following points should also be considered when using ProtecTIER with IBM Tivoli Storage Manager:

- ▶ Tivoli Storage Manager disk pools: For some large environments with several IBM Tivoli Storage Manager servers in place, you do not need to assign dedicated Tivoli Storage Manager disk storage pools to each server. With ProtecTIER, you can either share a virtual library or you can create virtual libraries for every server.
- ▶ LAN-free backups are easier: As ProtecTIER is a virtual tape library, it has the major advantage of presenting greatly increased tape resources to the backup server. This capability then positions you to be able to perform LAN-free backups to ProtecTIER without much regard for the limitations normally applied to these backups, such as tape drive availability. If you have many LAN-free clients already, then it is possible that your LAN-free backup windows were dictated not entirely by business needs but also by hardware availability. With ProtecTIER and its maximum of 256 virtual tape drives per ProtecTIER node, you can almost completely eliminate any hardware restrictions that you might have faced previously, and schedule your backups when they are required by your business needs.

- ▶ **Data streams:** You might be able to reduce your current backup window by taking full advantage of ProtecTIER's throughput performance capabilities. If tape drive availability has been a limiting factor for concurrent backup operations on your Tivoli Storage Manager server, you can define a greater number of virtual drives and reschedule backups to run at the same time to maximize the number of parallel tape operations possible on ProtecTIER systems.

**Note:** If you choose to implement this strategy, you might need to increase the value of the MAXSESSIONS option on your Tivoli Storage Manager server.

- ▶ **Reclamation:** You should continue to reclaim virtual storage pools that are resident on ProtecTIER. The thresholds for reclamation might need some adjustment for a period until the system reaches steady state (refer to "Steady state" on page 34 for an explanation of this term). When this point is reached, the fluctuating size of the virtual cartridges should stabilize and you can make a decision about what the fixed reclamation limit ought to be.
- ▶ **Number of cartridges:** This is a decision with several factors to be considered:
  - In ProtecTIER, the capacity of your repository is spread across all your defined virtual cartridges. If you define only a small number of virtual cartridges in ProtecTIER Manager, you might end up with cartridges that hold a large amount of nominal data each. While this might reduce complexity, it could also affect restore operations in that a cartridge required for a restore might be in use by a backup or housekeeping task. Preemption can resolve this issue, but it might instead be better to define extra cartridges so that your data is spread over more cartridges and drives to make the best use of your virtual tape environment.
  - **Reuse delay period for storage pool cartridges:** When deciding how many virtual cartridges to define, remember to consider using the current storage pool REUSEDELAY value. This is usually equal to the number of days that your Tivoli Storage Manager database backups are retained before they expire. The same delay period should apply to your storage pools that store data on ProtecTIER virtual cartridges and you might need to increase the number defined to ensure that you always have scratch cartridges available for backup.
  - **Collocation:** When using a virtual library, you should consider implementing collocation for your primary storage pools. If you begin a restore while 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. Using collocation means that all your data is contained on the same set of virtual cartridges. Because you do not have any of the restrictions of physical cartridges normally associated with this feature (such as media and slot consumption), you can enable the option quite safely.

Consider these points when determining how many virtual cartridges are to be created. Remember that you can always create additional virtual cartridges at any time.

- ▶ **Physical tape:** Depending on your data protection requirements, it might still be necessary to copy the deduplicated data to physical tape. This can be achieved by using standard Tivoli Storage Manager copy storage pools that have device classes directing data to physical libraries and drives.

### 7.3.4 LAN-free backup to disk with ProtecTIER

LAN-free backups add some complexity to an IBM Tivoli Storage Manager implementation. You must be careful about how the SAN-attached backup storage devices are used. Insufficient resource planning can create a storage device overload. Tape library sharing between multiple IBM Tivoli Storage Manager servers also requires proper planning. Instead of disk-sharing solutions, ProtecTIER is a *native LAN-free* backup to disk device. You can apply ProtecTIER between your backup system and disk system; no disk-sharing or file-sharing software is needed. You can set up the LAN-free environment as a real tape system. This reduces complexity and the risk to your backup environment, as shown in Figure 7-22.

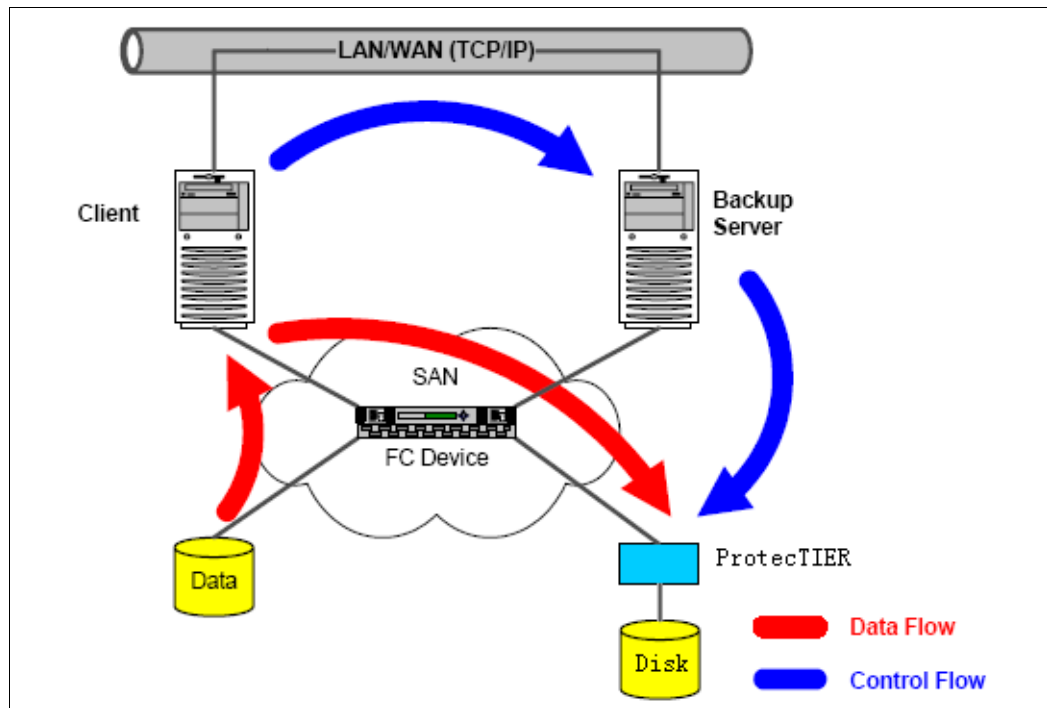


Figure 7-22 LAN-free backup to disk with ProtecTIER

Basically, the LAN-free environment consists of an IBM Tivoli Storage Manager server and client machines with both the backup archive and Tivoli Storage Manager for SAN clients installed. IBM Tivoli Storage Manager for SAN is referred to as the *storage agent*. The storage agent, sometimes described as a small IBM Tivoli Storage Manager server, is responsible for performing LAN-free operation upon client requests. Last but not least, both the server and Storage Agent need access to a shared tape library (real tape library or virtual tape library or disk system) connected to the SAN.

In addition, you need to implement the following items:

- ▶ Set up and configure the IBM Tivoli Storage Manager server for LAN-free.
- ▶ Configure devices on the client machine for LAN-free.
- ▶ Set up and configure the storage agent on a client machine.
- ▶ Customize the client for LAN-free.
- ▶ Operations and considerations when using LAN-free data transfer.

For more information about implementing Tivoli Storage Manager in SAN environments, refer to *IBM Tivoli Storage Manager Implementation Guide*, SG24-5416.

### 7.3.5 Moving data between real tape libraries and ProtecTIER systems

Sometimes you want to move data from the VTL to the real tape and move it back for DR, storage space management, or for other purposes.

In this section, we explain how to migrate data from a ProtecTIER system to the physical library and how to restore data from the physical library. To accomplish these tasks, we use a physical library connected to the IBM Tivoli Storage Manager server, as shown in Figure 7-23.




Figure 7-23 Moving data between the ProtecTIER and the real tape system

In Tivoli Storage Manager, the source and target of the data is a storage pool. You can create storage pools by using the device class of both the real tape library and the virtual tape library (such as the TS7650 and TS7650G). After storage pools are created, you can move data between them.

#### Migrating data in a storage pool

IBM Tivoli Storage Manager enables you to configure primary storage pools to provide the best combination of performance throughput and data retention. In most cases, keeping client data on tape or optical media is a requirement. However, making the backups direct to tape might not give the best performance, especially where there are many clients to back up concurrently, or many small files are being backed up.

Because of the limitations of storage media, IBM Tivoli Storage Manager provides the ability to configure storage pool hierarchies. A storage pool hierarchy allows different types of devices to be in a *chain*, with one pool overflowing to the next. A typical hierarchy consists of faster, smaller-capacity devices (disks) overflowing to slower, larger-capacity devices (tapes).



A client initially backs up to the disk storage pool. When the amount of data in the disk storage pool reaches a predefined high threshold, files are automatically migrated to the next storage pool in the hierarchy, that is, the tape storage pool. The client continues its backup operation without interruption. The migration process continues until the amount of data in the disk storage pool falls to a predefined low threshold, at which point migration stops. You can create the primary storage pool on ProtecTIER, create the next storage pool on the real tape library, and then migrate data through a schedule or manually (Figure 7-24). For more details about implementation, refer to *IBM Tivoli Storage Manager Implementation Guide*, SG24-5416.




Figure 7-24 Migrating data in the storage pool

**Note:** Migrated data from ProtecTIER to a physical tape pool will not be deduplicated.

### Backing up the storage pool

By backing up the storage pools, you can move data from one storage pool to another. The storage pool backups are usually (and should be) run by a schedule overnight. We provide examples of manually backing up the primary storage pools in this section.

In our example, the clients back up to the primary storage pool, TS7650\_LTO. We must back up to our copy storage pool, LTO\_2\_CPOOL. We run **stgpool1**, as shown in Example 7-24. Example 7-24 shows a number of logs that you can use to view the status of the jobs as they run.

#### Example 7-24 Query storage pool

```
tsm: FREEWAY>query stgpool
```

Storage Pool Name	Device Class Name	Estimated Capacity	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
ARCHIVEPOOL	DISK	2 G	0.0	0.0	90	70	LTO_2_POOL

```

BACKUPPOOL DISK          10 G   0.0  0.0  20  0  LTO_2_POOL
BIGPOOL    BIG           0.0 M  0.0  0.0  90  70
LTO_2_CPOOL LTO_2       0.0 M  0.0
LTO_2_POOL LTO_2       9,131 G  5.9  80.0  90  70
LTO_7_CPOOL LTO_7       0.0 M  0.0
LTO_7_POOL LTO_7      99,839 G  3.2  4.0  90  70
SPACEMGPOOL DISK        0.0 M  0.0  0.0  90  70
TS7650_LTO TS7650DEV   15,259 G  0.0  5.0  90  70

```

```

tsm: FREEWAY>backup stg TS7650_LTO LTO_2_CPOOL maxprocess=1 wait=no
ANS8003I Process number 2 started.

```

```

tsm: FREEWAY>query stgpool

```

Storage Pool Name	Device Class Name	Estimated Capacity	Pct Util	Pct Migr	High Mig	Low Mig	Pct	Pct
ARCHIVEPOOL	DISK	2 G	0.0	0.0	90	70		
BACKUPPOOL	DISK	10 G	0.0	0.0	20	0		
BIGPOOL	BIG		0.0 M	0.0	0.0	90	70	
LTO_2_CPOOL	LTO_2	76,294 G	0.0					
LTO_2_POOL	LTO_2	9,131 G	5.9	80.0	90	70		
LTO_7_CPOOL	LTO_7	0.0 M	0.0					
LTO_7_POOL	LTO_7	99,839 G	3.2	4.0	90	70		
SPACEMGPOOL	DISK	0.0 M	0.0	0.0	90	70		
TS7650_LTO	TS7650DEV	15,259 G	0.0	5.0	90	70		

```

tsm: FREEWAY>q act

```

Date/Time	Message
08/13/09 14:05:25	ANR2017I Administrator ADMIN issued command: BACKUP STGPOOL TS7650_LTO LTO_2_CPOOL maxprocess=1 wait=no (SESSION: 26)
08/13/09 14:05:25	ANR0984I Process 2 for BACKUP STORAGE POOL started in the BACKGROUND at 14:05:25. (SESSION: 26, PROCESS: 2)
08/13/09 14:05:25	ANR2110I BACKUP STGPOOL started as process 2. (SESSION: 26, PROCESS: 2)
08/13/09 14:05:25	ANR1210I Backup of primary storage pool TS7650_LTO to copy storage pool LTO_2_CPOOL started as process 2. (SESSION: 26, PROCESS: 2)
08/13/09 14:05:25	ANR1228I Removable volume BEL000L3 is required for storage pool backup. (SESSION: 26, PROCESS: 2)
08/13/09 14:05:25	ANR0512I Process 2 opened input volume BEL000L3. (SESSION: 26, PROCESS: 2)
08/13/09 14:05:25	ANR8337I LTO volume DFM015L3 mounted in drive DRIVE0 (/dev/rmt0). (SESSION: 26, PROCESS: 2)
08/13/09 14:05:25	ANR1340I Scratch volume DFM015L3 is now defined in storage pool LTO_2_CPOOL. (SESSION: 26, PROCESS: 2)
08/13/09 14:05:25	ANR0513I Process 2 opened output volume DFM015L3. (SESSION: 26, PROCESS: 2)
08/13/09 14:05:26	ANR1212I Backup process 2 ended for storage pool TS7650_LTO. (SESSION: 26, PROCESS: 2)
08/13/09 14:05:26	ANR0986I Process 2 for BACKUP STORAGE POOL running in the BACKGROUND processed 359 items for a total of 2,455,052 bytes with a completion state of SUCCESS at 14:05:26. (SESSION: 26, PROCESS: 2)
08/13/09 14:05:26	ANR0514I Session 26 closed volume BEL000L3. (SESSION: 26)

```
08/13/09 14:05:26 ANR0514I Session 26 closed volume DFM015L3. (SESSION: 26)
08/13/09 14:05:26 ANR1214I Backup of primary storage pool TS7650_LTO to copy
storage pool LTO_2_CPOOL has ended. Files Backed Up:
359, Bytes Backed Up: 2455052, Unreadable Files: 0,
Unreadable Bytes: 0. (SESSION: 26)
08/13/09 14:05:27 ANR8336I Verifying label of LTO volume DFM015L3 in drive
DRIVE0 (/dev/rmt0). (SESSION: 26, PROCESS: 2)
08/13/09 14:05:27 ANR8468I LTO volume DFM015L3 dismounted from drive DRIVE0
(/dev/rmt0) in library LTO_2. (SESSION: 26, PROCESS: 2)
.....
```

---

## 7.4 Symantec NetBackup

**Note:** Copying data from ProtecTIER to a tape pool will not deduplicate data on physical tape.

Symantec NetBackup (NBU) is an Open Systems Enterprise backup software solution. Its architecture has three main building blocks:

- ▶ Clients: The machines with the data that require backing up
- ▶ Media servers: The machines connected to backup devices
- ▶ Master server: The machine controlling the backups

Collectively, master, media, and clients are known as an *NBU Domain*. Typically, one master controls multiple media servers (typically 4–30) and backs up many clients (typically 10–1,000+). As the master server is the critical server in the domain, it is usually clustered and deploys other software available from the vendor (host-based volume manager for disk mirroring and cluster control through the Veritas Cluster server).

For detailed product information, visit the Symantec NetBackup websites at the following addresses:

- ▶ [http://www.symantec.com/enterprise/products/overview.jsp?pcid=1018&pvid=2\\_1](http://www.symantec.com/enterprise/products/overview.jsp?pcid=1018&pvid=2_1)
- ▶ <http://www.symantec.com/enterprise/support/index.jsp>

In a NetBackup environment, there is one master server, which holds the NetBackup database where all metadata and configurations are stored.

At least one media server is required. The media server has access to a storage unit and manages the storage unit. The master and media server 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, over LAN/IP, the backup data to a media server, 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 will be shared over several media servers, then an additional license, Shared Storage Option (SSO), is required. The TS7650G or TS7650 can eliminate or reduce the usage of SSO, because ProtecTIER can emulate many virtual tape drives, and sharing might no longer be required.

NetBackup often utilizes a traditional backup method with regular full backups and incremental or differential backups in between. NetBackup is therefore a good candidate for data deduplication.

## 7.4.1 Setting up NetBackup for backup and restore implementation

NBU 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

Most backup and restore deployments now use differential incremental backups, because they are smaller and faster. However, cumulative backups are now becoming more common in DR solutions.

From an operational standpoint, significantly more data is backed up than is ever restored. Typically, in an operational backup environment, a backup policy will back up the entire server daily. However, a typical restore might be for a single file. The product has been engineered with this in mind. There are configurations within the product to significantly improve backup performance (multiplexing, multiple data streams, fragment size, and buffer size) that have the consequence of making full restores slower. For example, if multiplexing is set to 5, then the backups from five clients will end up on the same tape. If the user is only restoring single files for a client, then the product will know which tape fragment to go to and modern high-speed tape drives with fast block locate make the process as fast as possible. In contrast, restoring an entire client involves multiple reads and tape skips so that restoring all five clients will prove time consuming, as restores are sequential.

## 7.4.2 Before you start

The implementations steps are:

1. Check the interoperability.
2. Perform the basic setup of a virtual cartridge library in ProtecTIER, as described in 5.6.1, “Creating libraries” on page 230.
3. Install the device driver, as described in *IBM Tape Device Drivers: Installation and User's Guide*, GC27-2130.
4. Check the device configuration and create a device table.
5. Implement the devices in Symantec NetBackup.

## 7.4.3 Setting up NetBackup for disaster recovery

When thinking of NBU for disaster recovery planning, the user should consider a number of key issues:

- ▶ NBU architecture: Does the NBU domain span across the primary and DR sites or are they two separate domains? This is an important consideration that has strong implications for DR.
- ▶ Classification of clients (RTO): When a company plans for DR, each server will be given a recovery time objective (RTO) depending on the importance of its application and the associated data to the business. Servers with short RTOs (typically less than 24 hours) will not use backup systems for DR. These servers will typically use clustering, volume mirroring, or some form of data replication to maintain business continuity. Servers with RTOs typically greater than 24 hours will use tape for DR. Servers will then be prioritized into RTO bands of typically 24, 36, 48, or 72 hours, depending on business requirements.

**Note:** Usually only production servers have DR setups. Test and development servers are usually out of scope for DR, although ProtecTIER makes DR protection affordable for all applications in any given environment.

- ▶ Classification of clients (RPO): In addition to RTO, there is the recovery point objective (RPO). This is the point in time to which the server must be recovered. For the majority of servers using a tape DR position, the RPO will be the point of the last complete backup before the disaster. For example, if a disaster strikes at 9:00 a.m., the RPO will be the previous night's backup.

To cater to these disaster recovery requirements, the following architectures have become common:

- ▶ Architect NBU with a single domain spanning both sites (NBU clustered): The master uses host-based replication to mirror the NBU databases and a clustering product to manage host failover. This is important, as it means that in the event of a DR event, the master's operations can seamlessly fail over to the remote site. As the NBU databases have been replicated, all of the backup information is known at the DR site and so restores can begin immediately.
- ▶ Cross-site backups: There are two main options:
  - Connect clients from one site through IP to media servers on the remote site. Backups then reside in the DR site library and are 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, because whole data is being transmitted.
  - Stretched tape SAN: The local client backs up to the local media server, which then sends the data across the SAN to the remote site. Backups then reside in the DR site library and are ready for restore. The downside is that large SAN pipes are required and backups are limited to the speed of the cross-site SAN, because whole data is being transmitted.

The downside of both options is that as normal backups are now resident in the DR library, any regular restores will be significantly slower, because data must come from a remote library.

- ▶ Multiplexing turned off: To achieve the best restore performance (to hit RTOs), NetBackup must be configured without multiplexing.
- ▶ Dedicated volume pools for RTO tiers or even clients: To achieve optimum restore times (and given sufficient media in libraries), having individual volume pools per client will achieve optimum restore performance. This way, there is no contention between media when doing restores. In the physical tape world where tape drives are limited, this is often impractical (however, it is worth noting for virtual environments). It should be stressed that this is not just a theoretical concept; systems in current production have implemented cross-site backups with client backups going to dedicated volume pools, although this was limited to 30 clients with low RTOs, because with separate volume pools, you need separate backup policies per client.

If the NBU 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 have been used by the primary site, the first operation is to get a copy of the NBU catalog from the primary site and load it into the master on the DR site. This process can either be done through disk replication or tape backup.

**Note:** NBU catalog backups are different from regular backups and need special handling to restore.

Not having the catalog available at the remote site means that every tape would have to 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.

### **Optimal ProtecTIER deployments with NBU for disaster recovery**

The following are key concepts that must be discussed with the NBU architects and senior administrators within the user's organization:

- ▶ In normal operation, back up to local VTL. This provides quick backups and quick restores.
- ▶ As ProtecTIER's VTL replication is at the cartridge level, and only the deduplicated data is being transferred over the wire, it will significantly reduce the bandwidth needed compared with traditional cross-site replication/backups.
- ▶ 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 multiplexing off for *all* backups requiring DR. Because MPX is done at either the storage unit level or backup policy level, this is easy enough. Note that it is a best practice to disable MPX for all backups going to ProtecTIER VTL.
- ▶ Use large fragment sizes, also configured at the storage unit level. This improves restore performance of entire file systems.
- ▶ Disable storage checkpoints. Storage checkpoints are a mechanism where pointers are added to the backup stream so that if the backup failed, a rerun of the backup would start from the last storage checkpoint, as opposed to from the start of the backup. Storage checkpoints have an adverse effect on the deduplication ratios.
- ▶ Disable software compression (if used), as this might reduce the efficiency of the ProtecTIER deduplication and affect its factoring ratio.

After the user's architects and the administrators have understood the basic concepts, they must apply them to their architecture, deciding whether to have one domain spanning both sites or two separate domains.

### **Single domain architecture**

In a single domain architecture, the same catalog is shared across sites and is always updated with the whereabouts of all cartridges. Within this architecture, the following items apply:

1. ProtecTIER replicates cartridges per the policies set by the user. Cartridges are copied onto a virtual shelf at the remote site ProtecTIER system.
2. Cartridges can also be *moved* using the replication policy *and* using the visibility control switch, so they will reside and be visible to the NBU application at the DR site (although the actual data will be available to ProtecTIER on both sites):
  - a. Eject (export) the cartridge from the primary library.
  - b. Inject (import) the cartridge to inventory at the remote site library.
  - c. This operation can be done by using NBU Vault or manually. Either way, it can be automated from within the NBU environment.
3. If a disaster occurs, the user either must inject the cartridges from the remote site shelf into the remote site library and inventory or, if the visibility switch control method was used, the user starts restoring or performing local backups at the DR site.

After the disaster situation has cleared and the primary site is back online, the user should use the failback procedure to move the main operation back to the primary site, including potential newly created cartridges from the DR site that will be replicated to the primary site.

### Multiple domains architecture

If the separate (multiple) domains approach is used, the following items apply:

- ▶ ProtecTIER replicates cartridges per the policies set by the user. Cartridges are copied to the virtual shelf at the remote site.
- ▶ The user should perform catalog backups to virtual tape at the end of its backup window and 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/replication cycle) the DR site will hold a full set of replicated cartridges with a matching NBU catalog, allowing for an RPO of one day.
- ▶ When disaster strikes, the user's *first* step is to get NBU catalog back on the DR site NBU server by restoring the cartridges containing the catalog.

The second step is to inject the cartridges from the remote shelf at the DR site ProtecTIER into the library and perform an inventory. After the NBU server is up and running with the DR repository, restores and local backup operations can resume at the DR site.

After the disaster situation has cleared and the primary site is back online, the user should use the failback procedure to move the main operation back to the primary site, including potential newly created cartridges from the DR site that will be replicated to the primary site.

## 7.4.4 Disaster recovery scenarios

ProtecTIER replication significantly reduces cross-site backup traffic because it only replicates deduplicated data, improves ease of operation (by enabling simple inject and inventory actions), and makes a recovery in the event of a disaster or DR test easy to plan and implement. Deploying ProtecTIER into a NBU environment makes the business more secure and removes significant obstacles for NBU architects and administrators.

The following sections provides a number of scenarios, detailing the necessary disaster recovery steps.

### Single domain environment master clustered

Two different scenarios are presented here:

- ▶ Scenario 1

All backups are complete, and all of the replication operations are complete as well. When disaster strikes, with the master clustered, the NBU catalog database at the remote site is up-to-date and no NBU recovery action is necessary.

Within ProtecTIER, the user should move tapes from the virtual shelf to import slots. Within NBU, the library must be inventoried. Remember to select the option to import tapes. After the inventory operation is complete, restores and local backups at the DR site can begin.

- ▶ Scenario 2

All backups are completed, but some or all of the replication operation is incomplete. When disaster strikes, with the master clustered, the NBU catalog database at the remote site is up to date. However, because replication was not complete, the user should roll back to the previous night's catalog and cartridges set (RPO of one day).

After the inventory operation is complete, restores and the local backups at the DR site can begin.

### Multiple domain environment master not clustered

If the multiple domain environment master is *not* clustered, the following scenarios apply:

► Scenario 1

All backups are completed, and all of the replication operations are completed as well. When disaster strikes, with the master not clustered, the catalog database at the remote site is *not* up to date, which means that the NBU catalog recovery action is necessary.

The first operation is to identify the latest backup catalog tape and load (import) it into the ProtecTIER library at the DR site. After the library has been inventoried, a standard NetBackup Catalog Recovery operation can begin. When recovery has been completed, restores and local backups at the DR site can begin.

► Scenario 2

All backups are completed, but some or all of the replication operation is incomplete. When disaster strikes, with the master not clustered, the catalog database at the remote site is *not* up to date, which means that the NBU catalog recovery action is necessary.

The first operation is to identify the previous night's NBU backup catalog tape and load (import) it into the ProtecTIER library at the DR site. After the library has been inventoried, a standard NetBackup Catalog Recovery operation of that previous night's catalog can begin. When recovery has been completed, restores (RPO of one day) and local backups at the DR site can begin.

**Note:** When working in a single-domain NBU environment (NBU master clustered) *and* using the visibility control switch option within ProtecTIER to *move* cartridges from the primary site directly into a DR site library, then the catalog is *always* up-to-date with the whereabouts of all cartridges in both the primary and DR repositories.

## 7.4.5 Determining what is available for restoration at the disaster recovery site

This section suggests ways for users to determine what catalog and data sets are complete or not complete, matched, and readily available to restore at the secondary/DR site.

### Which database copy at the remote site is valid

Before running a restore for disaster recovery, the user must verify that the list of associated cartridges is completely replicated to the remote site. Otherwise, an earlier full backup image must be used for recovery (usually the previous night's).

The easiest way to determine the time of the last full backup is if the user has a specific time each day when the replication backlog is zero (that is, there is no pending data to replicate).

If this is not the case, then the user 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.

The best practice for ensuring that a copy of the catalog is available at the remote site is to use the native replication function of ProtecTIER. Each day, the catalog should be backed up on a virtual cartridge following the daily backup workload so that it will be replicated to the remote 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 query each of the cartridges used for catalog backup to find the most recent sync dates marked on the cartridges. Assuming that there are multiple backup copies, you must find the latest backup that finished 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 remote site by performing the following actions:

- ▶ Each cartridge has a last sync time that displays the last time that the cartridge's data was fully replicated to the remote site. (The sync time is updated during the replication, and not just when the replication for this cartridge is finished.)
- ▶ The cartridge marked with the most recent last sync time date should be used to recover the backup application catalog.

### **Determine which cartridges at the remote site are valid for restore**

After the DR site NetBackup server is recovered, you must review the status of the replicated cartridges to ensure their replication consistency with NetBackup catalog. To achieve this goal, refer to 7.1.3, “Assessing cartridge status and synchronizing with the catalog” on page 347, where the ProtecTIER system Replicated Cartridges Status Report is explained in detail.

### ***Eject and inject commands from the backup application***

Although the process can be manually scripted to enable automation, the easiest way of using NBU CLI commands for automating this process is by using the vault service within the NBU software.

### ***How to eject a cartridge from a library***

You can eject a cartridge from a library by using the wizard command in the NBU GUI or by using the vault option.

If you are using the `vault` command, first run your vault policy by running the following command:

```
>/usr/opensv/netbackup/bin/vltrun <vault policy name>
```

At the end of the backup, eject the cartridge by running the following command:

```
>/usr/opensv/netbackup/bin/vltinject <vault policy name>
```

### ***How to insert a cartridge to a library***

To inject a cartridge, you only need to perform a simple robot inventory, which can be performed by selecting **Import from the load port** in the NBU GUI.

To automate this process, you can use CLI commands. For example, to update the Media Manager volume, run the following command:

```
>/usr/opensv/volmgr/bin/vmupdate -rt dlt -rn
```

## **7.4.6 NBU procedure to recover a master server from an existing DB copy**

You have two options for recreating the NBU backup catalog, that is, online and offline:

- ▶ The hot online catalog backup procedure can be found in the “Online, hot catalog backup method” section of the NBU Help menu.
- ▶ The offline, cold catalog backup procedure can be found the “Offline, cold catalog backup method” section of the NBU Help menu.

## 7.4.7 Configuration changes

The following configuration options in Symantec NetBackup should be checked and, if necessary, changed to provide the optimum performance of ProtecTIER:

- ▶ Ensure that multiplexing is turned off.
- ▶ The NUMBER\_DATA\_BUFFER file should be set to at least 32 and the SZ\_DATA\_BUFFER file should be set to at least 262144.

On an AIX system, these buffers can be configured by creating the files on the NetBackup media server by running the following commands:

- /usr/opensv/netbackup/db/config/SIZE\_DATA\_BUFFERS
- /usr/opensv/netbackup/db/config/NUMBER\_DATA\_BUFFERS

- ▶ Client compression should be disabled.
- ▶ Tape encryption (for images going to ProtecTIER) should be disabled.
- ▶ When initially creating the virtual library in ProtecTIER Manager, select V-TS3500 as the library type. This is required by Symantec for NetBackup support.
- ▶ When using Windows master or media servers, consider using NetBackup device drivers over native Windows drivers, as per Symantec recommendations.

## 7.5 EMC NetWorker

EMC NetWorker, formerly Legato NetWorker, is a centralized, automated backup and recovery product for heterogeneous enterprise data. The NetWorker Server runs on all major operating systems, such as AIX, Linux, Windows, SUN Solaris, and HP-UX.

The NetWorker Storage Node, which is a kind of LAN-free client with proxy node capability, runs on all major operating systems. The proxy node capability of the storage node can receive data from other NetWorker clients over the LAN and store the data directly to the storage device. Only the metadata will be handled by the NetWorker Server.

The NetWorker Client sends the backup data either to the NetWorker Server or to a NetWorker Storage Node. There are different clients available for the integration of special applications, such as NetWorker for IBM DB2.

A NetWorker Domain consists of one NetWorker Server and several NetWorker Clients. Several NetWorker Storage Nodes can exist in one NetWorker Domain. There is no data exchange or storage resource sharing outside one NetWorker Domain. In addition, if tape drives must be shared between one or more storage nodes and the NetWorker Server, additional licenses are required. Therefore, ProtecTIER might be a great solution for sharing physical tape resources.

## 7.5.1 EMC Legato NetWorker and ProtecTIER Replication in a LAN/WAN

Figure 7-25 shows a two-site disaster recovery configuration that also supports high availability. One of the ProtecTIER single-node or dual-node clusters is located at a production site and one ProtecTIER single-node or dual-node cluster is located at a remote, disaster recovery site. The ProtecTIER systems are connected through a WAN. Backup workloads are written to the ProtecTIER system at the production site and then are replicated to the ProtecTIER system at the disaster recovery site.

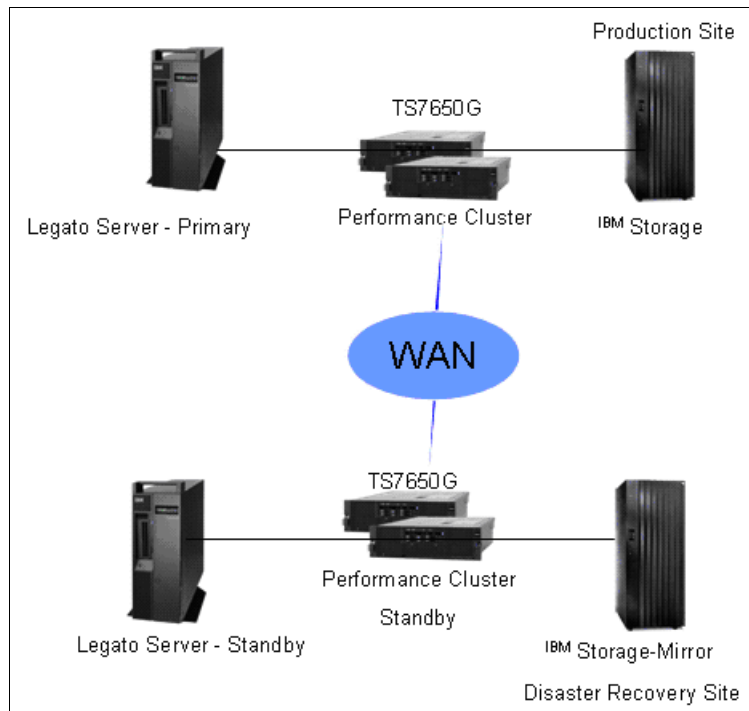


Figure 7-25 Two-site disaster recovery using IP replication

Asynchronous data replication is supported between the production environment and the disaster recovery site. After synchronization, only unique data *elements*, such as changes to index, metadata, and new data from Legato Backup Server at the production site, are replicated to the disaster recovery site. As a result, LAN or WAN bandwidth requirements are significantly reduced.

## 7.5.2 Implementation steps

Here is an overview of the different implementations steps:

1. Check the interoperability.
2. Perform the basic setup of a virtual cartridge library in ProtecTIER, as described in 5.6.1, “Creating libraries” on page 230.
3. Install the device driver, as described in *IBM Tape Device Drivers: Installation and User’s Guide*, GC27-2130.
4. Check the device configuration and create a device table.
5. Implement the devices in EMC NetWorker.
6. Perform additional setup and labeling.

Special considerations include:

- ▶ Use the 512 KB I/O block size for the virtual tape drives, which provides the best factoring ratio.
- ▶ If possible, use NetWorker Version 7.3 or later, which allows multiplexing to be completely disabled.
- ▶ Disable CDI on all virtual tape drives.
- ▶ Client compression should be disabled.
- ▶ Set parallelism to 1 on all virtual tape drives.

### 7.5.3 Replicating NetWorker database (bootstrap) backups

The EMC Legato NetWorker bootstrap contains the media database, the resource database, and the NetWorker's Client file index, and is critical to the recovery of the NetWorker Server. In most cases, the bootstrap's backups are usually written or saved to a local-attached physical tape drive connected to the NetWorker server, and a best practice from EMC is to have the bootstrap backups from NetWorker done at least two to three times a day, with one backup sent offsite for disaster recovery and the remaining copies kept onsite.

With the integration of the ProtecTIER replication, bootstrap backups can be sent to the ProtecTIER system. When bootstrap backup is done to the ProtecTIER system, a second copy is done immediately through the IP replication link to the disaster recovery site. This practice eases the system administration management of handling physical tapes for bootstrap backups.

### 7.5.4 Legato disaster recover procedures at the DR site

Should the production site become unavailable, user operations can be restarted by using the replication data from the production site and using the following manual procedures.

During disaster recovery, the NetWorker database and configuration files must be made available at the disaster recovery site during the disaster recovery procedures.

Legato's DR procedure, which use disaster recovery with ProtecTIER, requires the following high-level steps:

1. Ensure that the Legato standby server at the DR site has the exact same OS and patch level as the Legato server at the production site.
2. Ensure that the Legato standby server OS, which detects the ProtecTIER virtual tape library and tape drives, is available.
3. Ensure that the original directory location to which NetWorker server was installed is available.
4. Install the same version of the NetWorker Server software into its original location.
5. Ensure that the backup or clone volumes containing the NetWorker Server bootstrap and indexes are available. See 7.3, "Determining what is available for restore at the DR site" on page 382 to determine which cartridges at the remote site are valid for restore.
6. Ensure that the names of any links to NetWorker directories (that is, /nsr to /var/nsr) are available.
7. Reinstall any NetWorker patches that were installed prior to the disaster.
8. Configure the new NetWorker Server so that you can view the ProtecTIER virtual tape library and tape drives. Run **jbconfig** to accomplish this task.

9. From the ProtecTIER Replication Manager, move the cartridges from the virtual shelf to the I/O slots of the selected VTL library. For example, to insert a cartridge with bootstrap data, use the following command:

```
nsrib -d <volume name>
```

To eject a cartridge from the library, use the following command:

```
nrrib -w <volume name>
```

Full information about these commands can be found in the *EMC Legato NetWorker Commands Reference* document, found at the following address:

[ftp://ftp.legato.com/pub/infodev/.../networker\\_7.1/command.pdf](ftp://ftp.legato.com/pub/infodev/.../networker_7.1/command.pdf)

10. Inventory the selected Virtual Tape Library by running **nsrjb -i**. This action helps you determine whether the volumes required to recover the bootstrap are located inside the ProtecTIER virtual tape library.
11. If the user has replicated the bootstrap volumes from production to the disaster recovery site, use the following command:

```
nsrjb -lnv -S slot -f device name
```

Where *slot* is the slot where the first volume is located and *device\_name* is the path name for the first drive. You can obtain the *device\_name* by running **inquire**.

12. Run **scanner -B** to determine the save set ID of the most recent bootstrap on the media. For example, on Linux:

```
scanner -B /dev/nst0
```

**Note:** If you do not locate the save set ID of the most recent bootstrap on the most recent media, run **scanner -B** on the preceding media to locate the save set ID of the most recent bootstrap.

13. Record both the bootstrap save set ID and the volume label from the output.
14. Run **mmrecov** to recover the NetWorker server's bootstrap (media database and resource database).
15. Restart the NetWorker services on the Legato server.
16. Restore the indexes of the clients (only if needed).
17. Perform test backup and recovery with the standby ProtecTIER virtual tape library.
18. Begin normal operations at the DR site.

### 7.5.5 Determining which cartridges at the remote site are valid for restore

After the Legato server at the DR site is recovered with the required catalog, the user must review the status of the replicated cartridges to ensure their replication consistency with the Legato backup application catalog. For the process needed to achieve this goal, refer to 7.3, "Determining what is available for restore at the DR site" on page 382.

## 7.6 Backup, Recovery, and Media Services

The utilization of ProtecTIER in IBM i environments opens new possibilities for IBM System i® users to implement and deploy their backup solutions. With ProtecTIER, tape devices and media are virtual and reside on the disk space controlled by the system. The require space for backups is minimized by ProtecTIER's deduplication. ProtecTIER provides flexibility and easy management to users with complex backup environments, because they can define as many tape devices and cartridges as necessary, and they can share the ProtecTIER system amongst System i and open servers. The IT centers that keep their copies of cartridges at a remote DR site can employ the ProtecTIER replication function, which we discuss in detail in 7.6.1, “Advantages of ProtecTIER replication for an IBM i data center” on page 400.

### 7.6.1 Advantages of ProtecTIER replication for an IBM i data center

Many System i centers recognize the need to save data to a remotely connected tape library in addition to performing local backups. Users want to be able to restore from remote tapes directly to their local system in case the local tapes are damaged. During a more serious disaster at the primary site, users want the capability to restore data from this remote tape library by using a separate System i system at the remote site.

Some System i IT centers use Backup, Recovery, and Media Services (BRMS) to save to remotely attached tape drives, or to copy their backups from the local tape drive to the remote drive using BRMS functionality. However, this method of remote save has its drawbacks:

- ▶ The maximum distance for connecting a remote tape library to an IBM i system is 100 km (approximately 62 miles). Therefore, remote backups to such DR centers are limited by the distance between the sites.
- ▶ This way of connecting the tape library requires a Fibre Channel (FC) connection. If the distance exceeds 10 km, the customers must employ FC over IP with Converters, which adds latency to the save rate, thus significantly prolonging the duration of the backup operation.
- ▶ Both methods (saving to remote tape library and copying local tapes to remote ones) use the System i resources, such as main memory and CPU power, thus putting more load and constraints on the local backup operation.

Using ProtecTIER replication eliminates all these drawbacks:

- ▶ ProtecTIER replication is IP based (TCP protocol), so practically speaking there is no distance restriction.
- ▶ Replication is done directly over an IP network without any need for additional conversion devices.
- ▶ Copying cartridges/virtual tapes to the remote site is done exclusively within the ProtecTIER system so that it does not affect any host System i server resources.

In addition, ProtecTIER provides deduplication technology that drastically reduces the length of saves. Consequently, the required replication bandwidth is significantly lower compared to remote save or copying without deduplication.

## 7.6.2 Setting up BRMS for backups to ProtecTIER

We use BRMS to perform saves and restores of a database library to a ProtecTIER system.

**Note:** The BRMS implementation of virtual devices and media on ProtecTIER is no different from the BRMS setup of physical devices and media.

Each backup scenario with BRMS and ProtecTIER must be carefully discussed in advance to identify and handle possible issues, such as identical volume serial numbers of replicated cartridges, using BRMS locations, and so on.

BRMS can be accessed by:

- ▶ IBM i commands (text terminal commands and menus)
- ▶ The IBM i Navigator graphical interface
- ▶ The web interface of IBM Systems Director Navigator for IBM i

For more information about IBM i Navigator and IBM Systems Director Navigator for IBM i, refer to the IBM i Information Center at the following address:

<http://publib.boulder.ibm.com/infocenter/iserics/v6r1m0/index.jsp>

In our test, we use IBM i commands and text terminal menus for BRMS. To set up BRMS for backups to a ProtecTIER system, complete the following steps:

1. Initialize BRMS with the \*DEVICE option to recognize the new devices in the system. Run INZBRM \*DEVICE to perform this action.
2. Run GO BRMS to display the BRMS main menu.
3. In the BRMS main menu, select Option 1 = Media management and then Option 1 = Work with media classes to check whether the necessary media class is already defined in BRMS. If the needed cartridge density is not listed in the BRMS media classes, add it by selecting Option 1= Add. In our example, we need the media class ULTRIUM3, which is already defined in BRMS, as shown in Figure 7-26.

Work with Media Classes		15PF8	
Position to . . . . .	Starting characters		
Type options, press Enter.			
1=Add 2=Change 3=Copy 4=Remove 5=Display 6=Work with media			
Opt	Class	Density	Capacity Text
	FMT3592	*FMT3592	*DENSITY Entry created by BRM configuration
	FMT3592A1	*FMT3592A1	*DENSITY Entry created by BRM configuration
	FMT3592A1E	FMT3592A1E	*DENSITY Entry created by BRM configuration
	FMT3592A2	*FMT3592A2	*DENSITY Entry created by BRM configuration
	FMT3592A2E	FMT3592A2E	*DENSITY Entry created by BRM configuration
	FMT3592P	*FMT3592P	*DENSITY Entry created by BRM configuration
	SAVSYS		
	ULTRIUM1	*ULTRIUM1	*DENSITY Entry created by BRM configuration
	ULTRIUM2	*ULTRIUM2	*DENSITY Entry created by BRM configuration
	ULTRIUM3	*ULTRIUM3	*DENSITY Entry created by BRM configuration
	VRTUDF	*VRTUDF	*DENSITY Entry created by BRM configuration
More...			
F3=Exit F5=Refresh F12=Cancel			

Figure 7-26 Work with Media Classes

Refer to the IBM Software Technical Document to clarify the media densities and classes for a particular cartridge. You can find this document at the following address:

[http://techsupport.rchland.ibm.com/s\\_dir/slkbase.nsf/1ac66549a21402188625680b0002037e/b69de8ff8246ff2c8625688700566393?OpenDocument](http://techsupport.rchland.ibm.com/s_dir/slkbase.nsf/1ac66549a21402188625680b0002037e/b69de8ff8246ff2c8625688700566393?OpenDocument)

4. Enroll the media into BRMS by selecting Option 1 = Media Management and then Option 9 = Work with media libraries. Specify Option 8 = Work with MLB media at the tape library that contains the cartridges to be enrolled. The list of virtual cartridges in the tape library is displayed. Select Option 1 = Add MLN media to enroll the media to BRMS. See Figure 7-27.

```

Work with Media Library Media
ISPFE
Media library . . . . . : TAPMLB15
Position to . . . . .      Starting characters

Type options, press Enter.
 1=Add MLB media   2=Work with media   5=Initialize
 6=Change category 7=Eject   8=Mount   9=Demount

Media
Opt Volume  Type      Category  --BRM Information--  Media Class  Status
1 I00000 L3 *INSERT  Inserted              *NONE
1 I00001 L3 *INSERT  Inserted              *NONE
1 I00002 L3 *INSERT  Inserted              *NONE
1 I00003 L3 *INSERT  Inserted              *NONE
1 I00004 L3 *INSERT  Inserted              *NONE
1 I00005 L3 *INSERT  Inserted              *NONE
1 I00006 L3 *INSERT  Inserted              *NONE
1 I00007 L3 *INSERT  Inserted              *NONE
More...

F3=Exit F5=Refresh F9=System command F12=Cancel F13=Repeat

```

Figure 7-27 Work with Media Library Media

When adding a cartridge to BRMS, make sure that the media being added is in the correct media class and that the expiration date is set to \*NONE, which enables cartridges to be immediately available for use. See Figure 7-28.



```

Add Media Library Media to BRM (ADDMLMBRM)

Type choices, press Enter.

Media library .....> TAPMLB15   Name
Volume identifier .....> 'I00000

Add volume to BRM .....> *YES      *YES, *NO
Initialize media .....> *NO        *NO, *YES
Media class .....> ULTRIUM3      FMT3592, FMT3592A1, FMT3592...
Last moved date .....> *NONE      Date, *NONE
Move policy .....> *NONE          *NONE, OFFSITE
Expiration date .....> *NONE      Date, *PERM, *NONE

```

Figure 7-28 Add Media Library Media to BRM (ADDMLMBRM)

**Note:** In our testing, we initialized the media later (see step 9 on page 407). However, you can initialize the cartridge while adding it to the BRMS by specifying “Initialize media. \*YES on ADDMLMBRM”. The value “Initialize media” is shown in Figure 7-28.

The newly added cartridges now appear in the default category \*SHARE400 as inserted and expired (Figure 7-29).

```

Work with Media Library Media                                     I5PF8

Media library .....: TAPMLB15
Position to .....      Starting characters

Type options, press Enter.
 1=Add MLB media   2=Work with media   5=Initialize
 6=Change category 7=Eject   8=Mount   9=Demount

Media                                     --BRM Information--
Opt Volume Type Category Status          Media Class Status
I00008 L3 *SHARE400 Inserted             ULTRIUM3 *EXP
I00009 L3 *SHARE400 Inserted             ULTRIUM3 *EXP
I00010 L3 *SHARE400 Inserted             ULTRIUM3 *EXP
I00011 L3 *SHARE400 Inserted             ULTRIUM3 *EXP
I00012 L3 *SHARE400 Inserted             ULTRIUM3 *EXP
I00013 L3 *SHARE400 Inserted             ULTRIUM3 *EXP
I00014 L3 *SHARE400 Inserted             ULTRIUM3 *EXP
I00015 L3 *SHARE400 Inserted             ULTRIUM3 *EXP

Bottom

```

Figure 7-29 Work with Media Library Media

5. Create a new media policy for the virtual cartridges in ProtecTIER VTL. Select Option 11 = Policy administration from the BRMS main menu and then Option 7 = Work with media policies. Next, select Option 1 = Create add a new media policy for the tape library and media class, as shown in Figure 7-30.

```

                                Create Media Policy

Type choices, press Enter.

Media policy ..... FULLULT3   Name
Retention type ..... 2         1=Date, 2=Days,
                               3 Versions, 4=Permanent
Retain media ..... 35         Date, Number
Deleted library retention... *NONE   Number, *NONE
Move policy ..... *NONE       Name, *NONE, *ADSM, F4
Media class ..... ULTRIUM3    Name, *SYSPCY, *ADSM, F4
Storage location ..... TAPMLB15 Name, *ANY, F4 for list
Save to save file ..... *NO    *YES, *NO
ASP for save files ..... *SYSTEM Name, *SYSTEM, 1-32
Save file retention type ... 4    1=Date, 2=Days,
                               3=Permanent, 4=None
Retain save files ..... *NONE   Date, Number, *NONE
ASP storage limit ..... *SYS    *SYS, 1-99
Secure media ..... *NO        *YES, *NO, *ADSM

                                More...
F3=Exit F4=Prompt F5=Refresh F12=Cancel

```

Figure 7-30 Create Media Policy

6. Change the BRMS system policy to point to the new media policy and the relevant tape library by selecting Option 11 = Policy administration from the BRMS main menu. Next, select Option 1 = System policy and Option 1 = Display or Change system policy.

Insert the needed values at the Change System Policy panel, as shown in Figure 7-31.

```

V6R1M0                                Change System Policy                                15PFE

Type choices, press Enter.

Media policy ..... FULLULT3   Name, F4 for list
Devices ..... TAPMLB15      Name, F4 for list

Home location for media ..... TAPMLB15   Name, F4 for list
Media class ..... ULTRIUM3    Name, F4 for list
Sign off interactive users ..... *NO     *YES, *NO
Sign off limit ..... 30       0-999 minutes
Output queue ..... *PRTF     Name, *PRTF
Library .....                Name, *LIBL
Day start time ..... 0:00:00   Time
Media monitor ..... *YES      *YES, *NO
Shared inventory delay ..... 60  30-9999 seconds
Auto enroll media ..... *NO    *NO, *YES
Default usage ..... *YES      *NO, *YES

More...
F3=Exit F4=Prompt F5=Refresh F12=Cancel

```

Figure 7-31 Change System Policy

- Change the backup policy to use the created media class and tape library. Select Option 2 = Backup policy from the BRMS Policy Administration menu and then Option 1 = Display or change backup policy.

In the new backup policy, specify the type of backup for every day. Figure 7-32 shows changing the backup policy in our test environment.

```

Change Backup Policy

ISPFE

Type choices, press Enter.

Media policy for full backups . . . . . FULLULT3 Name, F4 for list
Media policy for
  incremental backups . . . . . FULLULT3 Name, F4 for list
Backup devices . . . . . TAPMLB15 Name, F4 for list

Default weekly activity . . . . . FFFFFF MTWTFSS(F/I)
Incremental type . . . . . *CUML *CUML, *INCR
Force full backup days . . . . . *NOMAX 0-365, *NOMAX
Sign off interactive users . . . . . *SYSPCY *YES, *NO, *SYSPCY
Sign off limit . . . . . *SYSPCY 0-999 minutes, *SYSPCY
Save journaled objects when
  saving changed objects . . . . . *NO *YES, *NO

More...
F3=Exit F4=Prompt F5=Refresh F9=System policy

```

Figure 7-32 Change Backup Policy

8. Create a backup control group to save to ProtecTIER. Select Option 2 = Backup, Option 1 = Backup planning and Option 2 = Work with backup control groups starting in the BRMS main menu. In the Work with Backup Control Groups menu, select Option 1 = Create and specify the name of new control group. In the next menu, specify the sequence number, the database library to save, and the save type (full or incremental) for each day. If you must add more items to the control group, press Enter, and an additional empty line is displayed. To finish, press F3 = Exit and then select Option 1 = Save and exit session. In our testing, we create control group PROTECTIER, where we perform the full save of the library QDEXDATA01 every day, as shown in Figure 7-33.

```

Create Backup Control Group Entries                                1SPFE

Group ..... : PROTECTIER
Default activity .... *BKUPCY
Text ..... *NONE

Type information, press Enter.

Backup List ASP Weekly Retain Save SWA
Seq Items Type Device Activity Object While Message Sync
          MTWTFSS Detail Active Queue ID

10 QDEXDATA01 *SYSBAS FFFFFFFF *NONE

F3=Exit F5=Refresh F10=Change item F11=Display exits
F12=Cancel F14=Display client omit status F24=More keys

Bottom

```

Figure 7-33 Create Backup Control Group Entries

- Initialize cartridges in the VTL for backups. Select Option 1 = Media management and Option 9 = Work with media libraries at the BRMS main menu. Select Option 9 = Work with media libraries at the relevant tape library. This action displays the list of cartridges in the library. Select Option 5 = Initialize to initialize the cartridges (Figure 7-34).

```

Work with Media Library Media
ISPFE
Media library . . . . . : TAPMLB15
Position to . . . . .   Starting characters

Type options, press Enter.
1=Add MLB media   2=Work with media   5=Initialize
6=Change category 7=Eject   8=Mount   9=Demount

Media                               --BRM Information--
Opt Volume Type Category Status      Media Class Status
5 I00000 L3 *SHARE400 Available      ULTRIUM3 *EXP
  I00001 L3 *SHARE400 Available      ULTRIUM3 *EXP
  I00002 L3 *SHARE400 Available      ULTRIUM3 *EXP
  I00003 L3 *SHARE400 Available      ULTRIUM3 *EXP
  I00004 L3 *SHARE400 Available      ULTRIUM3 *EXP
  I00005 L3 *SHARE400 Available      ULTRIUM3 *EXP
  I00006 L3 *SHARE400 Available      ULTRIUM3 *EXP
  I00007 L3 *SHARE400 Available      ULTRIUM3 *EXP

More...

```

Figure 7-34 Work with Media Library Media

At initialization, the correct media class and tape library are already inserted. You might want to change the Check for Active files field to \*NO to shorten the initialization time (Figure 7-35).

```

Initialize Media using BRM (INZMEDBRM)

Type choices, press Enter.

Device . . . . . > TAPMLB15   OPTVRT01, TAPMLB14, TAPMLB15..
New volume identifier . . . . . > 'I00000

Media class . . . . . > ULTRIUM3   FMT3592, FMT3592A1, FMT3592...
New owner identifier . . . . . *BLANK
Volume identifier . . . . . > 'I00000

Check for active files . . . . . *no   *YES, *FIRST, *NO
Code . . . . . *EBCDIC   *EBCDIC, *ASCII
End of media option . . . . . > *UNLOAD   *REWIND, *UNLOAD
Clear . . . . . *NO   *NO, *YES

Bottom

F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys

```

Figure 7-35 Initialize Media using BRM (INZMEDBRM)

IBM i/BRMS is now ready for backup to ProtectIER.

For more detailed information about this topic, refer to the ProtecTIER documentation found at the following address:

[http://www-03.ibm.com/systems/storage/tape/?cm\\_re=masthead\\_-\\_products\\_-\\_stg-tape](http://www-03.ibm.com/systems/storage/tape/?cm_re=masthead_-_products_-_stg-tape)

You can also refer to an IBM white paper for ProtecTIER attached to IBM i, found at the following address:

<http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101525>

## 7.7 Data types

The other factor that affects performance in a ProtecTIER environment is the data that is being targeted for backup.

Some data, such as databases and email, is highly compressible and also factors quite well. Other data, such as video or seismic data, cannot be compressed or factored well.

The following considerations are for the data type or application that sends backup data to your backup server.

### 7.7.1 Oracle database data

When using database backup tools such as Oracle Recovery Manager (RMAN), make sure that the multiplexing-like option is disabled. For example, prior to ProtecTIER being installed, RMAN by default sends a total of nine files requiring backup to three channels. This equates to three file streams, each containing three files ( $3 \times 3 = 9$ ) to three physical tape drives.

After ProtecTIER is implemented, RMAN must be altered to send nine file streams, each containing one file ( $9 \times 1 = 9$ ), to nine virtual tape drives in the virtual tape library. This is achieved by changing your RMAN backup commands to only send one file per stream/backupset using this option. For each backup command, set the value to "FILESERSET=1".



## IBM System Storage ProtecTIER native replication operation

In this chapter, we discuss ProtecTIER Replication. ProtecTIER with replication enables virtual tape cartridges to be replicated from the primary site to a secondary *or multiple* locations for enhanced disaster recovery (DR) and business continuity (BC) capabilities. By eliminating the need to transport physical tape cartridges, data can be recovered faster and more reliably, enabling users to get back online more rapidly in the event of a disaster or major system outage. The dramatic reduction in the required network bandwidth between the primary and secondary sites enabled by ProtecTIER's deduplication technology radically reduces the costs associated with electronically transmitting data to a remote location for disaster recovery purposes.

By dramatically reducing costs, ProtecTIER with replication enables IT organizations to easily expand the coverage of replication to all of the applications in their environment, as opposed to deploying replication for only a select few applications and tolerating significant risks of data loss and slow recovery for most other applications.

Figure 8-1 demonstrates the prohibitive costs of traditional replication, and contrasts it to the increased level of DR protection enabled by ProtecTIER. Figure 8-1 represents a generic IT environment in which the cost to protect 30% of their data with traditional replication is equal to the cost of covering 100% of the data with ProtecTIER replication.




Figure 8-1 Relationship between the cost and percent of environment protected

## 8.1 How replication works

Replication is an additional feature built into the VTL so that users can pick and choose some or all of their cartridges to be replicated to the DR site. Because ProtecTIER deduplicates data before storing it, only the changes, or unique *elements* of data, are transferred to the DR site over the replication link. This translates into substantial savings in the bandwidth needed for the replication link. Data transfer is started based on several trigger points, such as policy based transfer windows, or movement of the virtual tape to a VTL *export slot* (the VTL emulates import and export slots, and the opening/closing of the *library door* to *insert* or *eject* a cartridge). Data verification and validation is done at the DR site to ensure integrity of the transferred data prior to making the virtual cartridge/tape available.



Figure 8-2 shows an example of a deployment for two systems using replication.

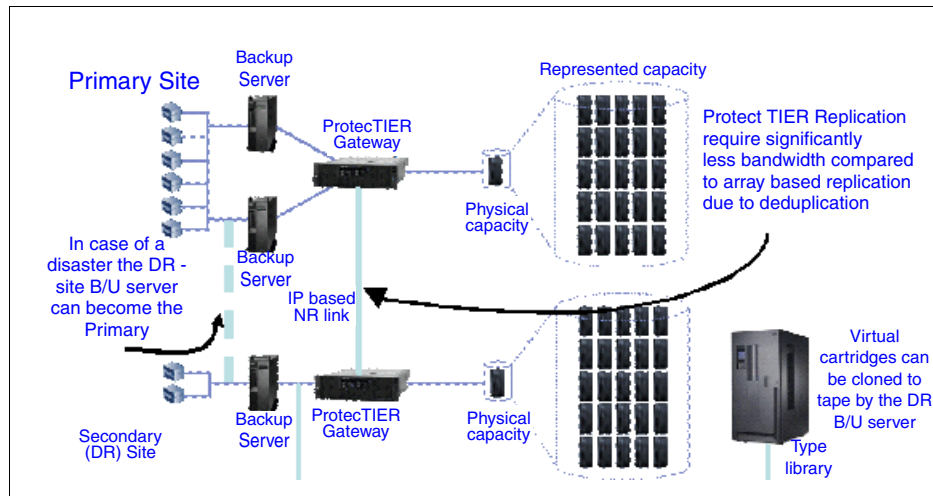


Figure 8-2 Typical deployment for two systems using replication

### 8.1.1 Replication features

ProtectTIER Replication is designed to provide the user with great flexibility, to work seamlessly with any of the leading backup applications, and to fit easily within the overall tape operations paradigm. Replication is policy based, allowing users to define several different policies for replicating cartridges from one system to another. The granularity of the replication policies allows users to set policies for an individual cartridge, a pool of cartridges, or an entire virtual tape library. Replication is performed asynchronously at the logical cartridge level, and replication progress can be tracked and monitored at the cartridge level through the ProtecTIER management GUI. Full data validation is performed at the secondary site to ensure enterprise-class integrity of the transferred data prior to making the virtual cartridge available. Once replicated to the DR site, users may choose to clone these cartridges to real physical tape using their backup application. In the event of a disaster, the DR site's TS7650, TS7650G, or TS7610 system can become the production site until the main site comes back online. At that point, the user may replicate or move the newly created tapes back to the main production site.

In the following sections, we discuss a few of the key features and design points that demonstrate the flexibility and synergy of the tape operations paradigm.

#### Virtual tape management framework

As a target of the backup application, a ProtecTIER system presents itself as a tape library (or many libraries) to the network. The backup application manages the cartridges within a ProtecTIER system as though they were real cartridges, including read, write, import/export, tracking media with barcodes, and many other operations. Because replication at the ProtecTIER level is transparent to the backup application, ProtecTIER's replication function is designed to allow synchronization with the backup application by way of normal tape management methodologies.

## Shelf

The ProtecTIER replication feature introduces the concept of a virtual shelf. As a general rule, replication always occurs from the source shelf to a destination shelf. As with physical tape shelves, there is a limit to the number of tapes that can be put on a virtual shelf. The limit is the result of subtracting the total number of tapes in all libraries from the total number of tapes supported in a single repository.

## Visibility switch control

Visibility switch control is the means by which ProtecTIER can determine *where* cartridges actually exist, because from a backup application standpoint, any specific cartridge or barcode can *exist* in only one location at a given time. ProtecTIER native replication will provide a virtual *shelf* (Figure 8-3) that is visible internally only at the ProtecTIER level (after it is exported by the backup application) and will provides more flexibility in managing tapes/cartridges and where they are kept, similar to keeping physical tapes on an actual shelf outside of the tape library.

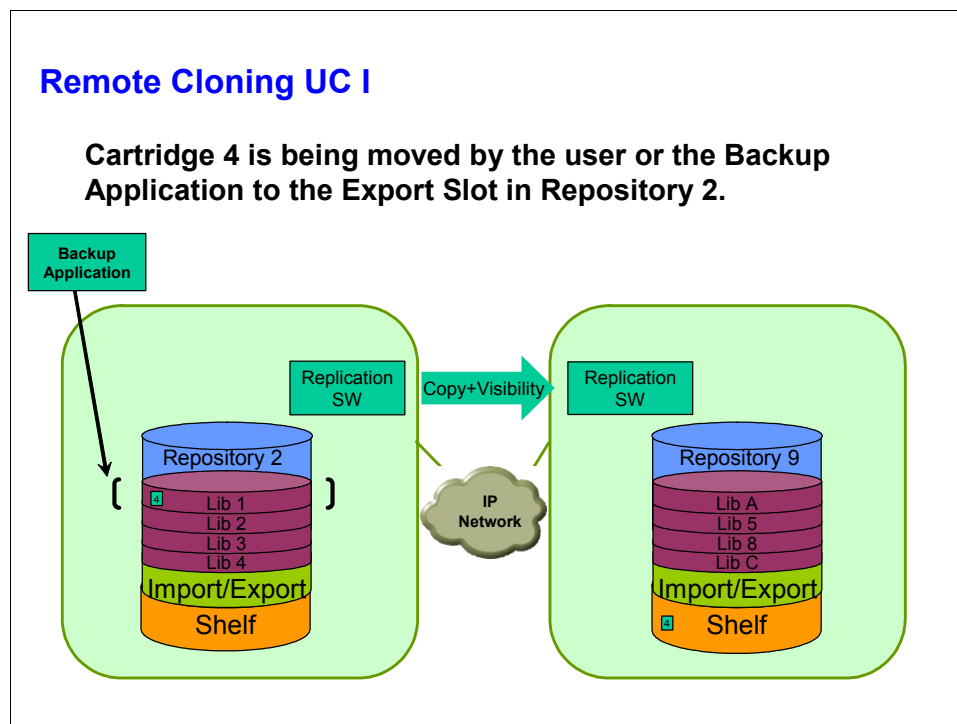


Figure 8-3 Virtual shelf

To ensure that a given cartridge is only visible to the backup application in one location despite the fact that a replica exists, ProtecTIER offers the *visibility control* function that allows the user to determine in which location the cartridge should be accessible to the backup application. This is achieved by using the import/export slots of the virtual libraries and exactly mimics the operation of physical tape management.

## Cartridge cloning and visibility control

A common use case for ProtecTIER users is to first replicate data from the primary site to the secondary site, and then to move the data from the disk-based repository onto physical tape cartridges for long-term retention. With the visibility control function, ProtecTIER makes this operation simple for users. After the cartridges complete replication to the secondary (DR) site, users may clone these cartridges to physical tape using their backup application tape copy function. This allows the backup application to remain in control of the end-to-end process and maintain its catalog of all cartridges, the associated data, and the location.

## Policy management

Replication policies allow for flexibility in controlling replication and enable a high degree of automation. Policies consist of rules that dictate the transfer of data from one repository to another, based on events such as writing new data to an existing cartridge that belongs to a policy. By setting policies for ranges of barcodes, users may implement differing degrees of protection for different applications. Users can also assign various priority levels to replication policies, which determine the order in which the data is transferred across the network.

## Replication-network management

ProtectTIER repositories belong to a replication grid, which is a framework for managing and monitoring the replication activities between ProtectTIER systems. Through ProtectTIER Manager, the user can monitor the status of the overall replication-network, the relationship between grid members, and the data throughput rate of replication. Further statistics about the cartridges involved in replication policies and the statistics of each repository are available in a single display from ProtectTIER Manager to enable ease of management and use.

## Recovery management

When the need to fail over to the secondart site arises, whether due to a full disaster or a lower-level disruption, the ProtectTIER system enables rapid recovery of the data and restoration of the production applications such that business operations can continue with minimal downtime. ProtectTIER is designed to enable rapid recovery of data from cartridges using the media server at the remote site. After data has been restored, the ProtectTIER system can be brought online as the production system and be used for backup and recovery until the primary site has been restored. During the time period that the disaster recovery site acts as the primary production site, its data can be protected by replicating it to another secondary site. ProtectTIER Manager should be used to *pair* the DR site (now acting as the primary) to the other secondary site. (For more details, see Chapter 11, “Native replication and disaster recovery” on page 575.) When ready, the user performs a failback operation to move production activity back to the primary site. At that time, the user may replicate any new data at the secondary site back to the primary and return the primary to its original status. All of these steps are enabled through the user interface.

### 8.1.2 Typical deployment

ProtectTIER native replication enables data replication capability across repositories, between ProtectTIER nodes connected to the same WAN. This WAN capability allows you to replicate data at any distance to remote cities, states, territories, and so on, for the purposes of disaster recovery. The deduplication software combined with replication reduces the amount of bandwidth required.

ProtectTIER native replication provides high availability in accessing backed up data for backup and restore operations across sites (see Figure 8-4). By having a current copy of critical data in a remote site along with a hot configuration of a TS7650, TS7650G, or TS7610 system and your backup application ready to restore on a remote machine, data restoration can commence in a matter of minutes.

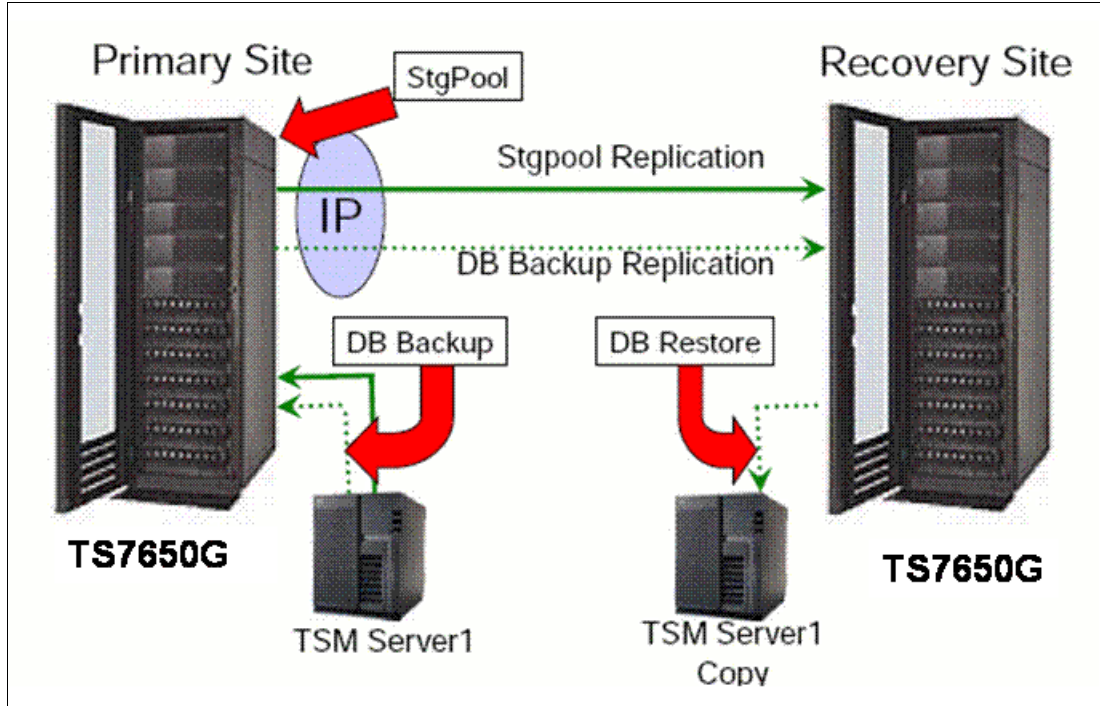


Figure 8-4 Backup applications using replication

### 8.1.3 ProtecTIER native replication Management Interface

The ProtecTIER native replication Management Interface provides an option-rich platform and is easy to use. The ProtecTIER native replication Management Interface:

- ▶ Defines policies and assigns priorities (Figure 8-5).
- ▶ Chooses which cartridges will be replicated.
- ▶ Schedules the replication window, that is, the time frame in which replication takes place for all policies.
- ▶ Defines cartridge *visibility* features:
  - Determines *where* virtual cartridges *exist* from ProtecTIER or a backup application standpoint.
  - ProtecTIER native replication emulates moving tapes in and out of VTL's import/export slots.
  - Allows users to use the VTL export/import slots through the backup application to change the visibility of the cartridges from one library to another.

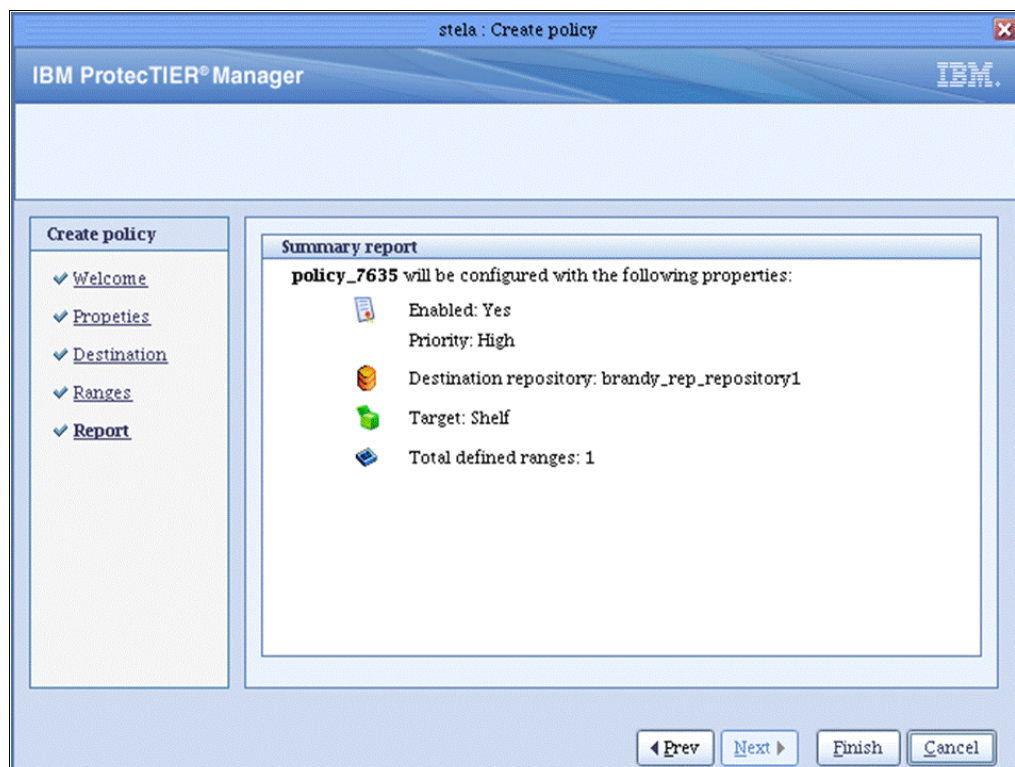


Figure 8-5 Replication policy

## 8.2 Normal operation concepts

This section describes a normal ProtecTIER operation sequence when replication is invoked.

### 8.2.1 Replication

When the backup application is writing to a cartridge that is part of a replication policy, ProtecTIER checks when it needs to be replicated and what its priority is so that the cartridge can be put in the correct order in the replication queue. Cartridges are always being replicated from the local library at the primary site to the virtual shelf of the repository at the secondary (DR) site. By definition, cartridges that are created at the primary (local) site repository are set to be read-write enabled, so that the backup application at the primary site has full control of them and their content. Cartridges that were replicated to the secondary (remote) site are set to be in a read-only mode. By default, only one cartridge instance is located in a library. The replica is located on the virtual shelf.

**Note:** At any time, through ProtecTIER Manager, the user can override the default *location* of any given cartridge and manually move the replica from the virtual shelf to a library in the secondary site repository.

Cartridges are marked as synced after the data finishes replicating from the primary to the secondary site, so that at the time of sync, the local cartridges and their remote replicas are exactly identical. Up to 128 cartridges can be replicated simultaneously. Before replication starts running, the system ensures that only unique, new data will be transferred over the wire. To achieve this objective, each side holds sync data per each of its cartridges. This sync data is used by the destination (secondary, remote site) to figure out which data (if any) should be replicated, as only new and unique *elements* of data are sent over the wire to the remote site. The replication mechanism has two types of data to transfer:

- ▶ Metadata, which is the data that describes the actual data and carries all the information about it
- ▶ User data, which is the actual backed-up data

Network failures, if and when they occur while the replication operation is being carried out, lead to retries of up to seven consecutive days to replicate any specific cartridge that did not finish its replication due to the line failure.

### 8.2.2 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 - primary site) repository and the destination (DR site) repository.
- ▶ Reads the unique replication data units upon requests from the remote ProtecTIER system based on what it is missing.
- ▶ Sends the unique data elements, 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 which the replication should start.
2. Receives many data units concurrently as part of the replication action.
3. Verifies CRC for all replicated data before it becomes available as part of the cartridge.
4. After the CRC check proves successful, the system moves each of the verified data elements into the cartridge scope and makes it available for the user.

After all verified pieces are inserted into the cartridge, it becomes available as a complete cartridge. However, as a replica in the destination (DR) repository, it is set as read only and cannot be used for backup purposes. This is an important factor in failover situations when the DR system might temporarily become the production site and can accept local backups. At that time, the user should create new tapes to accept this local backed-up data. These new local DR site tapes can be replicated to the primary site after it becomes available and ready for production again. During the failback process, which is when the user moves operations back to the primary site, the newly created cartridges from the DR site can be replicated to the primary site. Under these circumstances, the system grants read and write permissions to these replicated cartridges at the primary site, which becomes the *owner* of these tapes from that point on, just as though they were created there.

**Note:** The replication data transfer process requires that the replicated cartridge reside at the remote (DR) site ProtecTIER VTL shelf. As long as the data transfer is being performed, these cartridges cannot be moved from the shelf. If the cartridge is manually moved from the shelf, it may cause replication errors if it attempts to replicate.

### 8.2.3 Visibility switch control

This is an automated process of the ProtecTIER replication feature that transfers the visibility of a cartridge from its master to its replica and vice versa. Just like a *real* physical tape, a virtual cartridge can only reside in or be visible to the backup application in one location at a time. This is carried out by the ProtecTIER replication by using the VTL Import/Export slots. The system uses the export slots and eject operation as triggers to begin processing a tape move. The backup application can eject the cartridge into one of the export slots. As soon as that cartridge replication action is completed, the cartridge appears at one of the import slots of the secondary (remote, DR) site and can be imported into a library.

Replication visibility is an attribute defined in a replication policy and can be set by the user. A specifically defined library in the remote repository must be selected for the visibility transfer to be carried out. After the cartridge is ejected at the primary (local) site, it moves to the local virtual shelf. Then it is replicated as part of a policy to the remote site virtual shelf. After the replication is completed and verified, that replica moves to the respective library's import slot and can be imported into that library. At this point, that cartridge is only visible to the backup application at the remote site. A copy of the cartridge stays at the primary ProtecTIER system for fast recovery. However, it is hidden from the backup application. The way to move the cartridge back and make it visible to the backup application at the primary site is to eject it from the remote library and import it back into the local one (the same library it came from). Because the system keeps a hidden copy at the primary site, this move back is instantaneous.

## 8.2.4 Single domain and multiple domain backup application environments

Any backup environment can be set up in many topologies. From ProtecTIER's replication standpoint, there are two general setup methods for these environments:

- ▶ Single domain
- ▶ Multiple domains

The backup application catalog/DB has an entry for each cartridge used for backup. These entries include:

- ▶ Date when backup was performed
- ▶ List of files associated with the backup
- ▶ Retention period
- ▶ Other backup application-specific information

The backup application supports one catalog/DB per backup server instance. In many cases the primary and remote (secondary, DR) sites have two separate backup servers, each with its own DB or catalog.

To efficiently read replicated cartridges at the remote site, the remote backup server needs access to the actual catalog or DB of the primary backup server or an exact copy of it. There are two basic backup environment topologies, as discussed in the following sections

### 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 (NBU) and will not work with most deployments of IBM Tivoli Storage Manager.

### Multiple domain environment

A multiple domain approach is more widely used. This is where the backup application does not share the same catalog between the primary (local) and secondary (remote, DR) sites. This is the most common scenario with Tivoli Storage Manager environments. In this type of environment, each of the backup servers in both the primary and the secondary locations has its own backup catalog.

In a multiple domain environment, the backup server at the primary site manages the backup activity into the local ProtecTIER libraries.

An independent local backup server at the secondary remote location is used in case of a disaster (when the primary site is *lost*) to recover the replicated cartridges and potentially continue, temporarily, the production backup activity. In this type of environment, the general steps to recover from a disaster will include:

- ▶ The secondary (DR) site backup server should be recovered first with the catalog or DB from the primary site.
- ▶ Once up and running, this backup server should be used to restore data from replicated cartridges. It might also resume regular backup and restore production activity using new virtual cartridges.
- ▶ During the time frame in which the production activity is conducted at the remote (DR) site, new backups will register with the backup catalog or DB.



- ▶ After the primary site system has been recovered and is ready to become the production system again, the failback operation should be invoked. This replicates back all the needed data (including newly created cartridges) to the primary location, allowing for synchronization between the two ProtecTIER repositories. It is important to synchronize the backup application catalog during this failback action to complete the recovery operation.

For a more detailed discussion of specific backup applications, see Chapter 7, “Backup and restore applications” on page 343.

## Remote cloning

Remote cloning is the process of using a secondary site to clone cartridges. ProtecTIER replication enables users to offload tape cloning to their 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 onto physical tape cartridges for long-term retention. One of the advantages of performing this practice at the secondary site is to take the burden of cloning to physical tape from the production environment to the remote location, which is where tapes will most likely be kept after it is cloned. The remote cloning operation uses the cartridge replicas residing at the destination's ProtecTIER VTL shelf for cloning. The process imitates the commonly used physical process of transporting the physical cartridges from the primary site to a remote site either after being cloned or to clone them there for long-term archival purposes.

This feature is effective in single domain backup deployments because in these environments the backup application servers at both sites share the same catalog and can be connected to the ProtecTIER systems concurrently. Within these environments, the replication visibility switch control feature should be used and the cartridges to be cloned should be *moved* from the primary repository to the secondary repository and then be cloned to physical tapes. Because in a single domain environment the catalog is shared across the two sites, the backup application is fully aware of the locations and status of the cloned cartridges.

For a detailed discussion about this topic, refer to 8.2.4, “Single domain and multiple domain backup application environments” on page 418.

### ***Consideration for remote site tape cloning in a single domain environment***

More than 90% of restore requests use the most recent backup tapes. A best practice for users who will use the DR site physical tape cloning option in their operation is:

1. Back up to the local, primary site ProtecTIER system.
2. Replicate to the secondary, remote site ProtecTIER while leaving the cartridge visibility at the primary site.
3. Five to seven days after replication has completed, change the visibility of the relevant cartridges from the primary, local site to the secondary, remote site.
4. Run `/ault/clone` to clone the remote ProtecTIER tape copy to physical tape.
5. Change the visibility of these cartridges back to local ProtecTIER.

This practice improves the probability of a virtual tape being available at the primary site when it is most likely to be used in support of a tactical restore at the primary location.

## 8.3 Many to one replication

ProtectTIER many to one replication in VTL and OST configuration supports both single node and clustered ProtectTIER configurations within all platforms: TS7650, TS7650G, and the ProtectTIER Appliance Express.

### 8.3.1 ProtectTIER with Virtual Tape Library

In ProtectTIER V2.5, the grid members are divided into roles for replication: the sources, called *spokes*, and the targets, called *hubs*. There can be different spokes and hubs defined within a grid. Every spoke and hub has its own repository. The user has different possibilities to set up his replication. You can define which spoke will replicate with which hub within a grid. The ProtectTIER Replication Manager must be installed and configured on a node within the grid. One spoke can only be replicated with one hub. A hub can be the target of up to 12 spokes. There can be different hubs within a single grid. (see Figure 8-6)

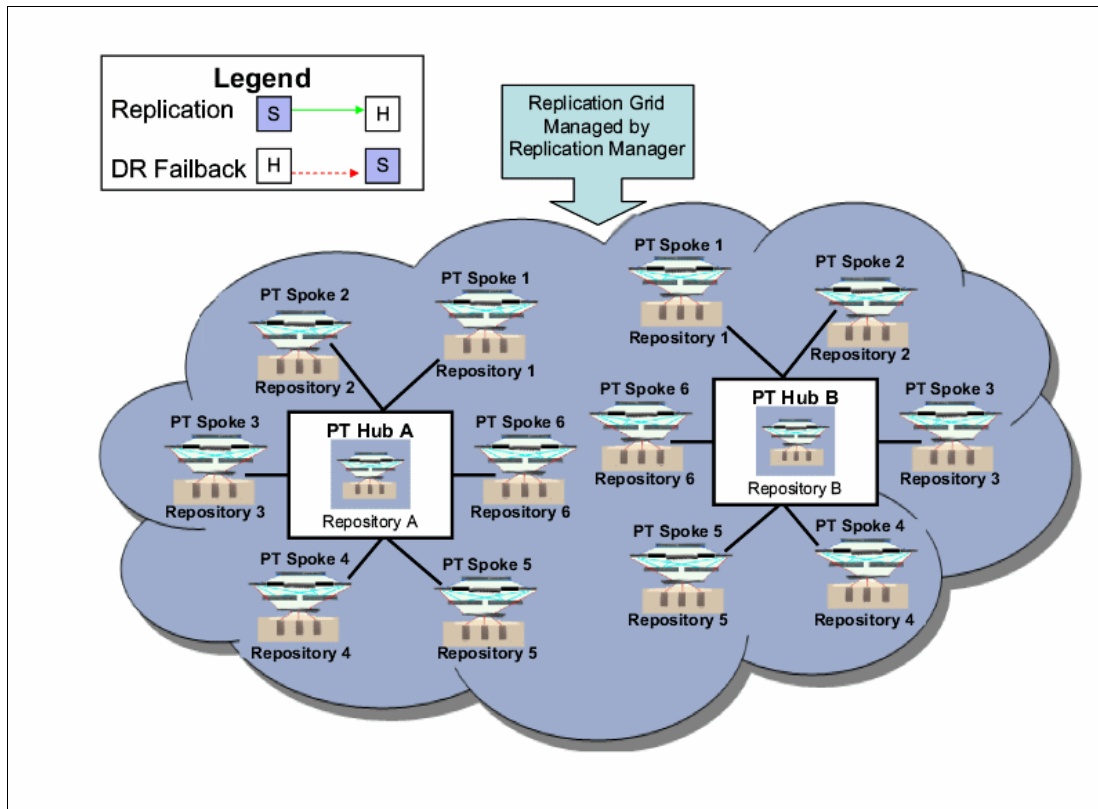


Figure 8-6 Replication Grid including two different hubs and connected spokes

Regarding spokes:

- ▶ A spoke is the source repository that runs the user backups. The backups are replicated to a defined hub.
- ▶ A spoke may not be target of a replication.
- ▶ A spoke may only replicate to a single hub. There is no ability to replicate with multiple hubs.

Regarding hubs:

- ▶ A hub is the target repository that may receive replications from up to 12 spokes.
- ▶ A hub can also be a target for local backups, but is not able to replicate them.
- ▶ There can be eight active hubs defined per grid.
- ▶ There can be a maximum of 12 spokes defined for a single hub.

In Figure 8-7, you can see the replication view of a many to one replication in ProtecTIER Manager.

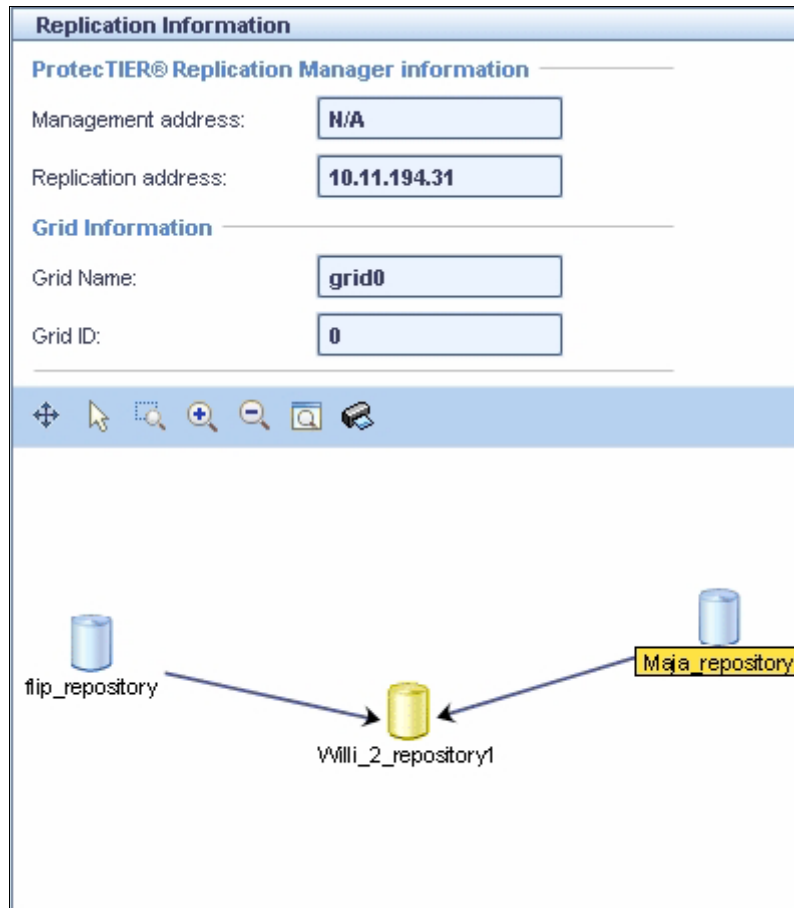


Figure 8-7 ProtecTIER Manager view of many to one replication

### 8.3.2 ProtecTIER with OpenStorage

ProtecTIER V2.5 adds the OpenStorage feature (OST). The OST is a Symantec NetBackup Application Program Interface (API), supporting direct communications between Symantec NetBackup and OST enabled ProtecTIER systems.

With OpenStorage, ProtecTIER can be integrated with NetBackup to provide the means for backup to disk without using a virtual tape library (VTL) emulation. Using a plug-in that is installed on an OpenStorage enabled media server, ProtecTIER can implement a communication protocol that supports data transfer and control between the backup server and the ProtecTIER server.

Therefore, to support the plug-in, ProtecTIER implements a storage server emulation. (Figure 8-8 shows the OST backup scheme.) If you want to configure the ProtecTIER server for OST, make sure that a repository has been created with OpenStorage as backup interface. For OST usage the repository is defined in logical storage units (LSU). The LSU is also defined in the OST API and is a container of storage. Up to 256 LSUs can be defined per ProtecTIER server and are identified by a name string that is unique within the storage server.

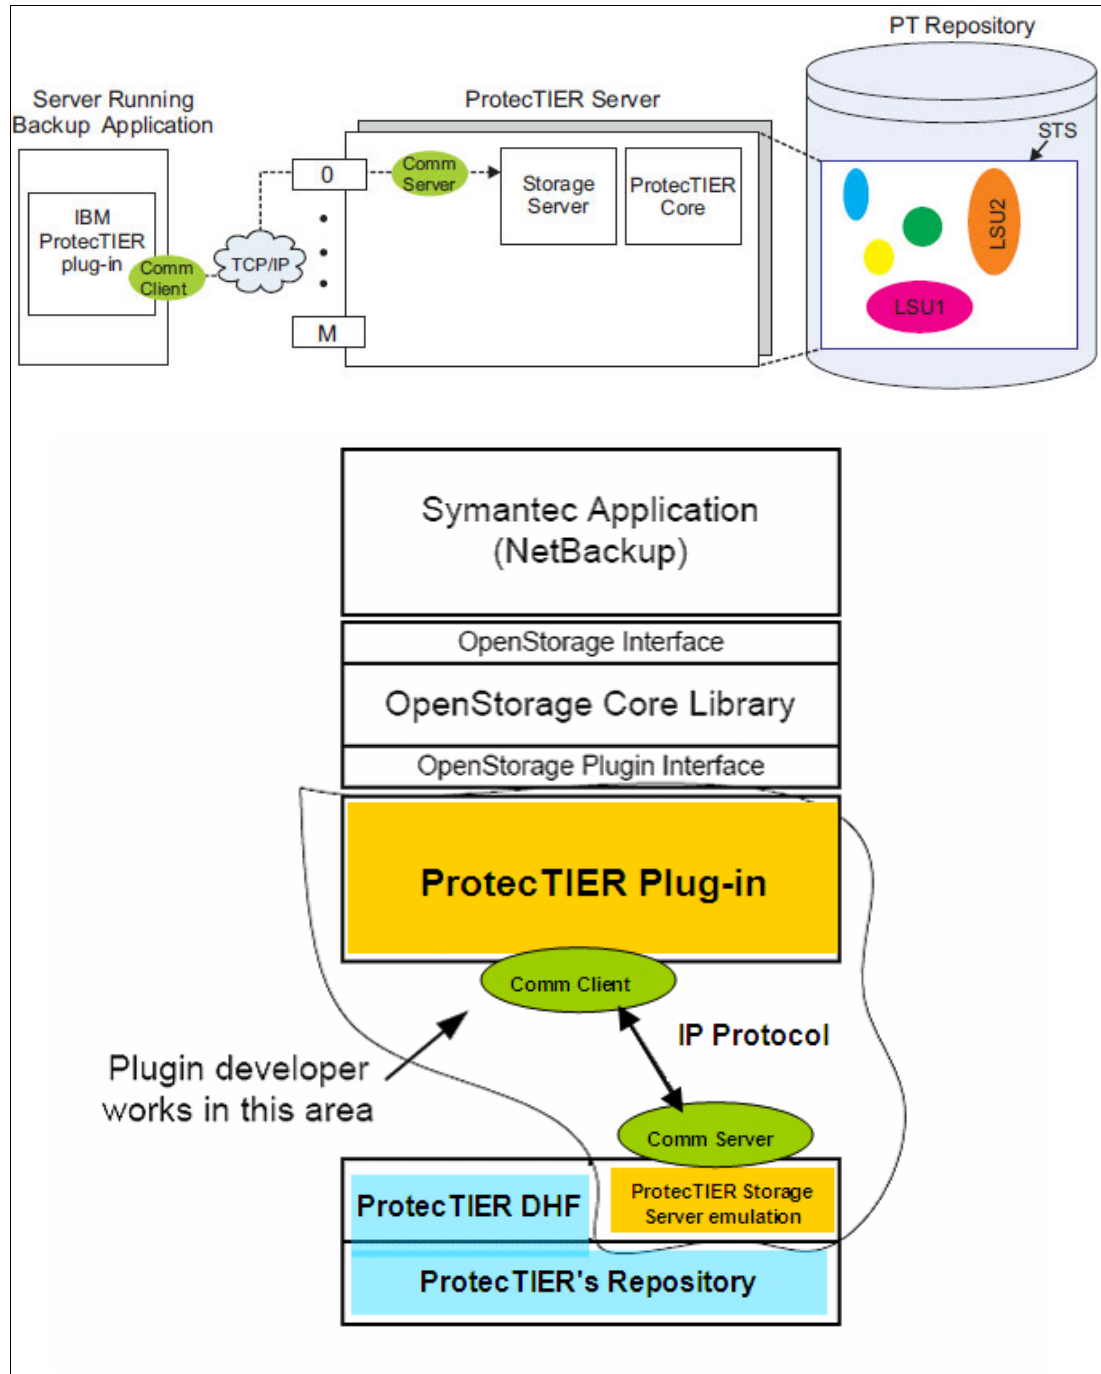


Figure 8-8 OST backup scheme

## Replication with OST

In an OpenStorage environment, replication policies and activities are managed by NetBackup. ProtecTIER Replication Manager is used to manage the replication grid, create application groups, and define the replication connections within the replication grid. Each repository, or grid member, can replicate to each other and replication is bidirectional. A maximum of 12 repositories (members) can replicate to each other at any given time, with up to 256 repositories existing in the entire grid.

## ProtecTIER Replication Manager tasks in OpenStorage

These tasks include:

- ▶ Managing the repositories in the replication grid that can be replicated to
- ▶ Maintaining the IP addresses of all repositories
- ▶ Updating of repositories leaving and joining the grid
- ▶ High-level monitoring and statistics of traffic in the replication grids

In Figure 8-9, you can see the ProtecTIER Manager for an OST enabled ProtecTIER system in a OST hub mesh replication group.

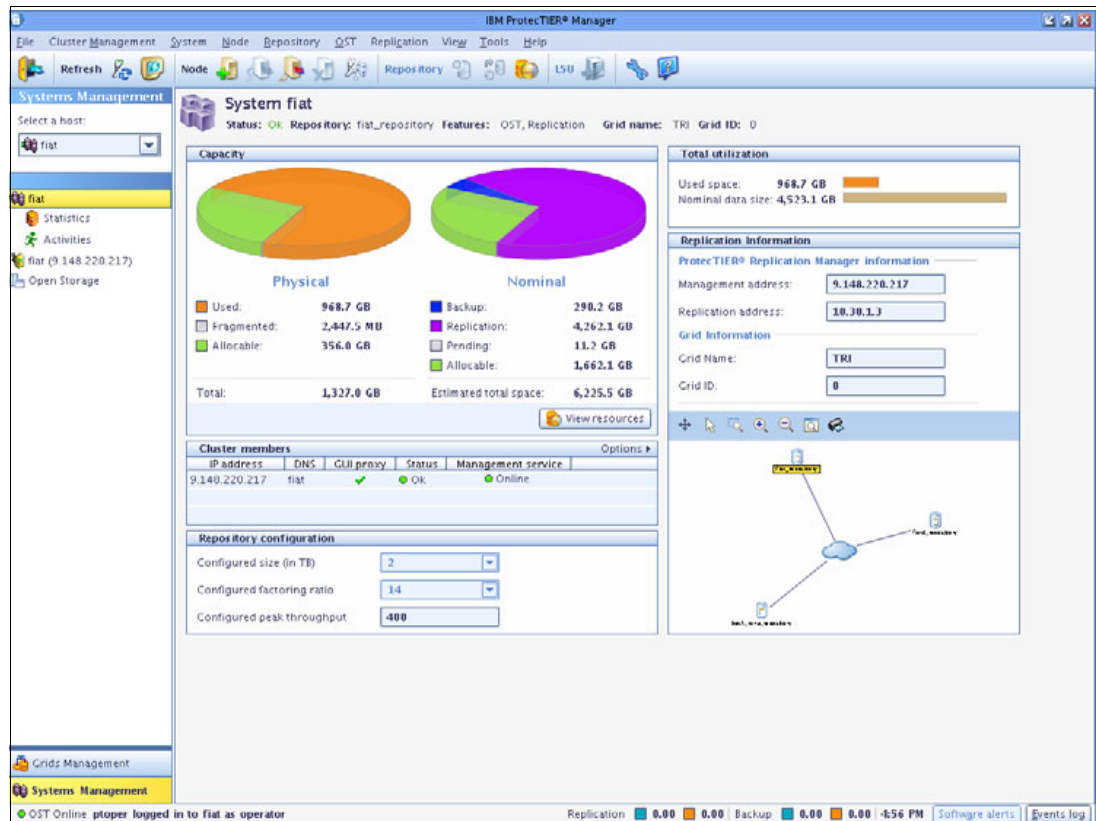


Figure 8-9 ProtecTIER Manager with OST enabled system

**Note:** For a detailed explanation of OST architecture, refer to Chapter 2, “IBM System Storage ProtecTIER architecture” on page 11. For host implementation, refer to Chapter 3, “Planning for deduplication and replication” on page 53.

### 8.3.3 Replication management

In this section, we describe the different topologies, prerequisites, components and features of a ProtecTIER replication grid. The section discusses ProtecTIER systems with OST features and systems with VTL features.

#### Replication topologies

In ProtecTIER V2.5, depending on whether it is a new installation or an upgrade, there are two different replication states with different topologies:

- ▶ VTL application only
  - With R2.3 only as a pair
  - With R2.4/R2.5 hubs and spokes and many to one replication implemented

- ▶ OST application only

With R2.5 and the OpenStorage (OST) feature, ProtecTIER supports implicit hubs, where up to 12 could be meshed to an OST Hub Mesh Group.

At this time, ProtecTIER may only function as a VTL or an OST system, not as both at the same time. In both of the applications, there are groups of repositories defined that enforce the supported application deployments. The application groups are defined and managed in the ProtecTIER Replication Manager. Each application group can remain separately in a replication grid (see Figure 8-10).

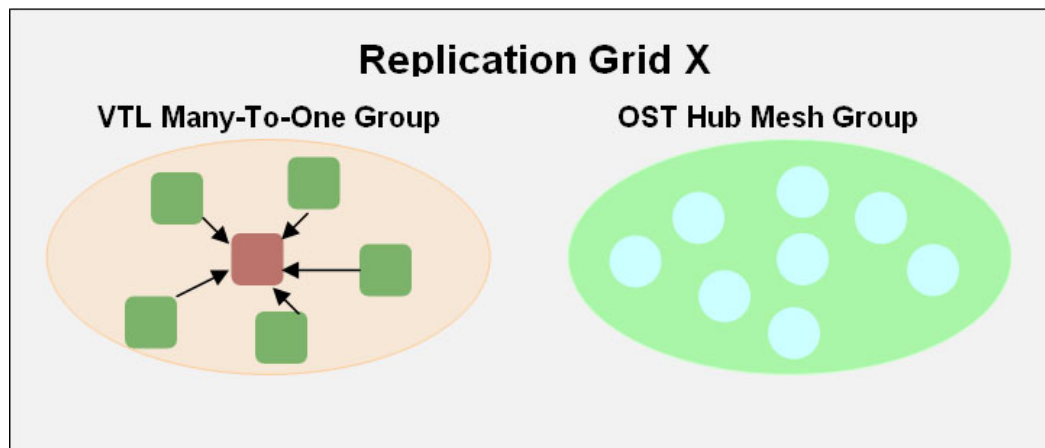


Figure 8-10 Replication grid topologies

#### Prerequisites

There are several prerequisites for replication:

- ▶ A fast network connection is needed.
- ▶ The management and replication networks need to be separated from each other.
- ▶ The ProtecTIER Replication Manager must be installed on a node of a cluster within the grid.
- ▶ The replication IP addresses must be in different subnets because of high network traffic during replication.
- ▶ For redundancy, the replication links should not be connected through the same network components.

- ▶ The reserved Ethernet ports for models 3958-DD3, DD4, SM1, and AP1 must be reserved. Refer to Chapter 2, “IBM System Storage ProtecTIER architecture” on page 11 for a list of the reserved ports.
- ▶ For replication, the following TCP ports have to be opened on all networks: 6202, 3501, 3503, 6520, 6530, 6540, and 6550.

## Replication grid

A replication grid is a logical set of repositories that share a common ID and can potentially transmit and receive logical objects through replication. The ProtecTIER Replication Manager is the component that manages the grid's configuration (for example, grid creation/deletion, repository membership in the grid, and so on). A replication grid defines a set of ProtecTIER repositories and actions between them and is configured using the ProtecTIER Replication Manager. The ProtecTIER Replication Manager is a software component that is installed on a ProtecTIER server node. In a single grid, there are a maximum of 24 repositories.

The following items are used in a replication grid:

- ▶ Virtual tape library (VTL)
 

The ProtecTIER virtual tape library (VTL) service emulates traditional tape libraries. By emulating tape libraries, ProtecTIER VTL enables you to transition 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 ProtecTIER actually stores data on a deduplicated disk repository.
- ▶ Replication grid member
 

A repository that is member of a replication grid.
- ▶ Replication grid ID
 

A number from 0 to 63 that identifies a replication grid within the whole organization.
- ▶ Replication pair
 

A pair of repositories within a replication grid that replicate from one to another.
- ▶ Replication policy
 

A policy defines rules for a set of objects from a source repository to be replicated with a target repository (for example, rules for VTL cartridges).
- ▶ Replication unique ID (RID)
 

A number that uniquely identifies every repository. The RID is a combination of the replication grid ID and the repository internal ID in the grid.
- ▶ Shelf
 

A shelf is a container of VTL cartridges within a single repository.
- ▶ Visible switching
 

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'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.

## ProtectTIER Replication Manager

The ProtectTIER Replication Manager is a software component that is installed on a ProtectTIER node. The installation of the ProtectTIER Replication Manager on a dedicated host that is not a ProtectTIER node is only available through RPQ. The ProtectTIER Manager can be installed on one or more workstations. The ProtectTIER Replication Manager manages the configuration of multiple replication grids in an organization. The Replication Manager should be able to discover all the members of the entire network it handles on both replication subnets. An agent on every node in each ProtectTIER server interacts with the server and maintains a table of its grid members. A single Replication Manager can manage up to 96 spokes and eight hubs. You can access, add, and remove ProtectTIER Replication Managers through the ProtectTIER Manager. For the ProtectTIER Replication Manager setup and configuration, refer to Chapter 5, “IBM System Storage TS7600 with ProtectTIER initial setup” on page 171.

Each ProtectTIER Replication Manager has a unique identity. A repository, after it has joined a replication manager, cannot join a replication grid managed by a different replication manager, even if it has left the grid. This action prevents data collision.

**Note:** Though not required, you should designate the target server (hub) as the replication manager. By doing so, the ProtectTIER Replication Manager remains available in a disaster recovery situation.

In Figure 8-11, you can see an example for the ProtectTIER Replication Manager view of a replication grid with ProtectTIER VTL configured systems.

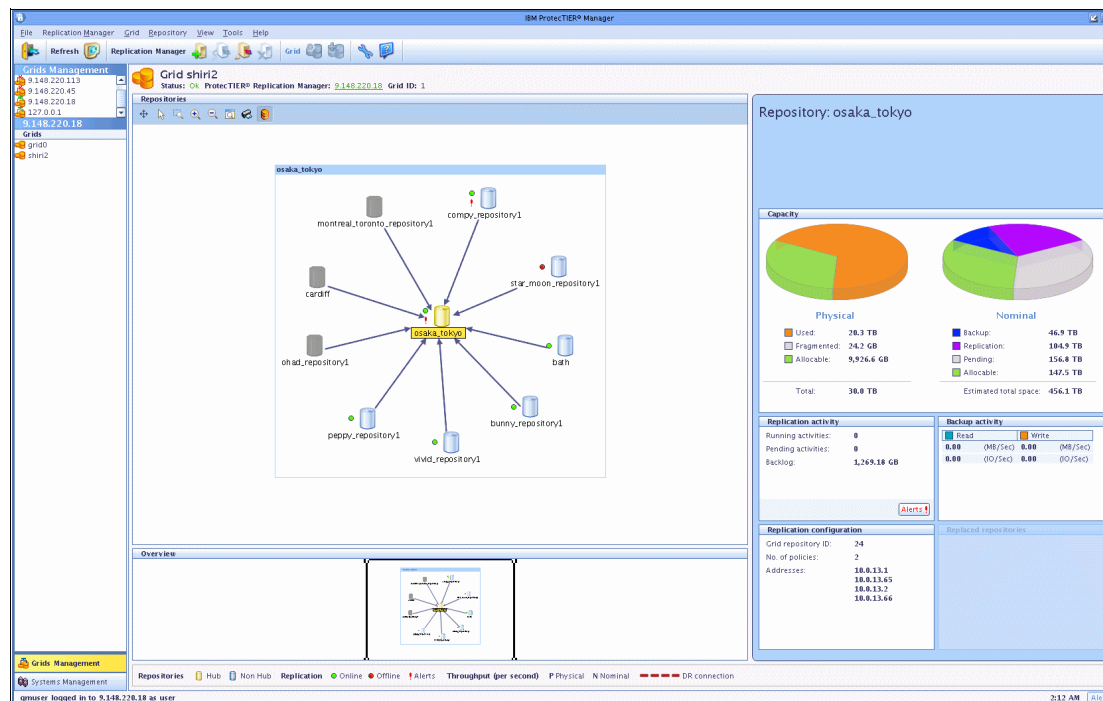


Figure 8-11 Many to one view in ProtectTIER Replication Manager in VTL configured system



To log in to a ProtecTIER Replication Manager, open the IBM ProtecTIER Manager and select **Grids Management** (Figure 8-12) and complete the following steps:

1. Select **Replication Manager** → **Login**. The login window opens.
2. Enter the gadmin user name and gadmin password.
3. Click **OK**. The Replication Manager window opens.

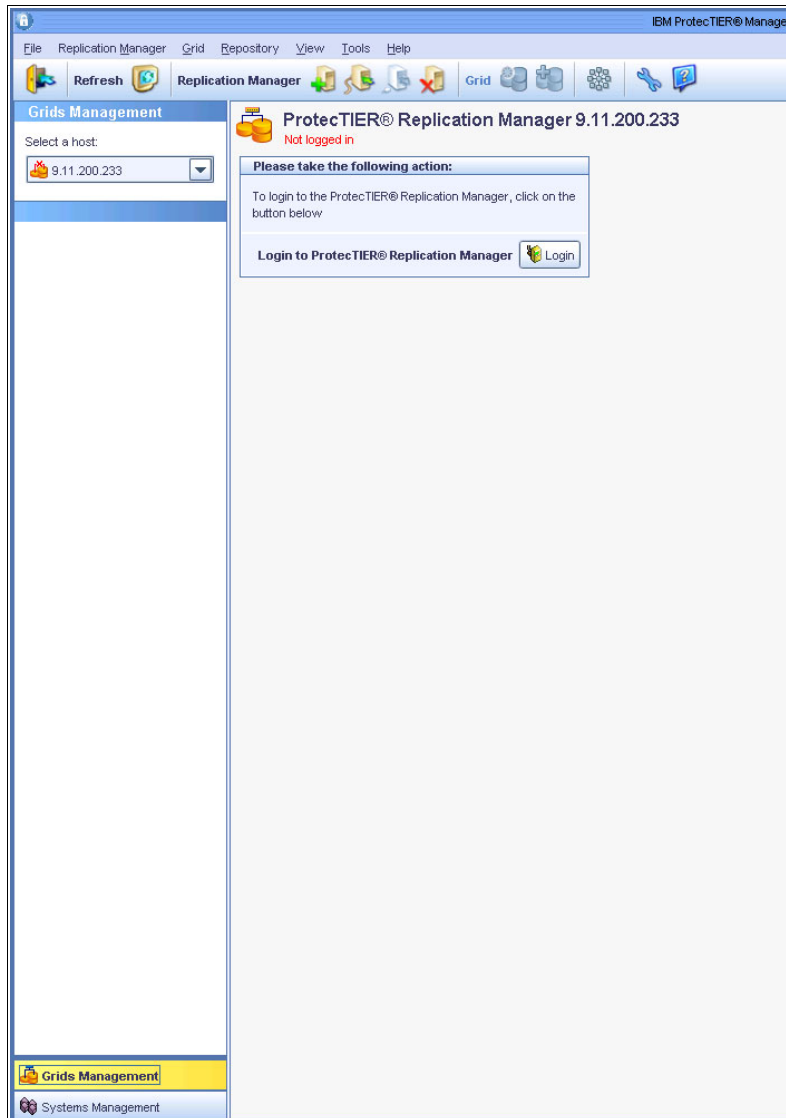


Figure 8-12 ProtecTIER Replication Manager login window

### Setting the replication rate limit

With ProtecTIER V2.5, you can set a replication rate limit. Setting the replication rate control (RRC) allows to limit the nominal and physical throughput (data flow rate) of replication. You can set up physical and nominal limits, which have no influence on one another. However, when using both methods, the physical settings will override the nominal ones.

The physical throughput limit restrains the amount of I/O and resources replication consumes on the local repository. Implicitly, this reduces the total load on the replication networks used by the repository (you can have two networks) and the amount of resources needed on the peer repository as well.

The nominal throughput directly affects the load on the destination repository. On the source repositories, the replication nominal rate does not necessarily compete with the backup. Setting the limit on a source repository guarantees that the backup gets the total possible throughput minus the nominal limit, but in many cases this is not needed.

The Replication Rate Limits window is divided into separate areas for physical and nominal throughput. These areas are divided further into various scenarios where you can limit the rate of replication, for example, replication rates for physical throughput during backup or restore, and replication rates when there is no backup or restore. The same options appear for nominal throughput (see Figure 8-13).

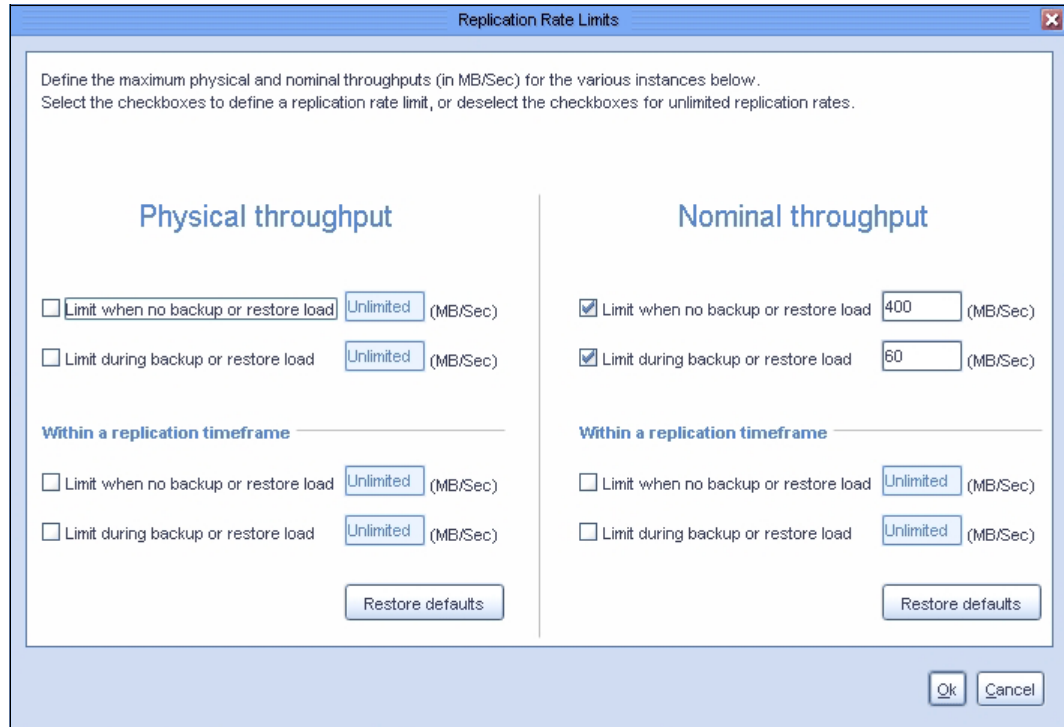


Figure 8-13 Replication rate limits window

### **Example for spoke replication rate**

Assuming a deployment where each spoke backs up 4 TB/day with a 16 hour replication window, and assuming a 10:1 deduplication ratio, 400 GB new/unique data must be transferred. 400 GB in 16 hours demands 7 MBps sustained for the 16 hour period. Each spoke will consume 7 MBps, equalling 38% of one OC3 link (20 MBps).

The same network infrastructure can be shared with other applications (under the assumption they have similar control mechanisms to ensure they do not consume ProtecTIER secured bandwidth)

### **Limiting the network interface (port) bandwidth consumption**

Bandwidth throttling is a way to control the speed at which replication activities operate, whereby the user can specify a maximum upper limit for the physical 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. In terms of physical ProtecTIER data transfer, after applying the TCP/IP impact, this translates into 175 - 190 MBps per single node and 350 - 380 MBps for a dual-node cluster. By default, there is no configured bandwidth limit; ProtecTIER will attempt to use as much bandwidth as it can.

If the physical network layer consists of a dark fibre or other high-speed network infrastructure, there is typically no reason to limit replication throughput. However, if ProtecTIER is running over a smaller network pipe that is shared by several applications simultaneously, the user may choose to restrict the maximum throughput used by ProtecTIER replication. This parameter is configurable per GigE port on all nodes in the replication grid, but it only applies to outgoing data, so it is only effective to 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, if you want to hold ProtecTIER replication to no more than 100 MBps and you are using all four available GigE ports of the source dual-node cluster, you must set each port's limit 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.

**Note:** 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.

### Replication modes: Scheduled and continuous

ProtecTIER offers two modes of operation for the replication activity: scheduled replication (with a predefined time window) and continuous replication, which runs concurrently with the backup operation.

The mode of operation is configured at the source system, and all the defined replication policies operate in one of these modes. In almost all cases, scheduled replication is a best practice approach, as it enables the user to accurately plan for performance and better ensure that SLAs are met.

When selecting the replication mode, the user is given two options:

1. **No Backup Precedence:** This is the dedicated window mode. With this option, the user can schedule a time window per day during which replication activities will occur. This mode of operation should be considered first, and for almost all customer 1:1 replication deployments (1 pair) use cases, it will produce the best performance.
2. **Precedence to Backup:** This is the continuous mode of operation, in which replication occurs concurrently with the backup operation. This mode is available and is the best practice for multi-site use-cases, as will be further described later in this section.

**Note:** For a detailed explanation of scheduled and continuous replication mode, refer to *Best Practices Guide for ProtecTIER V. 2.5 and TS7650G (Gateway) Attached Storage for TS7650*, GA32-0646.

Figure 8-14 shows the replication time frame window, where you can set up the specific replication time frame for your backups on a spoke repository.

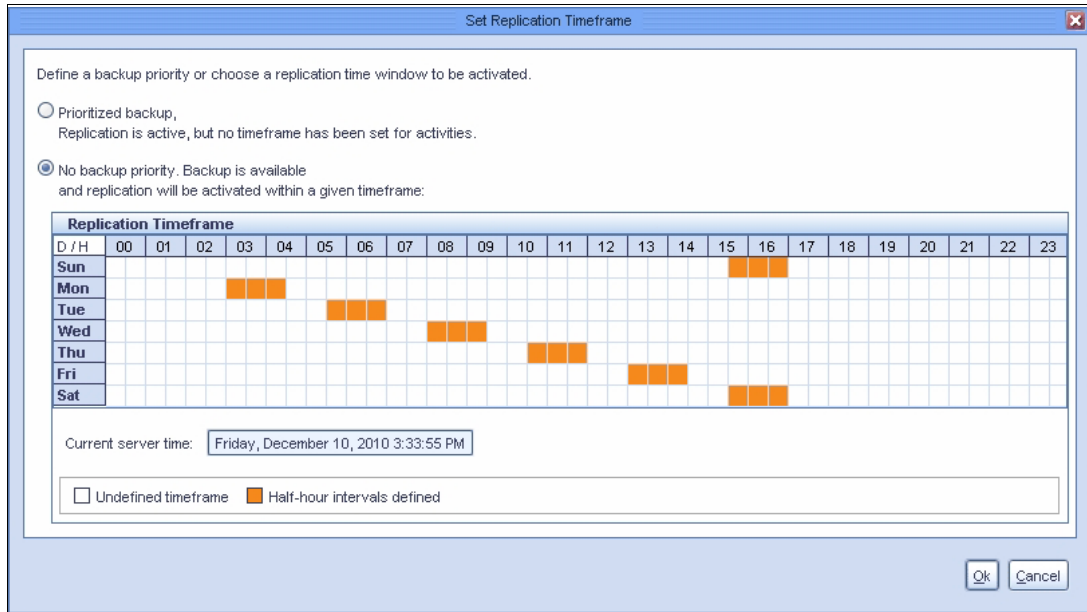


Figure 8-14 Replication time frame window on ProtectTIER Manager

### Reserve space feature

Local backups on a hub can receive priority when space in the repository is exclusively reserved for them. The replication from spokes connected to the hub cannot use the reserved space. By default, space is not reserved for specific backups. The space reservation can be set up through the graphical user interface (GUI). If there is space reserved for local backups, replication will fail when the minimum threshold is reached. The reserved space can be managed and changed during run time through ProtectTIER Manager. The space reservation can also be changed regardless of the actual stored data. The reserve space function will be handled by the ProtectTIER Manager as a repository management task.

Refer to Figure 8-15 for more details.

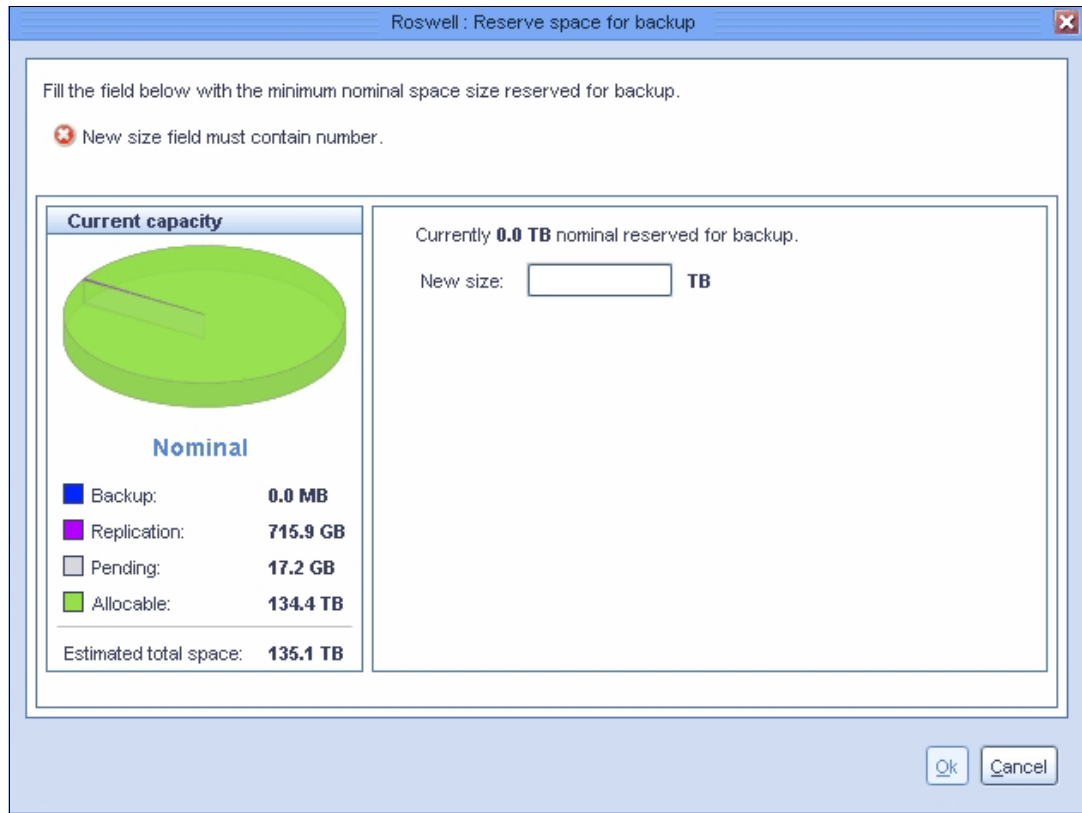


Figure 8-15 Reserve space menu in ProtecTIER Manager

### **Reserve space implications**

Consider the following implications:

- ▶ If the reservation definition is set too big, the replication fails because there is not enough space.
- ▶ If the reservation definition is set too low, the replication will occupy all the repository space and local backups will fail because of less free disk space.

**Note:** While planning the hub repository, it is essential to subtract the amount that is reserved for backup only from the total space to figure how much repository space is available for replication (or additional for backup if reserved space is exceeded).





# IBM System Storage ProtecTIER with Symantec OpenStorage

In this chapter, we describe the implementation and configuration of ProtecTIER systems with Symantec OpenStorage (OST). Throughout this chapter, we give you information about OST and how to install and maintain a ProtecTIER system with OST. This chapter discusses the following topics:

- ▶ Section 9.1, “OpenStorage setup” on page 434 explains the OST elements and OST environments.
- ▶ Section 9.2, “Configuring OpenStorage with ProtecTIER Manager” on page 438 explains how to add storage servers (STs) and logical storage units (LSUs) on to the ProtecTIER system.
- ▶ Section 9.3, “The ProtecTIER OpenStorage plug-in” on page 446 explains the installation and configuration of the OST plug-in.
- ▶ Section 9.4, “Configuring ProtecTIER OpenStorage on NetBackup” on page 452 explains how to configure the predefined ProtecTIER OST on NetBackup with GUI Tool.
- ▶ Section 9.5, “Replication settings for OpenStorage” on page 466 explains the replication settings for OST.
- ▶ Section 9.6, “Native replication with OpenStorage” on page 471 explains native replication on OST.
- ▶ Section 9.7, “The replication grid” on page 472 explains the replication grid.

## 9.1 OpenStorage setup

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

The OpenStorage function is supported starting with ProtecTIER V2.5. OpenStorage is an API that supports the Symantec NetBackup application. ProtecTIER V2.5 is able to support NetBackup because the OpenStorage function is supported. OST is mutually exclusive with VTL; if you install the ProtecTIER machine as OST, you cannot use the VTL function and vice versa.

### 9.1.1 OpenStorage elements

The OpenStorage environment is composed of several elements, which are a storage server (STS), a logical storage unit (LSU), ProtecTIER V2.5, the OST plug-in, and the ProtecTIER system.

#### Storage server

ProtecTIER V2.5 can support only one STS. An STS is a high level component defined by the OST API. An STS is container of LSUs. An STS is configured by the administrator within the storage appliance and correspondingly within the Data Protection Advisor (DPA). An STS can be created using ProtecTIER Manager.

#### Logical storage unit

An LSU is a storage container that contains images that consume storage. An LSU has capacity and used storage properties (nominal) that are reported to the DPA. LSUs partition the storage and are identified by a name string. An LSU is defined by the administrator within the storage appliance

Figure 9-1 shows the elements of an OpenStorage repository.




Figure 9-1 OpenStorage Repository elements



## OpenStorage plug-in

As mentioned at the beginning of this chapter, OST supports the NetBackup application. It can be implemented through an STS. STS is the vendor specific API. You have to install the plug-in software that supports each OS version. The OST plug-in is supported in AIX, Windows 2003, Windows 2008 server, and Linux. The ProtecTIER OST plug-in is a software component that is located on the NBU media server, and sends NBU requests to a ProtecTIER server. The plug-in uses a load balancing scheme to send requests between PT nodes and different interfaces in the nodes.

Table 9-1 shows which operating systems with which the OST plug-in is compatible.

Table 9-1 ProtecTIER OST plug-in compatibility

Platform	NBU version	Plug-in name
Windows 2003 32-bit	6.5.6	win32
Windows 2003 64-bit	6.5.6 and 7.0.0	win64
Windows 2008	7.0.0	win64
AIX V5.3 32-bit	6.5.6 and 7.0.0	aix32
AIX V5.3 64-bit	6.5.6 and 7.0.0	aix64
AIX V6.1 64-bit	7.0.0	aix64
Solaris5.10 32-bit	6.5.6 and 7.0.0	solaris32
Linux 5U4	7.0.0	linux64

## Hardware

ProtecTIER V2.5 can be installed on an IBM System Storage TS7650, TS7650G, or TS7610 machine type 3958-AP1, 3958-DD3, 3958-DD4, and 3959-SM1 system. The OST application can run only on DD4, AP1, and SM1 machines.

If you want to use the system as an OST machine, you have to choose between VTL and OST before ordering. The hardware configuration for OST is different than the one for VTL. The OST system must have an OST Ethernet card and does not have the front-end HBA card that is used to connect to the media server. For that reason, you cannot use VTL and OST at the same time

Figure 9-2 shows the OST on a DD4 system.




Figure 9-2 OST on DD4

Figure 9-3 shows the OST on an SMB system.




Figure 9-3 OST on SMB

## 9.1.2 The OpenStorage operating environment

There are two major components that comprise the OpenStorage operating environment and communicate through a TCP/IP network:

- ▶ The storage server
- ▶ The plug-in

The storage server is an entity that runs on the ProtecTIER servers and uses the major internal functionality of the ProtecTIER platform (such as DHF, clustering, and replication). The plug-in is a shared library (that is, a stateless software component) that resides on the NetBackup machine and is dynamically linked to the NetBackup application for data transfer to the ProtecTIER storage server emulation.

### Installing the ProtecTIER storage appliance

Installing the ProtecTIER storage appliance and the appropriate software is generally done either in manufacturing, or is the responsibility of a trained ProtecTIER specialist. Completing the ProtecTIER system setup tasks for new installations is a customer responsibility. The ProtecTIER storage appliance must meet the prerequisites to install and run the OST plug-in effectively.

## 9.1.3 OpenStorage network configuration

This section describes OpenStorage (OST) network configuration best practices that support independent IP configurations.

Each OpenStorage physical interface on a ProtecTIER node is configured with a corresponding virtual interface with a unique default IP address. During the system configuration, it is necessary to provide a new IP address to every virtual interface, each on a different subnet. This configuration can be modified to configure the port aggregation group. At this point in time, use individual ports rather than an aggregated group of ports.

## Information provided by the customer

It is necessary for the customer to provide different IP addresses, with each of them on a different subnet, for each ProtecTIER node. Different ProtecTIER nodes can have a port on the same subnet. On the first ProtecTIER node, you can configure the IP addresses shown in Example 9-1:

### *Example 9-1 Configured IP addresses*

---

```
192.168.151.1/24
192.168.152.1/24
192.168.153.1/24
192.168.154.1/24
192.168.155.1/24
192.168.156.1/24
```

---

In this case, the second ProtecTIER node can use the addresses shown in Example 9-2.

### *Example 9-2 Second node IP configuration*

---

```
192.168.151.2/24
192.168.152.2/24
192.168.153.2/24
192.168.154.2/24
192.168.155.2/24
192.168.156.2/24
```

---

In this example, the first network is 192.168.151.0, and 255 subnet addresses can be defined. Therefore, the first node is using an address in this subnet (192.168.151.1) and the second node can use a different address on the same subnet (192.68.151.2).

## Configuration of the ProtecTIER system

Each ProtecTIER system is configured by default upon installation with virtual interfaces, each of which contain a single physical interface. The virtual interfaces are configured with the default IP addresses that were assigned during manufacturing. First, it is necessary to modify the system and provide the actual customer host network IP address. Run **ptconfig** with the **applianceInterfaces** option to edit each virtual interface and configure it with a new IP address and netmask. For more information, refer to *IBM System Storage TS7600 with ProtecTIER User's Guide*, GC53-1156 or *IBM System Storage TS7610 ProtecTIER Deduplication Appliance Express User's and Maintenance Guide*, GA32-0779. Do not modify the Load Balancing (LB) method or reassign a physical interface to a different virtual interface. The default LB method and the default physical interface assignments are currently the only configuration supported.

## Configuration of the host

Each interface on the host side also needs to have its own IP address on a separate subnet. As with the ProtecTIER system, different hosts may use the same subnet. Assign IP addresses as necessary on the hosts.

## Routing the IP traffic

Static routes are a simple and effective way of instructing the host IP stack about how to route OST IP traffic destined for specific subnets. This is necessary whenever traffic to any specific subnet is required to 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 port on each ProtecTIER node to which it is connected. If possible, configure all IP addresses on the media servers on the same subnets that you have defined on the ProtecTIER nodes.

## Scheme

In this scheme, we see three media servers connected to a ProtecTIER cluster. Media server 2 has two Ethernet interfaces, one on the purple subnet and one of the red subnet. In this case there is no need to define static routes from media server 2 into ProtecTIER ports on the purple and red subnet, as they are on the same subnet.

## OpenStorage network configuration

All available OST interfaces on one ProtecTIER node are grouped by default in a “bond”, such that one virtual interface with one IP address is exposed for the server. This bond is used for high availability and for load balancing.

The user may modify this configuration and assign physical interfaces to the other virtual interfaces (Figure 9-4). The physical interfaces assigned to one virtual interface are load balanced.

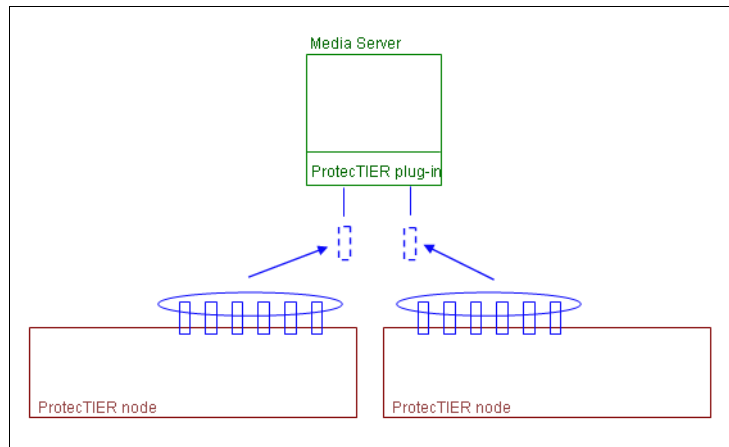


Figure 9-4 OST Network configuration using Virtual IP

Refer to “Chapter 5. IBM System Storage TS7600 with ProtecTIER initial setup” on page 171 for more details about the network configuration.

## 9.2 Configuring OpenStorage with ProtecTIER Manager

The configuration of OST can be separated into three parts:

- ▶ ProtecTIER storage configuration
- ▶ OST plug-in configuration
- ▶ NetBackup configuration

You have to prepare ProtecTIER V2.5 with OST before starting this section. In this section, The ProtecTIER OpenStorage plug-in is discussed.

## 9.2.1 Configuring the storage server

The following section describes how to configure the ProtecTIER system for use with OpenStorage using ProtecTIER Manager. Before you begin configuring the ProtecTIER system, make sure that a repository has been created with OpenStorage as the backup interface. Storage sizing for OST is no different than with a VTL machine.

**Note:** You can verify, during repository creation, that OST has been used by reviewing the Summary report of the Create repository wizard. The Backup interface field should appear as OST. However, this action is only available on the Gateway model; for the Appliance model, the repository is created at the manufacturing site.

After a repository has been created, use ProtecTIER Manager to create, configure, and monitor the storage server. The storage server (STS) is a high-level component defined by the OST API. In simple terms, it is a “container” of logical storage units (LSUs) and provides the means to access the logical storage units and their contents. Currently, only one STS can be defined for each OST storage appliance.

### Adding the storage server

This section describes how to add and configure a storage server. Complete the following steps to add the STS and then define the logon credentials. NetBackup uses these credentials so that the media server can log in to the storage server for storage access. (Refer to the relevant Symantec NetBackup OpenStorage documentation for more information about OpenStorage server credentials.)

1. Log in to ProtecTIER Manager.
2. From the OpenStorage menu, select **Storage server** → **Add Storage server**. The STS configuration window opens (Figure 9-5).




Figure 9-5 Add Storage Server

3. Enter a unique name for the new storage server in the STS Name field.

**Note:** The STS unique name cannot be modified after the STS has been created. Refer to Table 9-2 on page 440 for the field specifics when defining the logon credentials.

4. Enter a user name in the User name field.

5. Enter a password in the New Password field and reenter the password in the Verify password field, as shown in Figure 9-6.

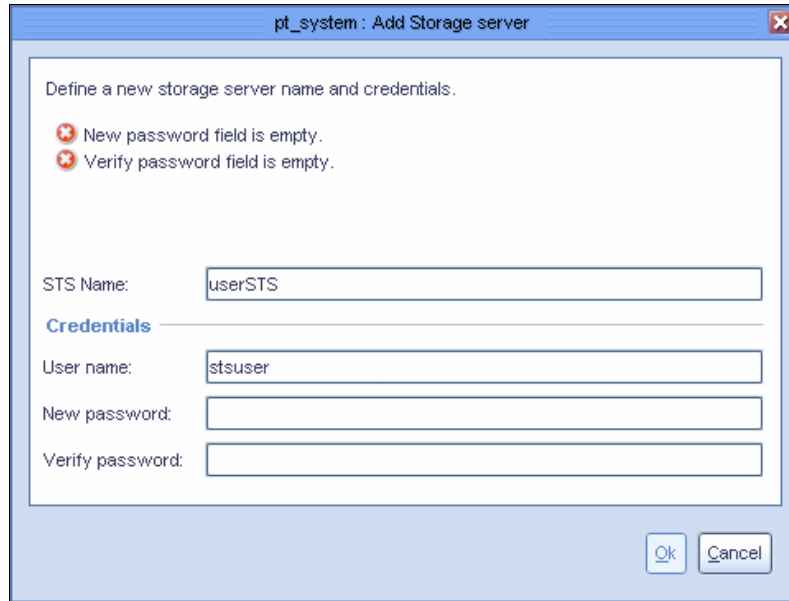


Figure 9-6 Add Storage server window

6. Click **Ok** to complete the task.

**Note:** Refer to Table 9-2 for the field specifics when defining the logon credentials.

Table 9-2 Logon credentials field specifics

Field name	Field length	Field type (alphanumeric or special characters)
STS Name	4 - 16 characters	alphanumeric, period(.), or underscore(_)
User Name	4 - 16 characters	ASCII
New Password	4 - 16 characters	All alphanumeric and special characters. The use of special characters may be unrecognized by NBU. Therefore, if special characters are used (for example, space, &, -, and others) use quotation marks (" ") before and after the string.

### Modifying the storage server credentials

This section describes how to modify the credentials of the selected storage server using ProtecTIER Manager. The STS unique name cannot be modified after the STS has been created. Complete the following steps:

1. Access ProtecTIER Manager.

- From the OpenStorage menu, select **Storage server** → **Modify Storage server credentials**. The STS modify credentials window opens (Figure 9-7).




Figure 9-7 Modify Storage server credential menu

- If you are modifying the user name, enter a new user name in the User name field.
- If you are modifying the password, enter the password that was previously defined in the Password field.
- Enter a new password in the New password field and reenter the password in the Verify password field.
- Click **Ok** to complete the task. (Figure 9-8).

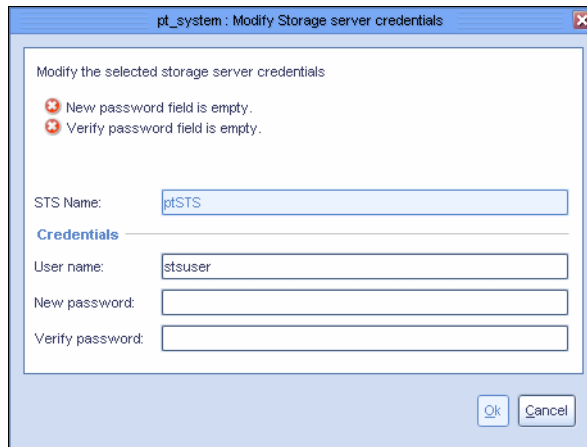


Figure 9-8 Modify Storage server credentials window

The STS credentials are modified on the ProtecTIER system and can be monitored through ProtecTIER Manager.

## 9.2.2 Configuring a logical storage unit

The following section describes how to configure a logical storage unit (LSU) on an STS.

Configuring logical storage units, or LSUs, on a storage server divides the appliance into one or more logical units of space. An LSU, like an STS, is also defined by the OST API and is a "container" of storage and images (which consume storage). Up to 256 LSUs can be defined per STS and are identified by a name string that is unique within the storage server.

In ProtecTIER for OST, an LSU's storage properties are defined in nominal terms. This means that the LSU's storage capacity is defined as a nominal percentage of the repository's overall nominal capacity, taking into consideration the configured percentages of the other LSUs that all share the repository's physical storage.

**Note:** The overall capacity of all LSUs together can be less than or equal to the repository's nominal capacity.

## Adding a logical storage unit

Complete the following steps to add an LSU to an STS:

1. From ProtecTIER Manager, select **OST** → **Logical Storage Unit management** (Figure 9-9). The Logical Storage Unit management window opens (Figure 9-10).



Figure 9-9 Logical Storage Unit Management Menu

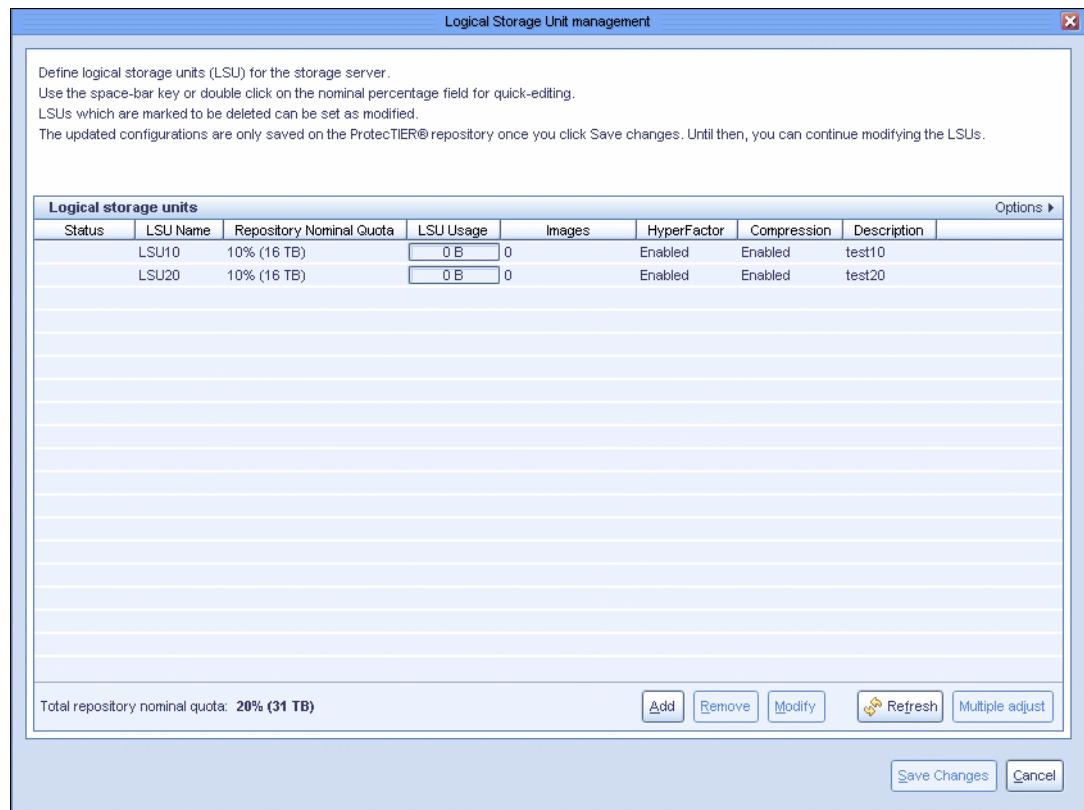


Figure 9-10 Logical Storage Unit management window



2. Click **Add** to configure a new LSU on the storage server. The Adding an LSU window opens. Refer to Table 9-3 for the field specifics when defining an LSU.

Table 9-3 LSU field specifics

Field name	Field length	Filed type (alphanumeric or special characters)
LSU Name	4 - 6 characters	alphanumeric, period (.), or underscore (_)
Description	0 - 1023 characters	All alphanumeric and special characters

3. Enter a unique name in the LSU Name field.

**Note:** After an LSU has been created and saved, the LSU name can no longer be modified.

4. Enter a description to describe the LSU in the Description field.
5. Enter the size of the LSU as the nominal percentage (up to two precision digits) of the repository in the Nominal percentage field.
6. Select the data factoring mode from the HyperFactor drop-down menu:
  - **Enabled** to turn on hyperfactoring
  - **Disabled** to turn off hyperfactoring
  - **Baseline** for no deduplication, but an index is built
7. Select either **On** or **Off** in the Compression field to enable or disable data compression.
8. Click **Ok** to add and save the LSU properties.

## Modifying a logical storage unit configuration

Complete the following steps to modify an LSU on an STS:

1. From ProtecTIER Manager, select **OST** → **Logical Storage Unit management**. The Logical Storage Unit management window opens (Figure 9-11).

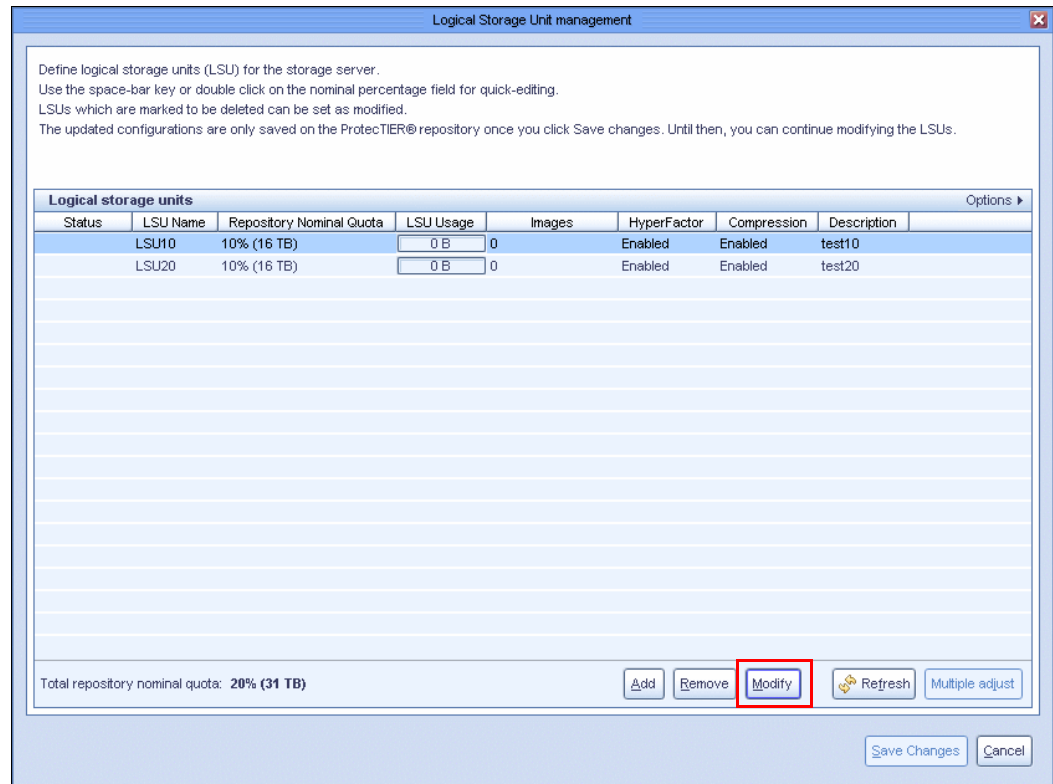


Figure 9-11 Logical Storage Unit Management window

2. Click **Modify** to change the LSU properties previously defined on the storage server. The Modifying an LSU window opens. Modify either one or all of the following items:

**Note:** After an LSU has been created and saved, the LSU name can no longer be modified.

- Enter a new description (up to 1023 characters) to describe the LSU in the Description field.
- Change the size of the LSU as the nominal percentage (up to two points of precision) of the repository in the Nominal percentage field.
- Modify the factoring mode option from the HyperFactor drop-down list by selecting one of the following options:
  - **Yes** to enable hyperfactoring
  - **No** to disable hyperfactoring
  - **Baseline** for no deduplication, but an index is built

- Modify the data compression option by selecting either **On** or **Off** in the Compression field to enable or disable data compression (Figure 9-12).

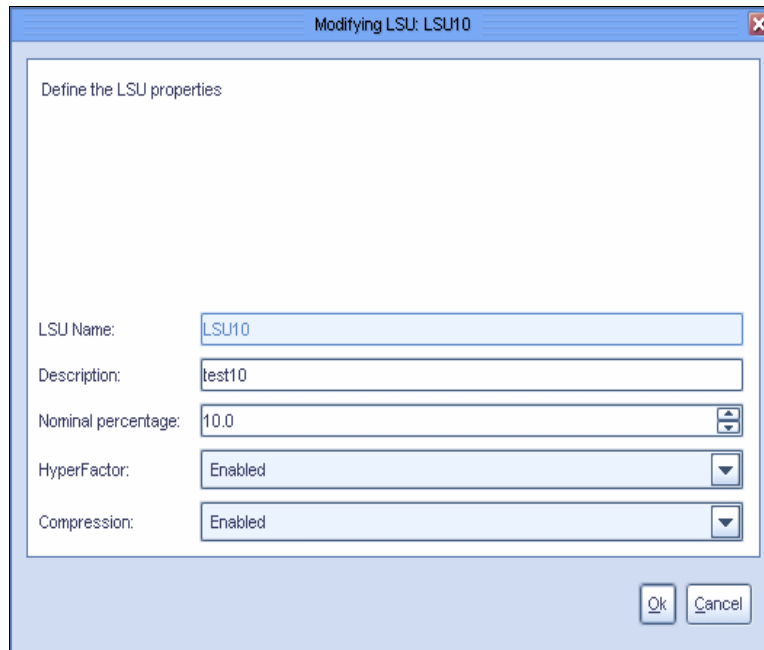


Figure 9-12 Logical Storage Unit Modify window

3. Click **Ok** to save the LSU properties.

### Managing the logical storage unit configuration

The following section describes how to manage the LSU configuration using ProtecTIER Manager.

The Logical Storage Unit management window contains the LSU configuration properties. Use the spacebar, or double-click the **Repository Nominal Quota** or **LSU usage** fields for “quick-editing” of the LSU storage capacity, and complete the following steps:

1. Select **OpenStorage** → **Logical Storage Unit management**. The Logical Storage Unit management window contains the following information (Table 9-4).

Table 9-4 Logical Storage Management window

Column	Description
Status	Displays if the LSU is new, or if it has been modified or deleted.
LSU Name	The unique name defined for the LSU.
Repository Nominal quota	The LSU's nominal percentage used of the repository's overall nominal capacity.
LSU usage	The actual amount of space (in MB) used from the LSU's nominal percentage.
Images	The number of images contained on an LSU.
Hyperfactor	The data factoring mode status.
Compression	Data compression is either enabled (on) or disabled (off).
Description	Describes the LSU.

2. Click **Refresh usage** to update the LSU usage column.
3. Click **Multiple adjust** to equally divide the repository nominal quota between all of the selected LSUs (that is, Available percentage / Number of selected LSUs), or set all of the selected LSUs with a selected percentage value.
4. Click **Save Changes** to save the updated configurations, or continue modifying the LSUs.

## 9.3 The ProtecTIER OpenStorage plug-in

This section describes the ProtecTIER OpenStorage plug-in and how to install and configure the plug-in as a NetBackup media server.

The ProtecTIER OpenStorage plug-in software is installed on media servers that access the ProtecTIER system. The plug-in allows the control of multiple storage servers over various IP network subnets. After the plug-in is installed, the storage server is configured in NetBackup.

### 9.3.1 Installing the OpenStorage plug-in on a NetBackup media server

This section describes the procedure for installing the ProtecTIER OpenStorage plug-in on the media server to communicate with the storage server. The OST plug-in has not supported a Linux media server until now.

#### Installing the OpenStorage plug-in on Windows

Complete following steps to install the OpenStorage plug-in on Windows:

1. From the IBM System Storage ProtecTIER OpenStorage Plug-in CD, access and launch the Windows plug-in installer by double-clicking the executable file. The InstallAnywhere wizard will guide you through the installation of the IBM ProtecTIER OST Plug-in for 32- or 64-bit.

2. Click **Next**. The installer prompts you to stop the NetBackup Remote Manager and Monitor Service if it is running (see Figure 9-13). To stop this service, select **Start** → **Run** and enter `services.msc` to stop the service and restart the installation, as shown in (Figure 9-14).

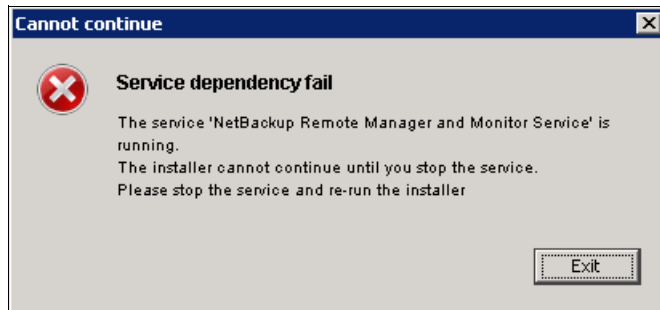


Figure 9-13 Message about NetBackup Service

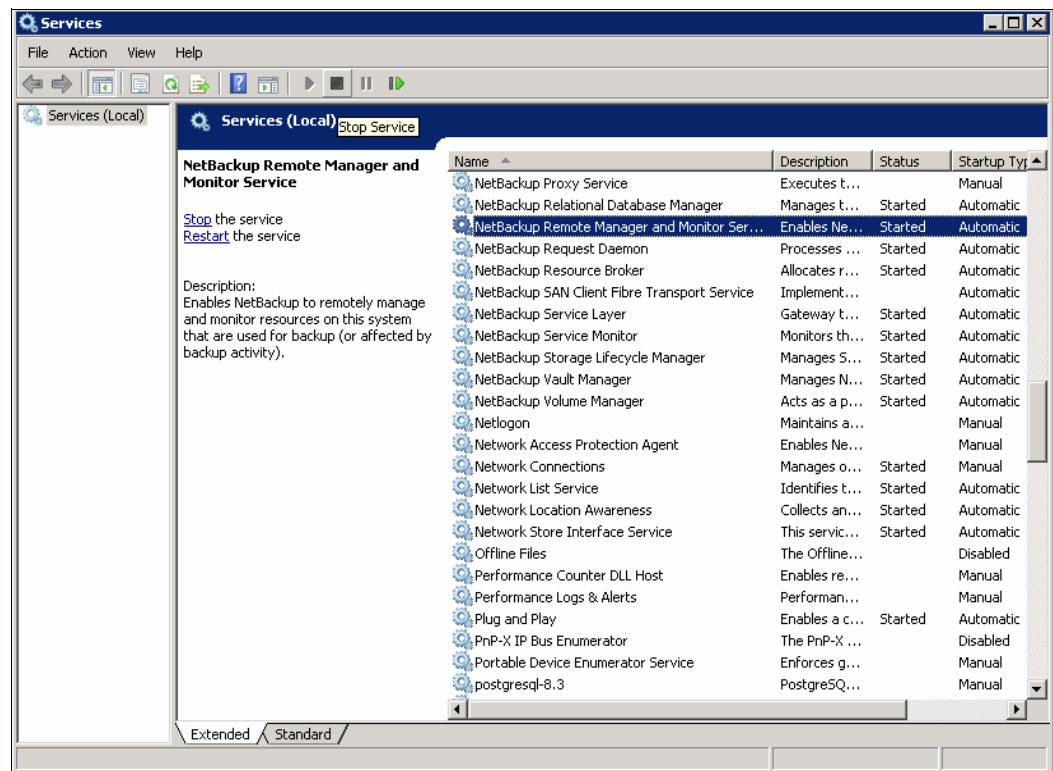


Figure 9-14 Stop the Netbackup Remote manager and Monitor service

3. Review the Pre-Installation Summary. Click **Install** to run the installation procedure.
4. When the installation is complete, you will be prompted to start the Netbackup Remote Manager and Monitor Service. Select **Start service** and click **Next**.
5. The OpenStorage plug-in installation is complete. Press **Done** to exit the installer.

### Installing the OpenStorage plug-in on an AIX server

Complete the following steps to install the OpenStorage plug-in on AIX:

1. Access the command-line interface on the media server.

2. Log in as root.
3. Run the rpm in install mode from the command line on the media server. The file name is IBM\_ProtectTIER\_OST\_plugin\_AIX64-1.1-2.aix6.1.ppc.rpm for 64-bit and IBM\_ProtectTIER\_OST\_plugin\_AIX32-1.1-2.aix6.1.ppc.rpm for 32-bit.
4. Enter the respective file name and press Enter. The following output is displayed:

```
[root@server tmp]# rpm -i IBM_ProtectTIER_OST_plugin_AIX64-1.1-2.aix6.1.ppc.rpm
Creating default configuration files...
Created config unit /etc/IBM/ost-plugin/plugin_network_cfg.xml
Created config unit /etc/IBM/ost-plugin/plugin_cfg.xml
done
```

### 9.3.2 Configuring the OpenStorage plug-in on a NetBackup media server

This section explains how to add the storage server credentials to the NetBackup media servers.

Complete the following steps using the `ostp_cli` executable to configure the server IP addresses on the plug-in, which were previously configured on the ProtectTIER server:

1. From the `/opt/IBM/ost_plugin_tools` directory, run `ostp_cli` to configure each of the server IP addresses on the ProtectTIER system.

**Note:** From Windows, the path to `ostp_cli` is `PROGRAMFILES\IBM\ost_plugin_tools\ostp_cli`, where `PROGRAMFILES` is the path to the program files directory, usually `C:\Program Files`. For a Windows 2008 installation, the plug-in must be run using the Administrator account. Select **Start** → **Programs** → **Accessories** → **Command Prompt**, right-click **Command Prompt**, and select **Run as Administrator**, as shown in Figure 9-15.

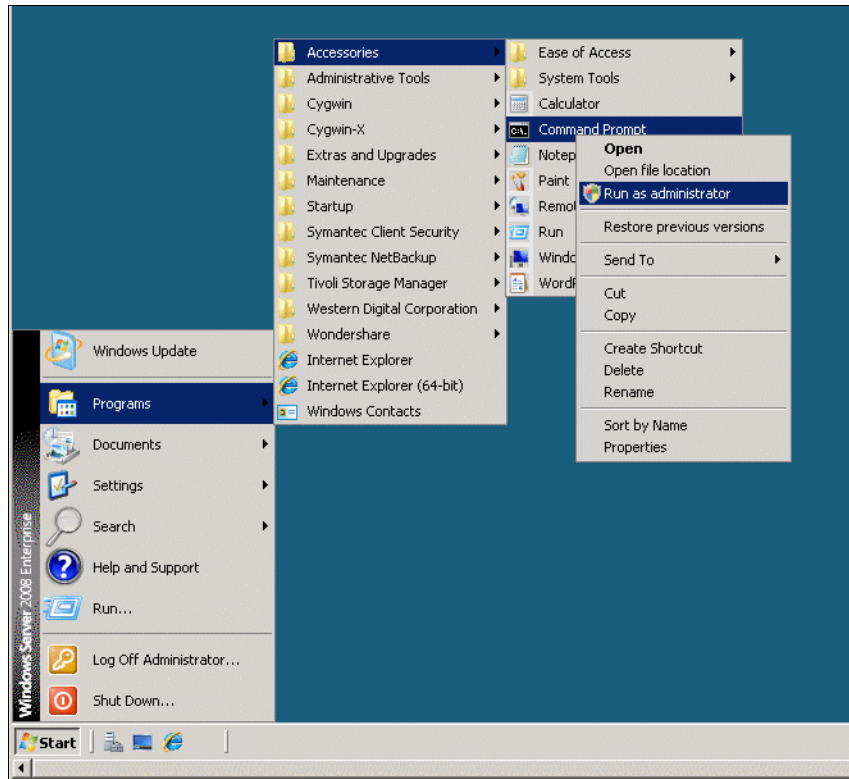


Figure 9-15 Start command shell from Windows server

2. Enter the `ostp_cli` command to add new IP addresses to the server's configuration as follows:

```
ostp_cli net_cfg_add_addr <sts_name> <address_set_index> <address:port>
```

Where:

- `<sts_name>` is the storage server name created on ProtecTIER Manager.
- `<address_set_index>` is a zero-based index of an address set. An address set is a collection, or group, of IP addresses. Addresses are grouped by ProtecTIER cluster node, that is, all OpenStorage addresses belonging to the same node in a ProtecTIER cluster should go to the same address set. Thus, there are two address sets for each `sts_name` with indexes "0" and "1".
- `<address:port>` is the OST IP address on the ProtecTIER Manager server and the port number. The default OST port number is 23979 (which is predefined on the system), as shown in Example 9-3.

*Example 9-3 OST plug-in command line setup on Windows media server*

```
C:\Program Files\ibm\ost_plugin_tools>ostp_cli net_cfg_add_addr ptSTS 0 192.168.
```

```

151.1:23979
Successfully added address

C:\Program Files\ibm\ost_plugin_tools>ostp_cli net_cfg_show
STS: ptSTS
address-set 0:
address 0: 192.168.151.1:23979

```

**Note:** The command should be run individually on each virtual IP address configured on the server side.

## Command options

Table 9-5 contains a list of **ostp\_cli** command options for configuring the plug-in on the media server.

**Tip:** If one or more of the configured IPs on the ProtecTIER server has failed, or is extremely slow, configuration commands, or backup, restore, or duplication jobs may fail on the AIX plug-ins. To identify the problematic IP(s), run the **ostp\_cli check\_comm** command to verify if one or more of the IPs are unreachable. If this is the case, complete the following steps:

1. Check that the cable is properly connected on both sides. If not, connect it properly.
2. For the short term, to continue working, remove this IP from the list of IPs configured to this server by running the following command:

```
ostp_cli net_cfg_remove_addr <sts_name> <address_set_index> <address_index>
```

3. For the long term, identify and repair the corrupted NIC, port, or cable.

Table 9-5 OST Plug-in configuration command options

Command	Parameters	Description
net_cfg_add_addr <sts_name> <address_set_index> <address:port>	<ul style="list-style-type: none"> <li>▶ &lt;sts_name&gt; is the storage server name.</li> <li>▶ &lt;address_set_index&gt; is a zero-based index of an address set.</li> <li>▶ &lt;address:port&gt; is the OST IP address on the ProtecTIER Manager server and the port number.</li> </ul>	Adds new IP addresses to the server's configuration.
net_cfg_remove_addr_set <sts_name> <address_set_index>	<ul style="list-style-type: none"> <li>▶ &lt;sts_name&gt; is the storage server name.</li> <li>▶ &lt;address_set_index&gt; is a collection of nodes or their IP addresses (0 or 1).</li> </ul>	Removes an address set from the addresses configuration. The last address set in an STS cannot be removed. Instead, remove the STS.
net_cfg_remove_addr <sts_name> <address_set_index> <address_index>	<ul style="list-style-type: none"> <li>▶ &lt;sts_name&gt; is the storage server name.</li> <li>▶ &lt;address_set_index&gt; is a collection of nodes or their IP addresses (0 or 1).</li> <li>▶ &lt;address_index&gt; is the OST IP address.</li> </ul>	Removes an IP address from the addresses configuration. The last IP address in an address set cannot be removed. Instead, remove the address set.



Command	Parameters	Description
net_cfg_modify_addr <sts_name> <address_set_index> <address_index> <new_address:port>	<ul style="list-style-type: none"> <li>▶ &lt;sts_name&gt; is the storage server name.</li> <li>▶ &lt;address_set_index&gt; is a collection of nodes or their IP addresses (0 or 1).</li> <li>▶ &lt;address_index&gt; is the OST IP address.</li> <li>▶ &lt;new_address:port&gt; is the modified IP address and port number of the node.</li> </ul>	Modifies an IP address in the addresses configuration.
net_cfg_rename_sts <sts_name> <new_sts_name>	<ul style="list-style-type: none"> <li>▶ &lt;sts_name&gt; is the storage server name.</li> <li>▶ &lt;new_sts_name&gt; is the modified storage server name.</li> </ul>	Renames an STS in the addresses configuration.
net_cfg_remove_sts <sts_name>	<sts_name> is the storage server name.	Removes an STS from the addresses configuration.
check_com	N/A.	Verifies if one or more of the IPs are unreachable.
net_cfg_show	N/A.	Displays the addresses' configuration.
cfg_show	N/A.	Displays the plug-in's general configuration.
cfg_change <param_name> <value>	<ul style="list-style-type: none"> <li>▶ &lt;param_name&gt; is the general plug-in configuration parameter.</li> <li>▶ &lt;value&gt; is the size of the respective parameter.</li> </ul>	Used to modify the plug-in's general configuration.
get_cfg_path	N/A.	This command is <i>not</i> for customer use.
get_log_path	N/A.	This command is <i>not</i> for customer use.
create_paths	N/A.	This command is <i>not</i> for customer use.
net_cfg_create	N/A.	This command is <i>not</i> for customer use.
cfg_create [create_mode]	N/A.	This command is <i>not</i> for customer use.

## General plug-in configuration

Table 9-6 contains a partial list of commonly-used parameters for modifying the plug-in's general configuration.

Table 9-6 General plug-in options

Parameter	Value	Description
total-archive-size-mb	200 - 1000 MB	The total space (in MB) dedicated to the logs and long-term statistics files.
target-archive-sizepercent	20 - 80%	Log target archive size. After the log size percentage is reached, delete the contents until the size is a fraction of the original size.
compression-enabled	0 = off, 1 = on	Enables or disables compression of non-active log files and non-active long-term statistics files
.log-verbosity	low, medium, or high	How much information to log.

Run the following command:

```
ostp_cli cfg_change <param_name> <value>
```

Now that you have finished configuring the ProtecTIER OST plug-in on the media server, refer to 9.4, “Configuring ProtecTIER OpenStorage on NetBackup” on page 452 in this document or the relevant Symantec NetBackup OpenStorage documentation for information about verifying that the plug-in is recognized by the NetBackup application and configuring NetBackup to work with the STS and its LSUs.

## 9.4 Configuring ProtecTIER OpenStorage on NetBackup

In this section, we describe the NetBackup software configuration. NetBackup is a Symantec product, so we only describe the OST configuration to use ProtecTIER. If you need more information about NetBackup, you can refer to the Symantec NetBackup OpenStorage documents.

Before you start the configuration, complete the following steps:

1. You must configure the hardware (set up the OST network connection)
2. You must install and configure ProtecTIER 2.5 with OST.
3. You must install and configure the OST plug-in on the NetBackup Media Server.

On Windows and UNIX servers, you can use the Graphic User Interface (GUI). In this section, we describe how to configure the storage server, storage unit, and disk pool using NetBackup GUI tools.

**Note:** We use Symantec NetBackup Software Version 7.0.

### Storage server settings

A storage server is defined on the ProtecTIER V2.5 system and in NetBackup Windows Server plug-in. Input the name of storage server and its parameters.

Figure 9-16 shows the NetBackup Administrator Console. Select **Master Server** from the left pane to access the StorageServer setting. Then complete the following steps:

1. Select the **Configure Disk Storage Servers** icon.

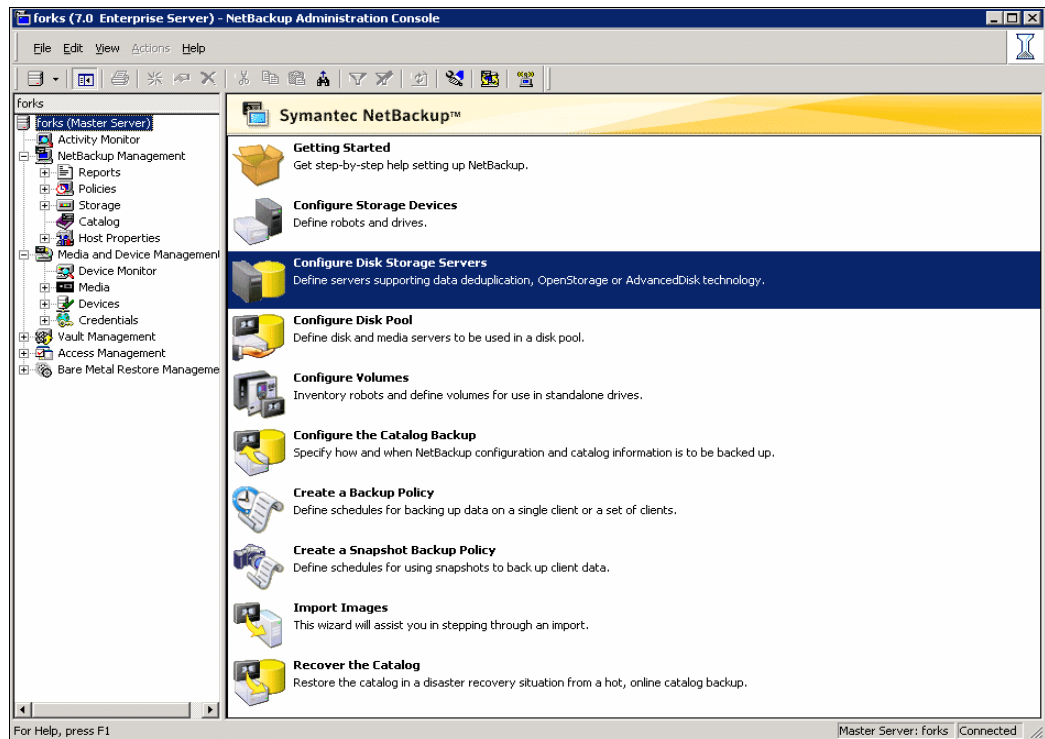


Figure 9-16 NetBackup Administrator Console window

2. Start the Storage Server configuration wizard shown in Figure 9-17 and click **Next**.

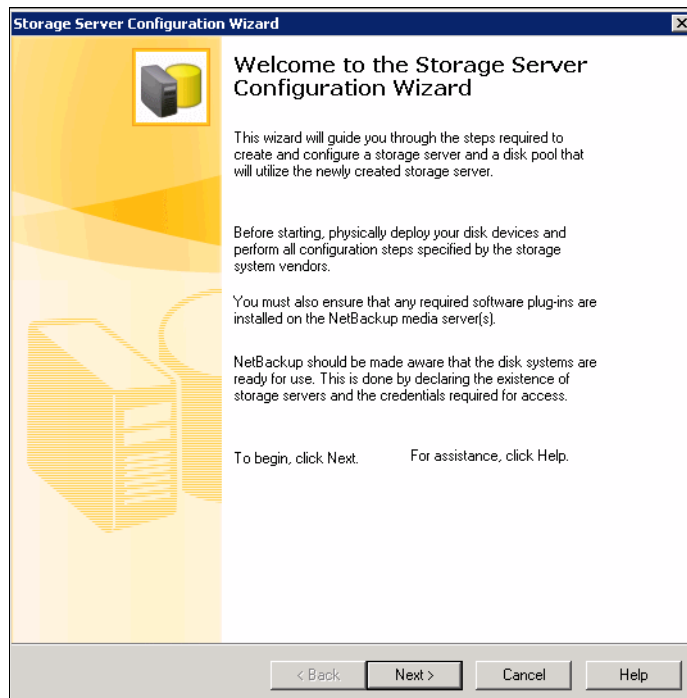


Figure 9-17 Start the Storage Server Configuration Wizard

3. For the Storage Type definition, select **OpenStorage** (Figure 9-18) and click **Next**.

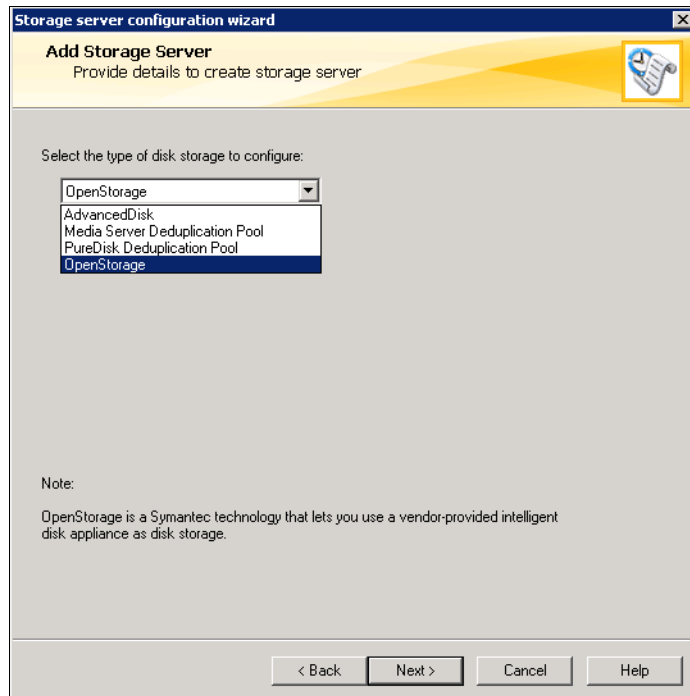


Figure 9-18 Storage type selection window

4. Input the following parameters into the fields shown in Figure 9-19:
  - Storage server type: <Vender type> input only: IBM-ProtectTIER
  - Storage server name : <STS name>
  - Select media server: <server name> Windows media server name
  - User name : <sts user name>
  - Password : <sts user password>

**Note:** STS name, media server, sts user name, and sts user password are predefined in ProtecTIER, as described in 9.2.1, “Configuring the storage server” on page 439.

Storage Server Configuration Wizard

**Add Storage Server**  
Provide details to create storage server

Storage server type: IBM-ProtectTIER

Storage server name: ptSTS

Select media server\*: forks

Enter credentials:

User name: stuser

Password: xxxxxxxx

Confirm password: xxxxxxxx

\* Select a media server that has the vendor's OpenStorage plug-in installed. NetBackup will query the storage server for its capabilities by sending the probe through the media server you specify.

< Back   Next >   Cancel   Help

Figure 9-19 Storage server parameters input window

5. Figure 9-20 shows a summary of the settings. Review them, and if you are satisfied, click **Next**; otherwise, click **Back** and make any necessary changes.

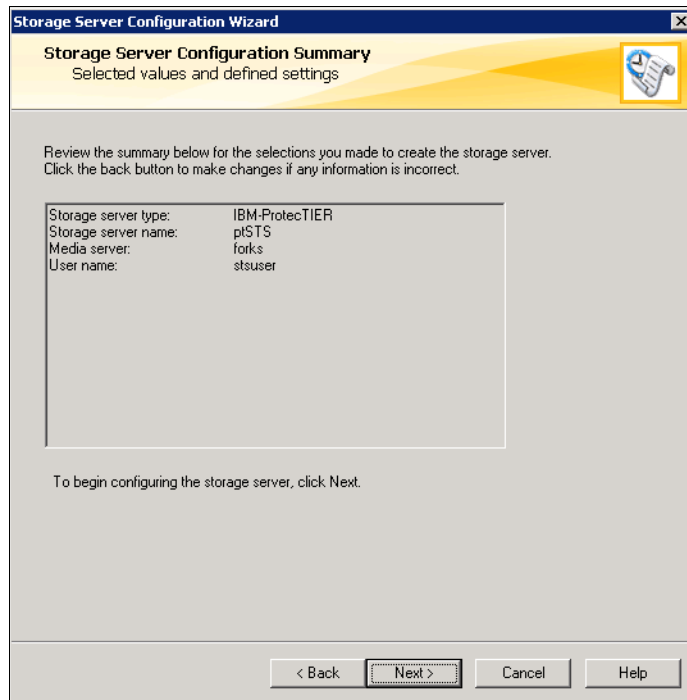


Figure 9-20 Storage server setting input parameter display window

6. Verify and check the parameters and configuration (Figure 9-21) and the result of verification (Figure 9-22).

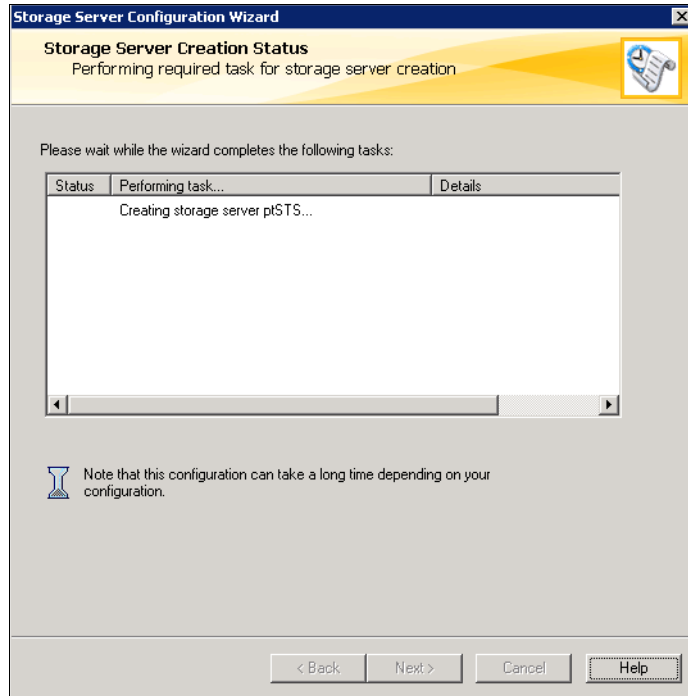


Figure 9-21 Input parameter validation and verification window

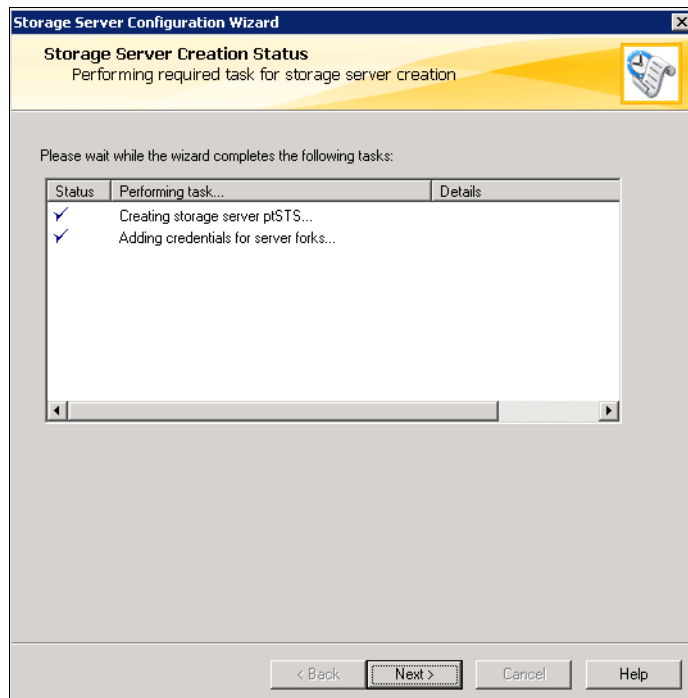


Figure 9-22 The results of verification

7. Complete the creation of STS by clicking **Next** (Figure 9-23).

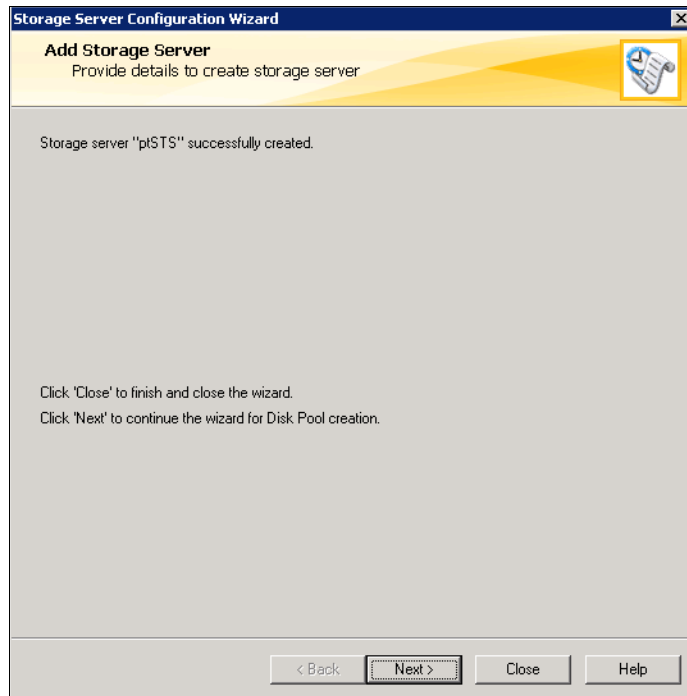


Figure 9-23 Complete the creation of STS

### Creating a disk pool

Symantec NetBackup has to have *storage unit* to back up data. You can create a *disk pool* before creating a storage unit. You can use the disk pool and the storage that is defined on NetBackup as the target of a backup policy. Therefore, if you want to back up to the ProtecTIER system, you have to create the disk pool as the STS parameter. However, you need to license a NetBackup feature to use a disk pool.



To create disk pool, complete the following steps:

1. Click the **Configure disk pool** icon from main window to open the window shown in Figure 9-24. Choose **IBM-ProtectTIER** and click **Next**.

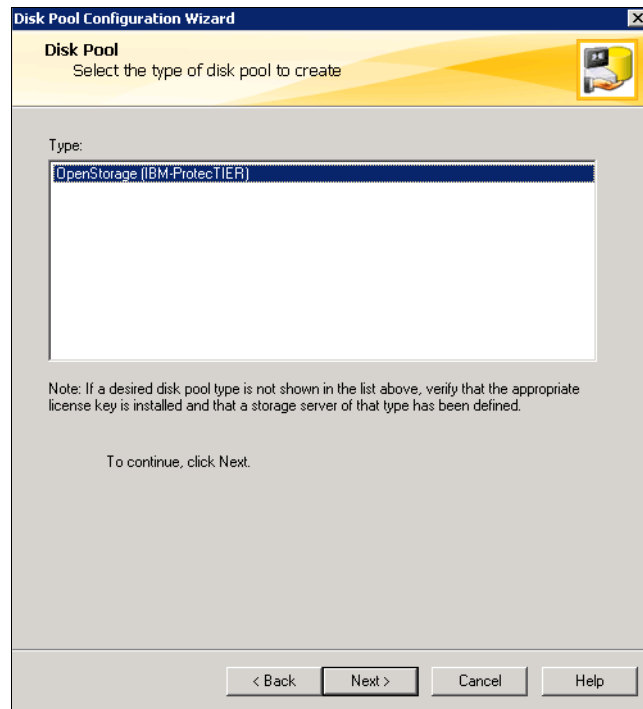


Figure 9-24 Select the storage type (select vendor)

- In the Select Storage Server window (Figure 9-25), you can choose the storage server that is predefined in the NetBackup software. Select the STS and then click **Next**.

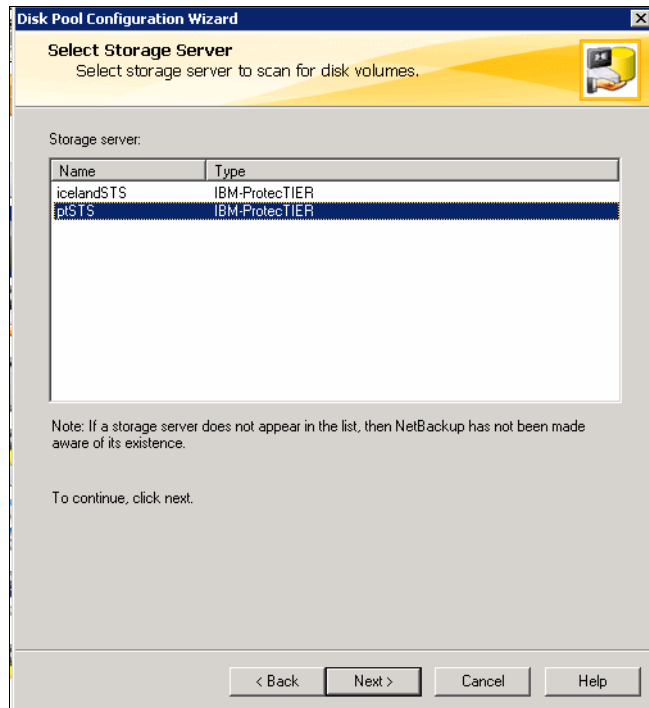


Figure 9-25 Select Storage Server window

- In the Select Volumes window, (Figure 9-26), you define the LSU in the ProtecTIER server. Select the LSUs to use for backup storage by selecting the LSUs' check boxes and then click **Next**.

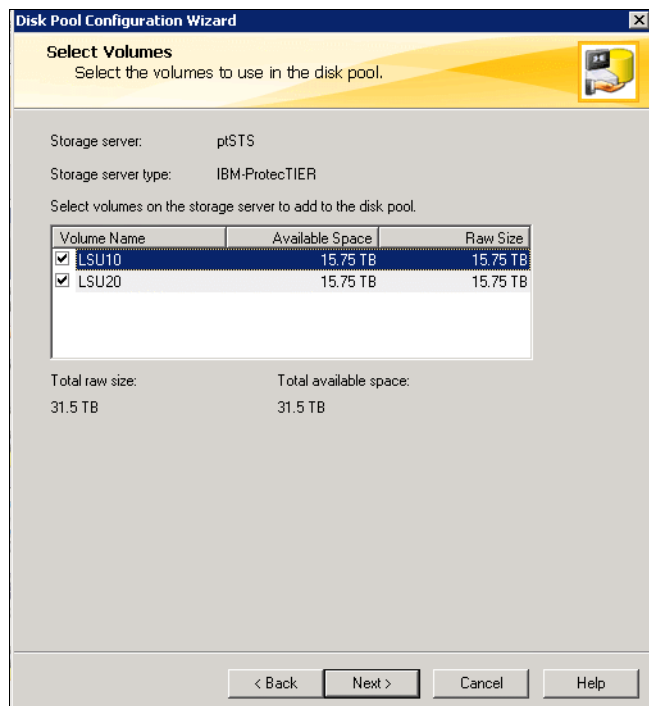


Figure 9-26 Select Volumes window

4. Define the disk pool name by entering the name of disk pool and clicking **Next** (Figure 9-27).

**Note:** The High water mark setting is a threshold that, when reached, signals NetBackup that the disk storage unit should be considered full. The default is 98%. As the disk storage capacity grows closer to the High water mark, NetBackup reduces the number of jobs sent to the storage unit. NetBackup does not assign new jobs to a storage unit that is considered full.

After the High water mark is reached, space is created on the disk storage unit until the Low water mark is met. The default is 80%. The Low water mark cannot be higher than High water mark. Refer to *Symantec NetBackup Administrator's Guide* for more details. This guide can be found at the following address:

<http://www.symantec.com/business/support/index?page=content&id=TECH52803>

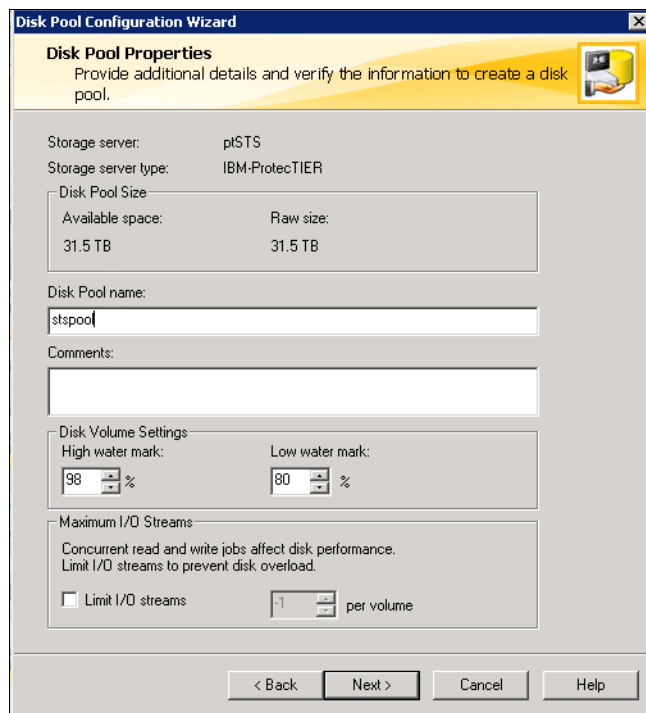


Figure 9-27 Disk Pool Properties window

5. Verify that the disk pool configuration is correct and then click **Next** (Figure 9-28).

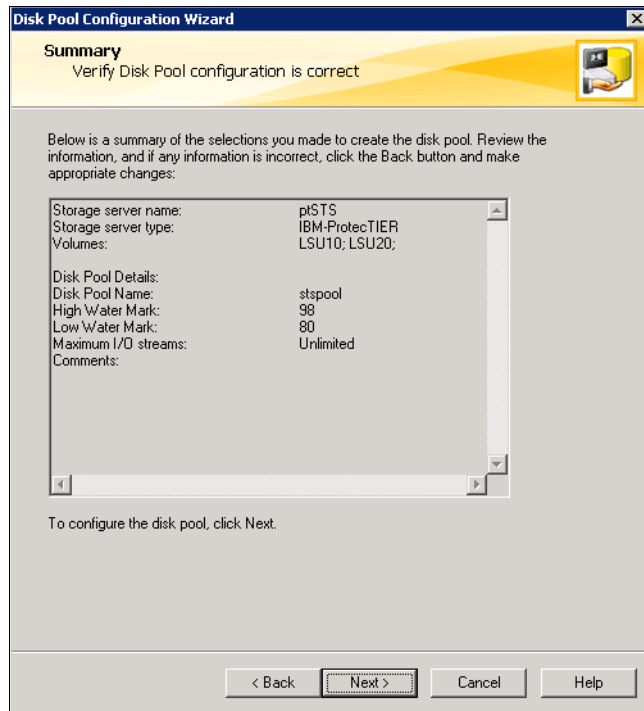


Figure 9-28 Summary window for disk pool configuration

6. Complete the disk pool creation by clicking **Next** (Figure 9-29).

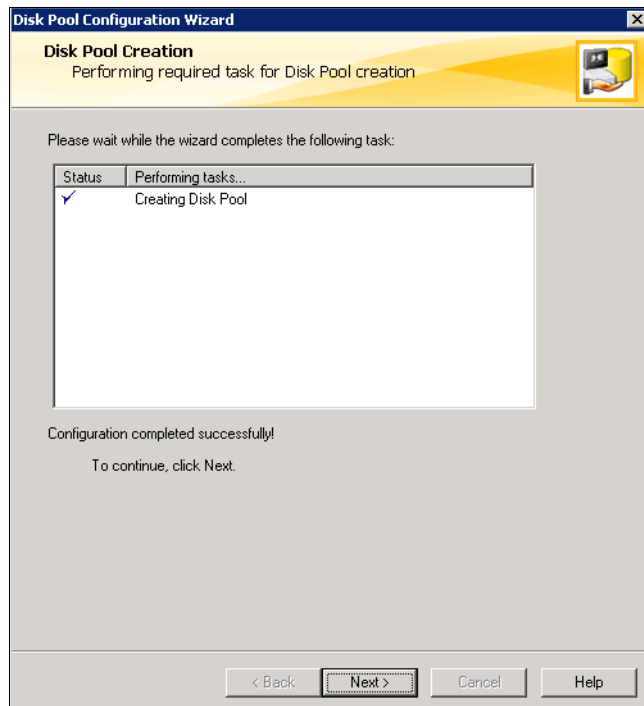


Figure 9-29 Disk pool creation complete

7. To go on to the next task, the creation of the storage unit, click **Next** (Figure 9-30).

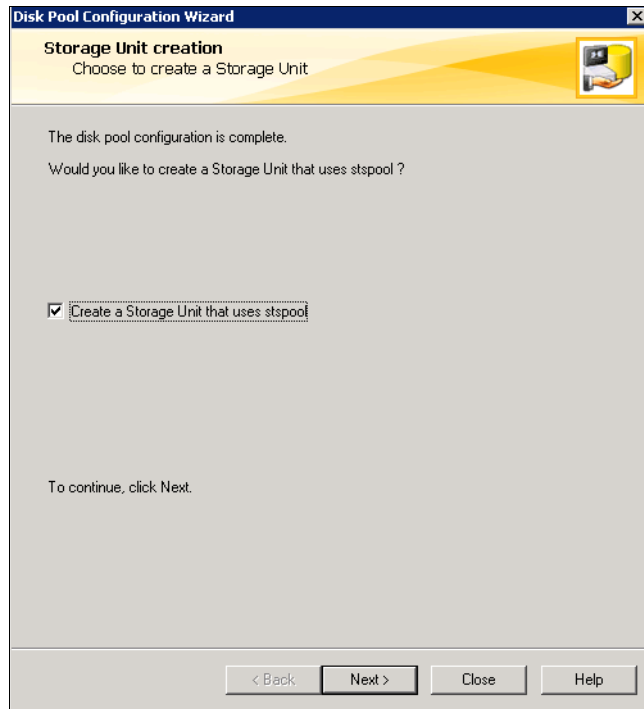


Figure 9-30 Continuing to the creation of the storage unit.

### Creating a storage unit

A storage unit is a label that NetBackup associates with physical storage. The label can identify a robot, a path to a volume, or a disk pool. In this section, we use a disk pool as our storage unit.

To create a storage unit, complete the following steps:

1. If you clicked Next in step 7 on page 463, the window shown in Figure 9-31 should be open; otherwise, open the window from the NetBackup Administration Console tree.

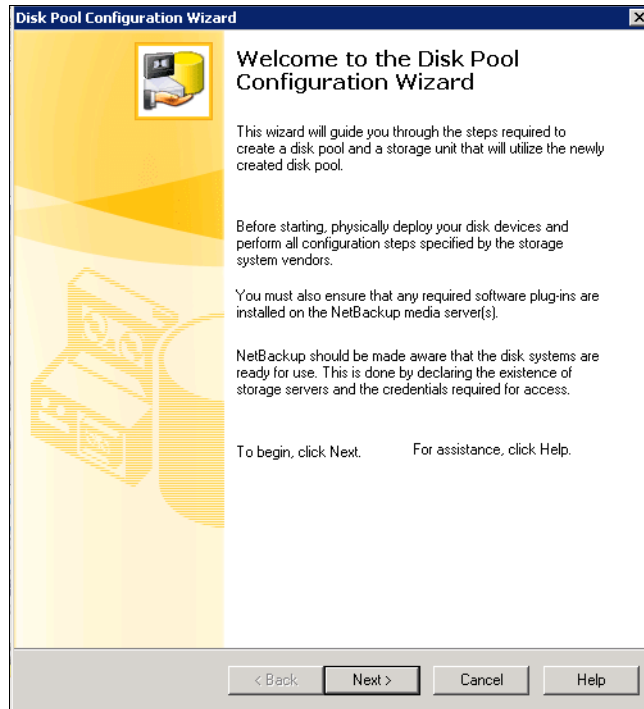


Figure 9-31 Start the creation of a storage unit

2. Input the storage unit's name and click **Next** (Figure 9-32).

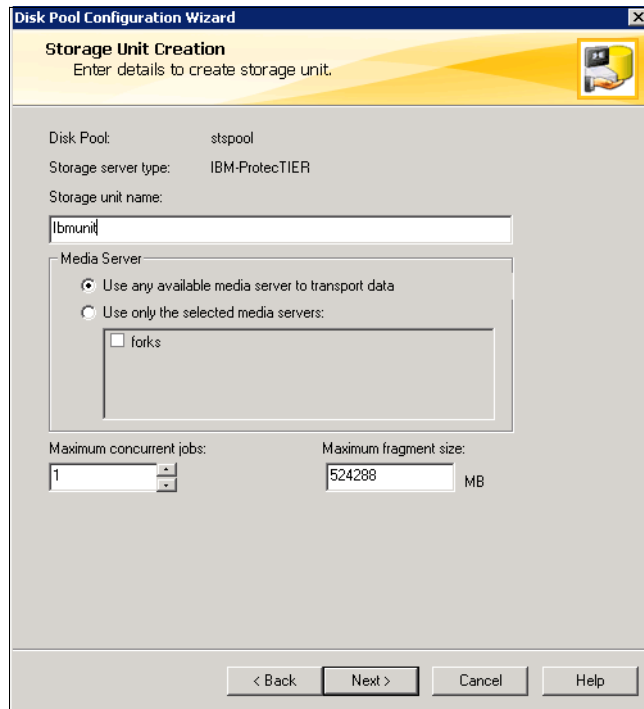


Figure 9-32 Storage unit creation window

**Note:** The Maximum concurrent jobs setting specifies the maximum number of jobs that NetBackup can send to a disk storage unit at one time. (The default is one. The job count can range from 0 to 256.). The Maximum concurrent jobs setting can be used to balance the load between disk storage units. The Maximum concurrent jobs setting depends on the available disk space and the server's ability to run multiple backup processes. Ask your system administrator or engineer what the proper value of this field is.

The default Maximum fragment size for a disk storage unit is 524,288 MB. To specify a Maximum fragment size other than the default, enter a value from 20 MB to 524,288 MB. Backups to disk are usually fragmented to ensure that the backup does not exceed the maximum size that the file system allows. For media server deduplication pools and PureDisk deduplication pools, Symantec recommends against specifying the largest fragment size allowed (512 GB). For best results, the default fragment size is set to 50 GB.

For more information about this topic, refer to the *Symantec NetBackup Administrator's Guide*, found at the following address:

<http://www.symantec.com/business/support/index?page=content&id=TECH52803>

3. Complete the storage unit creation (Figure 9-33).

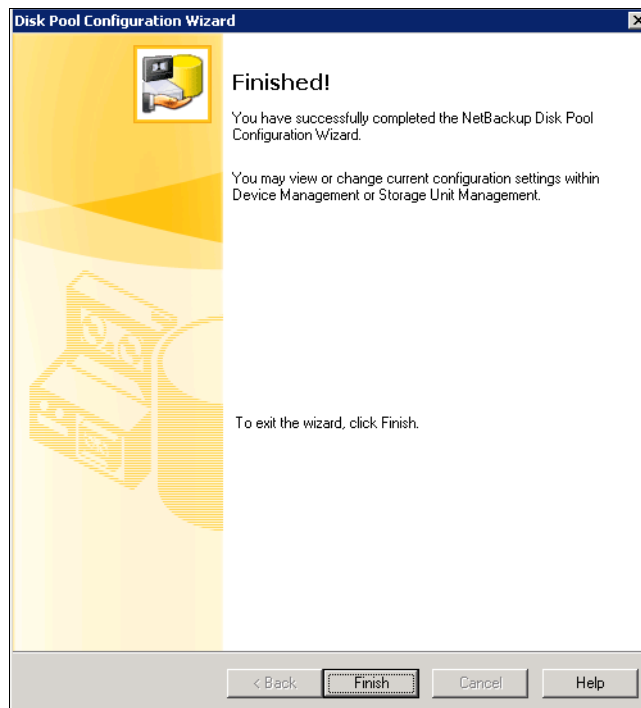


Figure 9-33 Create storage unit procedure completed

All the procedures to configure OST with NetBackup are complete, but you have to configure the backup policy and backup schedule to back up and restore data.

## 9.5 Replication settings for OpenStorage

This section describes the OST replication settings. For information about the replication initial setup, refer to 5.8, “Setting up native replication” on page 252.

### 9.5.1 Working with application groups in OpenStorage

This task explains how to add repositories to an application group in a grid running on an OpenStorage-enabled system.

Before any replication policies can be defined, a physical connection must be created between the repositories. The physical connections are created by defining application groups.

**Note:** Keep in mind when defining a connection in OpenStorage that the replication is bidirectional. A maximum of 12 repositories (members) can replicate to each other at any given time with up to 256 repositories existing in the entire grid.

#### Creating an application group

To create an application group, complete the following steps:

1. Log in to the Grid Manager and select **Repository** → **Application group management** (Figure 9-34).

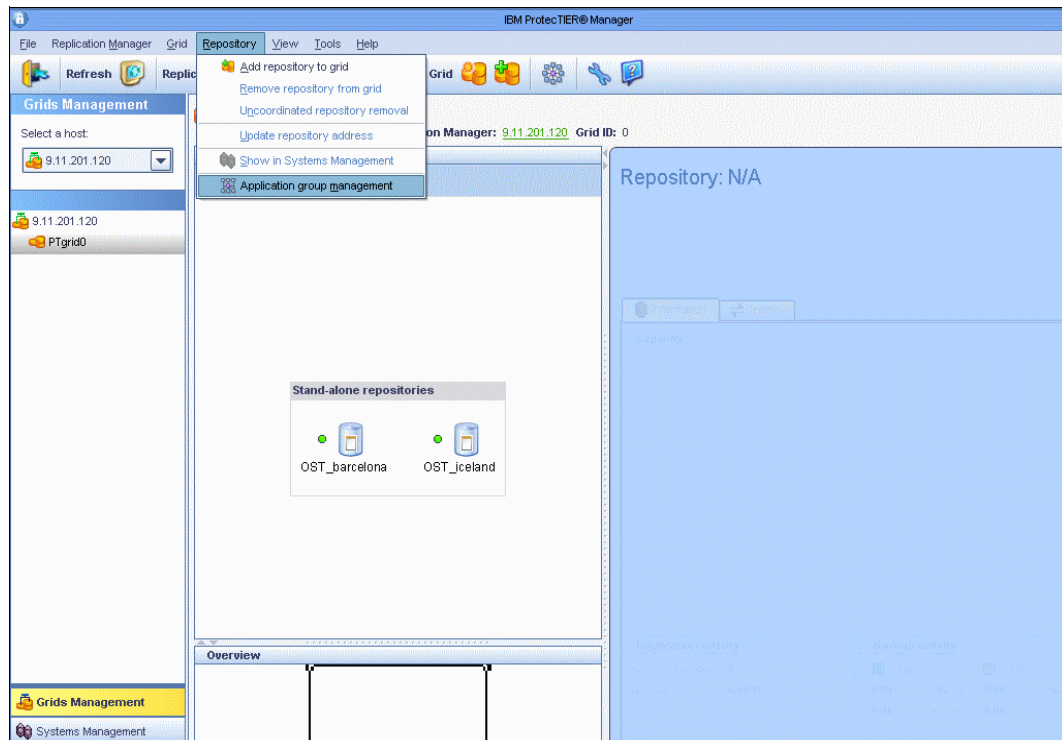


Figure 9-34 Grid Manager





4. The Application group settings and Preview window opens and shows the connected repositories. Enter a name for the application group in the Group name field and click **Finish**. The connections are shown in the Preview pane and the Group name is saved (Figure 9-37).

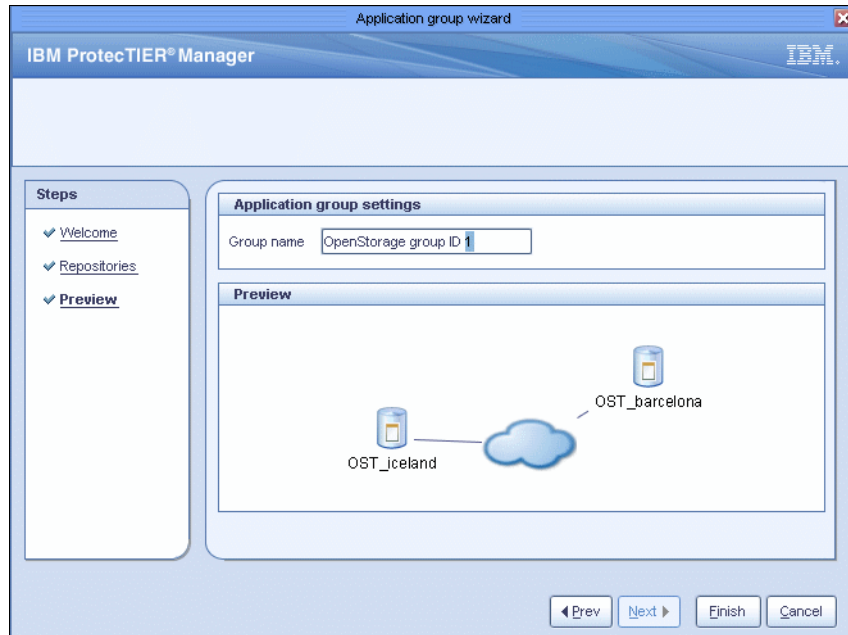


Figure 9-37 Application group wizard: Group settings and preview

**Note:** To add or remove a grid member from an existing group, click one of the group's members and click (or right-click) **Application group management**. A wizard opens, and you can add additional members to the group or remove existing group members.

5. Complete the creation of the application group (Figure 9-38).

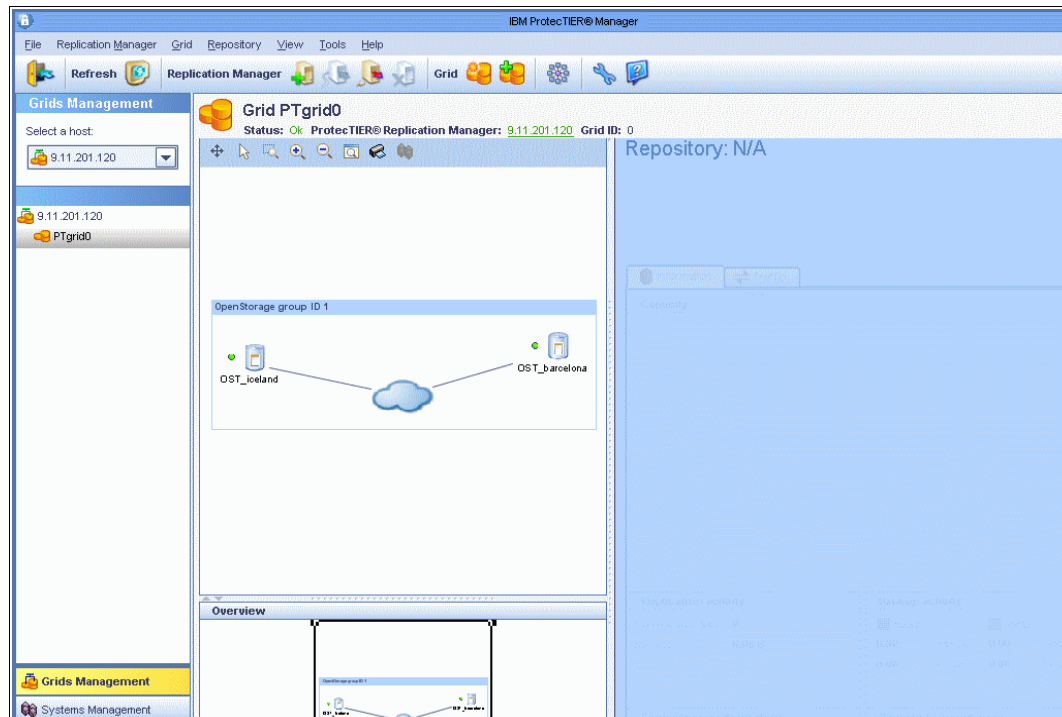


Figure 9-38 Application group

## Setting the replication rate limit for OpenStorage

Setting the replication rate control allows you to limit the nominal and physical throughput (data flow rate) of replication.

The values set for the physical and nominal limits have no explicit influence on one another, that is, the values set in the physical throughput might, but do not necessarily, impact those values set in the nominal throughput, and vice versa.

The physical throughput limit restrains the amount of I/O and resources replication consumes on the local repository. Implicitly, this reduces the total load on the replication networks used by the repository (you can have two networks) and the amount of resources needed on the peer repository as well.

The nominal throughput directly affects the load on the destination repository. On the source repositories, the replication nominal rate does not necessarily compete with the backup. Setting the limit on a source repository guarantees that the backup gets the total possible throughput minus the nominal limit, but in many cases this is not needed.

The Replication Rate Limits window is divided into separate areas for physical and nominal throughput. These areas are divided further into various scenarios where you can limit the rate of replication, for example, replication rates for physical throughput during backup or restore, and replication rates when there is no backup or restore. The same options appear for nominal throughput.

To accomplish this task, log in to System management on ProtecTIER Manager and select **Replication** → **Set replication rate limits** (Figure 9-39).

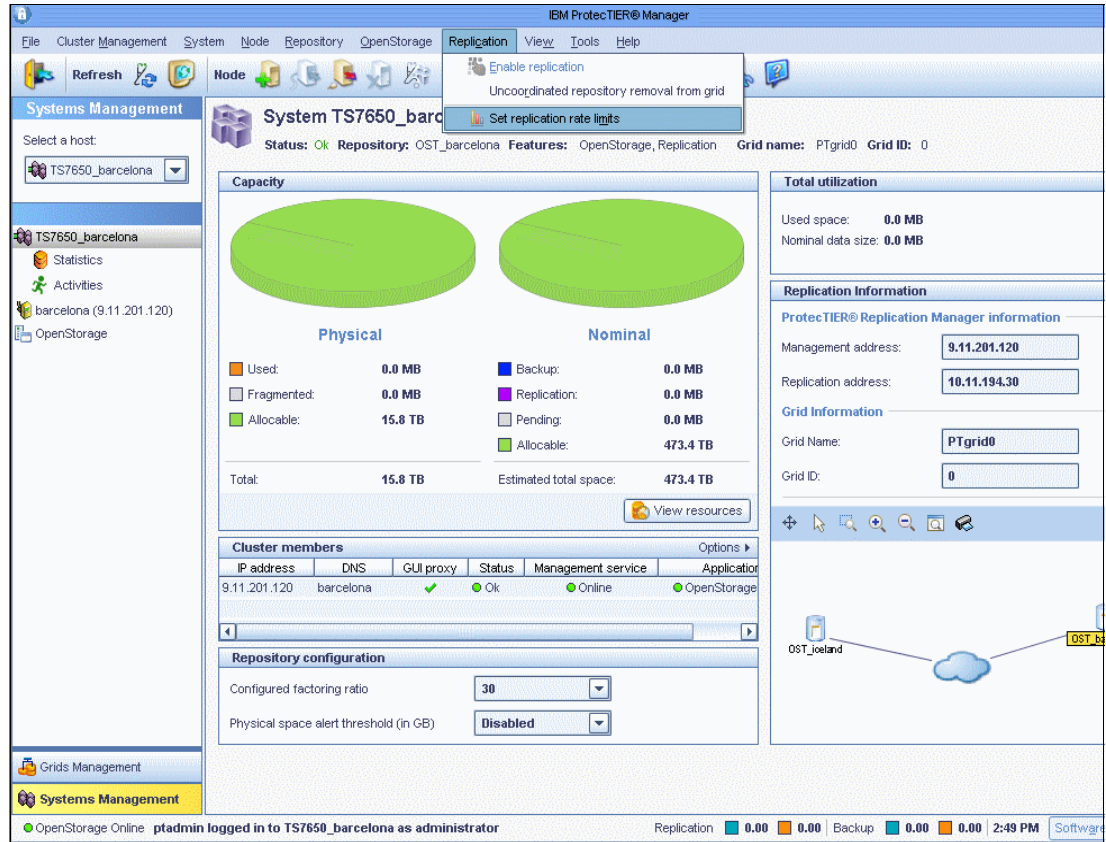


Figure 9-39 Set replication limits menu

To define the replication rate limits for these various scenarios of backup or restore, complete the following steps:

1. Select the check box next to the desired option and enter a value (in MBps). If a check box is not selected, the value will revert to an unlimited replication rate.
2. To return to the original replication rate limits that were defined during installation for the physical or nominal throughputs, click **Restore defaults** at the bottom of either or both areas. The values will default to their original settings at setup.

## 9.6 Native replication with OpenStorage

In order for a logical set of repositories to replicate from one to another, you must create a replication grid (Figure 9-40). The replication grid is remotely created and managed by the Replication Manager.

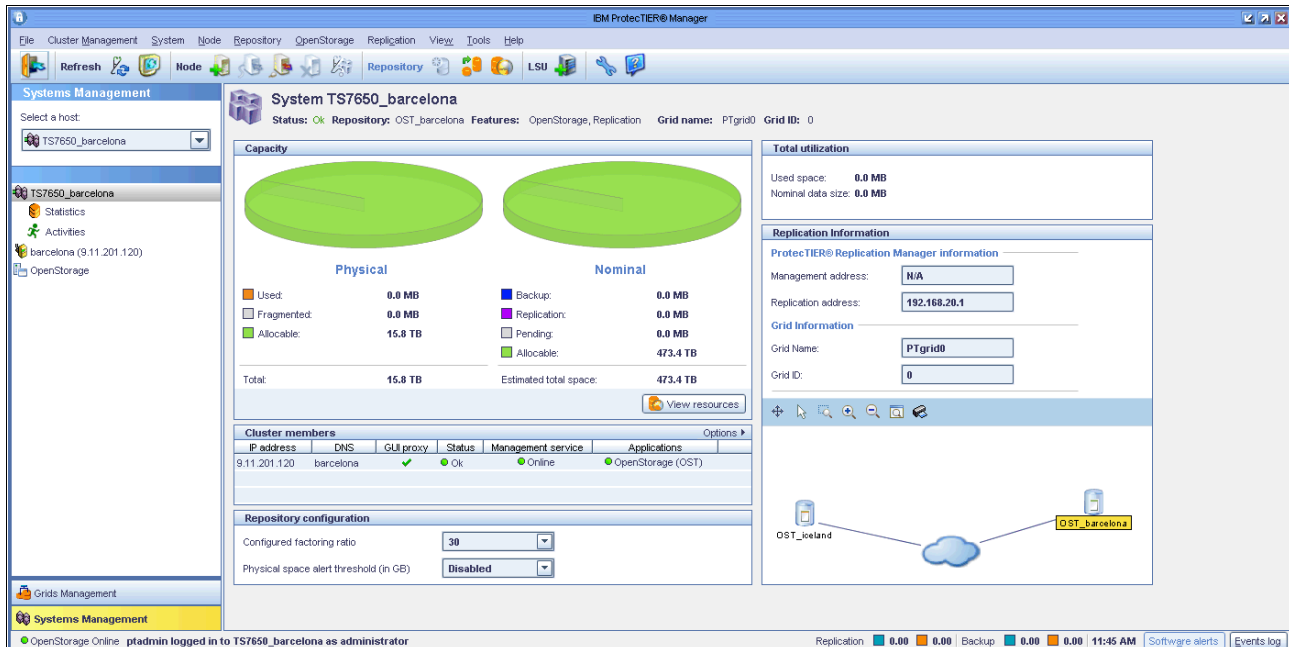


Figure 9-40 OST grid

The ProtectTIER Replication Manager is a server that remotely manages the replication grids within an organization. The ProtectTIER Manager connects to the ProtectTIER Replication Manager using the IP address of the ProtectTIER Replication Manager server. The ProtectTIER Replication Manager can be installed on a dedicated host (which requires an RPQ), or on a ProtectTIER node. If the ProtectTIER Replication Manager is installed on a ProtectTIER node, it can manage up to one grid with 24 repositories. If ProtectTIER Replication Manager is installed on a dedicated server, it can manage up to 64 grids with 256 repositories in each grid.

Each ProtectTIER Replication Manager has a unique identity. A repository, after it has joined a replication manager, cannot join a replication grid managed by a different replication manager, even if it has left the grid. This prevents data collision.

In an OpenStorage environment, replication policies and activities are managed by NetBackup. ProtectTIER Replication Manager is used to manage the replication grid, create application groups, and define the replication connections within the replication grid. Each repository, or grid member, can replicate to each other and replication is bidirectional. A maximum of 12 repositories (members) can replicate to each other at any given time, with up to 256 repositories existing in the entire grid

## 9.7 The replication grid

In an OpenStorage environment, management and configuration of a replication grid is done through the ProtecTIER Replication Manager through the Grids Management view of ProtecTIER Manager (Figure 9-41). A replication grid is composed of a set of repositories that share a common ID and can potentially transmit and receive logical objects through replication.

A maximum of 12 repositories (members) can replicate to each other at any given time, with up to 256 repositories existing in the entire grid. If ProtecTIER Replication Manager is installed on a ProtecTIER node, it can manage up to one grid with a total of 24 repositories.

Repositories do not need physical connections between them. However, all repositories do need a network connection to the ProtecTIER Replication Manager server.

The screenshot displays the IBM ProtecTIER Manager interface. The main window is titled "Grid PTgrid0" and shows a status of "OK". Below the grid name, there are two repositories: "OST\_iceland" and "OST\_barcelona". The "OST\_barcelona" repository is selected, and its details are shown in a large panel on the right. This panel includes a "Capacity" section with two 3D pie charts for "Physical" and "Nominal" capacity. The "Physical" capacity is 15.8 TB, and the "Nominal" capacity is 473.4 TB. Below the capacity section, there are sections for "Replication activity" and "Backup activity". The "Replication activity" section shows "Running activities: 0" and "Backlog: 0.00 B". The "Backup activity" section shows "Read: 0.00 (MB/Sec)" and "Write: 0.00 (MB/Sec)". At the bottom of the interface, there is a status bar that reads "gadmin logged in to 9.11.201.120 as administrator" and the time "11:48 AM".

Figure 9-41 Grids Management View

**Tip:** When working in the Grids Management view of ProtecTIER Manager, a repository can be monitored from the Systems Management view by selecting **Repository** → **Show in Systems Management**.



# Managing IBM System Storage ProtecTIER systems

This chapter describes the management and maintenance of the TS7650 and TS7650G systems, which includes:

- ▶ Managing nodes in ProtecTIER Manager
- ▶ Managing repositories
- ▶ Managing virtual libraries and cartridges
- ▶ Viewing the alerts and event log windows
- ▶ Viewing error messages
- ▶ Generating service reports
- ▶ Adding and removing cluster members
- ▶ Common maintenance tasks
- ▶ Uninstalling the ProtecTIER Manager
- ▶ Exchanging tape cartridges with shelf

# 10.1 Managing nodes in ProtecTIER Manager

In this section, we describe all tasks related to node management using the ProtecTIER Manager.

## 10.1.1 Adding a node and adding node subnetworks

To manage the ProtecTIER nodes, you must add them into the ProtecTIER Manager. The way to add one node into the ProtecTIER Manager is to select **Node** → **Add node**. The Node information window opens (Figure 10-1).




Figure 10-1 Add node window

Enter the IP address and leave the port and ping port fields as they are and click **OK**. If the node that you are adding is a member of a two-node cluster, then the other member of that two-node cluster will be added automatically (Figure 10-2).

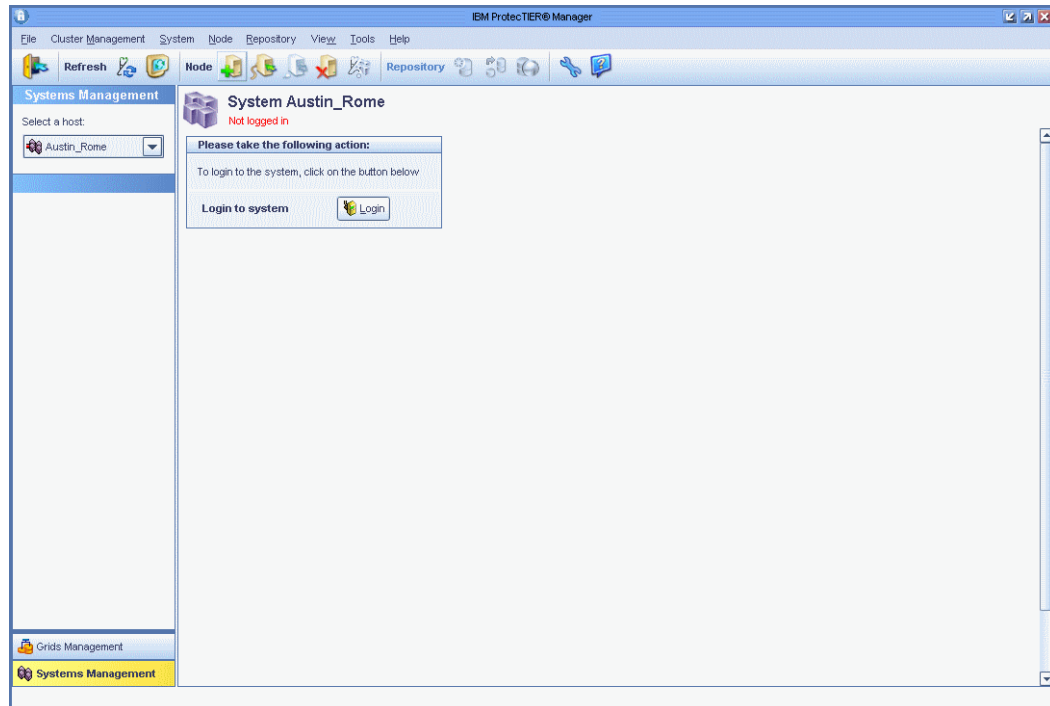


Figure 10-2 Adding a cluster by adding one member



**Note:** When adding a node that is one of the cluster's members, the other node of this cluster will be added automatically.

In addition to adding individual nodes to ProtecTIER Manager, you can add addresses for subnetworks to which nodes are connected. When ProtecTIER Manager restarts, it automatically detects all nodes on the added subnetworks of the TCP/IP network.

To add the addresses, complete the following steps:

1. From the menu bar, select **Tools** → **Preferences**. The Preferences window opens (Figure 10-3).

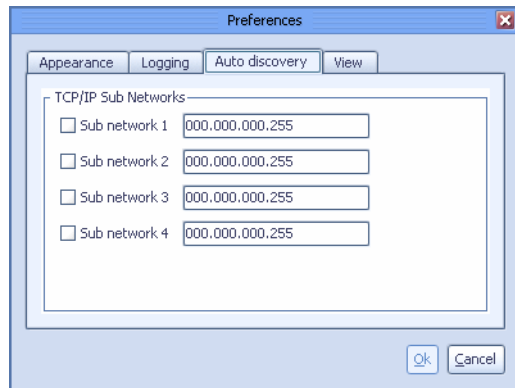


Figure 10-3 ProtecTIER Preferences window

2. For each subnetwork that you want to add, select a subnetwork check box and enter the subnetwork address in the corresponding field.
3. Click **OK**. The Preferences window closes and the subnetwork address is added to ProtecTIER Manager. When you restart ProtecTIER Manager, all nodes on the defined subnetwork addresses are automatically added to ProtecTIER Manager.

## 10.1.2 Removing a node from ProtecTIER Manager

Removing a node stops the instance of ProtecTIER Manager on your workstation from registering the node and being able to manage it. The node itself, and the two-node cluster with which the node is associated, is unaffected by removing nodes in this way.

**Note:** If you remove a node that is associated with a two-node cluster without removing the second node in the cluster, the first node might be detected automatically and added back to ProtecTIER Manager.

To remove a node, complete the following steps:

1. In the Nodes pane, select the node that you want to remove.
2. Click **Remove node**. A confirmation message box opens (Figure 10-4).

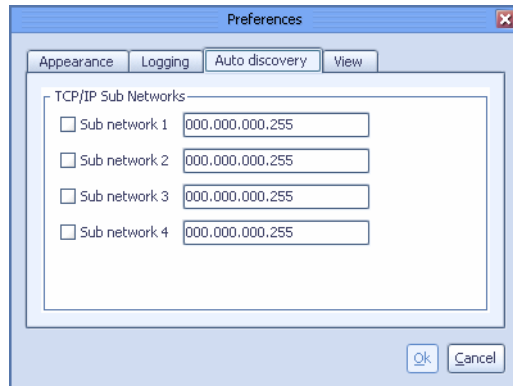


Figure 10-4 ProtecTIER Manager Remove node confirmation

3. Click **Yes**. The node is removed.

## 10.2 Managing repositories

In this section, we describe the tasks related to repository management.

### 10.2.1 Planning an expansion of the repository

Refer to 3.3.5, “Local repository sizing” on page 75, and 3.3.7, “Storage sizing” on page 80, for detailed information about how to plan your repository expansion.

### 10.2.2 Expanding existing file systems

The Repository Metadata storage requirements window might indicate that you must expand some of your existing file systems to accommodate the larger repository.

**Note:** You may also need to create additional file systems. For more information see *IBM System Storage TS7600 with ProtecTIER Installation Guide*, GC53-1155.

Complete the following steps for each file system that you want to expand:

1. Create a partition.
2. Unmount the file system.
3. Deactivate the logical volume.
4. Create a physical volume.
5. Add the physical volume to a volume group.
6. Expand the logical volume.
7. Reactivate the logical volume.
8. Remount the file system.

#### Creating a partition

Create a partition for each file system that you want to expand. For more information, see *IBM System Storage TS7600 with ProtecTIER Installation Guide*, GC53-1155.

To create a partition, you need new volumes that are assigned by a storage subsystem. You have to assign one volume to one partition.

**Note:** To recognize a new volume without rebooting the system, you can use the QLogic tools and run the `ql-dynamic-tgt-lun-disc.sh` script. You can download this tool from the QLogic home page found at the following address:

[http://driverdownloads.qlogic.com/QLogicDriverDownloads\\_UI/SearchByProduct.aspx?ProductCategory=39&Product=935&Os=65](http://driverdownloads.qlogic.com/QLogicDriverDownloads_UI/SearchByProduct.aspx?ProductCategory=39&Product=935&Os=65)

Complete the following steps:

1. Log in to a ProtecTIER node server. Open a Secured Shell (SSH) session to the node and log in using the root user ID and admin password.
2. From the command line, run `multipath -l` (Example 10-1). A list of multipath devices appears.

*Example 10-1 Multipath devices*

```
mpath1 (3600a0b800050e4be000016a24ca4d4ee) dm-1 IBM,1814      FASTT
[size=1.0G][features=1 queue_if_no_path][hwhandler=1 rdac][rw]
  \_ round-robin 0 [prio=0][active]
     \_ 5:0:0:1 sdd 8:48 [active][undef]
  \_ round-robin 0 [prio=0][enabled]
     \_ 3:0:0:1 sdh 8:112 [active][undef]
mpath0 (3600a0b800050e52e000014d54ca4d4d8) dm-0 IBM,1814      FASTT
[size=1.1T][features=1 queue_if_no_path][hwhandler=1 rdac][rw]
  \_ round-robin 0 [prio=0][active]
     \_ 3:0:0:0 sdc 8:32 [active][undef]
  \_ round-robin 0 [prio=0][enabled]
     \_ 5:0:0:0 sdb 8:16 [active][undef]
mpath9 (3600a0b800050e52e0000151b4ce6c709) dm-5 IBM,1814      FASTT
[size=550G][features=1 queue_if_no_path][hwhandler=1 rdac][rw]
  \_ round-robin 0 [prio=0][active]
     \_ 5:0:0:5 sdi 8:128 [active][undef]
  \_ round-robin 0 [prio=0][enabled]
     \_ 3:0:0:5 sdo 8:224 [active][undef]
mpath8 (3600a0b800050e52e000015184ce6c6f6) dm-4 IBM,1814      FASTT
[size=550G][features=1 queue_if_no_path][hwhandler=1 rdac][rw]
  \_ round-robin 0 [prio=0][active]
     \_ 3:0:0:4 sdn 8:208 [active][undef]
  \_ round-robin 0 [prio=0][enabled]
     \_ 5:0:0:4 sdg 8:96 [active][undef]
     \_ 5:0:0:4 sdg 8:96 [active][undef]
mpath7 (3600a0b800050e4be000016fa4ce6c6dc) dm-3 IBM,1814      FASTT
[size=550G][features=1 queue_if_no_path][hwhandler=1 rdac][rw]
  \_ round-robin 0 [prio=0][active]
     \_ 3:0:0:3 sdm 8:192 [active][undef]
  \_ round-robin 0 [prio=0][enabled]
     \_ 5:0:0:3 sdf 8:80 [active][undef]
mpath6 (3600a0b800050e4be000016f74ce6c644) dm-2 IBM,1814      FASTT
[size=550G][features=1 queue_if_no_path][hwhandler=1 rdac][rw]
  \_ round-robin 0 [prio=0][active]
     \_ 5:0:0:2 sde 8:64 [active][undef]
  \_ round-robin 0 [prio=0][enabled]
     \_ 3:0:0:2 sd1 8:176 [active][undef]
```

```

mpath12 (3600a0b800050e4be000016ff4ce6c792) dm-24 IBM,1814      FASST
[size=550G][features=1 queue_if_no_path][hwhandler=1 rdac][rw]
\_ round-robin 0 [prio=0][active]
  \_ 3:0:0:8 sdr 65:16 [active][undef]
\_ round-robin 0 [prio=0][enabled]
  \_ 5:0:0:8 sds 65:32 [active][undef]

```

---

Each device is named `mpathN`, where `N` is a number and can be found on the system as `/dev/mapper/mpathN`.

3. Run **parted /dev/mapper/<mpath device name>** to start the Parted Partition Manipulation utility on the selected device.

4. Run **mklabel gpt** to create a new disk label (partition table):

```

[root@Tuscany /]# parted /dev/mapper/mpath12
GNU Parted 1.8.1
Using /dev/mapper/mpath12
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) mklabel gpt
(parted) print

```

```

Model: Linux device-mapper (dm)
Disk /dev/mapper/mpath12: 591GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt

```

```

Number Start End Size File system Name Flags

```

(parted)

5. Run **mkpart primary <start> <end>** to create a primary partition beginning at `<start>` and ending in `<end>`, in megabytes. If the previous print command shows a size in GB, specify `<end>` with a GB suffix for a size in gigabytes:

```

(parted) mkpart primary 0 590000
(parted) print

```

```

Model: Linux device-mapper (dm)
Disk /dev/mapper/mpath12: 591GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt

```

```

Number Start End Size File system Name Flags
1 17.4kB 590GB 590GB primary

```

(parted)

6. A primary partition and a device named `/dev/mapper/mpathNp1` are created:

```

[root@Tuscany /]# ls /dev/mapper
control mpath10 mpath12 mpath6p1 mpath8p1 vg1-lv_vg1 vg5-lv_vg5
mpath0 mpath10p1 mpath12p1 mpath7 mpath9 vg2-lv_vg2 vg6-lv_vg6
mpath0p1 mpath11 mpath1p1 mpath7p1 mpath9p1 vg3-lv_vg3 vg7-lv_vg7
mpath1 mpath11p1 mpath6 mpath8 vg0-lv_vg0 vg4-lv_vg4

```

7. Run **quit** to exit the Parted Partition Manipulation utility.

8. Run the following command to present the new partition to the operating system:

```
kpartx -a /dev/mapper/mpathN
```

## Unmounting the file system

If you want to unmount the file system that you want to expand, run **umount /mnt/<file system name>** to unmount the file system. If you receive the device is busy message, you have to stop the vtfd service.

**Note:** If you are working with a two-node system, unmount the file system on both nodes.

## Deactivating the logical volume

To deactivate the logical volume corresponding to the file system that you want to expand, complete the following steps:

1. Run **lvchange -an /dev/<volume group name>/<logical volume name>**, where <volume group name> is the volume group corresponding to the file system that you want to expand.
2. Run **lvscan** to view the list of logical volumes. The logical volume should be deactivated, as shown in Example 10-2.

*Example 10-2 Deactivate the logical volume*

---

```
[root@Tuscany /]# lvchange -an /dev/vg1/lv_vg1
[root@Tuscany /]# lvscan
ACTIVE                '/dev/vg2/lv_vg2' [550.00 GB] inherit
ACTIVE                '/dev/vg5/lv_vg5' [550.00 GB] inherit
ACTIVE                '/dev/vg7/lv_vg7' [550.00 GB] inherit
ACTIVE                '/dev/vg0/lv_vg0' [1020.00 MB] inherit
ACTIVE                '/dev/vg6/lv_vg6' [550.00 GB] inherit
inactive            '/dev/vg1/lv_vg1' [1.09 TB] inherit
ACTIVE                '/dev/vg4/lv_vg4' [550.00 GB] inherit
ACTIVE                '/dev/vg3/lv_vg3' [550.00 GB] inherit
```

---

## Creating a physical volume

For each new partition created, create a physical volume by completing the following steps:

1. Run **pvcreate /dev/mapper/<device name>** to create the physical volume:

```
[root@Tuscany /]# pvcreate /dev/mapper/mpath12p1
Physical volume "/dev/mapper/mpath12p1" successfully created
```

2. Run **pvscan** to view the list of physical volumes. The physical volume that you created should appear in the list.

```
[root@Tuscany /]# pvscan
PV /dev/mapper/mpath9p1   VG vg2          lvm2 [550.00 GB / 0   free]
PV /dev/mapper/mpath6p1   VG vg5          lvm2 [550.00 GB / 0   free]
PV /dev/mapper/mpath10p1  VG vg7          lvm2 [550.00 GB / 0   free]
PV /dev/mapper/mpath1p1   VG vg0          lvm2 [1020.00 MB / 0   free]
PV /dev/mapper/mpath11p1  VG vg6          lvm2 [550.00 GB / 0   free]
PV /dev/mapper/mpath0p1   VG vg1          lvm2 [1.09 TB / 0     free]
PV /dev/mapper/mpath7p1   VG vg4          lvm2 [550.00 GB / 0   free]
PV /dev/mapper/mpath8p1   VG vg3          lvm2 [550.00 GB / 0   free]
PV /dev/mapper/mpath12p1          lvm2 [549.48 GB]
Total: 9 [4.85 TB] / in use: 8 [4.31 TB] / in no VG: 1 [549.48 GB]
```

## Adding the physical volume to a volume group

Add each new physical volume to the existing volume group corresponding to the file system that you want to expand by completing the following steps:

1. Run `vgextend <volume group name> /dev/<device name>` to add the physical volume to an existing volume group.

```
[root@Tuscany /]# vgextend vg1 /dev/mapper/mpath12p1
Volume group "vg1" successfully extended
```

2. Run `pvscan` to view the list of physical volumes. The physical volume path should include `/dev/<device name>`.

```
[root@Tuscany /]# pvscan
PV /dev/mapper/mpath9p1   VG vg2   lvm2 [550.00 GB / 0   free]
PV /dev/mapper/mpath6p1   VG vg5   lvm2 [550.00 GB / 0   free]
PV /dev/mapper/mpath10p1  VG vg7   lvm2 [550.00 GB / 0   free]
PV /dev/mapper/mpath1p1   VG vg0   lvm2 [1020.00 MB / 0   free]
PV /dev/mapper/mpath11p1  VG vg6   lvm2 [550.00 GB / 0   free]
PV /dev/mapper/mpath0p1   VG vg1   lvm2 [1.09 TB / 0   free]
PV /dev/mapper/mpath12p1  VG vg1   lvm2 [549.48 GB / 549.48 GB free]
PV /dev/mapper/mpath7p1   VG vg4   lvm2 [550.00 GB / 0   free]
PV /dev/mapper/mpath8p1   VG vg3   lvm2 [550.00 GB / 0   free]
```

## Expanding the logical volume

You can expand the logical volume corresponding to the file system that you want to expand by running the following command:

```
lvextend -L +<size in Gigabytes>G /dev/<volume group name>/<logical volume name>
/dev/<device name>
```

Example 10-3 shows the output of the command.

### Example 10-3 Expanding the logical volume

---

```
[root@Tuscany /]# lvextend -L +549G /dev/vg1/lv_vg1 /dev/mapper/mpath12p1
Extending logical volume lv_vg1 to 1.62 TB
Logical volume lv_vg1 successfully resized
[root@Tuscany /]# lvscan
ACTIVE          '/dev/vg2/lv_vg2' [550.00 GB] inherit
ACTIVE          '/dev/vg5/lv_vg5' [550.00 GB] inherit
ACTIVE          '/dev/vg7/lv_vg7' [550.00 GB] inherit
ACTIVE          '/dev/vg0/lv_vg0' [1020.00 MB] inherit
ACTIVE          '/dev/vg6/lv_vg6' [550.00 GB] inherit
inactive       '/dev/vg1/lv_vg1' [1.62 TB] inherit
ACTIVE          '/dev/vg4/lv_vg4' [550.00 GB] inherit
ACTIVE          '/dev/vg3/lv_vg3' [550.00 GB] inherit
```

---

## Reactivating the logical volume

To reactivate the expanded logical volume, run the following command and review the output:

```
lvchange -ay /dev/<volume group name>/<logical volume name>
[root@Tuscany /]# lvchange -ay /dev/vg1/lv_vg1
[root@Tuscany /]# lvscan
ACTIVE          '/dev/vg2/lv_vg2' [550.00 GB] inherit
ACTIVE          '/dev/vg5/lv_vg5' [550.00 GB] inherit
ACTIVE          '/dev/vg7/lv_vg7' [550.00 GB] inherit
ACTIVE          '/dev/vg0/lv_vg0' [1020.00 MB] inherit
```

```

ACTIVE          '/dev/vg6/lv_vg6' [550.00 GB] inherit
ACTIVE          '/dev/vg1/lv_vg1' [1.62 TB] inherit
ACTIVE          '/dev/vg4/lv_vg4' [550.00 GB] inherit
ACTIVE          '/dev/vg3/lv_vg3' [550.00 GB] inherit

```

## Remounting the file system

To remount the file system that you expanded, run **mount -a** to remount all the file systems. Afterwards, you have to start the vtfd service again.

**Note:** If you are working with a two-node system, remount the file system on both nodes.

## 10.2.3 Expanding the repository

Using the information generated during the repository expansion planning process, expand the repository by completing the following steps:

1. Log in to ProtecTIER Manager with administrator rights.
2. In the Repositories pane, select the repository that you want to expand.
3. From the menu bar, select **Repository** → **Increase capacity**. The Increase capacity wizard Welcome window opens (Figure 10-5).

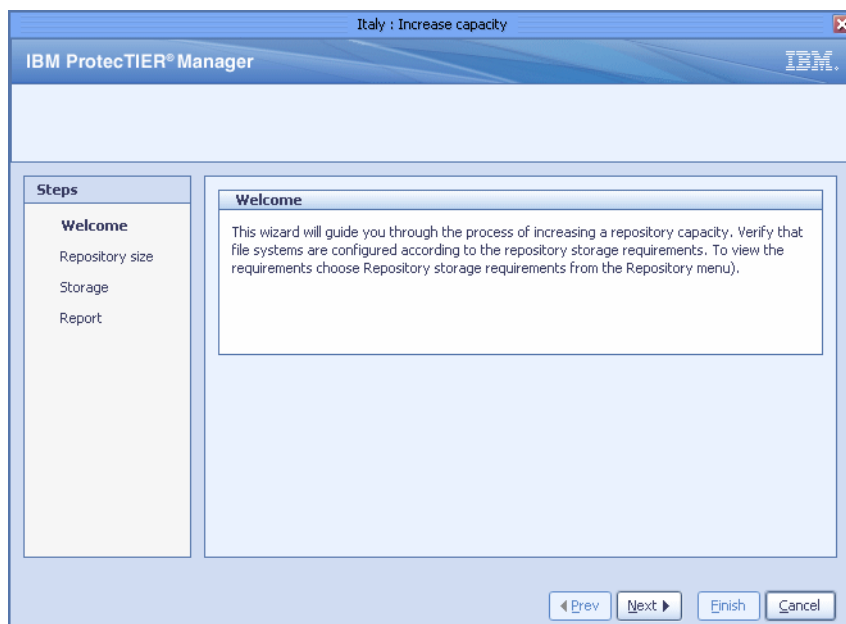


Figure 10-5 Increase capacity: Welcome window

Click **Next**. The Repository size window opens (Figure 10-6).

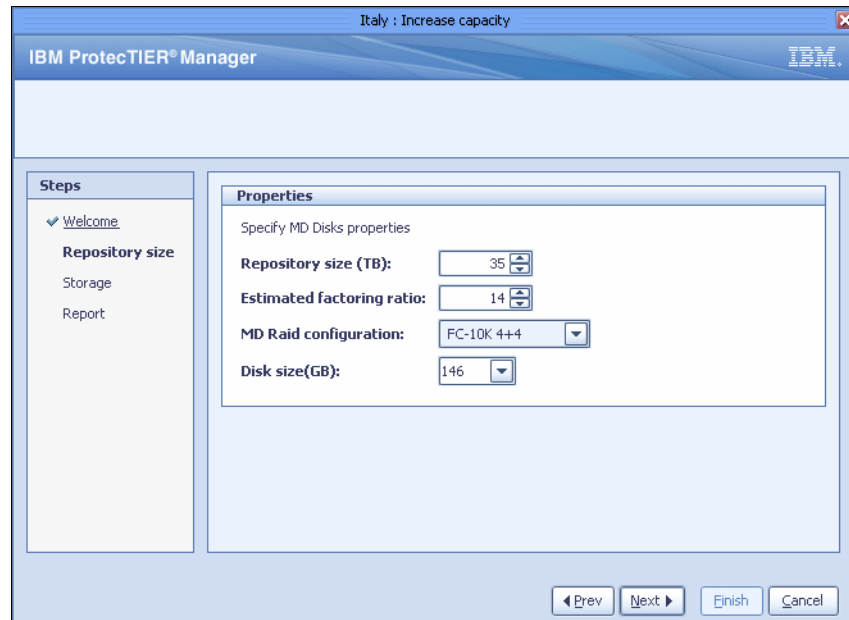


Figure 10-6 Increase capacity: Properties window

4. In the Repository size field, enter the size, in terabytes, to which you want the repository expanded.
5. In the Estimated factoring ratio field, enter the factoring ratio value that you estimate for the expanded repository.

Click **Next**. The Storage window opens (Figure 10-7).

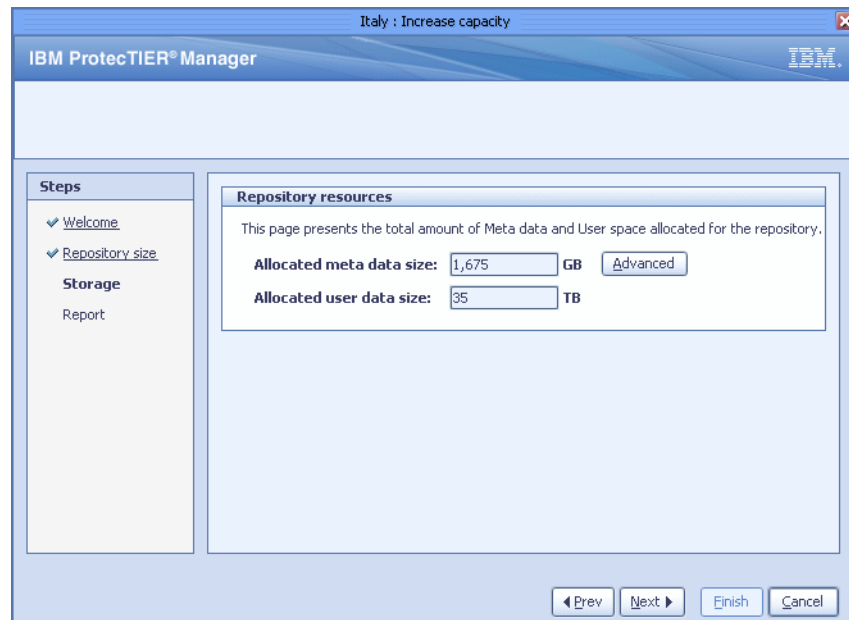


Figure 10-7 Increase capacity metadata window

6. The Allocated metadata size field value and the Allocated user data size field value are automatically generated based on the values entered in the Repository size window.



If you would like to change the default Metadata file system definitions, click **Advanced**. The Repository resources file system assignment window opens (Figure 10-8). Otherwise, skip to step 9 on page 484.

**Note:** If changes to the file system assignment are necessary, the Repository resources file system assignment window opens automatically

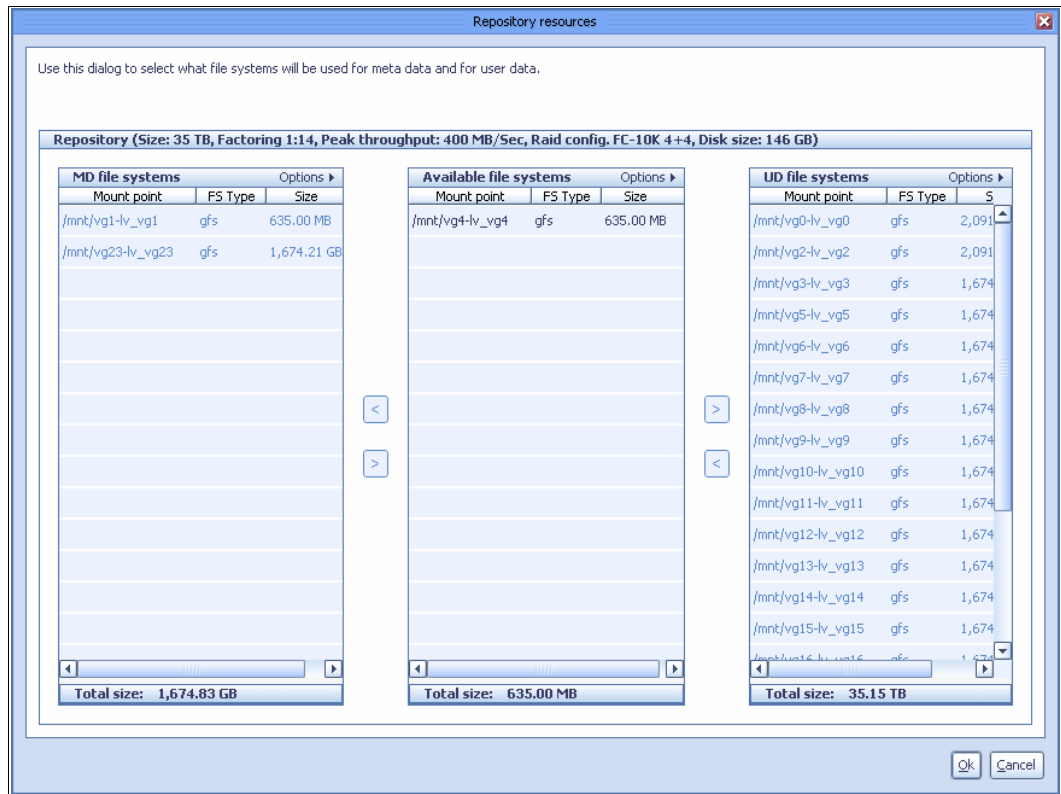


Figure 10-8 Repository resources: File system assignment window

7. Select the file systems that you want to add for Metadata or User Data in the Available file systems list and click the left-arrow button.

Optionally, remove available file systems from metadata use by selecting a file system in the MD file systems list and clicking the right-arrow button.

**Note:** File systems shown in grey are not movable.

8. Click **OK**. The Metadata resources window closes.

Click **Next**. The Report window opens (Figure 10-9).

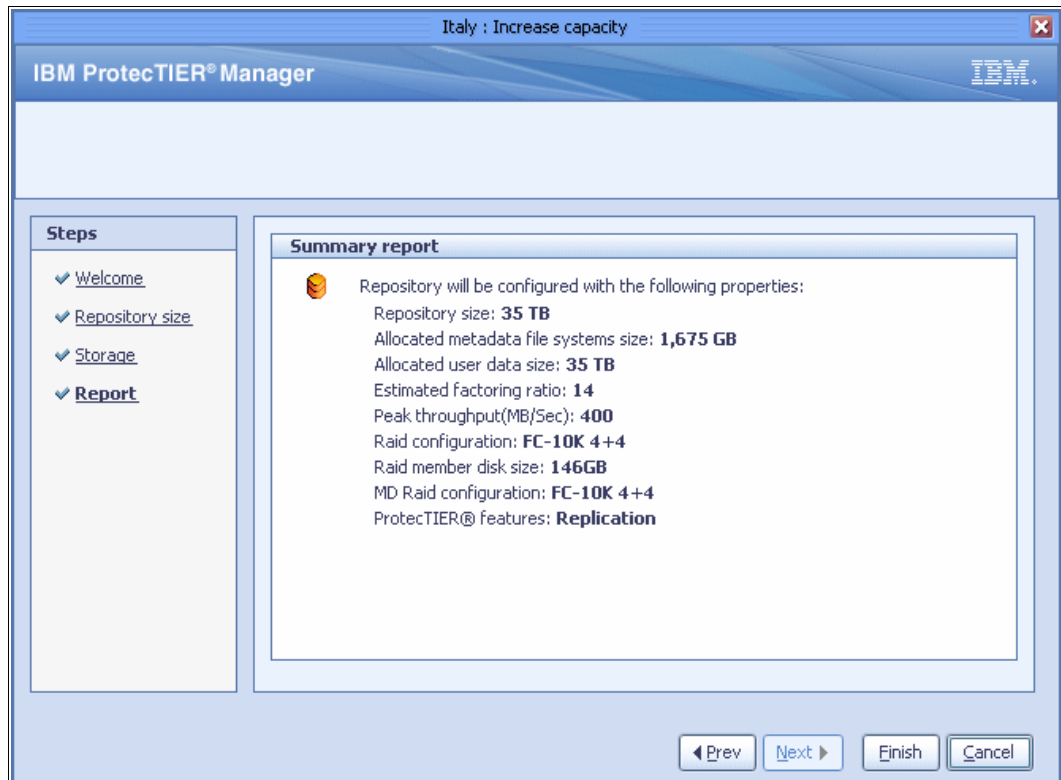


Figure 10-9 Increase capacity: Report window

9. Click **Finish**. A confirmation window opens (Figure 10-10).

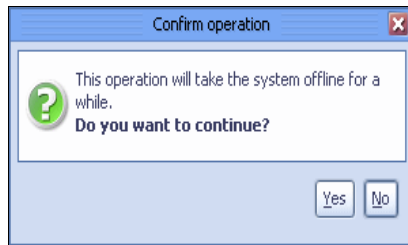


Figure 10-10 Increase repository confirmation window

10. Click **Yes**. The Increase repository wizard closes and the ProtecTIER system temporarily goes offline to increase the repository capacity. When the system goes back online, you can check the result, as shown in Figure 10-11.

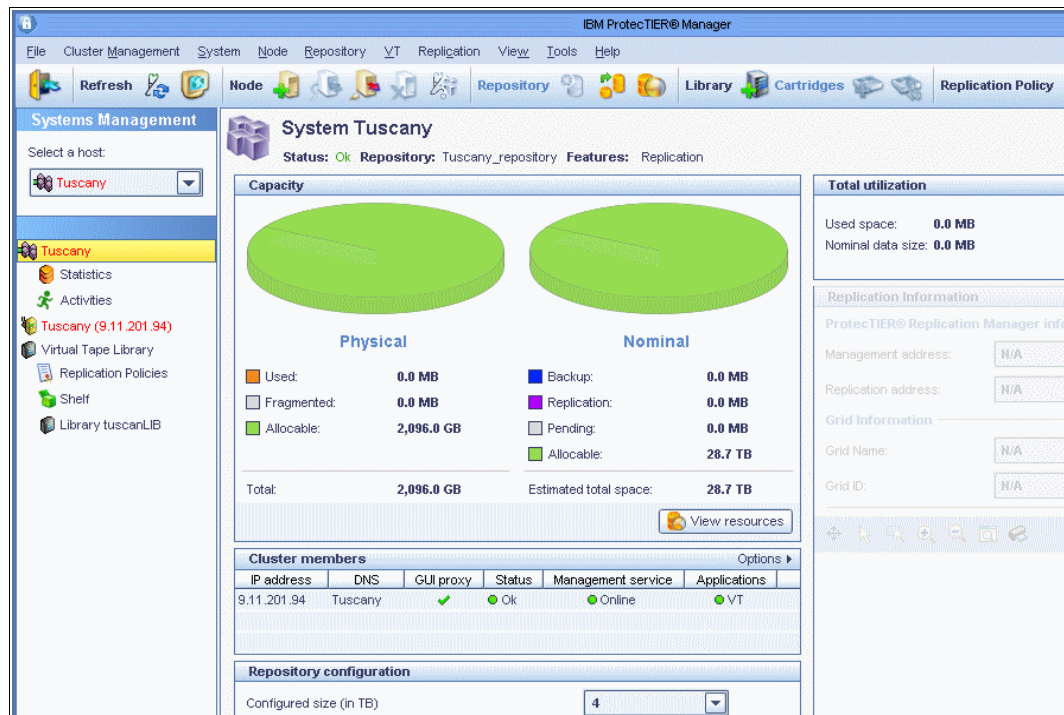


Figure 10-11 Increased repository capacity

## 10.2.4 Deleting the repository

In general, it is not necessary to delete the repository. However, it may be necessary in certain troubleshooting scenarios. For more information, contact IBM Support. In a two-node cluster where the repository was created on one of the nodes, you must remove the second node (cluster member) before you delete the repository. For more information, refer to 10.7.2, “Removing a cluster member” on page 528. In a replication environment, you also must remove the repository that you want to delete from the replication grid before deleting it.

**Attention:** Deleting the repository results in the loss of all the data contained in the repository. Do *not* delete the repository unless ordered to do so by IBM Support.

Complete the following steps:

1. Log in to ProtecTIER Manager with administrator rights.
2. From the menu bar, select **Repository** → **Delete repository** (Figure 10-12).

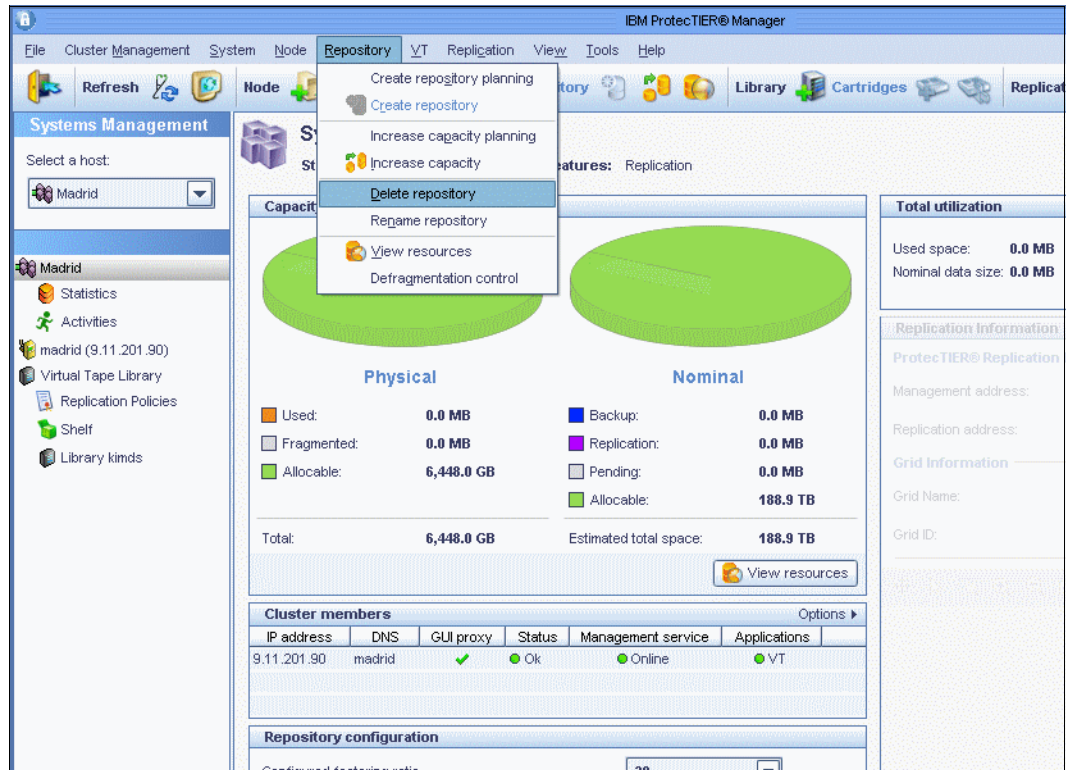


Figure 10-12 Delete repository window

3. A confirmation window opens (Figure 10-13).

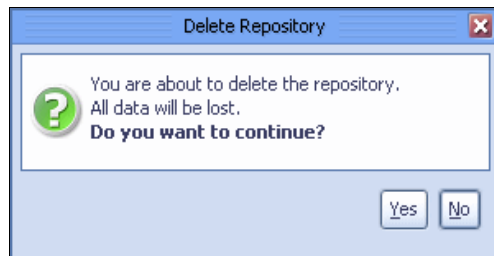


Figure 10-13 Delete repository confirmation window

4. Click **Yes**. The Confirm data loss window opens (Figure 10-14).

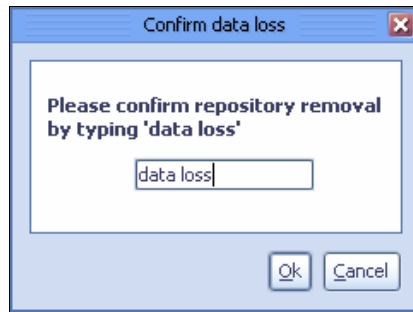


Figure 10-14 Delete repository data loss confirmation window

5. In the field, enter “data loss” and click **OK**. The ProtecTIER system temporarily goes offline to delete the repository (Figure 10-15).



Figure 10-15 Performing repository deletion window

The repository is deleted.

## 10.3 Managing virtual libraries and cartridges

In this section, we describe all tasks related to managing virtual libraries and cartridges.

### 10.3.1 Editing library parameters

The ProtecTIER system enables you to change the parameters of existing libraries, including changing the assignment of virtual devices, adding virtual tape drives, and increasing the library capacity.

**Note:** The Change dimensions wizard does not enable you to assign existing unassigned tape drives. Use the Re-assign devices wizard to assign unassigned tape drives. For more information, see 10.3.2, “Reassigning devices” on page 494.

Complete the following steps:

1. In the Services pane, select a virtual tape library.

- From the menu bar, select **VT** → **VT Library** → **Change dimensions**. The Change dimensions wizard welcome window opens (Figure 10-16).

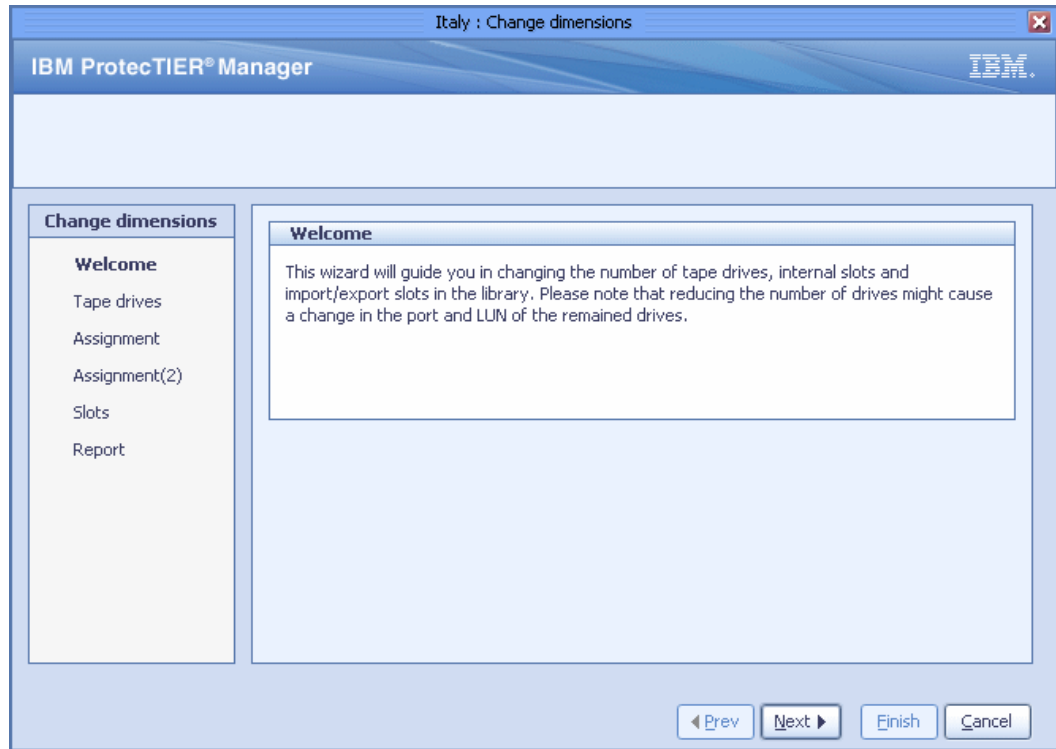


Figure 10-16 Change library dimension window

Click **Next**. The Tape drives window opens (Figure 10-17).

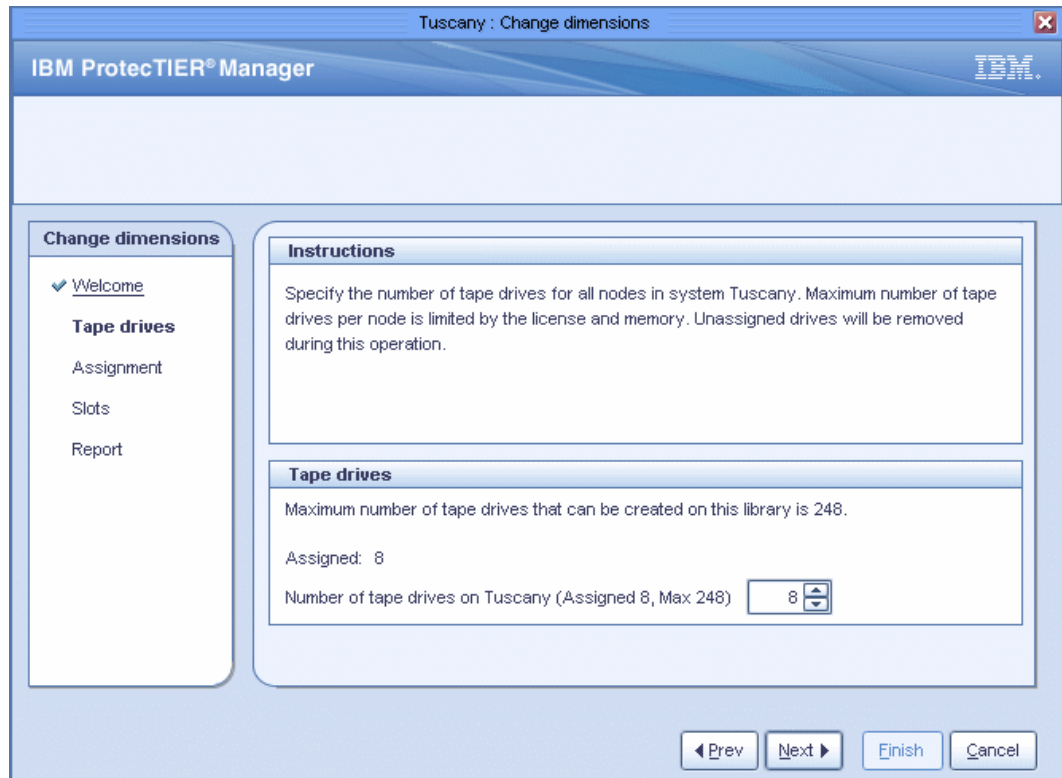


Figure 10-17 Change dimensions: Tape drives window

3. In the Number of tape drives field for each node, enter the number of tape drives that you want to have in the node.

Click **Next**. The Assignment window opens for the first node in the two-node cluster (Figure 10-18).

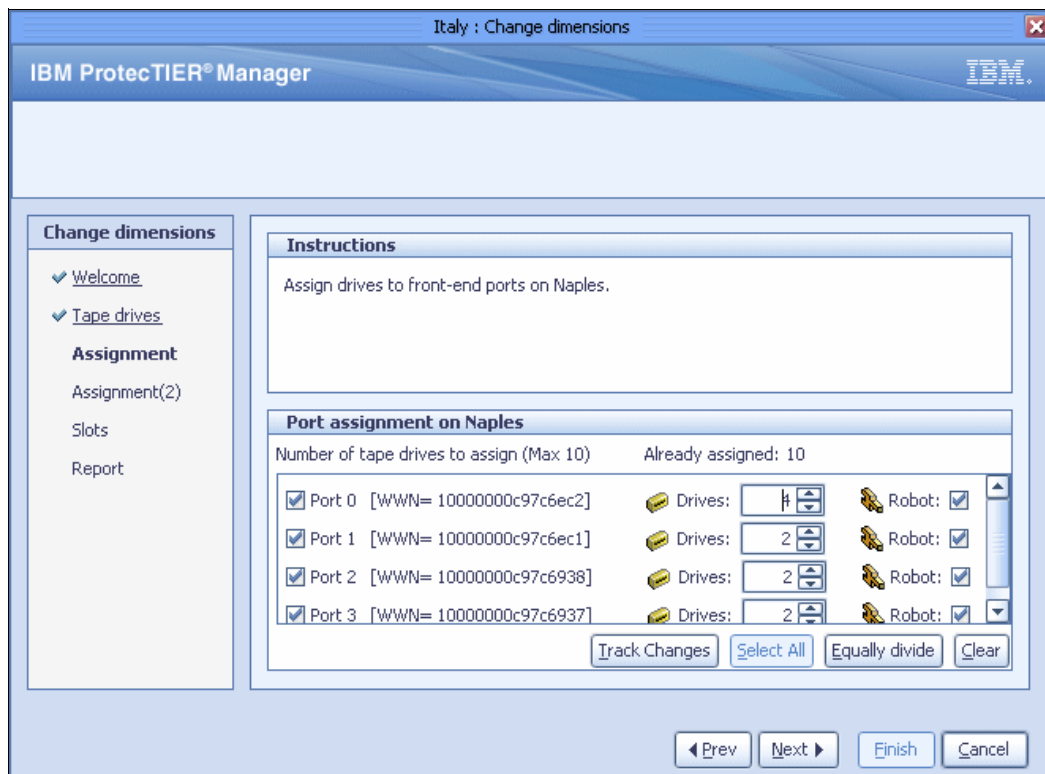


Figure 10-18 Change dimensions port assignment window

4. Select or clear the check boxes before each port to define which of the node's ports are assigned virtual tape drives.
5. In the Number of drives fields corresponding to each selected port, select the number of virtual tape drives that are assigned to each port.

**Note:** Optionally, click **Select All** to automatically select all possible ports. Click **Equally divide** to evenly divide the number of drives between the ports.

6. Select or clear the check boxes at the end of the line of each port to define which of the node's port are assigned virtual library robots, that is, open or close the control path of each port. If you want to share the virtual tape library or to implement path failover, refer to Chapter 6, "Host implementation for virtual tape libraries" on page 275.



7. Optionally, click **Track Changes** to display the modifications that the ProtecTIER system might need to make in response to the changes that you defined.
8. Click **Re-assign ports** to return to the Assignment window and continue assigning virtual tape drives. Click **Next**. If a second node exists in the cluster, the Assignment(2) window opens (Figure 10-19).

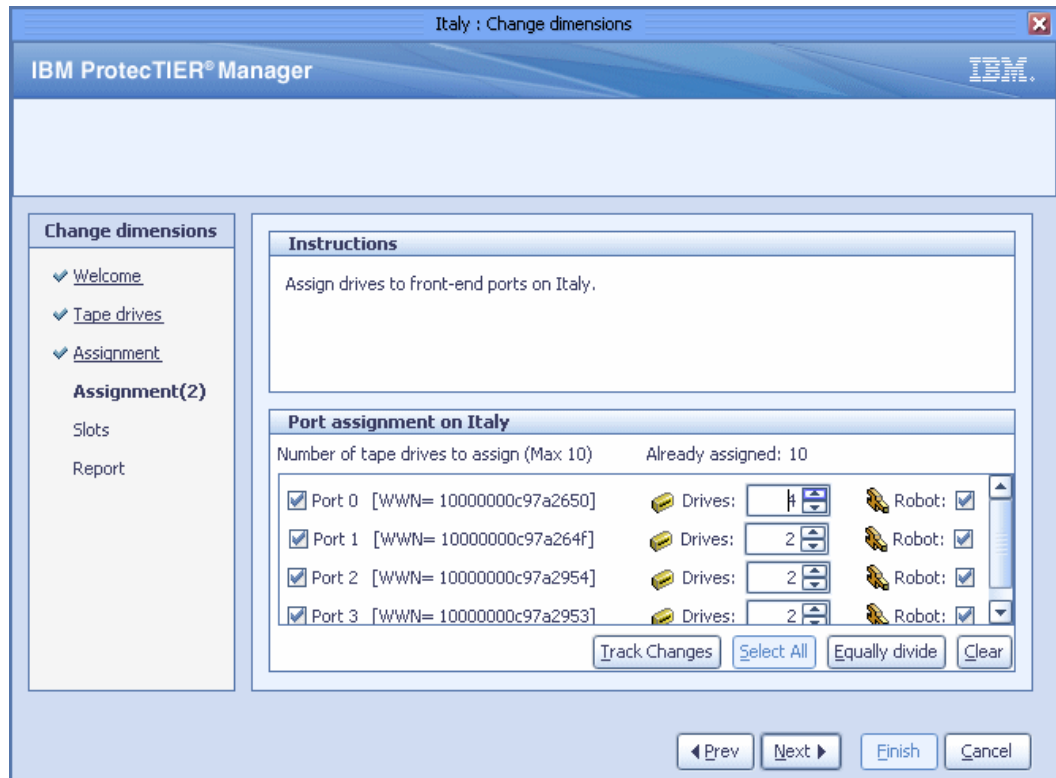


Figure 10-19 Change dimensions: Assignment(2) window

9. Repeat steps 9 through 11 for the Assignment(2) window.

10. Click **Next**. The Slots window opens (Figure 10-20).

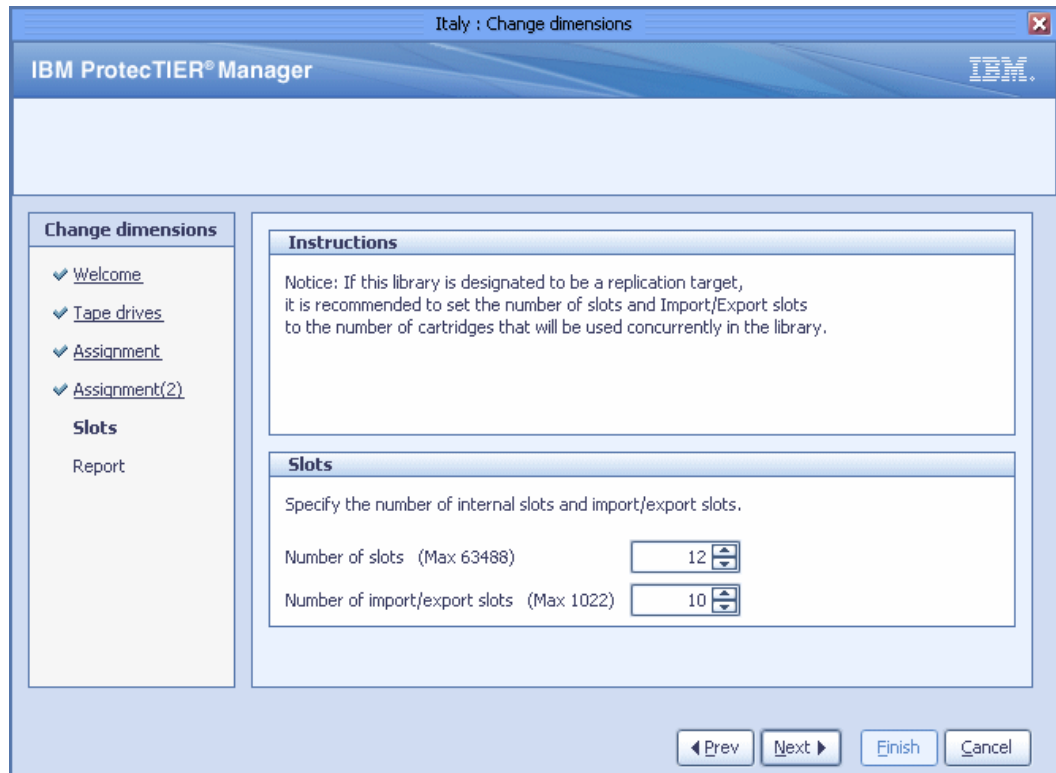


Figure 10-20 Change dimensions: Slots window

11. In the Number of slots field, enter the number of slots that you want in the library.

**Note:** The number of slots must be at least equal to the number of cartridges in the library.

12. In the Number of import/export slots field, enter the number of import/export slots that you want in the library.

Click **Next**. The Report window opens (Figure 10-21).

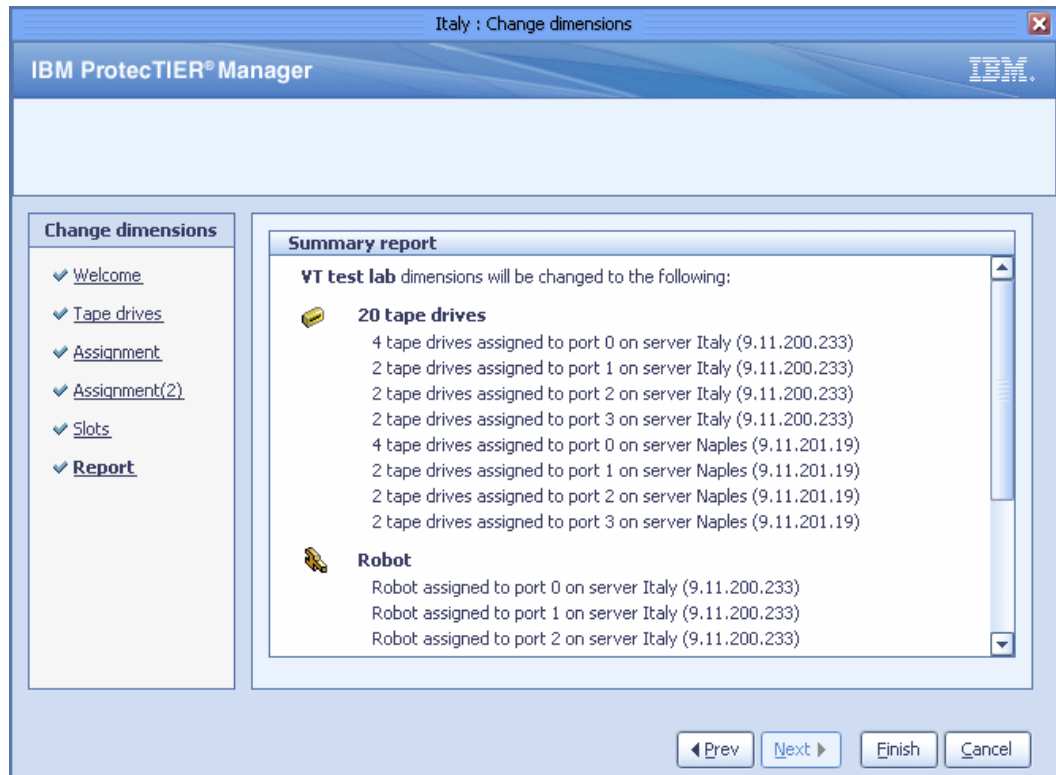


Figure 10-21 Change dimensions: Report window

13. Click **Finish**. The Confirm Operation window opens (Figure 10-22).

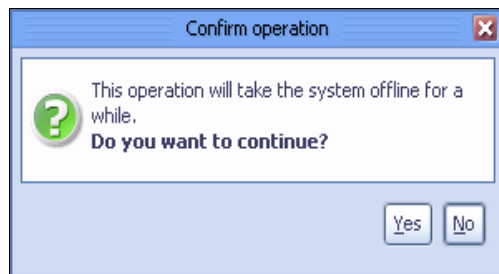


Figure 10-22 ProtecTIER change dimensions: Confirm operation window

14. Click **Yes**. The Change Dimension Wizard closes and the ProtecTIER system temporarily goes offline to update the library (Figure 10-23).



Figure 10-23 Change tape library dimensions window

### 10.3.2 Reassigning devices

The ProtecTIER system enables you to relocate the virtual robot and drives between nodes or node ports.

To reassign devices, complete the following steps:

1. From the menu bar, select **VT** → **VT Library** → **Re-assign devices** or click **Re-assign devices** on the General tab of the Library view. The Re-assign Devices wizard Welcome window opens (Figure 10-24).

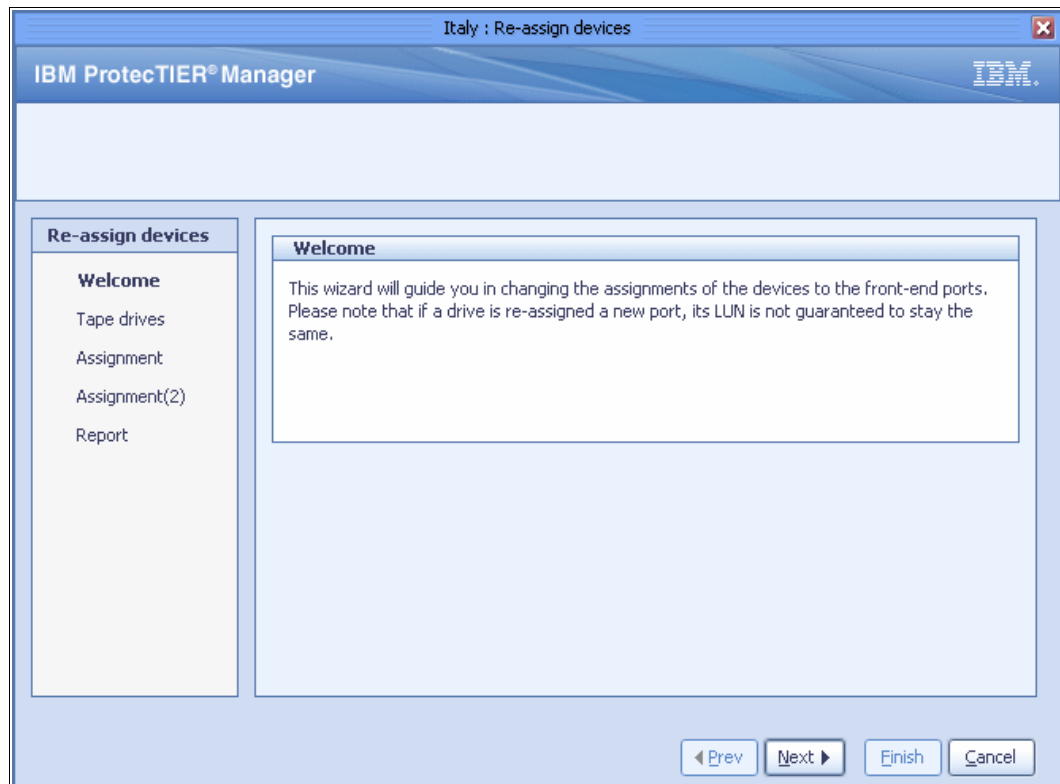


Figure 10-24 Re-assign devices: Welcome window

Click **Next**. The Tape drives window opens (Figure 10-25).

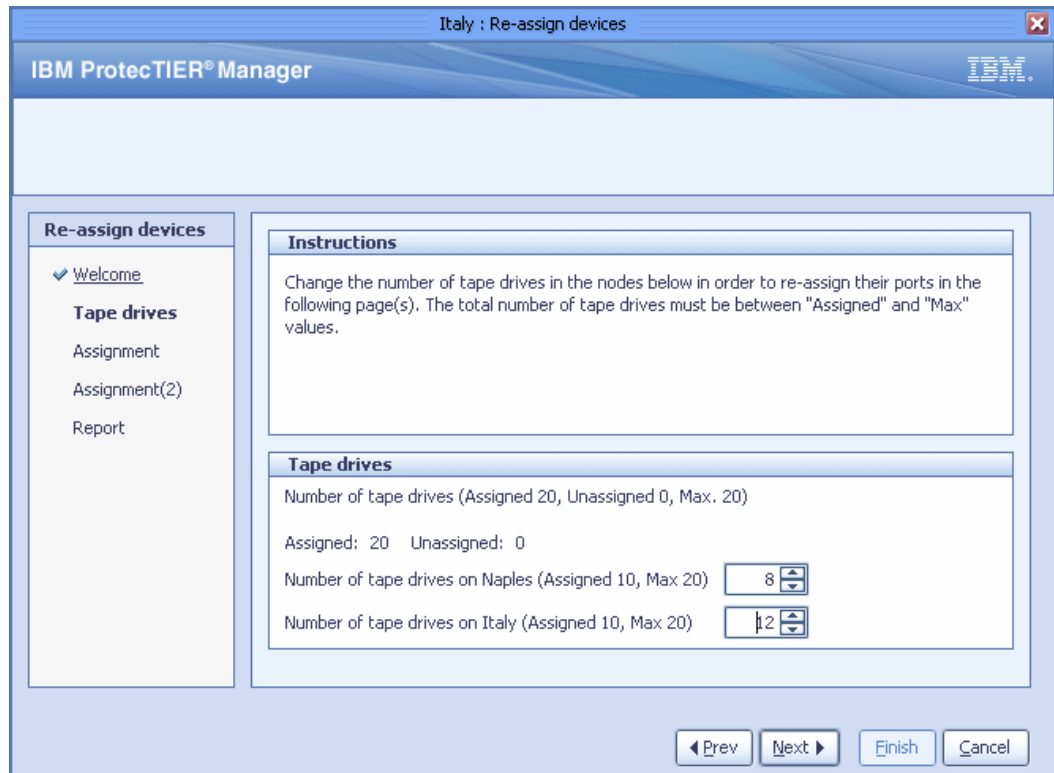


Figure 10-25 Re-assign devices: Tape drives window

2. In the Number of tape drives field, select the number of tape drives to assign to the node.

**Note:** If a node is unavailable, you can only remove drives from that node.

Click **Next**. The Assignment window opens (Figure 10-26).

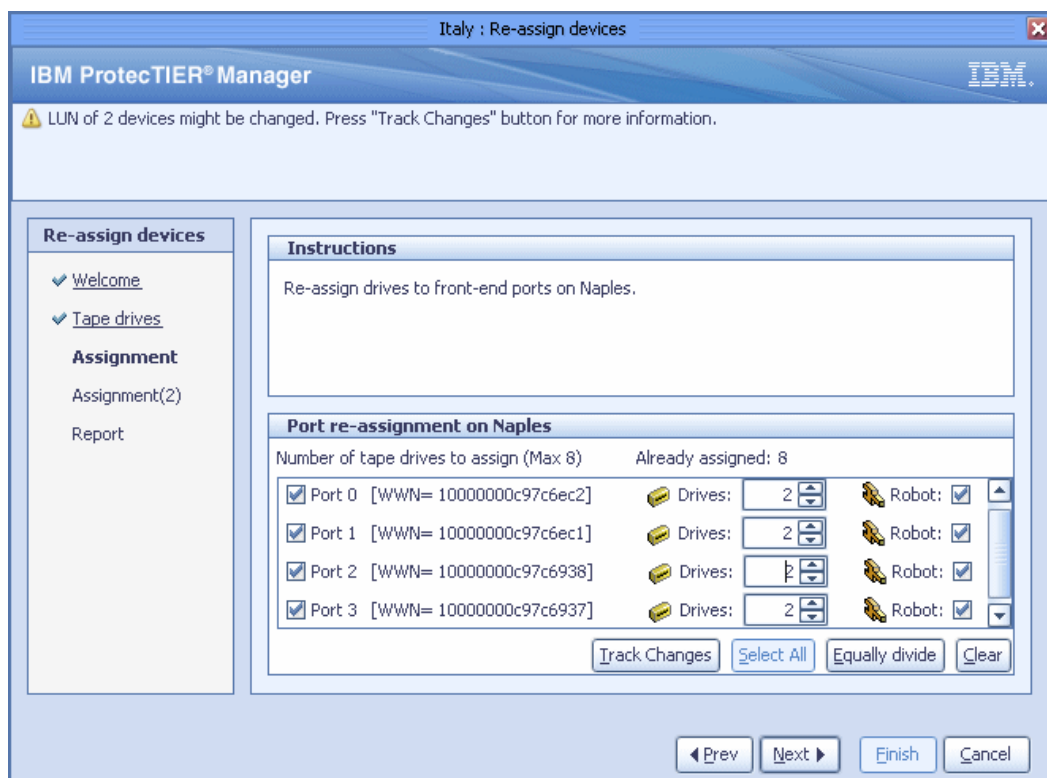


Figure 10-26 Re-assign devices: Assignment window

3. Select or clear the check boxes before each port to define which of the node's ports are assigned virtual tape drives.
4. In the Number of drives fields corresponding to each selected port, select the number of virtual tape drives that are assigned to each port.

**Note:** Optionally, click **Select All** to automatically select all possible ports. Click **Equally divide** to evenly divide the number of drives between the ports.

5. Select or clear the check boxes at the end of the line of each port to define which of the node's ports are assigned virtual library robots, that is, open or close the control path of each port. If you want to share the virtual tape library or to implement path failover, refer to Chapter 6, "Host implementation for virtual tape libraries" on page 275.

6. Optionally, click **Track Changes** to display the modifications that the ProtecTIER system might need to make in response to the changes that you defined (Figure 10-27).

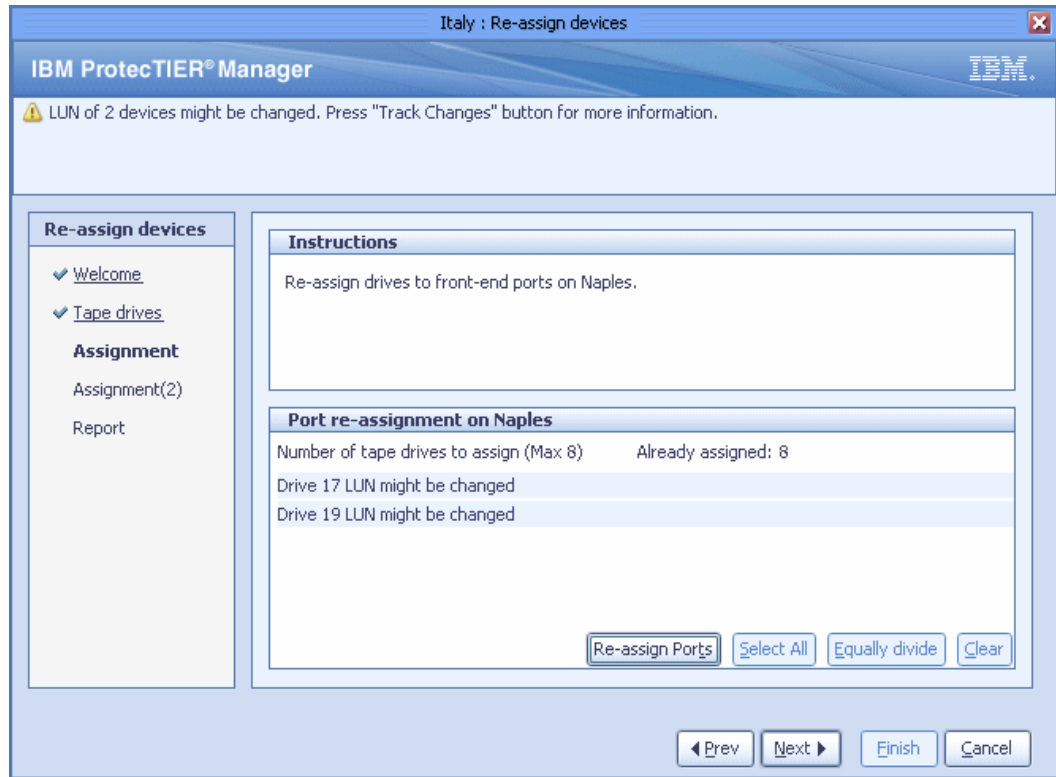


Figure 10-27 Re-assign devices: Track changes window

7. Click **Re-assign ports** to return to the Assignment window and continue assigning virtual tape drives.

Click **Next**. If a second node exists in the cluster, the Assignment(2) window opens (Figure 10-28).

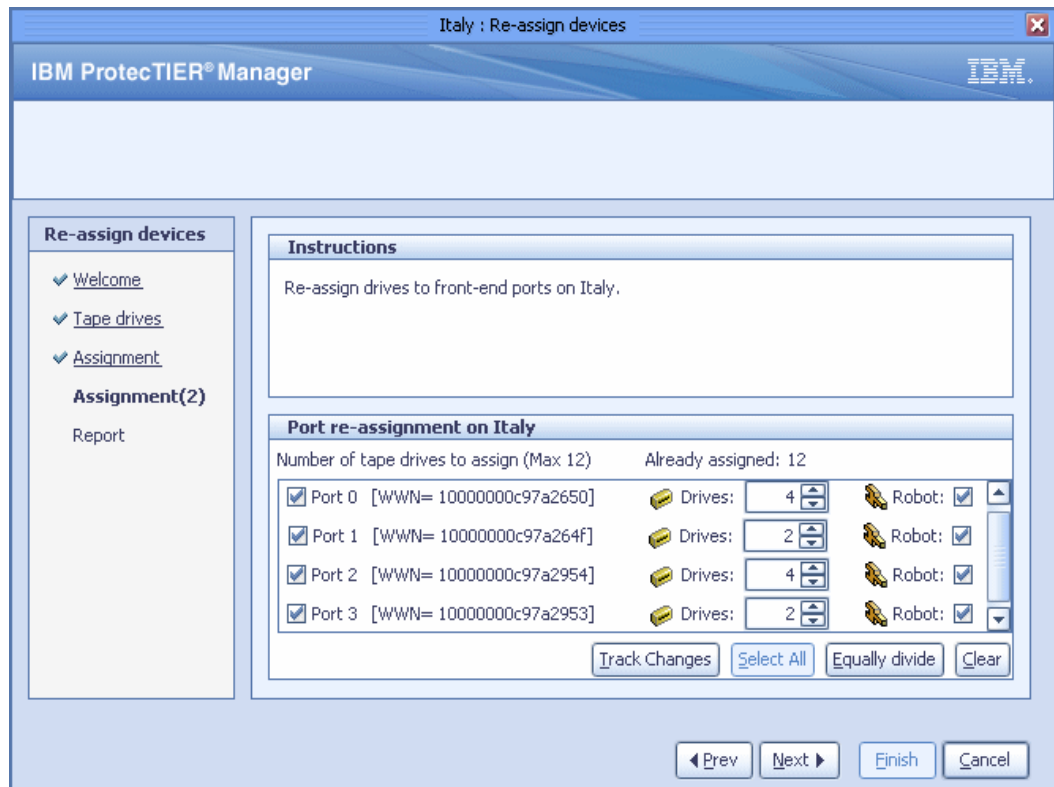


Figure 10-28 Re-assign devices: Assignment(2) window

- Repeat steps 5 through 10 for the Assignment(2) window.



Click **Next**. The Report window opens (Figure 10-29).

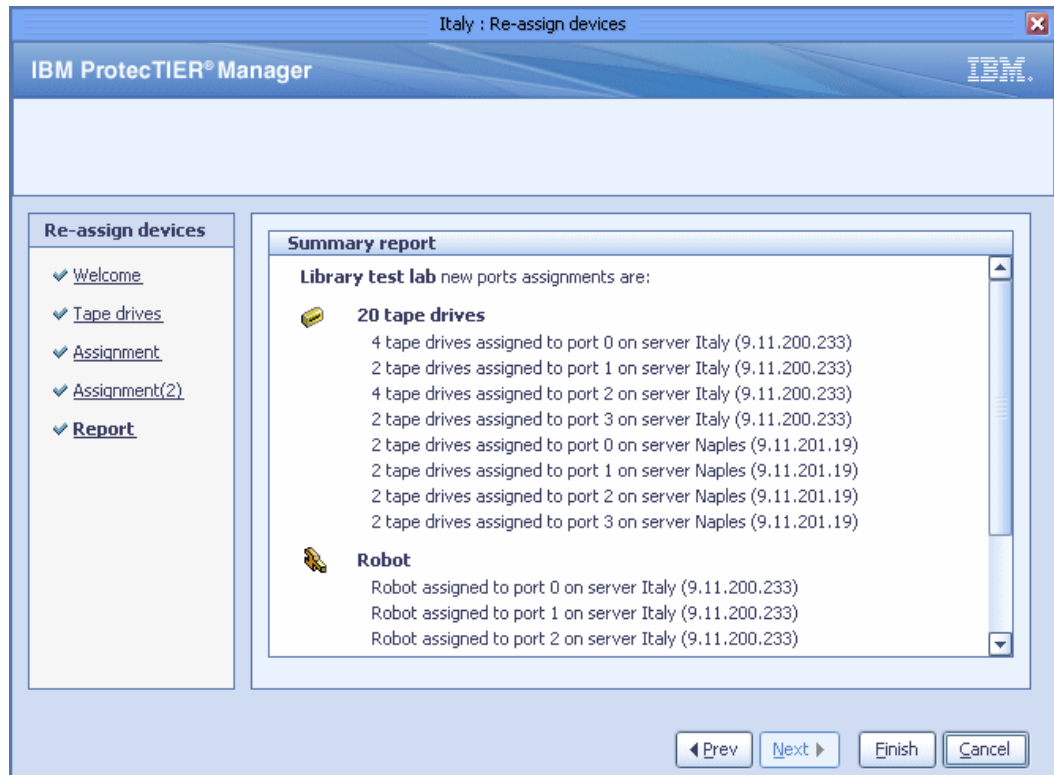


Figure 10-29 Re-assign devices: Report window

9. Click **Finish**. The Confirm Operation window opens (Figure 10-30).




Figure 10-30 Re-assign devices: Confirm operation window

10. Click **Yes**. The Re-assign devices wizard closes and the ProtecTIER system temporarily goes offline to update the library (Figure 10-31).

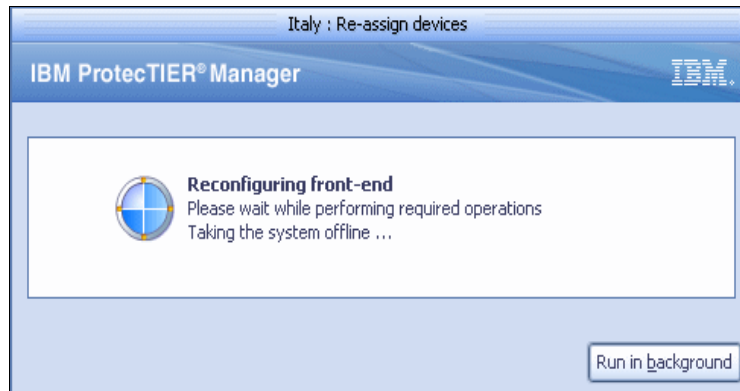


Figure 10-31 Re-assign devices: Reconfiguring the front-end window

### 10.3.3 Adding cartridges

The ProtecTIER system enables you to add and remove cartridges from your libraries.

If your virtual library has enough empty slots to accommodate the added cartridges, the adding cartridges process occurs online without disrupting backup. If the virtual library does not have enough empty slots, adding cartridges causes the ProtecTIER system to temporarily go offline to create more slots for the cartridges.

To add cartridges, complete the following steps:

1. From the menu bar, select **VT** → **VT** → **Add cartridges**. The Add cartridges wizard Welcome window opens (Figure 10-32).

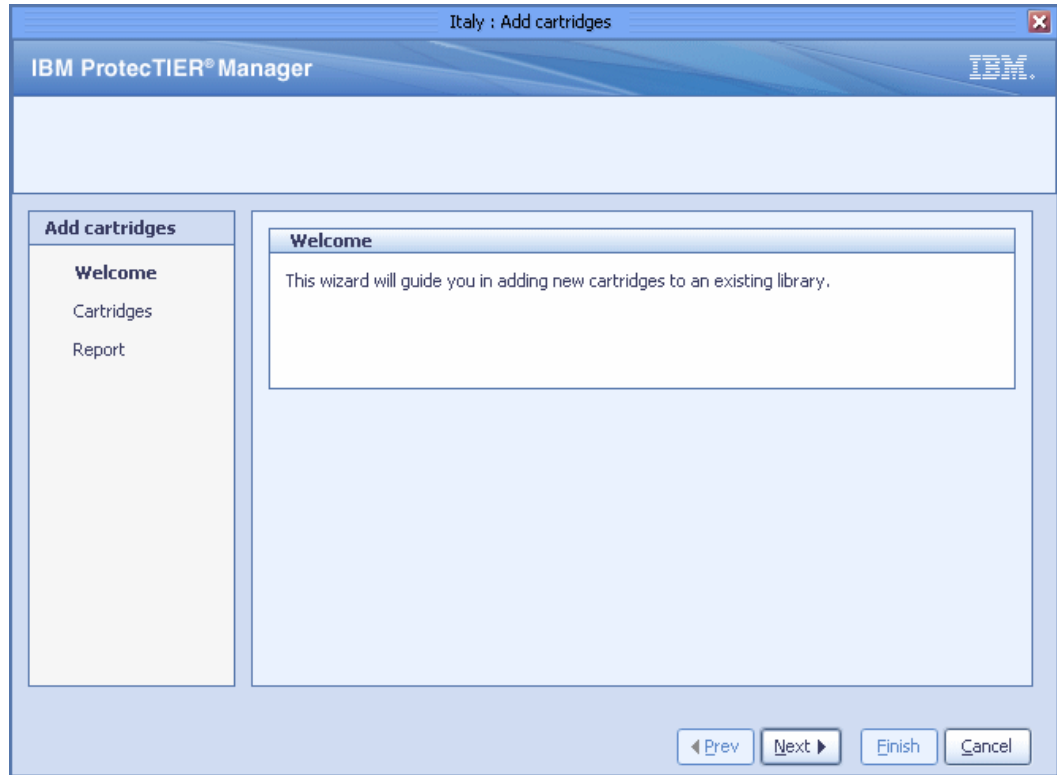


Figure 10-32 Add cartridges: Welcome window

Click **Next**. Then system validates the barcode seed (Figure 10-33).



Figure 10-33 Validating barcode seed window

Then cartridges window opens (Figure 10-34).

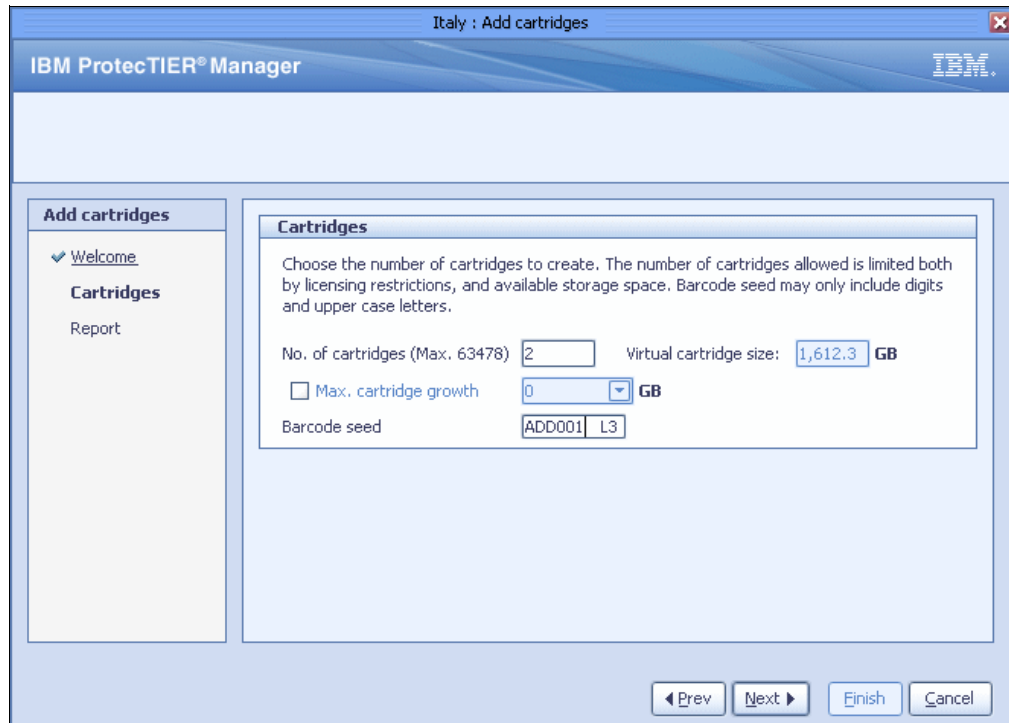


Figure 10-34 Add cartridges window

2. In the No. of cartridges field, enter the number of cartridges that you want to have in the library. The Virtual size field automatically displays the maximum possible size for virtual cartridges for your system, based on the number of cartridges entered, the total amount of available storage space in your repository, and the current HyperFactor ratio.

**Note:** Optionally, select the **Max. growth** check box. When selected, you can limit the maximum amount of nominal data that a cartridge can contain.

The value of the maximum number of cartridges possible on a system depends on the amount of storage space available on your system.

3. In the Barcode seed field, enter a value for the barcode seed. The default barcode seed is the continuation of the initial barcode seed assigned when the library was created.

**Note:** The barcode seed must contain only numbers and capital letters.

Click **Next**. The Report window opens (Figure 10-35).

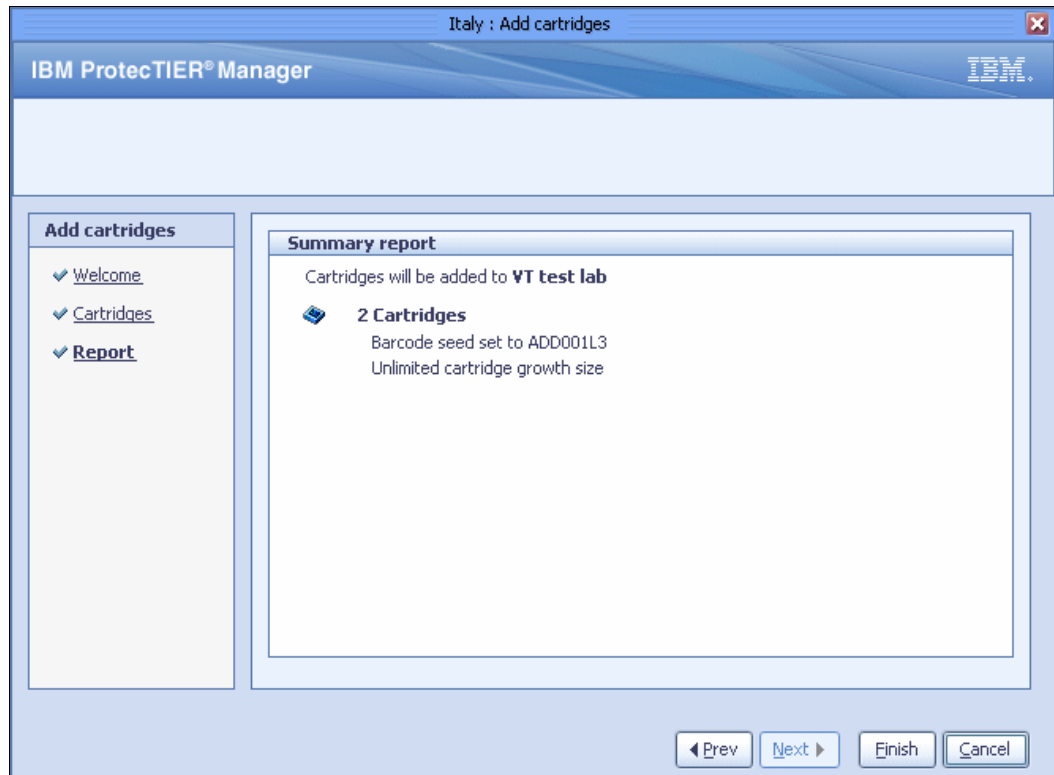


Figure 10-35 Add cartridges: Report window

4. Click **Finish**. The Add cartridges wizard closes and the cartridges are added to the library (Figure 10-36).

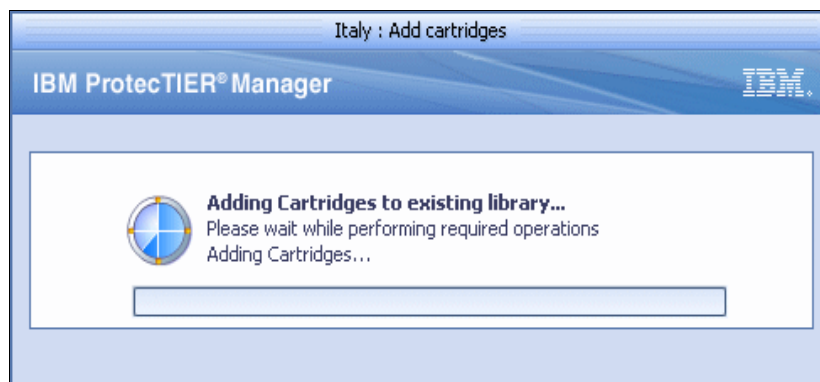


Figure 10-36 Add cartridges: Adding cartridges to an existing library window

If the necessary number of slots are not already available, the ProtecTIER system temporarily goes offline to create slots and then create the cartridges (Figure 10-37).




Figure 10-37 Confirm operation

Click **Yes** to continue (Figure 10-38).

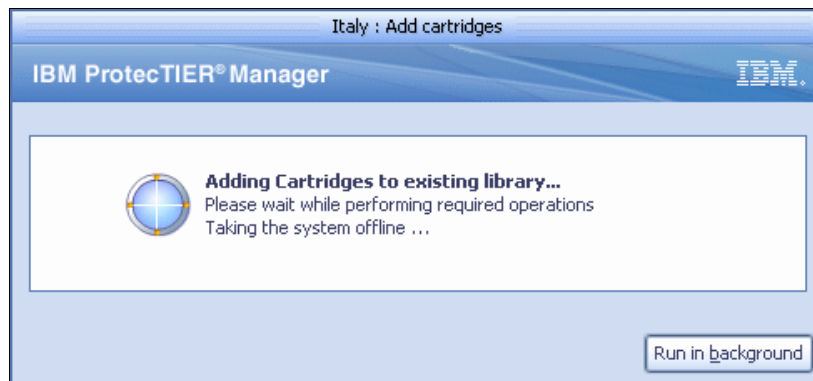


Figure 10-38 Add cartridges: System goes offline for a while

The slots and cartridges both are created (Figure 10-39).



Figure 10-39 Add cartridges: Slots and cartridges both are created

### 10.3.4 Deleting cartridges

You can delete virtual cartridges, and the results are the same as when you discard a physical tape. If the virtual volume was not in scratch status, you lose all data on that virtual cartridge and cannot recover it. If you plan to delete cartridges, check the status of the cartridges and make sure that you do not delete any cartridges that contain valid data.

**Attention:** Deleting a cartridge results in the loss of all data contained on that cartridge.

To delete a virtual cartridge, complete the following steps:

1. From the Services pane, select the library that contains the cartridge that you want to delete. The VT service monitoring window opens.
2. Click the **Cartridges** tab. The cartridges window opens (Figure 10-40).

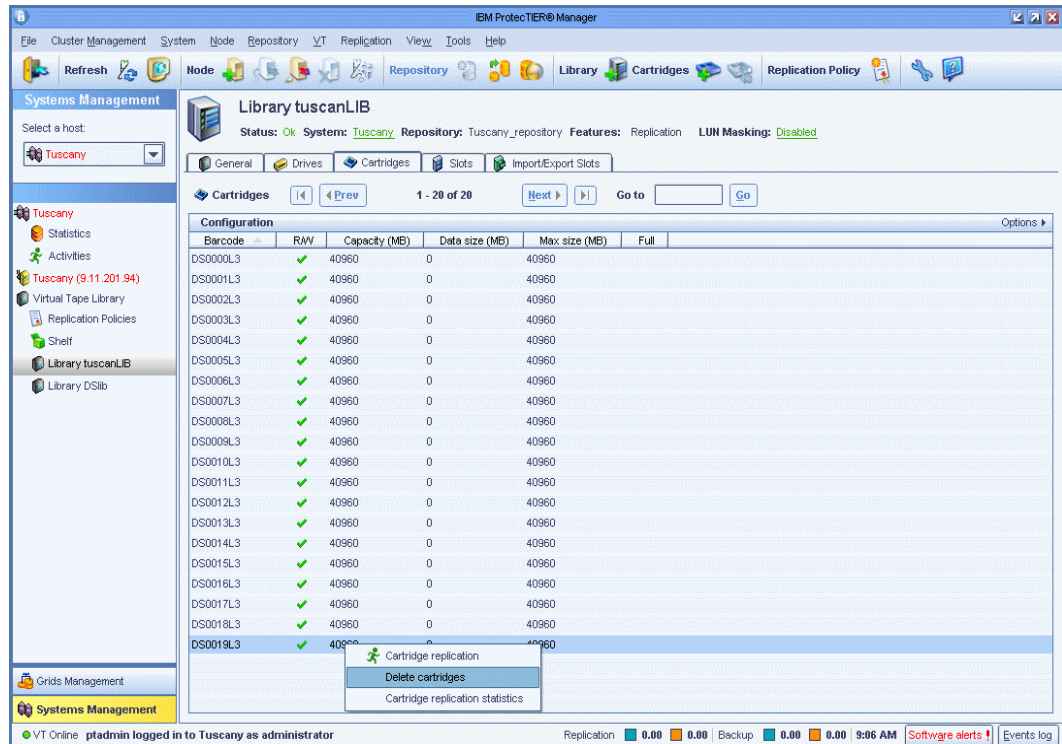


Figure 10-40 Cartridges window

3. Select one or more cartridges.
4. From the menu bar, select **VT** → **VT** → **Delete Cartridges**. A confirmation message box opens (Figure 10-41).

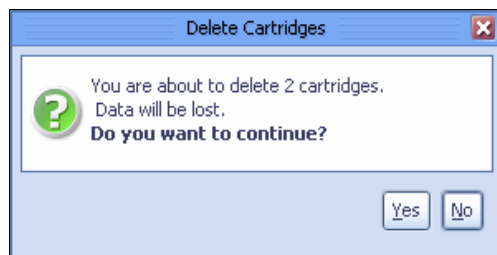


Figure 10-41 Delete cartridge confirmation window

5. Click **Yes**. The Data Loss Confirmation window opens. In the Confirm data loss window, enter “data loss” into the field and click **OK**. A confirmation message box opens (Figure 10-42).

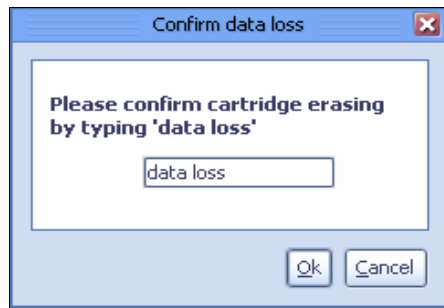


Figure 10-42 Delete cartridges: Confirm data loss window

6. Click **Yes**. While the cartridge is being deleted, the window shown in Figure 10-43 opens.




Figure 10-43 Delete cartridge confirmation window

7. The Delete cartridges window opens (Figure 10-44).



Figure 10-44 Delete cartridges window



After the deletion is complete, a confirmation window opens, as shown in Figure 10-45.

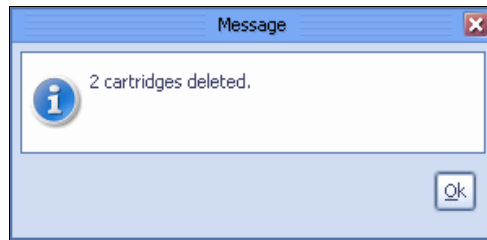


Figure 10-45 Delete cartridge message window

### 10.3.5 Switching cartridges to read-only mode

Complete the following steps to switch cartridges to read-only mode:

1. From the Services pane, select the library that contains the cartridge that you want to switch to read-only mode.
2. Click the **Slots** tab. The Slots window opens (Figure 10-46).

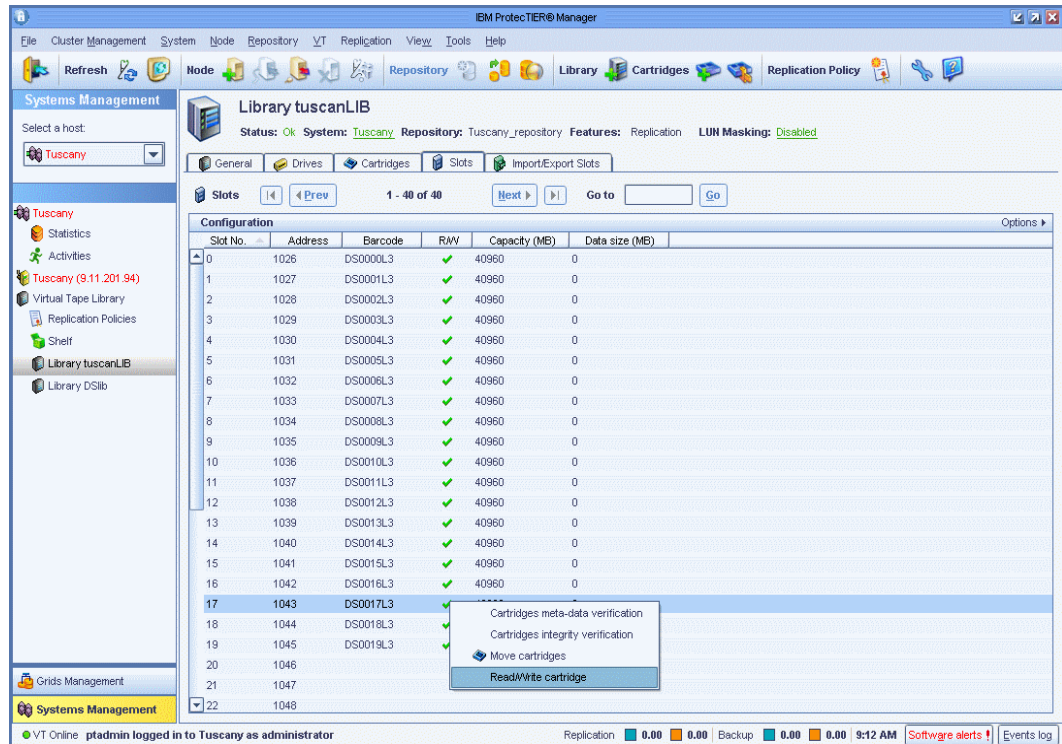


Figure 10-46 Slots window

- From the menu bar, select **VT** → **VT** → **Read/Write cartridge**, or you can right-click and select the **Read/Write cartridge** menu. A confirmation message box opens (Figure 10-47).




Figure 10-47 Read/Write cartridge window

- Click **Yes**. The cartridge switches to read-only mode (Figure 10-48).

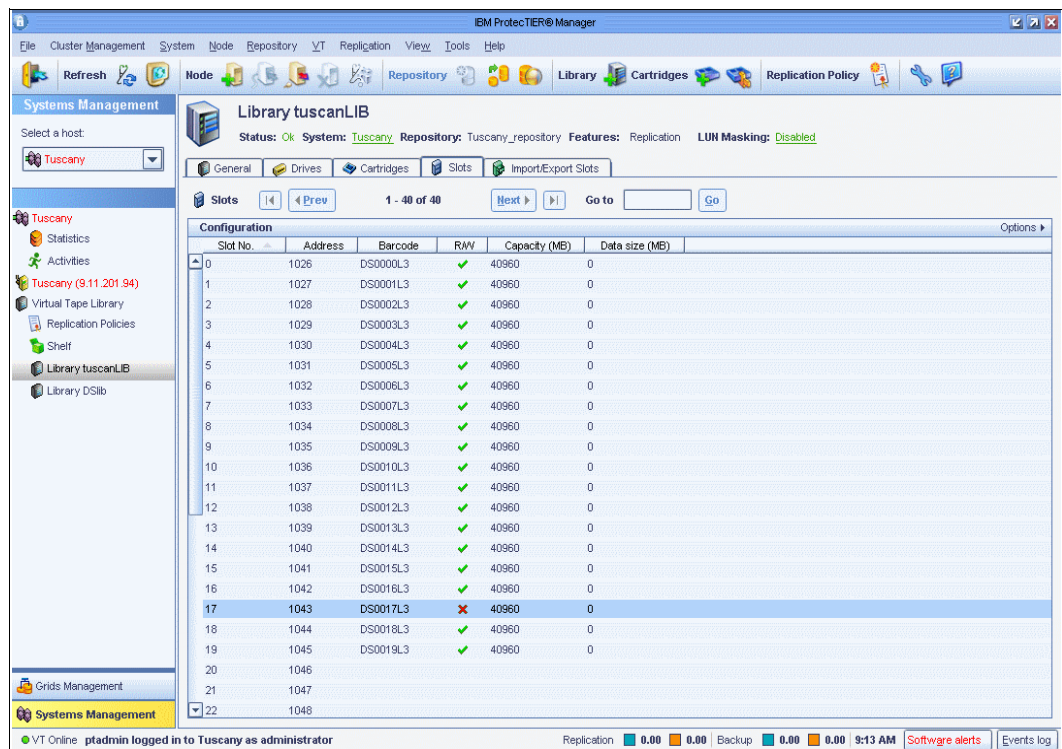


Figure 10-48 Read-only cartridge has a cross mark

**Note:** Clicking **Read/Write cartridge** for a read-only cartridge switches the cartridge to read/write mode.

## 10.3.6 Renaming libraries

The ProtecTIER system enables you to rename libraries after they have been created.

To rename a library, complete the following steps:

1. In the Services pane, select a library.
2. From the menu bar, select **VT** → **VT Library** → **Rename library** (Figure 10-49).

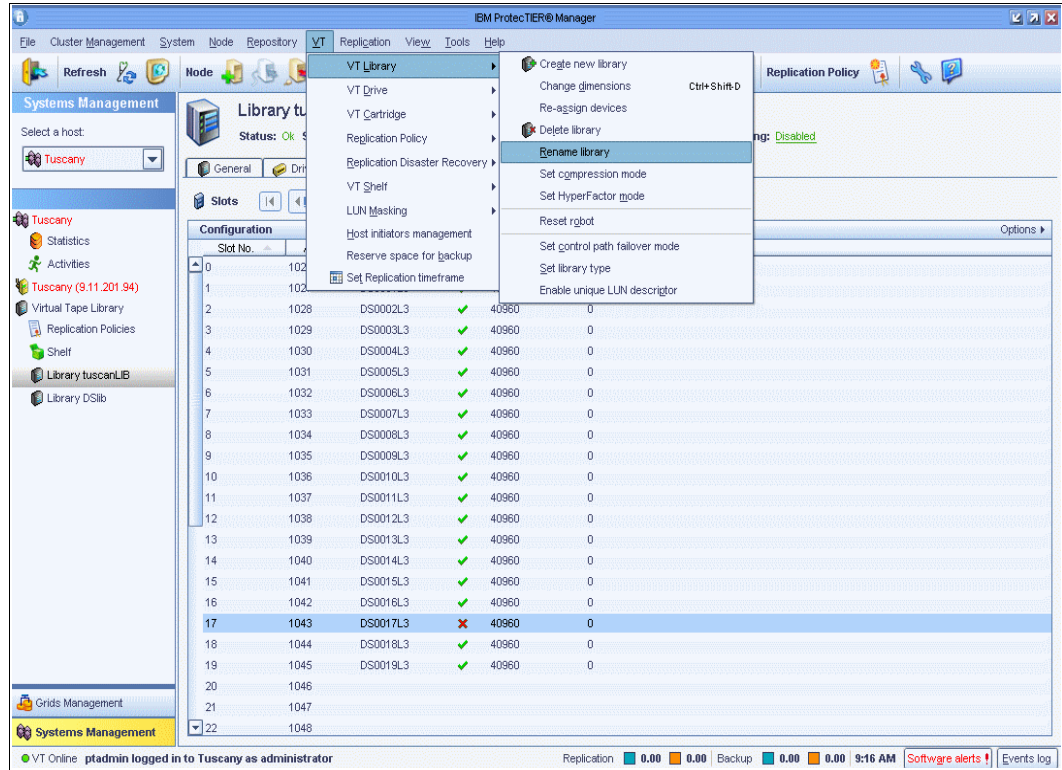


Figure 10-49 Select the virtual tape library to be renamed

The Rename library window opens (Figure 10-50).

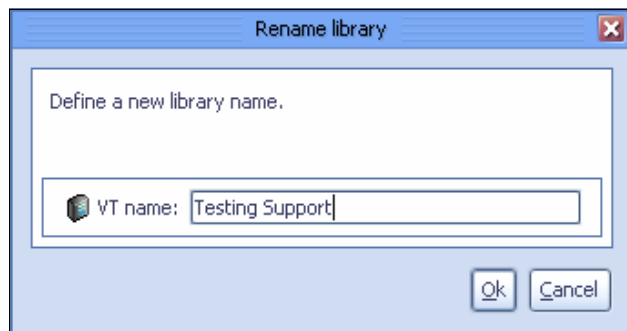


Figure 10-50 Rename library window

- Enter a new name for the selected library and click **OK**. The Rename library window closes and the library's name is changed (Figure 10-51).

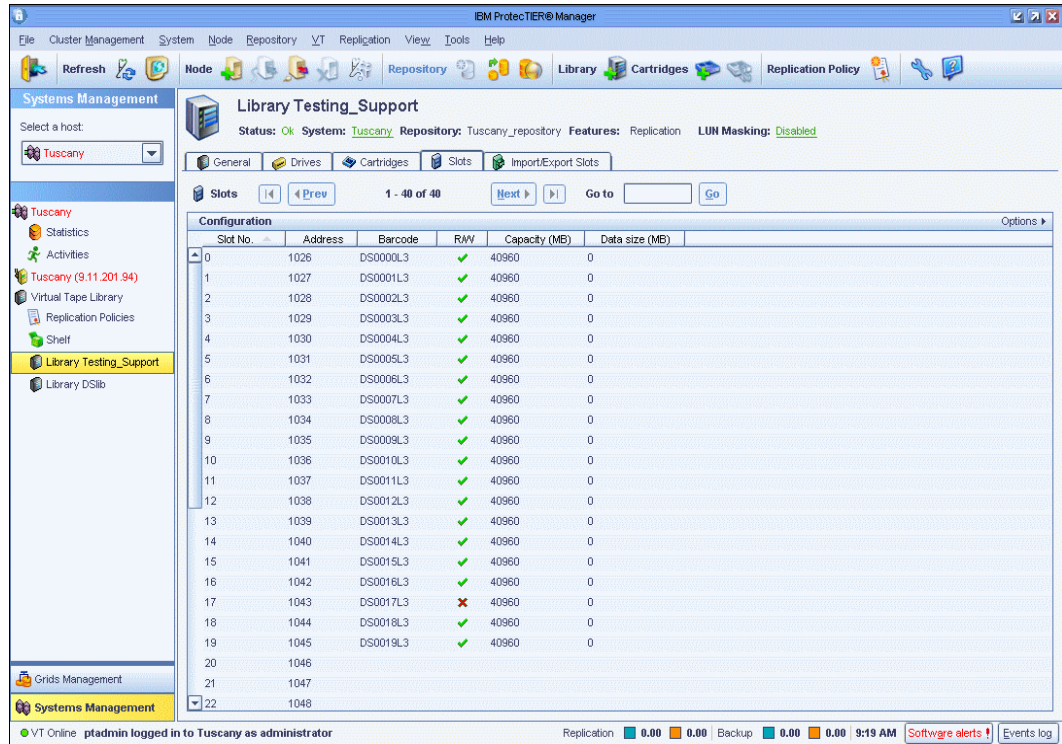


Figure 10-51 Virtual tape library is renamed

### 10.3.7 Deleting libraries

**Attention:** Deleting a library results in the loss of all data contained in that library.

To delete a library, complete the following steps:

1. In the Services pane, select a library (Figure 10-52).




Figure 10-52 Delete library

2. From the menu bar, select **VT** → **VT Library** → **Delete library**. A confirmation message box opens (Figure 10-53).

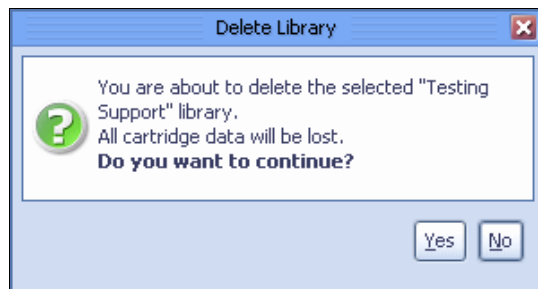


Figure 10-53 Delete Library confirmation menu

3. Click **Yes**. The Confirm data loss window opens (Figure 10-54).

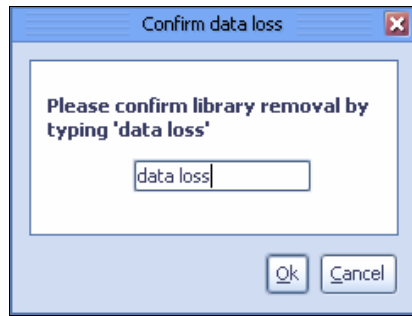


Figure 10-54 Delete library: Confirm data loss window

4. In the field of the Confirm data loss window, enter “data loss” and click **OK**. A confirmation message box opens (Figure 10-55).

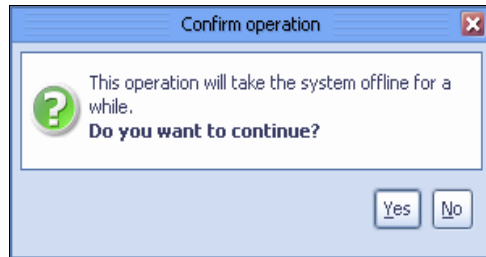


Figure 10-55 Delete library: Confirm operation window

5. Click **Yes**. The ProtecTIER system temporarily goes offline to delete the library (Figure 10-56).

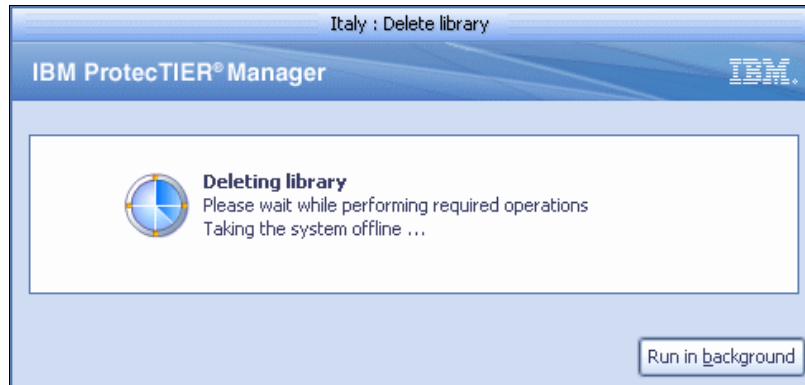


Figure 10-56 Deleting library window

**Note:** The tape cartridges whose visibility has been moved out of the deleted library to the shelf or to another tape library will be renamed.

## 10.4 Viewing the alerts and event log windows

In this section, we demonstrate some of the alerts and associated logs for the IBM ProtecTIER software.

### 10.4.1 Access alerts

In Systems Management, if an error occurs, the Alerts button on the bottom right of the View pane turns red and features a blinking exclamation mark (Figure 10-57). To simulate this situation, we create an error by removing an active member node of a two-node cluster. There are two type alerts, Software or Hardware, which will be displayed on the bottom right of view pane.

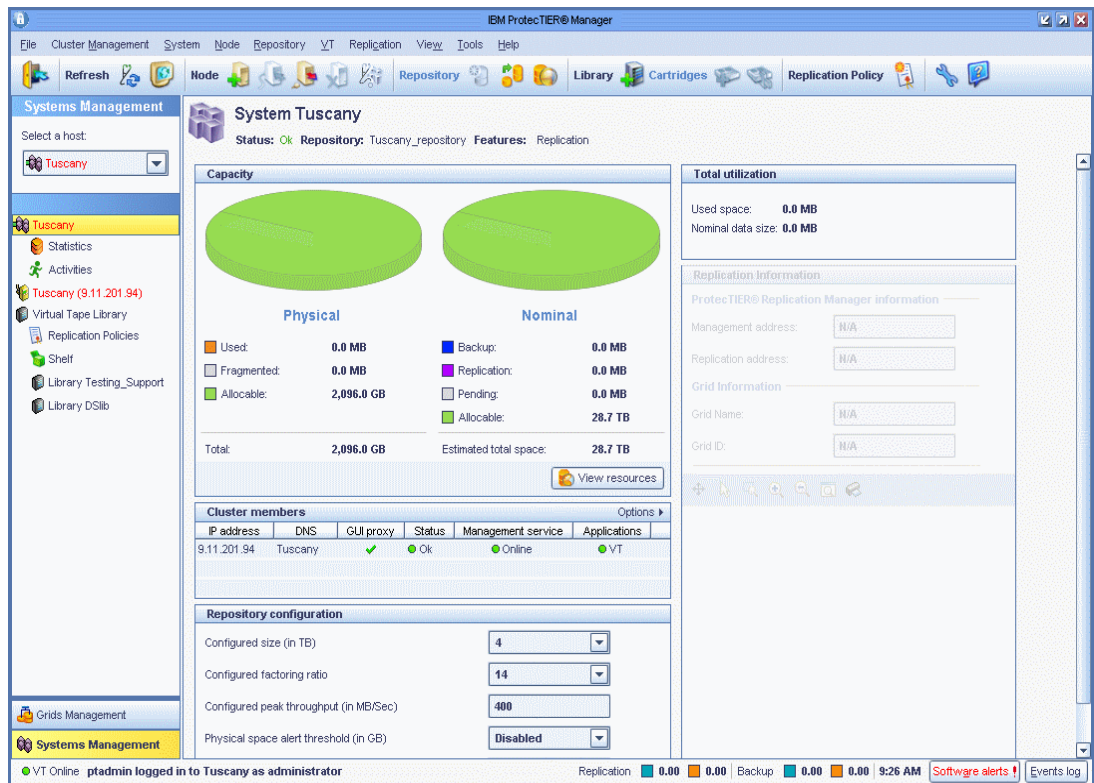


Figure 10-57 Alert button turns red

In this case, click **Software Alerts** to view the list of error events in the Alerts Log window (Figure 10-58).

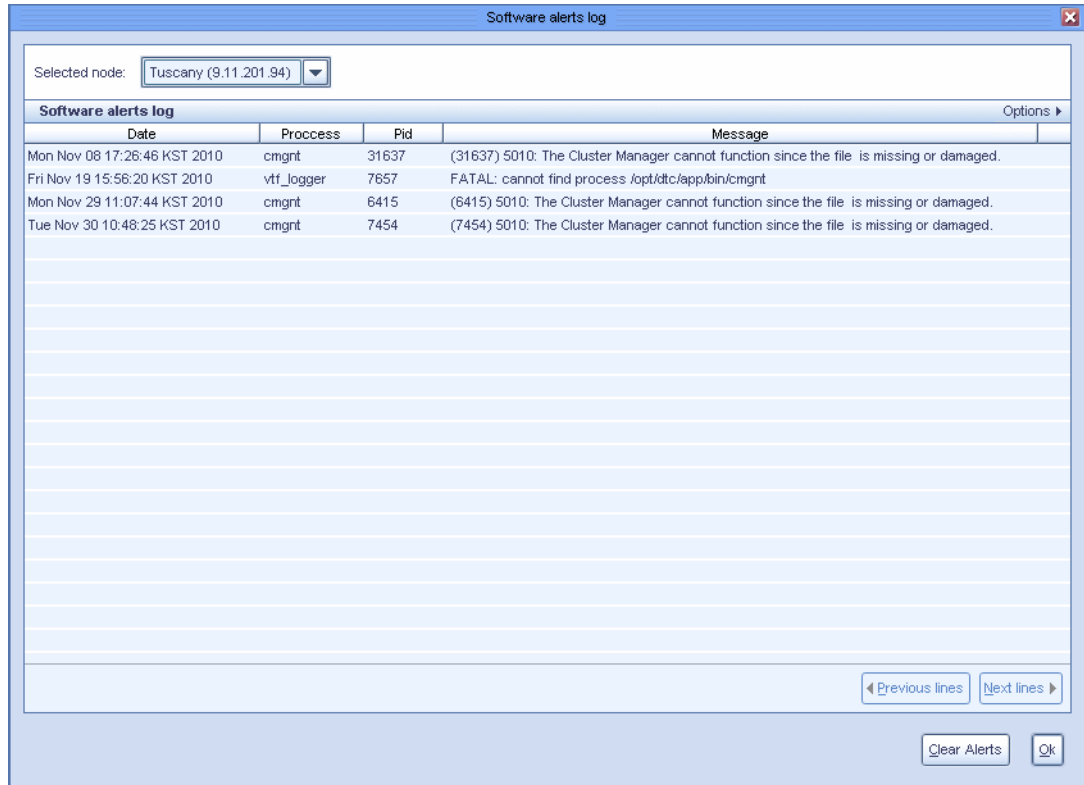


Figure 10-58 Alerts log

Click **Clear Alerts** to erase the alerts. This also causes the Alert button to turn grey again without a blinking exclamation mark.

In Grid Management, there is also an Alert button and an Alert window.



## 10.4.2 Access events

Click **Events Log** at the bottom right of the View pane to view all the events occurring in the ProtecTIER system (Figure 10-59).

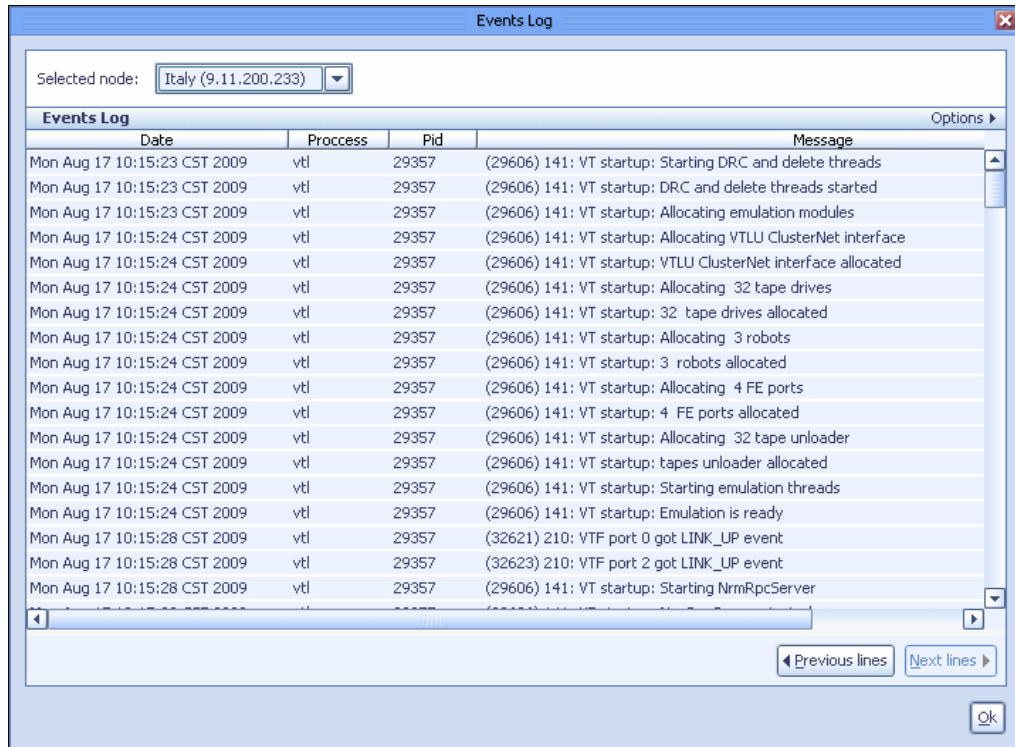


Figure 10-59 Event Log window

**Note:** The Alerts Log and Events Log windows only display information for one node at a time. In addition, the Alerts Log and Events Log windows only display up to 200 alert events at one time. Navigate to a previous set of 200 alert events by clicking **Previous lines**. Navigate to a later set of 200 alert events by clicking **Next lines**.

**Note:** You can view the error log directly on a ProtecTIER server by opening the `/pt_work/log/vtf_event.log` file. View the alerts log by opening the `/pt_work/log/vtf_error.log` file.

### 10.4.3 Grid Manager log

In Grids Management, you can access the manager log. Open Grids Management and select a grid manager from the Grids Management pane in the upper left of the window. Select **Tools** → **View manager log** (Figure 10-60).

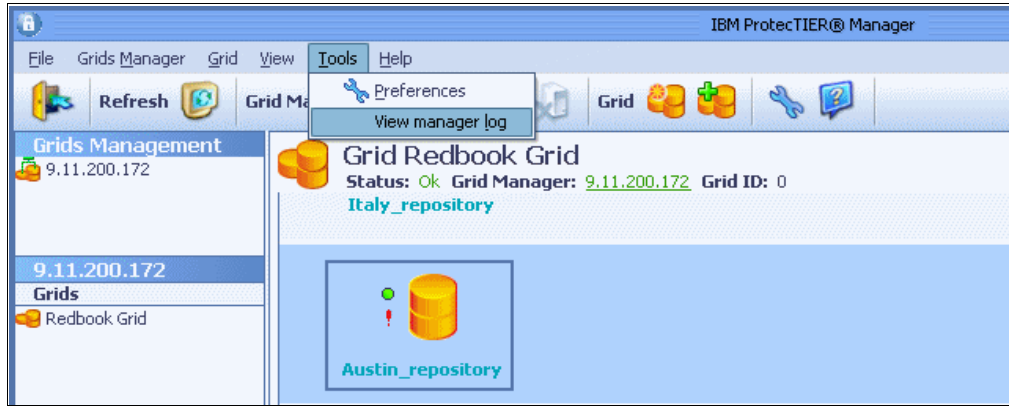


Figure 10-60 View manager log

The Manager Log Tail window opens (Figure 10-61).

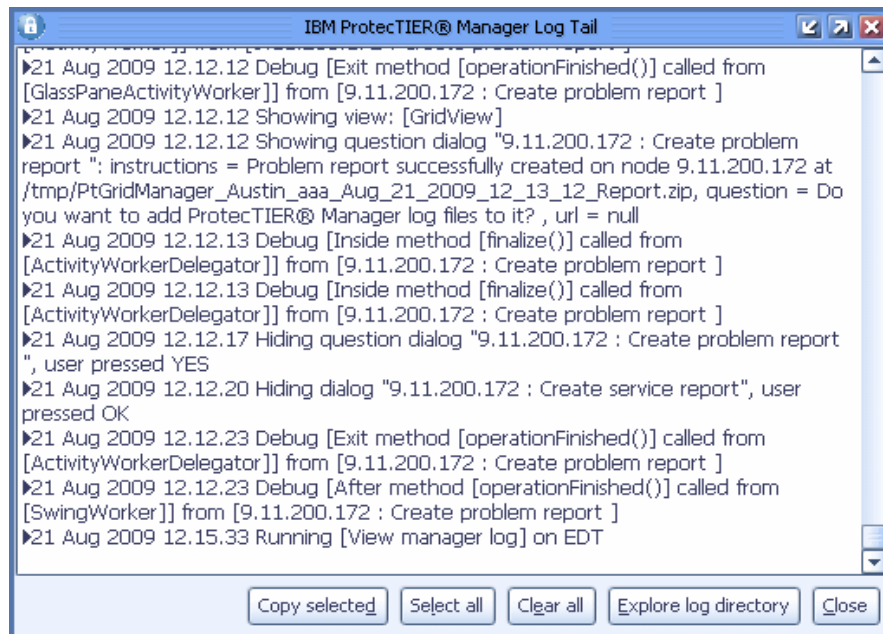


Figure 10-61 Manager Log Tail window

You can select and copy logs and pass them to a local file.

## 10.5 Wizard error messages

The ProtecTIER Manager wizards feature a message area to inform you about issues that relate to each wizard window. See Figure 10-62 and Figure 10-63.

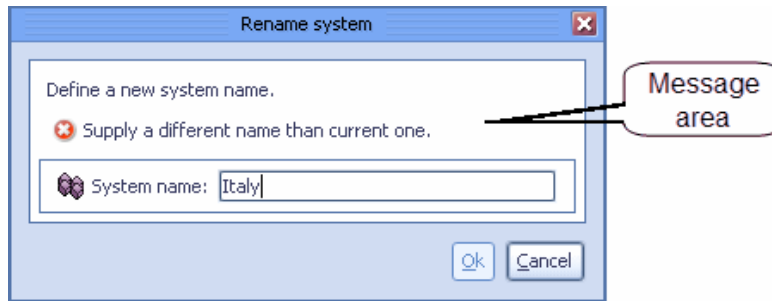


Figure 10-62 Wizard error messages window

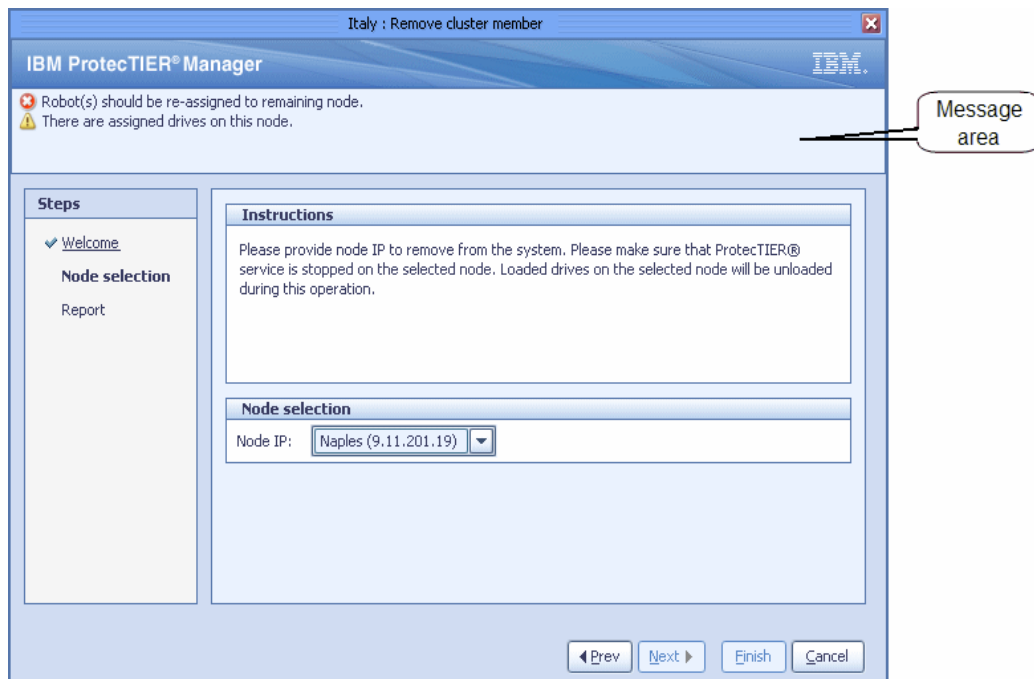


Figure 10-63 Wizard error messages window

Messages can be of the following types:

- ▶ You cannot continue the wizard until this problem is resolved.
- ▶ You can continue the wizard without resolving this issue, but it is not recommended

## 10.6 Generating a service report

IBM Support might request that you generate a service report detailing the status of your ProtecTIER system. Service report files then can be easily sent to IBM Support. The service reports can only be generated for one node at a time and they are not cluster-wide.

The service report contains the majority of the information needed to resolve problems, but not necessarily all. Core files are not collected.

**Note:** The reports are not automatically deleted by the log rotation script. The file system can fill up, as the reports tend to be large.

When a service report is created from the GUI, only an *incremental* is collected. A *full* report can only be generated from the nodes command line by running the following command:

```
/opt/dtc/app/sbin/report_problem
```

## 10.6.1 Creating a problem report in Systems Management

To create a problem report in Systems Management, complete the following steps:

1. On the Nodes pane, select a node.
2. From the menu bar, select **Node** → **Create problem report**. The Create problem report wizard Service report window opens (Figure 10-64).

Italy : Create problem report

**Fill in the service report fields to continue.**

Customer information

Customer site: IBM\_Tucson

Backup application

Host operating system: AIX5.3

Backup application name: TSM

Backup application version: 5.5

Problem description

Briefly describe the problem

Test problem report

Ok Cancel

Figure 10-64 Create problem report window

3. Click **OK**. The ProtecTIER system downloads the report files from the ProtecTIER system to the ProtecTIER Manager workstation and the Add ProtecTIER Manager log file window opens (Figure 10-65).

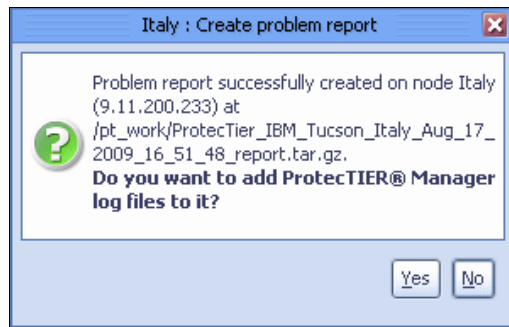


Figure 10-65 Create problem report: Add log file window

4. Optionally, click **Yes** to browse for additional files to add to the report if IBM Support requests additional log files (Figure 10-66).

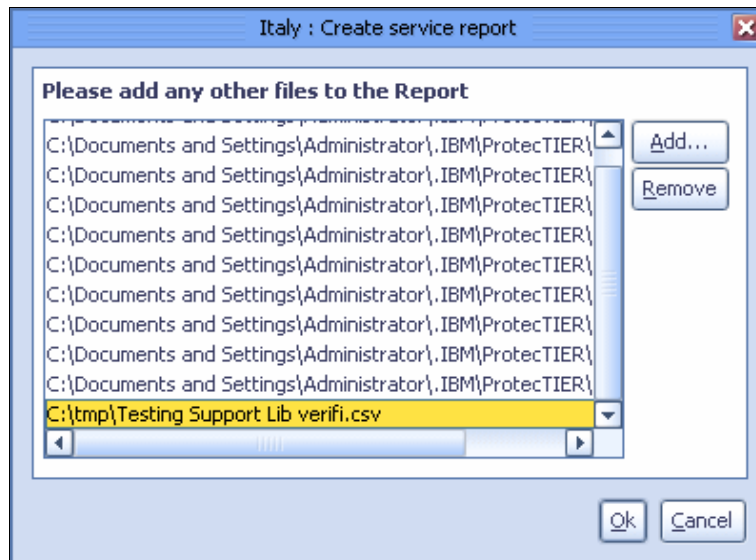


Figure 10-66 Adding additional files to the report

5. Click **No**. A standard save window opens, enabling you to save the report.zip file (Figure 10-67).

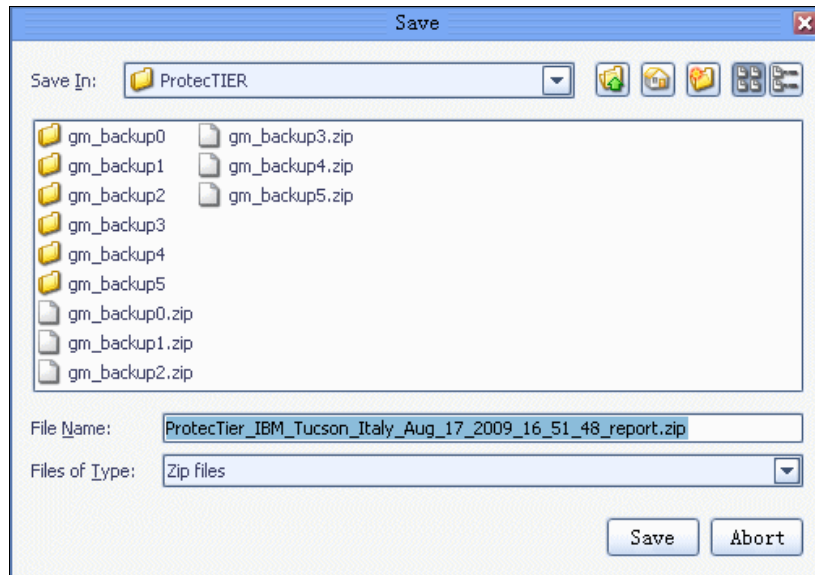


Figure 10-67 Save report.zip file window

6. Click **Save**. The report file is saved to the selected location (Figure 10-68).

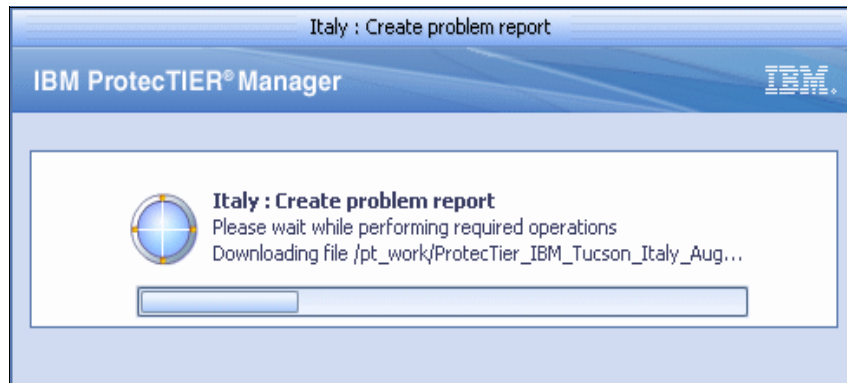


Figure 10-68 Create problem report window

7. You can get the package file from the path displayed by the system (Figure 10-69).



Figure 10-69 Package file path window

8. If the cluster contains a second node, repeat steps 1 through 5 for the second node in the two-node cluster.

**Note:** You can generate a service report directly on a ProtecTIER system by running `opt/dtc/app/sbin/report_problem`.

You can perform a system check on the server by running `sosreport`. The `sosreport` operation is time consuming and should only be used when you are directed to do so by IBM Support.

## 10.6.2 Creating a problem report in Grids Management

To create a problem report in Grids Management, complete the following steps:

1. Go to Grids Management.
2. Select a Grid Manager from the Grids Management pane.
3. Log in to the Grid Manager.
4. From the menu bar, select **Grids Manager** → **Create problem report** (Figure 10-70).

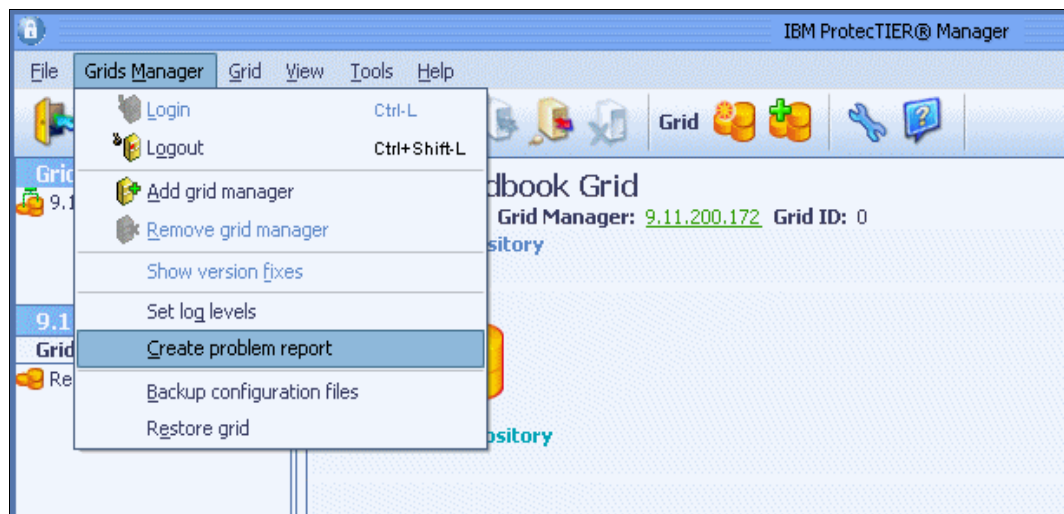


Figure 10-70 Create a problem report in Grids Manager

5. The Create problem report window opens (Figure 10-71). Give a customer name to the system.

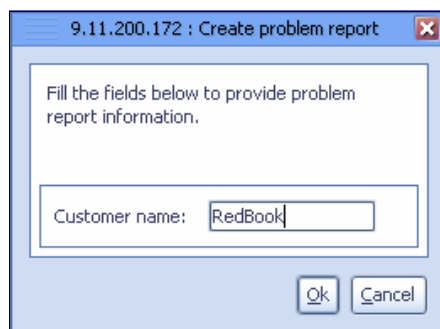


Figure 10-71 Create problem report window

6. Click **OK**. The Grids Management downloads the report files from the ProtecTIER system to the ProtecTIER Manager workstation and the Create problem report window opens (Figure 10-72).

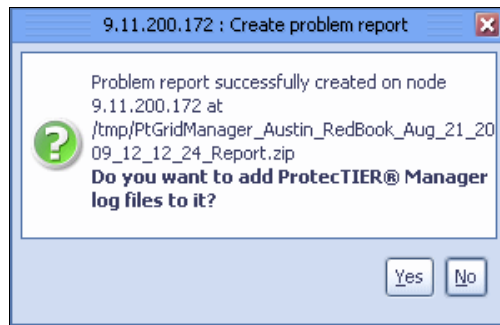


Figure 10-72 Create problem report: Add log file window

7. Optionally, click **Yes** to browse for additional files to add to the report if IBM Support requests additional log files (Figure 10-73).

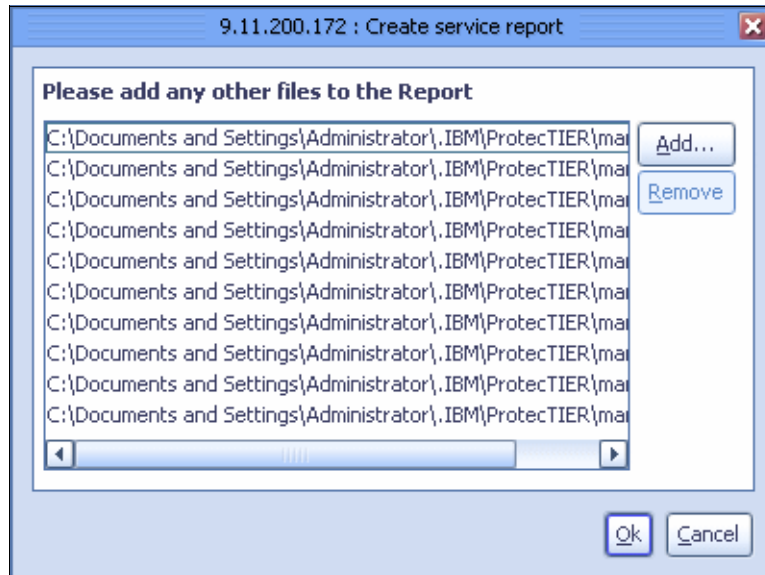


Figure 10-73 Adding additional files to the report



8. Click **No**. A standard save window opens, enabling you to save the report.zip file (Figure 10-74).

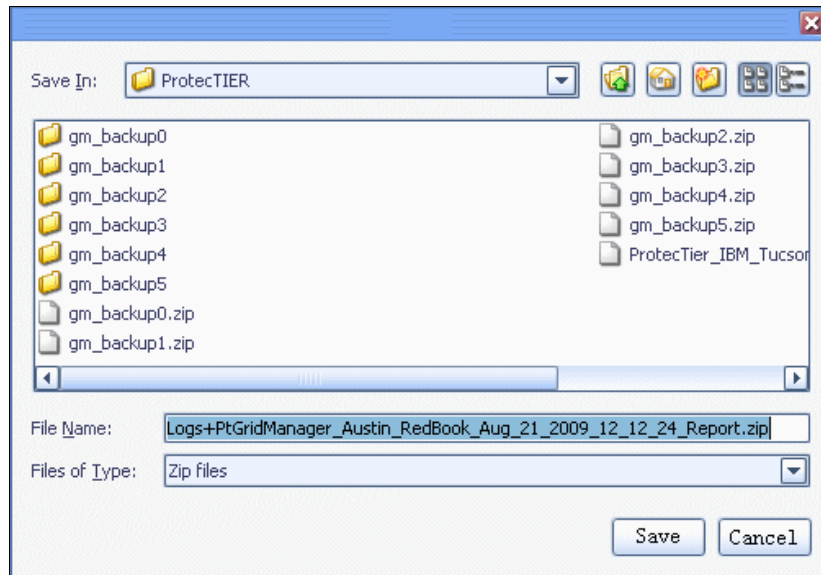


Figure 10-74 Save report.zip file window

9. Click **Save**. The report file is saved to the selected location (Figure 10-75).

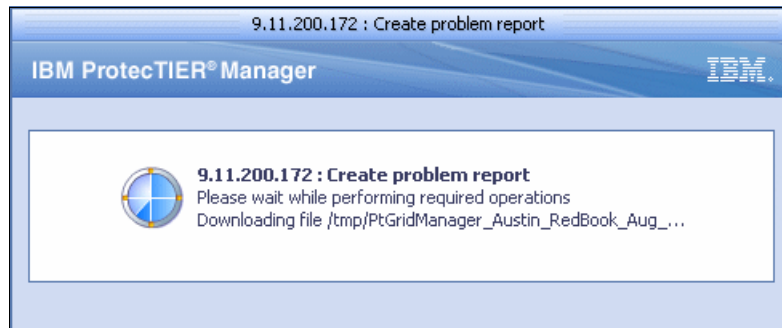


Figure 10-75 Create problem report window

10. You can get the package file from the path displayed by the system (Figure 10-76).

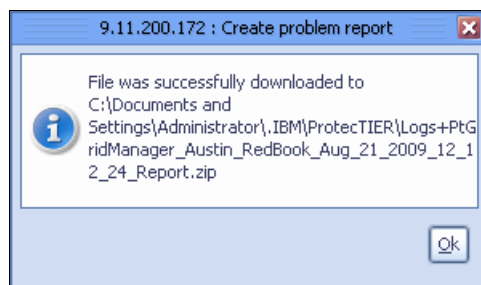


Figure 10-76 Package file path window

## 10.7 Adding and removing cluster members using the ProtecTIER Manager

This section describes how to add and remove cluster members using the ProtecTIER Manager.

### 10.7.1 Adding a cluster member

If you are using a two-node cluster, you can add a second node to an existing cluster after the cluster's first node has been added to ProtecTIER Manager. The following conditions must be met to be able to add a second node to a cluster:

- ▶ A single node cluster is established with an associated repository.
- ▶ The cluster-internal network between the nodes is functional.
- ▶ The Red Hat Cluster Suite is configured as a two-node cluster for each node.
- ▶ The new node is not already part of a cluster.
- ▶ The new node has no repository.
- ▶ ProtecTIER is online in the new node.
- ▶ The new node has full access to all the file systems of the repository on the first node.
- ▶ There is software version compatibility between the two-nodes.

For more information, refer to *IBM System Storage TS7600 with ProtecTIER Installation Guide*, GC53-1155.

To add a cluster member, the administrator permission level is needed. Complete the following steps:

1. From the Systems management pane, select the cluster to which you want to add a member.

2. Select **Cluster Management** → **Add Cluster member**. The Add cluster member wizard Welcome window opens (Figure 10-77).

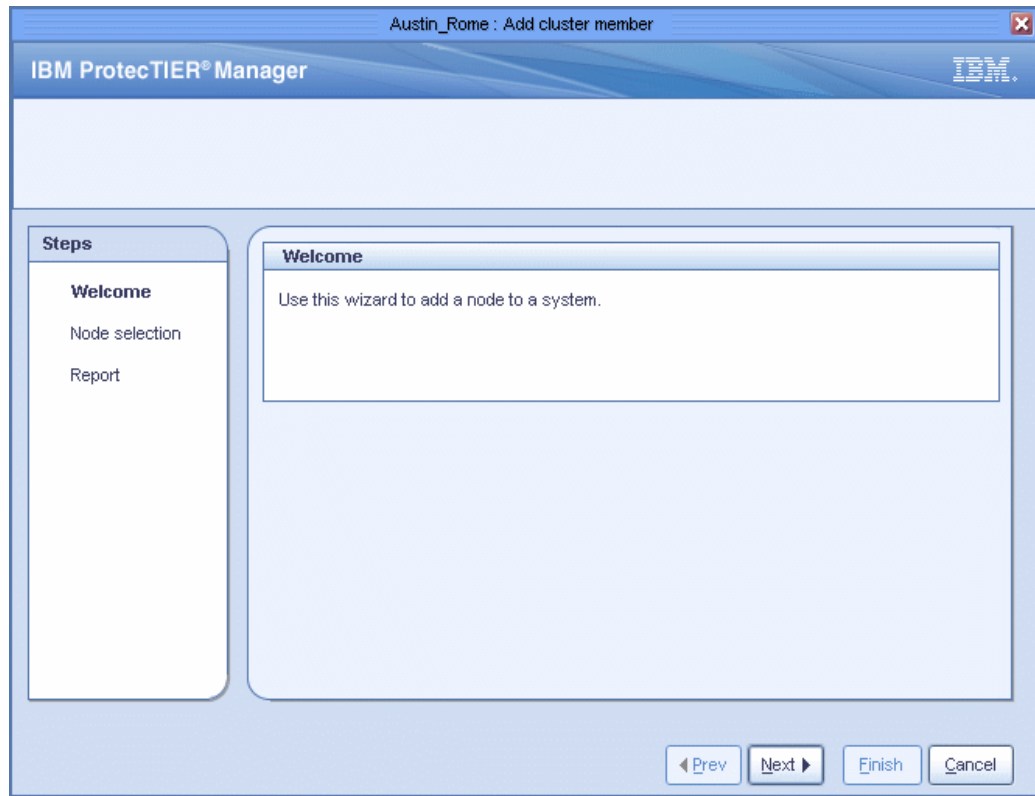


Figure 10-77 Add cluster member: Welcome window

Click **Next**. The Node Selection window opens.

3. In the Node IP field, enter the IP address of the node that you want to associate with the cluster or select the IP address from the drop-down list (Figure 10-78).

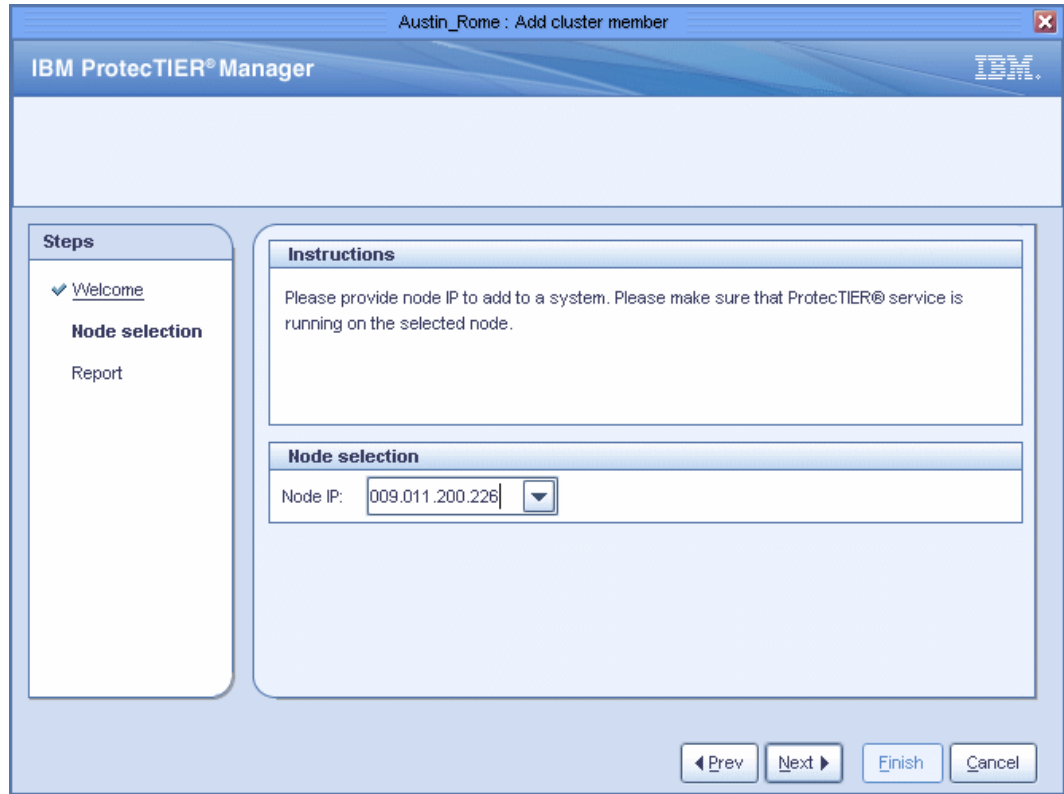


Figure 10-78 Add cluster member: Node selection window

Click **Next**. The validation node IP window opens (Figure 10-79).

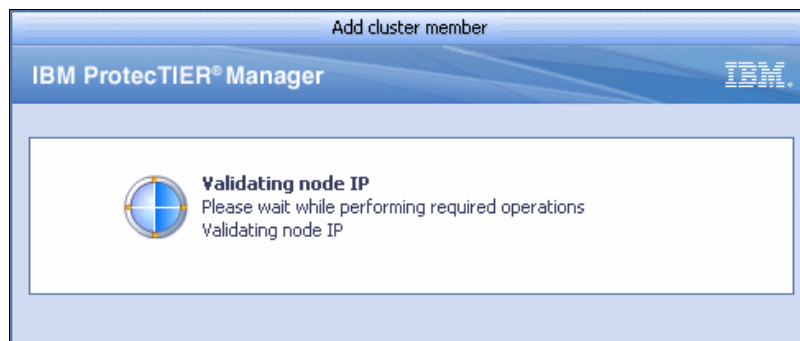


Figure 10-79 Add cluster member: Validating node IP window

4. After successful validation of the node IP, the Report window opens (Figure 10-80).

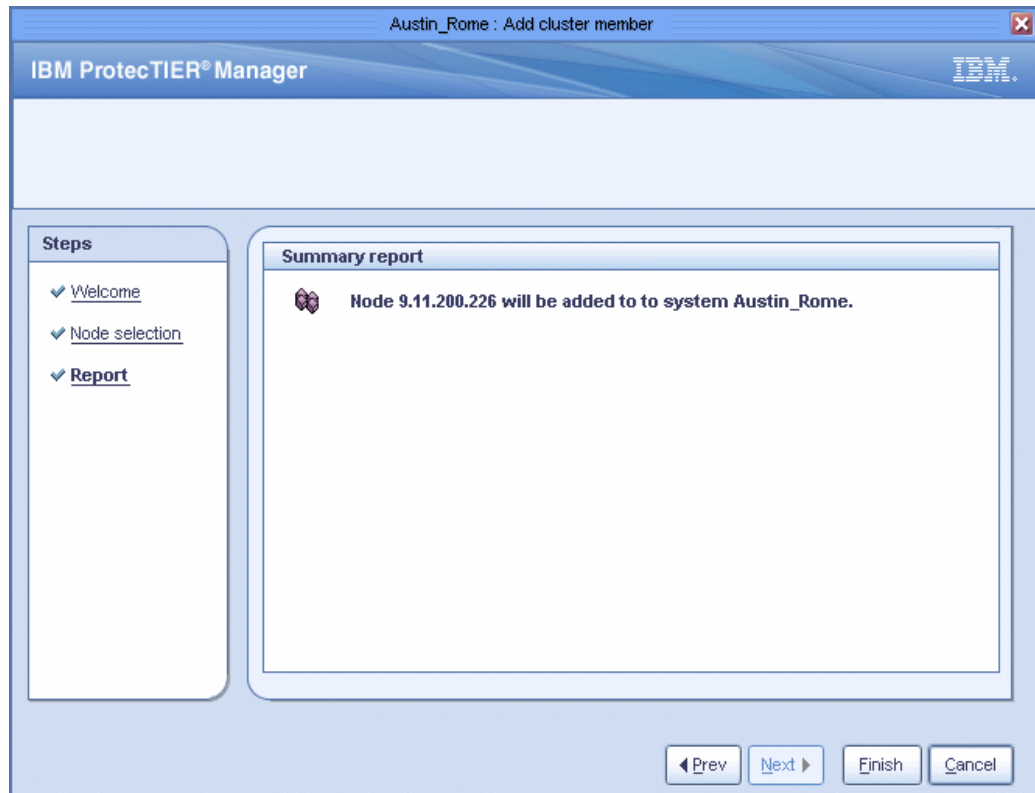


Figure 10-80 Add cluster member: Report window

- Click **Finish**. The Add cluster member wizard closes and the node is added to the selected cluster (Figure 10-81).

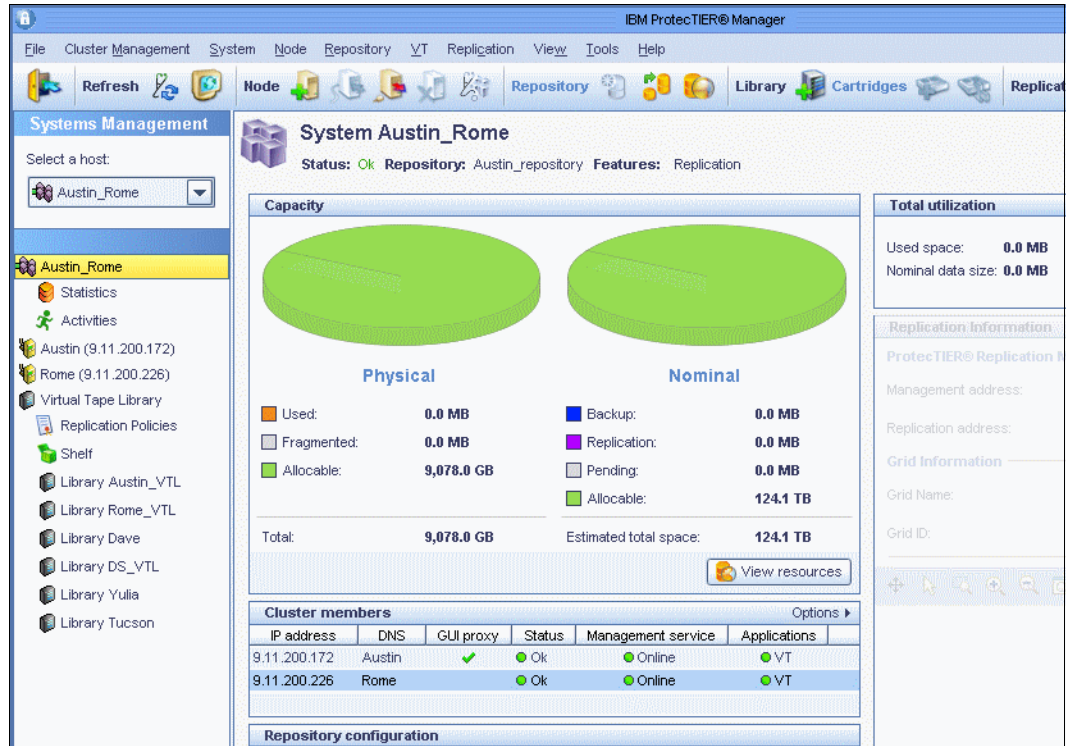


Figure 10-81 ProtecTIER Manager window with added cluster member

## 10.7.2 Removing a cluster member

You can only remove a node from a two-node cluster if the system has been stopped on that node. If there is only one node in the cluster, that node cannot be removed from the cluster. If you remove a node on which Virtual Tape (VT) service devices are assigned, the devices remain offline and unassigned until they are assigned to an active cluster member node. If there are cartridges loaded in drives on that node, the cartridges are automatically unloaded.

**Attention:** If you want to remove a cluster member to which *tape drives* are assigned, first reassign the tape drives on all libraries. Nodes that have robots assigned to them cannot be removed from a cluster until *the robot* is reassigned. For more information, refer to 10.3.2, “Reassigning devices” on page 494.

To remove a cluster member, the administrator permission level is needed. Complete the following steps:

1. From the nodes pane, select the node that you want to remove from the cluster (Figure 10-82).

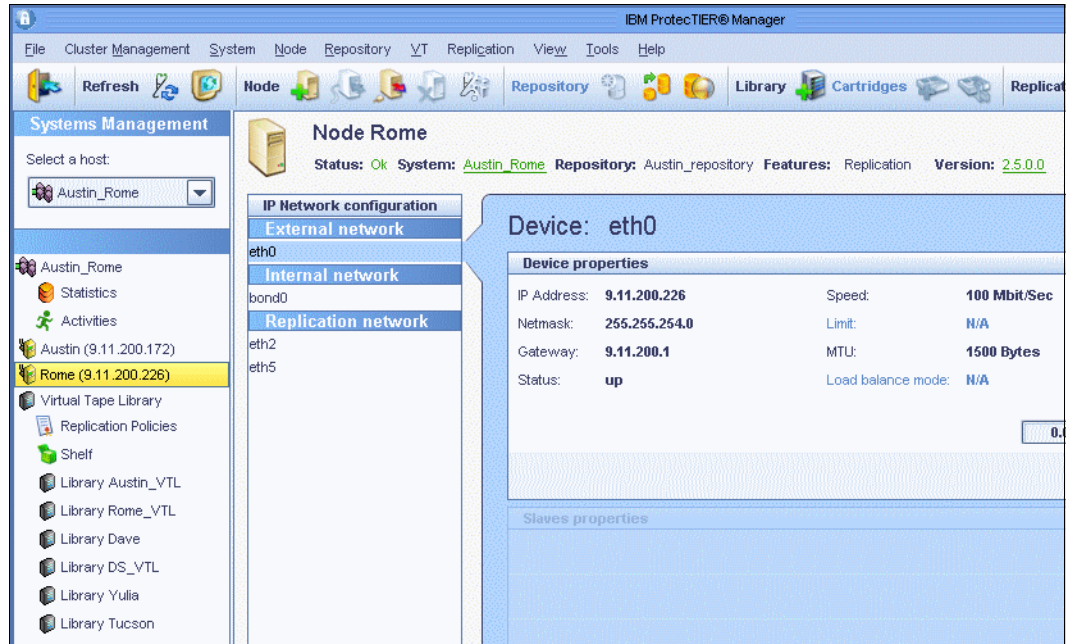


Figure 10-82 Remove cluster member select node window

2. Stop the server (you will be logged out from the system after that operation). For more information, refer to 10.8.1, “Starting and stopping the server” on page 535.

- In the Systems tab of the Navigation pane, select a two-node cluster from which you want to remove a member (Figure 10-83).

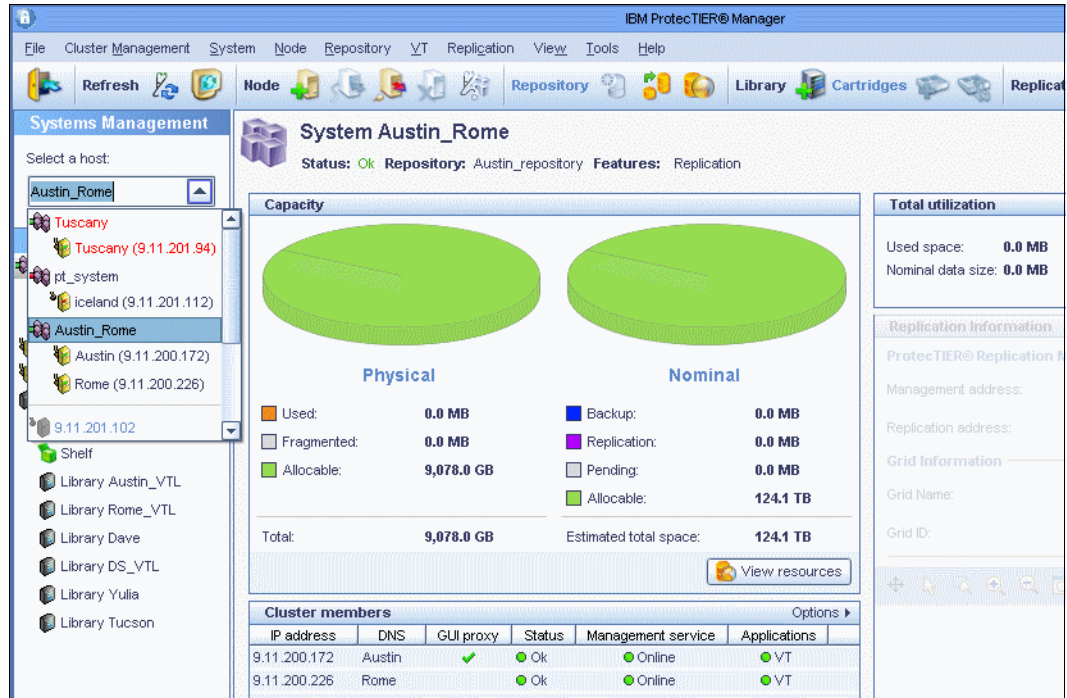


Figure 10-83 Remove cluster

- Log in to that system. For more information, refer to 5.3.2, “Logging in and out” on page 192.



- Verify that the selected node is offline (in the Nodes pane, the node appears in grey). In the Nodes pane, select the stopped node (Figure 10-84).

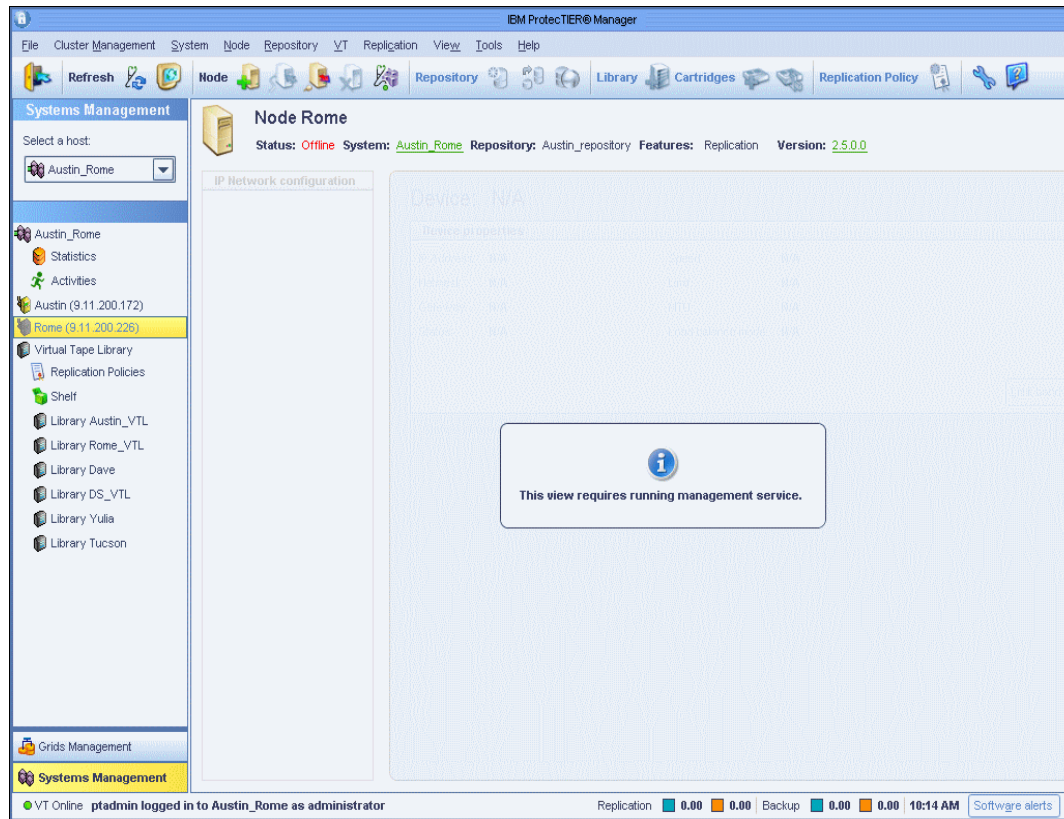


Figure 10-84 Remove cluster member select stopped node window

- From the menu bar, select **Cluster Management** → **Remove cluster member**. The Remove cluster member wizard Welcome window opens (Figure 10-85).

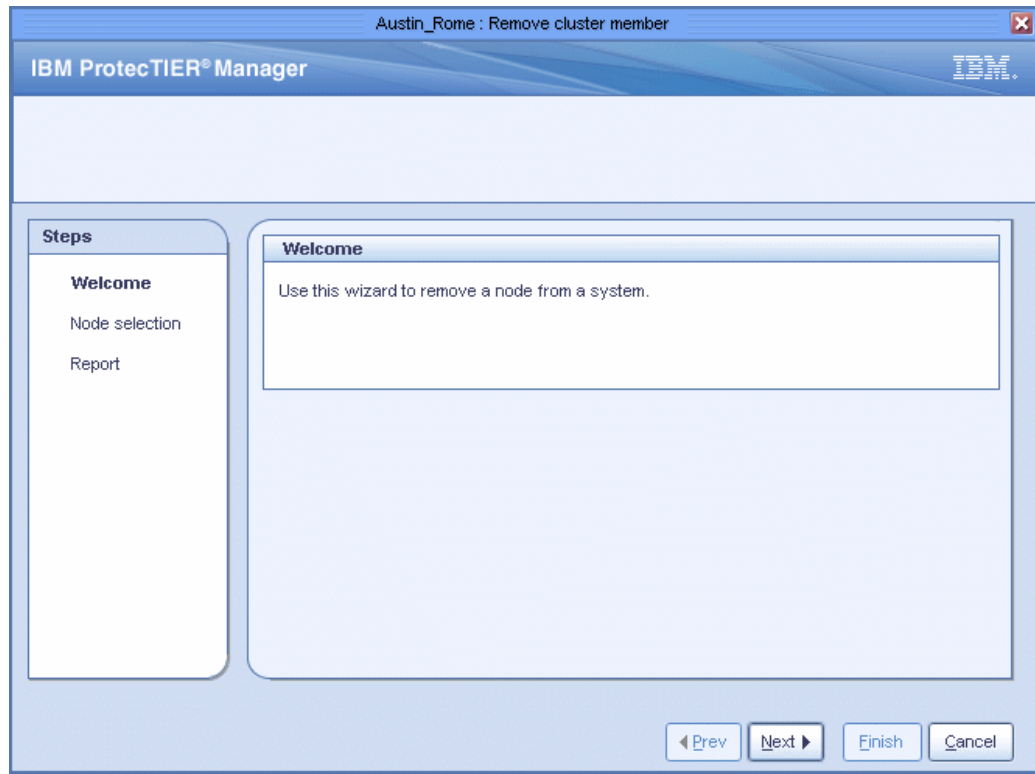


Figure 10-85 Remove cluster member: Welcome window

Click **Next**. The Node selection window opens (Figure 10-86).

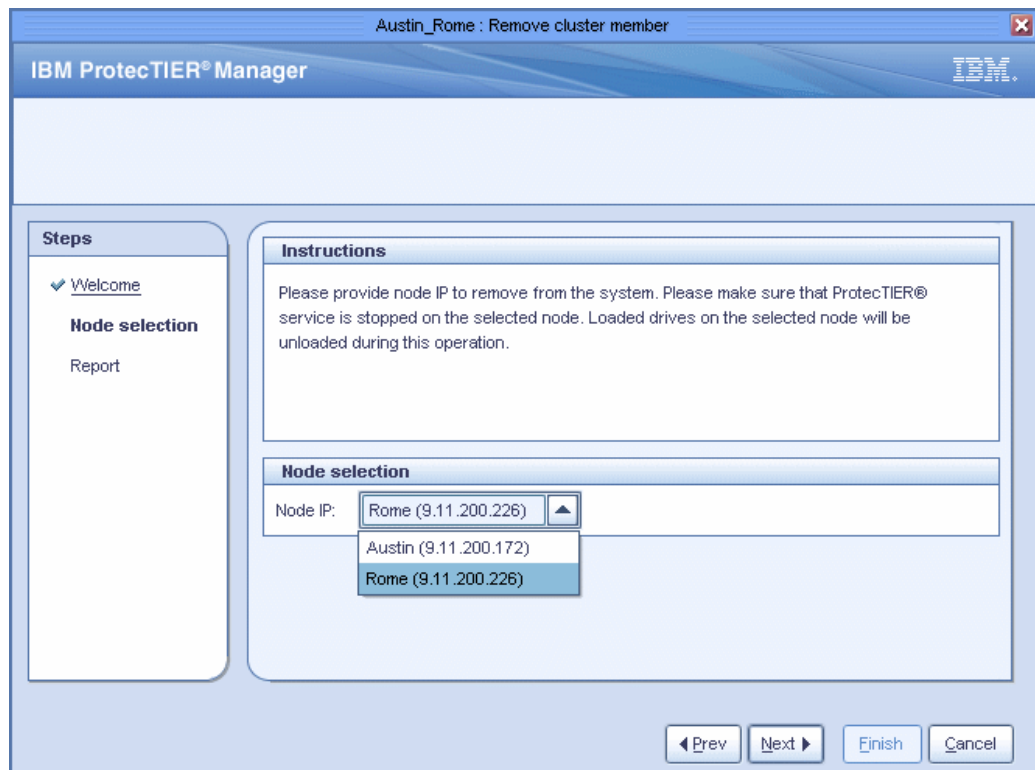


Figure 10-86 Remove cluster member: Node selection window

7. In the Node IP field, select the node that you want to remove from the two-node cluster.

Click **Next**. The Report window opens (Figure 10-87).

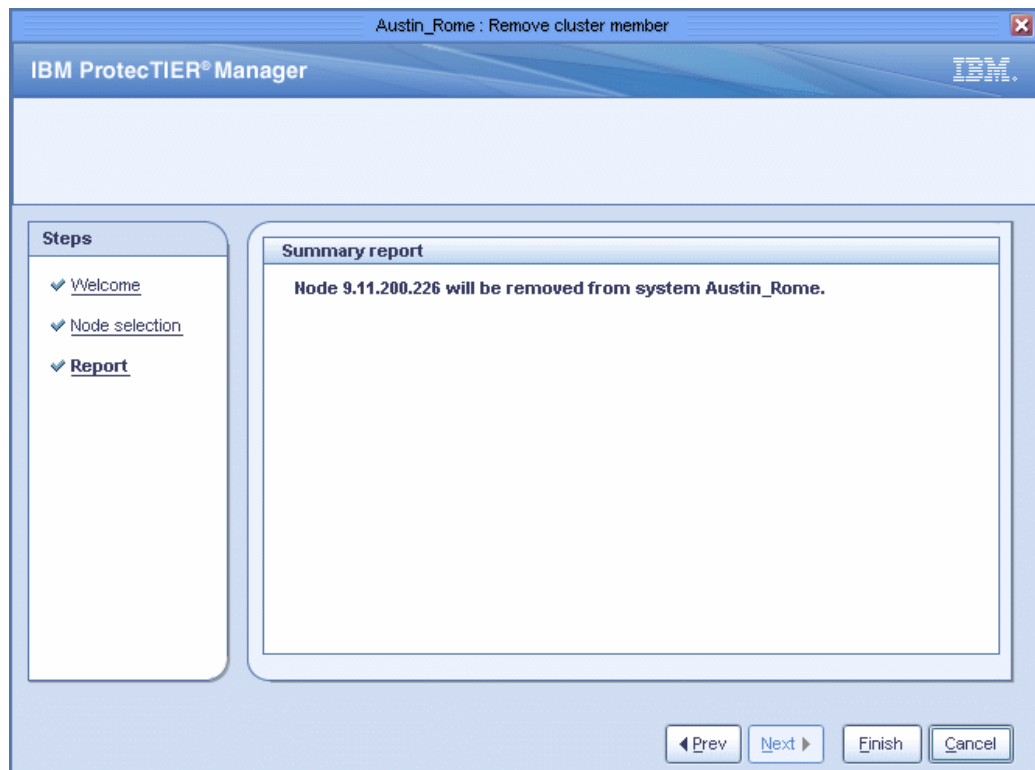


Figure 10-87 Remove cluster member report window

- Click **Finish**. The Remove cluster member wizard closes and the selected node is removed from the two-node cluster (Figure 10-88).

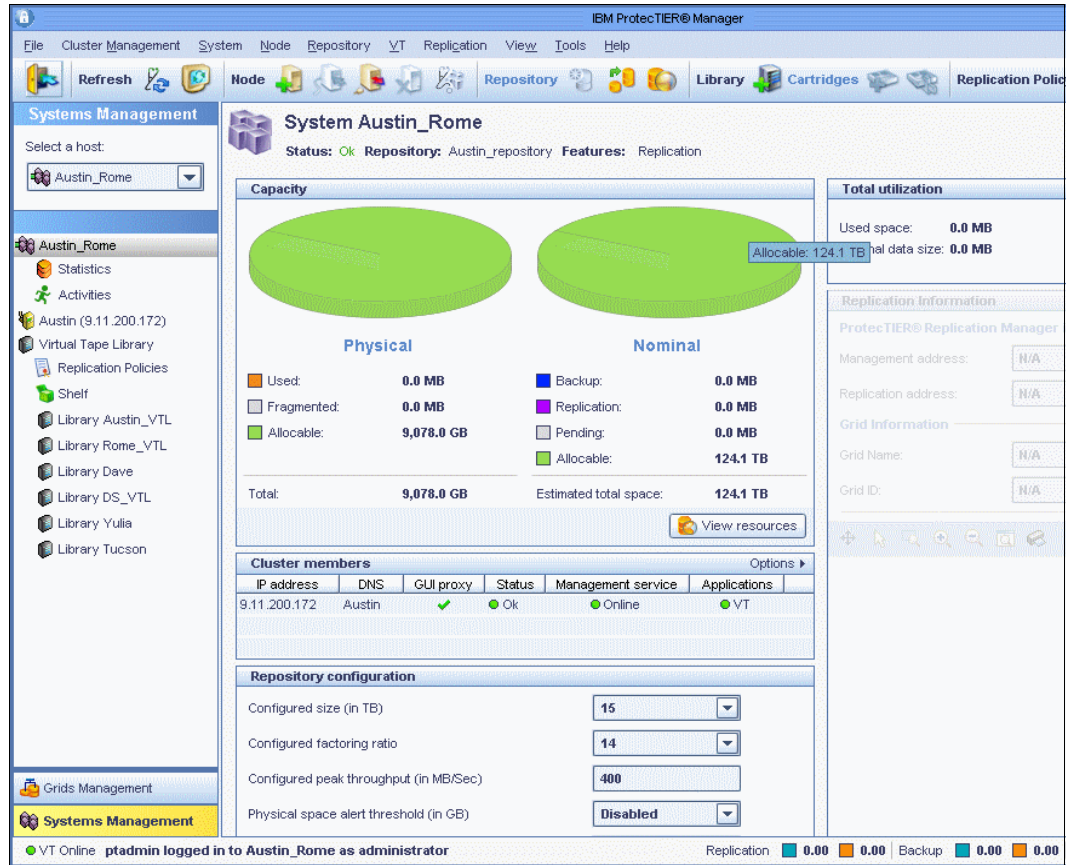


Figure 10-88 ProtecTIER Manager window without the removed cluster member

## 10.8 Common maintenance tasks

This section contains a list of miscellaneous tasks that IBM Support personnel might ask you to perform during the troubleshooting process.

**Note:** These tasks should not be performed unless you are directed to do so by IBM Support personnel.

### 10.8.1 Starting and stopping the server

To start and stop the server, the administrator permission level is needed.

**Note:** Starting and stopping the server only impacts the drives and robots assigned to the selected server.

#### Starting the server

To start the server, complete the following steps:

- From the Nodes pane, select the node on which you want to start the server.

2. Select **Node** → **Server** → **Start Server** (Figure 10-89).




Figure 10-89 Start server

3. Click **Start server**. A window with login information opens (Figure 10-90). This action happens even you are logged in already. Enter the user name and password.

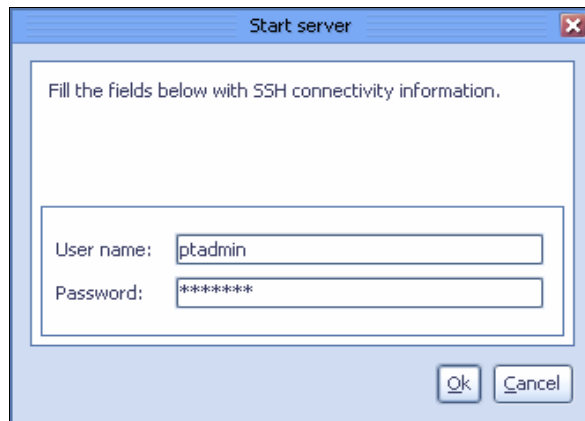


Figure 10-90 Start server window

The server starts (Figure 10-91).

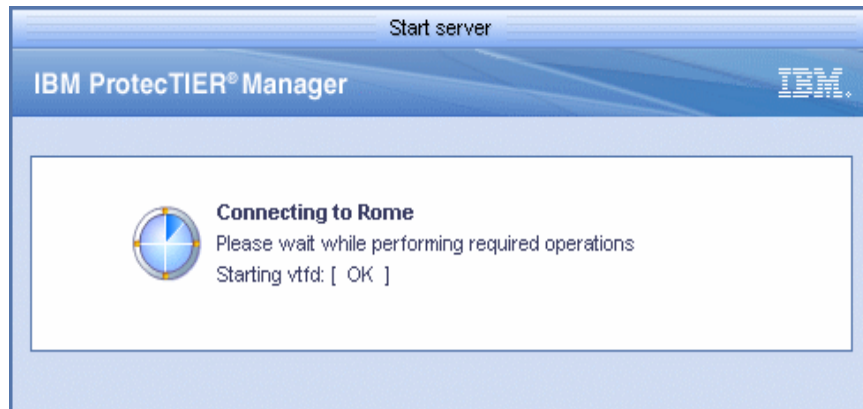


Figure 10-91 Start server: Performing start window

4. A message displays and states that the service will be started. Click **OK** (Figure 10-92).

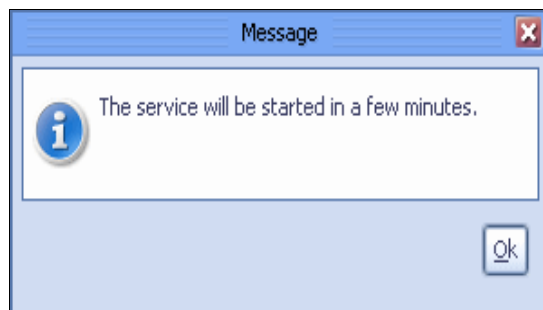


Figure 10-92 Start server message window

The server will be started and the window will be refreshed (Figure 10-93).

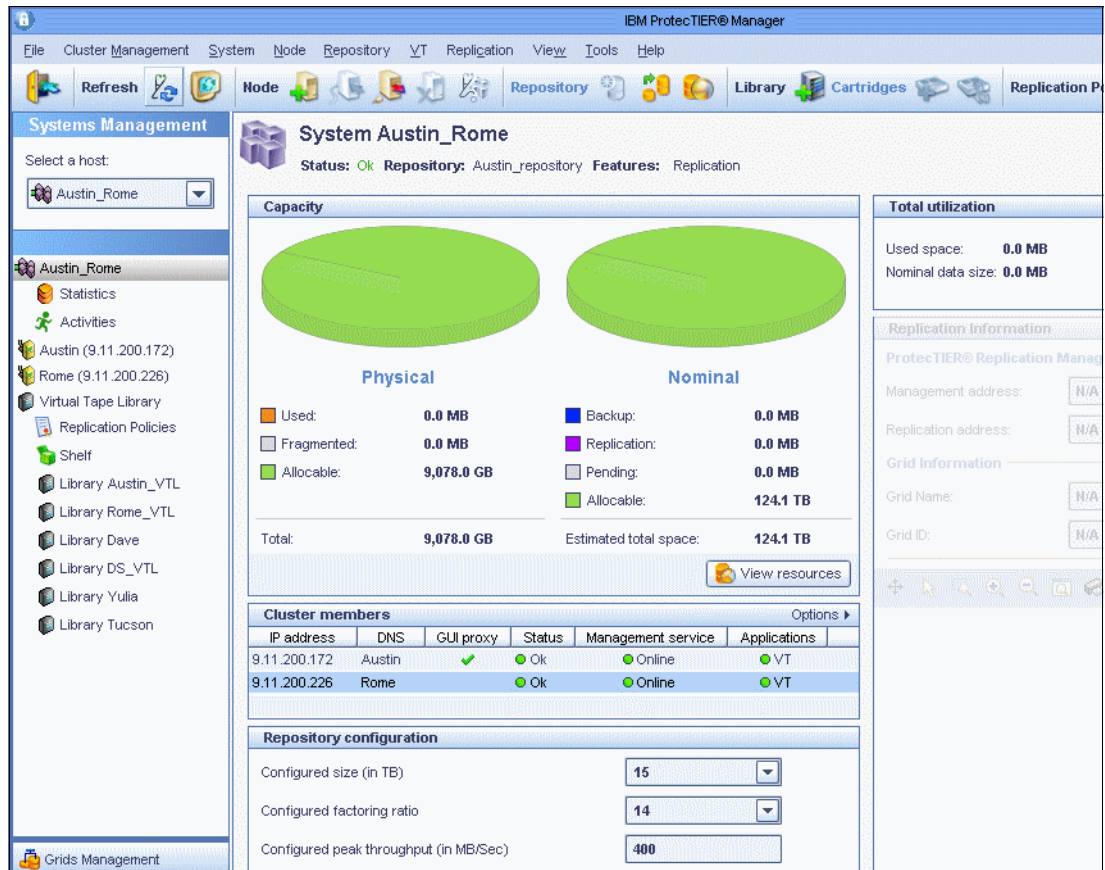


Figure 10-93 Start server refreshed window

## Stopping the server

To stop the server, complete the following steps:

1. From the Nodes pane, select the node on which you want to stop the server.



2. Select **Node** → **Server** → **Stop Server** (Figure 10-94).




Figure 10-94 Stop server

3. Click **Stop server**. A window with login information opens (Figure 10-95). This action happens even you are logged in already. Enter the user name and password.

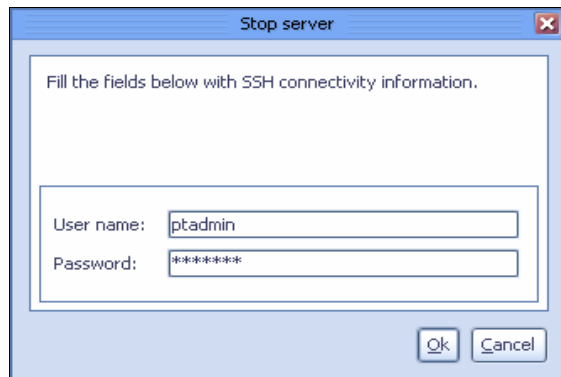


Figure 10-95 Stop server authentication window

After your authentication is verified, a shut down occurs. The server shuts down (Figure 10-96).

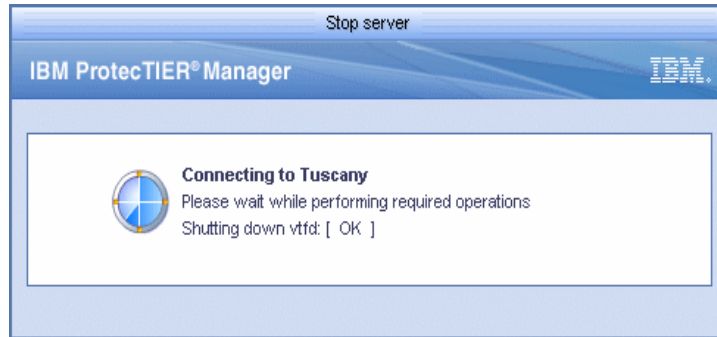


Figure 10-96 Stop server shutting down window

The server is stopped (Figure 10-97).

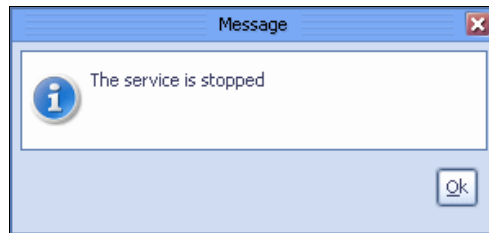


Figure 10-97 Stop server message window

**Note:** After stopping the server, you are logged off of the system.

## 10.8.2 Rebooting a node

To reboot a node, the administrator permission level is needed.

**Note:** Rebooting a node impacts *all* drives and robots assigned to the system to which the node belongs.

Complete the following steps:

1. From the Nodes pane, select the node that you want to reboot.

2. Select **Node** → **Server** → **Reboot** (Figure 10-98).




Figure 10-98 Reboot node

3. Click **Reboot node**. A confirmation window opens (Figure 10-99). Click **Yes** and log in again.

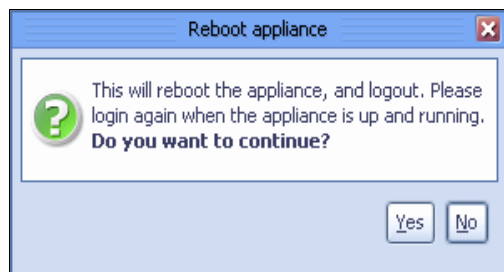


Figure 10-99 Reboot node confirmation window

4. You will be logged out and the node reboots without any further messages (Figure 10-100).

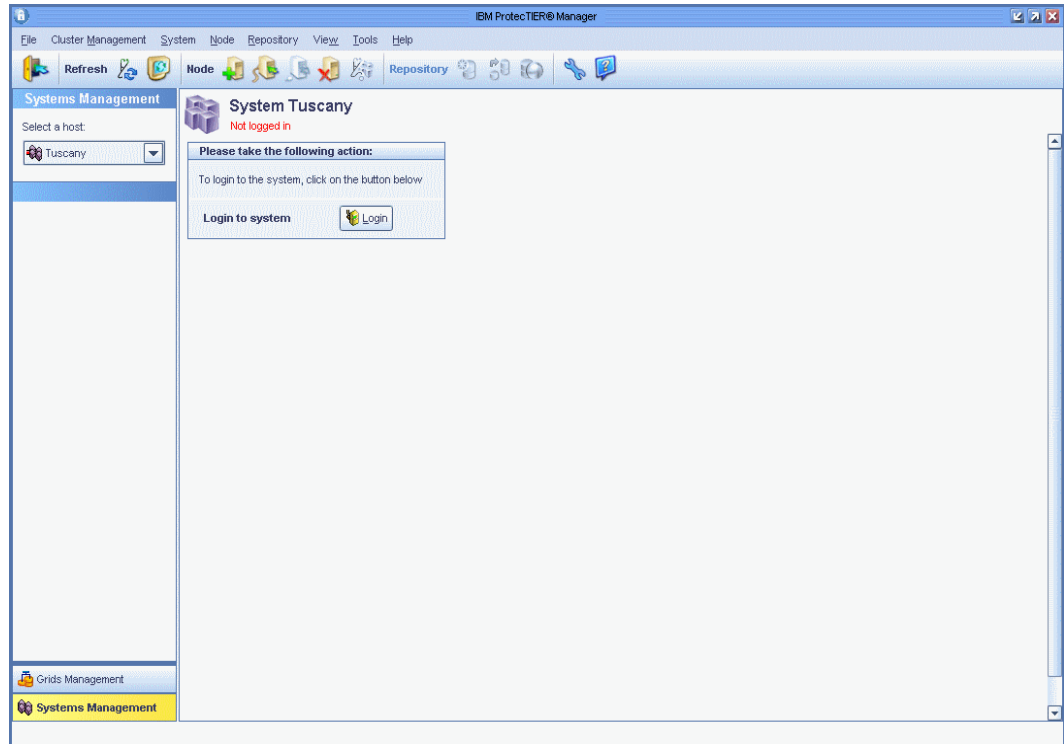


Figure 10-100 Reboot node window while rebooting

**Note:** You may refresh the Navigation pane from time to time to check the latest status of the reboot progress for that node.

### 10.8.3 Disabling defragmentation

The ProtecTIER system automatically defragments fragmented repository disk space as a background task at a rate that does not cause the system to slow down. Stop defragmentation to free the resources used by the defragmentation process by completing the following steps:

1. Select **Repository** → **Defragmentation control**. The Defragmentation control window opens (Figure 10-101).



Figure 10-101 Defragmentation control window

2. Select **Disable defragmentation** and click **OK**. The Defragmentation control window closes and defragmentation is disabled.

**Note:** Selecting **Enable defragmentation** in the Defragmentation control window resumes system defragmentation.

## 10.8.4 Disabling compression

Under normal circumstances, the ProtecTIER system compresses data. Stop compression on a specific virtual library to free the resources usually demanded by the compression process by completing the following steps:

1. Select **VT** → **VT Library** → **Set compression type**. The ProtecTIER compression mode window opens (Figure 10-102).

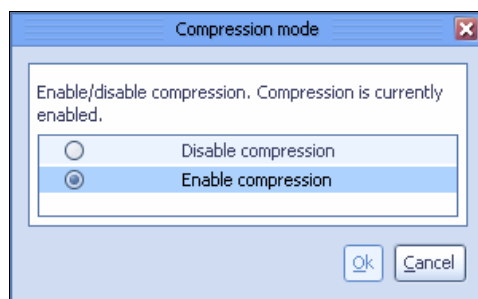


Figure 10-102 Compression mode window

2. Select **Disable compression** and click **OK**. The ProtecTIER compression mode window closes and compression is stopped.

**Note:** Selecting **Enable compression** in the ProtecTIER compression mode window resumes data compression.

## 10.8.5 Changing the HyperFactor mode

By default, ProtecTIER factors all new incoming data, detecting recurring data and storing only the data segments that have not previously been written to the repository. You can change the default HyperFactor mode for each library by completing the following steps:

1. Select **VT** → **VT Library** → **Set HyperFactor mode**. The ProtecTIER VT HyperFactor mode window opens (Figure 10-103).




Figure 10-103 Change HyperFactor mode window

2. Select one of the following options, as directed by IBM Support:
  - **HyperFactor enabled:** HyperFactor operates as normal.
  - **HyperFactor disabled:** HyperFactor stops. When you restart HyperFactor, the HyperFactor process proceeds as normal based on the data stored from before HyperFactor stopped.
  - **Baseline:** HyperFactor stops factoring incoming data and uses the newly stored non-factored data as the reference for factoring new data after HyperFactor is resumed.
3. Click **OK**. The ProtecTIER VT HyperFactor mode window closes.

## 10.8.6 Modifying the trace buffer

The ProtecTIER system stores runtime information in a cyclic memory buffer. IBM Support might direct you to dump the trace buffer for analysis, set the trace recording levels, or reset the trace buffer.

**Note:** You can only manage the trace buffer for one node at a time.

## Dumping the trace buffer

Dumping the trace buffer saves a file of the trace buffer contents to the ProtecTIER system. Complete the following steps:

1. In the Nodes pane, select a node.
2. Select **Node** → **Dump trace**. A confirmation message box opens (Figure 10-104).




Figure 10-104 Confirmation for dump trace

3. Click **Yes**. The trace buffer information is saved to the ProtecTIER system (Figure 10-105).

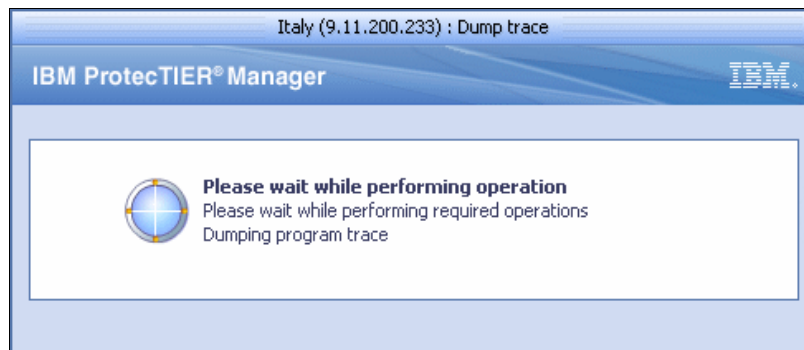


Figure 10-105 Perform dumping program trace

## Setting the trace levels

The ProtecTIER system traces and records many types of operation information at various levels of detail. IBM Support might direct you to reduce the level of detail traced for certain components of free system resources or to increase the level of detail for system components that are suspected to be problematic.

Complete the following steps:

1. In the Nodes pane, select a node.

2. Select **Node** → **Set trace levels**. The Set trace levels window opens (Figure 10-106).

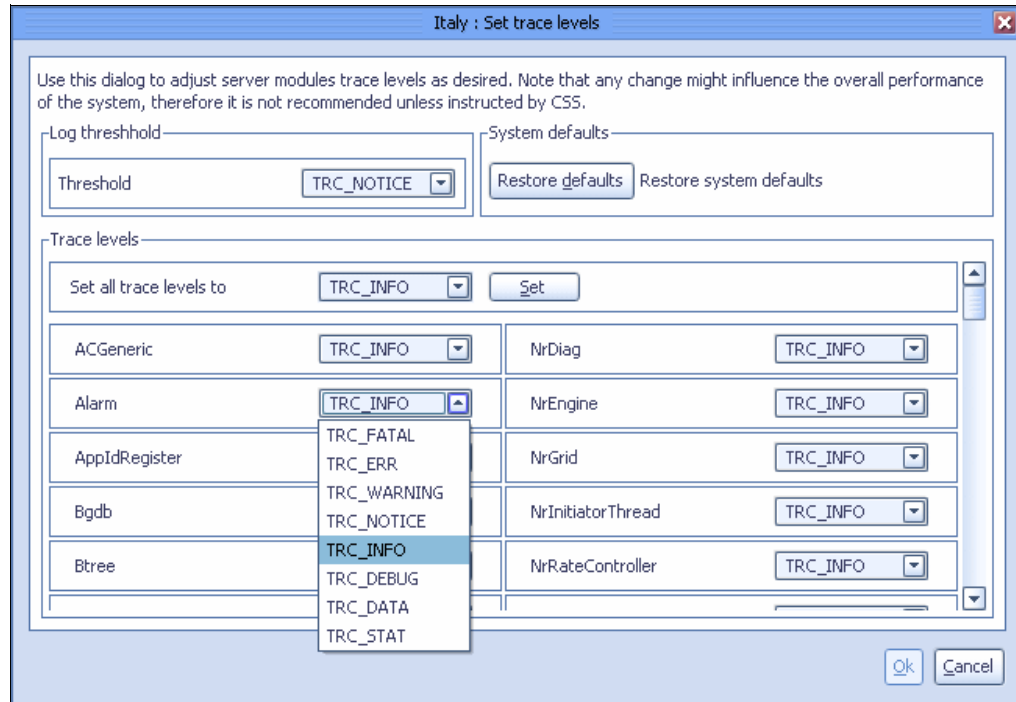


Figure 10-106 Set trace level window

3. Change the trace level settings, as directed by IBM Support.
4. Click **OK**. The Set trace levels window closes and the new trace levels are set.

### Resetting the trace buffer

Resetting the trace buffer empties the buffer. To reset the trace buffer, complete the following steps:

1. In the Nodes pane, select a node.
2. Select **Node** → **Reset trace**. A confirmation message box opens (Figure 10-107).




Figure 10-107 Reset program trace confirmation window

3. Click **Yes**. The trace buffer is reset.



## 10.8.7 Resetting drives

**Attention:** Resetting a tape drive while the backup application is accessing the library can harm the backup operations. Do not reset a device unless directed to do so by IBM Support.

If a virtual drive is locked, reset the device to break any existing SCSI reservations on the device by completing the following steps:

1. In the Services pane, select a library.
2. Click the **Drives** tab.
3. Select a drive (Figure 10-108).




Figure 10-108 Reset drive: Select drive

4. Select **VT** → **VT Drive** → **Reset drive**. A confirmation message window opens (Figure 10-109).

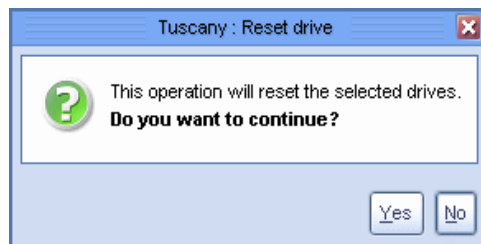


Figure 10-109 Reset drive confirmation window

5. Click **Yes**. The tape drive is reset.

## 10.8.8 Resetting robot

**Attention:** Resetting a robot while the backup application is accessing the library can harm the backup operations. Do not reset a device unless directed to do so by IBM Support.

If a virtual robot locks up, reset the device to break any existing SCSI reservations on the device by completing the following steps:

1. In the Services pane, select a library.
2. Select **VT** → **VT Library** → **Reset robot**. The Reset robot window opens (Figure 10-110).

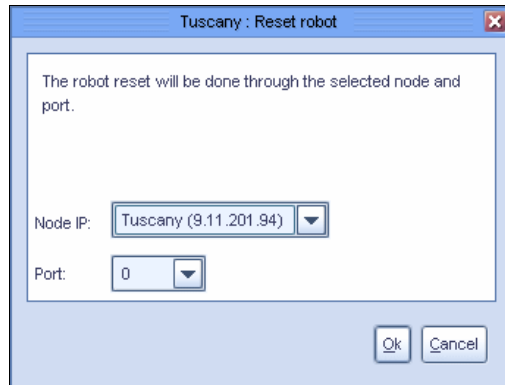


Figure 10-110 Reset robot window

3. Select **Node** and the port to which the robot is assigned.
4. Click **OK**. The tape robot is reset (Figure 10-111).

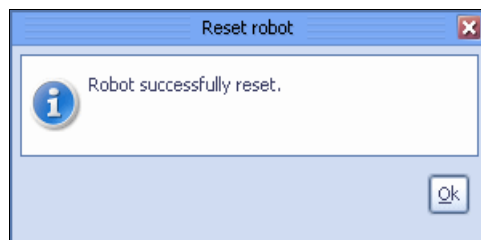


Figure 10-111 Reset robot confirmation window

## 10.8.9 Unloading and moving cartridges

**Attention:** Manual unloading and moving of cartridges is not detected by your backup application and can result in the loss of synchronization between your backup application and ProtecTIER. Furthermore, unloading or moving cartridges while the backup application is using the library can harm the backup operations. Do not unload or relocate a cartridge manually unless directed by IBM Support.

### Unloading cartridges

You must unload a cartridge from its virtual drive to relocate it to a slot. Complete the following steps:

1. In the Services pane, select a library.
2. Click the **Drives** tab.

3. Select a drive that contains a loaded cartridge (Figure 10-112).

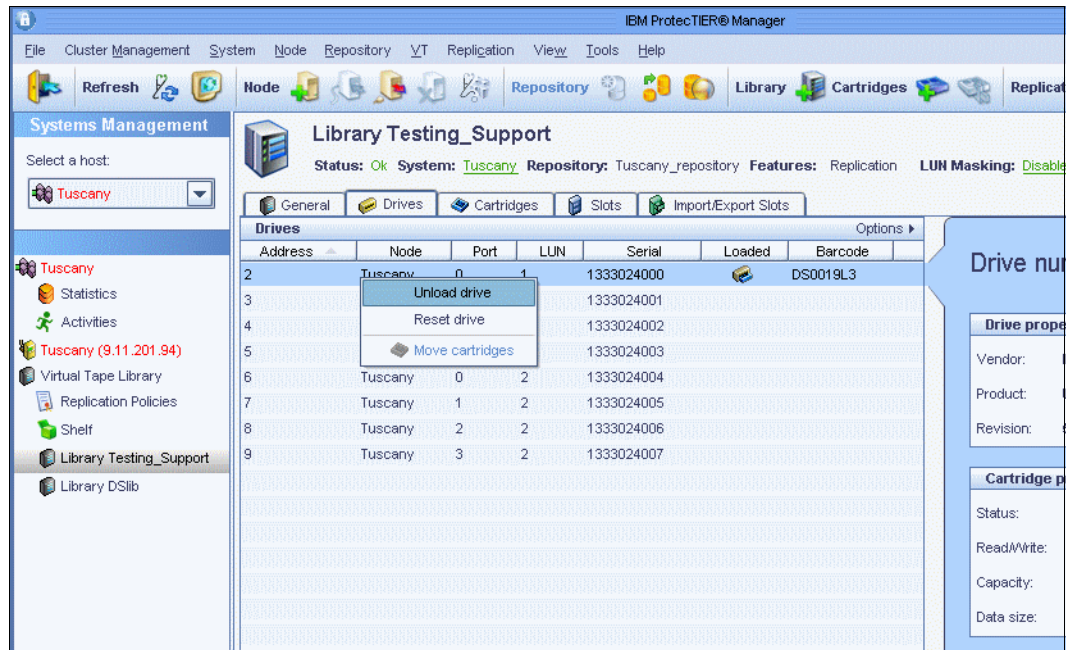


Figure 10-112 Select drive window

**Attention:** The next step causes the cartridge to unload immediately.

4. Select VT → VT Drive → Unload drive (Figure 10-113).

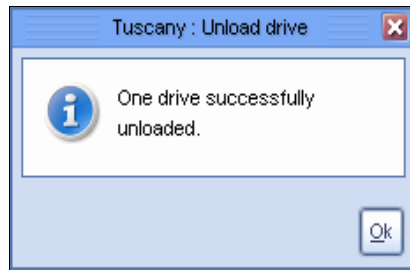


Figure 10-113 Drive unloaded

- The unloaded cartridge is still in the drive. You must relocate it (Figure 10-114).

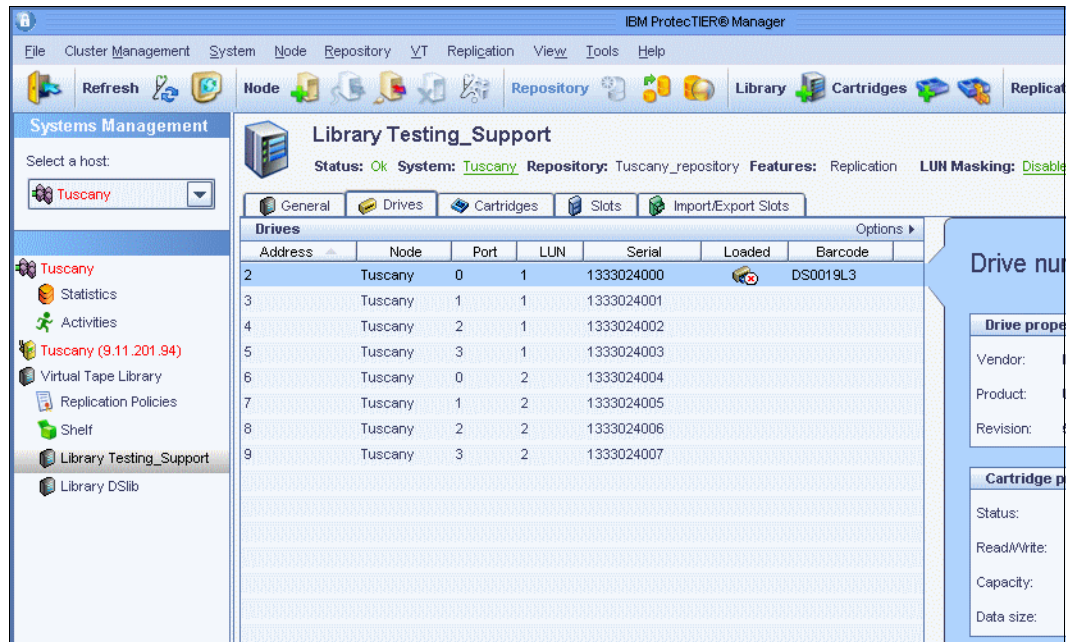


Figure 10-114 Unloaded cartridge still in the drive position

### Relocating cartridges inside the virtual library

You can change the locations of cartridges inside the virtual tape library. The source and target can be drives, slots, and islets. You can also change the logical location of the cartridges to the shelf (refer to 10.9, “Exchanging tape cartridges using the shelf” on page 565 for more information about this topic).

To relocate the cartridges inside the virtual library (for example, relocate a cartridge from a drive position to a slot), complete the following steps:

- In the Services pane, select a library.
- Click the **Drives** tab.

3. Select a drive that contains an unloaded cartridge (Figure 10-115).




Figure 10-115 Select a drive that contains an unloaded cartridge

**Note:** When moving virtual cartridges from a slot to another slot, select a slot that does not contain a cartridge.

4. Select **VT** → **VT Cartridge** → **Move cartridges**. The Move cartridge window opens (Figure 10-116).

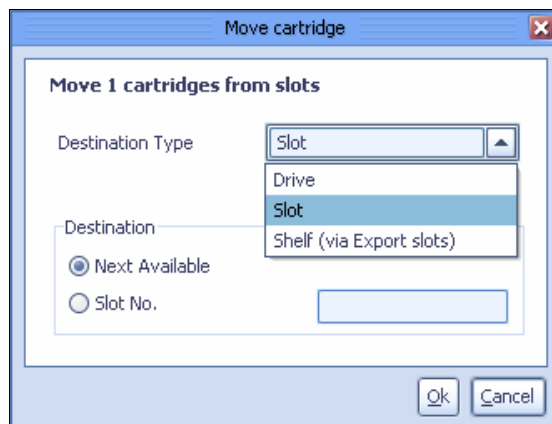


Figure 10-116 Move cartridge destination window

5. In the Destination Type field, select one of the following destinations:
  - **Drive**
  - **Slot**
  - **Shelf (via Export slots)**

6. In the Destination area, select one of the following options:
 

<b>Next Available</b>	The cartridge is placed in the next available location of the selected type.
<b>Slot/Drive No.</b>	The cartridge is placed in the slot or drive with the number specified in the field. The name of this field depends on your selection in the Destination Type field.
7. Click **OK**. The cartridge is relocated to the specified location (Figure 10-117).




Figure 10-117 Move cartridges window

### 10.8.10 Modifying port attributes

If one of your front-end Fibre Channel links is down, it might be because the attributes of that port do not fit the setup of the backup environment. Modify the port attributes from ProtecTIER Manager by completing these steps:

1. In the Nodes pane, select a node.
2. Select **Nodes** → **Port attributes**. The Port attributes wizard Welcome window opens (Figure 10-118).

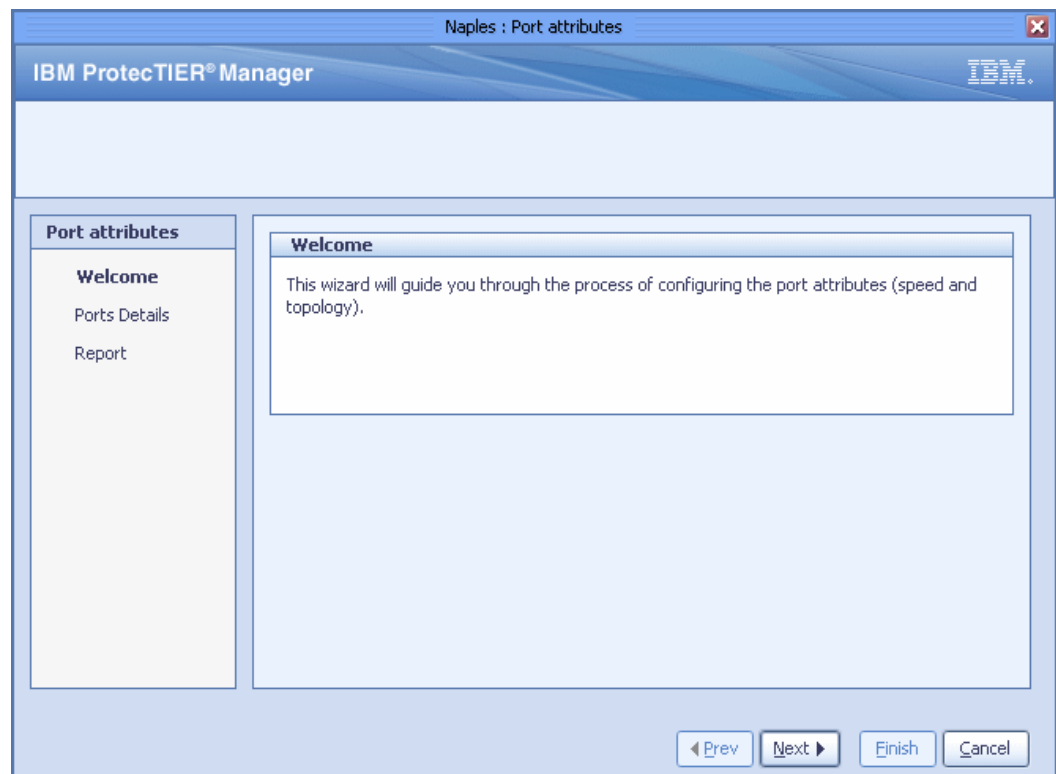


Figure 10-118 Port attributes: Welcome window

Click **Next**. The Port Details window opens (Figure 10-119).

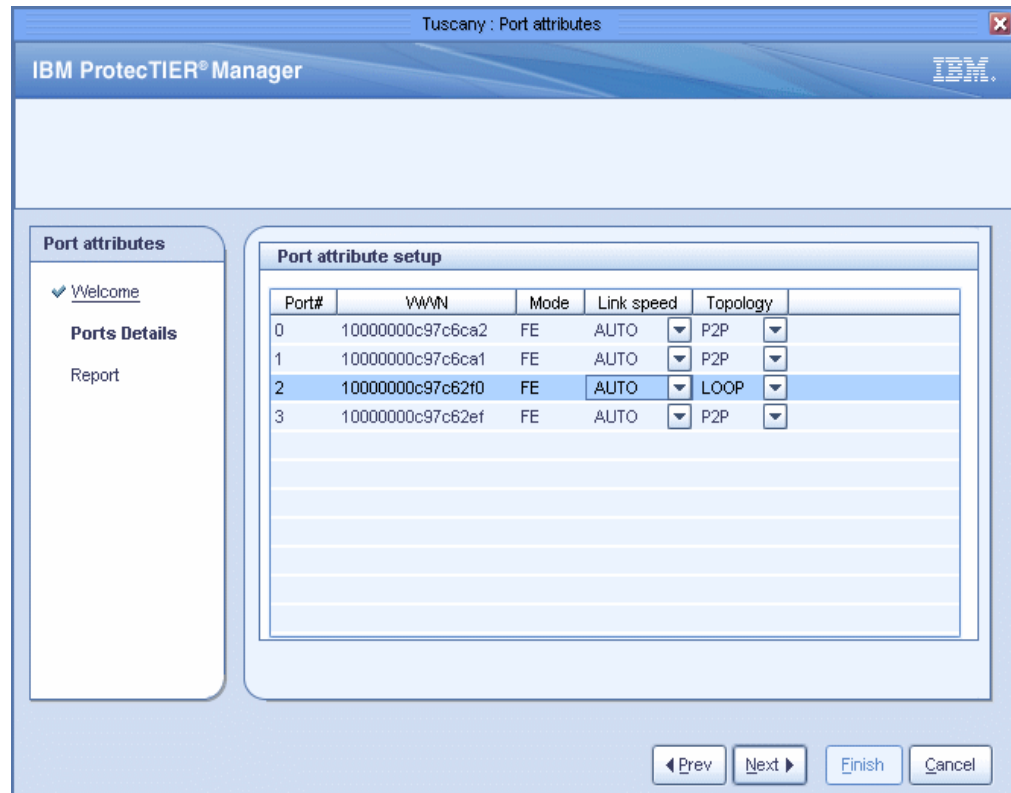


Figure 10-119 Port attribute setup window

- In the Link speed column, click the down arrow to select a link speed from the drop-down list. The options are:
  - **Auto**
  - **1 GB**
  - **2 GB**
  - **4 GB**
  - **8 GB**
- In the Topology column, click the down-arrow to select a network topology from the drop-down list. The options are:

<b>LOOP</b>	Fibre Channel-arbitrated loop
<b>P2P</b>	Point-to-point

Click **Next**. The Port attributes Report window opens (Figure 10-120).

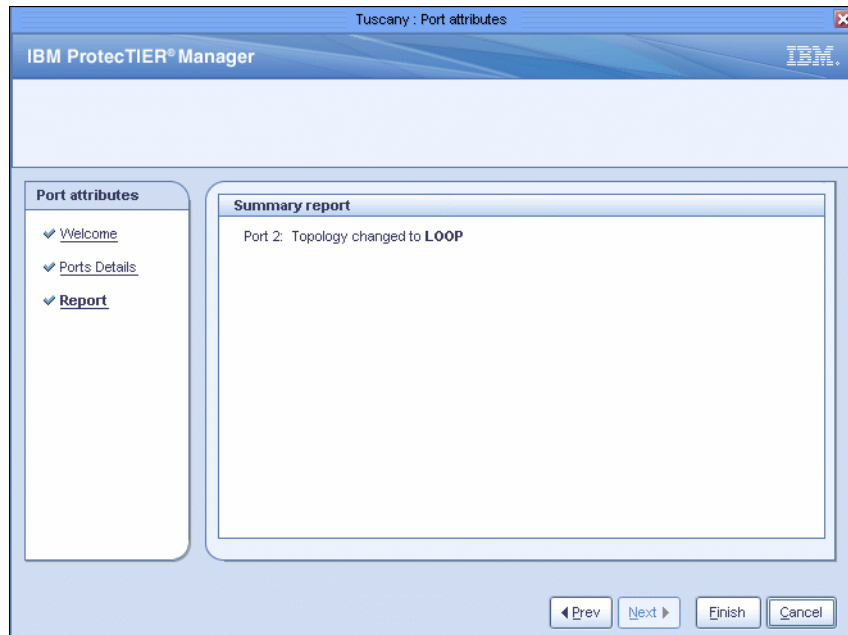


Figure 10-120 Port attributes: Summary report window

5. Click **Finish**. The new port attributes are set (Figure 10-121).

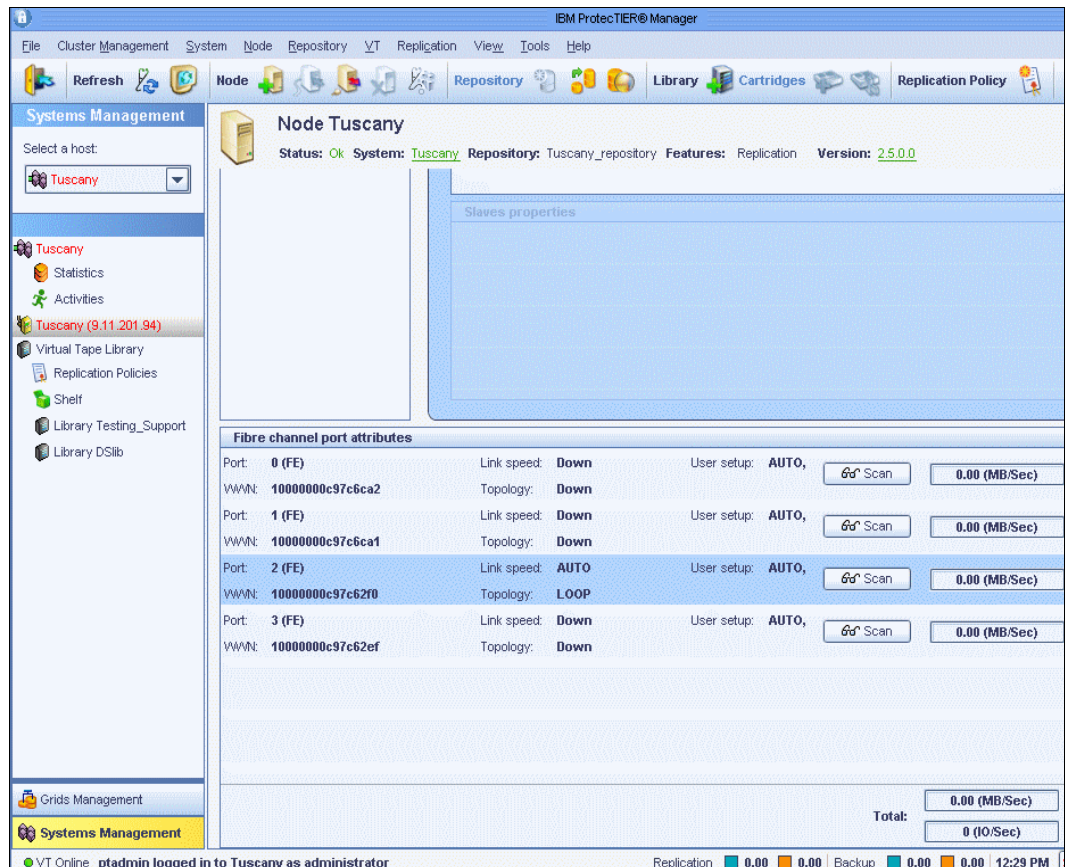


Figure 10-121 Port attribute window with new values



## 10.8.11 Checking the system

After any problematic file systems have been repaired, scan the ProtecTIER system for errors and attempt to repair them using the Check and recover wizard.

**Note:** The check and recovery process is time consuming and the ProtecTIER system goes offline for the duration of the process. Consider using the Check and recover wizard only for severe problems.

### Checking the system

To check the system, complete the following steps:

1. Select **System** → **Check and recover**. A confirmation message box opens (Figure 10-122).

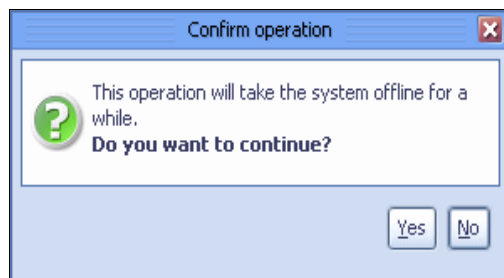


Figure 10-122 Confirmation window

2. Click **Yes**. The ProtecTIER system goes offline and scans itself (Figure 10-123).

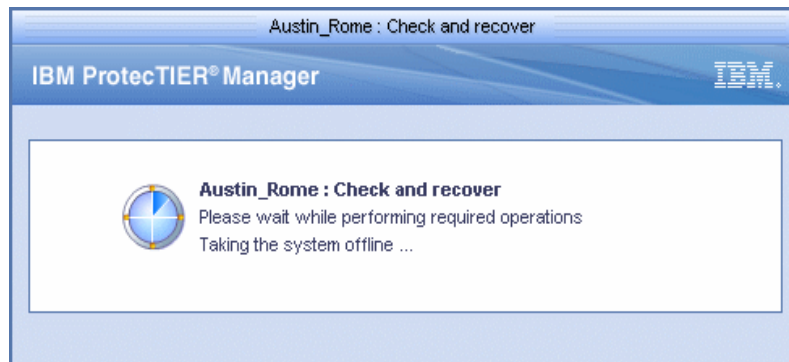


Figure 10-123 Check and recover window

3. The Check and recover window opens with the results of the scan (Figure 10-124).

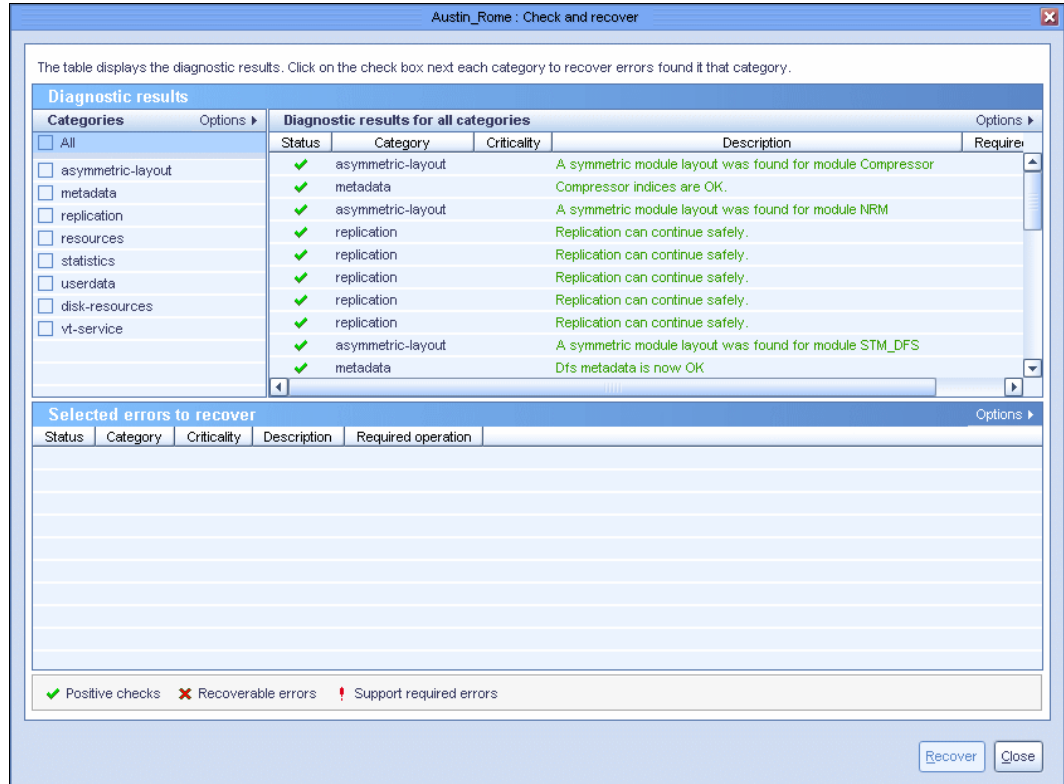


Figure 10-124 Check and recover result window

The Diagnostic results pane displays each element of the ProtecTIER system in one of the following lists:

- Positive checks** No errors
- ProtecTIER recoverable errors** Has errors that the ProtecTIER system might be able to repair
- Support required errors** Has errors that cannot be repaired by the ProtecTIER system without the assistance of IBM Support

4. In the Category sub-pane, filter the listed contents by selecting individual categories. Categories that contain errors of the type ProtecTIER recoverable errors have an empty check box.

5. You should save or print the report and click **Close** to close the window. The system then is brought online (Figure 10-125).

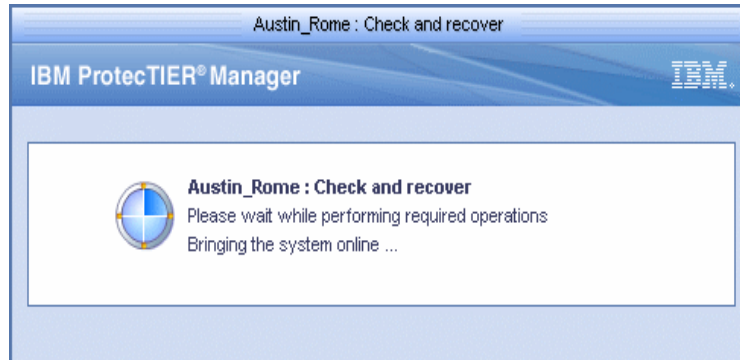


Figure 10-125 Check and recover window

### Repairing the system

If the checking process discovers errors of the type ProtecTIER recoverable errors, attempt to repair the errors using ProtecTIER Manager.

**Attention:** Make sure to allot sufficient downtime for system recovery. The recovery process can be lengthy.

Complete the following steps:

1. In the Categories sub-pane, select the check box for each category for which you want to repair errors. All list items of the type ProtecTIER recoverable errors in each selected category appear in the Selected errors to recover pane and are labeled with a red X.
2. Click **Recover**. ProtecTIER Manager attempts to repair the errors. Errors that are successfully repaired are labeled with a green check mark. Errors that cannot be repaired remain labeled with a red X.

**Note:** For assistance with errors that cannot be repaired, contact IBM Support.

### 10.8.12 Cartridge metadata verification

The cartridge metadata verification function verifies the metadata only. It reads the files with cartridge metadata information and verifies checksums and other parameters, such as barcode and cart\_id. It also returns how many filemarks are on the cartridge.

To perform a cartridge metadata verification operation, you must be logged in with administrator permissions and complete the following steps:

1. From the System Management pane, select the library that has the cartridge inserted on which the metadata should be verified (Figure 10-126).

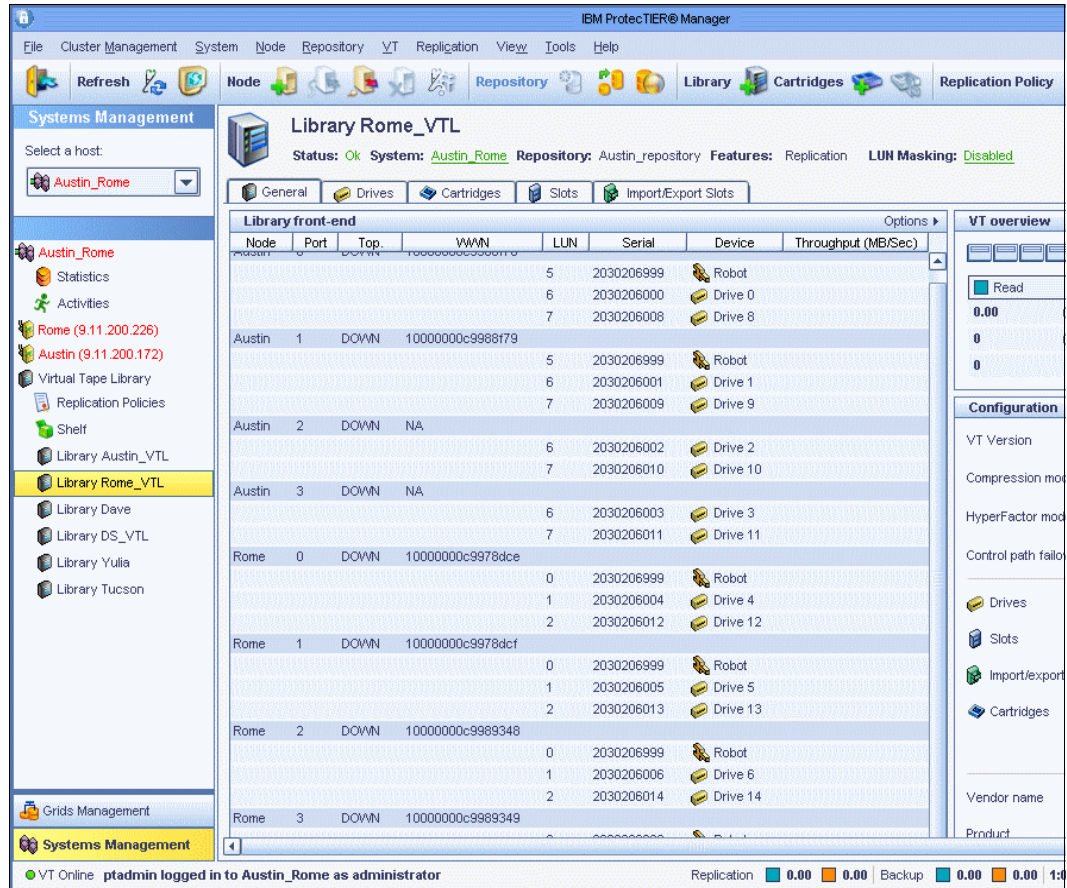


Figure 10-126 Check cartridge metadata verification select library window

2. Select the **Slots** tab. The library configuration window opens (Figure 10-127).

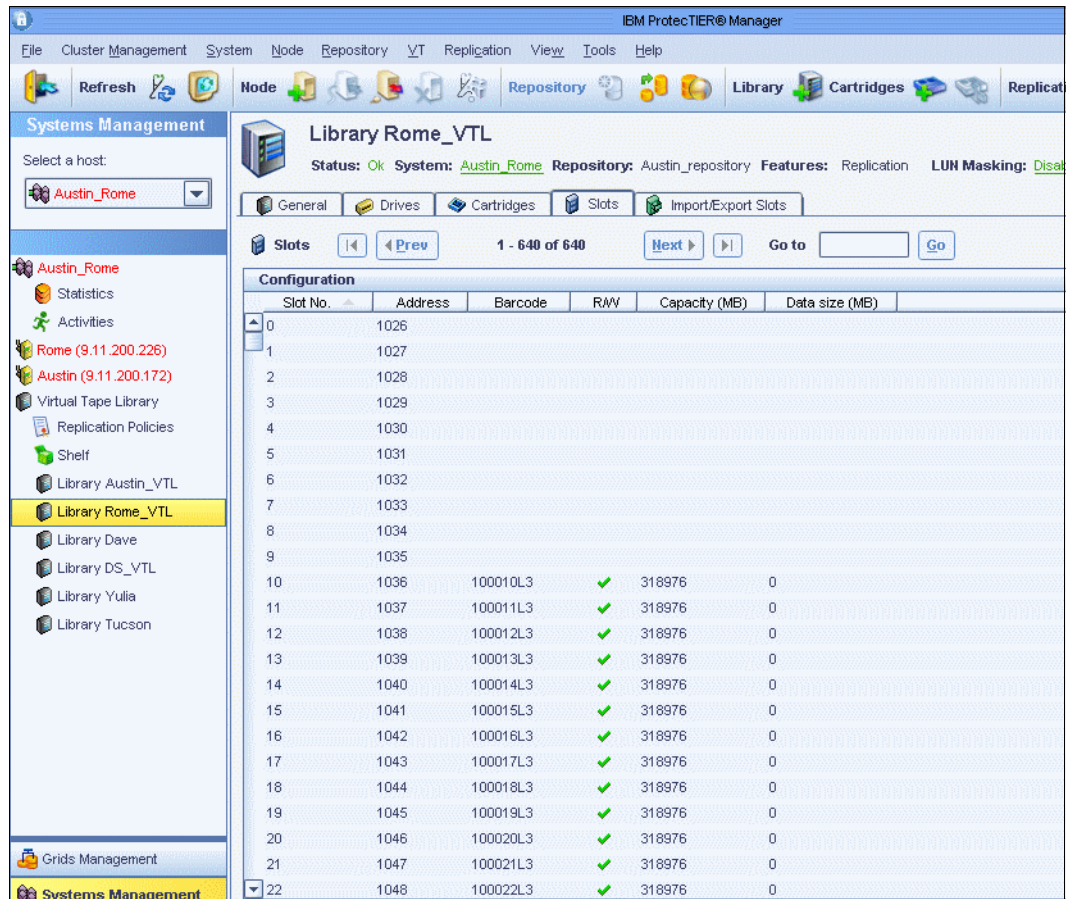


Figure 10-127 Library configuration window

- Select the cartridges that you want to verify the metadata and select **VT → VT cartridges → cartridges metadata verification** (Figure 10-128).

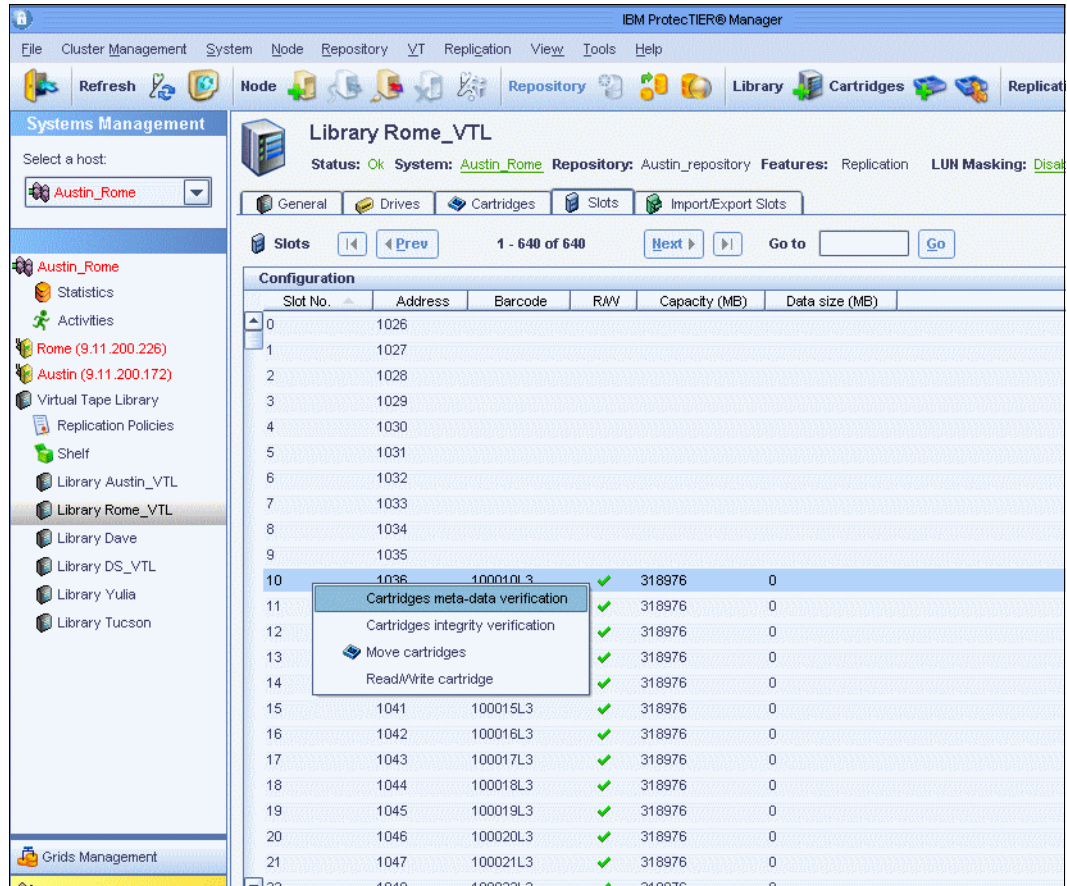


Figure 10-128 Metadata verification window

A confirmation window opens (Figure 10-129).

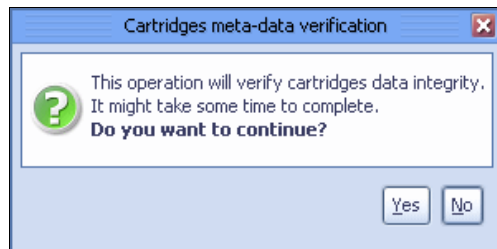


Figure 10-129 Cartridges metadata verification confirmation window



To perform a cartridge integrity operation, you must be logged in with administrator permissions and complete the following steps:

1. From the System management pane, select the library that has the cartridge inserted that should be verified (Figure 10-131).

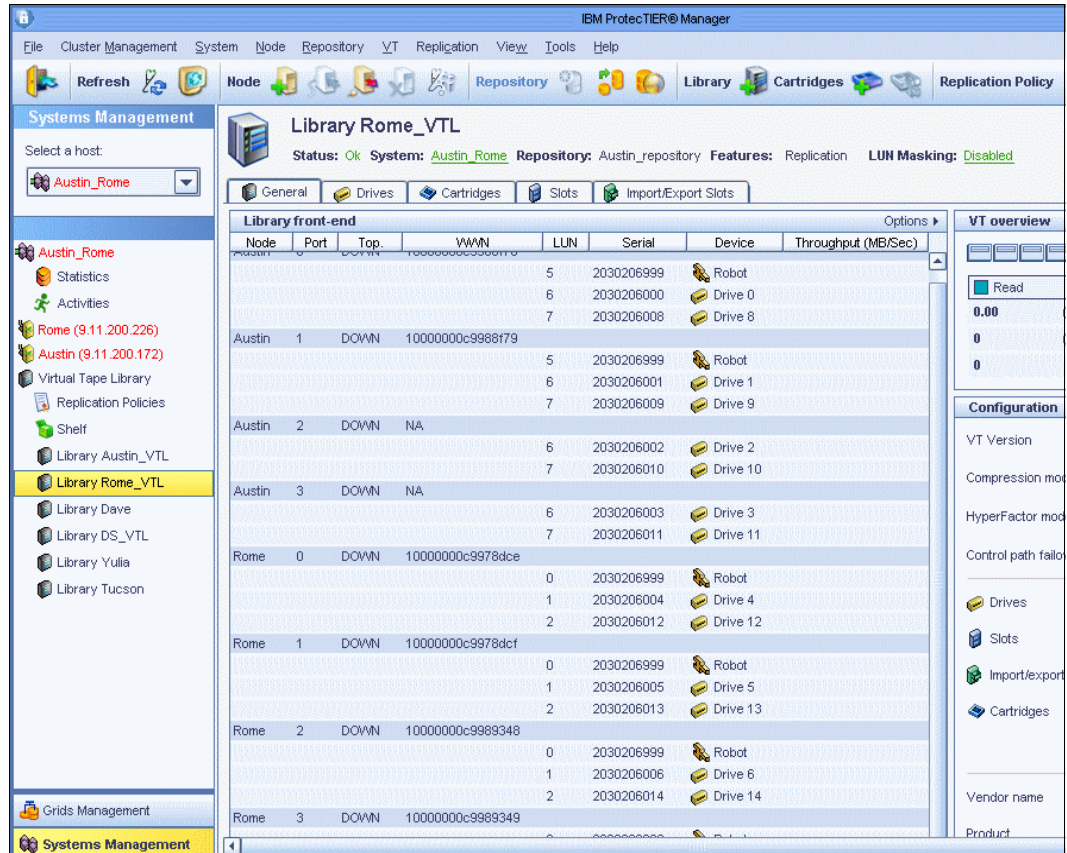


Figure 10-131 Check cartridge integrity select library window



2. Select the **Slots** tab. The library configuration window opens (Figure 10-132).

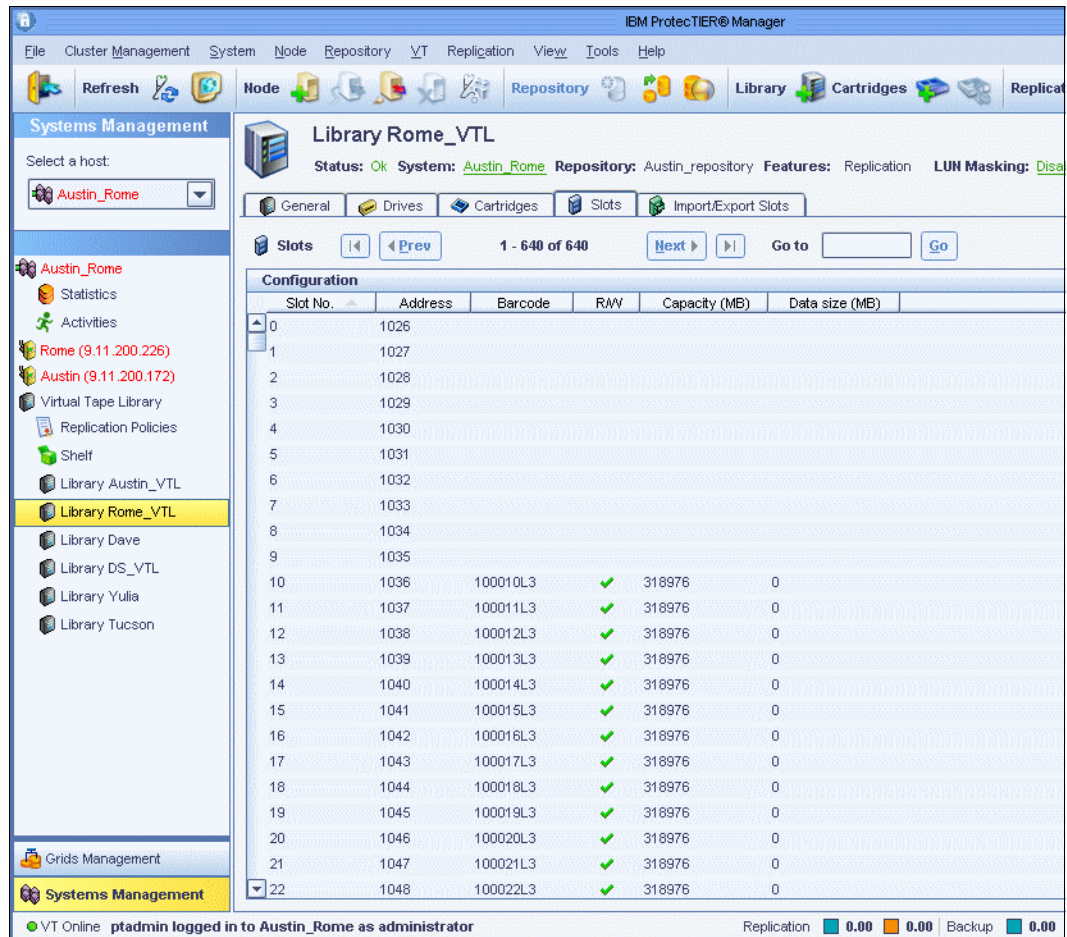


Figure 10-132 Library configuration window

- Select the cartridges that you want to verify and select **VT → VT cartridge → cartridges integrity verification** (Figure 10-133).

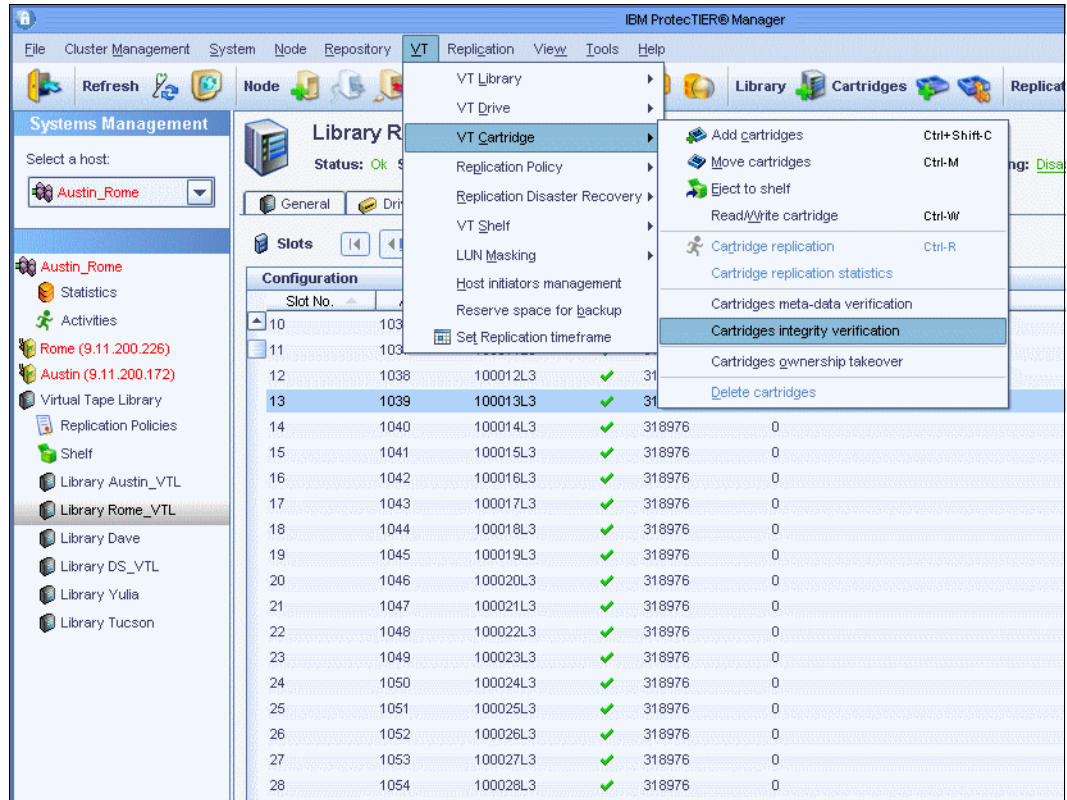


Figure 10-133 integrity verification window

- A confirmation window opens (Figure 10-134).



Figure 10-134 Cartridges integrity verification confirmation window

5. Click **Yes**. The cartridge integrity verification starts (Figure 10-135).

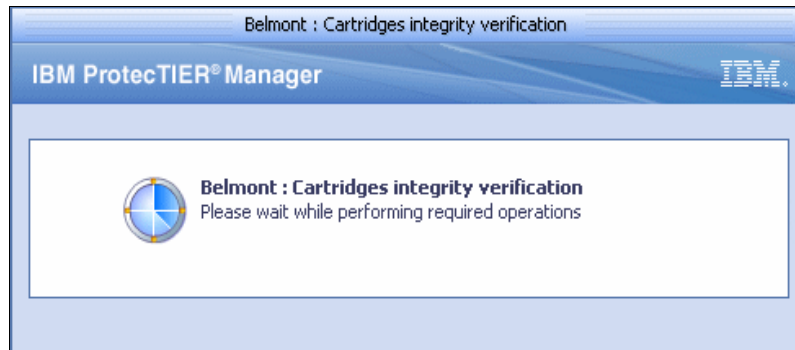


Figure 10-135 Performing cartridge integrity verification window

After the verification, a table with the results of the verification for the selected cartridges opens (Figure 10-136).

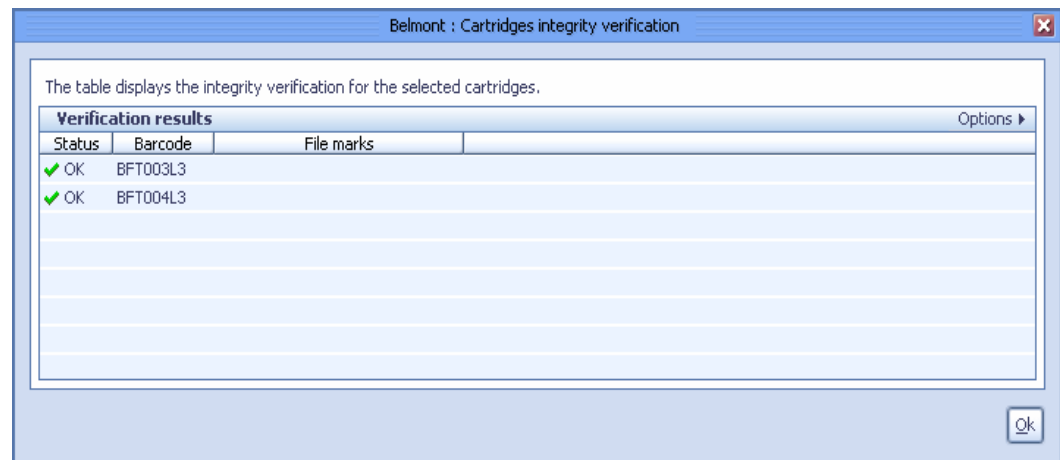


Figure 10-136 Integrity verification results window

6. Click **OK** to end the operation.

## 10.9 Exchanging tape cartridges using the shelf

The shelf is a virtual public area of a ProtecTIER system that can hold the virtual tape cartridges. The tape cartridges on the shelf do *not* belong to any virtual tape libraries. You can exchange tape cartridges between virtual tape libraries by using the shelf. In addition, you can use shelf replication. Refer to Chapter 8, “IBM System Storage ProtecTIER native replication operation” on page 409 for more information.

In this section, we demonstrate how to exchange tape cartridges between virtual tape libraries of the system by using the shelf (Figure 10-137). ProtecTIER system Italy is a two-node clustered system that consists of Italy and Naples. We are going to switch the visibility of one tape cartridge with barcode 000005L3 in the Italy system from the virtual tape library test lab to virtual tape library Italy Account.




Figure 10-137 Exchange tape cartridges in the ProtecTIER system

Complete the following steps:

1. From the Services pane, select the library test lab.
2. Select the **Slots** tab.

3. Select **000005L3** cartridge and select **VT → VT cartridge → Move cartridges**, or you can select the cartridge and right-click and select **Move cartridges** (Figure 10-138).

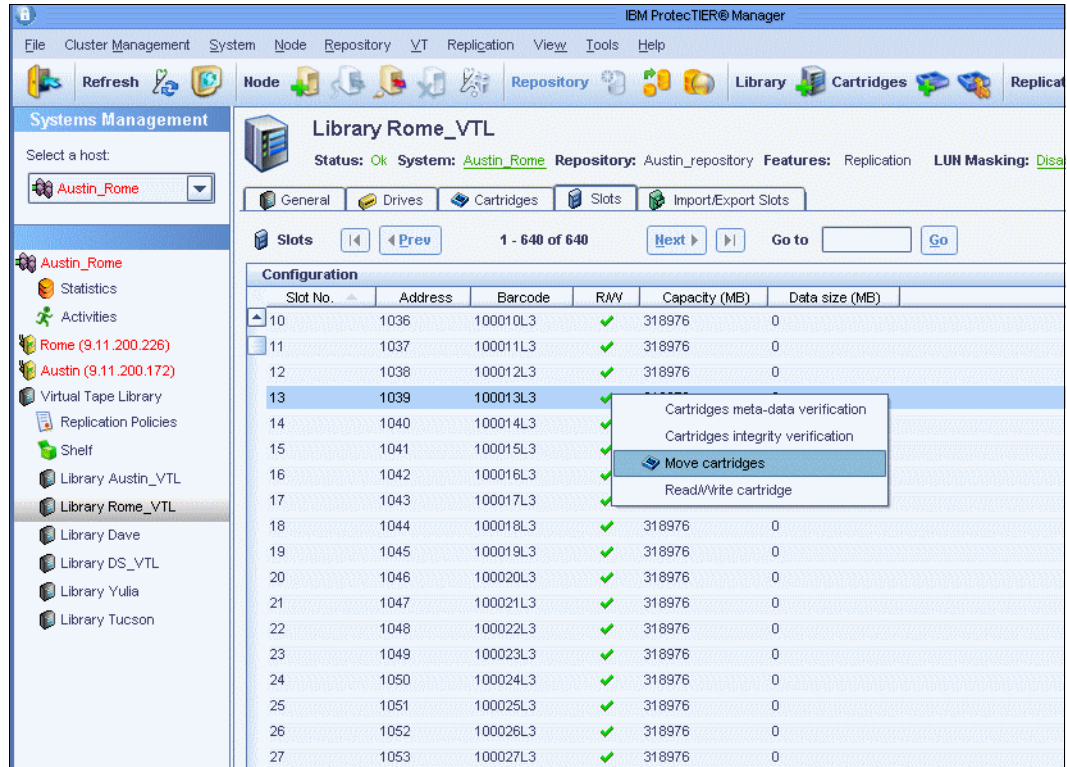


Figure 10-138 Select the cartridge that will be moved to shelf

4. The Move window opens. Select **Shelf (via Export slots)** for Destination Type (Figure 10-139).

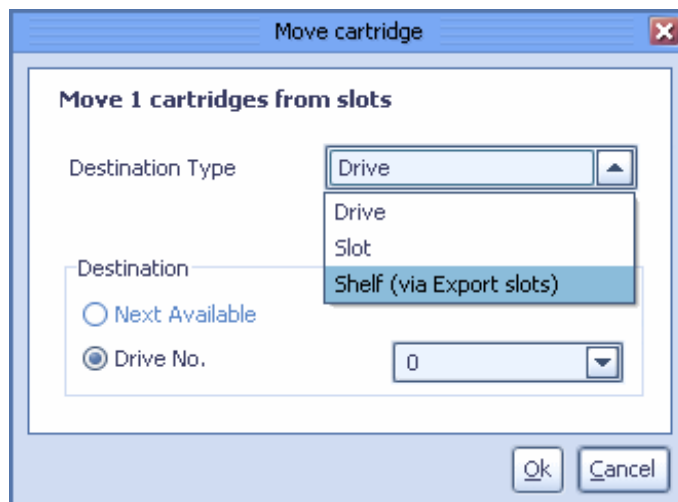


Figure 10-139 Select the destination type

5. Select **OK**. The cartridge is moved.

- From the Services pane, select **Shelf**. You can find the 100013L3 tape cartridge (Figure 10-140).




Figure 10-140 The cartridge has been moved into the shelf

7. Select the **00005L3** cartridge and select **VT → VT cartridge → Move cartridges**, or you can select the cartridge and right-click and select **Move cartridges** (Figure 10-141).

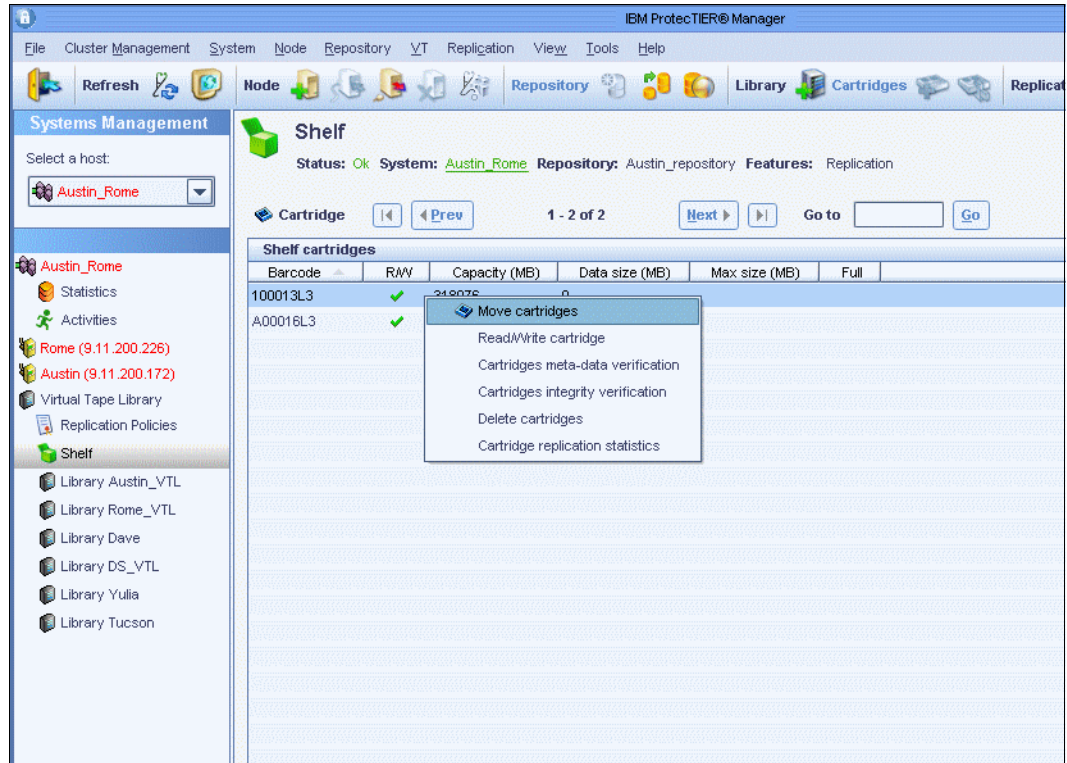


Figure 10-141 Select the cartridge to be moved

8. The Move cartridges window opens. Select **Italy Account** for the destination library (Figure 10-142).

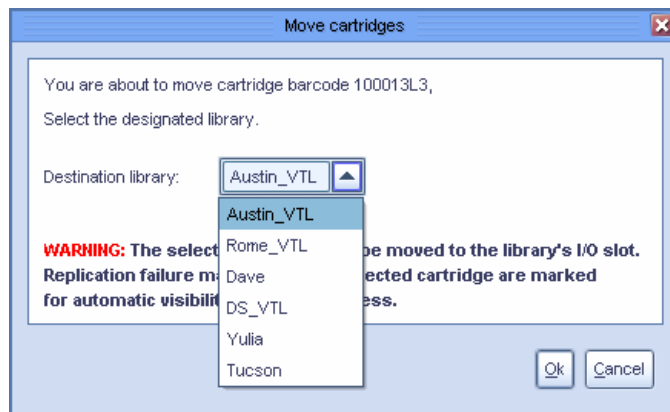


Figure 10-142 Select the destination library

- Click **OK**. The visibility of the tape is switched to the destination library Italy Account. You can select the library **Italy Account** from the Services pane, then select the **Cartridge** tab to check the visibility (Figure 10-143).

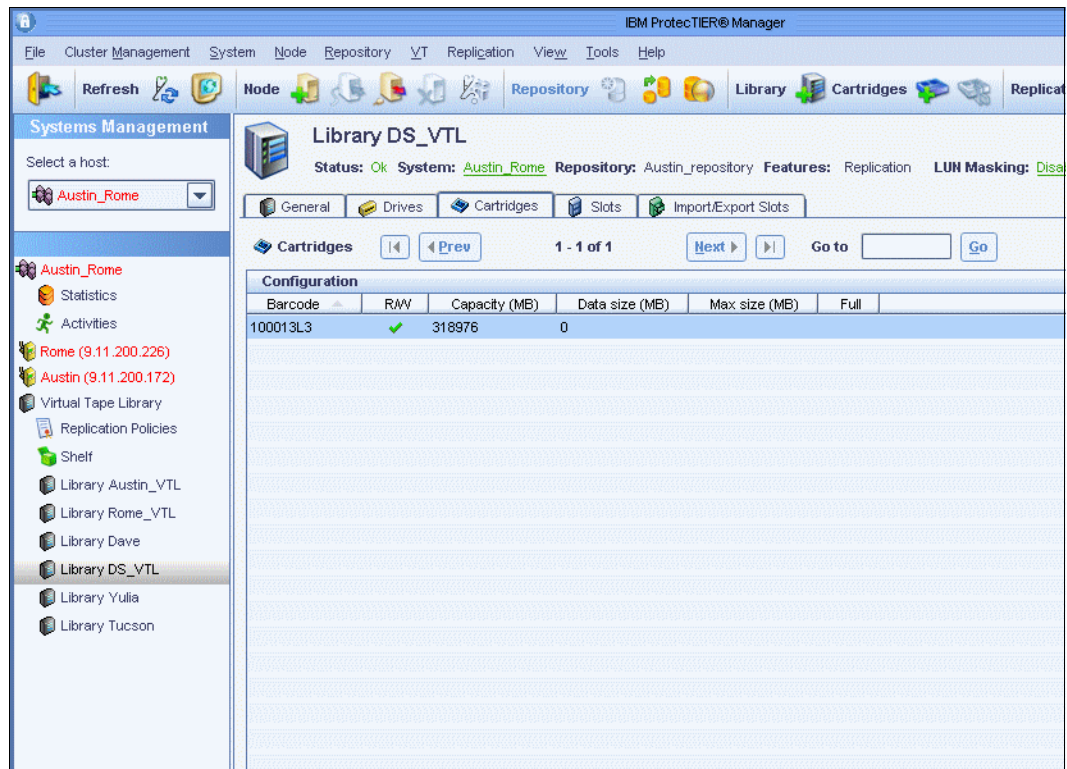


Figure 10-143 Visibility was moved to library Italy Account

## 10.10 Replication bandwidth throttling

Bandwidth throttling is a way to control the speed at which replication activities operate, whereby the user can specify a maximum upper limit for the physical network usage. The ProtecTIER hardware platform supports up to 200 MBps of physical TCP/IP network traffic through two Gigabit Ethernet interfaces per node (400 for a dual-node cluster). By default, there is no configured bandwidth limit. ProtecTIER attempts to use as much bandwidth as it can.

If the physical network layer consists of dark fibre or another high-speed network infrastructure, there is typically no reason to limit replication throughput. However, if ProtecTIER is running over a smaller network pipe that is shared by several applications simultaneously, the user may choose to restrict the maximum throughput used by ProtecTIER replication. This parameter is configurable per GigE port on all nodes in the replication grid, but it only applies to outgoing data, so it is only effective to 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, if you want to hold ProtecTIER replication to no more than 100 MBps and you are using all four available GigE ports of the source dual-node cluster, you must set each port's limit 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 port limit value is zero, indicating no limit. The command used to set limits is:

```
/opt/dtc/install/ptconfig -bandwidthLimit
```



Figure 10-144 shows that the Eth3 interface limit is modified to 25 MBps.

```
[root@Austin ~]# /opt/dtc/install/ptconfig -bandwidthLimit
Gathering information [ Done ]

Available interfaces:
=====
1. eth3
2. eth4
Please choose an interface (q to quit): 1

Please specify the traffic limit in MegaBytes (0 to disable traffic shaping) [0]: 25
Changing allowed Bandwidth for replication [ Done ]
Appending configuration changes [ Done ]

Available interfaces:
=====
1. eth3
2. eth4
Please choose an interface (q to quit): q
[root@Austin ~]#
```

Figure 10-144 Bandwidth throttling

## 10.11 Automation of daily operations

This section outlines the operational concepts of ProtecTIER with native replication in a typical user environment and introduces the available options to automate it. The focus of this discussion is ways to automate the *movement* of virtual cartridges from the primary site to the secondary (DR) site to clone it to physical tape for archiving and longer term retention.

In any user environment, the backup application manages the inject (import) of cartridges into a library. Cartridges that were moved from the virtual shelf into a VTL import/export slot *reside* there until moved again by the backup application into a specific library. In a typical ProtecTIER replication environment that is designed to support a basic DR strategy with no cloning to physical tape, cartridges are replicated by the system from the primary site to the remote site virtual shelf and stay there unseen by the backup application until they are manually moved by the user, under a DR scenario or DR test scenarios, from the shelf to the import/export slot of the remote site library. At that time, the remote site backup application server can take over the repository and enable the user to restore data or continue the backup operation at the secondary site.

Alternatively, if the user operation calls for creating physical tape copies at the secondary site (outside of the main production environment), the user can use ProtecTIER's visibility control switch when creating the replication policies. In this scenario, once replicated, cartridges are automatically moved from the remote site virtual shelf into a designated specific library's import/export slots, and the remote site backup application server are e made aware of it. This process can be automated from within the backup application environment, especially when it enables a single domain environment (shared catalog across sites). When using this approach, there are some conceptual topics to remember:

- ▶ The inject command is backup application dependent and is identical to the way physical cartridges are being injected to the library from the import/export slots.
- ▶ The number of cartridges in a virtual library can be significantly higher than in a physical library. If a large number of cartridges are requested to be imported to a library, consider scripting the inject command to ensure that cartridges are inserted to the library while freeing the import/export slots for other cartridges.

As mentioned, the automation process of *moving* cartridges between sites and performing clone-to-physical-tape operation at the secondary site is more suitable in single-domain backup application environments. Some of the major backup applications, such as NetBackup, Legato, and BRMS, allow for this type of environment.

The following sections provides examples of possible automation opportunity within a NetBackup backup application environment.

## Examples for the eject and inject commands for Symantec NetBackup

In this section, we provide examples of using eject and inject commands for Symantec NetBackup.

When using a vault profile, complete the following steps:

1. Create a vault profile for ejecting cartridges from a library.
2. Eject cartridges from the library using the vault profile by running the following command:

```
vltrun <vault profile name>
```

3. Inject a cartridge at a remote site to the library by running the following command:

```
vltinject <vault profile name>
```

When using eject/inject commands, complete the following steps:

1. Eject cartridges from the library by running the following command:

```
vmchange -res -multi_eject .... -ml <barcodeA:barcodeB:...:barcodeZ>
```

2. Inject cartridges into a library by running the following command:

```
vmchange -res -multi_inject .... -ml <barcodeA:barcodeB:...:barcodeZ>
```

You can also run an inventory command to inject cartridges as follows:

```
vmupdate -rt <robot type> -rn <robot #> -empty_map
```

This command scans the Imp/Exp slots and injects all available cartridges to the library.

## Scripting the inject commands

Vault, inject, and eject commands can be scripted to run periodically on the backup application host, which triggers automatic cartridge movement from import/export to the library whenever the relevant cartridge is located in the import/export slot, preventing running out of free import/export slots.

You should not script the inventory command, as it scans the robot and therefore might take a long time to complete on libraries with large numbers of cartridges.

Examples of vault scripts are:

```
▶ #!/bin/csh  
▶ while (1)  
▶ vltinject myVault  
▶ sleep 600  
▶ end
```

## 10.12 Updating the ProtecTIER Manager

To update the ProtecTIER Manager, you must uninstall the old version first and then install the new version. You are not required to add the nodes again to your new ProtecTIER Manager.





## Native replication and disaster recovery

Disaster recovery (DR) is the process of recovering production site data in a remote location that was the target for the replication operation from the primary site prior to the disaster occurrence.

In case of a disaster or a situation where the production (or primary) site has gone offline, the hub, or disaster recovery site, can take the place of the production site until the primary site comes back online. When the primary site comes back online, previously replicated and newly created tapes can be moved to the main production site using the failback process so that it can once again become the production/primary site for the user.

After the user enters the ProtecTIER Disaster Recovery (DR) mode, all incoming replication and visibility switching activities from the failed production site to the DR site (hub) are blocked. When that primary site is rebuilt or replaced, the user can then return the data to that primary site and continue the backup and replication operations.

If the user has a situation where they must stop using the primary repository, such as the repository is down or has been destroyed, and they want to use the DR repository as their backup target, the user must enter the DR mode at the DR repository to start working with the DR site as their primary site.

This chapter summarizes the different stages in handling a disaster scenario.

## 11.1 Moving to ProtecTIER DR mode

The main consideration is to keep the moving to DR mode simple for both the user and the operation to avoid potential complications and human errors when handling an emergency situation.

To move to ProtecTIER DR mode, complete the following steps:

1. Initiate the DR mode on the hub repository through the ProtecTIER Manager (Figure 11-1), which will block all replication activities (including visibility switch) at the DR site (and at the primary site, if it is still up in the case of a DR test). The repository must be defined as a hub in a grid environment.

In our example node, ChinChic is the primary (source or spoke) site and the rosswell node is the remote (target or hub) site. Entering the DR mode was done on the rosswell node. This is the site where the replica cartridges are located that we want to access.

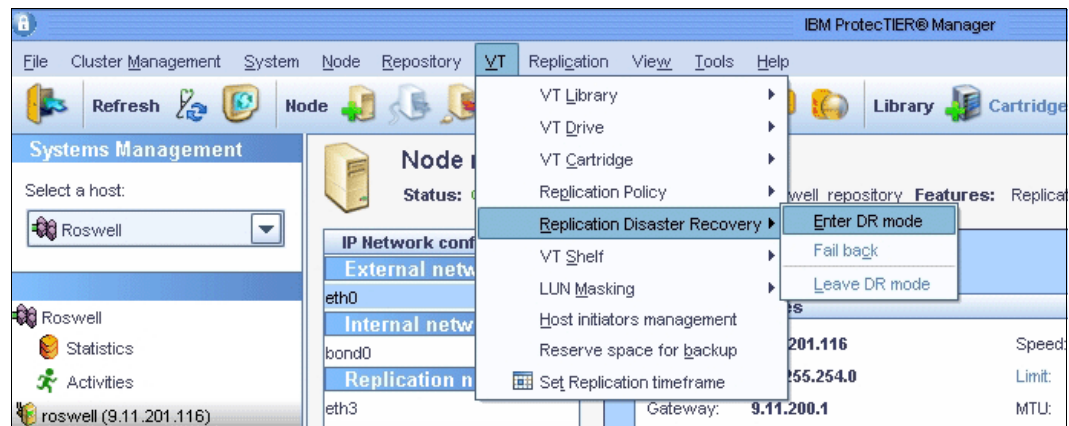


Figure 11-1 Moving to ProtecTIER DR mode

2. Select **VT** → **Replication Disaster Recovery** → **Enter DR mode**. The source repository selection window is displayed (Figure 11-2).

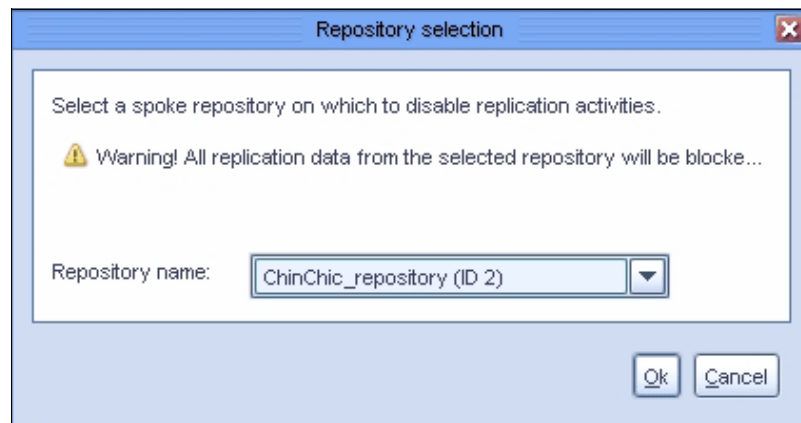


Figure 11-2 DR Repository selection

3. Select the spoke repository and click **OK**. A second confirmation window opens (Figure 11-3).

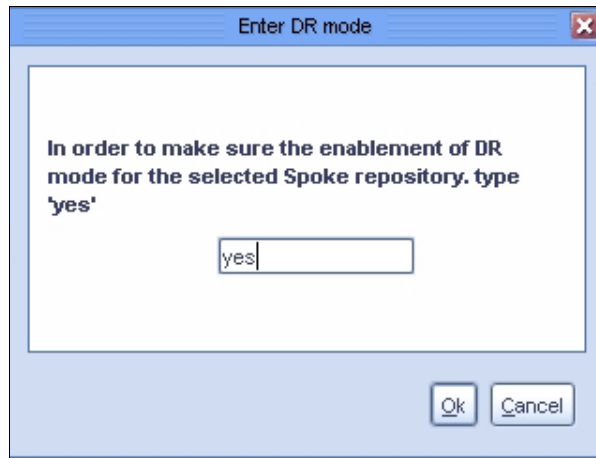


Figure 11-3 Second DR confirmation window

Enter yes and click **OK** to activate the DR mode. An automatic procedure is executed that blocks incoming replication to the DR site. In addition, visibility switching is blocked. No message is displayed that the DR mode entered was successful.

**Note:** If replication is not blocked, the safety of the data at the DR site cannot be guaranteed.

The DR mode is shown in the ProtecTIER Manager and ProtecTIER Replication Manager view pane. A dashed arrow for DR replication from hub to spoke is shown instead of the arrow shown for replication from the source to the hub repository (Figure 11-4).

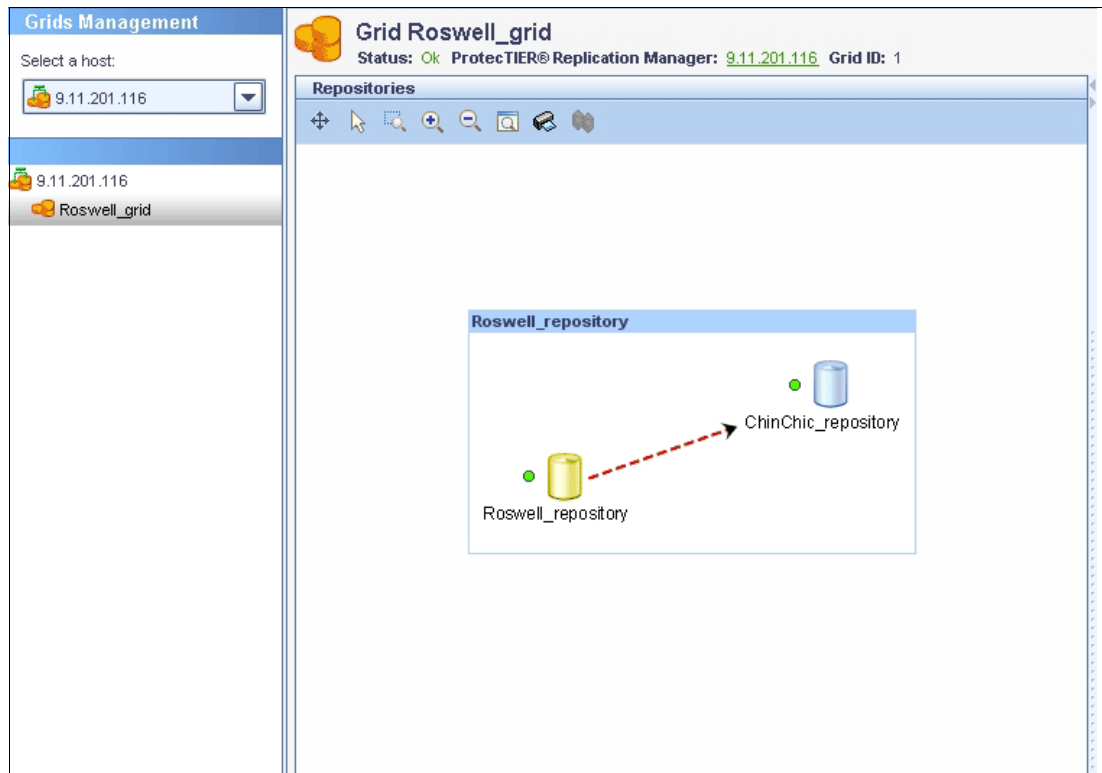


Figure 11-4 DR Mode on roswell

From the spoke point of view, you can see the hub DR mode (Figure 11-5).

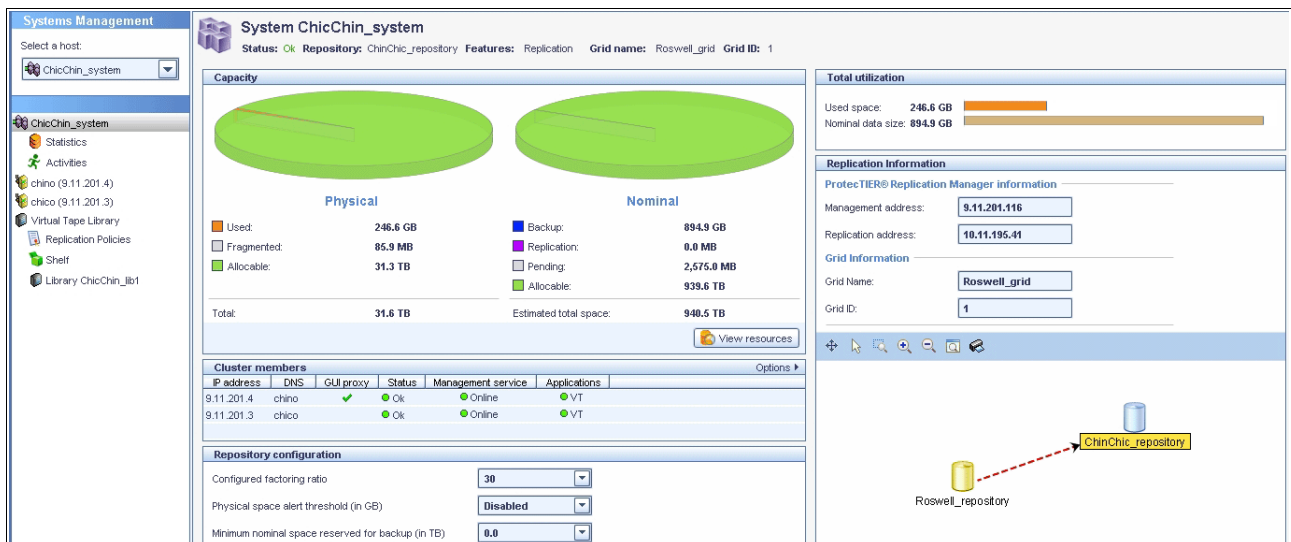


Figure 11-5 ProtecTIER spoke point of view for the hub in DR mode



The replicas of the cartridges that were created at the production site are in read-only mode at the DR site. To work with the DR site as the new backup target, the backup application must have read write (R/W) permission. Refer to 11.6, “Principality” on page 598 for a detailed description. If the user must run local backups at the DR site, they must use new locally created cartridges. For the cartridge status of replicated cartridges on hub repository refer to Figure 11-6. This figure shows the access mode as non-read/write. At the moment, the principality is owned by the source repository and the cartridges can only be accessed in read only mode.

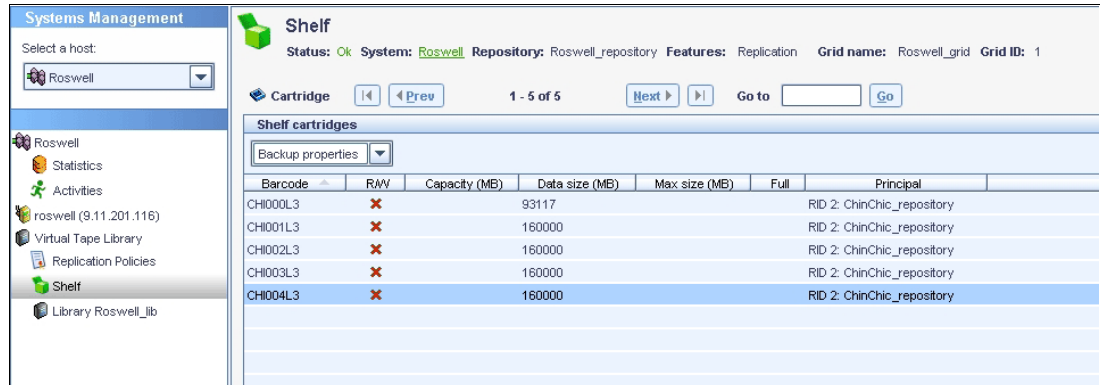


Figure 11-6 ProtecTIER hub replicated cartridges backup properties view

## 11.2 Failback for all data at the DR site

The next step is to create your failback policy. Failback is the procedure for replicating updated cartridges, new or old, back to the (restored) primary site from a DR site to bring it up to date in case production site was unavailable, or to a new production system.

If the primary repository was down and has been restored, you can return to working with the production site as the primary site and use the DR site as the remote, or secondary, site.

The DR site must be able to use a *one-time policy* that holds all cartridges to be replicated to the primary site. This special policy can also transfer the *principality attribute* of the relevant cartridges that were created at the DR site while the primary site was down, to the primary repository or to its substitute in case the original primary is *dead*.

**Note:** The failback procedure can only be initiated under DR mode.

**Note:** The cartridges that originally came from the primary site and have not been changed on DR site do not need to be replicated.

Before creating your failback policy, you must perform an eject operation on the cartridges to change the visibility of them from the library to the shelf before creating a failback policy for these cartridges. If you do not eject the cartridges, they will not be valid for failback.

**Note:** The failback policy, in addition to copying the cartridges, also takes care of transferring the principality, or ownership, of all the cartridges that belonged to the temporary primary repository at the DR site, to the restored primary repository at the production site. Refer to 11.6, “Principality” on page 598 for more information.

## 11.2.1 Splitting the failback according to priorities

Depending on the number of cartridges that have to be restored, it could make sense to perform the restore in several steps, and not with the full amount of cartridges at one time, because the cartridges are moved to the shelf for recovery and are not accessible during that time. If you only move some of them at a time, the rest will remain accessible to the host on the DR side. If you move all your cartridges at one time, the restore could take a long time, depending on the amount of cartridges and data. The cartridges will be accessible again when the restore of all cartridges has finished. Check your requirements and conditions before you create a failback policy. Count the amount of data and cartridges and check if you have a business need to access at least some of your data during the restore, or if you can waive all cartridges for the whole time the failback needs.

**Example of a partitioned restore:** You have 20 cartridges you want to restore. You set up a failback policy for the first five cartridges. These cartridges are moved to the shelf and are not accessible for the host during the failback. The other 15 cartridges are still accessible on DR side during that time. When the first five cartridges are restored, you create a second policy for the next five cartridges, and so on.

## 11.2.2 Creating a failback policy

To create a failback policy, complete the following steps:

1. Select **VT** → **Replication Disaster recovery** → **Fail back**. The Fail back wizard welcome window opens (Figure 11-7).

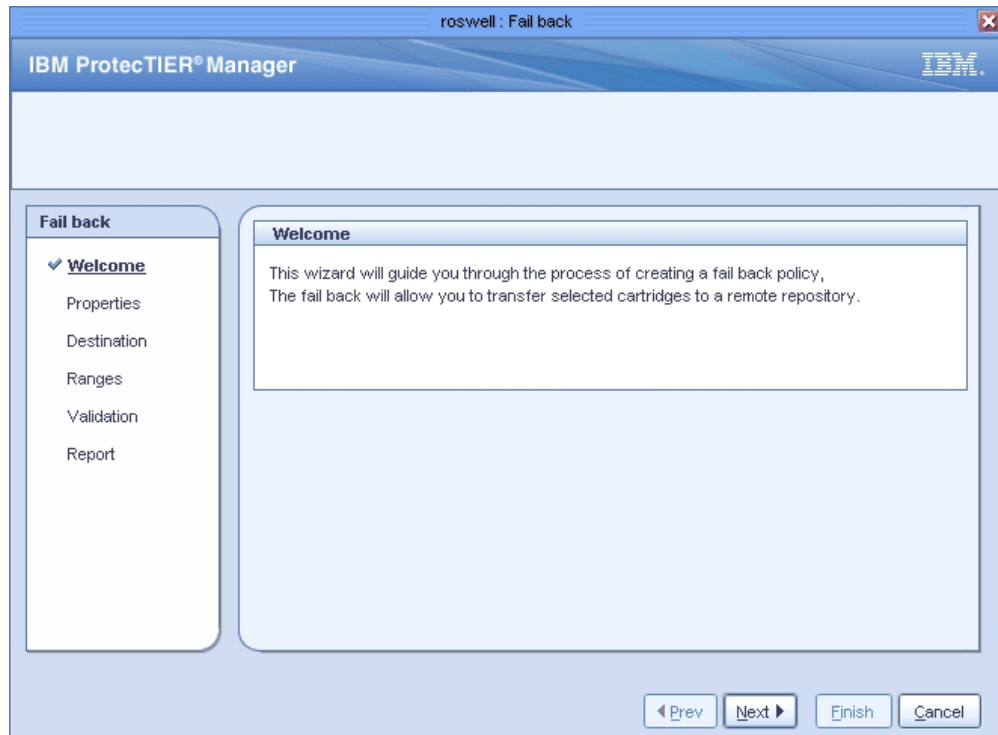


Figure 11-7 Create failback policy welcome window

2. Click **Next**. The Properties window opens. Enter a unique name for the failback policy in the Policy name field (Figure 11-8).

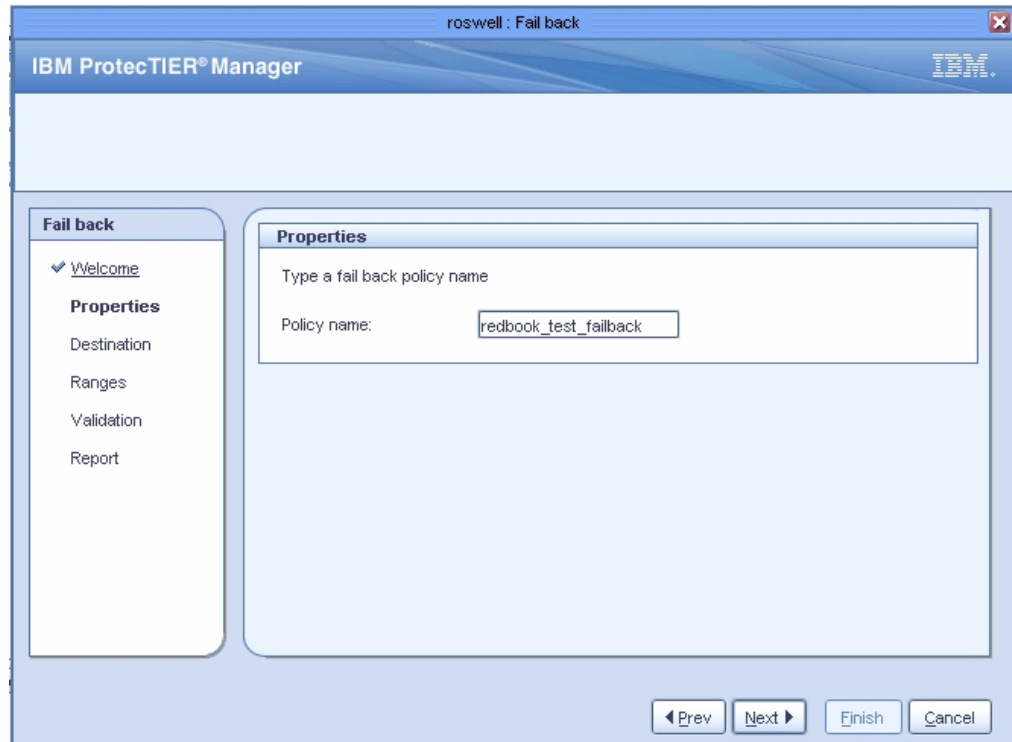


Figure 11-8 Create failback policy properties window

In our example, we entered redbook\_test\_failback.

3. Click **Next**. The replication failback destination window opens. Select the repository you want to select as destination for your failback policy (Figure 11-9).

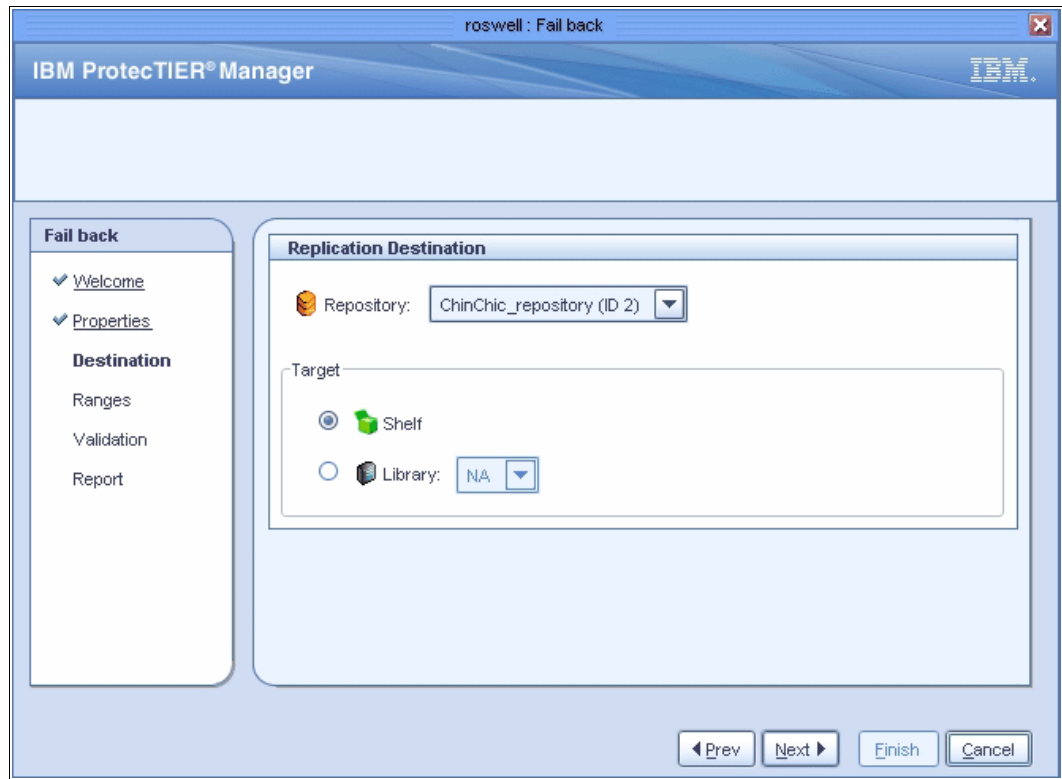


Figure 11-9 Failback destination properties window

4. Click **Next**. The Barcode ranges window opens (Figure 11-10). Enter the barcode range of the cartridges that you want to fail back to the old or replaced primary repository (refer to 11.2.1, “Splitting the failback according to priorities” on page 580). In our example, we use the barcode range CHI000–CHI004.

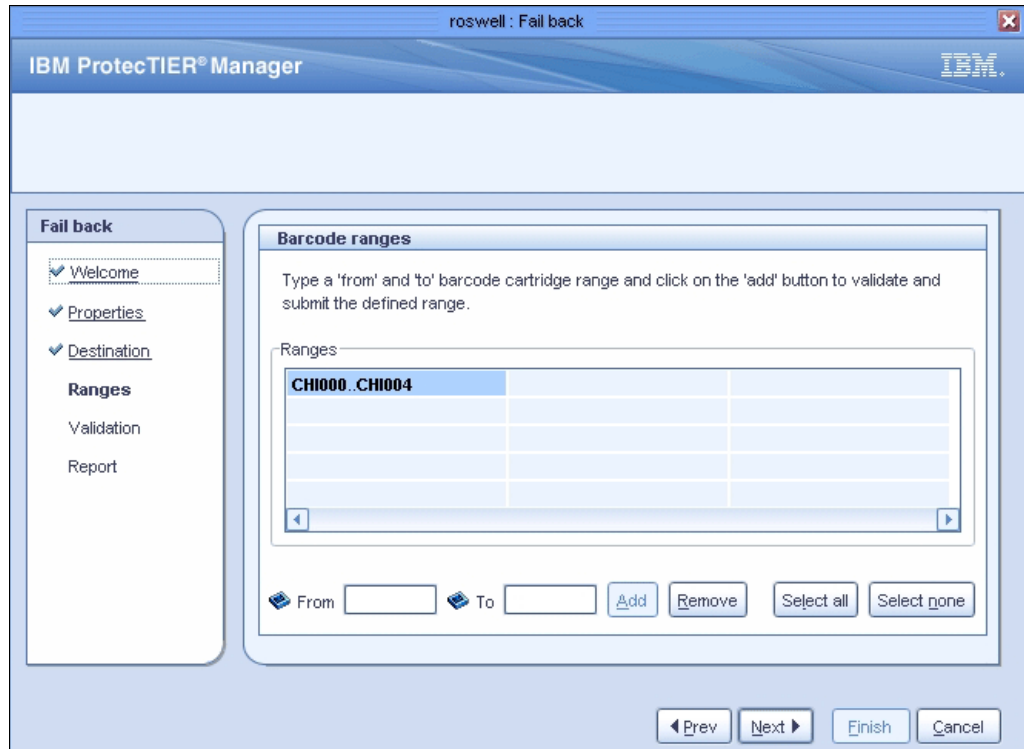


Figure 11-10 Select barcode range for failback policy window

5. Click **Next**. The failback policy validation window opens (Figure 11-11). Here you can see if all of your selected cartridges are valid for failback. If not, select **Previous** to check and change your fail back cartridge range.

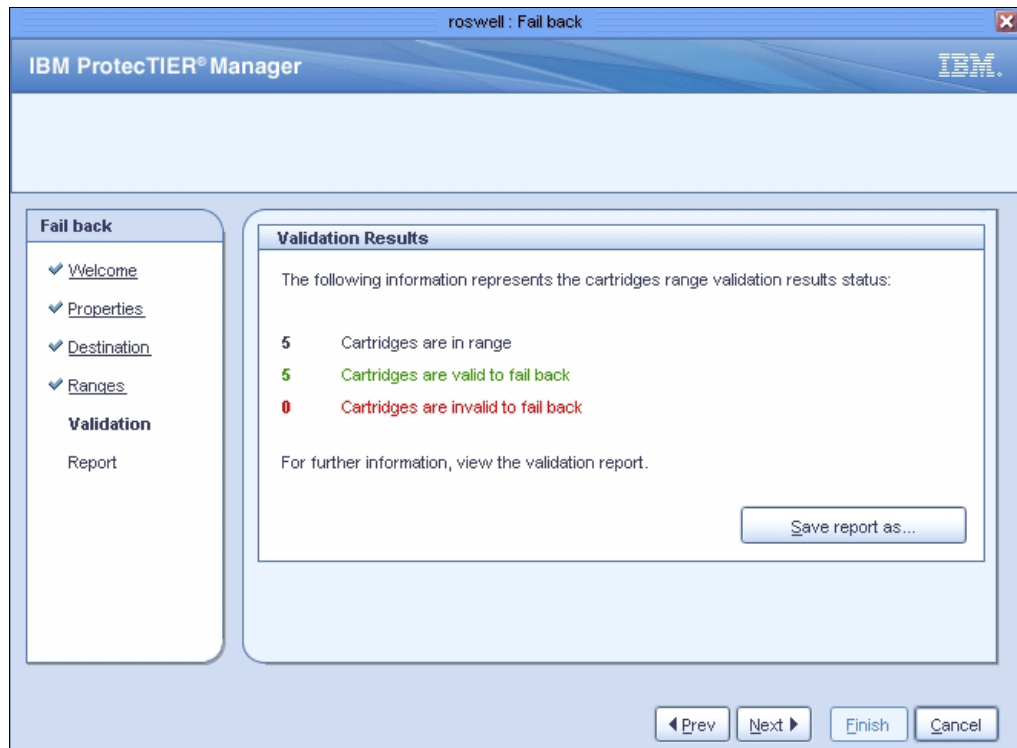


Figure 11-11 Create failback policy validation window

In our example, we select give cartridges for the failback scenario. Click **Save reports as** if you want to save the validation summary report.

- Click **Next** to view the validation summary report (Figure 11-12). Click **Finish** to create the failback policy.

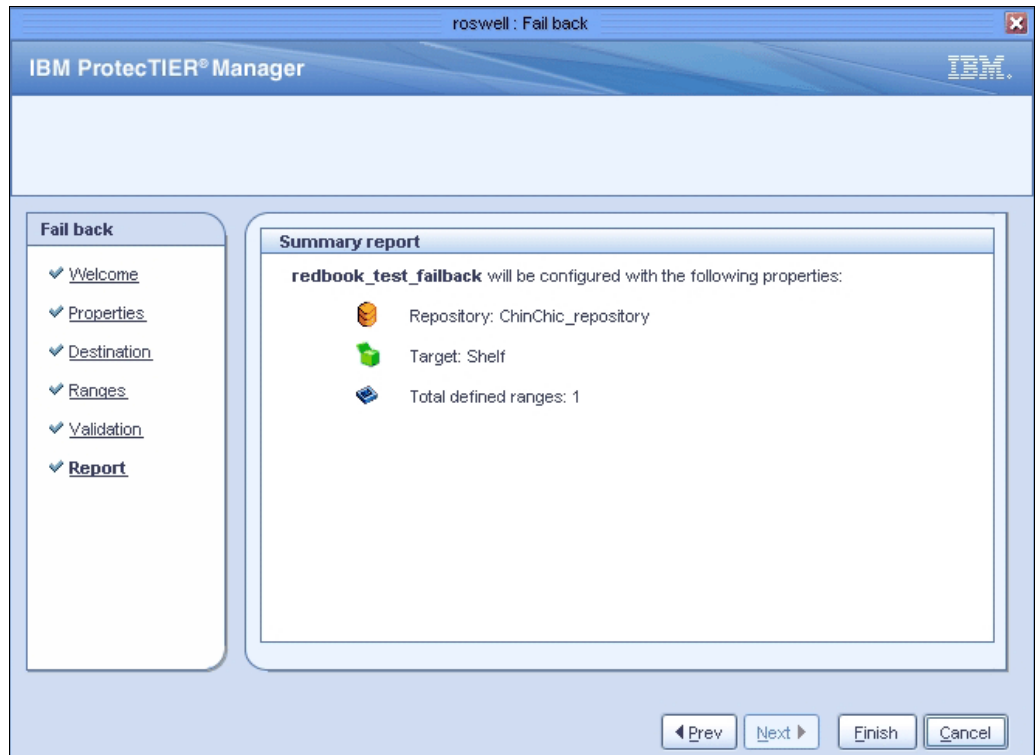


Figure 11-12 Create failback policy summary report

- No message is displayed after you create the failback policy. On the ProtecTIER view pane, click the DR repository, then **Replication Policies**. Here you see your newly created policy (Figure 11-13).

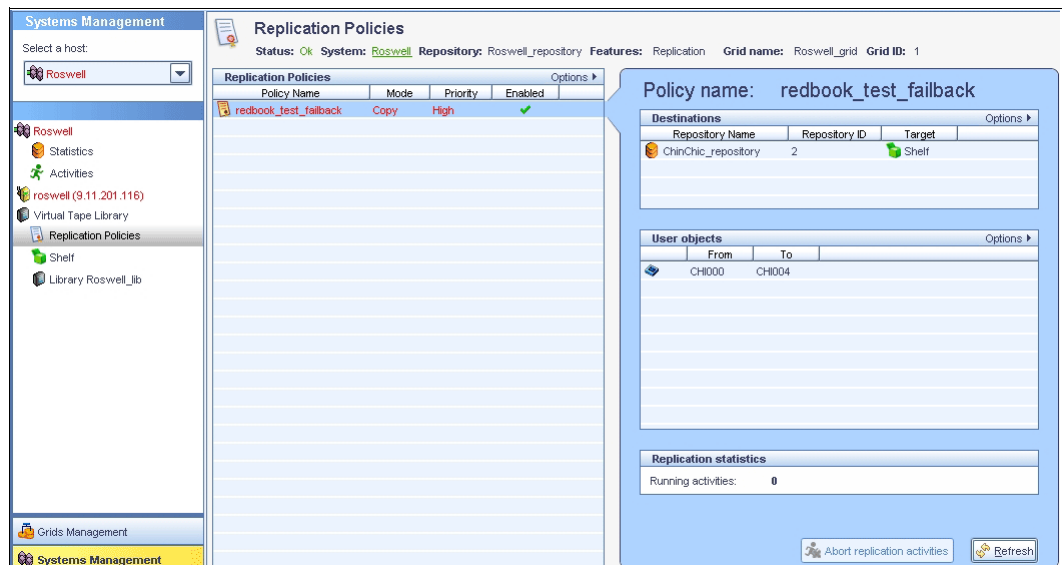


Figure 11-13 Failback policy

- Right-click your policy and click **Execute** (Figure 11-14). The failback policy starts. Depending on the number of cartridges and the amount of data, it can take some time for the failback procedure to finish. You can observe the progress of replication in the Activities view in your ProtecTIER view pane.

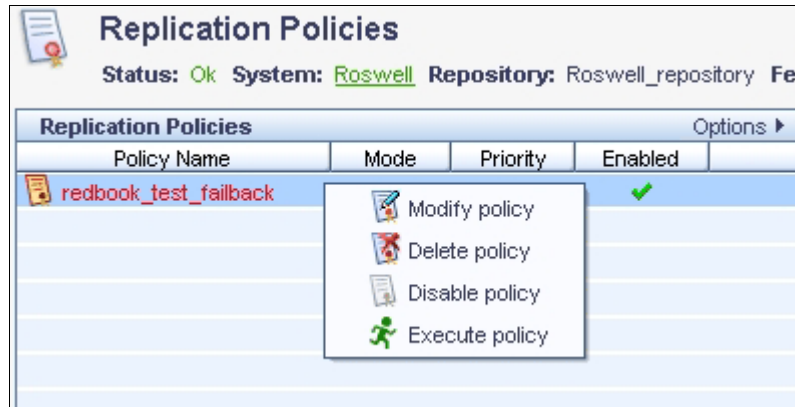


Figure 11-14 Execute replication policy

## 11.3 Recovery management

This section describes procedures that help to clarify and manage a DR situation on an IBM ProtecTIER system.

### 11.3.1 Dirty bit attribute

The dirty bit attribute (in-sync) should help you with consistency points during disaster recovery.

When a cartridge is fully synchronized between hub and spoke, the hub sets the dirty bit to off. During DR, that means that the cartridge is fully synchronized, not only at the consistency point, but also after the last replication took place. If a cartridge is out of sync during DR time, we need to explore the time that the cartridge was fully synchronized to determine which consistency point to which it belongs. The user can take advantage of some new ProtecTIER command-line interface (PTCLI) commands to best use the dirty bit (in-sync) attribute. The PTCLI commands are described in detail in 11.3.2, “ProtectTIER command-line interface” on page 587.

All the cartridges 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 the replica was updated, with 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 a consistent point in time. Otherwise, the cartridge is incomplete or in transit. If the cartridge has a consistent point in time, ensure that this time stamp is larger than the full backup image end time, which indicates that the cartridge contains all the required data for this recovery operation. Otherwise, the user will have to use a previous full backup image for recovery.

**Note:** 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 the DR administrator or operator(s).



The user 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 may finish replicating before the backup ends and they have a sync point earlier than the backup end time. As such, 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.

**Note:** Following an upgrade to V2.4 or V2.5 (backup or replication), all the known information on the cartridge (available in V2.3) will still be displayed, but the dirty bit “In-Sync” column will be empty until there is activity. Then the In-Sync will be populated accordingly, as shown in Figure 11-15 on page 587.

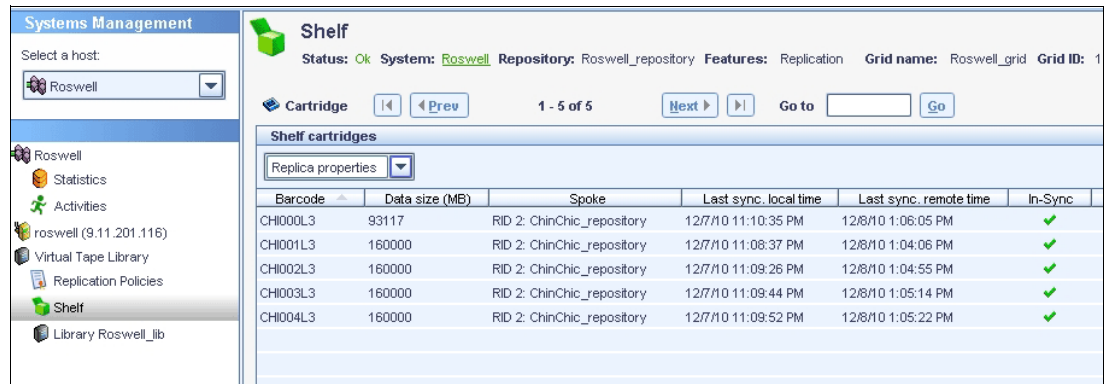


Figure 11-15 ProtecTIER VTL hub cartridge replica properties

### 11.3.2 ProtecTIER command-line interface

This section describes how to configure and monitor the ProtecTIER server through the command-line interface during a disaster recovery situation using the PTCLI.

The PTCLI is loaded during the installation of ProtecTIER software and ProtecTIER Manager software.

#### Creating a profile

It is necessary to create a profile before you work with the PTCLI. Open a command line or terminal window and change the directory to the installation folder of the ProtecTIER Manager. Activate the PTCLI by running `ptcli -p` followed by the full path of a file. The user file will be created (the folder must be already created).

Here is an example of this procedure (in this case, we are using a Windows workstation running the ProtecTIER Manager):

- ▶ The path is `C:\Program Files\IBM\ProtecTIER Manager>ptcli -p c:\ptcli\ptuser.`
- ▶ The user name is `ptadmin.`
- ▶ The password is what you set during the installation of ProtecTIER.

You should receive the following output:

```
<?xml version="1.0" encoding="UTF-8"?>
<response command="createprofile" status="success"/>
```

## Usage

You can perform the following tasks using the PTCLI:

1. Configure ProtecTIER (including a ProtecTIER repository and ProtecTIER virtual libraries).
2. Monitor ProtecTIER (including ProtecTIER repository and ProtecTIER virtual libraries statistics).
3. Filter and snapshot ProtecTIER virtual cartridges (mostly used for DR scenarios).

## Query types and cartridges set

A filter will always work on the specified set of cartridges. The set of cartridges is specified in the PTCLI command as `querytype` (you can view examples in “Useful queries in a disaster recovery situation” on page 589). The following query types can be used:

- ▶ All: All cartridges in the ProtecTIER repository
- ▶ Replica: All cartridges that were replicated into the repository
- ▶ Origin: All cartridges that were replicated from the repository

## Inventory command options for a DR scenario

The inventory command options are used to filter cartridges in a ProtecTIER repository using a variety of criteria. Also, these options can be used to move cartridges that match a certain criteria. Before filtering the cartridges, or moving cartridges using the PTCLI, you must first create a snapshot of the cartridges using the **InventoryRefresh** command. The snapshot will include the most updated properties of the cartridges at the time it is created.

**Note:** Any filter/move operation is executed using the snapshot's contents.

Running such a command without a previous refresh produces an error. Also, for larger repositories, a refresh operation may take a considerable amount of time and reduce ProtecTIER's performance during that time. Moving a cartridge using the PTCLI might fail if the snapshot is not up to date for the cartridge (for example, if the cartridge is moved or deleted after a snapshot is taken).

**Note:** Operations on libraries for which the snapshot is not up to date may have undesirable consequences.

## Disaster recovery with PTCLI

You should use the PTCLI in the following order:

- ▶ To figure out the consistency point at the hub, when all cartridges from a specific spoke were fully synchronized (replicated).
- ▶ For automatic bulk movement of cartridges from shelf to library and back.
- ▶ To figure out the new cartridges that were created at the hub during DR mode, and bulk move them to a shelf so they can be replicated back to the spoke.

Keep in mind the following consideration when using the PTCLI **snapshot** command:

- ▶ The PTCLI snapshot may take up to 15 minutes to be created and populated.
- ▶ The snapshot is a static one, so it reflects the state of all the cartridges only at the point in time that it was taken.

## Disaster recovery scenario

When a disaster occurs at a spoke while replication was running:

- ▶ A DR condition for the specific spoke is “declared” at the hub.
- ▶ The user turns to the DR site (the hub) with the goal of determining the last full backup so its data can be recovered.
- ▶ The user can use PTCLI to sort through the repository and determine which cartridges were in-sync at the time of the disaster.
- ▶ The user can use PTCLI to determine which cartridges were not in-sync at the time the last full backup at the spoke was completed.
- ▶ The user decides which cartridges to use at the DR site and uses PTCLI to move them (all or some) from the shelf to the library.

Using the latest PTCLI functionality in ProtecTIER V2.5:

- ▶ Results can be saved in a .csv file using the --output option.
- ▶ The resulting .csv file can be used as an input to a PTCLI `move` command.
- ▶ The resulting .csv file can be edited, and the user can remove lines (each line represents a cartridge).
- ▶ Users can also create their own barcodes file and use them as an input to a `move` command.

## Useful queries in a disaster recovery situation

The following queries can be useful if you run a DR through PTCLI. These examples are the commands to run on a Windows system running the ProtecTier Manager. If you are running them on Linux, you have to enter `./ptcli` followed by the same extensions. You have to replace the IP address with the IP address of your node and the path with the path to the login ptuser file. The ptuser file is the file you created in “Creating a profile” on page 587.

- ▶ Creating a snapshot (the first thing to do before running any other queries);  

```
ptcli InventoryRefresh --ip 9.11.200.233 --login c:\ptcli\ptuser
```
- ▶ Getting all in-sync cartridges:  

```
ptcli InventoryFilter --ip 9.11.200.233 --querytype replica --query "in_sync = true" --login c:\ptcli\ptuser --output c:\ptcli\non_dirty_cartr
```
- ▶ Getting all not in-sync cartridges (dirty bit flag):  

```
ptcli InventoryFilter --ip 9.11.200.233 --querytype replica --query "in_sync = false" --login c:\ptcli\ptuser --output c:\ptcli\dirty_cartr
```
- ▶ Getting all cartridges synchronized with a destination at a certain time range on a source:  

```
ptcli InventoryFilter --ip 9.11.200.233 --querytype replica --query "source_time_for_last_sync_point > datetime('2010-11-30 11:15:00')" --login c:\ptcli\ptuser
```
- ▶ Getting all cartridges replicated to repository 3 in grid 0:  

```
ptcli InventoryFilter --ip 9.11.200.233 --querytype origin --query "destination_repository_id = 3 and destination_grid_id = 0" --login c:\ptcli\ptuser
```
- ▶ Getting all cartridges in barcode range ZX000 and ZX999:  

```
ptcli InventoryFilter --ip 9.11.200.233 --querytype all --query "barcode > ZX000 and barcode < ZX999" --login c:\ptcli\ptuser --output c:\ptcli\barcode_file
```

- ▶ Moving all in\_sync cartridges to the shelf:

```
ptcli InventoryMoveFilter --ip 9.11.200.233 --querytype replica --query
"in_sync = true" --destination shelf --login c:\ptcli\ptuser
```

## 11.4 Disaster recovery with TS7650 OpenStorage and Symantec NetBackup

This section describes the utilization of IBM System Storage ProtecTIER systems in a NetBackup environment and discusses the ramifications of possible scenarios as it relates to disaster recovery (DR).

### 11.4.1 NetBackup background

NetBackup is an Open Systems Enterprise backup software solution. It has three main components:

- ▶ Clients: The machines with the data that needs backing up
- ▶ Media Servers: The machines connected to backup devices
- ▶ Master Server: The machine controlling the backups

Collectively, Master, Media, and Clients are known as a *NBU Domain*. Typically, one Master controls multiple Media Servers (typically 4 - 30) and back up many Clients (typically 10 - 1,000 +). As the Master server is the critical component in the domain, it is usually clustered, deploying other software available from the vendor (a host based volume manager for disk mirroring and cluster control through a Veritas Cluster server). For more details about the NBU domain scheme, refer to Figure 11-16. In this example, an IBM ProtecTIER server is the back-end storage.




Figure 11-16 NBU Domain with Client, Media, and Master server

## Setting up NetBackup for backup

NBU 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

Most backup and restore deployments now use differential incremental backups, because they are smaller and faster, but cumulative backups are now becoming more common when people are setting up DR.

## Setting up NetBackup for disaster recovery

When using NBU for disaster recovery planning, the user should consider a number of key issues:

1. NBU architecture: Does the NBU 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 on DR.
2. Classification of Clients (RTO): When a company plans for DR, each server is given a Recovery Time Objective (RTO) depending on the importance of its application and the associated data to the business. Servers with short RTOs (typically less than 24 hours) will 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 RTOs typically greater than 24 hours use tape for DR. Servers are then prioritized into RTO bands of typically 24, 36, 48, or 72 hours, depending on business requirements.
3. Classification of Clients (RPO): Running alongside RTO is the Recovery Point Objective (RPO). This is the point in time to which the server must be recovered. For the majority of servers using a tape DR position, 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 night's backup.

Given these disaster recovery requirements the following architecture has become common:

1. Architect NBU with a single domain spanning both sites (NBU Clustered). The Master uses host based replication to mirror the NBU databases and a clustering product to manage host failover. This is important as it means that in the event of a DR event, the Master operations can seamlessly fail over to the DR site. As the NBU databases have been replicated, all of the backup information is known at the DR site and restores can begin immediately.
2. Cross-site backups: Two main options:
  - Connect Clients from one site through IP to media servers on the DR site: Backups then reside in the DR site library and are ready for restore. The primary downside is that large IP pipes are required and backups are limited to speed of the cross-site network, because all of the data is being transmitted.
  - 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 then reside in the DR site library and are ready for restore. The downside is that large SAN pipes are required and backups are limited to the speed of the cross-site SAN, because all of the data is being transmitted.

The downside of both options is that now that normal backups are resident in the DR library, any regular restores will be significantly slower, because data has to come from a DR library.

3. Multiplexing turned off: To achieve the best restore performance (that is, meet RTOs) NetBackup needs to be configured without multiplexing.
4. Dedicated volume pools of for RTO tiers or even clients: To achieve optimum restore times (and given sufficient media in the libraries), having individual volume pools per client will achieve optimum restore performance. This way, there is no contention between media when doing restores. In the physical tape world where tape drives are limited, this is often impractical, but is more practical in a virtual environment. It should be stressed that this is not just a theoretical concept; systems in current production have implemented cross-site backups with client backups going to dedicated volume pools, although this was limited to 30 clients with low RTOs because the implication of separate volume pools is that you need separate backup policies per client.

If the NBU 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 son on, that have been used by the primary site, the first operation is to get a copy of the NBU catalog from the primary site and load into the Master on the DR site. This task can either be done through disk replication or tape backup.

**Note:** NBU 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 would have to 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.

### **Optimal ProtecTIER deployments with NBU for disaster recovery**

The following are key concepts that need to be discussed with the NBU architects and senior administrators within the user's organization:

1. For normal operation, back up to the local ProtecTIER system, as this provides quick backups and quick restores.
2. For ProtecTIER OST replication only, the deduplicated data is being transferred over the network, which will significantly reduce the bandwidth needed compared to traditional cross-site replication/backups.
3. Have servers for DR (usually production servers) split into their RTO classifications and plan for separate volume pools and backup policies.
4. For servers with low RTO requirements, consider individual volume pools and backup policies.
5. Turn multiplexing (MPX) off for *all* backups requiring DR. Because MPX is done at either the Storage Unit level or Backup Policy level, this is easy enough.
6. Use large fragment sizes. These fragment sizes are also configured at the Storage Unit level, which improves the restore performance of whole file systems.
7. Disable Storage Checkpoints. Storage checkpoints are a mechanism where pointers are added to the backup stream so that if the backup failed, a rerun of the backup would start from the last storage checkpoint, as opposed from the start of the backup. Storage Checkpoints have an adverse effect on the deduplication ratios.
8. Disable software compression (if enabled), as it might reduce the efficiency of the ProtecTIER deduplication and affect its factoring ratio.

After the user's architects and the administrators understand the basic concepts, they need to apply them to their architecture, deciding whether to have one domain spanning both sites, or two separate domains.

If single domain/same catalog is shared across sites, and is always updated with the whereabouts of all cartridges, the following tasks should occur:

1. ProtecTIER replicates the cartridges per the policies set by the user. The cartridges are copied onto a virtual shelf of the ProtecTIER system at the DR site.
2. Cartridges can also be moved using the replication policy *and* using the visibility control switch, so they will reside and be visible to the NBU application at the DR site (although the actual data will be available to ProtecTIER on both sites) by completing the following steps:
  - a. Eject the (export) cartridge from the primary library.
  - b. Inject (import) the cartridge to the inventory at the DR site library.This operation can be set / done using the NBU Vault or manually. Either way, it can be automated from within the NBU environment.
3. If disaster occurs, the user either needs to inject the cartridges from the DR site shelf into the DR site library and inventory or, if the visibility switch control method was used, the user begins restoring or performing local backups at the DR site.
4. After the disaster situation has been resolved and the primary site is back online, the user should use the failback procedure (as explained in *IBM System Storage TS7600 with ProtecTIER User's Guide for Enterprise Edition and Appliance Edition*, GC53-1156) to move their main operation back to the primary site, including potential newly created cartridges from the DR site that will be replicated to the primary site.

If the separate (multiple) domains approach is used, the following task should occur:

1. ProtecTIER replicates the cartridges per the policies set by the user. Cartridges are copied to the virtual shelf at the DR site.
2. The user should perform catalog backups to virtual tape at the end of its backup window and replicate the backups 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/replication cycle) that the DR site will hold a full set of replicated cartridges with a matching NBU catalog, allowing for an RPO of one day.
3. When disaster strikes, the user should get the NBU catalog back on the DR site NBU server by restoring the cartridge(s) containing the catalog. Then the user should inject the cartridges from the DR shelf on the ProtecTIER system at the DR site into the library and perform an inventory. After the NBU server is up and running with the DR repository, restores and local backup operations can resume at the DR site.
4. After the disaster situation has cleared and the primary site is back online, the user should use the failback procedure (as explained in *IBM System Storage TS7600 with ProtecTIER User's Guide for Enterprise Edition and Appliance Edition*, GC53-1156) to move their main operation back to the primary site, including potential newly created cartridges from the DR site that will be replicated to the primary site.

## 11.4.2 Disaster recovery scenarios

ProtectTIER replication significantly reduces cross-site backup traffic because it only replicates deduplicated data, improves ease of operation (by enabling simple inject and inventory actions), and makes recovery in the event of a disaster or DR test easy to plan and implement. Deploying ProtectTIER into a NBU environment makes a business more secure and reduces significant amounts of work for NBU architects and administrators.

The following sections provides a number of scenarios that detail the necessary disaster recovery steps.

### Single domain environment

In a single domain environment, there are two possible scenarios:

► **Master Clustered complete:**

All backups are completed, and all of the replication operation is completed as well. Disaster strikes the Master Clustered environment, the NBU catalog database at the DR site is up to date, so no NBU recovery action is necessary. Within ProtectTIER, the user should move tapes from the virtual shelf to import slots. Within NBU, the library needs to be inventoried; remember to select the option to import tapes. After the inventory operation is complete, restores can begin as well as local backups at the DR site.

► **Master Clustered incomplete:**

All backups are completed, but some or all of the replication operation is incomplete. Disaster strikes the Master Clustered environment, the NBU catalog database at the DR site is up-to-date, but because the replication was not complete, the user should roll back to the previous night's catalog and cartridges set (RPO of one day). After the inventory operation is complete, restores can begin, as well as local backups at the DR site.

### Multiple domain environment

In a multiple domain environment, there are two possible scenarios:

► **Master not Clustered complete:**

All backups are completed, and all of the replication operations are completed as well. Disaster strikes the Master not Clustered environment, the catalog database at the DR site is *not* up-to-date, which means that NBU catalog recovery action is necessary.

The first operation is to identify the latest backup catalog tape and load (import) it into the ProtectTIER library at the DR site. After the library is inventoried, a standard NetBackup Catalog Recovery operation can begin. When recovery is complete, then restores can begin as well as local backups at the DR site.

► **Master not Clustered incomplete:**

All backups are completed, but some or all of the replications operation is incomplete. Disaster strikes the Master not Clustered environment, the catalog database at the DR site is *not* up to date, which means that NBU catalog recovery action is necessary

The first operation is to identify the previous nights NBU backup catalog tape and load (import) it into the ProtectTIER library at the DR site. After the library is inventoried, a standard NetBackup Catalog Recovery operation of that previous night's catalog can begin. After recovery is completed, restores can begin (RPO of one day), as well as local backups at the DR site.



**Note:** When working in a single-domain NBU environment (NBU Master Clustered) *and* using the visibility control switch option within ProtecTIER to move cartridges from the primary site directly into a DR site library, the catalog is *always* up to date with the whereabouts of all cartridges in both the primary and DR repositories.

## **Determining what is available for restore at the disaster recovery site**

This section suggests ways for the users to determine what catalog and data sets are complete or not, matched, and readily available to restore at the secondary/DR site.

### ***Which database copy at the DR site is valid***

Before running a restore for disaster recovery, the user must verify that the list of associated cartridges is completely replicated to the DR site; otherwise, an earlier full backup image must be used for recovery (usually the previous night's).

The easiest way to determine the time of the last full backup is if the user has a specific time each day where the replication backlog is zero (there is no pending data to replicate).

If this is not the case, then the user 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.

The best practice for ensuring that a copy of the catalog is available at the DR site is to use the native replication function of ProtecTIER. Each day the catalog should be backed up on a virtual cartridge following the daily backup workload such that it will be replicated to the DR site at the end of each replication cycle.

If the catalog is backed up to a virtual cartridge, through the cartridge view of the library in ProtecTIER Manager, query each of the cartridges used for catalog backup to find the most recent sync dates marked on the cartridges. Assuming there are multiple backup copies, you need to find the latest backup that finished 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 has a Last Sync Time that displays the last time the cartridge's data was fully replicated to the DR site. (The sync time will be updated during the replication, not only when the replication for this cartridge is finished.)
- ▶ The cartridge marked with the most recent Last Sync Time date should be used to recover the backup application catalog.

### ***Determining which cartridges at the DR site are valid for restore***

After the DR site NetBackup server is recovered, you need to review the status of the replicated cartridges to ensure their replication consistency with NetBackup catalog. Use the available ProtecTIER system Replicated Cartridges Status Report to accomplish this task.

### ***Eject and inject commands from the backup application***

Although the process can be manually scripted to enable automation, the easiest way of using NBU CLI commands for automating this process is by using the Vault service within the NBU software.

### ***How to eject a cartridge from a library***

You can eject a cartridge by using a wizard of the NBU GUI or by using by using the Vault option. Complete the following steps:

1. If you are using a Vault command, first run your vault policy:

```
/usr/opensv/netbackup/bin/vltrun<vault policy name>
```

2. At the end of the backup, eject the cartridge using the following command:

```
/usr/opensv/netbackup/bin/vltinject<vault policy name>
```

### ***How to insert a cartridge into a library***

A simple robot inventory selecting the **Import from Load port** option can insert the cartridge into a library. For automation of this process, you can use CLI commands.

To update the Media Manager volume, run the following command:

```
/usr/opensv/volmgr/bin/vmupdate -rt dlt -r
```

### ***NBU procedure to recover a master server from an existing DB copy***

You have two options for re-creating the NBU backup catalog: online and offline.

- ▶ The hot online catalog backup procedure can be found under the “Online, hot catalog backup method” section of the NBU Help menu.
- ▶ The offline cold catalog backup procedure can be found under the “Offline, cold catalog backup method” section of the NBU Help menu.

**Note:** Consult the official NetBackup application documentation for more details.

## **11.5 Completing failback and leaving DR mode**

After the primary site is rebuilt or replaced, create the failback policy and wait for it to complete. When the data is returned to the primary site, you can exit DR mode and continue with backups and replication as usual.

The failback policy is a one-time policy. When failback is complete, the policy still remains in the Replication Policies list. Running or modifying this policy is not allowed, and if you try, an error message is displayed.

**Note:** Do not run any backup activities on the restored repository until the failback operation is complete.

When the failback operation is complete, complete the following steps to exit DR mode:

1. Select **VT** → **Replication Disaster Recovery** → **Leave DR mode**. The Leave DR mode window opens (Figure 11-17).

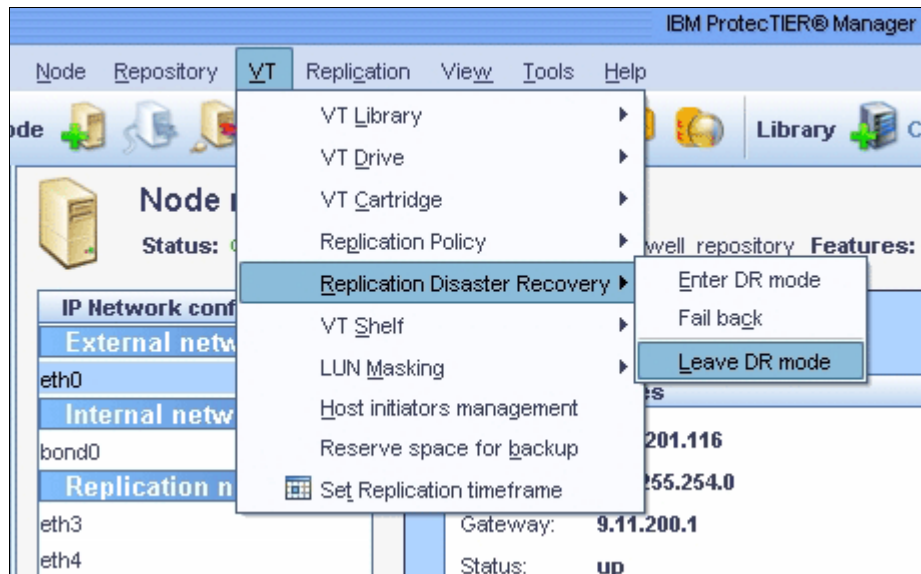


Figure 11-17 ProtecTIER leave DR mode

2. From the drop-down menu, select the spoke repository from which you want to exit DR mode (Figure 11-18).

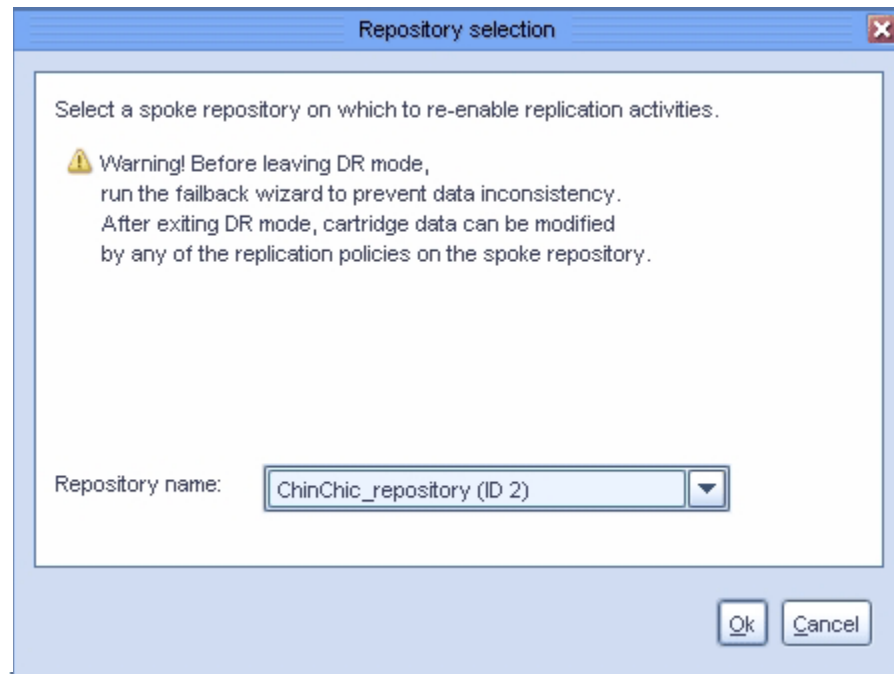


Figure 11-18 Select the spoke repository on which to leave the DR mode

3. Confirm your decision (Figure 11-19). After the repository is out of DR mode, any replication policy defined on the spoke repository can modify the cartridge data. You can begin working again with the restored repository.




Figure 11-19 Confirm selection

## 11.6 Principality

Principality is the privilege to write to a cartridge (set it to R/W mode). The principality of each cartridge belongs only to 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. The principality of a cartridge can be transferred only in one of three cases:

- ▶ The principality belongs to the DR repository.
- ▶ The principality belongs to the original primary repository and this site is the destination for the failback.
- ▶ The principality belongs to the original primary repository, but:
  - The original primary repository is out of the replication grid.
  - The target for the failback is a repository that was defined as a replacement repository through the ProtecTIER repository replacement procedure (see 11.7, “Repository replacement” on page 600).

You can see the principality in the backup properties view of the cartridge menu on a hub's shelf (Figure 11-20).

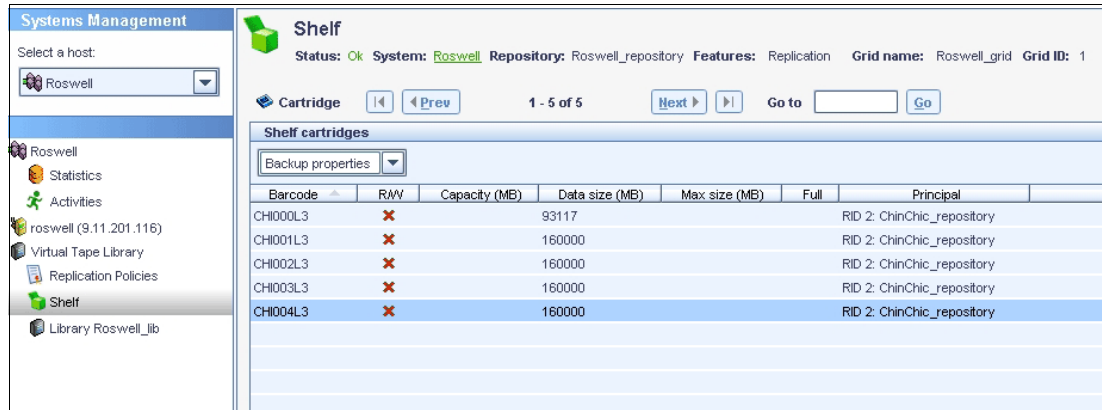


Figure 11-20 Principality from the ProtecTIER hub point of view

### 11.6.1 Taking over principality

This section describes how to initiate a principality takeover by a ProtecTIER DR site in case of a failing production ProtecTIER system.

#### About this task

Taking principality of the cartridges on a destroyed repository allows you to write cartridges previously belonging to the replaced repository.

#### Before you begin

Cartridge principality takeover is used to allow the local repository, or hub, to take control of cartridges belonging to a destroyed repository. The repository can only take principality of a cartridge if the repository is defined on the Replication Manager as the replacement of the destroyed repository. Select **VTL Repository takeover** from ProtecTIER Replication Manager before performing the task shown in "Procedure" (Figure 11-21).

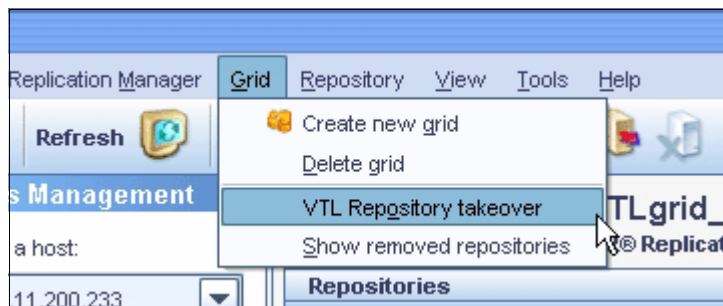


Figure 11-21 ProtecTIER VTL repository takeover

## Procedure

Complete the following steps:

1. Select **VT** → **VT Cartridge** → **Cartridge ownership takeover**. The Cartridge ownership takeover window opens (Figure 11-22).

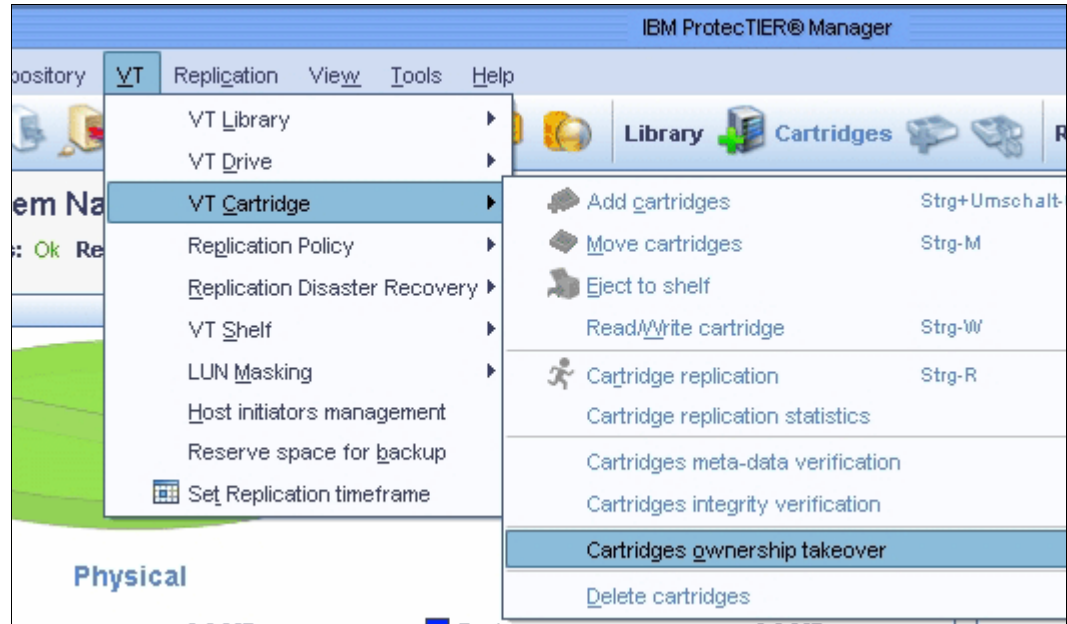


Figure 11-22 ProtecTIER VTL cartridge ownership takeover menu

2. Choose, from the drop-down list, a repository that was replaced by the current repository.

**Note:** Make sure that the replaced repository is inactive and unrecoverable to prevent conflict with the data on the cartridges.

3. Click **OK**. The principality of all cartridges belonging to the failing repository now belong to the current repository.

## 11.7 Repository replacement

If the user wants to fail back to a different or a rebuilt primary repository, complete the following steps:

1. Cancel the pairing of the original repositories in the ProtecTIER Replication Manager.
2. Take the original primary repository out of the replication grid.
3. If a new repository replaces the original one, then the new repository must be installed and join the replication grid. If it is an existing repository, it must be out of a replication pair.
4. Run the ProtecTIER repository replacement wizard and specify which repository is replacing which repository.

**Note:** This operation is *not* reversible.

If you run the repository takeover wizard, the ProtecTIER Replication Manager completes the following actions:

1. Deletes the current grid pair (if it exists) between the unrecoverable primary and the remote side (if you have not done that previously on your own).
2. Removes the unrecoverable repository from the grid.
3. Creates a connection between the replacing repository at the primary site and the new remote site.

### 11.7.1 Replacing a destroyed VTL repository

This section describes the tasks needed to initiate a ProtecTIER VTL repository takeover that replaces the destroyed repository at the production site with the repository at the disaster recovery (DR) site.

#### Before you begin

If the source repository at the production site (spoke) has been destroyed and cannot be repaired, you can replace it with a new repository. The replacing repository must be part of the grid.

**Note:** Be careful when replacing a repository that you do not accidentally replace a good repository.

#### Procedure

Complete the following steps:

1. Enter DR mode, as described in 11.1, “Moving to ProtecTIER DR mode” on page 576.
2. Run the Repository takeover wizard by selecting **Grid** → **Repository takeover** (Figure 11-23).

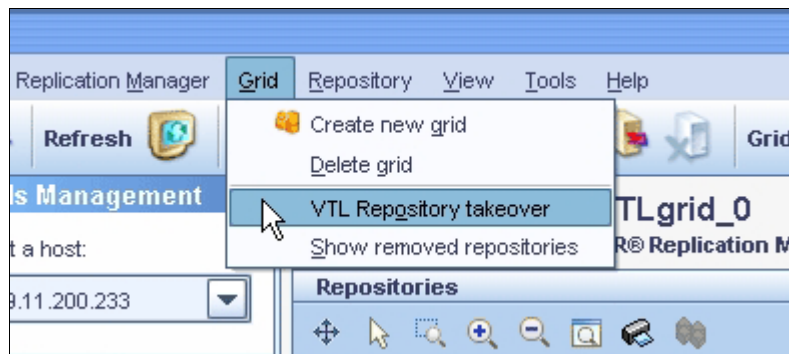


Figure 11-23 ProtecTIER VTL repository takeover

3. Run the Repository takeover wizard. From the Grids Management view, select **Grid** → **Repository takeover**.
4. Select **Unrecoverable repository** from the drop-down menu.
5. From the Replacing repository drop-down menu, select the new repository that you want to replace the unrecoverable repository.
6. Click **OK** to begin running the Repository takeover wizard.

## Results

When you run the Repository takeover wizard, the ProtecTIER Replication Manager will perform the following actions:

1. Delete the current grid pair, if a pair exists, between the unrecoverable primary site and remote site (if you have not done this previously on your own).
2. Remove the unrecoverable repository from the grid.
3. Create a connection between the replacing repository at the new primary site (spoke) and the remote site.

**Note:** After you finished that task, create a failback policy and copy the cartridges back to the primary site, as discussed in 11.2, “Failback for all data at the DR site” on page 579.

## 11.8 Restoring the ProtecTIER Replication Manager

This section describes how to analyze and restore an inactive ProtecTIER Replication Manager using the ProtecTIER Manager.

**Note:** After the ProtecTIER Replication Manager server is restored, you must make sure that the lost or inactive server is never reactivated. Reactivating a lost or inactive server may cause inconsistencies and may conflict with the restored server.

It is only possible to restore the ProtecTIER Replication Manager grid by grid. There are two possibilities to restore a ProtecTIER Replication Manager:

- ▶ Using an existing backup file, such as the following:

```
PtReplicationManager_<hostName>_<customerName>_<time>_ConfigurationBacku .zip
```

**Important:** Use the most updated file determined by the newest time value.

- ▶ Using an accessible repository that was once a member of the grid



In the ProtecTIER Manager, select **Replication Manager** → **Restore grid** (Figure 11-24).

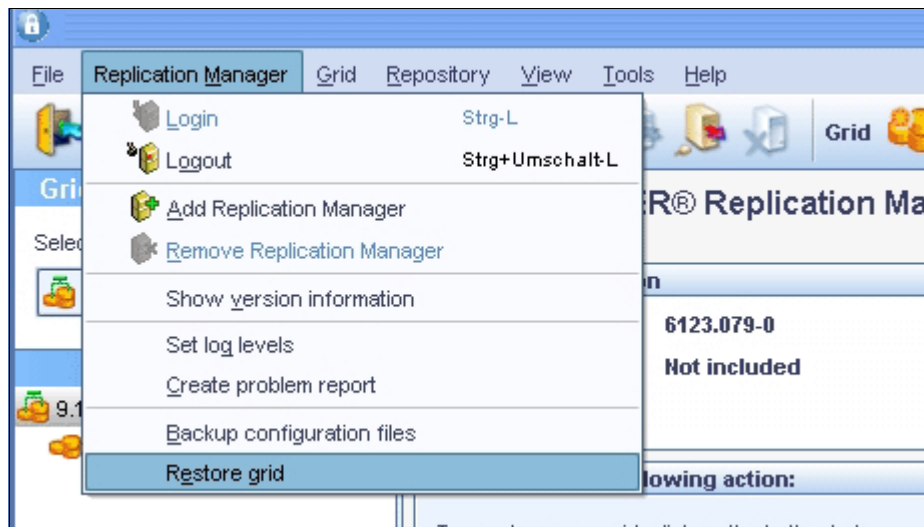


Figure 11-24 Open restore grid procedure in ProtecTIER Replication Manager

The following sections describe the different methods for restoring the processed grid using the Restore grid procedure. The ProtecTIER Replication Manager server automatically handles *grid analysis*. Grid analysis validates and checks the integrity of the grid in relation to its grid members. To begin restoring a grid, import a grid for analysis by selecting one of the following tabs appearing in the Restore grid window:

- ▶ **Default backup**
- ▶ **File**
- ▶ **IP Address**

### 11.8.1 Restoring from default backup

The following section describes how to restore the Replication Manager using the Default backup tab.

If ProtecTIER Replication Manager is installed on a ProtecTIER server, the entire Replication Manager is automatically analyzed and restored.

To import a grid for analysis and restore, complete the following steps:

1. Select **Replication Manager** → **Restore grid**. The Restore grid window opens.
2. Select the **Default backup** tab.
3. Click **Read file**. The grids are automatically retrieved from the backup configuration file that was found on the server and displayed. Before restoring a grid, you can preview the grid layout by clicking the **Grid preview** tab. If all the repositories are accessible and validated successfully, go on to the next step.
4. Click **Restore**. The current grid is restored and the next grid for analysis is displayed.
5. Continue restoring, grid by grid, until the entire Replication Manager is restored.

## 11.8.2 Restoring from file

The following section describes how to restore the Replication Manager using the File tab.

If you are restoring a grid from a backup compressed file, locate the grid data files in `/gm_work/data/grid.<x>.xml`, where `<x>` represents the grid ID number (Figure 11-25).

```
Last login: Thu Dec  9 20:17:31 2010 from 9.11.144
italy: /opt/dtc/ptadmin> cd /gm_work/data
italy: /gm_work/data> ls
grid.0.xml
italy: /gm_work/data> █
```

Figure 11-25 ProtecTIER grid xml backup file

To import a grid for analysis and restore, complete the following steps:

1. Select **Replication Manager** → **Restore grid**. The Restore grid window opens.
2. Select the **File** tab.
3. Browse to the backup file name from which you want to analyze and restore the grid, or enter a file name in the Backup .ZIP file name field.

**Note:** If you already extracted the compressed backup file, there is a folder structure already, as shown in Figure 11-26 on page 604, including the `/gm_work` and the `/data` folders, and you can select the grid XML file instead.




Figure 11-26 Folder structure for extracted backup .zip file

4. Click **Read file**.
5. Select a grid to analyze from the drop-down menu. The grid analysis window opens.
6. Click **Restore**. The current grid is restored and removed from the grid analysis list.
7. Select the next grid for analysis from the drop-down menu. Continue restoring, grid by grid, until the entire Replication Manager is restored.

## 11.8.3 Restoring from IP address

The following section describes how to restore a grid from the Replication IP address of a repository from the IP Address tab.

If you search for a backup file and cannot find one, then you can restore a grid on which the repository was previously a member. With a given replication IP address, the analysis phase is done with the respective repository's grid data in relation to the other grid members.

To import a grid for analysis and restore, complete the following steps:

1. Select **Replication Manager** → **Restore grid**. The Restore grid window opens.
2. Select the **IP Address** tab.
3. Enter the Replication IP address, the Port number, and the Ping port number of the accessible repository that was previously a member of the grid.
4. Click **Analyze grid**. The grid analysis window opens.
5. Click **Restore**. The current grid is restored.

**Note:** For further information and details about restoring a ProtecTIER Replication Manager and grid, refer to *IBM System Storage TS7600 with ProtecTIER User's Guide for Enterprise Edition and Appliance Edition*, GC53-1156.

## 11.9 Returning to normal operations

After a disaster situation ends, the primary site is either back online or is rebuilt and now has an empty repository. Failback is the procedure for replicating back updated cartridges, new or old, from the remote DR site to the original (or restored) production site to bring it up to date in case it was down or lost and rebuilt. If the primary repository was down and has been restored, you can return to normal operation at the production site as the primary site and use the DR site as the remote site, or secondary site, again. Define a failback policy on the remote repository and select all the cartridges that were used for backup during the time that the primary repository was down.

**Note:** The cartridges that originally came from the primary site do not need to be replicated because they were not changed at the DR site.

The failback policy procedure transfers the principality of all the cartridges that belonged to the temporary primary repository at the DR site to the restored primary repository at the production site.

The procedure to initiate the failback process using the ProtecTIER Failback wizard is as follows:

1. The user should define a policy with all the cartridges that must be transferred.
2. The ProtecTIER Manager/vtfd creates a policy in a failback mode and with the transfer principality option. This policy can be executed only manually and the system log will ignore runtime events for this policy.
3. The user should approve the execution of the policy. It is translated to manual execution of the policy in the VTL.
4. Cartridges are replicated only if they follow the principality rules described earlier.
5. Before initiating replication, cartridges will be ejected out of the library to the shelf at the remote site.
6. At this point, the user can close the failback wizard.

**Note:** The user cannot perform any editing actions on any failback policy through the normal view.

7. The system supplies and monitors information about the number of pending, running, and completed objects.
8. ProtecTIER Manager presents this information to the user. Thus, the user is able to determine when the failback process is complete.
9. The user is expected to delete the policy after the failback objects are replicated (Figure 11-27).

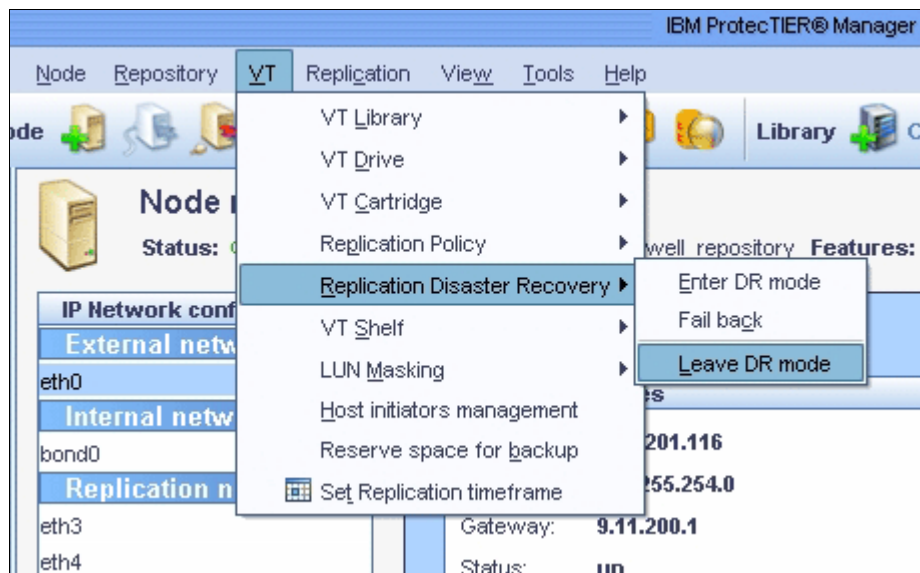


Figure 11-27 Leaving ProtecTIER DR mode

For more details, refer to Chapter 13, “Disaster recovery operations” and Appendix A. “Recovery procedures”, of *IBM System Storage ProtecTIER User Guide for Enterprise Edition and Appliance Edition*, GC53-1156.

## 11.10 Flushing the replication backlog after a long link outage/unavailability

If there is an unscheduled long network or remote site outage, the replication backlog might become too large for the system to catch up.

A replication backlog might be an indication that there is not enough available bandwidth allocated for the replication operation. Remember that for the system to support the organization-set SLAs, enough bandwidth should be planned and allotted for replication during the time frame that it runs (that is, the replication window) so that all the policies will be executed in time.

Consider the following items:

- ▶ There are several ways to *delete* a replication backlog:
  - Abort all replication tasks associated with a specific policy.
  - On ProtecTIER Manager, on the Replication policy view, select the policy and click **Abort Activities**.
  - Abort specific running replication activities.
  - On ProtecTIER Manager, on the Replication activities view, select a specific activity and click **Abort Activities**.
- ▶ Aborting replication tasks removes them from the pending/running tasks.
- ▶ These tasks rerun automatically whenever the specific cartridge is:
  - Appended
  - Ejected from the library
  - Selected for manual execution
- ▶ To prevent these replication tasks from rerunning, mark those cartridges as read only either on ProtecTIER or with the backup application. These cartridges will not be used for further backups, and therefore will not replicate the backlog. New/scratch sets of cartridges will be used for subsequent backups. They will not contain any backlog data that is not required to be replicated.

**Tip:** To resume I/O activity, use different barcodes, as the earlier data on the older cartridges will have to be replicated before the new and the backlog will be huge again. Using a different set of barcodes allows this new data to be replicated, without the need to replicate the data from the old cartridges.

### 11.10.1 Replication modes of operation: Visibility switch versus basic DR

The replication modes of operation that you select shape your disaster recovery plans. The following sections discuss how visibility switching and basic DR affect your disaster recovery plans.

#### Introduction

ProtecTIER Replication can be conceptually described as a two-phase process.

- ▶ The core replication tasks are executed to asynchronously move deduplicated cartridge data across the wire to the target site. This phase is entirely managed by the ProtecTIER and has only one *mode of operation* (that is, there is only one way for data to be physically transmitted to the other site). As soon as a cartridge replication job begins at the source, the respective replica cartridge will be created at the target's system's virtual shelf (unless it already exists). The replica cartridge grows on the target system's shelf as its incoming data is processed. When the replication jobs complete, phase 2 of the overall process comes into effect.
- ▶ This phase, configured through the replication policy, dictates the logical location of the replica cartridges. The options are that either the cartridge is copied to the target side or it is *ejected* from the source virtual library. When the source cartridge is ejected from the source system's virtual library, it automatically moves to that library's import/export slot. As soon as replication completes for the ejected volume (that is, source and replica cartridges are in sync), the replica can either:
  - a. Stay on the target shelf, if it was just copied
  - b. Automatically move to an available import/export slot at the target library

From here on, option (a) is referred as basic DR, and option (b) is referred as visibility switching.

## **Basic DR**

Basic DR can be compared to a disaster recovery strategy with physical tape libraries 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 are rapidly transferred to the DR location and imported into the standby library. The same notion exists with ProtecTIER when using the basic DR mode. Cartridges reside on the target system's virtual shelf ready to be imported to an existing or new virtual library. Granted, the time consumed moving physical cartridges can be far greater than virtual, so this is not an analogy of effort or convenience but that of process flow.

## **Visibility switching**

Visibility switching resembles a warm backup site practice, where physical cartridges are shipped from the source/primary location to the DR site and stored in the standby library's physical slots. When a disaster is declared, the necessary cartridges are immediately available for recovery at the DR site. Visibility switching emulates this process. Beyond serving as a warm DR site, there is a key functional value add that visibility switching provides to select a backup environment that supports a single domain between the primary and the secondary sites.

### ***Added value of visibility switching***

In a virtual world, the difference between the previously described modes is minimal in terms of operational impact. Importing cartridges from the virtual shelf to a library is fast, requires little effort, and can be done from anywhere (with network access to the system). Storing cartridges on the virtual shelf does not make the DR system significantly less reactive for recovery, so the RTO that a ProtecTIER Replication based DR solution can offer represents a significant improvement over a physical tape-based solution.

Furthermore, the major value add of visibility switching is more versatile management for backup applications that support cross-site distributed tape library management of their catalog (that is, within a single domain). Backup applications that can manage multiple sites through a universal catalog/database can use ProtecTIER's automated cartridge movement to easily move cartridges from site to site without using any interface other than the backup application itself. After initially configuring ProtecTIER's replication policy to use visibility switching, the backup application can eject a cartridge from a source library causing, pending completion of replication, the cartridge to appear at an import/export slot of a designated remote library. Likewise, cartridges can be moved back to the source site library using the same process. Having control of the cartridge movement through the backup application simplifies the process of cloning cartridges to physical tape at the target site if desired.

Cutting physical tape from replicas at the target site requires a few additional steps from ProtecTIER if visibility switching is not used. Due to the fact that cartridges can only be visible at one library at any given point in time (because backup applications cannot handle the same barcode at multiple sites), ProtecTIER does not permit a cartridge to be visible in two libraries, even though physically the data exists in both locations. To cut a physical copy at the remote site without visibility switching, an operator would have to import the replica into the target library after replication completes, and also import the source cartridge back into the primary library when the clone job completes and the replica is ejected back to the shelf.

Some of the known major backup applications that support single domain are:

- ▶ Symantec NetBackup
- ▶ Legato NetWorker
- ▶ IBM System i BRMS

Applications that do not support a single domain have no real value when using the visibility switching mechanism, as each library is managed by a separate entity with no shared knowledge of the replicated volumes' content and whereabouts. In these environments, the local backup server must process through a recent backup catalog/database that describes the content of the associated data volumes. Every set of replicated data cartridges imported into a target library must be preceded by a recent catalog/database update (typically stored on one of the ProtecTIER replicated volumes, but may also be replicated through other means).

After the local backup server at the target site has been updated with a recent catalog, the procedure for cutting physical tapes is the same as in a single domain without visibility switching. Cartridges must be moved from the shelf to the library and exported back to the shelf when finished.

One important note that applies mainly to multi-domain environments is that ProtecTIER will not replicate a cartridge if both source and target instances are in libraries. As such, if a cartridge is not exported to the shelf following a clone, but rather is left in the target library while the source cartridge is moved into a library, when new data is appended to the cartridge, replication will not start, and errors will be generated. In turn, when manually moving cartridges into a library, you should verify the state of the other site's cartridge.

## 11.10.2 Use cases to demonstrate features of replication/DR operation

This section includes three flow-chart use cases to demonstrate daily operation features available with ProtecTIER replication in different backup environments. The following tables show step-by-step, per site (primary and secondary), all the actions and activities that are required to perform backups, replication, cloning to physical tape at the remote site, and conducting a DR test.

### Single domain scenario

This use case is for backup applications that allow and support sharing of their catalog/DB across sites/locations (potentially in different geographies), such as IBM Tivoli Storage Manager, Symantec NetBackup, Legato, BRMS, and so on. This situation allows the backup servers on both the primary and remote/DR locations to share the same catalog/DB that makes the ProtecTIER visibility switch control effective and makes the operation of cloning cartridges to physical tape at the remote/DR site simple, as there is no need to bring online and recover the backup application every time that the user wants to do the cloning.

Table 11-1 shows the user operation flow when working in a single domain environment and conducts cloning to physical tape at the remote/DR site. The assumed environment consists of one master NetBackup server managing both sites, including the libraries in both locations (Lib A, A', B).

Table 11-1 User operation flow

NetBackup server/media	
Local site	Remote site
ProtecTIER 1 (local site).	ProtecTIER 2 (remote/DR site).

NetBackup server/media	
Local site	Remote site
Lib A.	Lib A' + Lib B (either virtual or physical).
Prerequisites	
Create repository.	Create repository.
Create libraries.	Create libraries.
Install the ProtecTIER Replication Manager SW module (on the local or remote ProtecTIER node) by performing the following steps: 1. Create Grid A. 2. Add a repository to Grid A. 3. Pair local and remote repositories.	
Create a replication policy to select the cartridges for replication and the specific remote library for the visibility switch.	
Using the Visibility feature use case	
Run regular backup activity to Lib A (to cartridges included in the replication policy).	
Create a vault policy related to Lib A to manage the export/import of the cartridges from and to this library vault_policy_local.	
Run a vault policy (vault_policy_local) to eject all required cartridges from Lib A.	
All ejected cartridges will move from lib A to the repository 1 (local) shelf.	
	Create a vault policy related to Lib A' to manage the export/import of the cartridges from and to this remote library vault_policy_Remote.
	Cartridges that were ejected from Lib A will automatically move from the repository 2 (remote) shelf to Lib A' import/export slots (pending completion of replication/data-sync for each cartridge).
	To enter the cartridges to the Lib A' slots, run a command to import them: 'vltinject <vault_policy_Remote>' This command must be run in a loop until all respective cartridges are imported into the library.
	Duplicate cartridges to physical library (Lib B).
	After the duplication/clone operation is completed, run the vault policy to eject all required cartridges using vault_policy_Remote.
	All ejected cartridges move from Lib A' to the repository 2 (remote) shelf.
Cartridges ejected from Lib A' will move from the repository 1 (local) shelf to Lib A Imp/Exp slots.	



NetBackup server/media	
Local site	Remote site
To move the cartridges into the library, run the following command to import them: 'vltinject <vault_policy_local>' This script must run in a loop until all respective cartridges are imported into the library.	
After all the cartridges are imported, they can be used for new backups.	
	Recovery at the remote site from duplicate/cloned cartridges.
	Restore cartridges from Lib B to the NBU backup server.
	The NBU catalog already contains all records of cartridges in Lib A' and Lib B and therefore can immediately be used for restore operations.

### DR test operation in a multiple (two) domain backup environment

This use case describes the option to perform a DR test to simulate a scenario of disaster at the primary site and the recovery operation from the replicated cartridges at the remote/DR site. This scenario assumes that:

- ▶ Backups may still run at the primary/local site (as this is just a DR test).
- ▶ The remote site has different/separate backup servers from the local site.
- ▶ Some or all of the cartridges might be part of a visibility-switch-enabled policy.

Backups can be running that the local primary site while recovery of the remote site is occurring from replicated cartridges. After the remote site is remote, site replication will resume between the two sites.

Table 11-2 gives an overview of this use case.

Table 11-2 NetBackup DR test simulation in two domain backup environment

NBU server/media 1	NBU server/media 2
Local site	Remote site
ProtectTIER 1 (local site).	ProtectTIER 2 (remote/DR site).
Lib A.	Lib A'.
	Lib B physical library.
Prerequisites	
Create repository.	Create repository.
Create libraries.	Create libraries.
Install the ProtectTIER Replication Manager SW module (on the local or remote ProtectTIER node) by completing the following steps: 1. Create Grid A. 2. Add a repository to Grid A. 3. Pair local and remote repositories.	

NBU server/media 1	NBU server/media 2
<b>Local site</b>	<b>Remote site</b>
Create a replication policy to select the cartridges for replication and the specific remote library for the visibility switch.	
<b>DR test use case</b>	
Run regular backup activity to Lib A (to cartridges included in the replication policy).	
	Enter DR mode using the designated ProtecTIER Manager wizard.
	Recover the backup application from the replicated catalog/database (either from the catalog backup cartridge or from other means).
	Rescan the backup application to <i>learn</i> the library dimensions.
Backups can continue running to Lib A while the system is in DR mode. However, no data will be replicated to the remote until exiting DR mode. As a result, the new backup data to be replicated is accumulated as a replication backlog/queue.	Move the required cartridges from the repository shelf to the required libraries (through ProtecTIER Manager).
	Use commands to import cartridges that are located in the import/export slots. This imports all available cartridges in these slots into the designated library.
	Remember that all cartridges in the library are in read-only mode.

NBU server/media 1	NBU server/media 2
Local site	Remote site
	<p>Restore any required data from these cartridges. To understand which catalog/DB image contains cartridges that completed replication to the remote site, complete the following steps:</p> <ol style="list-style-type: none"> <li>1. Find the required backup image for restore.</li> <li>2. Get a list of cartridges included in this image using the backup application reporting capabilities.</li> <li>3. Get the times that this image started and completed (that is, the backup start and completion time stamp).</li> <li>4. Depending on the number of cartridges, the user can take one of the following actions: <ol style="list-style-type: none"> <li>a. If there is a small number of cartridges: The user can inquire through the ProtecTIER Manager cartridge view for the last sync time of the cartridges from step 2.</li> <li>b. If there is a large number of cartridges: The user should run the cartridge status report through ProtecTIER Manager. For more details, refer to Chapter 6, "Host implementation for virtual tape libraries" on page 275.</li> </ol> </li> <li>5. If all cartridges report sync time later than the image time, all cartridges are in sync with the catalog image and it can be used for restore operations.</li> <li>6. If some cartridge's sync time is earlier than the image start backup time, it means that these cartridges had pending data that was not replicated, but was required by the catalog image, and therefore this image cannot be used for restore and a previous complete image should be used instead.</li> <li>7. Scan the catalog/DB for a previous image of the same data and perform steps 5 - 6 again.</li> </ol>
	<p>Before exiting DR mode, to move cartridges to their original location prior to entering the DR mode, the user must eject all cartridges in the library using the backup application GUI.</p> <ul style="list-style-type: none"> <li>▶ Cartridges that were on the shel" prior to entering the DR test mode return to the remote repository shelf.</li> <li>▶ Cartridges that were in library A' following a visibility switch operation move to the remote repository shelf and appear in library A import/export slots.</li> </ul>
	Exit DR mode.
Run a command in the backup application to rescan the import/export slots to import cartridges that were moved back to the library as part of the eject operation in lib A'.	
All pending replication data/cartridges resume replicating to the remote site automatically.	

## Using the visibility switch: Performing a clone-to-tape and recovery operation in a multiple (two) domain backup environment

This use case describes the option to use the visibility switch option from site 1 (primary/local) to site 2 (remote/DR) when each site has its own backup application server, each with its own catalog/database. The user should clone the replicated cartridges at the remote site into a physical copy for longer retention/recovery purposes. This situation allows for two options for the recovery operation at the remote site:

- ▶ From the replicated cartridges
- ▶ From the cloned cartridges (if the longer retention period is required or if both the local and the remote sites are in DR mode and the only available repository copy is the cloned one)

Because each of the sites has its own backup application catalog/DB, every recovery attempt at the remote site must be done with the respective catalog/DB replica that includes the complete set of replicated cartridges.

Table 11-3 gives an overview of this use case.

Table 11-3 NetBackup visibility switch option in two domain backup environment

NBU Server/media 1	NBU Server/media 2
<b>Local site</b>	<b>Remote site</b>
ProtecTIER 1 (local site).	ProtecTIER 2 (remote/DR site).
Lib A.	Lib A.
	Lib B Physical library.
<b>Prerequisites</b>	
Create repository.	Create repository.
Create libraries.	Create libraries.
Install the ProtecTIER Replication Manager SW module (on the local or remote ProtecTIER node) by completing the following steps: <ul style="list-style-type: none"> <li>▶ Create Grid A.</li> <li>▶ Add a repository to Grid A.</li> <li>▶ Pair local and remote repositories.</li> </ul>	
Create a replication policy to select the cartridges for replication and the specific remote library for the visibility switch.	
<b>DR use case (with cloning to physical tape)</b>	
Run regular backup activity to Lib A (to cartridges included in the replication policy).	
Create a vault policy related to Lib A to manage the export/import of the cartridges from and to this library (that is, vault_policy_local).	
Run a vault policy (vault_policy_local) to eject all required cartridges from Lib A.	
All ejected cartridges move from lib A to the repository 1 (local) shelf.	

NBU Server/media 1	NBU Server/media 2
Local site	Remote site
	Create a vault policy related to Lib A' to manage the export/import of the cartridges from and to this remote library (that is, vault_policy_Remote).
	Cartridges that were ejected from Lib A automatically move from repository 2 (remote) shelf to Lib A' import/export slots (pending the completion of data replication/sync for each cartridge).
	To move the cartridges into the Lib A' slots, run a command to import them: 'vltinject <vault_policy_Remote>' This command must be run in a loop until all respective cartridges are imported into the library.
	Recover the backup application with the replicated catalog (either located on one of the replicated cartridges or received at the remote site through other means).
	Use the Rescan option of the backup application to learn the dimensions and scan both libraries (Lib A' and Lib B), as the backup application catalog was rebuilt as part of the recovery operation.
	Restore any required data from these cartridges. To understand which catalog/DB image contains cartridges that completed replication to the remote site, complete the following steps: <ol style="list-style-type: none"> <li>1. Find the required backup image for restore.</li> <li>2. Get a list of cartridges included in this image using the backup application reporting capabilities.</li> <li>3. Get the times that this image started and completed (that is, the backup start and completion time stamp).</li> <li>4. Depending on the number of cartridges you may do one of the following: <ol style="list-style-type: none"> <li>a. If there is a small number of cartridges: The user can inquire through the ProtecTIER Manager cartridge view for the last sync time of the cartridges from step 2.</li> <li>b. If there is a large number of cartridges: The user should run the cartridge status report through ProtecTIER Manager. For more details, refer to Chapter 6, "Host implementation for virtual tape libraries" on page 275.</li> </ol> </li> <li>5. If all cartridges report sync time later than the image time, all cartridges are in sync with the catalog image and it can be used for restore operations.</li> <li>6. If a cartridge's sync time is earlier than the image start backup time, it means that these cartridges had pending data that was not replicated, but was required by the catalog image. Therefore, this image cannot be used for restore and a previous complete image should be used instead.</li> <li>7. Scan the catalog/DB for a previous image of the same data and perform steps 5 - 6 again.</li> </ol>

NBU Server/media 1	NBU Server/media 2
Local site	Remote site
	Duplicate cartridges from Lib A' to physical library (Lib B).
	After duplication/clone operation is complete, back up the catalog/database to cartridges located on Lib B (to be used during recovery from these cartridges).
	Eject all cloned physical cartridges from Lib B (including the catalog backup) and save them in a safe location for recovery purposes. The cloned (virtual) cartridges cannot be left in the library, as the next visibility switch iteration will run over the backup application catalog/database. Therefore, the cartridges used for cloning will be considered as scratch.
	After duplication/clone operation is completed, run the vault policy to eject all Lib A' required cartridges using vault_policy_remote.
	All ejected cartridges will move from Lib A' to the repository 2 (remote) shelf.
Cartridges ejected from Lib A' will move from the repository 1 (local) shelf to Lib A import/export slots.	
To move the cartridges into the library, the user must issue a command to import them: 'vltinject <vault_policy_local>' This script must run in a loop until all respective cartridges are imported into the library.	
After cartridges are <i>imported</i> , they can be used for new backups.	
	Perform a recovery at the remote site from duplicate/cloned (physical) cartridges.
	Every cloned <i>box</i> consists of two types of cartridges: ▶ The backup application catalog consistent with this complete set of cartridges. ▶ The cloned cartridges.
	For every box of cartridges that requires recovery, complete the following steps: 1. Import all cartridges from the box into library B. 2. Recover the backup application from the catalog/database located on one of the cartridges. 3. Rescan the library dimensions. 4. Restore cartridges from Lib B to the backup server. 5. After the restore operation completes, eject all required cartridges from Lib B to the box. 6. Continue with the next box.
	Perform full/selective recovery at the remote site from replicated cartridges.
	Recover the backup application from the latest complete catalog/database located on one of the virtual cartridges in Lib A'.

<b>NBU Server/media 1</b>	<b>NBU Server/media 2</b>
<b>Local site</b>	<b>Remote site</b>
	<p>Restore cartridges from Lib A' to the NBU backup server:</p> <ul style="list-style-type: none"> <li>▶ If selective restore is required, scan the catalogue/DB for the cartridge containing the exact file.</li> <li>▶ If full recovery is required, restore all the required cartridges.</li> </ul>







# Monitoring and reporting of the IBM System Storage TS7600 with ProtecTIER

The ProtecTIER Manager application is the universal monitoring tool for the ProtecTIER software installed on the IBM System Storage TS7600 with ProtecTIER systems. In this chapter, we describe what to look for in the ProtecTIER Manager user interface and what actions to take when problems are encountered. We also describe other means of monitoring the system.

We cover the following topics:

- ▶ Managing ProtecTIER Manager user accounts
- ▶ Monitoring the ProtecTIER software
- ▶ Monitoring the ProtecTIER VTL service
- ▶ Reporting on ProtecTIER activity
- ▶ Monitoring replication policies and activities
- ▶ ProtecTIER Replication Network Performance Validation Utility

## 12.1 ProtecTIER Manager user management

The ProtecTIER Manager has three default user accounts that come with predefined default passwords (Table 12-1).

Table 12-1 Default user names and passwords in ProtecTIER Manager

Permission level	Default user name	Default password
Administrator	ptadmin	ptadmin
Operator	ptoper	ptoper
Monitor	ptuser	ptuser

Change the passwords or replace and delete these default user accounts to secure the system against unauthorized access.

The three permission levels have different capabilities within ProtecTIER Manager. After you are logged in, depending on your user ID's permission level, only certain options from menus and toolbar buttons are available to you. These capabilities are indicated by the color of the option's text. If the option is available to you, the text appears in black and is selectable. If it is not available to you, the text appears in blue and is not selectable. For example, in Figure 12-1, from the Repository menu, only the View Resources option is available to this user account.

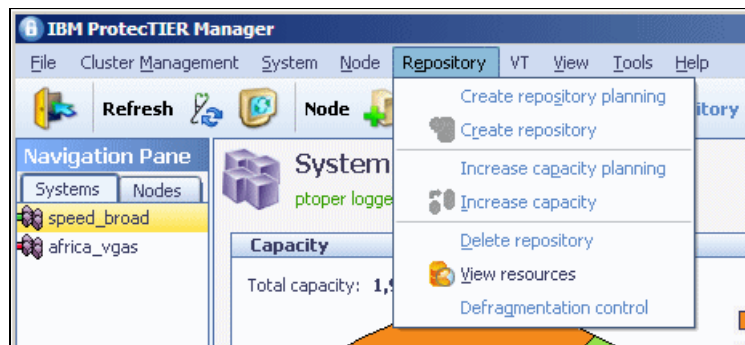


Figure 12-1 Example of available and unavailable options for a user account in ProtecTIER Manager

The ProtecTIER Manager application supports the permission levels shown in Table 12-2.

Table 12-2 Permission levels

Permission level	Authority
Administrator	Has full access to ProtecTIER Manager.
Operator	Has restricted access to ProtecTIER Manager: <ul style="list-style-type: none"> <li>▶ Toggles cartridges between read/write and read-only modes. For more information, refer to 10.3.5, "Switching cartridges to read-only mode" on page 507.</li> <li>▶ Sets the HyperFactor mode for libraries. For more information, refer to 10.8.5, "Changing the HyperFactor mode" on page 544.</li> <li>▶ Resets virtual tape drives and robots. For more information, refer to 10.8.8, "Resetting robot" on page 547.</li> <li>▶ Unloads and changes the visibility of cartridges from virtual tape drives. For more information, refer to 10.8.9, "Unloading and moving cartridges" on page 548.</li> </ul>

Permission level	Authority
Monitor	Can only access ProtecTIER Manager monitoring windows.

To make changes to ProtecTIER Manager user accounts, after you have successfully logged in to ProtecTIER Manager as an administrator, select **System** → **User Management** (Figure 12-2).




Figure 12-2 The User Management option on the System menu in ProtecTIER Manager

The window shown in Figure 12-3 opens. From here, you may add, remove, or change user accounts.

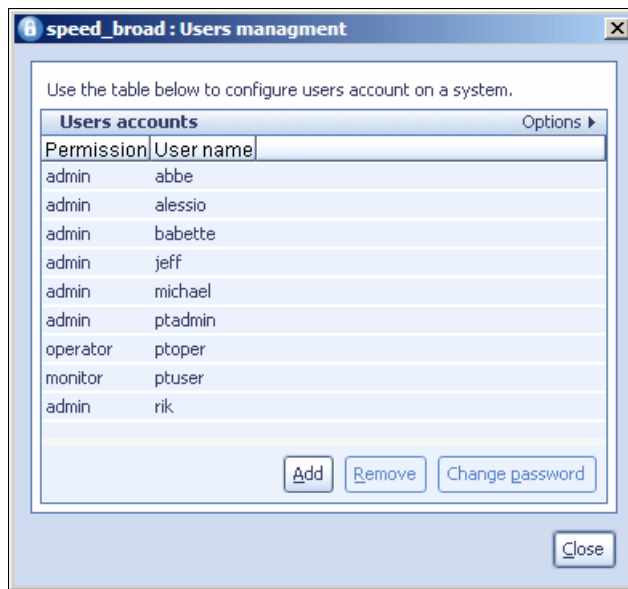


Figure 12-3 User Management window in ProtecTIER Manager

From here you can choose an action and, if required, an existing user ID on which to perform the operation. As an example, click **Add** to create a new user account. The window shown in Figure 12-4 opens, where you must specify a user name and password. You must also select a permission level from the drop-down menu. Click **Close** to close the window.

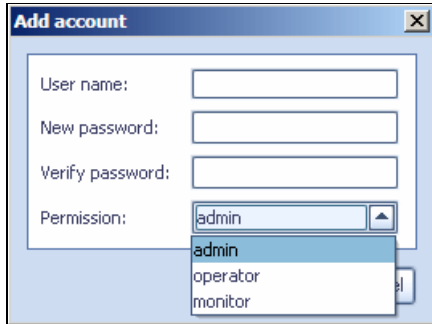


Figure 12-4 Add account window in ProtectTIER Manager

**Note:** Only one administrator can be logged into a ProtectTIER two-node cluster at a time. You should log out at the end of each session by clicking **Logout**.

There is no restriction on the number of operator or monitor user accounts that can be logged in concurrently.

If you log in with administrator level permission while another administrator is already logged in, a message box opens, as shown in Figure 12-5.

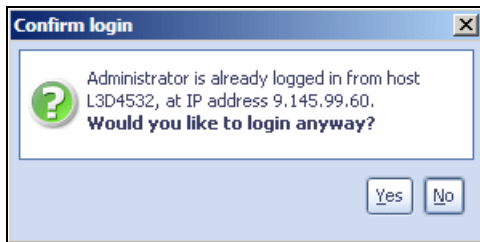


Figure 12-5 Confirm login window in ProtectTIER Manager

Click **Yes** to force the other administrator to log out or **No** to cancel your administrator login attempt.

## 12.2 Monitoring ProtecTIER

In this section, we describe the information about the ProtecTIER systems that is available through the ProtecTIER Manager application. Figure 12-6 shows the first window that you will see after starting the ProtecTIER Manager application on your workstation. We describe the different parts of this window in detail in the following sections.

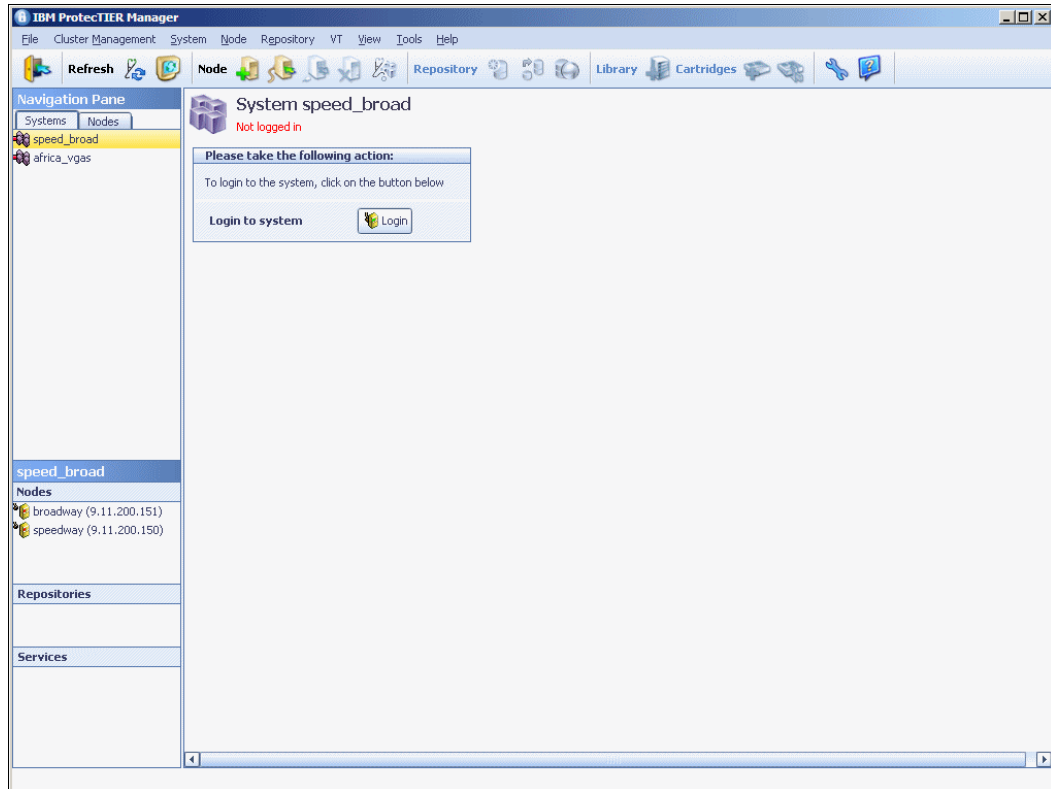


Figure 12-6 Initial login window in ProtecTIER Manager

Click **Login** to begin your ProtecTIER Manager session. A dialog box that asks for your user ID and password to authenticate the session opens (Figure 12-7).




Figure 12-7 User ID and password window in ProtecTIER Manager

Enter your user ID and password. If desired, you can have ProtecTIER Manager save your user ID and password for future logins by checking the **Save password** check box. After your user ID and password have been successfully authenticated, the Systems window opens (Figure 12-13 on page 628).

## 12.2.1 The status line

At the bottom of every window in ProtecTIER Manager there is a brief system status displayed on a single line (Figure 12-8.)




Figure 12-8 Status line in ProtecTIER Manager

It lists the following information, from left to right:

- ▶ The current state of the virtual tape task, as indicated by text and color:
  - Online: Green, static for available
  - Offline: Red, blinking for unavailable
- ▶ The user ID name that you have logged in with, the system you that are on, and the user authority level that you are using.
- ▶ The current read and write throughput rates in MBps.
- ▶ The current system time for the ProtecTIER (supplied from the TS7650G or TS7650 server running Linux).
- ▶ The Alerts and Events Log buttons.

### The Alerts button

The Alerts button turns red and blinks if there are any alert or error messages to be viewed. Click **Alerts** to view these messages in the Alerts Log window (Figure 12-9.)

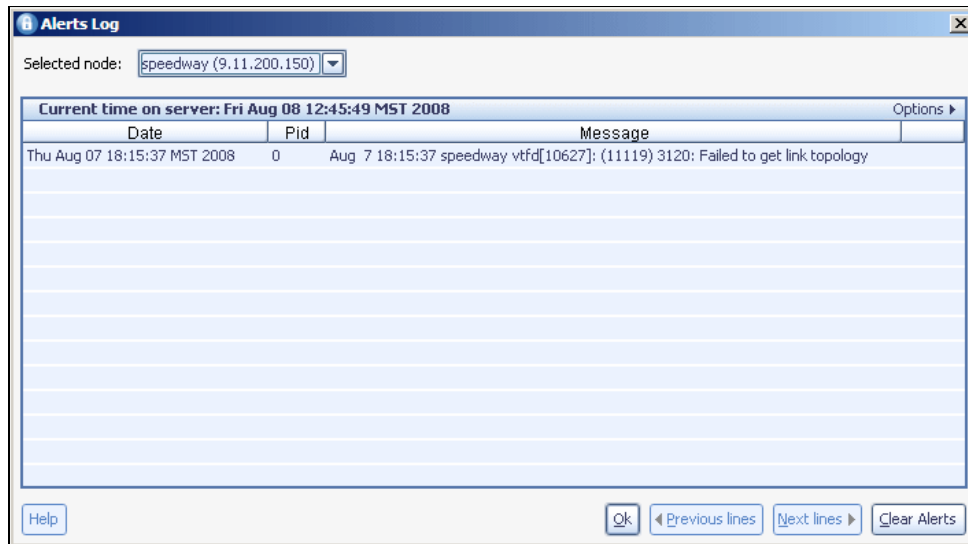


Figure 12-9 Alerts window in ProtecTIER Manager

After you have finished reading the messages, you can click **Clear Alerts** to acknowledge and remove the messages from the view. These messages still remain in the ProtecTIER Manager events log, along with all information messages. If you choose to exit without clearing them, click **OK** to close the window. When you return to the main ProtecTIER Manager window, the Alerts button will have stopped blinking and will be greyed out.

### The Events Log button

Click **Events Log** to view the log of all actions that you have performed during the current session, displayed in a window (Figure 12-10.) Click **OK** to close the window.

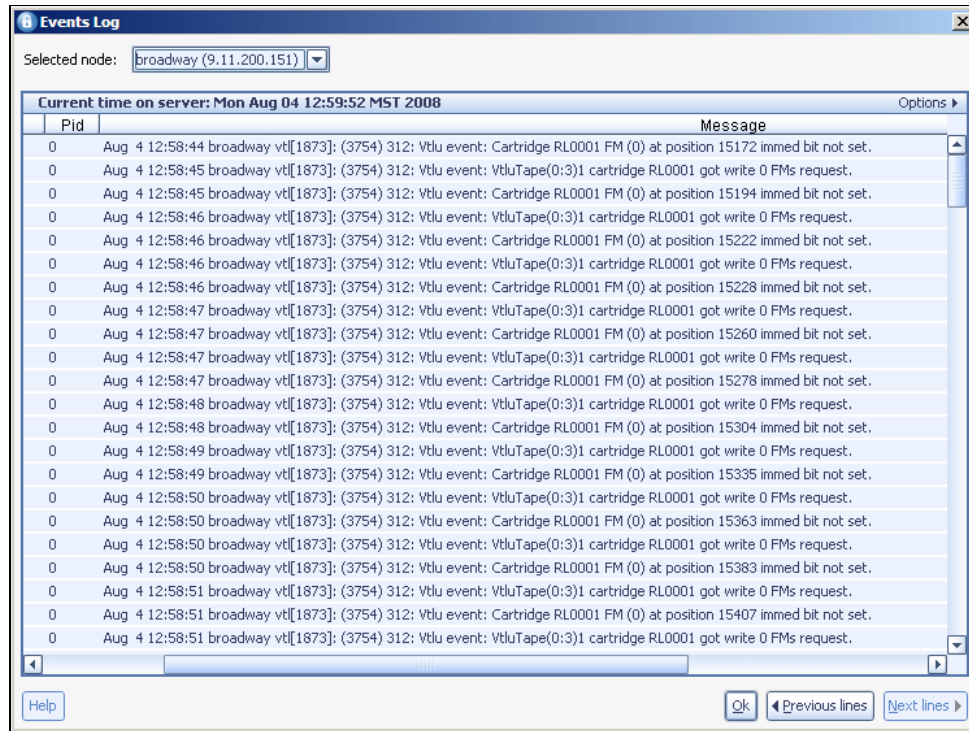


Figure 12-10 Events Log window in ProtecTIER Manager

Both the Alerts Log and the Events Log windows only display information for one node in a system at a time. The node whose messages are currently being displayed is shown in the Selected Node field. To change to another node in the same system, click the down arrow in the Selected Node field to open a drop-down menu, where you can select a new node by name.

In addition, the Alerts Log and Events Log windows display only up to 200 messages per page. You can navigate to other pages of messages by clicking **Previous Lines** or **Next Lines**.

## 12.2.2 The Navigation pane

The Navigation pane is located on the left side of the ProtecTIER Manager window (Figure 12-11) and is constant across all windows. It runs the entire height of the window and contains navigational features that can be used to move quickly to any desired view.




Figure 12-11 Nodes pane: General tab, Systems window in ProtecTIER Manager

At the top of the Navigation pane are two tabs, Systems and Nodes. These lead to the Systems window (see 12.2.3, “The Systems window” on page 628) and the Nodes window (Figure 12.2.4 on page 633).

### The Nodes pane

The Nodes pane is on the left side of the window below the Navigation Pane (Figure 12-11). Its heading is the name of the system that you are currently viewing (the system is called `speed_broad` in this example.) The Nodes pane contains the names and IP addresses of nodes that are defined to the system.



## The Repositories and Services panes

The Repositories pane and the Services pane are both on the left side of the window below the Nodes pane (Figure 12-12.) The Repositories pane contains the repository that is defined to the system that you are currently viewing. (There is only ever one repository per system.) This is explained in more detail in 12.2.5, “The Repository window” on page 637.



Figure 12-12 Repositories and Services panes in ProtecTIER Manager

The Services pane contains the names of virtual tape libraries (VTLs) defined to the system that you are currently viewing. (There can be just one or several, up to a maximum of 16 VTLs per system.) These libraries are explained in more detail in 12.3.1, “The Library window” on page 642.

## 12.2.3 The Systems window

When you first log in to ProtecTIER Manager, you see the General tab of the Systems window (Figure 12-13).

**Note:** The example figures from here on display a two-node clustered system.

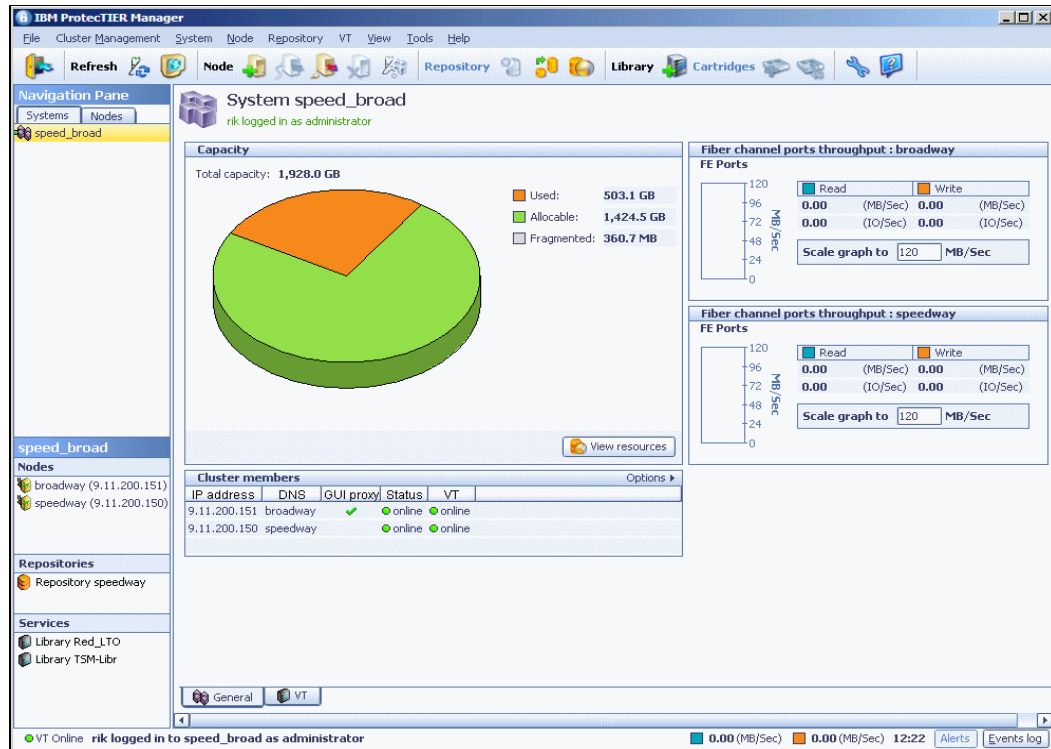


Figure 12-13 General tab: Systems window in ProtecTIER Manager

### The General tab

In the General tab of the Systems window, the name of the system that you logged in to is automatically selected in the Systems tab of the Navigation pane on the left side. It is also displayed along the top line under the menu and tool bars.

Other systems that ProtecTIER Manager detected are also displayed in this tab. On the left side of the icon for each system is a small dot that is colored green if you are logged in to that system and colored red if you are not.

You can change to another system by clicking a different system name. You must log in to each system separately.

When you select the **Nodes** tab in the Navigation pane, ProtecTIER Manager displays all the nodes of which it is aware (Figure 12-21 on page 633). Again, there is a small part of the icon for each node that is colored red or green depending on whether you are logged into that system.

The General tab of the Systems window has a short summary of the system as a whole. We provide more detailed information about each component later in this chapter.

The important points to note on the General tab of the Systems window are:

- ▶ The repository summary
- ▶ The Fibre Channel port status
- ▶ The cluster member and virtual tape library status

### The Capacity pane

The Capacity pane (Figure 12-14) is a graphical summary of the current state of the repository of the system that you are currently viewing. It is exactly the same as the view in the Repository window and is explained in more detail in “The Capacity pane” on page 637.

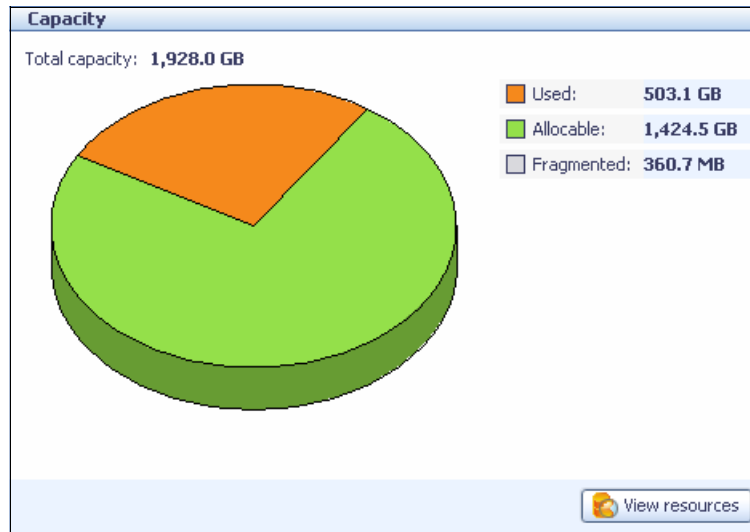


Figure 12-14 Capacity pane: General tab, Systems window in ProtectIER Manager

### The Fibre Channel Port Throughput pane

The Fibre Channel Port Throughput pane (Figure 12-15) displays the current Fibre Channel port throughput (both read and write) for each node. It is exactly the same as the view in the Nodes window and is explained in more detail in “The Fibre Channel Ports Throughput pane” on page 636.

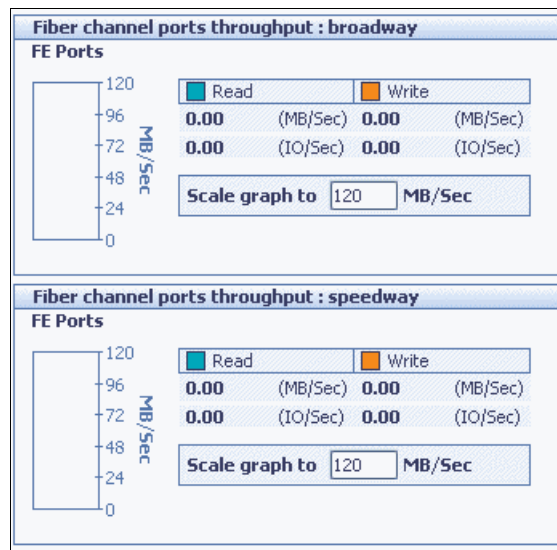


Figure 12-15 Fibre Channel Port Throughput pane: General tab, Systems window in ProtectIER Manager

## The Cluster Members pane

The Cluster Members pane (Figure 12-16) displays the Internet Protocol (IP) address and Domain Name Server (DNS) names for each node in the system that you are currently viewing. The GUI proxy column indicates which node is currently being used by ProtecTIER Manager to monitor the two-node cluster. The Status column indicates whether the ProtecTIER Software is online or offline on the node and the VT column indicates whether the VT service for that node is online or offline.

Cluster members						Options ▾
IP address	DNS	GUI proxy	Status	VT		
9.11.200.151	broadway	✓	● online	● online		
9.11.200.150	speedway		● online	● online		

Figure 12-16 Cluster Members pane: General tab, Systems window in ProtecTIER Manager

## The VT tab

In the Systems window, there is a second tab that can be selected for viewing, labelled VT. When you click the **VT** tab next to the General tab in the bottom pane of the window, the window shown in Figure 12-17 appears.

The screenshot shows the IBM ProtecTIER Manager interface. The main window displays the 'VT' tab for the system 'speed\_broad'. The 'Library front-end' table lists nodes and their associated drives and libraries. The 'Libraries performance' panel shows read and write speeds. The 'Libraries configuration' panel shows settings for drives, slots, and cartridges.

Node	Port	Top	WWN	LUN	Device	Library	Throughput (MB/Sec)
speedway 0	FC-AL		10000000c971d90a				
				0	Drive 0	Red_I TO	
				1	Drive 2	Red_I TO	
				2	Drive 4	Red_I TO	
				3	Drive 6	Red_I TO	
				4	Drive 8	Red_I TO	
				5	Drive 10	Red_I TO	
				6	Drive 12	Red_I TO	
				7	Drive 14	Red_I TO	
				8	Robot	TSM-Libr	
				9	Drive 0	TSM-Libr	
				10	Drive 2	TSM-Libr	
				11	Drive 4	TSM-Libr	
				12	Drive 6	TSM-Libr	
broadway 0	FC-AL		10000000c971d798				
				0	Robot	Red_I TO	
				1	Drive 1	Red_I TO	
				2	Drive 3	Red_I TO	
				3	Drive 5	Red_I TO	
				4	Drive 7	Red_I TO	
				5	Drive 9	Red_I TO	
				6	Drive 11	Red_I TO	
				7	Drive 13	Red_I TO	
				8	Drive 15	Red_I TO	
				9	Drive 1	TSM-Libr	
				10	Drive 3	TSM-Libr	
				11	Drive 5	TSM-Libr	
				12	Drive 7	TSM-Libr	

Figure 12-17 VT tab, Systems window in ProtecTIER Manager

In the VT tab, Systems window, the name of the system that you logged in to is automatically displayed on the Systems tab of the Navigation pane on the left side. It is also displayed along the top line under the menu and tool bars. You can change to another system by clicking a different system name in the Navigation pane at any time.

The important points to note in the VT tab of the Systems window are:

- ▶ All defined tape devices for this system
- ▶ Total current front-end Fibre Channel port throughput rates
- ▶ VTL configuration totals

**The Library front-end pane**

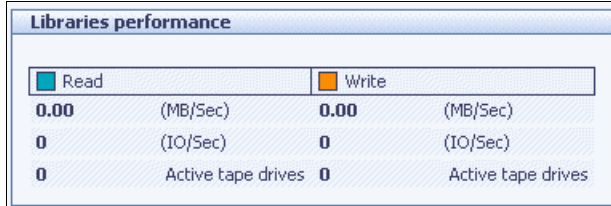
The Library front-end pane (Figure 12-18) lists all the virtual tape devices (robots and drives) across all the VTLs defined to the system that you are currently viewing, sorted in the order of node, front-end port number, LUN, device, and library name. It is the same as the view in the General tab of the Library window (except that is across ALL VTLs). The columns are explained in more detail in “The Library front-end pane” on page 643.

Library front-end							Options ▶
Node	Port	Top.	WWN	LUN	Device	Library	Throughput (MB/Sec)
speedway 0 FC-AL 10000000c971d90a							
				0	Drive 0	Red_LTO	
				1	Drive 2	Red_LTO	
				2	Drive 4	Red_LTO	
				3	Drive 6	Red_LTO	
				4	Drive 8	Red_LTO	
				5	Drive 10	Red_LTO	
				6	Drive 12	Red_LTO	
				7	Drive 14	Red_LTO	
				8	Robot	TSM-Libr	
				9	Drive 0	TSM-Libr	
				10	Drive 2	TSM-Libr	
				11	Drive 4	TSM-Libr	
				12	Drive 6	TSM-Libr	
broadway 0 FC-AL 10000000c971d798							
				0	Robot	Red_LTO	
				1	Drive 1	Red_LTO	
				2	Drive 3	Red_LTO	
				3	Drive 5	Red_LTO	
				4	Drive 7	Red_LTO	
				5	Drive 9	Red_LTO	
				6	Drive 11	Red_LTO	
				7	Drive 13	Red_LTO	
				8	Drive 15	Red_LTO	
				9	Drive 1	TSM-Libr	
				10	Drive 3	TSM-Libr	
				11	Drive 5	TSM-Libr	
				12	Drive 7	TSM-Libr	

Figure 12-18 Library front-end pane: VT tab, Systems window in ProtecTIER Manager

### The Libraries Performance pane

The Libraries Performance pane (Figure 12-19) shows the cumulative usage and performance figures for all the VTLs defined to the system that you are currently viewing. This display includes the read and write throughput rates in MBps and IOPS and how many total tape drives are active for each type of operation (read and write).

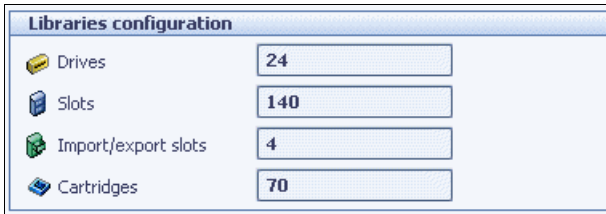


Libraries performance			
Read		Write	
0.00	(MB/Sec)	0.00	(MB/Sec)
0	(IO/Sec)	0	(IO/Sec)
0	Active tape drives	0	Active tape drives

Figure 12-19 Libraries Performance pane: VT tab, Systems window in ProtecTIER Manager

### The Libraries Configuration pane

The Libraries Configuration pane (Figure 12-20) shows the cumulative totals for library elements, such as drives, slots, I/E slots, and cartridges, across all the VTLs defined to the system that you are currently viewing.



Libraries configuration	
Drives	24
Slots	140
Import/export slots	4
Cartridges	70

Figure 12-20 Libraries Configuration pane: VT tab, Systems window in ProtecTIER Manager

## 12.2.4 The Nodes window

You can access the Nodes window by clicking the **Nodes** tab in the Navigation pane on the left side of the window. This window lists all the nodes currently defined to ProtecTIER Manager. Select the node that you want to display and click its name. If you are not yet logged in to that system, ProtecTIER Manager displays the node login window (Figure 12-6 on page 623.) If you are logged in already, the window shown in Figure 12-21 opens.

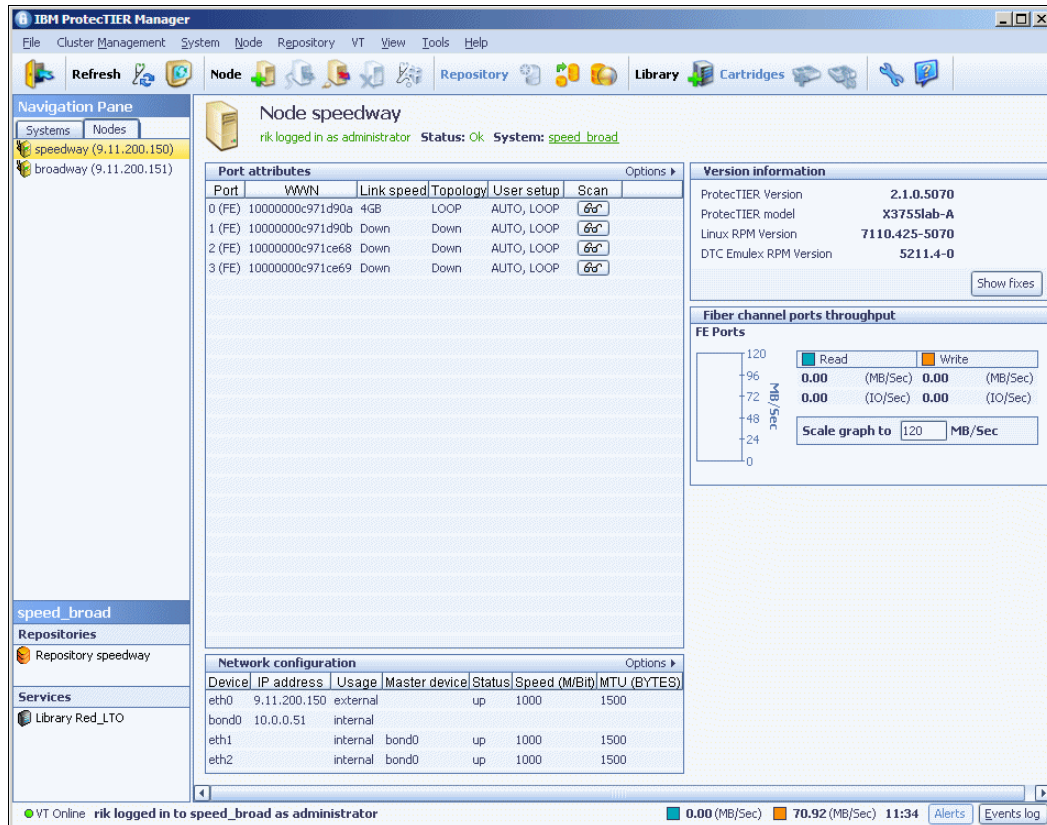


Figure 12-21 Nodes window in ProtecTIER Manager

In the Nodes window, the name of the node that you have logged in to is displayed along the top line under the menu and tool bars.

The important points to note in the Nodes window are:

- ▶ The port attributes for the four front-end (FE) ports
- ▶ The ProtecTIER version information
- ▶ The Fibre Channel ports throughput display
- ▶ The network configuration

## The Port Attributes pane

The current state of health for each of the four FE ports associated with the node is displayed in the Port Attributes pane. The worldwide name (WWN) for each port is displayed along with the link speed, topology (and state of the link), and the user setup (Figure 12-22).


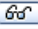

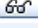
Port attributes						Options ▾
Port	WWN	Link speed	Topology	User setup	Scan	
0 (FE)	10000000c971d90a	4GB	LOOP	AUTO, LOOP		
1 (FE)	10000000c971d90b	Down	Down	AUTO, LOOP		
2 (FE)	10000000c971ce68	Down	Down	AUTO, LOOP		
3 (FE)	10000000c971ce69	Down	Down	AUTO, LOOP		

Figure 12-22 Port Attributes pane: Nodes window in ProtecTIER Manager

The Link speed column shows the transmission speed of the port. Possible values are:

- ▶ AUTO: A transmission speed that is auto-negotiated between the two ports depending on the combined highest possible link speed.
- ▶ 1 GB: A fixed transmission speed of 1 Gbps.
- ▶ 2 GB: A fixed transmission speed of 2 Gbps.
- ▶ 4 GB: A fixed transmission speed of 4 Gbps.
- ▶ DOWN: There is no Fibre Channel connection.

The Topology column displays the Fibre Channel topology of the port. Possible values are:

- ▶ LOOP: Fibre Channel Arbitrated Loop connection.
- ▶ P2P: Peer-to-peer connection.
- ▶ DOWN: There is no Fibre Channel connection.

The User setup column is the user-assigned link speed and topology. Possible values are a combination of the Link Speed and Topology column values above, separated by a comma.

There is also the Scan button (marked with a pair of glasses icon) displayed in the far right column for each port. Clicking this icon opens the Scan Port dialog box. Scanning the port displays a numbered list of the WWNs of the remote ports detected by the port. This is useful during the initial setup or when diagnosing problems.





## The Fibre Channel Ports Throughput pane

The Fibre Channel (FC) Ports Throughput pane displays the rate of data movement and I/O operations for both read and write operations for the node (Figure 12-26). The data movement rate is also displayed graphically for each front-end Fibre Channel port on the node. The bars will be colored blue (for read) or orange (for write). There is enough space for four bars to be displayed at once in the bar graph, one bar for each FC port.

You can change the scale of the graph by editing the value in the Scale Graph To field to see the throughput rates in finer or fewer details.

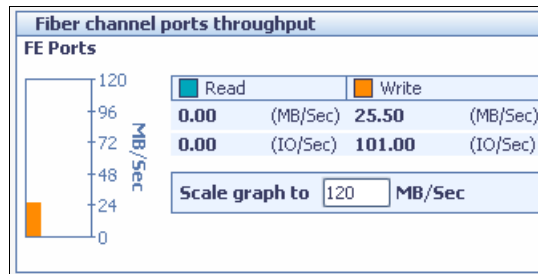


Figure 12-26 Fibre Channel Ports Throughput pane: Nodes window in ProtecTIER Manager

## The Network Configuration pane

The Network Configuration pane displays information about the setup of the network interface cards (NIC) for the node (Figure 12-27.)

Network configuration						Options ▾
Device	IP address	Usage	Master device	Status	Speed (M/Bit)	MTU (BYTES)
eth0	9.11.200.150	external		up	1000	1500
bond0	10.0.0.51	internal				
eth1		internal	bond0	up	1000	1500
eth2		internal	bond0	up	1000	1500

Figure 12-27 Network Configuration pane: Nodes window in ProtecTIER Manager

The column values are explained in Table 12-3.

Table 12-3 Column definitions for network configuration values in ProtecTIER Manager

Column	Definition
Device	The devices in the NIC: <ul style="list-style-type: none"> <li>▶ Eth0 is the node port that communicates with the ProtecTIER Manager workstation.</li> <li>▶ Eth1 and Eth2 are the node ports used in the two-node cluster-internal network.</li> <li>▶ Bond0 is the virtual bond master device to which Eth1 and Eth2 are enslaved.</li> </ul> Bond devices are defined as part of the installation process.
IP Address	IP address of the device.
Usage	Indicates whether the device is used for the two-node cluster-internal network or to communicate with the ProtecTIER Manager workstation and the external network.
Master Device	The master device or bond device, if any, to which the device is enslaved.
Status	Indicates whether the device is functioning properly.

Column	Definition
Speed	The supported speed of data transfer across the device in megabits per second.
MTU	Configured maximum transmission unit for the device.

## 12.2.5 The Repository window

You can access the Repository window by clicking a repository name in the Repositories pane on the left side of any window (Figure 12-28). The repository for a two-node cluster system with two nodes is named after the first node in the two-node cluster that created the repository.

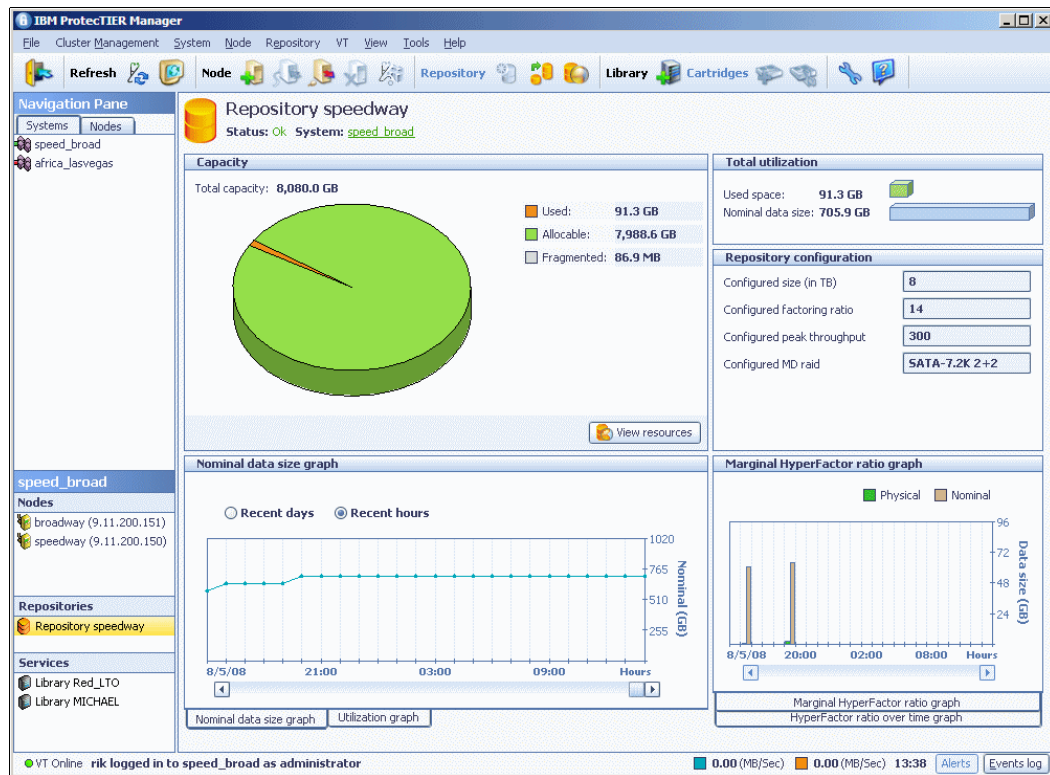


Figure 12-28 Repository window in ProtecTIER Manager

### The Capacity pane

The repository's *state of health* is displayed in a pie chart in the Capacity pane near the center of the window. This provides a visual indication of the current capacity status. It also shows three repository capacity totals:

- ▶ Used
- ▶ Allocable
- ▶ Fragmented

#### Used

This total represents the amount of space presently occupied in the repository *after* any data deduplication has taken place. It is the real consumption level of the space allocated to the repository.

### **Allocatable**

This total represents the amount of unoccupied space in the repository.

### **Fragmented**

This total represents the amount of data in the repository that is in pieces smaller than 1 MB. ProtecTIER runs housekeeping tasks in the background to constantly reassemble these small fragments into larger pieces of free space to make the best use of all available space.

**Repository Space Usage:** The Repository view represents the space used on all the virtual cartridges defined, regardless of the logical state (active or expired) of the data in the backup application that created the data. After a virtual cartridge has been filled, the amount of data that it contains is included in the *Used* portion of the pie chart in the repository view. There is no reconciliation between the backup application and the ProtecTIER repository to delete expired data from the repository. This emulates the process used for physical cartridges. When a backup application logically expires data, physically the data is still there on the cartridge media surface. Refer to 2.6.2, “Steady state” on page 34 for more information.

### **The View Resources button**

The View Resources button just below the pie chart graphic displays the layout of the repository at the Linux file system level in a window (Figure 12-29). It shows the mount point names, the file system type (always GFS, meaning global file system), the file system size, and the current usage type of the file system (either user data or metadata). Nothing can be changed from this window and it is intended for informational purposes only.

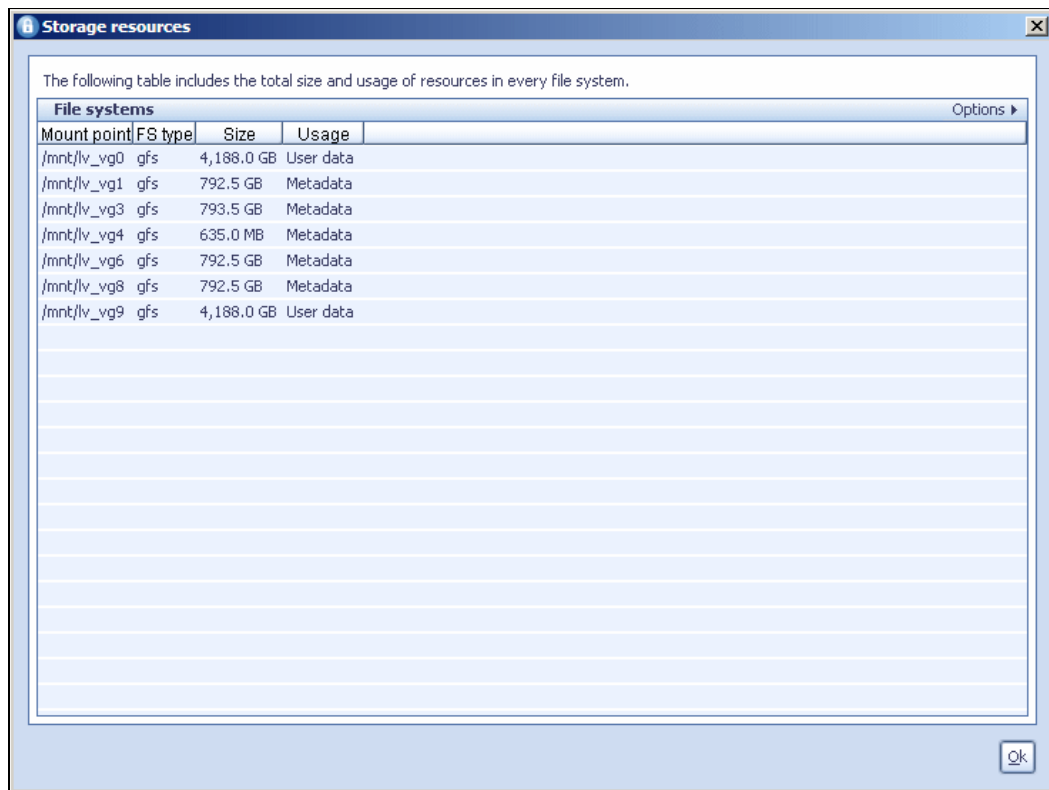


Figure 12-29 View resources window in ProtecTIER Manager

## The Total utilization pane

The Total Utilization pane of the Repository window is an at-a-glance view of the nominal data received compared to the physical data stored, with totals (Figure 12-30).




Figure 12-30 Total utilization pane: Repository window in ProtecTIER Manager

## The Repository configuration pane

The Repository configuration pane of the Repository window lists important repository information in a single place (Figure 12-31).

Figure 12-31 Repository configuration pane: Repository window in ProtecTIER Manager

It contains the information shown in Table 12-4.

Table 12-4 Field definitions for repository configuration values in ProtecTIER Manager

Field	Definition
Configured size (in TB)	The physical repository size in terabytes (user data only).
Configured factoring ratio	The estimated HyperFactor factoring ratio that was used to create the repository.
Configured peak throughput	The expected maximum peak throughput specified when the repository was created.
Configured MD raid	The disk technology (SATA or FC), disk speed (revolutions per minute), and RAID 10 group for the metadata LUNs.

## The ProtecTIER Manager data graph pane

Depending on your screen's resolution size, the following graphs appear in tabs or are displayed directly, one above the other. This section describes the tabbed display.

Two tabs showing two different graphical views of the data sent to and stored by the ProtecTIER are displayed beneath the Capacity pane in the Repository window. The graphs are called:

- ▶ The Nominal Data Size graph
- ▶ The Utilization graph

### The Nominal Data Size graph

By clicking the **Nominal Data Size graph** tab, you can display how much data has been sent to the ProtecTIER systems over time (Figure 12-32).




Figure 12-32 Nominal Data Size graph tab: Repository window in ProtecTIER Manager

This graph represents the amount of data written to the ProtecTIER systems by the backup application without any compression or deduplication. You can display data figures from recent hours or recent days using the radio buttons in the top left corner of the pane. There is also a scroll bar to view earlier historical data, if available.

### The Utilization graph

By clicking the **Utilization graph** tab, you can display how much actual data has been stored by the ProtecTIER systems over time (Figure 12-33).




Figure 12-33 Utilization graph tab: Repository window in ProtecTIER Manager

This graph represents the amount of data written to the disk repository by the ProtecTIER system after deduplication and compression has taken place. It is likely to be significantly less than the nominal data amount. You can display data figures from recent hours or recent days using the radio buttons in the top left corner of the pane. There is also a scroll bar to view earlier historical data, if available.

### The HyperFactor graph pane

Depending on your screen's resolution size, the following graphs will appear in tabs or be displayed directly, one above the other. The following sections describe the tabbed display.

Two tabs showing two different graphical views of the ProtecTIER HyperFactor algorithm at work are displayed beneath the Repository configuration pane in the Repository window. The graphs are called:

- ▶ The Marginal HyperFactor ratio graph
- ▶ The HyperFactor ratio over time graph

**The Marginal HyperFactor ratio graph**

By clicking the **Marginal HyperFactor Ratio graph** tab, you can see how well the ProtecTIER is performing (Figure 12-34).




Figure 12-34 Marginal HyperFactor ratio tab: Repository window in ProtecTIER Manager

This pane provides a visual comparison through a bar graph of data written by the ProtecTIER (physical, green) versus data received by the ProtecTIER (nominal, grey) over the previous few hours, summarized for each hour. It is possible to immediately see whether the data sent recently has factored well.

**The HyperFactor ratio over time graph**

By clicking the **HyperFactor Ratio Over Time graph** tab, you can see the recent performance of ProtecTIER (Figure 12-35).




Figure 12-35 HyperFactor ratio over time tab: Repository window in ProtecTIER Manager

This pane shows the deduplication ratio through a line graph. The ratio is a measure of the amount of data received by the ProtecTIER system from the backup application to the amount of data written by the ProtecTIER system out to the disk repository over the previous few hours, summarized for each hour.

## 12.3 Monitoring the ProtecTIER virtual tape libraries service

In this section, we describe the monitoring of the ProtecTIER virtual tape libraries.

### 12.3.1 The Library window

You can access the Library window by clicking a library name in the Services pane on the left side of any window (Figure 12-36). The General tab of the Library window appears.

All of the tabs in the Library window display a status line across the top of the page beneath the menu and tool bars that shows the name of the selected library, its status, the system it is defined to, and the repository to which it is attached.

All of the tabs except the General tab have an area beneath this status line containing page navigation options for when multiple pages are displayed. On this tab, you can also sort the display according to any of the columns by clicking the column heading. Note though that this only sorts the data contained on the current page, not the entire data set.

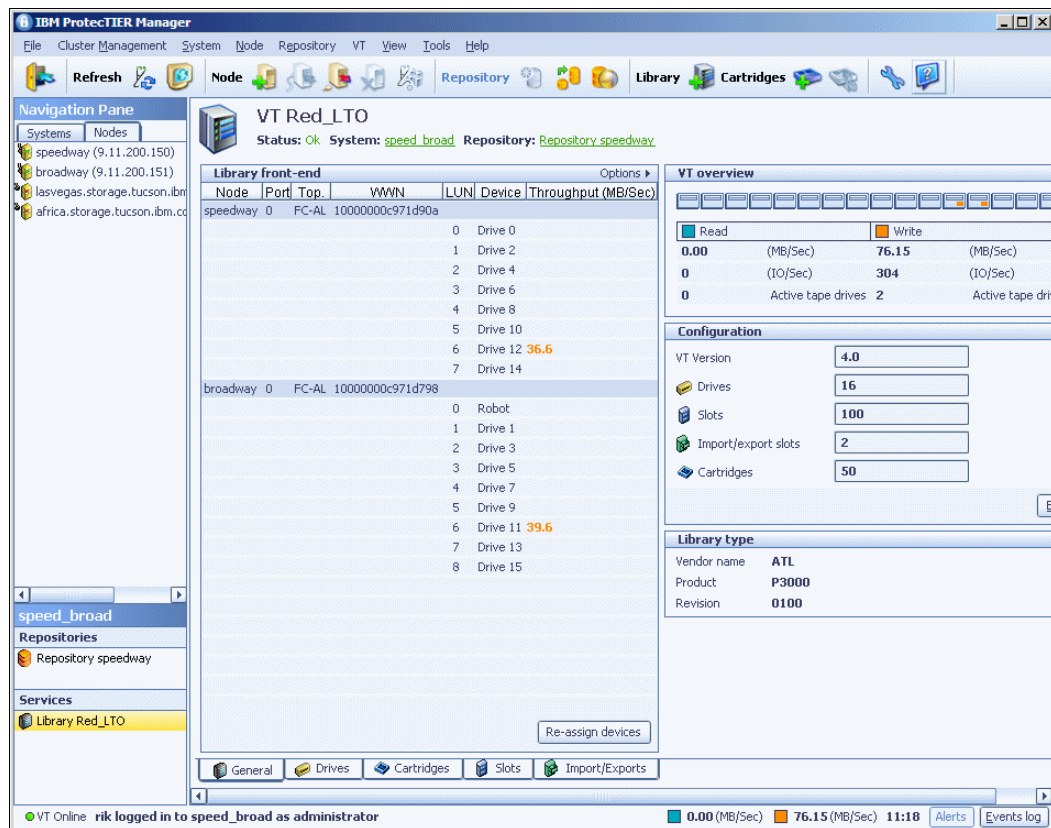


Figure 12-36 General tab: Library window in ProtecTIER Manager

There are several tabs along the bottom of the Library window. These tabs are discussed in detail in the next few sections.



## 12.3.2 The General tab

The General tab gives a summary view of the selected library (Figure 12-36 on page 642). There are four panes in this window.

### The Library front-end pane

The Library front-end pane shows the virtual robot and tape drives for the selected library, sorted by node, then by the Linux LUN number (Figure 12-37).

Node	Port	Top	WWN	LUN	Device	Throughput (MB/Sec)
speedway 0	FC-AL		10000000c971d90a	0	Drive 0	
				1	Drive 2	
				2	Drive 4	
				3	Drive 6	
				4	Drive 8	
				5	Drive 10	
				6	Drive 12	36.6
				7	Drive 14	
broadway 0	FC-AL		10000000c971d798	0	Robot	
				1	Drive 1	
				2	Drive 3	
				3	Drive 5	
				4	Drive 7	
				5	Drive 9	
				6	Drive 11	39.6
				7	Drive 13	
				8	Drive 15	

Figure 12-37 Library front-end pane, General tab, Library window in ProtecTIER Manager

There are several columns of information. Some columns apply to the node in general and others apply to individual drives. They are explained in Table 12-5.

Table 12-5 Column definitions for Library front-end pane: General tab, Library window in ProtecTIER Manager

Column	Definition
Node	The node on which the virtual device is assigned.
Port	The port within the node on which the virtual device is assigned.
Top	The Fibre Channel topology of the port. Possible values are: <ul style="list-style-type: none"> <li>▶ P2P (point-to-point)</li> <li>▶ FC-AL (Fibre Channel-arbitrated loop)</li> <li>▶ DOWN (There is no Fibre Channel connection.)</li> </ul>
WWN	The World Wide Name of the port.

Column	Definition
LUN	The logical unit number of the robot or tape drive relative to the port.
Device	The name of the robot or tape drive.
Throughput	The rate of data transfer across the device.

If a drive is currently active, the Throughput (MBps) column shows a value in one of two colors:

- ▶ Orange for a WRITE operation
- ▶ Blue for a READ operation

This pane also contains the Reassign Devices button. You can use this button to change the assignments of the devices to the front-end ports across both nodes, as explained in detail in 10.3.2, “Reassigning devices” on page 494.

### The VT overview pane

The VT overview pane shows a summary of current drive activity and updates dynamically. The drives are represented graphically and display an orange or blue icon if they are being used (the colors are the same as above). Hovering your cursor over the graphic displays the drive number, the current read/write rate of the drive in MBps, and the percentage of time that the tape drive is idle during backup operations due to low backup application data transfer rates (Figure 12-38).

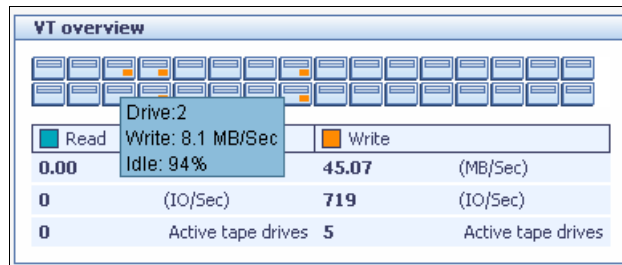
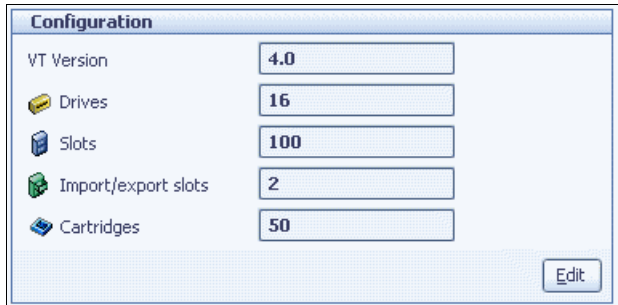


Figure 12-38 VT overview pane: General tab, Library window in ProtecTIER Manager

Summary totals are listed for both read and write operations and for total active tape drives in the selected library.

## The Configuration pane

The Configuration pane (Figure 12-39) displays the current number of drives, slots, import/export slots, and cartridges in the selected library.



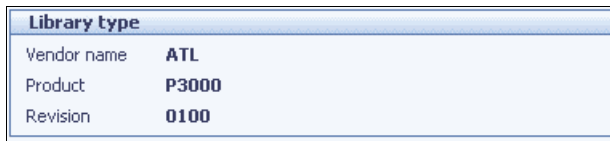
Configuration	
VT Version	4.0
Drives	16
Slots	100
Import/export slots	2
Cartridges	50
<input type="button" value="Edit"/>	

Figure 12-39 Configuration pane: General tab, Library window in ProtecTIER Manager

The pane also contains the Edit button. You can use this button to change the dimensions of the selected library, as explained in detail in 10.3.1, “Editing library parameters” on page 487.

## The Library type pane

The Library type pane (Figure 12-40) displays the definition for the emulated virtual library type, including the vendor name, product name, and revision number.



Library type	
Vendor name	ATL
Product	P3000
Revision	0100

Figure 12-40 Library type pane: General tab, Library window in ProtecTIER Manager

### 12.3.3 The Drives tab

The Drives tab displays detailed information about the virtual tape drives in the selected library (Figure 12-41).

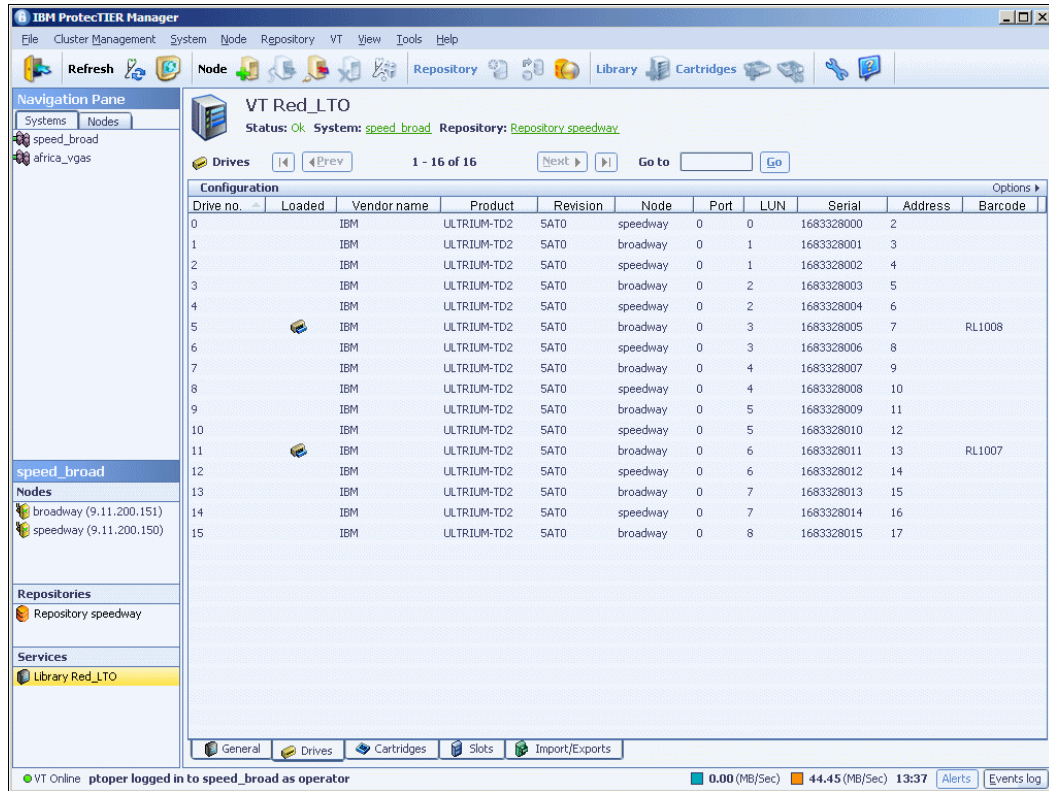


Figure 12-41 Drives tab: Library window in ProtecTIER Manager

The column values are explained in Table 12-6.

Table 12-6 Column definitions for Drives tab: Library window in ProtecTIER Manager

Column	Definition
Drive No.	The drive number in ProtecTIER.
Loaded	Whether the drive is loaded with a virtual cartridge. If the drive is loaded with a cartridge, an icon is displayed.
Vendor Name	The vendor whose product the virtual drive emulates.
Product	The product name for the product that the virtual drive emulates.
Revision	The revision number for the product that the virtual drive emulates.
Node	The node to which the drive is assigned.
Port	The port on the node to which the drive is assigned.
LUN	The drive's logical unit number relative to the port. Note that this is for ProtecTIER only. The backup server's LUN numbers for the drives might be different.
Serial	The drive's serial number.

Column	Definition
Address	The drive's address within the library.
Barcode	If the drive is loaded with a cartridge, this column displays the cartridge's barcode.

### 12.3.4 The Cartridges tab

The Cartridges tab displays detailed information about the cartridges in the selected library (Figure 12-42).

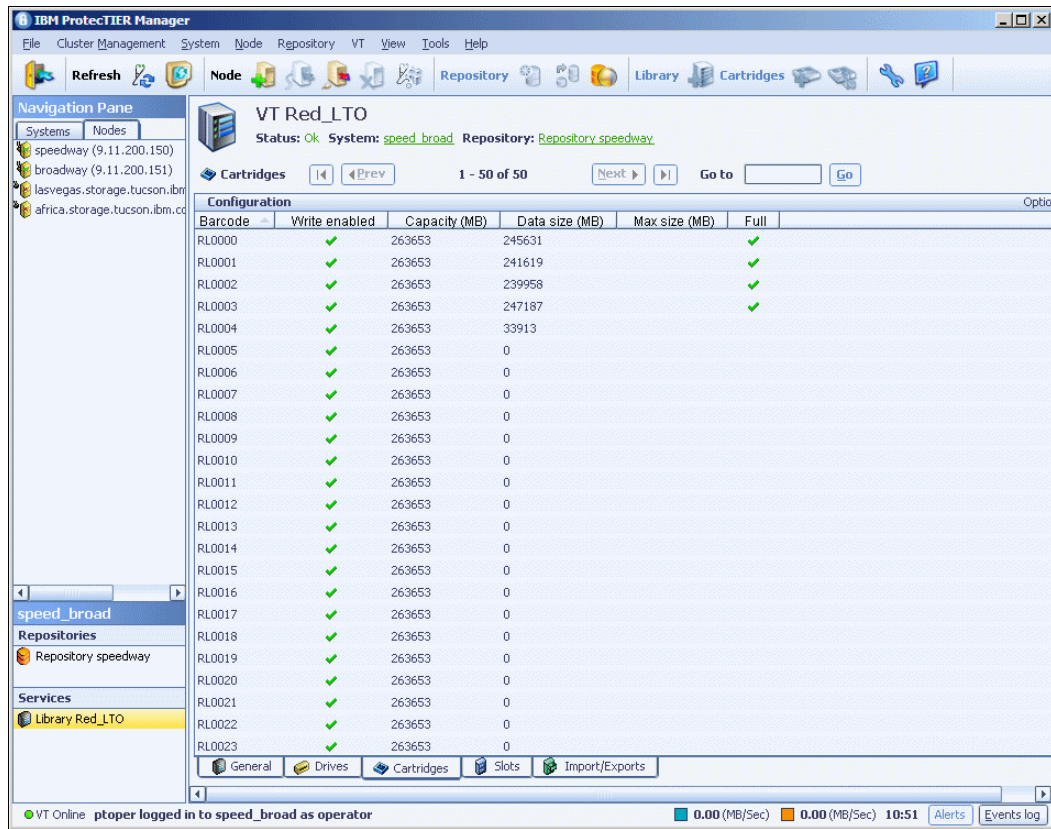


Figure 12-42 Cartridges tab: Library window in ProtecTIER Manager

The column values are explained in Table 12-7.

Table 12-7 Column definitions for Cartridges tab: Library window in ProtecTIER Manager

Column	Definition
Barcode	The cartridge's barcode.
Write enabled	Whether the cartridge is write enabled. If the cartridge is write enabled, a green tick icon is displayed. Otherwise, a red cross icon is displayed.
Capacity (MB)	The cartridge's estimated data capacity in megabytes. This value varies over time depending on the HyperFactor ratio and the number of cartridges configured in the system.
Data size (MB)	The amount of nominal data, in megabytes, currently stored on the cartridge.

Column	Definition
Max size (MB)	The maximum (fixed) amount of nominal data at which the ProtecTIER will announce early warning for this cartridge to the backup application.
Full	Indicates whether the cartridge has reached the early warning threshold. If so, the cartridge is regarded as full and a green tick icon is displayed.

### 12.3.5 The Slots tab

The Slots tab displays detailed information about the slots in the selected library (Figure 12-43), some of which is repeated from the Drives tab. The information includes slot number, element address, cartridge barcode, estimated capacity, and data size (amount of nominal data stored) if the slot contains a cartridge.

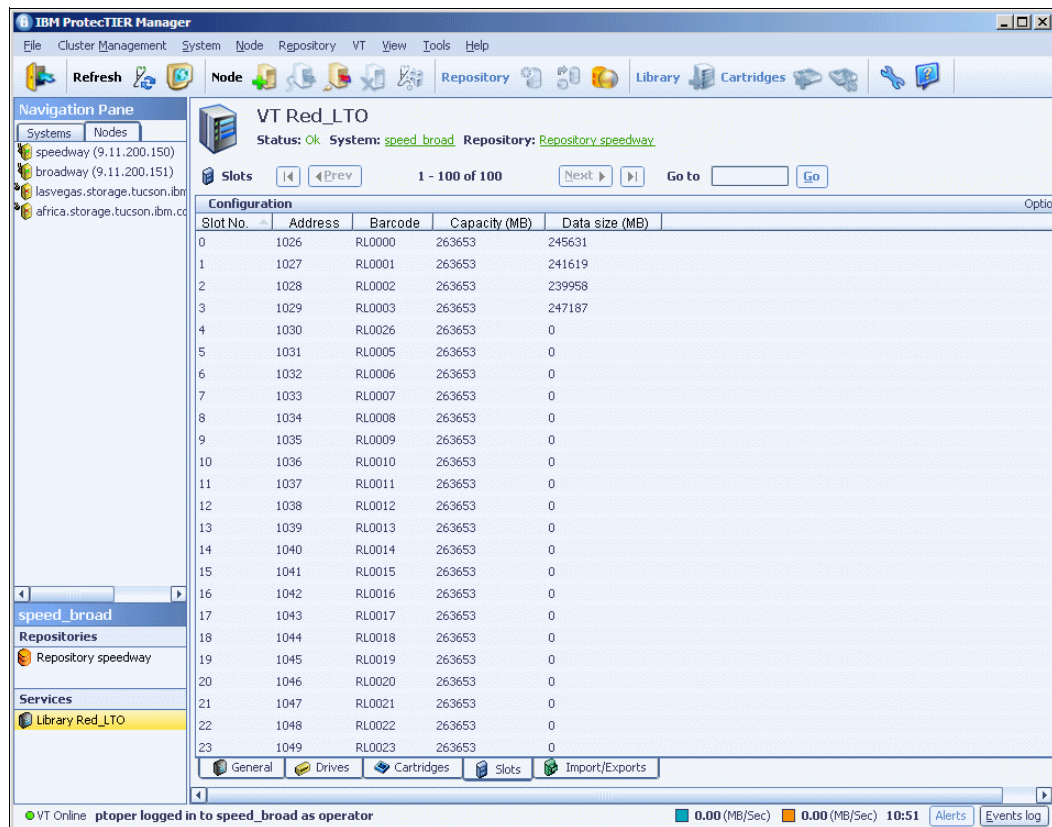


Figure 12-43 Slots tab: Library window in ProtecTIER Manager

## 12.3.6 The Import/Exports tab

The Import/Exports (I/E) tab displays detailed information about the import/export slots in the selected library (Figure 12-44), some of which is repeated from the Drives and Slots tabs. The information includes I/E slot number, element address, cartridge barcode, estimated capacity, and data size (amount of nominal data stored) if the slot contains a cartridge and shows what is the direction of the cartridges. Some cartridges are in the *in* direction and some cartridges are in the *out* direction.

Sporadically, cartridges appear at the I/O slot of the library. Therefore, it is a best practice that a script injects periodically the cartridges from the I/O slot into the library slots.

IBM ProtecTIER Manager

File Cluster Management System Node Repository VT Replication View Tools Help

Refresh Node Repository Library Cartridges

Systems Management

Systems Nodes

- beast\_woody
- star\_moon
- mir\_pluto
- lublin

beast\_woody

Nodes

- beast (168.159.150.126)

Repositories

- beast\_woody\_repository1

Services

- Library Lib1
- Shelf

VT Lib1  
Status: Ok System: beast\_woody Repository: beast\_woody\_repository1

Import/Export Slots 1 - 104 of 104

Slot no.	Address	Barcode	Capacity (MB)	Data size (MB)	Imp/Exp
0	64514	000009	1008	0	
1	64515	000010	1008	0	
2	64516	000011	1008	0	
3	64517	000012	1008	0	
4	64518	000013	1008	0	
5	64519	000014	1008	0	
6	64520	000015	1008	0	
7	64521	000016	1008	0	
8	64522	000017	1008	0	
9	64523	000018	1008	0	
10	64524	000308	1008	0	
11	64525				
12	64526				
13	64527	000292	1008	0	
14	64528	000293	1008	0	
15	64529	000294	1008	0	
16	64530	000295	1008	0	
17	64531	000296	1008	0	
18	64532	000297	1008	0	
19	64533	000298	1008	0	
20	64534	000299	1008	0	
21	64535	000300	1008	0	
22	64536	000301	1008	0	
23	64537	000302	1008	0	
24	64538	000303	1008	0	
25	64539	000304	1008	0	

Figure 12-44 Import/Exports tab: Library window in ProtecTIER Manager

## 12.3.7 Monitoring the shelf

Through ProtecTIER Manager, a view of the shelf can be displayed and you can access the cartridges in the shelf (Figure 12-45). From the shelf, cartridges can be relocated to a library's import slot, replicated to another repository, or deleted. Cartridges can also be automatically moved to a library's import slot through the visibility switching process. A cartridge whose visibility is moved to an export slot is automatically displaced to the shelf.

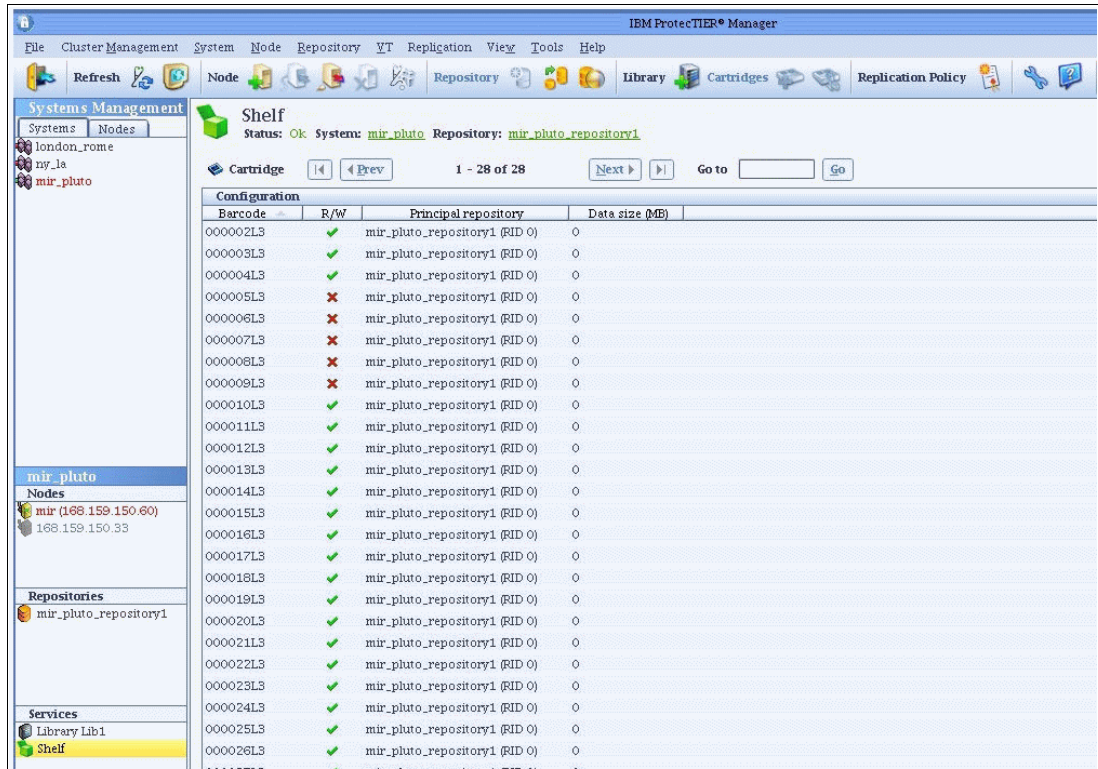


Figure 12-45 Shelf display

### Fields

The fields on this window are:

- ▶ Barcode: The barcode of the cartridge
- ▶ R/W: Displays if the cartridge can be modified
- ▶ Principal repository: On which repository the cartridge can be modified
- ▶ Data size (MB): The nominal size (in MB) of the cartridge

## 12.4 Reporting on ProtecTIER activity

In the section, we discuss the two types of reports that can be generated by the user for ProtecTIER activities. Both reports produce spreadsheet format files that can be viewed and analyzed.



## 12.4.1 The ANALYZE\_SESSIONS utility

You can gather and view ProtecTIER backup statistics to see how the system is performing. These statistics show throughput, compression, and HyperFactor rates for backup data sent to the system. The ProtecTIER provides a utility called *analyze\_sessions* to perform this data collection task. The *analyze\_sessions* utility analyzes and summarizes the contents of the current, cumulative compressor log. A compressor log is generated each time that the Linux *vtfd* daemon is started and accumulates entries until the *vtfd* daemon is shut down. Repeated runs of the utility between restarts will report mostly identical data each time, plus any new session data since the last execution of the utility. When the *analyze\_sessions* utility is run after a restart of the ProtecTIER software, it only reports on new session data since the restart.

The recommended use of the *analyze\_sessions* utility is for the deliberate testing of a specific workload or type of data. The best practice would be isolating a desired workload to be the only data being sent to the ProtecTIER for a certain period with a five-minute break on either side of it, to separate it in a single session from any other data. You would then run the *analyze\_sessions* utility and examine the output to determine how well the data might be suited to data deduplication processes.

### Running the utility

To run this command and gather the data, execute the following steps:

1. Log on to one of the TS7650 or TS7650G node as root.
2. Change to the `/opt/dtc/app/utis/` directory.
3. Run `./analyze_sessions`.

This creates an output file in comma-separated-values (CSV) format in the following directory:

```
/pt_work/<server name>--<date>--<time>.csv
```

Figure 12-46 gives an example of how to run the analyze\_sessions utility on a ProtecTIER node and what output will be returned.

```
login as: root
root@9.11.200.150's password:
Last login: Wed Aug 06 13:00:28 2008 from 10.0.0.52
[root@speedway ~]# cd /opt/dtc/app/utis/
[root@speedway utis]# ./analyze_sessions

analyze_sessions: this program analyzes the logs on the current
ProtecTIER server. Its output is the change rate between each
session and data that was available in the repository at the time.

This program and its output contain proprietary information of
International Business Machines Corporation.

(c) Copyright International Business Machines Corporation 2008. All rights
reserved.

Cutoff date: 2008-5-6 13:00:50
Processing file: 2/2
Read 1MB, throughput: 7.35MB/s
Output file is /pt_work/speedway-2008-08-06-13-00-50.csv
[root@speedway utis]#
```

*Figure 12-46 Running the analyze\_sessions utility on ProtecTIER Linux system*

## Examining the output

You can either view the generated CSV format file on the ProtecTIER system or you can use ftp commands or similar utilities to download the file to a Windows workstation to open it. When you open the CSV file in Microsoft Excel or an equivalent product, you will see data similar to that shown in Figure 12-47.

Name	Total data (TB)	Total data (KB)	System change rate	Factoring ratio	start time	end time
Grand totals						
all	0.74	791260000	18.67%	5.35738	29/07/2008 18:22	5/08/2008 19:44
By session (summary)						
2008-7-29 18:22:54 to 2008-7-29 18:22:54	0.00	826810	100%	1	29/07/2008 18:22	29/07/2008 18:22
2008-7-29 18:29:41 to 2008-7-29 18:34:14	0.00	489220	35.34%	2.82926	29/07/2008 18:29	29/07/2008 18:34
2008-8-4 11:56:26 to 2008-8-4 12:13:09	0.01	8388610	99.82%	1.00183	4/08/2008 11:56	4/08/2008 12:13
2008-8-4 12:26:03 to 2008-8-4 12:41:38	0.01	9437180	10.43%	9.58874	4/08/2008 12:26	4/08/2008 12:41
2008-8-4 12:52:20 to 2008-8-4 13:07:06	0.01	9259260	0.13%	760.404	4/08/2008 12:52	4/08/2008 13:07
2008-8-4 13:14:42 to 2008-8-4 13:53:30	0.02	16777200	38.56%	2.59367	4/08/2008 13:14	4/08/2008 13:53
2008-8-4 14:02:13 to 2008-8-4 14:19:06	0.04	40894500	100%	1	4/08/2008 14:02	4/08/2008 14:19
2008-8-4 14:39:15 to 2008-8-4 14:39:15	0.00	340736	100%	1	4/08/2008 14:39	4/08/2008 14:39
2008-8-4 15:14:05 to 2008-8-4 16:17:16	0.21	227044000	25.99%	3.84737	4/08/2008 15:14	4/08/2008 16:17
2008-8-4 16:22:24 to 2008-8-4 18:04:55	0.12	132392000	13.30%	7.51867	4/08/2008 16:22	4/08/2008 18:04

Figure 12-47 Example of the analyze\_sessions utility output

In ProtecTIER terms, a *session* is defined as a sustained period of data throughput (for example, data sent to the PT server from a backup application) preceded by and followed by an idle period of at least five minutes.

In the data, there is a Grand Totals row summarizing all the sessions in the output file, followed by detail rows for each individual session. There are several columns for each row, containing statistical information. See Table 12-8 for an explanation of the fields.

Table 12-8 Recent backup session statistics definitions

Statistic	Description
Total data (TB)	Total amount of data backed up during the session, in terabytes
Total data (KB)	Total amount of data backed up during the session, in kilobytes
System change rate	Percentage of data in the backup session recognized as changed relative to the previous backup session
Factoring ratio	The ratio of the quantity of actual backed-up data over the total amount of physical data
Start time	Start time of the backup session
End time	End time of the backup session

## 12.4.2 Long-term statistical data report

In ProtecTIER Manager, there is an option to create an extremely detailed statistical report in CSV format. This report can then be downloaded to your workstation. The report gathers and compiles a large amount of information from the ProtecTIER system and puts it into a user viewable format. The data can be difficult to understand in the proper context and is usually only generated when requested by IBM Support in response to a problem situation.

You should consider not running this report, as it is difficult to derive useful user information from it.

### Creating and downloading the report

To create and download the report, complete the following steps:

1. From the menu bar, select **Node** → **Create and download long-term statistics** (Figure 12-48).

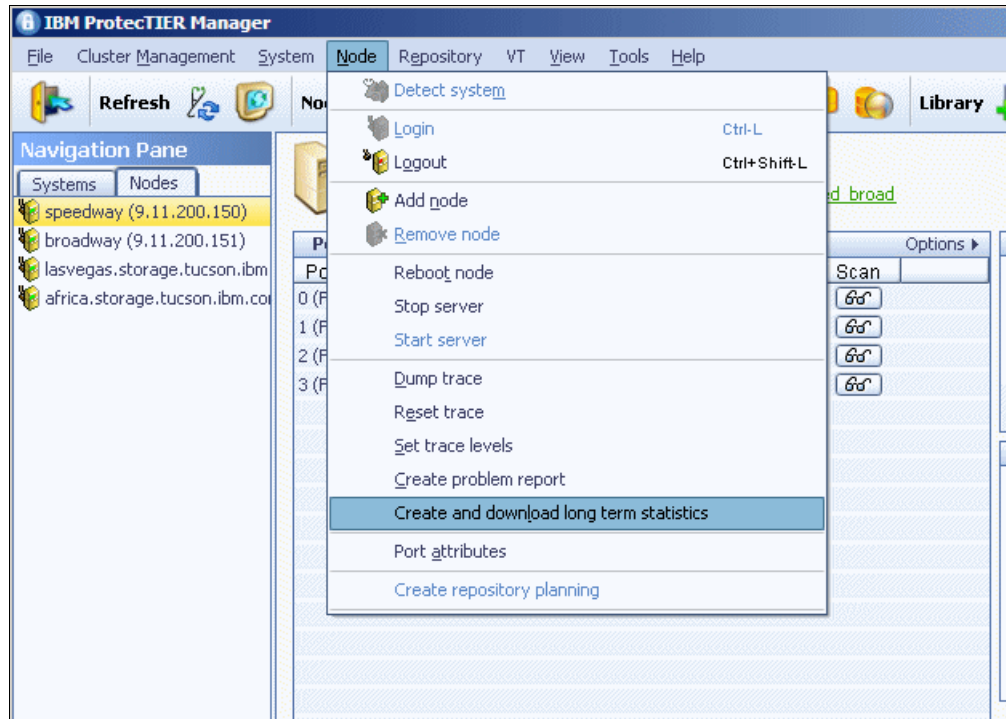


Figure 12-48 Creating long-term statistics report in ProtecTIER Manager

ProtectTIER Manager displays a message box while it creates the report on the ProtectTIER node (Figure 12-49).

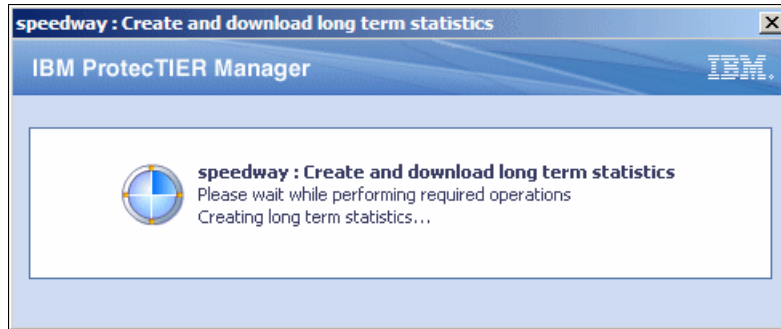


Figure 12-49 Message window for long-term statistics report creation in ProtecTIER Manager

ProtectTIER Manager displays a dialog box asking whether you want to download the completed report from the ProtectTIER node server to your workstation (Figure 12-50).

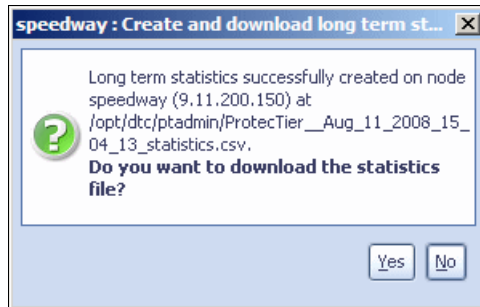


Figure 12-50 Download window for long-term statistics report in ProtecTIER Manager

If you click **No**, the task ends and you are returned to ProtecTIER Manager. If you want to use the file at a later time, you must retrieve it from the ProtectTIER node through manual methods.

If you click **Yes**, ProtecTIER Manager displays a dialog box asking you to specify to which location on your workstation you want to save the statistics file (Figure 12-51).

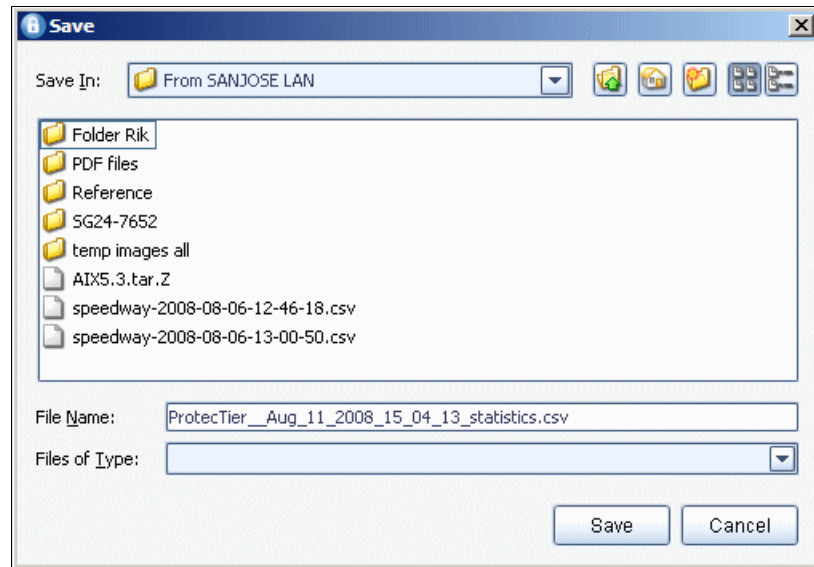


Figure 12-51 Save window in ProtecTIER Manager

2. Make your selection and click **Save**. If you do not want to save the report at this time, click **Cancel**. If the report is saved successfully, ProtecTIER Manager displays a message box confirming this action (Figure 12-52).

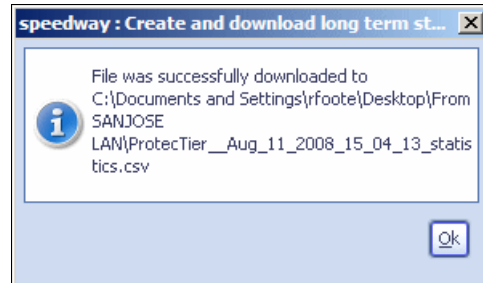


Figure 12-52 Download confirmation message window in ProtecTIER Manager

3. Click **OK** to continue.

The file is now ready to open using a spreadsheet program (for example, Lotus Symphony) to view the detailed data. It is also ready to forward to IBM Support at their request to assist with diagnosis in problem situations.

## 12.5 Monitoring replication policies and activities

You can monitor the replication policies and activities of your repository through the Systems Management view. Click the repository appearing in the left navigation pane. At the bottom left of the work pane, you can now find the Replication Policy and activities tabs, as shown in Figure 12-53.

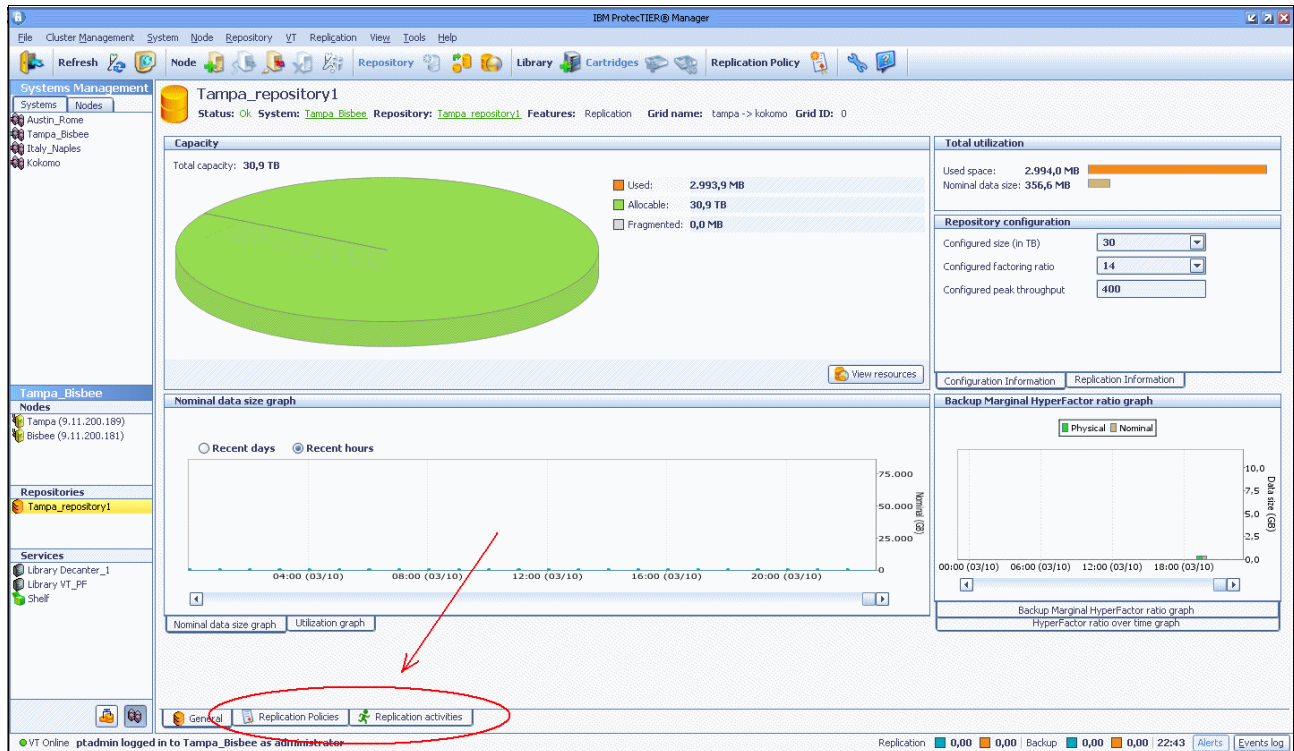


Figure 12-53 Replication Policies and Replication activities tabs

## 12.5.1 Replication Policies window

Click the **Replication Policies** tab at the bottom of the work pane to see the replication policies that are defined on the repository (Figure 12-54). Select a policy from the list to view the policy's details.

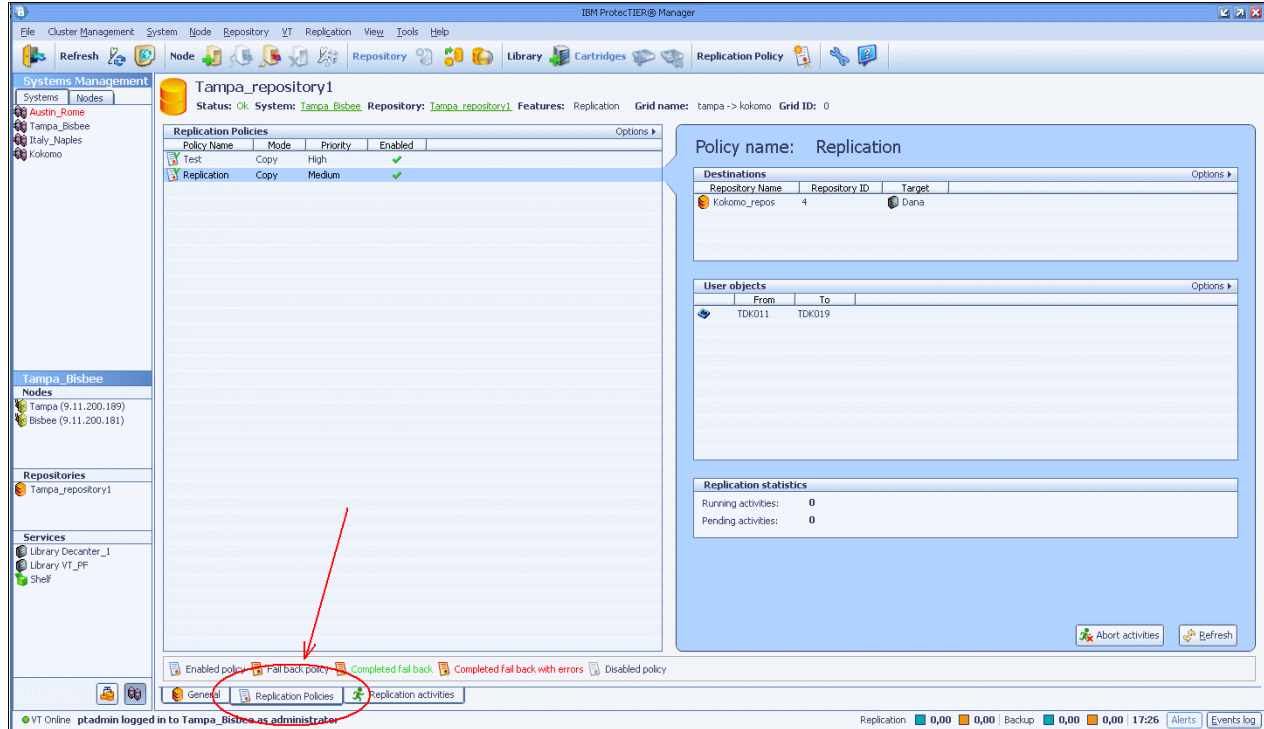


Figure 12-54 Replication policies tab

When you select a policy, you can see the following policy details as defined by the Create policy wizard (Figure 12-55):

- ▶ Policy name: The name that you define for the policy.
- ▶ Mode: Cartridges are either copied or moved.
- ▶ Priority: The policy's level of importance:
  - High
  - Normal
  - Low
- ▶ Enabled: The policy can be enabled or disabled for execution.

**Note:** In this release, *copy* is the only mode of operation.

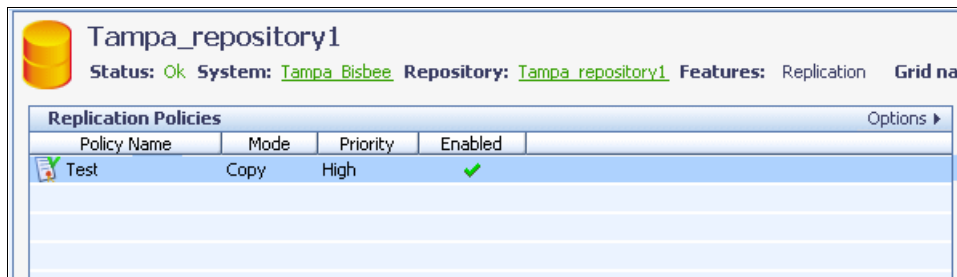


Figure 12-55 Policy details



In addition, on the right side of the Replication Policies view, you can find additional helpful information related to the selected policy. The window on the right, as shown in Figure 12-56, contains the following information:

- ▶ Policy Name.
- ▶ Destination Repository Name.
- ▶ Destination Repository ID.
- ▶ Target: The replication target for the cartridges on the destination side. This is the shelf or the name of the VTL.
- ▶ User objects: The range of cartridges chosen to replicate with that policy.
- ▶ Running activities: The number of currently running replication activities per policy. One activity is always the replication of one cartridge, which means that the maximum number of running activities per policy can never be higher than the number of cartridges in the policy.
- ▶ Pending activities: The number of currently pending replication activities. One activity is also always the pending replication of one cartridge. This number can never be higher as the number of cartridges in the policy.

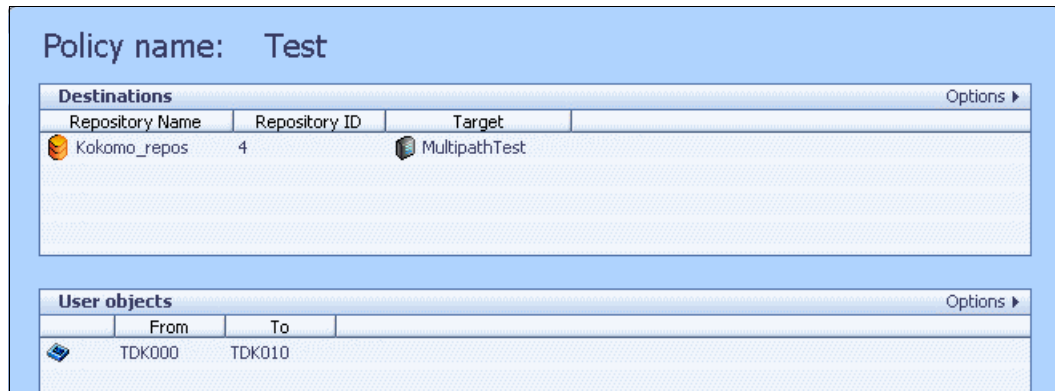


Figure 12-56 Additional policy information window

## 12.5.2 Replication Activities window

The ProtecTIER GUI has a dedicated display showing running replication activities. Click the **Replication activities** tab at the bottom of the work pane, as shown in Figure 12-57. The list of cartridges undergoing replication activity is displayed. The activities have information about the current replication state and about the deduplication ratio. The activities window also shows system-wide backlog and ETA for the entire replication and the number of pending activities. Pending activities cannot be monitored.

The screenshot shows the IBM ProtecTIER Manager interface. The main window is titled "mir\_pluto\_repository1" and displays a table of replication activities. The table has the following columns: Object ID, Policy, Throughput (Nom.), Dedup, Priority, Progress, and Time left. The data rows are as follows:

Object ID	Policy	Throughput (Nom.)	Dedup	Priority	Progress	Time left
A00009L3	Policy name 1	46 KB/Sec	5:1	High	[Progress bar]	0d 0h 2m
A00011	Policy DEDUP	43 KB/Sec	5:1	Normal	[Progress bar]	0d 0h 12m
B00010	Policy ABC	50 KB/Sec	5:1	Normal	[Progress bar]	0d 0h 3m
B00510	Policy ABC	42 KB/Sec	5:1	Normal	[Progress bar]	0d 0h 3m
C00070	Policy ABC	45 KB/Sec	5:1	Normal	[Progress bar]	0d 0h 3m
D00010	Policy ABC	45 KB/Sec	5:1	Normal	[Progress bar]	0d 0h 3m

Below the table, there are buttons for "Select all", "Select none", and "Abort activity". To the right of the table, there is a panel for "Object ID: A00009L3" with "Additional activity information" including Location (Library Lib1), Activity direction (Outgoing), Destination (NA), Replicated (Nominal) (4,896.0 MB), Start time (12:13:14, 08/09/2009), Time elapsed (0d 0h 9m), and Number of retries (0). At the bottom of the window, the "Overall statistics" section shows:

- Pending activities: 16
- Running activities: 3
- Physical In: 0 MB/Sec
- Physical Out: 271 MB/Sec
- Nominal In: 0 MB/Sec
- Nominal Out: 1,359 MB/Sec
- Backlog: 1,176.00 MB
- Replication time left: 0d 1h 12m

The "Replication activities" tab at the bottom of the work pane is highlighted with a red circle and a red arrow pointing to it.

Figure 12-57 Select Replication activities pane

**Note:** The Replication activities view is *ad hoc*. As such, if new data is written to a cartridge while it is being replicated, the progress of the displayed data details will appear as though it has been restarted.

When you select a replication activity, you can see the following activity details (Figure 12-58):

- ▶ Object ID: The barcode of the cartridge
- ▶ Policy: The name of the policy that contains the respective cartridge in its definition
- ▶ Throughput: The nominal throughput of replicating data in MBps
- ▶ Dedup: Displays the savings in the bandwidth needed for replicating this cartridge
- ▶ Priority: The replication policy's priority in the queue
- ▶ Progress: The progress of replicating the latest changes to their destination
- ▶ Time left: The amount of time remaining for the replication process to complete

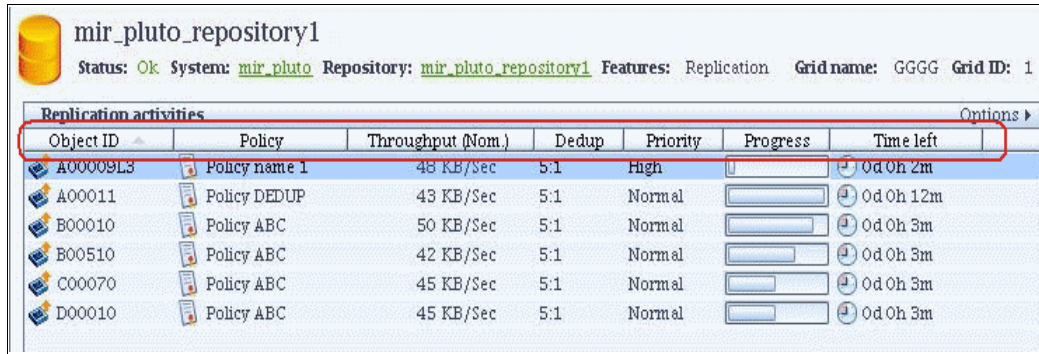


Figure 12-58 Replication activities details

In addition, on the right side of the replication activities view is additional helpful information related to the selected activity. The window on the right, as shown in the Figure 12-59, contains the following additional information:

- ▶ Object ID: The barcode of the cartridge
- ▶ Location: The cartridge source and name of the library
- ▶ Activity direction: The replication direction
- ▶ Destination: Name of the destination repository (This information is not available with that release. It shows NA (not applicable).)
- ▶ Replicated (Nominal): The nominal data replicated
- ▶ Start time: The replication start time
- ▶ Time elapsed: The already elapsed replication time
- ▶ Number of retries: Counts the replication retries

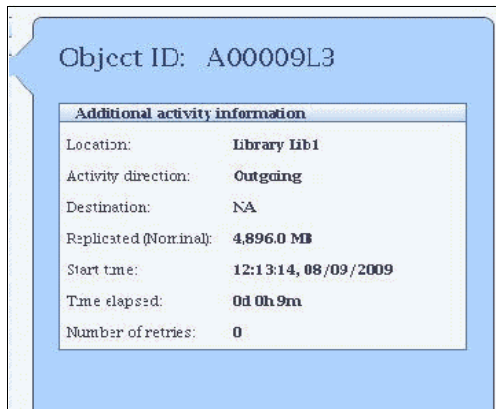


Figure 12-59 Additional activity information window

## 12.5.3 Cartridge replication status information

In the cartridge view (Figure 12-61 on page 663), a user can see, among the other attributes of a cartridge, a few replication-related attributes. A single or a small set of cartridges can be chosen for querying their attributes (like regular cart view). This query actually polls the most current information of the cartridge. After selecting the cartridge from the list, a blue replication information window, belonging to that cartridge, will open at the right side.

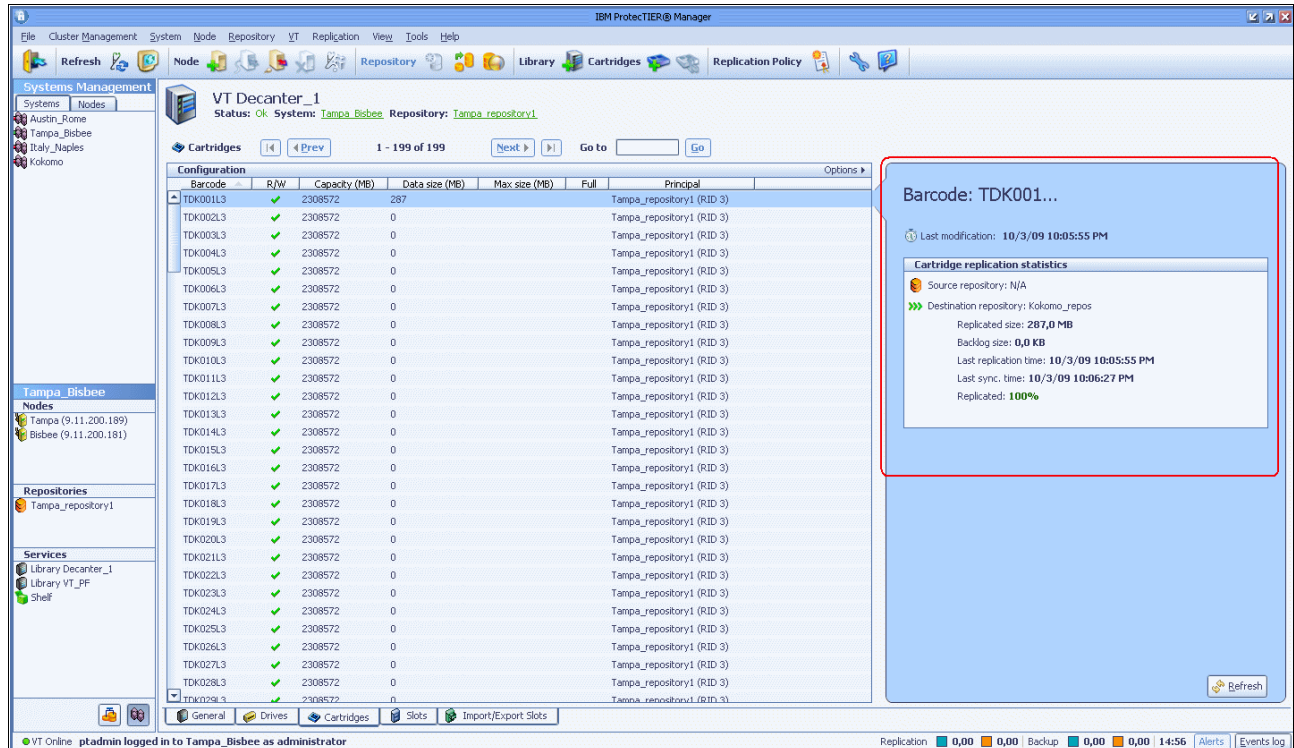


Figure 12-60 Cartridge replication attributes

The user can find attributes such as:

- ▶ Whether the cartridge is a part of the replication
- ▶ Where the cartridge is written to
- ▶ Whether this cartridge can change visibility
- ▶ Whether this cartridge is a replica
- ▶ Synchronization percentage towards a destination cartridge
- ▶ Last time of replication
- ▶ How much data is not replicated

To represent the current replication status of the cartridge we use a colored arrow (red, orange, or green) (Figure 12-61 through Figure 12-63). Beside the arrow color and the barcode, the window shows this additional information:

- ▶ Source repository: Not available at that time (NA)
- ▶ Destination repository: Primary replication side (repository)
- ▶ Replicated size: Current replicated data size in MB
- ▶ Backlog size: Backlog data size in MB
- ▶ Last replication time: Last time of data replication
- ▶ Last sync. time: Last synchronization of the source and target
- ▶ Replicated: Current replication status in percentage

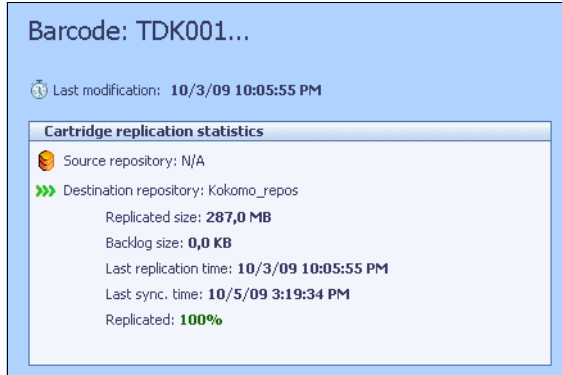


Figure 12-61 100% replicated cartridge (green color)

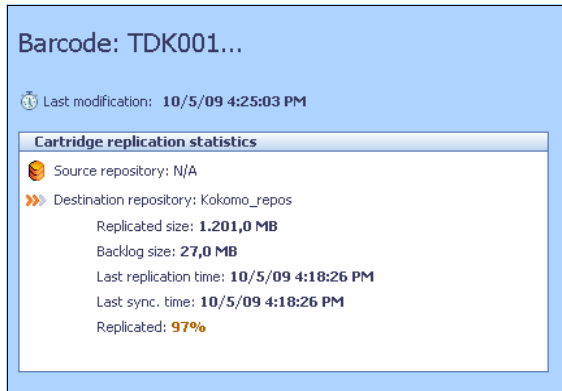


Figure 12-62 97% replicated cartridge (orange color)



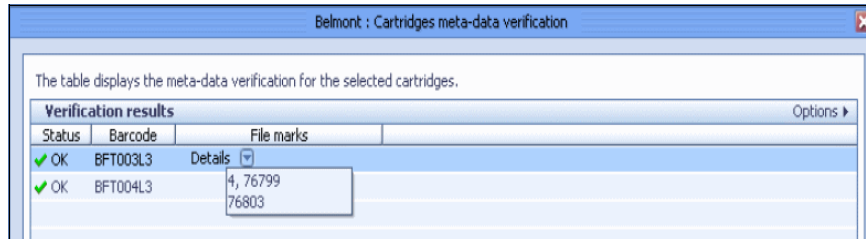
Figure 12-63 8% replicated cartridge (red color)

## 12.5.4 Cartridge verification after replication

After the replication from a virtual cartridge to the destination location completes, you can verify the metadata and the integrity of that cartridge by using two different verification tests. Through the comparison of the cartridge test results from the source and destination location, you can verify the data consistency.

### Cartridge metadata verification

The Cartridge metadata verification function (Figure 12-64) verifies the metadata only. It reads the files with cartridge metadata information and verifies checksums and other parameters, such as barcode and cart\_id. It also returns how many filemarks are on the cartridge.

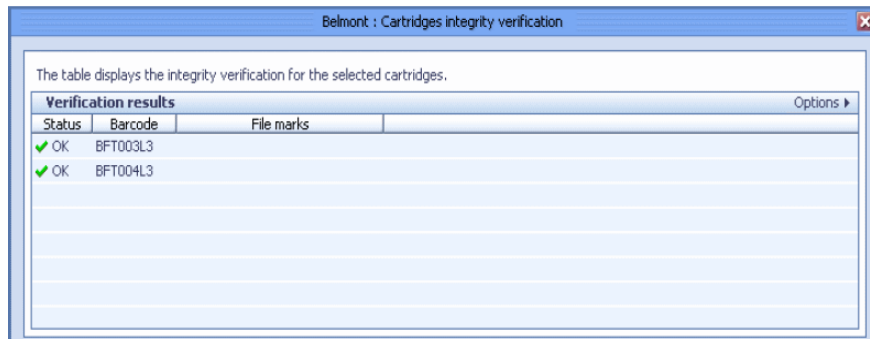


Status	Barcode	File marks
✓ OK	BFT003L3	Details
✓ OK	BFT004L3	4, 76799 76803

Figure 12-64 Cartridge metadata verification

### Cartridge integrity verification

The cartridge integrity verification function (Figure 12-65) reads all of the data that is referenced by the virtual cartridge metadata files into memory and, block by block, verifies all the structures of the data, making sure that there are no issues with block length, CRCs, and so on. This is essentially the same as performing a restore from the virtual tape cartridge, so this verification function takes approximately the same amount of time that a restore would take.



Status	Barcode	File marks
✓ OK	BFT003L3	
✓ OK	BFT004L3	

Figure 12-65 Cartridge integrity verification

For more detailed information about the tests, refer to 7.1.3, “Assessing cartridge status and synchronizing with the catalog” on page 347.

## 12.5.5 Replication long-term statistics

You can also collect node Replication long-term statistics by using the ProtecTIER GUI. Select **Node** → **Create and download long-term statistics**, as shown in Figure 12-66.

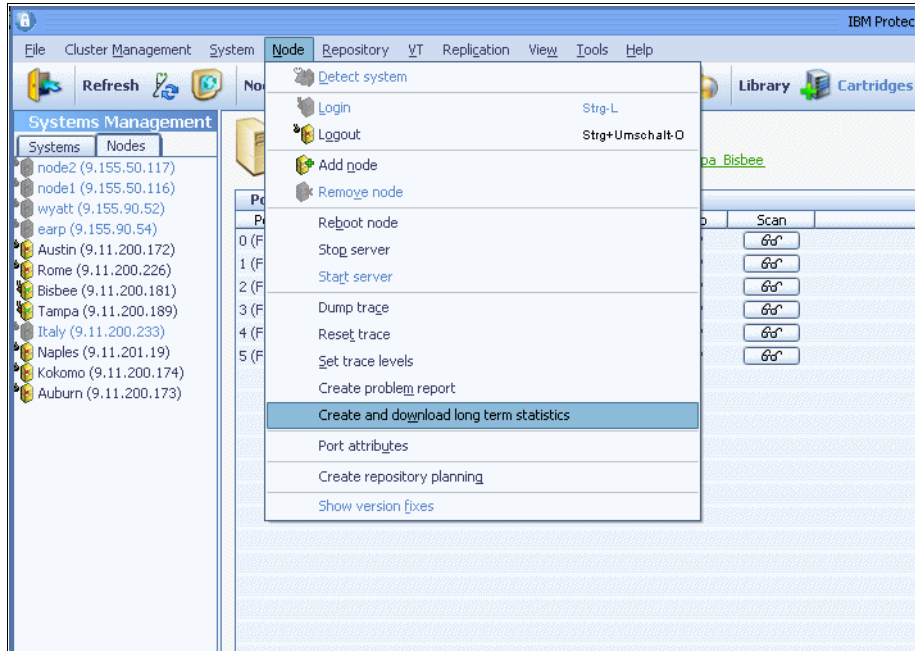


Figure 12-66 Create and download long-term statistics

In the next window, select the **ProtecTIER Replication Statistics Report**, as shown in Figure 12-67.

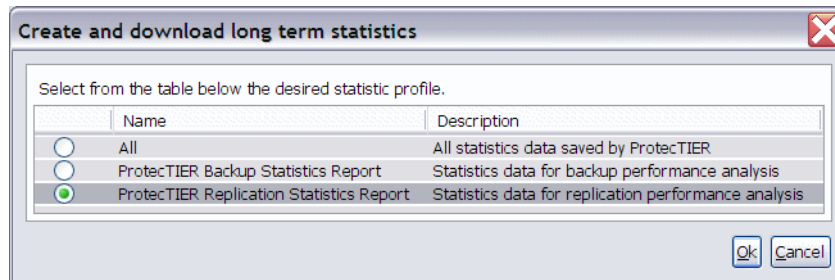


Figure 12-67 Selecting statistics for download

The system creates and downloads the long-term statistics (Figure 12-68).

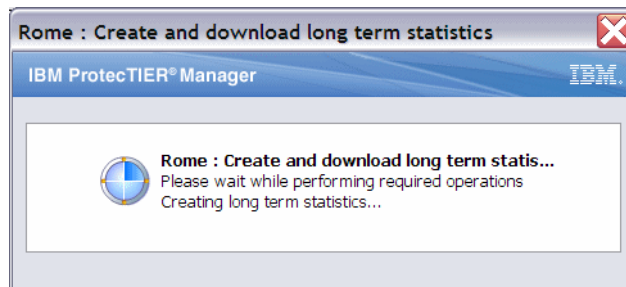


Figure 12-68 Downloading statistics

Now you can download the statistics (Figure 12-69).

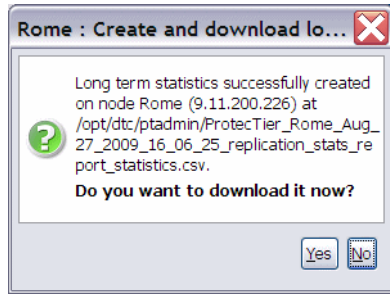


Figure 12-69 Download selection

You can select where you want to download your statistics at the server or where you are currently running PT Manager, including you workstation (Figure 12-70).

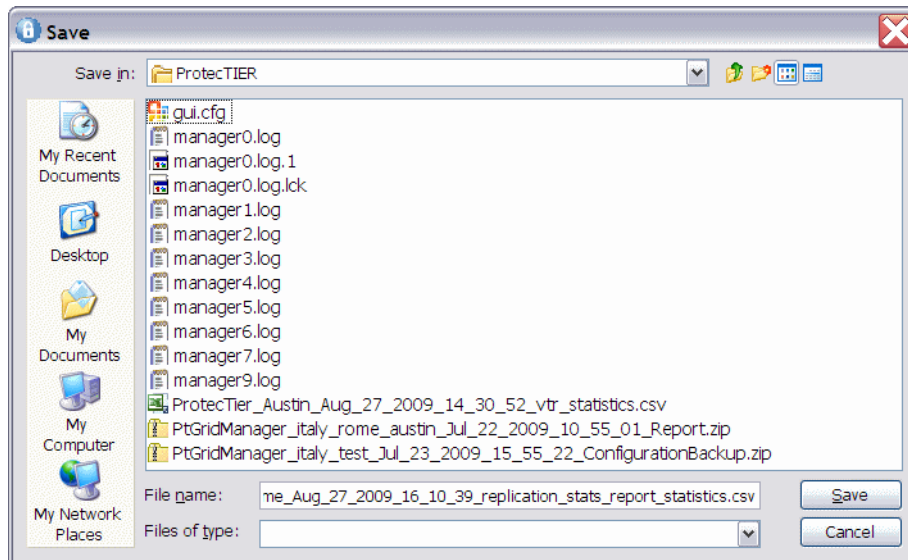


Figure 12-70 Selecting file location

Import the CSV file into you spreadsheet program. The following fields are available:

- ▶ samplingTimeSeconds
- ▶ SamplesPerPoll
- ▶ readBlocksBlocks.avg
- ▶ writeBlocksBlocks.avg
- ▶ incBlocksBlocks.avg
- ▶ HilnThr.avg
- ▶ nrmHilnThr.max
- ▶ nrmHiOutThr.avg
- ▶ nrmHiOutThr.max
- ▶ nrmHiNRunActv.avg
- ▶ nrmRunInNSuccActv.avg
- ▶ nrmRunOutNSuccActv.avg
- ▶ nrmRunInNAbrtActv.avg
- ▶ nrmRunOutNAbrtActv.av
- ▶ nrmRunInNFidActv.avg
- ▶ nrmRunOutNFidActv.avg
- ▶ TimReplicate.avg



- ▶ nrmTimAcceptReplication.avg
- ▶ nrmRunInMDPhysB.avg
- ▶ nrmRunInMDPhysB.std
- ▶ nrmRunInMDPhysB.min
- ▶ nrmRunInMDPhysB.max
- ▶ nrmRunOutMDPhysB.avg
- ▶ nrmRunOutMDPhysB.std
- ▶ nrmRunOutMDPhysB.min
- ▶ nrmRunOutMDPhysB.max
- ▶ nrmRunInMDNomB.avg
- ▶ nrmRunInMDNomB.std
- ▶ nrmRunInMDNomB.min
- ▶ nrmRunInMDNomB.max
- ▶ nrmRunOutMDNomB.avg
- ▶ nrmRunOutMDNomB.std
- ▶ nrmRunOutMDNomB.min
- ▶ nrmRunOutMDNomB.max
- ▶ nrmRunInUDPhysB.avg
- ▶ nrmRunInUDPhysB.std
- ▶ nrmRunInUDPhysB.min
- ▶ nrmRunInUDPhysB.max
- ▶ nrmRunOutUDPhysB.avg
- ▶ nrmRunOutUDPhysB.std
- ▶ nrmRunOutUDPhysB.min
- ▶ nrmRunOutUDPhysB.max
- ▶ nrmRunInUDNomB.avg
- ▶ nrmRunInUDNomB.std
- ▶ nrmRunInUDNomB.min
- ▶ nrmRunInUDNomB.max
- ▶ nrmRunOutUDNomB.avg
- ▶ nrmRunOutUDNomB.std
- ▶ nrmRunOutUDNomB.min
- ▶ nrmRunOutUDNomB.max
- ▶ rmHiBacklog.avgnrmTimVerifyDU.avg
- ▶ nrmRunValidateDataUnitBytes.avg
- ▶ nrcNetIncBytesPerSec.avg
- ▶ nrcNetIncBytesPerSec.max
- ▶ nrcNetOutBytesPerSec.avg
- ▶ nrcNetOutBytesPerSec.max
- ▶ VerifyDataUnitTime.mean
- ▶ VerifyDataUnitTime.max
- ▶ AppendDataUnitTime.mean
- ▶ AppendDataUnitTime.max
- ▶ GetSyncPointForCartTime.mean
- ▶ GetSyncPointForCartTime.max

## 12.5.6 Remote Cartridges report

Through ProtecTIER Manager, generate the cartridge's .csv file report by selecting **Replication** → **Create** and download statistics (Figure 12-71). The .csv file is created in the /pt\_work directory and a message is displayed to download the file (Figure 12-72) by default to the C:\Documents and Settings\Administrator\IBM\ProtecTIER directory on the workstation or server from which you are running PT Manager (Figure 12-73 on page 669). You can open the file with any spreadsheet application that you desire, such as Lotus Symphony.

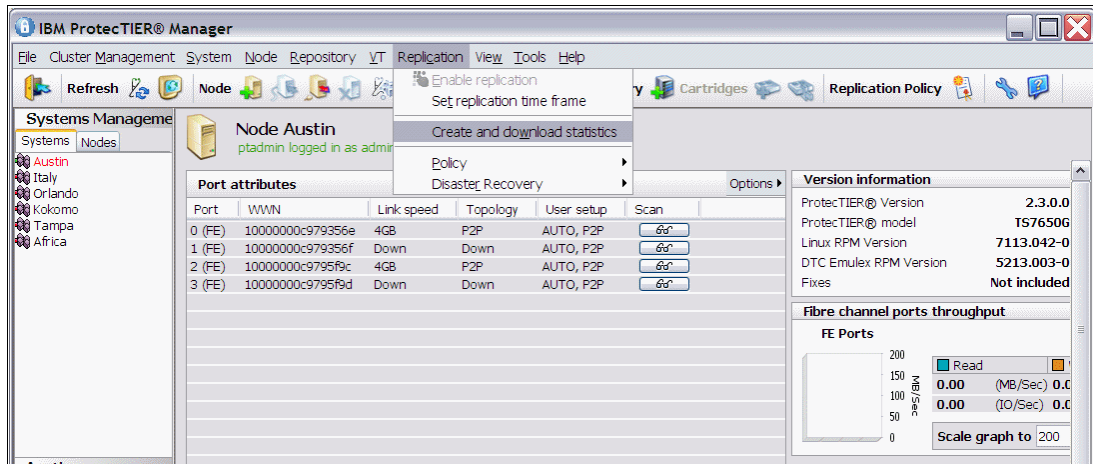


Figure 12-71 Replication Statistics

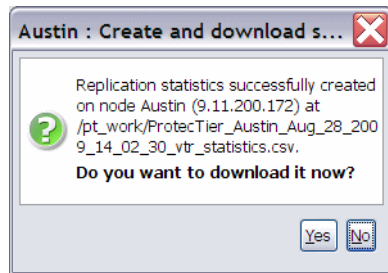


Figure 12-72 Create Statistics

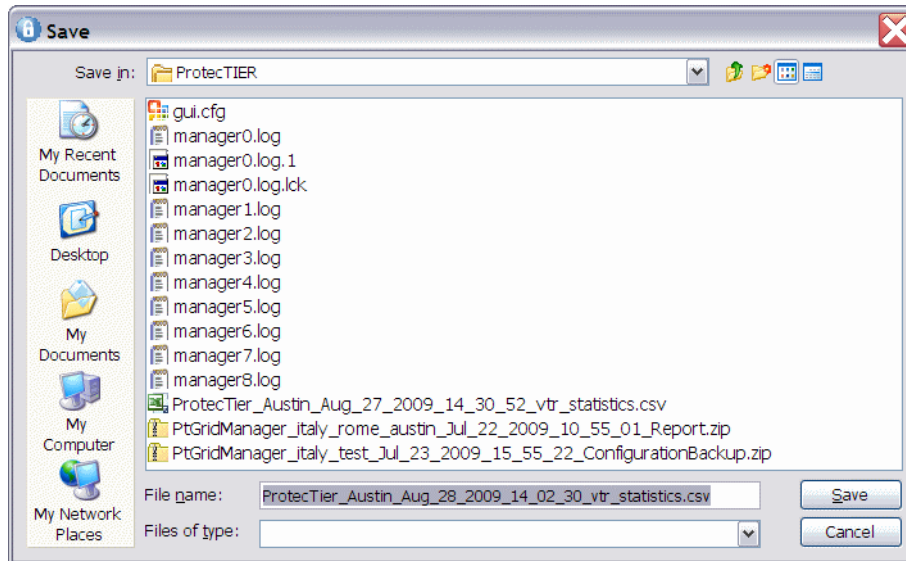


Figure 12-73 Download location

The report (Figure 12-74), which includes statistics for all the cartridges at the remote site, indicates whether the replica cartridge data is consistent with the associated sync time stamp (that is, the current data on the tape represents what it was at the specified sync time).

	A	B	C	D	E	F	G
1	cart unique id	barcode	nominal size	last update time	last sync	last sync point source time	last sync point destination time
2	1000	0003E8	1,048,576.00	05/07/2009 02:31	21	05/07/2009 02:31	05/07/2009 02:31
3	1001	0003E9	2,097,152.00	05/07/2009 03:30	21	05/07/2009 03:31	05/07/2009 03:31
4	1002	0003EA	3,145,728.00	05/07/2009 04:31	21	05/07/2009 04:31	05/07/2009 04:31
5	1003	0003EB	4,194,304.00	05/07/2009 05:31	21	05/07/2009 05:31	05/07/2009 05:31
6	1004	0003EC	5,242,880.00	05/07/2009 06:31	21	05/07/2009 06:31	05/07/2009 06:31
7	1005	0003ED	1,048,576.00	05/07/2009 07:31	21	05/07/2009 07:31	05/07/2009 07:31
8	1006	0003EE	2,097,152.00	05/07/2009 08:31	21	05/07/2009 08:31	05/07/2009 08:31
9	1007	0003EF	3,145,728.00	05/07/2009 09:31	21	05/07/2009 09:31	05/07/2009 09:31

Figure 12-74 CSV file

Compare the last update time, which represents the last time that the replica was updated, with 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 a consistent point in time. Otherwise, the cartridge is incomplete or in-transit. If the cartridge has a consistent point in time, ensure that this time stamp is larger than the full backup image end time. This indicates that the cartridge contains all the required data for this recovery operation. Otherwise, you 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 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 they get a sync point earlier than the backup end time.

As such, 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 require 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 processing the cartridge list to find a complete set of DR tapes, you must keep track of 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, and, as such, be unreliable. Thus, use the source date and time, rather than the destination sync time, when comparing cartridge states to the backup catalog/database. The destination sync time should only be used to determine which cartridges are whole.

In addition, there could be a time difference between the source backup server and the source ProtecTIER system. Your administrator should be aware of the discrepancy, measure it regularly, and communicate the delta to the 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 (that is, it will appear as a previous, old backup).

If there is uncertainty regarding the time differences, compare the nominal size of the cartridge to the catalog/DB value as an additional (not a substitute) layer of verification.

## 12.6 Network replication performance validation

The `pt_net_perf_util` utility's objective is to test maximal replication performance between two future ProtecTIER repositories by emulating the network usage patterns of the ProtecTIER native replication component. This utility does not predict replication performance, but it might discover performance bottlenecks.

The requirements of this utility are:

- ▶ Red Hat Enterprise Linux 5.2
- ▶ Standard external utilities expected to be in the current path:
  - ping
  - netstat
  - getopt
  - echo

The `pt_net_perf_util` utility and the `iperf` and `nuttcp` tools that it uses are installed as part of the ProtecTIER software installation. To test the replication performance, use one of the following tools:

- ▶ `iperf 2.0.4`:  
`/usr/local/bin/iperf`
- ▶ `nuttcp 6.1.2`:  
`/usr/local/bin/nuttcp-6.1.2`

This utility has two modes of operation, client and server. The server must be started before the client. Before running the utility, shut down all other programs on both the client and the server ProtecTIER systems. The *client* is the ProtecTIER system that transmits the test data and the *server* is the ProtecTIER system that receives the data (also known as the *target* server). Based on the data sent by the client and received by the server, the script outputs key network parameter values, which indicate certain attributes of the network. The goal of these tests is to benchmark the throughput of the network. The most important benchmark is the direction that replication will actually take place, that is, the target should be tested as the server because the flow of data will be *to* that server from the client. However, it is also important to also test the reverse direction to measure the bandwidth performance during disaster recovery failback. Network bandwidth is not always the same in both directions.

In the following procedure, the goal is to test network performance between two machines on a WAN, *server1* and *server2*. Each test runs for five minutes. Because there are five tests, the process takes a about of 25 minutes.

1. Start the server mode of the utility on *server1* by entering the commands in Example 12-1 on the command line.

*Example 12-1 Start the server mode using iperf*

---

```
cd /opt/dtc/app/sbin
./pt_net_perf_util -s
```

---

**Note:** The above commands use the iperf tool. To use the nuttcp tool instead, add -n to the command.

Enter one of the series of commands in Example 12-2 to use **nuttcp**.

*Example 12-2 Start the server mode using nuttcp*

---

```
cd /opt/dtc/app/sbin
./pt_net_perf_util -sn
or
cd /opt/dtc/app/sbin
./pt_net_perf_util -s -n
```

---

2. Start the client mode of the utility on *server2* by entering the following command on the command line:

```
./pt_net_perf_util -c server1 -t 300
```

**Note:** This step uses the iperf external utility. To use nuttcp instead, add -n to the command.

Enter the following command to use nuttcp:

```
./pt_net_perf_util -c server1 -t 300 -n
```

3. The utility automatically performs the tests in sequence. The *client* output (*server2* in the example) looks similar to Example 12-3.

**Note:** In the sample output in Example 12-3 the test ran for only 5 seconds instead of 300.

*Example 12-3 Client output*

---

```
*** Latency
PING 9.5.53.33 (9.5.53.33) 56(84) bytes of data.
--- 9.5.53.33 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4001ms
rtt min/avg/max/mdev = 0.257/0.406/0.484/0.079 ms
*** Throughput - Default TCP
[ 3] 0.0- 5.0 sec 56.6 MBytes 94.8 Mbits/sec
*** Throughput - 1 TCP stream(s), 1MB send buffer
[ 3] 0.0- 5.0 sec 57.0 MBytes 95.0 Mbits/sec
*** Throughput - 16 TCP stream(s), 1MB send buffer
[SUM] 0.0- 5.8 sec 65.0 MBytes 94.3 Mbits/sec
*** Throughput - 127 TCP stream(s), 1MB send buffer
```

```
[SUM] 0.0-11.3 sec 127 MBytes 94.1 Mbits/sec
Number of TCP segments sent: 230536
Number of TCP retransmissions detected: 21 (0%)
Done.
```

---

Refer to 12.6.1, “Interpreting the test results” on page 672 for information about interpreting the results of the tests.

## 12.6.1 Interpreting the test results

The utility performs five foreground tests (“Test 1: Latency” on page 672 through “Test 5: Throughput (127 streams, 1 MB send buffer)” on page 673) and one background test (“Test 6: TCP retransmissions versus total TCP segments sent” on page 674). The example outputs in the following section are from the client side, with the script using `iperf` (not `nuttcp`) in tests 2 - 5. Each of the first five tests ran for 30 seconds (`-t 300`), while the last test monitored TCP performance during that time.

### Test 1: Latency

This test checks the nominal network link latency and packet loss (Example 12-4).

*Example 12-4 Nominal network link latency and packet loss check*

---

```
*** 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
```

---

Interpreting the results:

- ▶ The average round-trip-time (rtt) was 78.4 ms and there was 0% packet loss.
- ▶ The latency in WAN topologies might vary, but should never exceed 200 ms. Contact your network administrator if latency reports more than 200 ms, as that might significantly decrease replication throughput.
- ▶ Higher latency values will cause a major deterioration in replication throughput.
- ▶ Packet loss should be 0%. Any other value implicates a major network problem.

### Test 2: Throughput (default settings)

This test checks maximal TCP throughput using a single data stream with default TCP settings (Example 12-5).

*Example 12-5 Maximal TCP throughput check with default TCP settings*

---

```
*** Throughput - Default TCP
[ 3] 0.0-120.1 sec 2.41 GBytes 173 Mbits/sec
```

---

**Note:** 1 MByte = 1,048,576 bytes. 1 Mbit/sec = 1,000,000 bits/sec.

Interpreting the results:

The test ran for 120.1 seconds, transferred 2.41 GB, with an average throughput of 173 Mbps.

### Test 3: Throughput (single stream, 1 MB send buffer)

This test checks maximal TCP throughput using a single data stream with a 1 MB send buffer (Example 12-6).

*Example 12-6 Maximal TCP throughput check with 1 MB send buffer*

---

```
*** Throughput - 1 TCP stream(s), 1MB send buffer  
[ 3] 0.0-120.0 sec 2.51 GBytes 180 Mbits/sec
```

---

Interpreting the results:

- ▶ The test ran for 120.0 seconds and transferred 2.51 GB, with an average throughput of 180 Mbps.
- ▶ There might be an improvement here for high-latency links.

### Test 4: Throughput (16 streams, 1 MB send buffer)

Example 12-7 shows an example result.

*Example 12-7 Throughput with 16 streams, 1 MB send buffer*

---

```
*** Throughput - 16 TCP stream(s), 1MB send buffer  
[SUM] 0.0-121.4 sec 5.91 GBytes 418 Mbits/sec
```

---

Interpreting the results:

- ▶ The test ran for 121.4 seconds, transferred 5.91 GB, with an average throughput of 418 Mbps.
- ▶ The extra streams yielded higher utilization of the connection.
- ▶ The Mbps reported in this test is the maximum replication performance that your system will achieve if your backup environment uses 2 - 3 cartridges in parallel.

### Test 5: Throughput (127 streams, 1 MB send buffer)

Example 12-8 shows an example result.

*Example 12-8 Throughput with 127 streams, 1 MB send buffer*

---

```
*** Throughput - 127 TCP stream(s), 1MB send buffer  
[SUM] 0.0-126.1 sec 8.08 GBytes 550 Mbits/sec
```

---

Interpreting the results:

- ▶ The test ran for 126.1 seconds and transferred 8.08 GB, with an average throughput of 550 Mbps.
- ▶ TCP takes a while to reach its maximal throughput. Longer testing times, such as 300 seconds or more, produce more accurate results.
- ▶ The throughput value given by this test is the potential physical replication throughput for this system. It is directly affected by the available bandwidth, latency, packet loss, and retransmission rate.
- ▶ The Mbps reported in this test is the maximum replication performance that your system might achieve. If this number is lower than anticipated, contact your network administrator.

## Test 6: TCP retransmissions versus total TCP segments sent

Example 12-9 shows an example result.

*Example 12-9 TCP retransmissions versus total TCP segments sent*

---

Number of TCP segments sent: 1619061

Number of TCP retransmissions detected: 201038 (12%)

---

Interpreting the results:

- ▶ A total of 1619061 TCP segments were sent during the five tests, out of which 201038 were lost and retransmitted.
- ▶ The retransmission rate imposes a direct penalty on the throughput, as the retransmission of these packets takes up bandwidth. The retransmission can be caused by the underlying network (for example, packet dropping by an overflowed router) or by the TCP layer itself (for example, retransmission due to packet reordering).
- ▶ Segment loss can be caused by each of the network layers.
- ▶ TCP retransmission larger than 2% might cause performance degradation and unstable network connectivity. Contact your network administrator to resolve this issue and reduce it to approximately 0%.

You might want to run these tests again to test the reverse throughput in the network. To run the tests in reverse, change *server1* to the client and *server2* to the server and repeat the procedures.

## 12.7 Managing and reporting TS7650 and TS7650G using the command-line interface

This section describes how to query the system using the command-line interface (CLI) and how to receive various statistics about the IBM System Storage TS7600 with ProtecTIER system. This section provides valuable insights for the administrator about the performance, capacity, configuration, and operation of the system, and can be accessed by other management applications.

### 12.7.1 Command-line interface

This section describes how to query the ProtecTIER system through the CLI using **ptcli**. The **ptcli** utility is loaded during the installation of ProtecTIER software. Make sure that you are in this directory when running **ptcli** commands from the command prompt.

#### Usage

Run **ptcli** to query the system through the CLI, receive various statistics about the ProtecTIER system, and issue commands that are available in the interface. This information can provide valuable insights for the administrator about the performance, capacity, configuration, and operation of the system, and can be accessed by other management applications. The **ptcli** command is issued from the command line as follows:

```
./ptcli <--parameter> <--server options> <--parameter options> <variables>  
<-processing options>
```



## Parameters and options

The parameters and options available with the `ptcli` command are:

### ► AddCartridges

Adds cartridges to the library:

- name *<NAME>*  
Specify the library name (taken from the Libraries output).
- cartridges *<NUM>*  
Specify the number of cartridges.
- maxcartsize *<NUM>* cartridges  
Specify the maximum cartridge growth size in megabytes. This is optional. The default is no limit.
- seed *<SEED>*  
Specify the barcode seed.

### ► AddClustermember

Adds a node to the cluster:

`addip <IP>`

Specify the external IP address of the node to add to the cluster.

### ► AddLibrary

Creates a new library:

- name *<NAME>*  
Specify the library name (taken from the Libraries output).
- slots *<NUM>*  
Specify the number of slots. This is optional. The default is 0.
- imexp *<NUM>*  
Specify the number of imports and exports. This is optional. The default is 8.
- libtype *<NAME>*  
Specify the library type (taken from the LibrariesTypes output).
- robot *<X, Y>*  
Specify a list of robot assignments in the form of X, Y (for example, X=node external IP address and Y=port).
- drivemodel *<NAME>*  
Specify the tape drive model (taken from the DriveModels output).
- drives *<X, Y, Z>*  
Specify a list of drive assignments in the form of X, Y, Z (for example, X=number of drives, Y=node external IP address, and Z=port).

### ► Cartridge info

Prints information about specific cartridges in the library. The list is sorted by barcode.

- name *<NAME>*  
Specify the library name (taken from the Libraries output).

- from *<NUM>*  
Specify the number of cartridges before the first printed cartridge in the list of cartridges (taken from the Number of Cartridges output).
- count *<NUM>*  
Specify the maximum number of cartridges in the output.
- ▶ Create Repository  
Creates a new repository on a node:
  - name *<NAME>*  
Specify the repository name.
  - system *<NAME>*  
Specify the system name.
  - peak *<NUM>*  
Specify the maximum number of cartridges in the output.
  - ratio *<NUM>*  
Specify the deduplication ratio.
  - raid *<X, Y, Z>*  
Specify the metadata raid configuration in the form of X, Y, Z (for example, X=Type, Y=Members, and Z=Disk size in gigabytes, taken from the RAID Configurations output).
  - metadata *<MOUNT\_POINTS>*  
Specify a list of GFS mounted file systems that will be used to store ProtecTIER repository metadata. This is optional.
  - userdata *<MOUNT\_POINTS>*  
Specify a list of GFS mounted file systems that will be used to store ProtecTIER repository user's data. This is optional.
- ▶ DriveModels  
Prints information about the supported tape drive models:  
libtype *<NAME>*  
Specify the library type.
- ▶ Libraries  
Prints the list of libraries on the repository.
- ▶ LibraryInfo  
Prints information about a specific library in the repository:  
name *<NAME>*  
Specify the library name (taken from the Libraries output).
- ▶ Librarytypes  
Prints information about the supported library types.
- ▶ NodeVtlStatistics  
Prints the statistics history about the local host:

hours <NUM>

Specify the number of statistics hours included in the output. By default, the output should include four chronicles (statistics records) per hour (if uptime is greater than or equal to hours).

- ▶ NumberofCartridges  
Prints the number of cartridges in a library.  
name <NAME>  
Specify the library name (taken from Libraries output).
- ▶ RaidConfigurations  
Prints information about the supported RAID configurations.
- ▶ Rpositorystatistics  
Prints the repository statistics.
- ▶ ServerVersion  
Prints the server version.

## Server options

Server options are used for connection to the management server.

**Note:** Administrator privileges are required for configuration operations.

- ▶ ip <IP>  
Specify the IP address of the management server.
- ▶ port <PORT>  
Specify the port of the management server. This is optional. The default is 3501.
- ▶ username <NAME>  
Specify the user name for the login.
- ▶ password <PASSWORD>  
Specify the password for the login.
- ▶ force  
Force the login if another administrator is already logged in. This is optional.

## Processing options

You can specify the following processing options at the end of the command string:

- ▶ d <debug level>  
Specify the debug level to print log messages to the console. Choices for the debug level are:
  - SEVERE
  - WARNING
  - ALL
- ▶ f <logfile>  
Specify the name (with full path) to use for the log file name. The defaults are c:\Documents and Settings\hanoch\cli0.log.

- ▶ `h`  
Prints this message and exits.
- ▶ `v`  
Prints the version information and exits.

## ptcli responses

`ptcli` responses always appear in a well-structured printed XML document.

The XML always has the same root element, *response*, with attribute, *status*, which may have the value *success* or *failed*. The response element displays the reply as a child element on success and, in this case, returns an exit code of zero to the shell. If status is failed, the response element has a single child element, *error*, with attribute *description* (a human readable string) and a non-zero *code* value. The same value is returned to the shell.

### Example of failure

Example 12-10 is an example of calling to DriveModels without `--libtype`.

*Example 12-10 Failed response for calling to DriveModels without --libtype*

---

```
<?xml version="1.0" encoding="UTF-8" ?>
<response status="failed">
<error code="3" description="Missing required option --libtype" />
</response>
```

---

Example 12-11 is an example of calling to a server without running VTFD.

*Example 12-11 Failed response for calling to DriveModels without --libtype*

---

```
<?xml version="1.0" encoding="UTF-8" ?>
<response status="failed">
<error code="3" description="Management server 168.159.150.60 is not
responding." />
</response>
```

---

### Example of success

Example 12-12 shows the success of the configuration operations. The XML for all configuration operations is the same. The only difference is the name.

*Example 12-12 Successful response for configuration operations*

---

```
<?xml version="1.0" encoding="UTF-8" ?>
<response status="success">
<command host-name="beast" host-time="1243147115"
host-time-string="24-May-2009 09:38" local-time="1243146789"
local-time-string="24-May-2009 09:33" name="AddLibrary" />
</response>
```

---

Example 12-13 shows the success of the monitoring operations:

```
./ptcli RepositoryStatistics --ip 168.159.151.62 --username ptoper --password
ptoper
```

*Example 12-13 Successful response for RepositoryStatistics*

---

```
<?xml version="1.0" encoding="UTF-8" ?>
```

```
<response status="success">
<repository-statistics current-factoring-ratio="1.00"
host-name="mir" repository-allocable-physical-storage-mb="724992"
repository-free-physical-storage-mb="724992"
repository-total-nominal-storage-mb="0"
repository-total-physical-storage-mb="724992"
repository-used-physical-storage-mb="0"
host-time="1243147611" host-time-string="24-May-2009 09:46"
local-time="1243147285" local-time-string="24-May-2009 09:41"/>
</response>
```

---

Example 12-14 shows the results of the following command:

```
./ptcli Libraries --ip 168.159.150.60 --username ptoper --password ptoper
```

*Example 12-14 Successful response for libraries*

---

```
<?xml version="1.0" encoding="UTF-8" ?>
<response status="success">
<libraries-list host-name="mir" num-returned="2"
host-time="1243147611" host-time-string="24-May-2009 09:46"
local-time="1243147285" local-time-string="24-May-2009 09:41">
<library name="Lib1" unique-id="2199023256354" />
<library name="Lib2" unique-id="2199023256454" />
</libraries-list>
</response>
```

---

Example 12-15 shows the results of the following command:

```
/ptcli LibraryInfo --ip 168.159.151.62 --username ptoper --password ptoper --name
Pisa_Parma_24X7 without active drives
```

*Example 12-15 Successful response for LibraryInfo without active drive*

---

```
<?xml version="1.0" encoding="UTF-8" ?>
<response status="success">
<library-info host-name="parma" library-name="Pisa_Parma_24X7"
library-unique-id="2199023265953" num-total-drives="30"
host-time="1243147611" host-time-string="24-May-2009 09:46"
local-time="1243147285" local-time-string="24-May-2009 09:41">
<active-drives num-returned="0" />
</library-info>
</response>
```

---

Example 12-16 shows the results of the following command:

```
./ptcli LibraryInfo --ip 168.159.151.62 --username ptoper --password ptoper --name
Pisa_Parma_24X7 with active drives
```

*Example 12-16 Successful response for LibraryInfo with active drives*

---

```
<?xml version="1.0" encoding="UTF-8" ?>
<response status="success">
<library-info host-name="parma" library-name="Pisa_Parma_24X7"
library-unique-id="2199023265953" num-total-drives="30"
host-time="1243147611" host-time-string="24-May-2009 09:46"
local-time="1243147285" local-time-string="24-May-2009 09:41">
<active-drives num-returned="12">
```

```

<drive lea="0" loaded-cartridge-barcode="001025" node-id="1"
read-rate-mb-per-second="0" write-rate-mb-per-second="4" />
<drive lea="2" loaded-cartridge-barcode="001024" node-id="1"
read-rate-mb-per-second="0" write-rate-mb-per-second="5" />
<drive lea="8" loaded-cartridge-barcode="001028" node-id="1"
read-rate-mb-per-second="0" write-rate-mb-per-second="3" />
<drive lea="10" loaded-cartridge-barcode="001027" node-id="1"
read-rate-mb-per-second="0" write-rate-mb-per-second="2" />
<drive lea="14" loaded-cartridge-barcode="001022" node-id="1"
read-rate-mb-per-second="0" write-rate-mb-per-second="5" />
<drive lea="18" loaded-cartridge-barcode="001021" node-id="1"
read-rate-mb-per-second="0" write-rate-mb-per-second="8" />
<drive lea="7" loaded-cartridge-barcode="001018" node-id="2"
read-rate-mb-per-second="0" write-rate-mb-per-second="6" />
<drive lea="17" loaded-cartridge-barcode="001019" node-id="2"
read-rate-mb-per-second="0" write-rate-mb-per-second="6" />
<drive lea="19" loaded-cartridge-barcode="001026" node-id="2"
read-rate-mb-per-second="0" write-rate-mb-per-second="4" />
<drive lea="21" loaded-cartridge-barcode="001020" node-id="2"
read-rate-mb-per-second="0" write-rate-mb-per-second="5" />
<drive lea="25" loaded-cartridge-barcode="001030" node-id="2"
read-rate-mb-per-second="0" write-rate-mb-per-second="0" />
<drive lea="27" loaded-cartridge-barcode="001029" node-id="2"
read-rate-mb-per-second="0" write-rate-mb-per-second="4" />
</active-drives>
</library-info>
</response>

```

---

Example 12-17 shows the results of the following command:

```

./ptcli NumberOfCartridges --ip 168.159.150.60 --username ptadmin --password
ptadmin --force --name Lib1

```

*Example 12-17 Successful response for NumberOfCartridges*

```

<?xml version="1.0" encoding="UTF-8" ?>
<response status="success">
<library host-name="mir" name="Lib1" num-configured-cartridges="20"
host-time="1243147611" host-time-string="24-May-2009 09:46"
local-time="1243147285" local-time-string="24-May-2009 09:41"
unique-id="2199023256354" />
</response>

```

---

Example 12-18 shows the results of the following command:

```

./ptcli CartridgeInfo --ip 168.159.150.60 --username ptadmin --password ptadmin
--force --name Lib1 --from 0 --count 200

```

*Example 12-18 Successful response for CartridgeInfo*

```

<?xml version="1.0" encoding="UTF-8" ?>
<response status="success">
<cartridges-information host-name="mir" library-name="Lib1"
library-unique-id="2199023256354" num-returned="20"
host-time="1243147853"
host-time-string="24-May-2009 09:50" local-time="1243147527"
local-time-string="24-May-2009 09:45">

```

```

<cartridge barcode="000000" early-warning-reported="false"
nominal-size-mb="0" />
<cartridge barcode="000001" early-warning-reported="false"
nominal-size-mb="0" />
<cartridge barcode="000002" early-warning-reported="false"
nominal-size-mb="0" />
<cartridge barcode="000003" early-warning-reported="false"
nominal-size-mb="0" />
<cartridge barcode="000004" early-warning-reported="false"
nominal-size-mb="0" />
<cartridge barcode="000005" early-warning-reported="false"
nominal-size-mb="0" />
<cartridge barcode="000006" early-warning-reported="false"
nominal-size-mb="0" />
<cartridge barcode="000007" early-warning-reported="false"
nominal-size-mb="0" />
<cartridge barcode="000008" early-warning-reported="false"
nominal-size-mb="0" />
<cartridge barcode="000009" early-warning-reported="false"
nominal-size-mb="0" />
<cartridge barcode="000010" early-warning-reported="false"
nominal-size-mb="0" />
<cartridge barcode="000011" early-warning-reported="false"
nominal-size-mb="0" />
<cartridge barcode="000012" early-warning-reported="false"
nominal-size-mb="0" />
<cartridge barcode="000013" early-warning-reported="false"
nominal-size-mb="0" />
<cartridge barcode="000014" early-warning-reported="false"
nominal-size-mb="0" />
<cartridge barcode="000015" early-warning-reported="false"
nominal-size-mb="0" />
<cartridge barcode="000016" early-warning-reported="false"
nominal-size-mb="0" />
<cartridge barcode="000017" early-warning-reported="false"
nominal-size-mb="0" />
<cartridge barcode="000018" early-warning-reported="false"
nominal-size-mb="0" />
<cartridge barcode="000019" early-warning-reported="false"
nominal-size-mb="0" />
</cartridges-information>
</response>

```

---

Example 12-19 shows the results of the following command:

```

./ptcli NodeVtlStatistics --ip 168.159.150.60 --username ptadmin --password
ptadmin --force --hours 1

```

*Example 12-19 Successful response for NodeVtlStatistics*

---

```

<?xml version="1.0" encoding="UTF-8" ?>
<response status="success">
<statistics host-name="mir" num-returned="1" host-time="1243147853"
host-time-string="24-May-2009 09:50" local-time="1243147527"
local-time-string="24-May-2009 09:45">
<chronicle n-samples-per-poll="30" sample-date="06-05-2009 15:08:15"

```

```
sample-time-seconds="30">
<vtl-overall num-active-tapes="0" num-reading-tapes="0"
num-writing-tapes="0" read-io.average-bytes="0" read-io.max-bytes="0"
read-io.min-bytes="0" valid="true" write-io.average-bytes="0"
write-io.max-bytes="0" write-io.min-bytes="0" />
</chronicle>
</statistics>
</response>
```

---

Example 12-20 shows the results of the following command:

```
./ptcli ServerVersion --ip 168.159.150.60 --username ptofer --password ptofer
```

*Example 12-20 Successful response for server version*

---

```
<?xml version="1.0" encoding="UTF-8" ?>
<response status="success">
<version-info host-name="mir" host-time="1243147853"
host-time-string="24-May-2009 09:50" local-time="1243147527"
local-time-string="24-May-2009 09:45">
<version build="7.113.2" pt-version="2.2.0.TESTING10013" />
</version-info>
</response>
```

---

Example 12-21 shows the results of the following command for library types:

```
./ptcli LibraryTypes --ip 168.159.150.60 --username ptofer --password ptofer
```

*Example 12-21 Successful response for library types*

---

```
<?xml version="1.0" encoding="UTF-8" ?>
<response status="success">
<version-info host-name="mir" host-time="1243147853"
host-time-string="24-May-2009 09:50" local-time="1243147527"
local-time-string="24-May-2009 09:45">
<version build="7.113.2" pt-version="2.2.0.TESTING10013" />
</version-info>
</response>
```

---

Example 12-22 shows the results of the command for drive models:

```
./ptcli DriveModels --ip 168.159.150.60 --username ptofer --password ptofer
--libtype TS3500
```

*Example 12-22 Successful response for drive models*

---

```
<?xml version="1.0" encoding="UTF-8" ?>
<response status="success">
<drive-models host-name="beast" host-time="1243147853"
host-time-string="24-May-2009 09:50" local-time="1243147527"
local-time-string="24-May-2009 09:45" num-returned="1">
<drive-model name="ULT3580-TD3" />
</drive-models>
</response>
```

---



Example 12-23 shows the results for the RAID configurations command:

```
./ptcli RaidConfigurations--ip 168.159.150.60 --username ptadmin --password ptadmin (Requires administrator privileges)
```

*Example 12-23 Successful response for RaidConfigurations*

---

```
<?xml version="1.0" encoding="UTF-8" ?>
<response status="success">
<raid-configurations host-name="beast" host-time="1243148247"
host-time-string="24-May-2009 09:57" local-time="1243147921"
local-time-string="24-May-2009 09:52" num-returned="18">
<raid-configuration members="2+2" type="FC-10K" />
<raid-configuration members="3+3" type="FC-10K" />
<raid-configuration members="4+4" type="FC-10K" />
<raid-configuration members="5+5" type="FC-10K" />
<raid-configuration members="6+6" type="FC-10K" />
<raid-configuration members="8+8" type="FC-10K" />
<raid-configuration members="2+2" type="FC-15K" />
<raid-configuration members="3+3" type="FC-15K" />
<raid-configuration members="4+4" type="FC-15K" />
<raid-configuration members="5+5" type="FC-15K" />
<raid-configuration members="6+6" type="FC-15K" />
<raid-configuration members="8+8" type="FC-15K" />
<raid-configuration members="2+2" type="SATA-7.2K" />
<raid-configuration members="3+3" type="SATA-7.2K" />
<raid-configuration members="4+4" type="SATA-7.2K" />
<raid-configuration members="5+5" type="SATA-7.2K" />
<raid-configuration members="6+6" type="SATA-7.2K" />
<raid-configuration members="8+8" type="SATA-7.2K" />
</raid-configurations>
</response>
```

---

**Note:** Responses might vary based on your installation.





# Part 4

# Appendixes





# A

## **Installation and implementation checklists**

In this appendix, we summarize checklists and worksheets that will help you during the planning and installation of the IBM System Storage TS7600 with ProtecTIER.

## Customer installation responsibilities

You are responsible for preparing the installation site prior to the installation of the IBM System Storage TS7600 with ProtecTIER. All physical planning for the IBM System Storage TS7600 with ProtecTIER is a customer responsibility. In summary, the customer is responsible for:

- ▶ Providing the site preparation: Cabling and wiring for connections to the host, cooling and heating, telephone service, safety, and security.
- ▶ Providing network connections: Cables and switches.
- ▶ Providing SNMP catcher, host clients, and email services for RSA alerts.
- ▶ Providing all the necessary IP addresses for the installation of the IBM System Storage TS7600 with ProtecTIER.
- ▶ For TS7650G installations where customer-supplied frames, disk arrays, and expansions are used, these items must be fully installed and operational before the installation of the TS7650G can begin. This is done by your IBM SSR.
- ▶ Providing a workstation where the ProtecTIER Manager software can be installed, as well as appropriate adapters and cables. Fibre Channel cables are required to attach the IBM System Storage TS7600 with ProtecTIER to various server adapters.

## Customer information work sheet

This work sheet (Table A-1) allows you to provide basic information about your company, administrator, and general system information.

IBM service representatives use the information that is provided on the company information work sheet to customize your IBM System Storage TS7600 with ProtecTIER. When you use any of the remote support features, the TS3000 System Console sends this information to IBM so that the IBM Support center can contact you.

You must complete this work sheet for all installations that include a Tape Storage Service Console (TSSC).

*Table A-1 Company Information work sheet*

Item or setting	Instructions	Your information
Company name	Provide the name of your company. IBM service representatives use this information to identify your company when they receive Call Home reports from your IBM System Storage TS7600 with ProtecTIER. Ensure that the company name that is provided here is consistent with all other machines that correspond to your IBM customer account.	
Customer number	Provide the customer number that is assigned by IBM to your company.	

Item or setting	Instructions	Your information
<b>Administrator information: Provide information about your storage system administrator in the following section.</b>		
Administrator name	Provide the name of the individual at your site who service representatives can contact about IBM System Storage TS7600 with ProtecTIER service matters.	
Administrator email address	Provide the email address that can be used to contact the administrator.	
Administrator telephone number	Provide the primary telephone number for service personnel to use to contact the IBM System Storage TS7600 with ProtecTIER administrator. Include the area code and the country code, if appropriate.	
Alternate telephone number	Provide an alternate or off-shift telephone number that IBM service representatives can use to contact the IBM System Storage TS7600 with ProtecTIER administrator. Include the area code and the country code, if appropriate.	
Fax number	Provide the primary fax number that IBM service representatives can use when they must fax documents to the IBM System Storage TS7600 with ProtecTIER administrator. Include the area code and the country code, if appropriate.	
Alternate fax number	Provide an alternate fax number that service personnel can use when they must fax documents to the IBM System Storage TS7600 with ProtecTIER administrator. Include the area code and the country code, if appropriate.	
Administrator's mailing address	Provide the mailing address for the administrator. Specify the complete address, including the street address, building (if appropriate), city or locality, state or province, and postal or zip code.	

Item or setting	Instructions	Your information
<b>IBM System Storage TS7600 with ProtecTIER system information: We provide basic information about your TS7650 or TS7650G and the TS3000 in the following section.</b>		
TS7650 or TS7650G location	If different from the administrator's address, provide the full address where the IBM System Storage TS7600 with ProtecTIER unit is located. Include the street address, building (if appropriate), city or locality, state or province, and postal or zip code.	

## Customer network settings work sheets

In this section, we provide a worksheet to be used when planning your IBM System Storage TS7600 with ProtecTIER. For a smooth implementation, make sure that all required IP addresses are available before starting the implementation. Do *not* start implementation when one of the required IP addresses is missing.

### TSSC network IP scheme

The TSSC IP address range changes from frame to frame and each new frame increments by a multiple of 10. For example, the first stand-alone frame range would be 10 and the second frame range would be 20. By default, the first stand-alone or two-node cluster uses IP address range 172.31.1.110. The next stand-alone or two-node cluster uses address range 172.31.1.20, and so on. Table A-2 shows the TSSC IP addresses for a 3958-DD3.

**Note:** In a two-node clustered system, node A and node B are installed in the same frame.

Table A-2 TSSC ports for a 3958-DD3

Component	Port	IP addresses: Stand-alone server or server A in a two-node cluster	IP addresses: Server B in a two-node cluster
Server	Eth5	172.31.1.x0	172.31.1.x5
Server	RSA	172.31.1.x1	172.31.1.x6

Depending on the address ranges available on the TSSC, the IBM SSR may have to use ranges other than those shown here.



## IP address worksheet

Use this worksheet to specify the IP addresses assigned to the IBM System Storage TS7600 with ProtecTIER components. IBM service representatives use the information provided to define the IP addresses of components supported by the TSSC. When the TSSC sends Call Home information to IBM through a VPN or modem or sends you notices about serviceable events, these settings will be included in the information to identify and provide important information about the TSSC that sent a service request.

We created the following tables explaining which ports are used for a stand-alone and a two-node clustered IBM System Storage TS7600 with ProtecTIER.

Table A-3 shows the required IP addresses for a stand-alone IBM System Storage TS7600 with ProtecTIER.

*Table A-3 Factory-default server IP addresses for a stand-alone TS7650G or TS7650*

IBM System Storage TS7600 with ProtecTIER stand-alone	Component	Port	IP address
NODE A: the server located in the lower part of the rack	Server	eth0	External customer IP
	Server	eth1	10.0.0.51
	Server	eth2	NA
	Server	eth3	NA
	Server	eth4	NA
	Server	eth5	172.31.1.xx TSSC

Table A-4 shows the required IP addresses for a clustered IBM System Storage TS7600 with ProtecTIER.

*Table A-4 Factory-default server IP addresses for a two-node clustered TS7650 or TS7650G*

IBM System Storage TS7600 with ProtecTIER two-node cluster	Component	Port	IP address
Node A: the server located in the lower part of the rack	Server	eth0	External customer IP
	<b>Note:</b> By default, the TS7650 and TS7650G use the IP address range 10.0.0.50 through 10.0.0.59 for the power control network. The server IP addresses do not change from frame to frame.		
	Server	eth1	10.0.0.51
	Server	eth2	10.0.0.51
	Server	eth3	N/A
	Server	eth4	N/A
	Server	eth5	172.31.1.xx
Network power switch	N/A	10.0.0.50	

IBM System Storage TS7600 with ProtecTIER two-node cluster	Component	Port	IP address
Node B: the server located in the upper part of the rack	Server	eth0	External customer IP
	<b>Note:</b> By default, the TS7650 and TS7650G use the IP address range 10.0.0.50 through 10.0.0.59 for the power control network. The server IP addresses do not change from frame to frame.		
	Server	eth1	10.0.0.52
	Server	eth2	10.0.0.52
	Server	eth3	N/A
	Server	eth4	N/A
	Server	eth5	172.31.1.xx
	Network Power Switch	N/A	10.0.0.50

**Note:** The IP addresses for the internal Network Power Switch are automatically set up during pt\_config.

## Customer IP addresses

Table A-5 shows the customer IP addresses.

Table A-5 Customer IP addresses

Node A (the server located in the lower part of the rack)	Port	IP address	Network address	Default gateway
	eht0			
Node B (the server located in the upper part of the rack)	Port	IP address	Network address	Default gateway
	eht0			

## Customer and replication IP addresses

Table A-6 shows the customer and IP Replication addresses.

Table A-6 Customer and IP replication addresses

The default Gateways for eth3 and eth4 should be different. Otherwise, the VLANs are meaningless. Provide a routing path from the IP address on the eth3-server1 to the IP address of eth3-server2, and a routing path form the IP address on eth4-server1 to the IP address of the eth4-server2.
Source site

The default Gateways for eth3 and eth4 should be different. Otherwise, the VLANs are meaningless. Provide a routing path from the IP address on the eth3-server1 to the IP address of eth3-server2, and a routing path form the IP address on eth4-server1 to the IP address of the eth4-server2.					
Node A (the server located in the lower part of the rack)	<b>Port</b>	<b>IP address</b>	<b>Network mask</b>	<b>Default Gateway</b>	<b>Dedicated VLAN</b>
	eth3				
	eth4				
Node B (the server located in the upper part of the rack)	<b>Port</b>	<b>IP address</b>	<b>Network mask</b>	<b>Default Gateway</b>	<b>Dedicated VLAN</b>
	eth3				
	eth4				
Destination Site					
Node A (the server located in the lower part of the rack)	<b>Port</b>	<b>IP address</b>	<b>Network Mask</b>	<b>Default Gateway</b>	<b>Dedicated VLAN</b>
	eth3				
	eth4				
Node B (the server located in the upper part of the rack)	<b>Port</b>	<b>IP address</b>	<b>Network Mask</b>	<b>Default Gateway</b>	<b>Dedicated VLAN</b>

A virtual LAN, commonly known as a VLAN, is a group of hosts with a common set of requirements that communicate as though they were attached to the broadcast domain, regardless of their physical location. A VLAN has the same attributes as a physical LAN, but it allows for end stations to be grouped together even if they are not located on the same network switch. Network reconfiguration can be done through software instead of physically relocating devices.

## Host names and DNS settings

Table A-7 shows the host names and DNS settings worksheet.

Table A-7 Host names and DNS settings

Item or setting	Instructions	eth0	eth1 (if applicable)
Source host name	Record the console or host name that you want to assign to the management console workstation (for example, ezr1). The console name and the domain are used to identify the IBM System Storage TS7600 with ProtecTIER to the network.	IP address #1 (client)	IP address #1 (client)
		IP address #2 (service)	IP address #2 (service)

Item or setting	Instructions	eth0	eht1 (if applicable)
Domain name	Provide the domain name that you are assigning to the TSSC (for example, medina.xyz.nl).		
Ethernet settings: Complete the LAN Adapter Details section when the TSSC connects to your LAN.			
Media speed (Ethernet)	Check <b>Auto detection</b> or the media speed of the Ethernet adapter.	<ul style="list-style-type: none"> <li>_ Auto detection</li> <li>_ 10 Mbps Half Duplex</li> <li>_ 10 Mbps Full Duplex</li> <li>_ 100 Mbps Half Duplex</li> <li>_ 100 Mbps Full Duplex</li> <li>_ 1000 Mbps Half Duplex</li> <li>_ 1000 Mbps Full Duplex</li> </ul>	<ul style="list-style-type: none"> <li>_ Auto detection</li> <li>_ 10 Mbps Half Duplex</li> <li>_ 10 Mbps Full Duplex</li> <li>_ 100 Mbps Half Duplex</li> <li>_ 100 Mbps Full Duplex</li> <li>_ 1000 Mbps Half Duplex</li> <li>_ 1000 Mbps Full Duplex</li> </ul>
TCP/IP interface network mask	Record the dotted decimal network mask that you want to apply to the TCP/IP address (for example, 127.123.546.0).		
DNS settings: Complete this section if you plan to use a domain name server (DNS) to resolve network names.			
Name server (DNS) internet address 1	Provide the dotted decimal address of the name server that the TSSC will access (for example, 5.127.42.250).		
Name server domain 1	Provide the domain name of the name server (for example, medina.xyz.nl).		
Name server (DNS) internet address 2 (Optional)	Provide the dotted decimal address of the second name server that this workstation can access (for example, 5.127.42.252). Although this is optional, you can specify a second name server when you configure a backup or a secondary server.		
Name server domain name 2	If you have a second name server, provide the domain name of the second name server (for example, medina2.xyz.nl).		
Routing settings: Complete the following section if you want to specify a default gateway for routing.			
Gateway address	Confirm and record the dotted decimal or symbolic name address of the gateway (for example, 8.127.155.254 or londongate).		

## Replication settings worksheet

Use Table A-8 to determine and record the information that you need to implement replication on your IBM System Storage TS7600 with ProtecTIER.

Table A-8 Replication policy information

Replication policy information			
Policy name			
Policy priority level	Low	Medium	High
Policy enabled Y/N			
Destination repository name			
Destination target	Shelf	Library	
		Library name	
Barcodes for range of cartridges to be replicated			
Replication metadata reserved space			
Replication Management information			





# B

## **Western Telmatic Inc. Network Power Switch**

In this appendix, we discuss the Western Telmatic Inc. (WTI) Network Power Switch. We also describe how the 3958-DD3 and the 3958-AP1 use the Red Hat Cluster Suite for clustering two servers.

The Red Hat Cluster Suite checks the health of the nodes in the cluster and prevents data corruption in a clustered environment by *fencing* the peer in the cluster using the WTI Network Power Switch. For example, if there are two nodes in a cluster, node 1 and node 2, and node 2 stops responding to Red Hat Cluster Suite heartbeat messages, the Red Hat Cluster Suite running on node 1 automatically opens a telnet connection to the WTI Network Power Switch and toggles the power supplies connected to node 2. This action results in node 2 rebooting, preventing node 2 from modifying data on the shared storage when it cannot synchronize the modifications with its peer (Figure B-1).

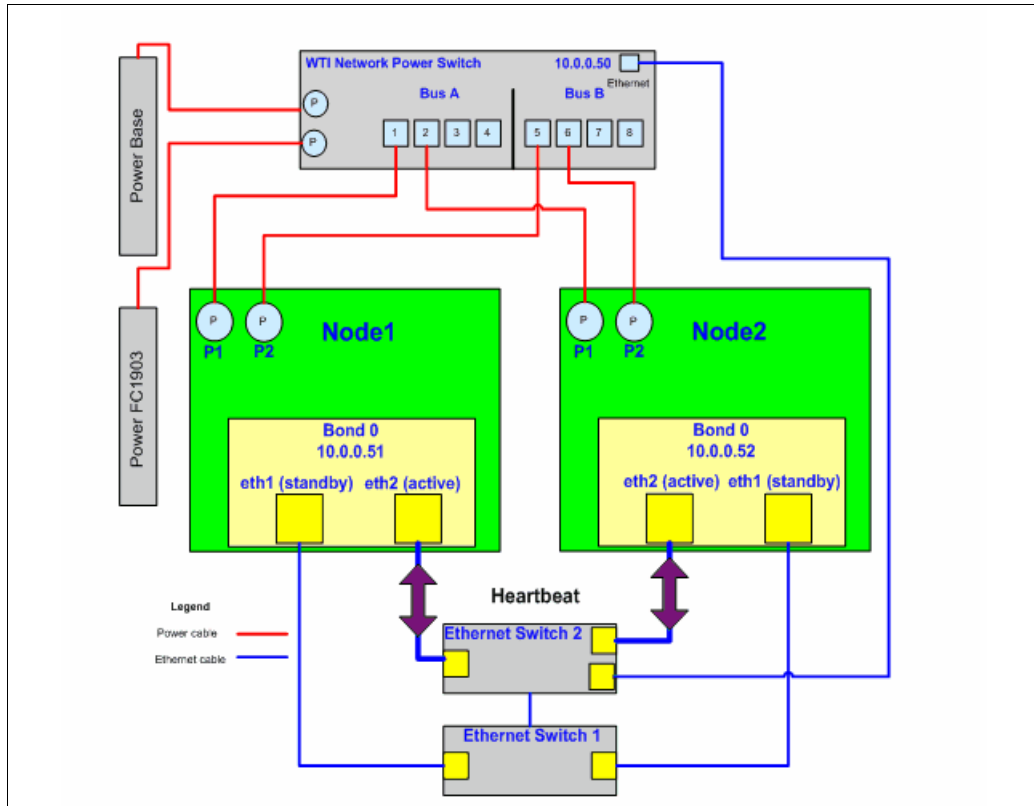


Figure B-1 Power fencing

If one of the nodes is fenced, the view pane of the ProtectTIER Manager also shows the fenced status (Figure B-2).

Cluster members							Options ▾
IP address	DNS	GUI proxy	Status	Management service	VT		
9.11.201.19	Naples	✓	● Ok	● Online	● Online		
9.11.200.233	Italy		● Fenced	● Offline	● Offline		

Figure B-2 Fenced status



In the rare case where the fenced status is not changing automatically from Fenced to OK after the reboot of the node, check the status of the ports of the WTI switch. Use a program (for example, PuTTY if using Windows) or else you can use SSH natively on UNIX systems to make an SSH connection to one of the ProtecTIER nodes. After you are connected to your node, open a telnet connection to the WTI switch. The default IP address is 10.0.0.50 and the default password is *password*. All ports should have the status ON (Figure B-3).

```
[root@Austin ~]# telnet 10.0.0.50
Trying 10.0.0.50...
Connected to 10.0.0.50 (10.0.0.50).
Escape character is '^]'.

Enter Password: *****

Internet Power Switch v1.41h   Site ID: (undefined)

Plug | Name           | Password | Status | Boot/Seq. Delay | Default |
-----+-----+-----+-----+-----+-----+
 1 | internal_node1 | (undefined) | ON | 5 Secs | ON |
 2 | internal_node2 | (undefined) | ON | 5 Secs | ON |
 3 | (undefined)   | (undefined) | ON | 0.5 Secs | ON |
 4 | (undefined)   | (undefined) | ON | 0.5 Secs | ON |
 5 | internal_node1 | (undefined) | ON | 5 Secs | ON |
 6 | internal_node2 | (undefined) | ON | 5 Secs | ON |
 7 | (undefined)   | (undefined) | ON | 0.5 Secs | ON |
 8 | (undefined)   | (undefined) | ON | 0.5 Secs | ON |
-----+-----+-----+-----+-----+

"/H" for help.

IPS>
```

Figure B-3 WTI power switch status

Figure B-3 on page 699 shows the status of all the power ports. Port 1 and port 5 are connected to node 1 and port 2 and port 6 are connected to node 2. A useful function is the /H function, which shows you all the options that you can do on the WTI switch, as shown in Figure B-4.

```

"/H" for help.

IPS> /h

Internet Power Switch v1.41h   Site ID: (undefined)

Display                               Configuration
/H      Display Help Screen         /G      View/Set General Parameters
/S      Display Plug Status          /P [n]  View/Set Plug Parameters
/SN     Display Network Status       /C      View/Set Serial Parameters
                                               /N      View/Set Network Parameters
Control
/D      Set Plugs to Default         /T      View/set Telnet Parameters
/Boot <n> Boot Plug n                /W      View/Set Web Server
/On <n>  Turn On Plug n              /E      Save Parameters
/Off <n> Turn Off Plug n             /R      Recall Parameters
/X      Exit/Disconnect              /DL     Download Parameters to File

                                               Utilities
+-----+                               /I      Reset Network Interface
| [n] = optional plug name or number |   /U      Upgrade Firmware
| <n> = required plug name or number |
| n+n or n n = plug n and plug n   |
| n:n = plug n through plug n     |
| * = all plugs                    |
| ,y = bypass "Sure? (y/n)"       |
+-----+

IPS>

```

Figure B-4 Help function for WTI switch

In this example, we show how to set the a port from status off to status on. Figure B-5 shows that node 1 port (plug) number 1 has the status OFF.

```

Internet Power Switch v1.41h   Site ID: (undefined)

Plug | Name           | Password | Status | Boot/Seq. Delay | Default |
-----+-----+-----+-----+-----+-----+
 1 | internal_node1 | (undefined) | OFF | 5 Secs | ON |
 2 | internal_node2 | (undefined) | ON | 5 Secs | ON |
 3 | (undefined) | (undefined) | ON | 0.5 Secs | ON |
 4 | (undefined) | (undefined) | ON | 0.5 Secs | ON |
 5 | internal_node1 | (undefined) | ON | 5 Secs | ON |
 6 | internal_node2 | (undefined) | ON | 5 Secs | ON |
 7 | (undefined) | (undefined) | ON | 0.5 Secs | ON |
 8 | (undefined) | (undefined) | ON | 0.5 Secs | ON |
-----+-----+-----+-----+-----+

"/H" for help.

IPS>

```

Figure B-5 WTI port 1 is OFF

From the CLI, enter / on one on the command-line prompt (Figure B-6) and press Enter. You will see the question Sure? (Y/N). Answer that question with Y and the port status will be changed from OFF to ON.

```
Escape character is '^]'.
Enter Password: *****
Internet Power Switch v1.41h   Site ID: (undefined)

Plug | Name           | Password | Status | Boot/Seq. Delay | Default |
-----+-----+-----+-----+-----+-----+
 1 | internal_node1 | (undefined) | OFF | 5 Secs | ON |
 2 | internal_node2 | (undefined) | ON | 5 Secs | ON |
 3 | (undefined) | (undefined) | ON | 0.5 Secs | ON |
 4 | (undefined) | (undefined) | ON | 0.5 Secs | ON |
 5 | internal_node1 | (undefined) | ON | 5 Secs | ON |
 6 | internal_node2 | (undefined) | ON | 5 Secs | ON |
 7 | (undefined) | (undefined) | ON | 0.5 Secs | ON |
 8 | (undefined) | (undefined) | ON | 0.5 Secs | ON |
-----+-----+-----+-----+-----+

"/H" for help.
IPS> /on 1
Plugs to be turned on:
Plug 1: internal_node1
Sure? (Y/N): 
```

Figure B-6 Changing the port status

Figure B-7 displays the status of the WTI after turning on port 1.

```
[root@Austin ~]# telnet 10.0.0.50
Trying 10.0.0.50...
Connected to 10.0.0.50 (10.0.0.50).
Escape character is '^]'.
Enter Password: *****
Internet Power Switch v1.41h   Site ID: (undefined)

Plug | Name           | Password | Status | Boot/Seq. Delay | Default |
-----+-----+-----+-----+-----+
 1 | internal_node1 | (undefined) | ON | 5 Secs | ON |
 2 | internal_node2 | (undefined) | ON | 5 Secs | ON |
 3 | (undefined) | (undefined) | ON | 0.5 Secs | ON |
 4 | (undefined) | (undefined) | ON | 0.5 Secs | ON |
 5 | internal_node1 | (undefined) | ON | 5 Secs | ON |
 6 | internal_node2 | (undefined) | ON | 5 Secs | ON |
 7 | (undefined) | (undefined) | ON | 0.5 Secs | ON |
 8 | (undefined) | (undefined) | ON | 0.5 Secs | ON |
-----+-----+-----+-----+

"/H" for help.
IPS>
```

Figure B-7 WTI status



# 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

You can search for, view, or download books, papers, Technotes, draft publications and additional materials, as well as order hardcopy Redbooks, at the IBM Redbooks website:

[ibm.com/redbooks](http://ibm.com/redbooks)

Note that some of the documents referenced here might be available in softcopy only.

- ▶ *IBM Midrange System Storage Implementation and Best Practices Guide*, SG24-6363
- ▶ *IBM System Storage DS8000: Architecture and Implementation*, SG24-6786
- ▶ *IBM System Storage DS8000: Copy Services in Open Environments*, SG24-6788
- ▶ *IBM System Storage Tape Library Guide for Open Systems*, SG24-5946
- ▶ *Implementing IBM Tape in Linux and Windows*, SG24-6268
- ▶ *Implementing IBM Tape in UNIX Systems*, SG24-6502

## Other publications

These publications are also relevant as further information sources:

- ▶ *IBM System Storage ProtecTIER Software Upgrade and Replication Enablement Guide*, GC53-1196
- ▶ *IBM System Storage ProtecTIER User's Guide for Enterprise Edition and Appliance Edition*, IBM form number GC53-1156
- ▶ *IBM System Storage TS7600 with ProtecTIER Installation Guide*, GC53-115
- ▶ *IBM System Storage TS7600 with ProtecTIER Installation Roadmap Guide for the TS7650 (3958 AP1)*, GC53-1194
- ▶ *IBM System Storage TS7600 with ProtecTIER Installation Roadmap Guide for the TS7650G (3958 DD3)*, GC53-1154
- ▶ *IBM System Storage TS7600 with ProtecTIER Introduction and Planning Guide*, GC53-1152
- ▶ *IBM System Storage TS7600 with ProtecTIER Introduction and Planning Guide for the TS7650 (3958 AP1)*, GC53-1193
- ▶ *IBM System Storage TS7600 with ProtecTIER Introduction and Planning Guide for the TS7650G (3958 DD3)*, GC53-1152
- ▶ *IBM System Storage TS7600 with ProtecTIER Problem Determination Guide*, GC53-1157
- ▶ *IBM Tape Device Drivers: Installation and User's Guide*, GC27-2130
- ▶ *IBM Tape Device Drivers Programming Reference*, GA32-0566

## 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 Interoperability Guide:  
<http://www.ibm.com/systems/support/storage/config/ssic/index.jsp>
- ▶ IBM Support website:  
<http://www.ibm.com/systems/support/storage/config/ssic/index.jsp>
- ▶ 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/products/overview.jsp?pcid=1018&pvid=2_1)  
<http://www.symantec.com/enterprise/support/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>
- ▶ XIV:  
<http://www.xivstorage.com/>

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks, at this website:

[ibm.com/redbooks](http://ibm.com/redbooks)

## Help from IBM

IBM Support and downloads

[ibm.com/support](http://ibm.com/support)

IBM Global Services

[ibm.com/services](http://ibm.com/services)

# Index

## Numerics

- 1U 30, 131
- 2 rack unit (2U) 122
- 220V 162
  - power distribution units 162
- 36U 142
- 3850 X5 Gateway 19
- 3958 DD1 18
- 3958 DD3 18
- 3958 DD4 17
- 3958 Model DD3. 131–132, 135
- 3958-AP1 43, 121–128, 131–133, 139–140, 156–159, 162, 166–168, 435, 697
  - front-end ports 167
- 3958-DD1 xx, 19–20, 43–44, 123, 142, 147–148, 162, 219, 255
- 3958-DD3 xx, 19–20, 122–126, 128–129, 131, 140–143, 148, 150, 153, 156–163, 166–168, 219, 425, 435, 690, 697
  - nodes 148
- 3958-DD4 18–20, 121–124, 127–129, 131, 135, 137, 140–142, 144–147, 151–155, 157–159, 161–164, 166–168, 435
- 3958G-DD3 43
- 3958-SM1 125
  - features 124
- 3959-SM1 29, 122
- 4X octo-cores 123
- 5639-XXB 137–138, 159

## A

- abstracting 74
- AC power to switch 166
- access alerts 513
- access events 515
- active blocks 35
- ACTIVE Tivoli Storage Manager pool 351
- active-active 128
- Add cartridges option 49, 501
- Add cartridges wizard 501, 503
- add or remove a grid member 468
- AddCartridges 675
- AddClustermember 675
- added value of visibility switching 608
- Adding an LSU window 248, 443
- AddLibrary 675, 678
- administrator email address 689
- AIX 355
- Alerts button 513, 624–625
- allocatable space 35
- Alloctable repository capacity total 637–638
- analog phone line 50
- ANALYZE\_SESSIONS utility 651
- Appliance xx, 11–12, 17, 26–31, 43, 55–56, 82, 90,

- 124–126, 128, 132, 135, 139, 159, 222, 234, 255, 420, 437, 593, 605–606
- Appliance Server 132
- application groups
  - creating 466
- Architect NBU 391, 591
- asynchronous data replication 397
- Atape drivers 298–299
- audit/fix 380–381
- automation of daily operation 571
- available models 27

## B

- back-end physical disks 345
- backup application catalogs
  - recovering 347
- backup applications 18, 21, 24, 39, 54, 57, 60, 73, 77, 116–117, 119, 171, 230, 235, 547–548, 561, 638, 640, 642, 644, 648, 653
  - data structure 7
  - eject and inject commands 383, 395
  - multiplexing features 73
  - read command 39
  - standpoint 412
  - tape cartridges 39
- Backup interface field 241, 439
- backup servers 14, 17, 31, 36, 69, 102, 115, 149, 156–158, 240, 276, 280–281, 314, 331, 344–346, 350, 383, 408, 418, 421, 434, 586, 609, 611, 616–617, 646, 670
  - considerations 344
- backup streams
  - multiple 72
- Backup, Recovery, and Media Services 331, 400
- backup/recovery throughput 55
- backups 4, 38, 414
  - administrator 115–119
    - intuitive way 117
  - architecture 57
  - block-level database 345
  - cross-site 391, 591–592
  - cumulative 390, 591
  - data 5, 7, 10, 35, 54, 58, 74, 651
  - database 73, 381, 408
  - default 603
  - differential 60, 390, 591
  - differential incremental 4
  - environment 38
  - existing infrastructure 57
  - first baseline 74
  - full 4, 58–60, 78–80, 118, 347–349, 382, 390, 394, 586, 589, 595, 669
    - image 348–349, 586, 669
  - full-plus-incremental 350

- incremental 4, 31, 59–60, 78–80, 350, 389–390, 406, 591
- LAN-free 57
  - to disk with ProtecTIER 385
- monthly full 60
- native LAN-free 385
- policies 5, 41, 56, 60, 69, 77–78, 91, 217, 391–392, 592
  - factoring ratio 77
- sessions 276
- storage pools 387
- Symantec NetBackup (NBU) catalog 391
- weekly 59
- backup-to-disk 13, 36, 102, 115, 240, 421, 434
- bandwidth validation tool 108
- bandwidth 413
- Barcode field 650
- Barcode ranges window 269, 583
- Barcode seed field 236
- barcode seeds 502
- barcodes 113, 348, 647
- Basic DR 608
- basic DR 608
- basic input/output system (BIOS) 161
- beginning of tape (BOT) 39
- bidirectional 273, 423, 466, 471
- big files 73
- binary differential 8, 42
- bit-by-bit comparison 7
- block groups 35
- BRMS
  - setting up for backups to ProtecTIER 401
- built-in compression 347
- byte-level 40
  - changes 8

## C

- cache controllers 69, 162
- cache frames 141–142, 147–148, 152, 162
- cache modules 162
- cache size 59
- call home 131
- capacity fluctuations 118
  - managing 117
- capacity management 116, 118
  - implications 118–119
- Capacity pane 629, 637, 639
- capacity planning 59, 77, 142, 148
- capacity savings 38
- capacity sizing 54, 59, 75–76, 100
- Cartridge info option 675
- Cartridge integrity verification function 561, 664
- Cartridge replication menu 50
- Cartridge replication statistics menu 50
- cartridges 415
  - adding 500, 503
  - capacity 117
  - deleting 504
  - ejection 392
  - injecting 392

- level 392, 411
  - maximum number 502
  - metadata verification 557, 664
  - nominal capacities 117
  - number 117–118, 384
  - planning 89
  - relocating inside the virtual library 550
  - size 32, 118–119, 236
  - slot 504
  - sync point 350, 587, 669
  - unique ID 348
  - unloading 548
  - virtual 21, 24, 28, 34–35, 39, 48–49, 124, 126–127, 236, 415, 502, 504, 638, 646
    - maximum possible size 236
    - prior use 35
- Cartridges integrity verification menu 50
- Cartridges ownership takeover menu 50
- Cartridges tab 505, 647
- catalog 110, 113–114, 117, 344, 347–350, 381–382, 391–395, 399, 412, 418–419, 571, 587, 592–596, 608–609, 612–616, 669–670
- CDI
  - disabling 398
- chain 386
- Change dimensions menu 48, 488
- Change dimensions wizard 487–488
- change rate 60, 77, 119, 652–653
- checkin libvolume command 383
- checking the system 555
- CHECKOUT LIBVOLUME 365
- CHECKOUT LIBVOLUME command 365
- checkout libvolume command 383
- chunks 5–6, 8, 40
- classification of clients 390–391
- Clear Alerts 514, 625
- CLI 383, 395, 595, 674, 701
- client compression 351, 396, 398
- Client Node Configuration Wizard 375
- Client Node Configuration Wizard GUI 374
- clone-to-tape and recovery operation 614
- cloning 112, 412, 419, 571, 608–609, 614, 616
- Cluster Connection 136
- Cluster Connection Kit 133, 136, 148
- cluster internal network 149
- Cluster Management 525
- Cluster Management menu 532
- cluster members 194, 473, 524, 629
  - adding 524
- Cluster Members pane 630
- Cluster Members section 630
- clustered configuration 20–22, 25, 28, 68, 128, 133, 136, 141–142, 148, 165, 211, 219
  - two-node 21–22, 142, 148
- clustering 591
- clustering support 18, 124, 126–127
- clusters
  - single node 219, 524
  - two-node 20, 27
- CM Virtual Tape Facility Daemon 46



- collocation 384
- command queuing 81–82
- common disk storage pools 139
- common ID 16, 45, 425, 472
- complex fabrics 276
- compliance 10
- compression 38, 42, 73, 78, 80, 345
  - disabling 543
  - standard 38, 345
- Compression field 248, 250–251, 443, 445
- compression rate 80
- CompressionRate 80
- conceptualization 36
- Configure Disk Storage Servers icon 453
- connectivity prerequisites 159
- content aware 5, 7
- contiguous blocks 35
- control path failover 25, 296, 313, 323
- copy groups 365
- corrupted NICs 450
- CPF 25, 48, 303, 313
- CPUs 9
- Create new library menu 48, 230
- Create new library wizard 230, 238
- Create Repository option 676
- Create repository planning menu 216
- Create repository planning option 75
- Create repository planning wizard 216, 222
- Create repository wizard
  - summary report 241
- cross-site 608
- CSV files 348, 653, 666
- customer environment 57
- customer information work sheet 688
- customer installation responsibilities 161, 688
- customer profiles for open systems 55

**D**

- D2D2T 56–57
- daily network workload 94
- data
  - change
    - rate 4, 41, 76–77, 80, 117–119
    - rate decrease 77
  - element 4
  - landscape 37
  - online 10
  - proliferation 56
  - protection 18
  - protection environment 75
  - replication 413
  - restoration 414
  - retention 4
  - saving and printing 195
  - segments 109
  - type 59–60, 69, 77–78
  - virtual 20, 27
- data collision 14, 273, 471
- data deduplication 1, 4–5, 9–10, 29, 37–38, 51, 56, 124, 126–127

- basic concepts 1
- benefits 10
- components 9
- concept 3
- cycle 9
- hash-based 5–6, 40
- IBM solution 1
- method 3
- process 651
- processing 8
- ratio 4–5, 8
- software 10
- technology 26
- data elements
  - similar 8, 40, 42
- data factoring 40, 248, 251, 443, 445
- data flows 41–42, 316, 427, 469
- data path failover 297, 313, 326
- Data Protection Survey 59, 77–78
- data retention period 41
- data scrubbing 42
- data sharing 30
- data streams 7, 9, 42, 89, 99, 112, 119, 252, 345–346, 384, 672–673
- data transfer 36, 89–90, 92, 94, 102, 115, 162, 240–241, 276, 376, 385, 416–417, 421, 428, 434, 436, 637, 644
- data type 60, 72
- data-matching rates 345
- DDMs 26, 222
- Dedicated Replication Timeframe window 110
- Dedicated volume pools 391, 592
- deduplication 1, 3–10, 12, 17–18, 20–21, 26–27, 33, 37–38, 40–42, 51, 53–54, 56, 58, 89, 91–94, 105, 124, 126–128, 138–139, 250, 254–256, 345–347, 352, 383, 389, 392, 400, 409, 413, 428, 443–444, 465, 592, 637, 640, 651, 660, 676
  - assumptions 78
  - rate consideration 296
  - ratio 4–5, 8, 42, 54–55, 642
  - server 10
  - software 10
  - solution 57
  - technology 17
- Deduplication Gateway 18, 128, 131
- default user name 193, 620
- defragment 35
- defragmentation 35
  - disabling 542
- Delete cartridges option 50
- Delete library menu 48, 296, 511
- delta data 42
- deployment 20, 27, 31
  - for two systems using replication 411
  - planning guidelines 105
- Description field 248, 250, 443–444
- device drivers 161, 297–298, 309, 328, 396
- DHF 241, 436
- Digital Linear Tape (DLT) 57
- DIMMs 31, 125–126, 129
- dirty bit attribute 586

- disaster recovery 15, 123, 393, 575, 580, 588–590, 594
  - cartridges
    - restoring 595
  - high availability 57
  - inventory command options 588
  - mode 575–579, 588, 596–598, 601, 606, 608, 612–614
    - completing failback and leaving 596
  - sites 58, 89, 101, 105, 108, 110, 112–114, 254, 264, 382, 391–395, 397–400, 410–411, 413, 416–417, 419, 575–577, 579, 589, 591–595, 599, 605, 608–609, 611, 614
  - solutions 57
  - tapes 350, 670
  - test operation in a multiple (two) domain backup environment 611
- disk arrays 8, 10, 14–15, 20, 42, 59, 70–72, 128, 141–143, 145–146, 149–150, 153, 157, 162, 166–168, 171, 209, 211, 215, 219
  - Ethernet connections 145
  - existing data 211, 215
- disk controllers 26, 122, 134
- disk drives 82, 85, 143, 150
  - modules 26, 222
  - SATA 82
- disk expansion modules 26, 122
- disk pools
  - configuration 462
  - creating 458, 462
- disk repositories 42
- disk size for TS7610 82
- disk space 15, 70, 223
- disk storage 10, 76
- Disk Storage Pools 74, 346–347
- disk subsystems 38
- disk-based data storage system 13
- disk-based targets 18
- DLT7000 351, 383
- DNS 630, 693–694
- domain name servers (DNS) 630
- DriveModels 675–676, 678, 682
- drives
  - resetting 547
  - speed 81
- Drives tab 547–548, 550, 646, 648
- DS4700 142, 147
  - disk arrays 146, 154
  - storage 122
- DS5020 cache subsystem 126
- DTC VTF 18, 138–139, 156
- dual in-line memory modules 31
- dual-port FC expansion 164–165
- duplicate blocks 345–346
- duplicate data 4, 8, 18, 40

## E

- early warning (EW) 116–118, 648
  - signal 116
- effective capacity 37
- effective data transfer rate 276

- EIA-310-D standards compliance 141
- Eject to shelf menu 50
- electronic media 161
- element numbers 279, 354, 357, 360, 363, 367, 370
- EMC Legato NetWorker bootstrap 398
- EMC NetWorker 396–397
  - implementation steps 397
- emulation 18, 26, 124, 126–127, 138–139
- Emulex LPe12002 126, 129, 155
- Enable Unique LUN descriptor option 48
- encryption 299, 346
  - features 73, 346
- end of cartridge signal 117
- end time 653
- enterprise level data protection 125–126
- enterprise scale in-band factoring 68
- environment 60
- environment configuration survey 58
- environmental costs 10
- equilibrium state 116
- Eth0 network interface 149
- Eth1 19, 149, 636
- Ethernet 13, 19, 29, 31, 44, 46, 92, 94, 125–127, 129, 131–133, 135–136, 145–146, 153–154, 157, 160, 162, 166, 428, 570, 694
  - adapters 155
  - connections 145, 153
  - interfaces 438
  - ports 44
  - switches 22, 29, 128, 142, 147–148, 161, 171
- Events Log 515, 624–625
- exact chunk 40
- Excel 348–349, 653
- EXP520 142, 148
- EXP810 expansions 148
- expansion 128
  - cards 165
  - enclosures 143, 150

## F

- fabric 31, 33, 59, 156, 162, 165–168, 276, 279
  - switches 69, 128, 166
- factoring process 57
- factoring ratio 15, 34, 41, 54, 72–73, 76–78, 80, 117–119, 217, 222–223, 344, 392, 398, 482, 592, 639
- failback procedure 393, 579, 586, 593
- FC3447 133
- feature codes 129, 131
  - 3958-SM1 131
- fenced status 699
- fencing 171, 698
- Fibre cables 166
- Fibre Channel (FC) 15, 26, 56, 58, 81–82, 128, 133, 143, 149–150, 156, 161–162, 164, 166, 168, 171, 217, 222, 279, 296, 329–332, 400, 552, 629, 631, 633–634, 636, 643, 688
  - arbitrated-loop 162, 553, 634, 643
  - cables 162
  - connections 143, 150
  - connectivity 144, 151

- director 162
- disks 82
- expansion 164
- ports 18, 124, 126–127, 149, 156
  - back-end 149
  - switches 159
- file match 7
- file systems 15, 35, 55, 59, 86, 168, 190, 208, 211, 215, 222–223, 476, 479, 483, 518, 638
  - current usage type 638
  - expanding existing 476
  - prerequisites 159
  - remounting 481
  - shared 148
- filemarks 557
- files
  - small 60, 74
- firewalls 157
- firmware 128, 166
- fragmentation 9
- Fragmented repository capacity total 637–638
- frames 16, 30, 91, 94, 101, 108–110, 112, 122, 128, 134, 136, 139, 141–142, 147–148, 152, 154, 161–162, 255, 264, 268, 271–273, 344, 378, 392, 415, 418, 430, 592, 606, 690–692
  - configuration 146
- free blocks 35
- front-end connections 14
- front-end ports 163, 167
- fsCreate parameters 215
- fstab 168
- FullCapacity 78, 80
- FullCapacityVersions 78, 80
- FullChangeRate 80
- FullFrequency 78–79
- FullPhysicalCapacity 79–80
- FullRetention 78–79

## G

- Gateway hardware 159
- Gateway servers 128, 145–146, 148, 153–154, 159, 163, 171
- general plug-in configuration 452
- General tab 494, 626, 628–631, 642–645
- GFS 34, 209–210, 215, 638, 676
- GIF images 345
- GIF reader 345
- Global File System (GFS) 34, 159, 638
- global replication time frames 271
- GO BRMS 401
- government regulations 10
- graphical user interface (GUI) 13, 46, 55, 140, 146, 154, 179, 189–190, 194, 199–200, 264, 281–282, 284–285, 296, 301, 334, 339, 374, 383, 395, 411, 430, 433, 452, 518, 561, 596, 613, 630, 660, 665
- Grid Manager 45
  - log 516
- grid topology 46, 253

## H

- hard disk drives (HDD) 31, 59, 88, 127, 129, 133–134
- hardware 30, 43, 121, 124, 140, 158–159, 162, 175, 283, 311, 313, 339–340, 435, 452, 513
- hash 5
  - collision 6
  - matches 6
- hash-based techniques 8, 40
- hashing 5
  - algorithm 5–6
- HDS AMS1000 20
- heartbeat 698
- high availability 25
- high redundancy 156, 167
- High water mark setting 461
- high-speed transfer 276
- Host Bus Adapter (HBA) 19, 58, 69, 125, 128, 131, 133, 136, 157, 162, 167, 276–277, 279–282, 314, 322–323, 325, 327, 329–330, 339, 345–346, 435
  - WWN 276
- host configuration 437
- Host Initiators 288–291, 293, 295
- HP-UX 11.x 158
- Hub Mesh Group 424
- hubs 16, 103, 105–108, 257, 421, 426, 575, 587, 599
  - total capacity 107
- HyperFactor 5, 7–8, 11, 26, 40–42, 48, 54, 92, 124, 126–127, 138–139, 143, 236, 248, 250, 345–346, 443–444, 502, 544, 620, 639–641, 651
  - data 70
  - data deduplication 8
  - mode 48
  - ratio 236, 502, 641, 647
  - technology 38

## I

- I/O 35
  - disk 40, 81, 276, 279–280
  - operation 561, 636
  - pattern 72
- IBM 3958-DD3 128
- IBM AIX 277, 282, 296–301, 305, 352, 356–357, 396, 435, 447, 450
- IBM AIX 5L 158
- IBM DB2 72, 345, 396
- IBM i Navigator 401
- IBM i V5R4 158
- IBM Informix 72, 345
- IBM Interoperability Matrix 314, 331, 344
- IBM Lotus Domino 60
  - email 72
- IBM MIB definition file 173
- IBM Path Failover 124, 126, 138–139
- IBM Software Technical Document 402
- IBM solutions 38
- IBM Storage Controller 26
- IBM Storwize 345
- IBM Support 48, 191, 485, 517, 519, 521–522, 535, 544, 546–548, 556–557, 654, 656

- IBM System i 400, 609
  - BRMS 609
- IBM System Service Representative (SSR) 121–122, 131, 160–161, 171, 688, 690
  - installation responsibilities 160
  - tasks 160
- IBM System Storage DS4000 122
- IBM System Storage DS4700
  - disk arrays 141
  - dual-controller 148
- IBM System Storage DS8300 142, 147
- IBM System Storage ProtecTIER Enterprise Edition V2.3 124–125, 127
- IBM System Storage TS7600 with ProtecTIER
  - management 118
- IBM System Storage TS7600 with ProtecTIER software 138
- IBM System Storage TS7650 ProtecTIER Deduplication Appliance 17, 26–27
- IBM System x servers 26, 131, 145, 153
- IBM System x3850 M2 server 128
- IBM System x3850 M2 Type 7141 18
- IBM System x3850 M2 Type 7145-PBR 128
- IBM System x3850 M2 Type 7233 14, 18
- IBM System x3850X5 129
- IBM Systems Director Navigator for IBM i 401
- IBM Tape Device Drivers 297–298, 311, 313–314, 316, 319, 325, 390, 397
  - determining the device names 367
  - device names 367
- IBM Tivoli Storage Manager 5, 60, 159, 232, 280, 312, 350–355, 360–368, 370, 373, 375, 377–386, 418, 609
  - Backup Archive Client 375
  - client compression 383
  - configuration 370
  - database 378, 380
  - databases 351–352, 378–380, 382
  - disk pools 383
  - environment 351, 380
  - servers 351
  - users 60
  - V5.3 351
- IBM Tivoli Storage Manager in SAN environments 385
- IBM Tivoli Storage Manager for SAN clients 385
- IBM TotalStorage DS4000 EXP810 Expansion Enclosures 142
- IBM TotalStorage DS4000 EXP810 Expansions 148
- IBM TotalStorage Service Console 50
- IBM Ultrium 3 tape drives 18, 124, 126–127
- implementation worksheets 158
- import/export (I/E) slots 32, 353, 415, 492, 645, 649
- Import/Exports tab 47, 649
- incremental
  - disk cost 78
  - physical capacity 80
  - report 518
  - versions 78
- IncrementalCapacity 79
- IncrementalCapacityVersions 79–80
- IncrementalFrequency 79

- IncrementalPhysicalCapacity 80
- IncrementalRetention 79
- index save 42
- inject
  - commands
    - scripting 572
- inline data 18, 124, 126–127
  - deduplication 9
- inline method 9
- inline performance 7
- input/output 9
- InstallAnywhere 446
- installation
  - on Windows 175
  - planning 158, 256
  - tasks 121
  - using SMIT 301
  - worksheets 158
- in-sync 586–587
- Integrated Management Module (IMM) 127, 129, 145, 153, 161
- Interoperability 314, 331, 344
- Inter-Switch Links (ISL) 69, 167, 276
- IOPS 35, 81–82, 276, 280
- iostat 74
- IP Address column value 636
- IP Address tab 603–605
- IP addresses 140, 158–159, 161, 172–173, 191, 526, 626, 688
  - customer 692
  - customer and replication 692
  - defining 172
  - worksheet 691
- IP traffic
  - routing 438
- iperf 95–96, 98, 670–672
- iSCSI technology 168
- ISLs 276, 280
- ISVs 159

## K

- Keyboard/Video/Monitor (KVM) 134, 136, 142, 145–147, 153–154, 161, 171

## L

- Lab Services 160, 222
- LABEL LIBVOLUME 364–365
- LAN 57, 146, 154, 158, 383, 385, 389, 396–397, 693–694
- LAN-free clients 280, 383
- latency 97, 99, 671–672
- learning algorithm 117
- Legato Backup Server 397
- Legato disaster recover procedures at the DR site 398
- Legato NetWorker 396–399, 609
- Lempel-Ziv (LZ) 4, 42
- Lempel-Ziv-Haruyasu (LZH) 42
- libraries
  - definition 123, 353–354, 367

- deleting 296, 510
- emulation 25, 353
- physical 34
- related operations 296
- renaming 509
- shared 241, 436
- slots 114, 353, 649
- Libraries Configuration pane 632
- Libraries option 676, 679
  - output 675–677
- Libraries Performance pane 632
- Library front-end pane 631, 643
- Library type pane 645
- Library window 365, 509, 631, 642–643, 645–647
- LibraryInfo 676, 679
- Librarytypes 676
- Light Speed 5
- Linear Tape Open (LTO) 57
- Linux 140
  - Red Hat 175
- Linux 5.2 Advanced Platform 159
- load balancing 22–23, 29, 167, 234, 280, 305, 326, 333, 435, 438
- Load Balancing (LB) method 437
- logical storage unit (LSU) 246–251, 422, 433–434, 441–446, 460
  - adding 247, 442
  - configuration 246
    - managing 250, 445
    - modifying 249, 444
  - management 247, 249, 251, 442, 444–445
- logical unit number (LUN) 48, 74, 166–167, 209, 214, 275, 283–286, 293–296, 339, 354, 366–367, 631, 643–644, 646
- logical volumes 209, 476, 479
  - expanding 480
  - GFS file systems 209
  - reactivating 480
- Long-term statistical data report 654
- long-term statistics
  - creating and downloading 654, 665
- loop environment 163
- Low water mark setting 461
- LTO2 326, 353
- LTO3 30, 32, 326, 353–355, 370
- LUN Masking 283–286, 291, 293–296
  - enabling or disabling 283
  - groups
    - adding 284

## M

- M5015 SAS/SATA controller 125–126, 129
- Machine Code 161
- manual intervention 20, 27
- Master Server 590
- Max. growth check box 502
- Maximum concurrent jobs 465
- Maximum concurrent jobs setting 465
- maximum transmission unit (MTU) 637
- MD5 5

- MDADM 162, 167
- Media Manager volumes 395, 596
- media servers 56, 58, 69, 113, 162–165, 167, 389, 438, 452, 590–591
  - maximum number 165
- medium changer devices 309, 316, 325, 370
- Memory Resident 8
- Memory Resident Index 7–8, 40, 42, 143
- Message-Digest Algorithm 5 5
- metadata 9, 14–15, 34–35, 39, 42, 50, 54–55, 70, 168, 222–224, 476, 482–483, 557, 561
  - available file systems 483
  - file 35
  - file systems 15, 42, 70
  - logical unit number (LUN) 35, 72
  - resources dialogue 224, 483
  - window 223–224
- Microsoft Windows 2003 158
- mission-critical data 89
- mode of operation
  - best practices 112
- mount virtual tapes 24
- Move cartridges menu 551, 567, 569
- Move cartridges option 49
- Move cartridges window 551, 569
- multicore virtualization 18, 126–127
- multipath software 162
- multiple domain environment 594
  - master not clustered 394
- multiple domains architecture 393
- multiplexing (MPX) 72, 346, 390–392, 396, 398, 408, 592

## N

- NAS 56, 74, 346–347
- native Windows drivers 351, 383, 396
- navigation 192
- Navigation pane 192, 196–197, 530, 542, 626, 628, 630, 633
  - Nodes tab 628
  - Systems tab 630
- NDMP 57, 74, 345–347
- NetBackup Media Server 452
- network
  - connections 145, 153, 161, 424, 452, 472
  - failures 94, 416
  - reconfiguration 693
  - replication performance validation 670
- Network Configuration pane 636
- Network Power Switch 29, 128, 158, 692, 697–698
- Network Replication xx
- NetWorker Client 396
  - file index 398
- NetWorker database
  - replicating (bootstrap) backups 398
- NetWorker Domain 396
- NetWorker Server 396, 398
- NetWorker Storage Node 396
- NIC xx, 450, 636
- Node column 643, 646
- Node information window 474

- Node IP field 526, 533
- Node menu 474, 518, 536, 539, 541, 545–546, 548, 654, 665
- Node Selection window 525
- Node selection window 533
- nodes 13–14, 209, 366
  - adding 474
  - rebooting 540
  - redundancy 25
  - redundant 22
  - second 475, 485, 491, 498, 521, 635
  - single 14, 20–21, 27–28, 108, 141, 192
    - configuration 21, 126, 128, 138–139, 141–143, 149
  - subnetworks 474
  - two-node clustered 165
- Nodes pane 191–192, 219, 476, 518, 531, 535, 538, 540, 545–546, 552, 626–627
- NodeVtlStatistics 676, 681
- nominal capacity 15, 76, 115, 117–119
  - certain amount 76
  - system wide shift 117
- Nominal Data Size graph 639–640
- Nominal percentage field 248, 250, 443–444
- nominal throughput 428, 469, 661
- NominalCapacity 78
- normal operation concepts 416
- normal operations
  - returning 605
- NTutil 368, 370
- NUMBER\_DATA\_BUFFER 396
- NumberOfCartridges 677, 680
- nuttcp 95–96, 98, 670–672

## O

- OEM 47
- offline cold catalog backup method 395
- offsite 10
  - storage 10
- OFFSITE\_TAPE 351
- one-time policy 579, 596
- one-unit 131, 145, 153
- online hot catalog backup method 395
- ONSITE TAPE pools 351
- open systems support 56
- OpenStorage (OST) 13, 16, 19, 29–31, 37, 44, 90, 102–103, 115, 124, 126, 128, 130–131, 135, 137, 145–146, 153, 160, 163, 199, 208, 217, 225, 241–243, 245–247, 249, 251–252, 421–424, 433–436, 438–439, 441–442, 444–452, 454, 465–466, 469, 471–472, 590, 592
  - configuration 129, 144, 151–152, 154–155, 420, 452
  - defining a connection 466
  - Ethernet card 435
  - plug-in 433, 438
  - setting up replication 273
- operating systems (OS) 10, 50, 58, 85, 128, 143, 156, 158–159, 162, 175, 195, 346
- Oracle 5, 345, 350, 408
- Oracle Recovery Manager 408

- ostp\_cli command options 450
- ostp\_cli executable 448
- oversubscription 165, 167

## P

- P3000 18, 123, 138–139, 156, 297, 383
  - libraries 351
- parallelism 398
  - scheme 91
- partitions
  - creating 476
- Password field 244, 440–441
- patch cords 162–163, 168
- Path Failover 128, 167
- paths
  - primary and alternate 303, 307, 327
- peer-to-peer (P2P) 634, 643
- per terabyte basis 138–139
- performance 11, 18
  - requirements 57, 72, 101
  - sizing 55, 100
- permission levels 193, 620
- physical capacity 15, 55, 76, 78–80, 88, 117, 119
- physical data-transfer barrier 94
- physical disks 38
- physical tapes 10, 39, 57, 60, 90, 113, 116, 277, 345, 351–352, 355, 377, 384, 389, 391, 396, 398, 408–409, 411–412, 417, 419, 504, 571, 592, 608–609, 614
- physical throughput limit 427, 469
- physical volumes 209, 476, 479
  - adding to a volume group 480
- PhysicalCapacity 79
- point-in-time 349, 381, 391
- policies 5, 16, 41, 76, 89, 91, 113, 124, 127, 138–139, 254, 256, 264, 272, 375, 392–393, 404, 411, 413, 415, 593, 606, 657
  - enabling and disabling 272
  - execution 109
  - fallback
    - creating 580
    - management 413
    - per respository 46
    - running 272
- Port Attributes pane 634
- ports 144, 150, 157, 346, 636
- Positive checks list 556
- post processing 4, 9
- post-processing method 9
  - disadvantages 9
- Power On Self Test (POST) 174
- power outlets 143, 150
- primary (source) repository sizing 106
- principality
  - taking over 599
- principality attribute 579
- problem reports
  - in Grids Management 521
  - in Systems Management 518
- problematic IP 450
- processing options 677

- progressive backup methodology 350
- Progressive Incremental 60
- ProtecTIER 8–9, 15, 17, 26, 34–35, 41, 53–55, 57–58, 69–70, 73, 77, 80, 118, 121–123, 128, 140, 148, 156, 158–162, 165–166, 168, 171, 173–179, 185–188, 190, 192–193, 195, 224–225, 231, 473–474, 476, 481, 485–487, 491, 494, 497, 500, 504, 509, 512, 515, 517, 519, 521, 524, 542–544, 548, 552, 555–557, 561, 573, 619–620, 622–626, 628–633, 635–649, 652–655, 703
  - activity
    - reporting on 619
  - bandwidth 276
  - cluster setup 149
  - clustered systems 43–44, 46
  - clusters 44
  - command-line interface 586–587
  - configuration 437
  - configuring to work with the OpenStorage environment 241
  - connecting hosts 276
  - deduplicated data 20
  - environment 72
  - hardware 414
  - HyperFactor 40–41
  - IBM i
    - prerequisites 330
  - implementation 20
  - installing 241, 436
  - IP replication in the Tivoli Storage Manager environment 352
  - models 12
  - monitoring 619, 623
  - native replication 8, 13–14, 42, 252, 273, 412–415, 471
    - Management Interface 415
    - technology 26
  - nodes 413
  - OST storage appliances 439
  - prerequisites 158
  - recoverable errors 556–557
  - servers 14, 44, 158, 162, 166, 168, 174, 230, 515, 519, 522, 545
    - configuration 54
    - loop technology 162
    - nodes 45
    - ports 167
  - service modes 13
  - software 26, 33–35, 40–42, 50, 158, 619, 622, 630, 635, 638–642, 646, 648, 650
    - servers 42
  - specialist 160
  - specialists 160, 241, 436
  - storage pools 346, 383
  - Symantec NetBackup (NBU)
    - optimal deployments for disaster recovery 592
  - system Replicated Cartridges Status Report 382, 395
  - systems 35, 44, 56–57, 72–74, 77–78, 109, 174, 230, 500, 544, 555–556
  - Tivoli Storage Manager attachment on Windows 2008 366
    - unique I/O pattern 72
    - upgrading the existing system 255
    - versus traditional tape libraries 116
- ProtecTIER Enterprise Edition V2.1 32, 124, 126, 128
- ProtecTIER IP Replication 378, 380
- ProtecTIER Manager 13–14, 22, 24, 29, 31, 46–47, 50, 55, 75, 110–111, 113–114, 122, 140, 149, 159–160, 171, 174–175, 178–179, 182, 187–196, 199–200, 241–242, 244, 246–247, 249–250, 252–253, 256–258, 264–265, 272–273, 281–282, 297, 332, 336, 338, 347–348, 353–355, 366, 382, 384, 395–396, 413, 416, 421, 423, 426–427, 430–431, 434, 439–442, 444–445, 449–450, 470, 472–475, 517, 519, 522, 524, 557, 573, 576, 578, 587, 595, 602–603, 605–607, 612–613, 615, 619–630, 632–636, 638–643, 645–647, 650, 654–656, 668
  - account dialog box 622
  - Add log files window 519
  - applications 13, 46, 140, 159, 174, 620, 623
  - console 162
  - different abilities 620
  - Library window 642–649
  - login dialog box 622
  - long-term statistics report 654–655
  - managing nodes 473
  - network configuration values 636
  - Nodes window 634–636
  - password dialog box 623
  - repository planning 216
  - restart 475
  - software 140, 159, 172
  - user account 619, 621
  - User Management 621
  - View pane 196
  - window 178, 625
  - workstation 140, 149, 162, 190, 475, 636
- ProtecTIER Replication 92, 158, 397, 607–608
  - Network Performance Validation Utility 256, 619
  - Performance Validation Utility 95
- ProtecTIER Replication Manager xx, 13–16, 45, 157, 160, 172, 252–253, 255–260, 263, 273, 399, 420, 423–427, 471, 578, 599–603, 605, 610–611, 614
  - restore from file 604
- ProtecTIER Repository 15, 20, 54, 70, 76–78
  - subsequent physical storage requirements 77
- ProtecTIER V2.5 434–435, 438, 452
- ProtecTIER VT HyperFactor mode window 544
- ProtecTIER VT name field 231
- pt\_net\_perf\_util utility 95, 670
- PTCLI 586–589
- ptcli
  - parameters and options 675
  - responses 678
- ptcli responses 678

## Q

- quad port 133, 136
- query types and cartridges set 588

## R

- R/W mode 598
- rack cabinets 148
- RAID 0 85, 88
- RAID 1 85–86, 88, 125–126, 129
- RAID 10 15, 71, 85, 88, 217, 222, 639
- RAID 5 86–88
  - arrays 86
- RAID 6 87
- RaidConfigurations 677, 683
- rail kit 122
- RAM 40, 140, 143
- random reads 35, 72
- RAS package 160
- rate-control function 101, 111
- Read/Write cartridge menu 508
- Read/Write cartridge option 50, 508
- real tape libraries 73, 116, 119, 157, 279, 332, 344, 346, 351, 360, 385–387
  - moving tape between them and ProtecTIER systems 386
- real-time compression 38, 345
- rear view 144, 150
- Re-assign device option 48
- Reclamation 380, 384
- reclamation
  - disabling 381
- Reclamation considerations 380
- Recommended configuration changes 383, 396
- Recovering the data 348
- Recovery Management 586
- Recovery management 413
- recovery point objective 344, 377, 391
- recovery point objective (RPO) 344, 377–378, 391, 393–394, 591, 593–594
- Red Hat
  - native 167
- Red Hat Cluster Suite 524, 697–698
- Red Hat Enterprise 47, 50, 140, 159, 167, 175, 324
- Red Hat Enterprise 3 160
- Red Hat Enterprise 4 140
- Red Hat Enterprise Linux 5 Advanced Platform 64-bit license 167
- Red Hat Enterprise Linux 5.2 50, 670
- Red Hat Enterprise Linux 5.2 Advanced Platform 64-bit 124, 126, 128
- Red Hat Enterprise Linux 5.4 50, 140
- Red Hat ES3 158
- Red Hat ES4 158
- Red Hat Global File Systems 159
- Red Hat PID 159
- Redbooks website 704
  - Contact us xvii
- Redundant Array of Independent Disks (RAID) 15, 35, 55, 71, 73, 80, 84, 87–88, 158
  - group 73
  - levels 84–85, 87
- Redundant Power supplies 127, 129
- rehydration processes 345
- Remote Cartridges report 668
- remote cloning 419
- remote sites 255, 347, 349, 391–395, 398–400, 413, 416–417, 419, 601–602, 605, 608–609, 611, 614–616
  - valid database copies 382, 394
- Remote Supervising Adapter (RSA)
  - alerts 688
- Remote Supervising Adapter II 127, 129
- Remote Supervising Adapter II (RSAII) 127, 129, 145, 153
- Removing Host Initiators 295
- Rename library menu 509
- Rename library option 48
- Rename library window 509–510
- replication 139, 410, 416
  - adding to an existing production system 254
  - backlog 344, 347, 382, 394, 595, 606–607, 612
    - flushing 606
  - bandwidth efficient 7
  - bandwidth throttling 570
  - capacity 107
  - cartridge status 662
  - continuous 95, 109, 111
  - data transfer 416
  - definition 16
  - destinations 264–265
  - failure 431
  - grids 14, 16, 43–45, 160, 226, 252–253, 255, 273, 413, 423–426, 429, 471–472, 485, 570, 598, 600
    - definition 16
    - IDs 425
    - members 425
      - definition 16
  - hub incoming 106
  - IP 44
  - IP addresses 604–605, 692
    - restore from 604
  - licensing 43
    - TS7610 SMB 43
    - TS7650 43
    - TS7650G 43
  - long-term statistics 665
  - many to one 420
  - mode of operation 75, 106, 109
  - modes 429
  - modes of operation 607
  - OST 423
  - pairs 425
    - definition 16
  - parallelism schemes 91
  - policies 91, 108–110, 171, 254–255, 263–264, 273, 411, 413, 423, 425, 429, 466, 471, 571, 596, 658
    - definition 16
    - setting up 264
  - policies and activities
    - monitoring 619, 657
  - ports 16, 44, 104
  - prerequisites 424
  - rate limits
    - return to the original rate limits 470
  - repository sizing 75



- settings worksheet 695
- setup 272
- throughput control 252
- time frame
  - definition 16
- topologies 424
- unique IDs 425
- visibility 417
- window 415
- Replication activities tab 657, 660
- Replication activities view 607
- Replication Activities window 660
- Replication Destination window 268
- Replication grid ID
  - definition 16
- Replication IP Address field 262
- Replication menu 108
- Replication Policies list 596
- Replication Policies menu 585
- Replication Policies tab 658
- Replication Policies view 659
- Replication Policies window 658
- Replication policy view 607
- Replication Rates Limits window 428, 469
- replication/DR operation
  - features 609
- replication-network management 413
- repositories 21, 28, 34, 41–42, 46, 80
  - adding to the grid 259
  - creating 219
  - deleting 227, 485
  - destination 108
  - expanding 481
  - planning an expansion 476
  - remote 112
  - renaming 225
  - single 29
- Repository configuration pane 639, 641
- Repository Nominal Quota 250, 445
- Repository resources window 223
- Repository Space Usage 638
- repository takeover wizard 601
- reserve space
  - implications 431
- Reserve space feature 430
- Reset drive option 49, 547
- Reset robot option 48, 548
- Reset robot window 548
- resource virtualization optimization 30
- Restore grid procedure 603
- Restore grid window 603–605
- retention period 15, 77–80
- REUSEDELAY parameter 381
- revolutions per minute (RPM) 26, 59, 81
- RHELAP 50, 140, 159
- RMAN 5, 73, 408
- robots 22–23, 25, 29, 235, 354, 366–367, 528, 548, 572, 644
- routes
  - static 257–258, 438

- rpm
  - commands 316–319
  - packages 316, 318, 324, 326, 448
- Rpositorystatistics 677
- RTO
  - classifications 392, 592
  - tiers 391, 592

## S

- SAN 69, 162, 165
  - configuration 162–163, 297
    - rules 166
  - connectivity 69
  - design 276
  - environment 162, 165
  - environments 167, 297
  - Stretched Tape 391, 591
  - tools 165
  - zoning 281
- SAN Fabric 162, 167
  - configuration 276
  - mode 156
- SAN-attached backup storage devices 385
- SAS Drives 125–126, 129
- SATA CD/DVD-ROM 127, 129
- SATA hard disk drives 31
- Save password check box 624
- scalability 6–8, 18, 53–54, 141
- scratch 34–35, 39, 91, 361, 365, 372, 380–381, 383–384, 504, 607, 616
  - pool 116
- SCSI 168, 222, 279, 303, 305, 311, 314, 318, 354, 358–359, 361, 367, 370, 547–548
- Second Server 134, 137
- secondary (target) repository sizing 107
- Secure Hash Algorithm (SHA-1) 5
- seek time 81
- Selected Node field 625
- separate (multiple) domains approach 393, 593
- Server options 677
- servers 10, 25–26, 126
  - Outgoing Email 205
  - starting 535
  - stopping 538
- ServerVersion 677, 682
- Services pane 487, 505, 507, 509, 547–548, 558, 562, 566, 568, 570, 627, 642
  - library name 642
- Set compression mode 48
- Set control path failover mode 48, 297
- Set library type option 48
- Setting replication rate limit 427
- Setting the replication rate 427, 469
- Setting the trace levels 545
- SHA-1 5–6
- Shared Storage Option (SSO) 389
- shelf 16, 90, 412, 425, 551, 567–568, 650, 695
  - monitoring 650
- single domain architecture 392
- single domain environment 594

- master clustered 393
- single domain scenario 609
- single fabric 156
- Single Point of Failure (SPOF) 156
- sites
  - assessment 76
  - customer 173
  - disaster recovery 58, 89, 101, 105, 108, 110, 112–114, 254, 264, 382, 394, 397–399, 410, 576–577, 579, 589, 591, 593–595
  - FTP 298
  - installation 161, 688
  - primary 8, 58, 393, 400, 409, 412, 416, 571, 596, 601–602, 605, 611
  - production 15, 417, 575, 579
  - remote 10, 15, 89–91, 93–94, 114, 252, 268, 349, 391–392, 395, 414, 418, 572, 606, 611, 613, 615–616, 669
  - second 94
  - secondary 254, 411, 413, 419
  - target 607
  - user 172
- size reduction 42
- sizing inputs 54
- Slots tab 507, 559, 563, 566, 648
- SMB 11–12, 29, 121, 140, 436
- SMIT 301
- SNMP 161
  - compatibility 173
- software
  - compatibility 345
  - compression 392, 592
  - currency 345
  - prerequisites 158
- space reclamation 116, 119
- SQL Litespeed data 408
- SSH 699
- stand-alone configuration 128, 141–142
  - IBM System Storage DS4700 143, 149
- standard backup activity 35
- statistics
  - creating and downloading 348–349, 668
- status line 624, 642
- steady state 35–36
- stgpool command 381, 387
- storage
  - arrays 59, 72–73, 165, 217, 222
  - blocks 35
  - checkpoints 392
  - media
    - limitations 386
  - performance optimization 84
  - pools
    - hierarchy 386
    - migrating data 386
  - requirements 5, 10
    - reduction 10
  - sizing 439
  - unit
    - creating 463, 465

- storage servers (STS) 13, 36, 102, 115, 240–247, 249–250, 422, 433–436, 439–444, 446, 448–452, 455, 458, 460
  - configuration 242
  - deleting 245
  - modifying credentials 244, 440
  - settings 452
- storage system
  - administrator 689
- storage system administrator 689
- Storage Type definition 454
- subnet addresses 437
- Sun Solaris 8 158
- Support required errors list 556
- SUSE 10 158
- Symantec 25, 232, 251, 396, 433, 452, 465
  - Symantec NetBackup (NBU) 113, 124, 126, 128, 242–243, 246, 389–390, 392–396, 418, 421, 434, 439–440, 452, 458, 460–461, 465, 572, 590–596, 609, 611, 614, 617
    - architecture 390, 591
    - background 590
    - catalog backups 592
    - clustered 391
    - environment 389, 590
    - GUI tools 452
    - recovering a master server from an existing DB copy 395
    - Remote Manager 447
    - setting up for backup and restore 390
    - setting up for backups 591
    - setting up for disaster recovery 390, 591
    - Vault 392
    - websites 389
  - Symantec NetBackup (NBU) Domain 389
  - Symantec NetBackup Administration Console 464
  - Symantec NetBackup Catalog Recovery 394, 594
  - sync time 349–350, 382, 395, 587, 595, 613, 615, 669–670
  - synchronization
    - considerations 113
    - time stamps 348–349
- System Storage Interoperation Center (SSIC) 122, 128, 158–159, 316, 330–331
- systems
  - renaming 198
  - repairing 557
- Systems window 624, 626, 628–632
  - General tab 628
  - VT tab 631
- system-wide nominal capacity 117

## T

- tape cartridges
  - exchanging with shelf 473, 565
- tape devices 167, 277, 279, 293, 317–319, 356, 361, 370, 400, 631
- tape drives 18, 21–24, 28–29, 48, 487, 489–491, 495–497, 528, 620, 632, 643–644, 646
  - backup application 116

- tape libraries
  - management 117
  - shared 385
  - zoning 345
- tape pools 91, 255, 381, 387, 389
- Tape Storage Service Console 688
- tapes
  - and medium changer devices
    - configuration 302
  - block sizes 346
  - encryption 396
- tapeutil 357
- TCP
  - ports 157, 425
  - protocol 400, 416
  - retransmissions versus total TCP segments sent 100, 674
- terabyte capacity ordering 43
- throughput
  - 16 streams 99
  - considerations 68
  - default settings 99
  - single stream 99
  - single stream, 1 MB send buffer 673
- tier levels 138–139
- tight integration 102, 115
- Total utilization pane 639
- trace buffers
  - dumping 545
  - modifying 544
  - resetting 546
- traditional tape libraries 13, 32, 36, 102, 115–116
- treat disk as disk 102, 115
- TS3000 System Console (TSSC) 18, 22, 26, 29, 50, 128, 131–133, 135, 141–142, 145–148, 153–154, 156, 160–162, 171–173, 688, 690–691, 694
  - IP address range 690
  - switches 142, 147
- TS3500
  - tape libraries 18, 124, 126–127
  - virtual libraries 32
- TS7610 29–32, 55, 73, 80, 82, 115–116, 118–119, 122, 124, 132, 173, 194–195, 199–200, 208, 234, 256, 411, 414, 435, 437
  - configuration wizard setup 200
  - start message 194
- TS7610 ProtecTIER Deduplication SMB Appliance 11–12, 121
- TS7650 17, 26, 28–29, 43–44, 50–51, 55–56, 59, 73, 80, 82, 90, 115–116, 118–119, 121–122, 124, 126, 128, 131–132, 136, 138–139, 143, 149, 159, 161–162, 169, 171, 173, 199, 208, 222, 234, 254, 256, 297, 330, 346, 386, 389, 411, 414, 429, 435, 473, 590, 624, 651, 674, 690–692
  - servers 29, 69
  - systems performance 74
- TS7650 Appliance
  - system console components 26
- TS7650 ProtecTIER 89
  - Deduplication Appliance 27–28, 53
  - Deduplication Gateway 19
  - enterprise-level systems nodes 68
  - technology 73
- TS7650 ProtecTIER Deduplication Appliance 11–12, 17, 26–27
- TS7650 ProtecTIER systems 50, 55, 58, 68, 115–116, 118–119
  - front-end ports 69
- TS7650-AP1
  - appliance 27
- TS7650G 15, 17–18, 20–22, 25, 43–44, 50–51, 56–57, 59, 73, 75, 80–81, 90, 103, 115–116, 118–119, 121–125, 127–128, 131, 137–138, 140–143, 149, 160–161, 169, 171–173, 190, 199, 208, 210, 234, 254, 256, 297, 330, 346, 386, 389, 411, 414, 429, 435, 473, 624, 635, 651, 674, 688, 690–692
  - back-end ports 69
  - systems 116–117, 119
    - cartridge capacities 116
- TS7650G Gateway models 17
- TS7650G ProtecTIER 58
- TS7650G ProtecTIER Deduplication Gateway 11–13, 17–18, 20–21, 23–24, 27–28, 46–47, 50, 124–128, 140
  - configuration 21–22, 28–29
  - order 20
  - solution 18
- TS7650G servers 17–18, 24, 59, 128, 131, 144–145
- typical deployment 413

## U

- ULTRIUM3 363, 371, 401
- uncompressible 4
- understanding the existing environment 57
- Unload drive option 49, 549
- user data 14, 34–35, 41, 54, 70, 76, 157, 168, 174, 223, 482–483, 638–639
  - segment 34
  - storage space 224
- User name field 244–245, 439
- utilities
  - running 651
- Utilization graph 639–640

## V

- Verify password field 244–245, 440–441
- Veritas Cluster server 389, 590
- Version Information pane 635
- View Resources button 638
- virtual
  - drives 49
  - LANs 693
  - medium changers 48
  - resources 22
  - scratch tapes 35
  - shelf 46
- virtual drives 20, 26, 32, 47, 49, 74, 219, 276, 279, 297, 344, 365–366, 384
- virtual libraries 12, 21–24, 28–29, 39, 46–48, 171–172, 230, 237, 240, 487, 500, 543

- robots 24
- single 23
- space utilization 39
- supported 48
- TS7650G ProtecTIER Deduplication Gateway cluster configuration 24
- virtual robots 23–25, 235, 353, 494, 548, 643
- Virtual Tape (VT) service 528, 630
- virtual tape drives 18, 23–25, 27, 32, 73–74, 123–124, 126–127, 138–139, 235, 276, 278, 336, 345, 355, 383, 389, 398, 408, 490, 496, 620
- Virtual Tape Facility Daemon (VTFD) 46, 96, 678
- virtual tape libraries (VTL) 10, 13–14, 16–17, 19, 23–25, 31–33, 40, 42, 46, 56–57, 68, 73, 89–90, 92, 103, 110, 113–114, 123, 126, 129–131, 134, 137, 144–146, 151–155, 157, 163–165, 194, 199, 208, 225, 252, 273, 275, 277, 279, 283–284, 297, 334–336, 338–339, 344–345, 352–355, 360, 366, 383, 385–386, 392, 404, 407–408, 410–411, 415, 417, 419–421, 424–426, 434–435, 439, 487, 490, 550, 566, 571, 601, 605, 619, 629, 631, 642, 659
  - concepts 38
  - data storage 39
  - drives 156
  - export slots 410
  - export/import 415
  - implementation 352
- virtual tape management framework 411
- virtual tapes
  - cartridges 39, 381, 561, 664
  - emulation 18, 124, 126–127
- visibility control switch 255, 264, 392, 394, 412, 417, 571, 593, 595
- visibility features 415
- visibility switch 113, 264, 392, 419, 576, 593, 609–610, 612–614, 616
  - versus basic DR 607
- visibility switching 17, 608
- VLAN 693
- vmstat 73–74
- VT
  - column 630
  - menu 284–285, 295, 488
  - monitoring service window 505
  - overview pane 644
  - tab 630–631

## W

- WAN 16, 44, 89, 92, 95–96, 99, 397, 413, 416, 671–672
- Western Telmatic Inc. (WTI) 697–701
  - switch 699
- Windows 2000 140, 160
- Windows 2003 140
- Windows 7 140, 160
- Windows Server 2003 313–314
- Windows Server 2008 313–314, 366
- Windows XP 140
- work sheet 688
- workloads 5, 37, 345–346
- WWN of ProtecTIER Front End ports 276

## X

- X Window System 140
- x86 159
- XML 678

## Z

- zone 69, 157, 166–167, 276, 281, 345, 350, 670
- zoning 166



**Redbooks**

# IBM System Storage TS7650, TS7650G, and TS7610

(1.0" spine)

0.875" <-> 1.498"

460 <-> 788 pages







# IBM System Storage TS7650, TS7650G, and TS7610



**Redbooks®**

**Many to one native replication**

**Benefits of deduplication**

**Version 2.5 introduced**

This IBM Redbooks publication describes the IBM solution for data deduplication, the IBM System Storage TS7650G ProtecTIER Deduplication Gateway, and the IBM System Storage TS7650 ProtecTIER Deduplication Appliance. This solution consists of IBM System Storage ProtecTIER Enterprise Edition V2.3 software and the IBM System Storage TS7650G Deduplication Gateway (3958-DD1 and DD3) hardware, as well as the IBM System Storage TS7650 Deduplication Appliance (3958-AP1). They are designed to address the disk-based data protection needs of enterprise data centers.

We introduce data deduplication and IP replication and describe in detail the components that make up IBM System Storage TS7600 with ProtecTIER. We provide extensive planning and sizing guidance that enables you to determine your requirements and the correct configuration for your environment. We then guide you through the basic setup steps on the system and on the host, and describe all operational tasks that might be required during normal day-to-day operation or when upgrading the TS7650G and TS7650.

This publication is intended for system programmers, storage administrators, hardware and software planners, and other IT personnel involved in planning, implementing, and using the IBM deduplication solution, as well as anyone seeking detailed technical information about the IBM System Storage TS7600 with ProtecTIER.

## **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](http://ibm.com/redbooks)

SG24-7652-03

ISBN 0738435848