

IBM ProtecTIER Implementation and Best Practices Guide

Karen Orlando

Mara Miranda Bautista

Emmanuel Barajas Gonzalez

Carlos A. Martinez Vazquez



Storage



International Technical Support Organization

IBM ProtecTIER Implementation and Best Practices Guide

June 2016

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

Fourth Edition (June 2016)

This edition applies to IBM ProtecTIER Version 3.4

Contents

Notices	xiii
Trademarks	xiv
IBM Redbooks promotions	xv
Preface	xvii
Authors	xviii
Now you can become a published author, too!	xix
Comments welcome	xix
Stay connected to IBM Redbooks	xx
Summary of changes	xxi
June 2016, Fourth Edition	xxi
Part 1. General best practices	1
Chapter 1. IBM ProtecTIER basics	3
1.1 Terminology	4
1.1.1 ProtecTIER Appliance terminology	5
1.1.2 ProtecTIER Gateway terminology	5
1.1.3 ProtecTIER replication terminology	6
1.2 ProtecTIER fundamental concepts	9
1.2.1 ProtecTIER with deduplication disabled	11
1.3 ProtecTIER models	11
1.3.1 ProtecTIER appliance	11
1.4 Remote support and call home	13
1.5 Simple Network Management Protocol (SNMP)	14
1.5.1 Alert filtering	15
1.5.2 Space thresholds	15
1.5.3 Email alerts	17
1.6 ProtecTIER command-line interface	18
1.7 Performance considerations for repository creation	19
1.7.1 Example 1: Calculate performance throughput	19
1.7.2 Example 2: Calculate performance throughput	20
Chapter 2. Deduplication considerations	21
2.1 HyperFactor data deduplication	22
2.1.1 HyperFactor, deduplication, and bandwidth savings	23
2.2 ProtecTIER HyperFactor deduplication processing	24
2.3 Components of a ProtecTIER system	24
2.3.1 ProtecTIER server	25
2.3.2 HyperFactor deduplication algorithm	26
2.3.3 Disk storage subsystem	26
2.4 Benefits of ProtecTIER HyperFactor	27
2.4.1 Flexibility	27
2.4.2 High availability	27
2.4.3 High performance, low storage requirements, and lower environmental costs ..	27
2.5 General ProtecTIER deduplication considerations	28
2.5.1 Rethinking your overall backup strategy	28

2.5.2	Data reduction technologies should not be combined	29
2.5.3	Data streams must be in order	29
2.5.4	Data organization in your ProtecTIER repository	30
2.5.5	The dynamics of the ProtecTIER repository	30
2.5.6	ProtecTIER repository usage	31
2.5.7	Compression	33
2.5.8	Encryption	33
2.5.9	Database logs and other data types with high data change rates	34
2.5.10	Multiplexing	34
2.5.11	Tape block size	34
2.5.12	File size	34
2.6	Data types	35
2.6.1	Candidates for a high factoring ratio	35
2.6.2	Candidates for a low factoring ratio	35
Chapter 3.	Networking essentials	37
3.1	Network terminology	38
3.2	General configuration considerations	39
3.3	Bonding and teaming	39
3.3.1	The three bonding modes of ProtecTIER	39
3.4	Preferred ProtecTIER bonding configuration	43
3.4.1	VLANs	45
3.4.2	IP addresses	47
3.4.3	Routing the IP traffic	47
Chapter 4.	Virtual Tape Library guidelines	49
4.1	ProtecTIER Virtual Tape Library introduction	50
4.2	General preferred practices for the virtual tape library	50
4.3	Setting up the virtual library and cartridges	51
4.3.1	Creating libraries	51
Chapter 5.	ProtecTIER File System Interface: General introduction	65
5.1	ProtecTIER FSI network overview	66
5.1.1	ProtecTIER network	66
5.1.2	Network configuration considerations	67
5.1.3	Connecting a ProtecTIER server to the network	67
5.1.4	Replication	70
5.1.5	Disaster recovery: Test	71
5.1.6	Disaster recovery: Event	71
5.1.7	General FSI suggestions	71
5.2	File System Interface guidelines for NFS	72
5.2.1	ProtecTIER NFS authentication and security management	72
5.2.2	Understanding root squash	75
5.3	File System Interface guidelines for CIFS	76
5.3.1	Mounting the NFS export in a UNIX system	76
5.3.2	ProtecTIER authentication and user management	77
5.4	FSI file system scalability	78
Chapter 6.	Host attachment considerations for Virtual Tape Library	79
6.1	General suggestions	80
6.2	Device driver specifications	80
6.2.1	AIX specifications to work with VTL	81
6.2.2	Solaris specifications to work with VTL	82
6.2.3	Linux specifications to work with VTL	82

6.2.4	Windows specifications to work with VTL	82
6.2.5	IBM Tape Device Driver	82
6.2.6	Control path failover and data path failover	83
6.2.7	Persistent device naming	85
6.3	LUN masking for VTL systems	87
6.3.1	LUN masking methods and preferred practices	88
6.3.2	LUN masking configuration steps	89
Part 2.	Back-end storage subsystems	97
Chapter 7.	Back-end storage overview	99
7.1	Overview	100
7.1.1	ProtectTIER Planner tool	102
7.2	Dependencies from a back-end storage subsystem view	103
7.3	Dependencies from a ProtectTIER view	103
7.4	Smart storage subsystems	104
7.5	Key rules for a ProtectTIER server	105
7.6	Storage arrays configuration	106
7.6.1	General requirements	106
7.6.2	RAID considerations	106
7.6.3	LUNs	108
7.6.4	Expanding the repository	108
7.7	Storage area network fabric	109
7.7.1	Two Fibre Channel paths to each storage controller	109
7.7.2	Dedicated zones	109
7.7.3	Front-end zones	110
7.7.4	Back-end zones	110
7.7.5	SAN paths	110
Chapter 8.	IBM FlashSystem 900 with System Storage TS7650G ProtectTIER Deduplication Gateway	111
8.1	Introduction to flash technology	112
8.2	IBM FlashSystem overview	113
8.3	General FlashSystem considerations	114
8.4	Deploying the IBM FlashSystem 900 with IBM ProtectTIER	115
8.4.1	Creating user data and metadata volumes	115
8.4.2	Mapping the FlashSystem Volumes to the ProtectTIER System	117
8.4.3	Creating file systems and building the ProtectTIER repository	119
8.5	General preferred practices integrating FlashSystem 900 with ProtectTIER	122
8.5.1	ProtectTIER Metadata Planner preferred practices for FlashSystem 900	123
Chapter 9.	IBM Storwize Family and SAN Volume Controller	125
9.1	IBM Storwize V3700 and V5000 overview	126
9.2	General considerations: V3700, V5000, and V7000	127
9.2.1	Configuration steps: ProtectTIER repository	127
9.2.2	Creating empty user data and metadata storage pools	127
9.2.3	Creating the MDisk arrays or discovering unmanaged MDisks	129
9.2.4	Creating volumes with a sequential virtualization type	132
9.2.5	Creating a host connection for the ProtectTIER nodes by using the GUI	136
9.2.6	Mapping volumes to a host	140
9.2.7	Creating file systems and building the ProtectTIER repository	142
9.2.8	Expanding the repository	145
9.3	IBM System Storage SAN Volume Controller, IBM Storwize V7000, and V7000 Unified Storage	145

9.3.1 Storage virtualization introduction	145
9.3.2 Terminology	146
9.4 General notes	149
9.5 Firmware level	150
9.6 Fibre Channel connection topology	151
9.7 User data and metadata pool: General suggestions	154
9.7.1 Metadata pool	154
9.7.2 User data pool	155
9.8 Configuration steps	155
Chapter 10. IBM XIV Storage System	157
10.1 XIV Storage System hardware	158
10.2 Fibre Channel switch cabling and zoning	159
10.2.1 Zoning configuration	160
10.3 Configuring XIV Storage System for ProtecTIER server	160
Chapter 11. IBM System Storage DS8000	167
11.1 DS8000 series overview	168
11.1.1 Disk drives	169
11.1.2 RAID levels	170
11.2 General considerations	171
11.2.1 Planning tools	171
11.2.2 Metadata	171
11.2.3 User data	172
11.2.4 Firmware levels	172
11.2.5 Replication	172
11.3 Rotate extents: Striping and when to use it	173
11.3.1 When not to use rotate extents	174
Chapter 12. Backup management introduction	179
12.1 Introduction	180
12.2 General suggestions	180
12.2.1 Interoperability	180
12.2.2 Software compatibility	180
12.2.3 Software, backup application, and operating system	180
12.2.4 Tape library zoning	180
12.2.5 Compression	182
12.2.6 Encryption	182
12.2.7 Multiplexing	182
12.2.8 Tape block sizes	182
12.2.9 Type of data that is backed up	182
12.3 General advice for backups	183
12.4 ProtecTIER integration with backup applications	183
12.5 Backup application terminology cross-reference	184
12.6 Backup application catalog	184
12.7 Remote cloning of virtual tapes	185
Chapter 13. IBM Spectrum Protect	187
13.1 IBM Spectrum Protect VTL	188
13.2 IBM Spectrum Protect: Preferred options	188
13.2.1 LAN-free backups with the ProtecTIER product	189
13.2.2 Data streams	190
13.2.3 Reclamation	190
13.2.4 Collocation	191

13.2.5 Migration	191
13.2.6 Backing up your IBM Spectrum Protect database	191
13.2.7 Physical tape.	192
13.2.8 Avoiding mount conflicts	192
13.2.9 Multiple streams from the client with the resourceutilization parameter	192
13.2.10 Accommodating increased sessions.	194
13.2.11 IBM Spectrum Protect storage pool selection.	194
13.2.12 Technical overview	195
13.2.13 Advantages of an IBM Spectrum Protect environment with ProtecTIER	196
13.2.14 IBM Spectrum Protect configuration with VTL	196
13.2.15 Updating to a VTL library type	199
13.2.16 Defining and deleting IBM Spectrum Protect libraries with many drives.	199
13.3 IBM Spectrum Protect: FSI	200
13.3.1 Setting up backup and restore on IBM Spectrum Protect.	201
13.3.2 Parameters for best performance with ProtecTIER FSI	205
Chapter 14. Symantec NetBackup and BackupExec	207
14.1 NetBackup overview	208
14.2 General suggestions for NetBackup	208
14.3 NetBackup in a VTL environment	210
14.4 NetBackup in an FSI environment	210
14.4.1 NetBackup in an FSI-CIFS environment.	210
14.4.2 NetBackup in an FSI-NFS environment	214
14.5 Symantec BackupExec in an FSI environment.	219
Chapter 15. EMC NetWorker.	221
15.1 Overview	222
15.2 EMC NetWorker in a VTL environment.	223
15.2.1 General suggestions	223
15.2.2 Suggestion if a ProtecTIER server is used as a VTL	223
15.3 EMC NetWorker in an FSI environment	224
15.3.1 Creating a Windows user for EMC NetWorker	224
15.3.2 Setting up for backup and restore.	225
15.3.3 General configuration suggestions	225
15.3.4 Setting the information to be backed up	226
15.3.5 Setting the time for the backup	227
15.3.6 Performing a restore	227
15.3.7 Parameters for best performance with ProtecTIER FSI	227
Chapter 16. HP Data Protector	229
16.1 HP Data Protector with ProtecTIER	230
16.1.1 HP Data Protector architecture with ProtecTIER	230
16.2 HP Data Protector in a VTL environment	231
16.2.1 Enabling the robotic bar code reader	232
16.2.2 Increasing the tape block size.	232
16.2.3 Enabling the lock name.	234
16.2.4 Disabling compression, encryption, and CRC checksum	235
16.2.5 Hosts multipath support	236
16.2.6 Load balancing	238
16.2.7 Using a mirroring functionality.	238
16.2.8 Troubleshooting logs.	239

Chapter 17. IBM i and Backup, Recovery, and Media Services	241
17.1 IBM i overview	242
17.1.1 Integrated file system	242
17.1.2 Integrated database	242
17.1.3 Object-based architecture	242
17.1.4 Libraries	243
17.1.5 Backup considerations in IBM i	243
17.2 Integration of IBM i and ProtecTIER in a VTL environment	243
17.2.1 Backup considerations with ProtecTIER	243
17.2.2 Suggested ProtecTIER and IBM i configuration	244
17.3 Configuration of BRMS for ProtecTIER	246
17.3.1 BRMS overview	246
17.3.2 Suggested configurations of BRMS	247
17.4 Deploying ProtecTIER with BRMS for disaster recovery	248
17.4.1 BRMS available at the production site and DR site	248
17.4.2 No BRMS at the DR site	250
17.5 Use DUPMEDBRM to duplicate media from ProtecTIER to physical media	253
17.5.1 Enable append to multiple parallel volumes during DUPMEDBRM	253
17.5.2 Support multiple batch jobs for the DUPMEDBRM command	254
17.6 Device configuration	255
17.7 BRMS movement of replicated volumes: considerations	255
Chapter 18. Commvault	257
18.1 Commvault introduction	258
18.1.1 Commvault components	258
18.2 Commvault with ProtecTIER VTL	260
18.2.1 Commvault configuration	260
18.2.2 Data multiplexing	267
18.2.3 Hardware compression	269
18.2.4 Data encryption	269
18.2.5 Alternative data paths	270
18.3 Commvault FSI	271
18.3.1 Setting up backup and restore in a CIFS environment	271
18.3.2 Parameters for best performance with ProtecTIER FSI-CIFS	273
18.3.3 Setting up backup and restore in an NFS environment	277
18.3.4 Parameters for best performance with ProtecTIER FSI-NFS	281
Chapter 19. Veeam FSI-CIFS	283
19.1 Setting up backup infrastructure	284
19.1.1 Creating backup repository	284
19.1.2 Adding VMware backup proxy	285
19.1.3 Adding a VMware Server	286
19.2 Setting up backup and restore	286
19.2.1 Creating a new backup job for backup	286
19.2.2 Performing backup	287
19.2.3 Performing restore	287
19.3 Parameters for best performance with ProtecTIER-CIFS	288
19.3.1 Disabling data de-duplication and compression	288
19.3.2 Setting data block size for backup to SAN or local storage	289
19.3.3 Setting Incremental backup mode	289
19.4 Advanced backup repository settings	290
19.4.1 Setting data blocks alignment	290
19.4.2 Setting data blocks decompression	290
19.5 Summary of parameters for best performance	291

Part 3. Application considerations	293
Chapter 20. Application considerations and data types	295
20.1 IBM Domino	296
20.1.1 Common server	296
20.1.2 Existing backup and disk space usage	296
20.1.3 Domino attachments and object service	298
20.1.4 Applying the DAOS solution	300
20.1.5 ProtecTIER considerations	302
20.2 Microsoft Exchange	305
20.2.1 Defragmentation	305
20.2.2 Suggestions for Microsoft Exchange	305
20.2.3 Microsoft Exchange 2010	306
20.3 Microsoft SQL Server	306
20.3.1 Integrating the ProtecTIER server with Microsoft SQL Server backup	306
20.3.2 Index defragmentation	308
20.3.3 Suggestions for Microsoft SQL Server	308
20.3.4 LiteSpeed for SQL Server	309
20.4 DB2	309
20.4.1 Combining DB2 compression and ProtecTIER deduplication	309
20.4.2 Upgrading the DB2 database to improve deduplication	310
20.4.3 DB2 DEDUP_DEVICE setting	312
20.4.4 Example of DEDUP_DEVICE setting	313
20.4.5 Excluding logs from the DB2 database backup	314
20.4.6 DB2 suggested settings without DEDUP_DEVICE	314
20.4.7 Example of DB2 command using sessions, buffers, and parallelism	314
20.5 Oracle	315
20.5.1 Suggested RMAN settings	315
20.5.2 Mounting NFS Oracle Server to ProtecTIER NAS	317
20.5.3 Using ProtecTIER NAS to run RMAN incremental merge backups	319
20.5.4 Using ProtecTIER NAS to store Oracle Data Pump exports	321
20.5.5 Using ProtecTIER NAS to store Oracle database files and offline redo logs	321
20.5.6 Other suggestions for RMAN	322
20.6 SAP	322
20.6.1 SAP introduction	323
20.6.2 Data protection for SAP	323
20.6.3 Integration of Tivoli Storage Manager for ERP with SAP	325
20.6.4 Tivoli Storage Manager for ERP for Oracle database	326
20.6.5 Tivoli Storage Manager for ERP for DB2	329
20.6.6 SAP BR*Tools for Oracle using BACKINT	330
20.6.7 SAP BR*Tools for Oracle using RMAN with Tivoli Storage Manager	332
20.6.8 SAP BR*Tools for Oracle: Using RMAN to configure DB2 to use Tivoli Storage Manager	333
20.6.9 Preferred practices for Tivoli Storage Manager for ERP with ProtecTIER	334
20.7 VMware	335
20.7.1 Technical overview	335
20.7.2 Settings and tuning for VMware and Tivoli Storage Manager	336
20.7.3 Backup solutions	337
20.7.5 Configuring the ProtecTIER server	341
20.7.6 Installing the tape driver on the Tivoli Storage Manager server and the Tivoli Storage Manager storage agent	342
20.7.8 Tivoli Storage Manager server configuration	347
20.7.9 Tivoli Storage Manager client installation	350

20.7.10	Disabling compression and deduplication on Tivoli Storage Manager	351
20.7.11	Configuring a full VM backup through the vStorage API.	353
20.7.12	VMware Guest OS backup to ProtecTIER	353
Part 4.	Replication and disaster recovery	359
Chapter 21.	ProtecTIER replication	361
21.1	ProtecTIER IP replication	362
21.2	Native replication.	362
21.2.1	One-to-one replication	363
21.2.2	Many-to-one replication	363
21.2.3	Many-to-many replication	363
21.2.4	VTL replication	363
21.2.5	FSI replication	364
21.2.6	Replication grid	364
21.2.7	Replication topology group	365
21.3	Replication policies	365
21.4	Visibility switching	366
21.5	Principality.	367
21.6	Replication Manager	367
21.7	Initial synchronization	369
21.8	Replication schedules	370
21.8.1	Continuous replication	370
21.8.2	Scheduled replication	371
21.8.3	Centralized Replication Schedule Management.	372
21.8.4	Replication rate control	373
21.8.5	Setting replication rate limits	375
21.8.6	Limiting port bandwidth consumption	376
21.9	Replication backlog	377
21.9.1	SNMP alerts for replication backlog	378
21.10	Replication planning	379
21.10.1	Replication throughput barriers	380
21.10.2	Calculating the replication data transfer	381
21.10.3	Calculating replication bandwidth	381
21.10.4	Ports for replication in firewalled environments.	382
21.11	Bandwidth validation utility	382
21.11.1	Using the bandwidth validation utility to test the data flow	383
21.11.2	Interpreting the results	385
21.11.3	Repository replacement	386
21.12	Planning ProtecTIER replication	387
21.12.1	Deployment planning scenario: many-to-many.	388
21.12.2	Many-to-one replication	391
21.13	The backup application database backup.	393
Chapter 22.	Disaster recovery deployment with backup applications.	395
22.1	Disaster recovery operations	396
22.2	ProtecTIER replication overview	396
22.3	Disaster recovery operations with VTL	397
22.3.1	Replication data transfer.	397
22.3.2	Managing cartridges after replication	397
22.3.3	Cartridge replication requirements	399
22.3.4	Importing/exporting slots allocation in VTL	399
22.3.5	Import/export slots searching	399
22.3.6	Automation of daily operation	399

22.3.7 Gauging the replication completion status	400
22.4 Disaster recovery operations with FSI	402
22.4.1 Replication destination directory (RDD)	403
22.4.2 ProtecTIER FSI cloning	403
22.4.3 Preferred replication practices for FSI.	404
22.5 Entering ProtecTIER DR mode	404
22.5.1 Working at the disaster recovery site	405
22.5.2 Inventory command options for a VTL disaster recovery scenario	405
22.5.3 Commonly used disaster recovery queries.	406
22.5.4 Returning to normal operations.	406
22.6 The backup application catalog.	408
22.6.1 ProtecTIER replication with IBM Spectrum Protect.	409
22.6.2 Recovering the backup application catalog.	409
22.6.3 IBM Spectrum Protect reclamation and housekeeping.	411
22.7 Single domain and multiple domains	411
22.7.1 Single domain environment	412
22.7.2 Multiple domain environment	412
22.8 Deploying replication with specific backup applications	412
22.8.1 Recovery point objective (RPO)	412
22.9 IBM Spectrum Protect.	413
22.9.1 Scenario 1: Replication complete	413
22.9.2 Scenario 2: Replication incomplete.	414
22.9.3 Scenario 3: Auditing tapes and fixing the IBM Protect database	414
22.9.4 Reclamation considerations	415
22.9.5 Determining what is available for restoration at the DR site	416
22.10 Symantec NetBackup deployment with ProtecTIER replication	418
22.10.1 Scripting the inject/eject commands	418
22.10.2 Scripting the inventory commands	418
22.10.3 Setting up NetBackup for backup and restore	418
22.10.4 Setting up NetBackup for disaster recovery	418
22.10.5 Cross-site backups	419
22.10.6 ProtecTIER disaster recovery with Symantec NetBackup	420
22.10.7 Single domain versus two separate domains	420
22.10.8 Disaster recovery scenarios	421
22.10.9 Determining what is available for restore at the disaster recovery site.	422
22.10.10 Eject and inject commands from NetBackup software	423
22.11 EMC NetWorker deployment with ProtecTIER replication	424
22.11.1 Cloning physical tapes with ProtecTIER replication	425
22.11.2 Disaster recovery with ProtecTIER replication	427
22.12 Commvault	430
22.12.1 Prerequisites	430
22.12.2 Running the Commvault backup operation.	430
22.12.3 Commvault resources	432
Part 5. Appendixes	433
Appendix A. Hints and tips	435
Remote access with IBM Endpoint Manager for Remote Control (Assist On-site)	436
Assist On-site (AOS) description.	436
AOS terminology.	437
Available session options	437
AOS in ProtecTIER (SM2 and DD6)	438
Operating system recovery functionality	440

OS Recovery functionality description	440
The OS Recovery menu	441
Launching an OS Recovery image creation	442
Recovering the operating system using an image	442
Online fsck command	444
The Linux fsck tool	445
Execution of fsck	445
Online fsck procedure	445
Dedup estimator tool	446
Previous solution	446
Tool functions	447
Local directory cloning for FSI	449
Cloning a local directory in FSI	450
Local directory cloning benefits	450
Commands to send files to and receive files from IBM Support	451
Commands to upload files from a ProtecTIER node to ECUREP	451
Receiving files from IBM support	453
Graphical performance tool	453
ProtecTIER licensing tips	454
Capacity representation in the ProtecTIER Manager GUI	454
The role of the metadata space in the licensing process	455
Reconfigure an IBM Spectrum Protect instance attached to ProtecTIER after a reboot . .	455
Appendix B. ProtecTIER compatibility	457
IBM System Storage Interoperation Center (SSIC)	458
Independent Software Vendor Support Matrix	459
Appendix C. ProtecTIER parsers	461
The ProtecTIER parsers	462
Terminology	462
How metadata from the backup application hinders deduplication	462
ProtecTIER parser functionality	464
ProtecTIER parsers support	464
What workloads benefit from the ProtecTIER parsers	465
Background information: Causes of low deduplication ratios	466
Estimating the benefit of a parser	466
Environments that benefit from parsers	467
Experience from one user site	469
Use the analyze_sessions utility to monitor the benefit of a ProtecTIER parser	469
Appendix D. Managing cartridge sizes with ProtecTIER	473
Effects of dynamic cartridge sizes	474
The mechanism behind fluctuating cartridge sizes	474
Glossary	477
Related publications	481
IBM Redbooks	481
Other publications	482
Publications common to the TS7650 Appliance and TS7650G	482
TS7620 Appliance Express publications	482
Online resources	483
Help from IBM	484

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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 grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

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 jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

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

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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 actual people or business enterprises 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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

AIX®	HyperSwap®	Notes®
DB2®	IBM®	POWER®
DB2 Connect™	IBM FlashCore™	ProtecTIER®
DB2 Universal Database™	IBM Spectrum™	Redbooks®
Domino®	IBM Spectrum Control™	Redbooks (logo)  ®
DS4000®	IBM Spectrum Protect™	Sametime®
DS8000®	IBM Spectrum Storage™	Storwize®
Easy Tier®	Informix®	System i®
eServer™	iNotes®	System Storage®
FlashCopy®	Lotus®	Tivoli®
FlashSystem™	Lotus Notes®	XIV®
HyperFactor®	MicroLatency®	z/OS®

The following terms are trademarks of other companies:

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

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

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

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

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

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

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

Find and read thousands of IBM Redbooks publications

- ▶ Search, bookmark, save and organize favorites
- ▶ Get personalized notifications of new content
- ▶ Link to the latest Redbooks blogs and videos

Get the latest version of the Redbooks Mobile App



Download
Now

Android



Promote your business in an IBM Redbooks publication

Place a Sponsorship Promotion in an IBM® Redbooks® publication, featuring your business or solution with a link to your web site.

Qualified IBM Business Partners may place a full page promotion in the most popular Redbooks publications. Imagine the power of being seen by users who download millions of Redbooks publications each year!



ibm.com/Redbooks

About Redbooks → Business Partner Programs

THIS PAGE INTENTIONALLY LEFT BLANK

Preface

This IBM® Redbooks® publication provides best practice guidance for planning, installing, configuring, and employing the IBM TS7600 ProtecTIER® family of products. It provides the latest best practices for the practical application of ProtecTIER Software Version 3.4. This latest release introduces the new ProtecTIER Enterprise Edition TS7650G DD6 model high performance server. This book also includes information about the revolutionary and patented IBM HyperFactor® deduplication engine, along with other data storage efficiency techniques, such as compression and defragmentation.

The IBM System Storage® TS7650G ProtecTIER Deduplication Gateway and the IBM System Storage TS7620 ProtecTIER Deduplication Appliance Express are disk-based data storage systems: The Virtual Tape Library (VTL) interface is the foundation of ProtecTIER and emulates traditional automated tape libraries.

Notes:

- ProtecTIER GA Version 3.4 was released with only the Virtual Tape Library (VTL) interface support. File System Interface (FSI) support was added to ProtecTIER PGA 3.4 Version. For details, see the announcement letter:

http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/6/872/ENUSAP16-0076/index.html&lang=en&request_locale=en

- ProtecTIER Version 3.4 no longer supports OpenStorage (OST).

For your existing ProtecTIER solution, this guide provides best practices and suggestions to boost the performance and the effectiveness of data deduplication with regards to your application platforms for your VTL and FSI (systems prior to version 3.4).

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

This Redbooks publication provides expertise that was gained from an IBM ProtecTIER System Client Technical Specialist (CTS), Development, and Quality Assurance teams.

This planning should be done by the Sales Representative or IBM Business Partner, with the help of an IBM System CTS or IBM Solution Architect.

Note: The following ProtecTIER products are withdrawn; they cannot be directly ordered and are not supported by ProtecTIER v3.4:


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

For more information, search at the following web page:

<http://www.ibm.com/common/ssi/index.wss>

Authors

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




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



Mara Miranda Bautista is a Level 3 Support Engineer for the ProtecTIER team in Guadalajara, at IBM Mexico Software Lab. Mara has eight years of experience in the storage systems field. She joined IBM in 2011, and has worked with the Level 3 Support team ever since. Mara holds a B.A. in Computer Science from Universidad Autonoma de Baja California in Mexico.



Emmanuel Barajas Gonzalez has a bachelor's degree in Electronics and Computer Science and a master's degree in Information Technologies. He joined IBM in 2008 and has participated in several projects that have involved both Development and Support of IBM products related to massive storage of data for both IBM Mainframe and Open Systems. Emmanuel has been awarded 30 patents for his innovations in many areas of the technology and with different applications. Additionally Emmanuel has worked in on several initiatives related to big data and specifically in Apache Hadoop. He is currently working on DataScience and Machine Learning projects in order to generate new and more efficient implementations of current processes.



Carlos A. Martinez Vazquez is a Test Engineer at the IBM Mexico Software Lab, where he joined in 2010. With eight years of experience in the Quality Assurance field, he currently works as the Team Lead of the Quality Assurance team in IBM Guadalajara. He began his career at IBM working on ProtecTIER products as a Test Engineer, where he performed testing for Virtual Tape Library (VTL), OpenStorage (OST), File System Interface (FSI), Simple Network Management Protocol (SNMP), among other ProtecTIER functionalities. After four years on the testing team Carlos moved to Level 3 support where he spent one year, then returned in 2015 to the Quality Assurance team as its team leader. Carlos holds a bachelor's degree in Computer Science from the Universidad de Guadalajara, and he an ISTQB Certified Tester with an Agile testing specialization.

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

Eva Balogh, Jiang Wen Cao, Thomas Chandler, Javier A. Guzman Cortazar, Erik Franz, Rosane Goldstein Langnor, Benji Lee, Camilo Augusto Espinosa Martinez, Jesus Arturo Serrano Preciado, James Reichardt, Ren Ting Tong, Mervyn Venter, Ming Yang

IBM Systems

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

- ▶ Third edition, published August 2014:
Mara Miranda Bautista, Jose Roberto Mosqueda Mejia, Rosane Goldstein Langnor, Karen Orlando,
- ▶ Second edition, published May 2013:
Mathias Defiebre, Dan Riedel, Daniel Wendler, Karen Orlando
- ▶ First edition, published August 2012:
Mathias Defiebre, Adriana Pellegrini Furnielis, Jane Lau, Libor Miklas, Angela Pholphiboun, Karen Orlando

Now you can become a published author, too!

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

- ▶ Use the online **Contact us** review Redbooks form found at:
ibm.com/redbooks
- ▶ Send your comments in an email to:
redbooks@us.ibm.com
- ▶ Mail your comments to:
IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:
<http://www.facebook.com/IBMRedbooks>
- ▶ Follow us on Twitter:
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>

Summary of changes

Since the previous edition of this book, technical changes have been made.

Summary of Changes
for SG24-8025-03
for IBM ProtecTIER Implementation and Best Practices Guide
as created or updated on August 16, 2016.

June 2016, Fourth Edition

This fourth edition reflects the information that is implemented for the release of IBM ProtecTIER, which is based on release 3.4.

This edition includes new and changed information. It also includes minor corrections and editorial changes that are not identified.

New information

ProtecTIER introduces a new hardware refresh feature, the TS7600 3958 DD6 model family, which includes the IBM System Storage TS7650G ProtecTIER Deduplication Gateway and the IBM System Storage TS7620 ProtecTIER Deduplication Appliance Express.

The DD6 model is a higher performance server available in March 2016. The new ProtecTIER Enterprise Edition v3.4 TS7650G TS7600 3958 DD6 Model integrates the service interface into the offering. The TS7600 DD6 model family, together with IBM ProtecTIER Enterprise Edition v3.4, introduces these items:

- ▶ A clustered system architecture that integrates both nodes into a single chassis, eliminating the external switches required to cluster previous gateway models.
- ▶ An integrated service interface, eliminating the need for a separate TS3000 Service Console for the TS7650G and the TS7620 ProtecTIER Appliance Express DD6.
- ▶ The new DD6 hardware has two, 10 Gb Ethernet (GbE) interfaces per system as an upgrade from 1 Gb. The cables are made of copper, which reduces costs and improves performance.
- ▶ A significant reduction in the amount of rack space required for ProtecTIER TS7650G gateway systems.
- ▶ The new performance tool analyzes ProtecTIER system performance variables and displays a graphic report to the user.
- ▶ The FSI application interface is supported with ProtecTIER PGA 3.4 version.

This Redbooks publication also includes the following updates:

- ▶ Best practices and examples for IBM ProtecTIER TS7650G DD6 ProtecTIER models
- ▶ A new hints and tips chapter that provides guidance and practical examples for the following topics:
 - Remote access with IBM Endpoint Manager for Remote Control (Assist On-site)
 - Operating system recovery functionality
 - Online file system check (FSCK)
 - Dedup estimator

- Local directory cloning for File System Interface (FSI)
- Commands to send/receive files to/from IBM Support
- Graphical performance tool
- ProtecTIER licensing tips
- How to reconfigure a IBM Spectrum™ Protect (formerly IBM Tivoli® Storage Manager) instance attached to ProtecTIER after a reboot

Changed information

This edition includes the following changed information.

- ▶ Removed support as of ProtecTIER V3.4 for OpenStorage (OST). Systems that are OST enabled must not be upgraded to v3.4.
- ▶ Added support for IBM Spectrum Protect™ (as of version 7.1.3, formerly Tivoli Storage Manager)
- ▶ This initial release of ProtecTIER V3.4 supports VTL only. FSI interface support will be added in the next ProtecTIER PGA version. This book has been updated to reflect this.
- ▶ Updated best practices to support latest versions of IBM storage products and applications.
- ▶ The TS7650G DD6 and the TS7620 ProtecTIER Appliance Express do not support the IBM TS3000 System Console (TSSC), however all the functionalities provided by the TSSC are now embedded in the code of the DD6.



Part 1

General best practices

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

This part describes configuration options regarding host attachment and contains the following chapters:

- ▶ Chapter 1, “IBM ProtecTIER basics” on page 3
- ▶ Chapter 2, “Deduplication considerations” on page 21
- ▶ Chapter 3, “Networking essentials” on page 37
- ▶ Chapter 4, “Virtual Tape Library guidelines” on page 49
- ▶ Chapter 5, “ProtecTIER File System Interface: General introduction” on page 65
- ▶ Chapter 6, “Host attachment considerations for Virtual Tape Library” on page 79

Note: ProtecTIER GA Version 3.4 was released with only the Virtual Tape Library (VTL) interface support. File System Interface (FSI) support was added to ProtecTIER PGA 3.4 Version. For details, see the announcement letter:

http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/6/872/ENUSA P16-0076/index.html&lang=en&request_locale=en



IBM ProtecTIER basics

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

This chapter introduces the ProtecTIER engine for deduplication, which is IBM HyperFactor. It also describes remote support and call home capabilities, and the Simple Network Management Protocol (SNMP) functionality, for the TS7650G DD5 and the TS7650G DD6 ProtecTIER models.

This chapter describes the following topics:

- ▶ Terminology used in the ProtecTIER environment
- ▶ ProtecTIER fundamental concepts
- ▶ ProtecTIER models
- ▶ Remote support and call home
- ▶ Simple Network Management Protocol (SNMP)
- ▶ ProtecTIER command-line interface
- ▶ Performance considerations for repository creation

1.1 Terminology

These terms are used in this book and are common for all products in the ProtecTIER family:

ProtecTIER	When used alone, this expression points to the IBM patented deduplication solution based on HyperFactor. Depending on the context, it can mean the family of products, a specific device, or just the deduplication engine.
TS7600	When used alone, this term signifies the IBM Spectrum Storage™ family of virtualization solutions that operate on the ProtecTIER platform, including the ProtecTIER appliances and gateways.
Factoring ratio	The factoring ratio is the ratio of nominal capacity to physical capacity in the ProtecTIER repository. For example, if you have 100 terabytes (TB) of user data (nominal capacity) and it is stored on 10 TB of physical capacity, your factoring ratio is 10:1.
HyperFactor	The patented IBM algorithm that eliminates data duplicates from the incoming backup data streams. The factoring ratios are the result of HyperFactor processing and compression.
VTL	Virtual Tape Library. The ProtecTIER VTL service emulates traditional tape libraries. By emulating tape libraries, you can use ProtecTIER VTL to migrate to disk backup without having to replace your entire backup environment.
Shelf	A container of VTL cartridges in a ProtecTIER repository. This container is analogous to a shelf or a rack where physical tapes are kept outside of the automated tape library (ATL) in cartridge slots. This container is only applicable with the VTL configuration option, and is not available in FSI configuration option.
FSI	File System Interface. This configuration option enables ProtecTIER to present disk repository storage as a virtualized hierarchy of file systems.
FSI-CIFS	<p>FSI Common Internet File System. ProtecTIER emulates Windows file system behavior and presents a virtualized hierarchy of file systems, directories, and files to Windows CIFS clients.</p> <p>When configured for FSI-CIFS, ProtecTIER emulates a network-attached storage (NAS) backup target capable of using both HyperFactor and ProtecTIER native replication bandwidth reduction techniques for storing and replicating deduplicated data.</p>
FSI-NFS	<p>FSI Network File System. ProtecTIER emulates UNIX file system behavior and presents a virtualized hierarchy of file systems, directories, and files to UNIX-based clients using the NFS protocol.</p> <p>When configured for FSI-NFS, ProtecTIER emulates an NAS backup target that can use both HyperFactor and ProtecTIER native replication bandwidth reduction techniques for storing and replicating deduplicated data.</p>
Metadata	Metadata is the information used to track the user data sent from the backup servers, including where it is stored on the disk.
User data	User data consists of backup files and data sets that are stored in the ProtecTIER repository. It is the data that the backup applications are storing on disk.

Repository	The repository is the physical disk that holds the ProtecTIER factored data. Two types of file systems make up the ProtecTIER Repository: Metadata User data
Front-end	The connection from ProtecTIER to the backup server.
Back-end	The connection from ProtecTIER to the attached disk storage subsystem that acts as a repository.
Node and server	A single ProtecTIER system. It can be either a TS7650G or TS7620 or one of the previous models, and is represented as a node from the ProtecTIER Manager software. Stand-alone node or dual-node cluster configurations are available. This book uses the terms <i>node</i> and <i>server</i> interchangeably.
IBM Assist On-site (AOS)	A web-based tool that enables a remote support representative in IBM to view or control the management node desktop. On ProtecTIER V3.4 this tool is embedded in the ProtecTIER node. See the AOS web page: http://www.ibm.com/support/assistsite

1.1.1 ProtecTIER Appliance terminology

The System Storage TS7620 ProtecTIER Appliance Express is the IBM self-contained virtualization solution that includes an embedded, pre-configured disk storage repository. The following terms are specific to the ProtecTIER Appliance Express:

3959 SM2	The TS7620 ProtecTIER Appliance Express has a base unit that has two capacity versions (6 TB and 12 TB), and it can have up to two expansion units providing more capacity (23 TB and 35 TB).
3959 EXP	The 3959 EXP expansion drawer enhances the capacity and improves the performance of the TS7620 ProtecTIER Appliance Express. The base unit with one field expansion drawer offers 23 TB repository capacity. The base unit with two field expansion drawers offers 35 TB of repository capacity.

1.1.2 ProtecTIER Gateway terminology

ProtecTIER TS7650G is the IBM virtualization solution that does not include a disk storage repository. The customer can choose a solution from various storage options to build the back-end disk repository.

Two types of servers can be used in the TS7650G:

3958 DD5	This server is available since May 2012, and is included with ProtecTIER Version 3.2 or higher. This server is based on the IBM System x7145 model. When used as a server in the TS7650G, its machine type and model are 3958 DD5. Use this machine type and model for service purposes.
3958 DD6	This server, available since March 2016, was introduced with ProtecTIER Version 3.4. This server is a 2u platform based on two Intel Xeon Processor E5-2695 v2 (12 cores, 2.4 GHz); it includes 128 GB RAM (8x 16 GB DDR3 at 1866 MHz).

Other terms to be aware of are as follows:

Disk controller	<p>The customer must choose a disk controller for use with the TS7650G. A list of compatible controllers is on the IBM Tape storage web page:</p> <p>http://www.ibm.com/systems/storage/tape/library.html</p> <p>Click the TS7650G ProtecTIER model, and the detailed list of specifications and description for that model is shown.</p>
Disk expansion unit	<p>The customer must choose a disk expansion unit for use with the TS7650G. A list of compatible expansion units is on the IBM Tape storage web page:</p> <p>http://www.ibm.com/systems/storage/tape/library.html</p> <p>Click the TS7650G ProtecTIER model, and the detailed list of specifications and description for that model is shown.</p>

1.1.3 ProtecTIER replication terminology

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

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

Replication terminology

These terms define replication in a ProtecTIER context, and apply to VTL and FSI environments:

ProtecTIER Replication Manager	<p>A software component that is installed on a ProtecTIER server. The ProtecTIER Replication Manager remotely manages the configuration of the Replication Grid (for example, grid creation and deletion, and repository membership in the grid).</p> <p>An agent on each ProtecTIER server interacts with the ProtecTIER Manager, and maintains a table of its grid members.</p>
Replication grid	<p>A set of repositories that share a common Grid ID and can potentially transmit and receive logical objects through replication from and to the members of the grid. A replication grid includes up to 24 ProtecTIER repositories and the connections between them.</p> <p>The replication grid is configured using the ProtecTIER Replication Manager.</p>

Replication group	Also known as a <i>topology group</i> . A replication group defines the relationships between nodes in a grid, and determines which nodes are allowed to replicate to other nodes. Currently, these are types of Topology groups: VTL many-to-one, VTL bidirectional (many-to-many), and FSI bidirectional.
Many-to-many topology group	Bidirectional replication among multiple ProtecTIER nodes. Each node can define multiple replication targets, with up to 8 nodes per group in an FSI topology, and up to 4 nodes per group in a VTL topology.
Replication grid ID	A number, 0 - 63, that identifies a replication grid. The ID numbers are not recycled.
Replication grid member	A repository that is a member in a replication grid.
Replication pairs	Two repositories in a replication grid that replicate from one to another.
Repository unique ID (RID)	A number that uniquely identifies the repository. The RID is created from the replication grid ID and the repository internal ID in the grid.
Replication time frame	A scheduled period for replication to take place for all policies.

VTL terminology only

The following terms apply only to VTL:

Many-to-one topology group	For VTL, this configuration is also known as a <i>Hub</i> . Up to 12 ProtecTIER systems (spokes) can replicate to a single ProtecTIER node (hub). If the hub is a TS7620 ProtecTIER Appliance Express, the maximum number of spokes that can be connected is four.
Hub	For VTL, this item is a replication target. It receives replicated data from up to 12 spokes in a many-to-one replication group.
Spoke	For VTL, this item is a backup source that can replicate only to a single hub in many-to-one replication groups.
Replication policy for VTL	A replication policy is defined on a ProtecTIER VTL server, and is made up of rules that define a set of objects (for example, VTL cartridges) from a source repository to be replicated to one or more target repositories.
Visibility	This term represents whether an application backup server can see or has visibility to a VTL cartridge. This construct is unique to VTL technology. ProtecTIER ensures that a tape is accessible by backup server only in one place at a time.
Visibility switching	The automated process that transfers the visibility of a VTL cartridge from its origin to its replica and vice versa. Visibility switching is defined in the replication policy. The process is triggered by moving a cartridge to the source library import/export (I/E) slot.

	<p>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.</p>
Principality/ownership	<p>An attribute that is set at the repository where an individual cartridge can be updated or written to by a backup application. A cartridge at its principal repository can be in read/write (RW) or read-only (RO) mode. At other sites, it is always RO. Each cartridge has enabled principality/ownership for one site only.</p>
Dirty bit	<p>The dirty bit attribute (in-sync) helps identify consistency points during disaster recovery (DR). When the dirty bit is off for a cartridge at the Hub, this cartridge is fully synchronized with the spoke.</p>
Disaster recovery	<p>DR is the process of recovering production site data at a remote location. It includes a way to indicate to a remote repository that the production site went down, and notifies an administrator to initiate a data recovery process.</p>
Failover	<p>Failover is a process of enabling the production at a remote site when there is a critical event or disaster at the primary site. It can be initiated intentionally if the primary site is under threat of a catastrophe and it is beneficial to perform takeover at the remote site with full control.</p>
Failback	<p>A process that is initiated from the remote site when the source site is again available for operation. The process ensures that the paired repositories are resynchronized using the least amount of bandwidth and maintaining the most recent copies of backups.</p>

FSI terminology only

The following terms apply only to FSI:

Replication policy for FSI	<p>A replication policy that is defined on a ProtecTIER FSI server and is made up of rules that define a set of objects (for example, FSI directories) from a source repository to be replicated to one or more target repositories.</p>
Remote destination directory	<p>Applies only to FSI replication. It is a dedicated directory at the remote replication destination. Used to replicate a file system's directories, and all objects that are contained in those directories, recursively.</p>
Cloning	<p>Cloning creates a space-efficient, writable point-in-time copy of a replication destination directory (RDD). Cloning an RDD is used for DR testing without disrupting ongoing replication to the RDD.</p>

Important: ProtecTIER GA Version 3.4 was released with only the Virtual Tape Library (VTL) interface support. File System Interface (FSI) support was added to ProtecTIER PGA 3.4 Version. For details, see the announcement letter:

http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/6/872/ENUSA/P16-0076/index.html&lang=en&request_locale=en

1.2 ProtecTIER fundamental concepts

ProtecTIER is a Data Protection and Retention (DP&R) product that appears to the backup servers as one of two standard interfaces in V3.4:

Virtual Tape Library	A VTL is a standard tape library with a robot, cartridges, and tape drives with support for <i>LAN-free</i> data streams from the hosts.
File System Interface	The ProtecTIER FSI emulates file system behavior and presents a virtualized hierarchy of file systems, directories, and files to clients.

Important: The OpenStorage (OST) interface is not supported by ProtecTIER V3.4.

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

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

Figure 1-1 on page 10 shows the relationship between HyperFactor (HF) processing and the respective storage savings.

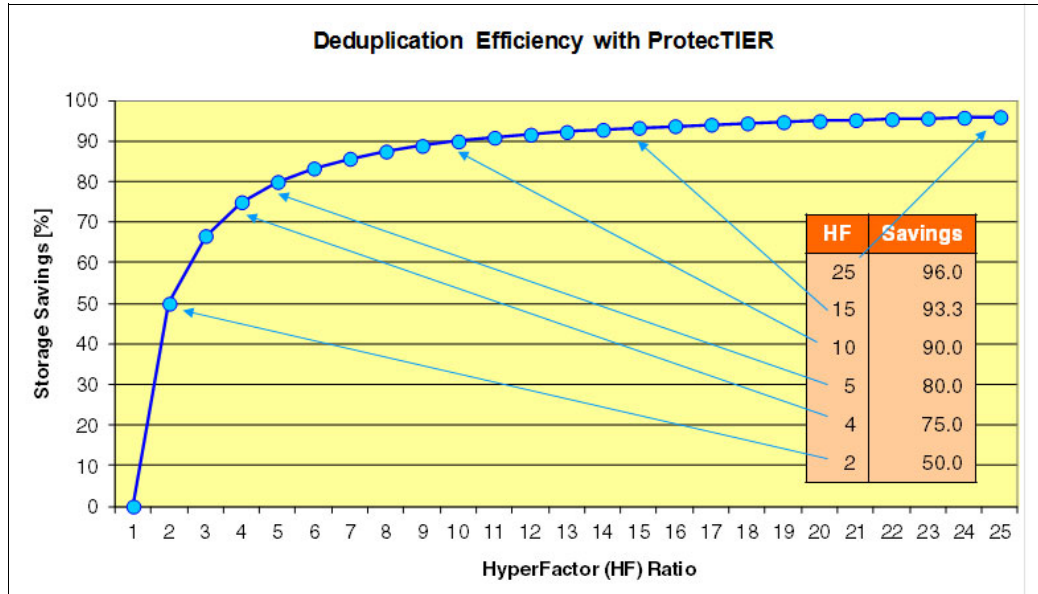


Figure 1-1 Storage savings with ProtecTIER HyperFactor

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

Two key elements affect your factoring ratio and the effectiveness of your ProtecTIER environments:

Data retention period

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

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

Data change rate

The average percentage of data that is received from the backup application that changed from the previous backup iteration. This parameter is important when you size the backup repository and planning for optimal backup windows of various applications or systems.

Examples of typical applications to be sized are progressive forever incremental backups and online database backups (IBM Lotus®, IBM Domino®, Oracle, SAP, IBM DB2®, Microsoft SQL, MaxDB, and so on). The data change rate might vary from platform to platform (from 1% to greater than 25%).

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

For more information about the concepts of IBM ProtecTIER deduplication products, see *IBM System Storage TS7600 with ProtecTIER Version 3.3*, SG24-7968, or see the Tape storage web page:

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

1.2.1 ProtecTIER with deduplication disabled

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

For certain scenarios, especially in small to medium business environments, clients might want to use existing storage area networks (SANs) and disk storage subsystems for backup purposes by using a VTL, but without deduplication techniques.

This requirement is valid for backup solutions with short-term retention periods, where the high frequency and amount of data that is changed on the hosts daily makes restore performance from physical tape not optimal. This situation includes, for example, the backup of large file servers with millions of small files.

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

Note: Excluding the capacity license in ProtecTIER appliances (SM2) is not possible where minimal available capacity must be always chosen.

1.3 ProtecTIER models

This section provides an overview of the existing models of ProtecTIER servers currently marketed. For more technical details about each model, see the following web page:

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

1.3.1 ProtecTIER appliance

The ProtecTIER appliance is designed and built with the repository preconfigured on internal disk drives. The available ProtecTIER appliance is the 3959 SM2 (IBM System Storage TS7620 ProtecTIER Deduplication Appliance Express).

IBM System Storage TS7620 ProtecTIER Deduplication Appliance Express

Available in four configuration options, the TS7620 model 3959-SM2 is an integrated server and storage hardware platform that is included with IBM ProtecTIER deduplication software, preinstalled. The appliance is targeted at small or medium business (SMB) backup environments.

The available repository capacities are as follows:

- ▶ TS7620-SMALL 6 TB / 5.5 TiB (Capacity 1)
- ▶ TS7620-MEDIUM 12 TB / 10.5 TiB (Capacity 2)
- ▶ TS7620 with 1 EXP 23 TB / 21.0 TiB (Capacity 3)
- ▶ TS7620 with 2 EXP 35 TB / 31.5 TiB (Capacity 4)

Clients can choose between VTL or FSI (for the PGA release of ProtecTIER v3.4) configuration as a backup target.

The TS7620 is shown in Figure 1-2.



Figure 1-2 IBM System Storage TS7620 ProtecTIER Deduplication Appliance Express

The TS7620 features and hardware summary include the following items:

- ▶ VTL or FSI support.
- ▶ Emulates up to 12 virtual tape libraries (in its larger configuration).
- ▶ Emulates up to 256 tape drives.
- ▶ A single 3U integrated appliance in 6 TB (5.5 TiB) and 12 TB (10.5 TiB) physical capacities.
- ▶ Up to two additional expansion units for 23 TB (21.0 TiB) and 35 TB (31.5 TiB) physical capacities.
- ▶ 48 GB of RAM.
- ▶ At the time this publication was written, official performance benchmarks are still pending for FSI.

The ProtecTIER Deduplication Gateway products offer clients the capability to use their existing IBM or third-party disk storage subsystem and SAN infrastructure as a repository. For a list of supported disk storage subsystems, see the IBM SSIC and ProtecTIER ISV Support Matrix. For details, see Appendix B, “ProtecTIER compatibility” on page 457.

IBM System Storage TS7650G ProtecTIER Deduplication Gateway

The TS7650G is available in two server models: the 3958-DD5 (Figure 1-3 on page 13) and the 3958-DD6 (Figure 1-4 on page 13). They can be configured as a stand-alone (one 3958-DD5 or one 3958-DD6), or as a cluster (two 3958-DD5 or two 3958-DD6; they cannot be combined in a cluster). The disk storage array attaches to the TS7650G through Fibre Channel (FC) connections and holds the repository of deduplicated backup data.



Figure 1-3 IBM System Storage TS7650G ProtecTIER Deduplication Gateway 3958-DD5



Figure 1-4 IBM System Storage TS7650G ProtecTIER Deduplication Gateway 3958-DD6

The TS7650G 3958-DD6 features and hardware summary include the following items:

- ▶ VTL or FSI support
- ▶ Emulates up to 16 virtual tape libraries per repository
- ▶ Emulates up to 256 tape drives per node
- ▶ Scales to 1 PB of physical back-end storage
- ▶ Server requires only 2U of rack space
- ▶ 128 GB of RAM
- ▶ One solid-state drive (SSD) of 128 GB
- ▶ For VTL performance, up to the following MBps (using 128 drives):
 - 2,500 MBps backup
 - 3,200 MBps restore
- ▶ At the time of writing, official performance benchmarks are still pending for FSI.

Note: The performance numbers published in this book are the official numbers:

<http://www.ibm.com/systems/storage/tape/ts7650g/index.html>

Depending on the environment's configuration, performance numbers might be exceeded.

1.4 Remote support and call home

The IBM Call Home feature reports failures that are detected by the ProtecTIER servers. In previous ProtecTIER models (such as TS7650G DD5 and TS7650 Appliance), the Call Home capability used the IBM TS3000 System Console (TSSC) to send the Call Home alerts; in the TS7650G DD6 and in the TS7620 ProtecTIER Appliance Express this feature is embedded in

the ProtecTIER server. The TS7650G DD6 and the TS7620 ProtecTIER Appliance Express use the ProtecTIER management Ethernet network to send the Call Home alerts to IBM Service Center (home).

When a failure is detected, Call Home sends detailed error information to the IBM Service Center (home).

Reminder: The TS7650G DD6 and the TS7620 ProtecTIER Appliance Express do not support TSSC.

An IBM Service Support Representative (SSR) can then prepare an action plan to handle the problem before he travels to the affected installation.

The ProtecTIER server might also periodically send support information (such as configuration, code versions, and error logs) to IBM. Doing so increases the speed of the problem determination and fault resolution process. Call Home is enabled and tested by SSRs during the initial system installation.

When the reliability, availability, and serviceability (RAS) software on the ProtecTIER server detects a problem, it initiates a Call Home operation to create a problem management record (PMR) in RETAIN. If the error indicates a problem with a field-replaceable unit (FRU), an SSR can pre-order the affected unit for optional replacement at the site.

Important: RAS software on the ProtecTIER server must be configured before setting up Call Home.

1.5 Simple Network Management Protocol (SNMP)

SNMP is not configured by default on ProtecTIER; enable it whenever the infrastructure where the ProtecTIER is installed allows it.

SNMP, when configured on a ProtecTIER system, can be used to send a problem notification to designated recipients in the event of hardware or software degradation or failure. It can also be used for warnings and events in the ProtecTIER system according to the configuration parameters.

SNMP notifications, or traps, can be sent even if the ProtecTIER Manager interface is unavailable. Using SNMP traps requires that you have the following items:

- ▶ SNMP trap receiver software installed on an SNMP trap server. Follow the instructions from the manufacturer to install and configure the SNMP trap receiver software.
- ▶ The file name and location of the Management Information Base (MIB) file for the SNMP trap receiver. On the ProtecTIER server, the file name is `IBM-TS7600-SNMP-MIBV2.mib`, which is located in the `/usr/share/snmp/mibs` directory. The full path to the file is `/usr/share/snmp/mibs/IBM-TS7600-SNMP-MIBV2.mib`.
- ▶ The `IBM-TS7600-SNMP-MIBV2.mib` file must be copied onto the SNMP trap receiver, and the trap receiver software must point to the directory location of the MIB file for translation of the trap messaging.
- ▶ SNMP trapping enabled on one or more of your ProtecTIER servers.

For further details about how to enable and configure the SNMP traps for a specific model, review the relevant User's Guide. For the general configuration process, see the *User's Guide for VTL Systems for IBM System Storage TS7600 with ProtecTIER V3.4*, GA32-0922.

The following sections have useful information and tips about several options that you can define when configuring your SNMP traps.

1.5.1 Alert filtering

ProtecTIER SNMP traps are grouped in three categories:

- ▶ Informational
- ▶ Warnings
- ▶ Errors

You can filter the traps by selecting only the ones that you want to receive.

As shown in Figure 1-5, you can click the **Alert Filtering By Severity** option in the left pane of the Configuration wizard and select the alerts that you want to receive.

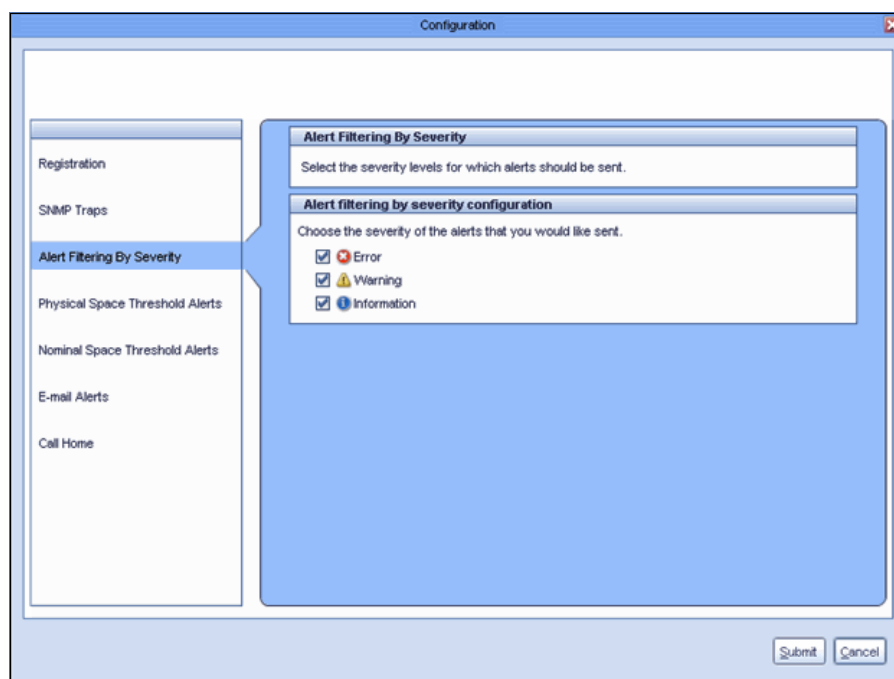


Figure 1-5 Alert Filtering by Severity section from the SNMP Configuration wizard

1.5.2 Space thresholds

ProtecTIER also has two special SNMP alerts for monitoring the free space in the repository; both operate equally:

- ▶ Physical Space Threshold Alerts: Monitors the physical space.
- ▶ Nominal Space Thresholds Alerts: Monitors the nominal space.

Configure these alerts as follows (see Figure 1-6):

1. Click either the **Physical Space Threshold Alert** or the **Nominal Space Threshold Alert** option in the left pane of the Configuration wizard.
2. Activate the threshold by selecting the **Send an alert when the repository enters or leaves the limited physical/nominal space state** check box.

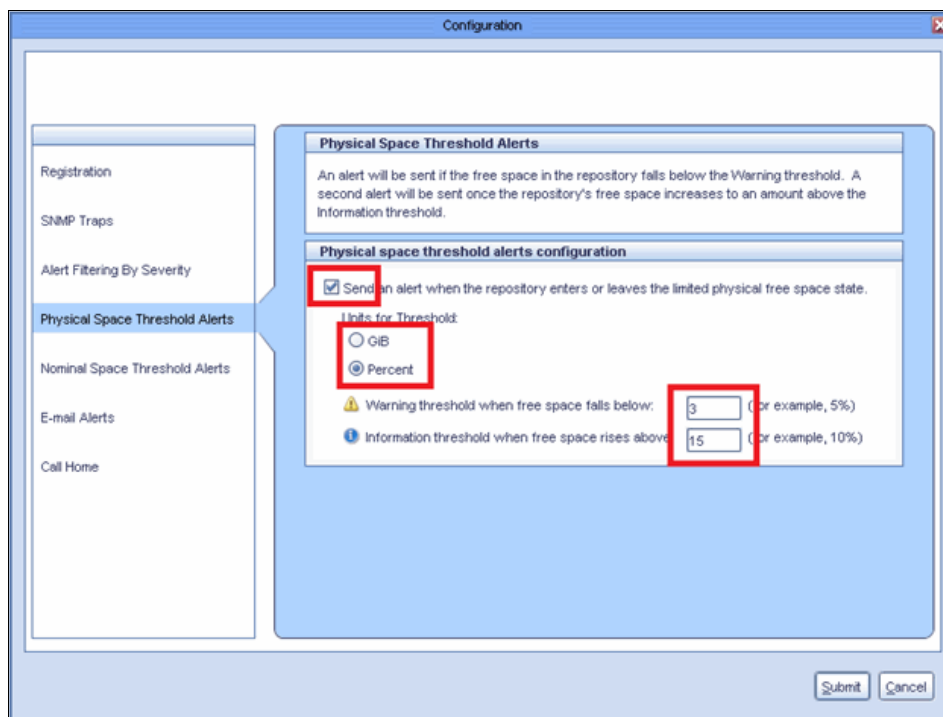


Figure 1-6 Physical Space Threshold Alerts section from the SNMP Configuration wizard

3. After the threshold is active, you can choose which unit you prefer for the threshold, gibibyte (**GiB**) or **Percent**.
4. The next step is to configure the warning threshold. This warning trap will be sent when the free space falls below the configured value.
5. Finally, you need to configure the informational threshold. This informational trap will be sent when the space is released, and only if the warning threshold was previously passed.

Figure 1-7 shows the relationship between both thresholds and the description about the criteria to send these SNMP traps.

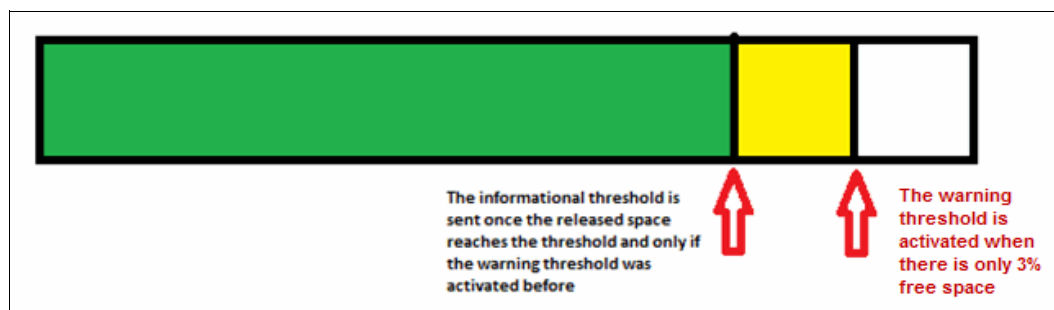


Figure 1-7 SNMP Warning and Informational thresholds description

1.5.3 Email alerts

ProtectTIER systems can be configured to send the whole set of SNMP traps by email. An important note is that the alerts sent using SNMP traps and email are exactly the same. The only difference is the means by which they are monitored. SNMP traps are monitored using SNMP trap receivers, and email alerts are sent directly to an email address.

Enable email alerts as follows (see Figure 1-8):

1. Select the **E-mail Alerts** option in the left pane of the Configuration wizard.
2. Select the **Enable E-mail alerts** check box and add the correct SMTP server IP address in the **SMTP server address** text box. In this case, keep the **Port** value of 25, which is by default the port for SMTP.

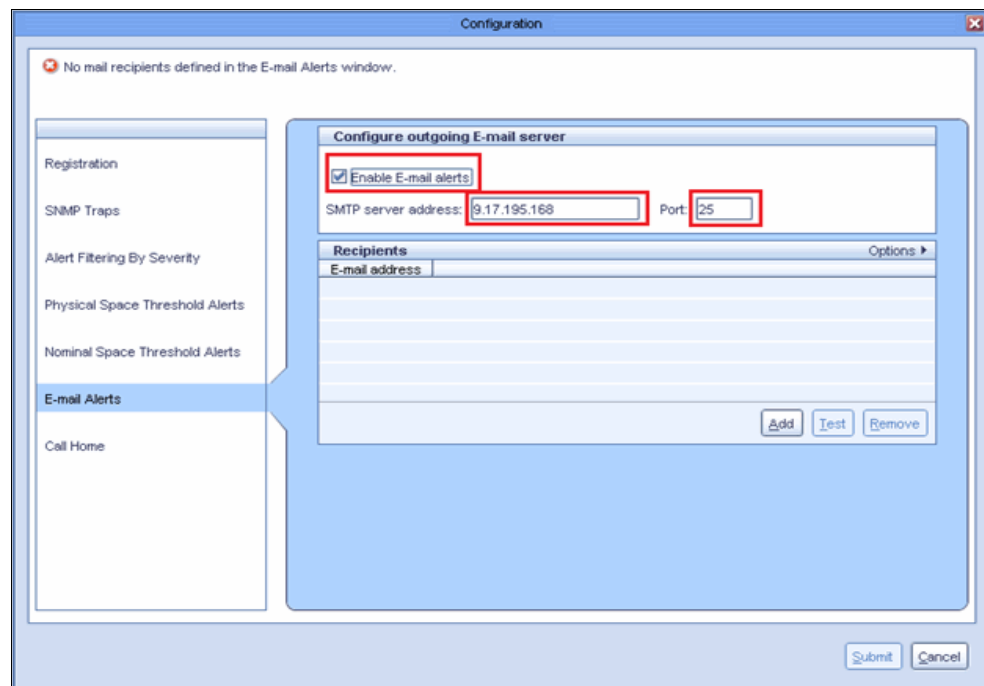


Figure 1-8 E-mail Alerts section from the SNMP Configuration wizard

3. Click **Add** to open the Add a new mail recipient window.
4. Enter a valid email address in the **New mail recipient** text box and click **Add** to add the destination email address.

5. After the destination email address is added, and if everything is correctly configured, you receive a test mail at your destination. Figure 1-9 is an example.

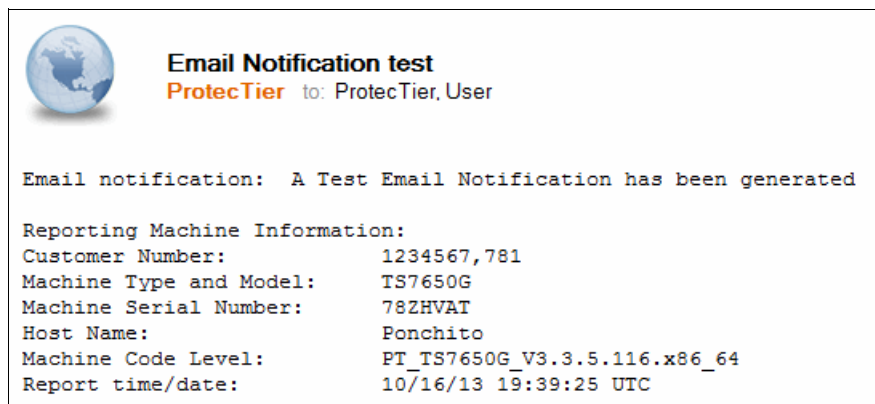


Figure 1-9 Email notification test template

6. Click **Submit** to save your configuration and close the Configuration wizard window.

1.6 ProtecTIER command-line interface

This section provides a brief overview of the ptcli. The ptcli is run from the workstation or on the ProtecTIER node. The ptcli is a low-level entry point into the ProtecTIER system. It provides information that can be formatted and queried to provide wanted data, and direct the output to a file for persistent storage.

This information provides the administrator with data that gives insight into the long-term operation of the system's performance, capacity, and configuration, and the effect of operational changes on the system. The information can be used as input to the administrator's other management applications.

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

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

For more details about how to install and use the ptcli, see the command-line interface section of *IBM System Storage TS7600 with ProtecTIER Version 3.3*, SG24-7968.

1.7 Performance considerations for repository creation

Select the performance and the size of a repository based on current and future needs. This planning should be done by the Sales Representative or IBM Business Partner, with the help of an IBM System Client Technical Specialist (CTS), and IBM Solution Architect.

The performance is restricted by the System Peak Throughput, which is selected when the repository is being created.

Important: The System Peak Throughput *cannot be changed* after the repository is created. For this reason, be sure the number is accurately calculated. As input, use the amount of data that the ProtecTIER is expected to handle concurrently on a daily basis for all the operations (primary and maintenance).

The primary operations are considered to be backups, restores, incoming replication, and outgoing replication. The maintenance operations are considered to be deletes and defragmentation (*defrag*) activities.

This means that the System Peak Throughput should be based on the backups, restores, incoming replication, outgoing replication, deletes, and defrag workloads. You must consider whether these operations will be done at the same time. Will the system be running backups, restores, replication continuously? Or will there be time frames where no backups, restores, and restores are performed? These time frames with no backups, restores, and replication are known as *maintenance* or *idle* windows.

The maintenance operations are a function of the primary operations. Consider that when a repository reaches the *steady state*, the same amount of data that is backed up or received from replication must be deleted. Although no specific formula exists for calculating how much fragmented space is generated when data is deleted, account for at least 35%.

Note: Although no formula exists to calculate how much fragmented space is generated when data is deleted in ProtecTIER, the 35% was selected based on the percentages in which the Maintenance Rate Controller (MRC) assigns resources for the deletes and for the defrags. By default, MRC assigns 65% of the maintenance resources for deletes, and 35% of the maintenance resources for defrags.

The following examples are of what is accounted for by the system performance. Although for some sizings (like the PreSales sizing) the Maintenance Rate Control (MRC) jobs are implicit, these examples explicitly illustrate the overall operations of the ProtecTIER.

1.7.1 Example 1: Calculate performance throughput

If the daily workload requirements are to back up 5 TB, restore 2 TB, and replicate all data, consider for maintenance operations 5 TB of deletions, and 1.75 TB of defrags. These total 18.75 TB, which, when converted, represent a sustained performance required of 225 MBps.

Another way that this requirement can be expressed is from a time frame constraint perspective, that is, having the backups and restores complete in a specific time frame. For example, complete the 5 TB backups and 2 TB restores in 6 hours. To achieve that, a sustained throughput of at least 340 MBps is needed. If the main operations are completed in 6 hours, the system is left with 18 hours to do MRC operations at 340 MBps, which will suffice.

1.7.2 Example 2: Calculate performance throughput

If the daily workload requirements are to back up 20 TB, restore 15 TB, replicate all data, and receive 10 TB of replication, consider for maintenance operations 30 TB of deletes (20 TB from backups, and 10 TB from incoming replication), and 10.5 TB of defrags. These total 105.5 TB, which, when converted, represent a sustained performance required of 1,281 MBps.

Another way that this requirement can be expressed is from a time frame constraint perspective, that is, having the backups, restores, and replication complete in a specific time frame. For example, complete the 20 TB backups, 15 TB restores, and 30 TB of replication in 12 hours. To achieve that, a sustained throughput of at least 1578 MBps is needed. If the main operations are completed in 12 hours, the system is left with 12 hours to do MRC operations at 1578 MBps, which will suffice.



Deduplication considerations

This chapter describes the IBM System Storage TS7600 ProtecTIER family data deduplication concepts, methods, and system components. This chapter also elaborates on the benefits of data deduplication with IBM HyperFactor, and on general ProtecTIER deduplication considerations. This chapter also describes data type candidates for high factoring ratios, and describes data types that can have a negative effect on factoring ratios.

This chapter describes the following topics:

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

2.1 HyperFactor data deduplication

Data deduplication is the process of storing only a single instance of data that is backed up repetitively. It reduces the amount of space that is needed to store data on disk. Data deduplication is not a storage device; it is a function of a system, for example, a Virtual Tape Library (VTL), an OpenStorage (OST) application programming interface (API) for versions of ProtecTIER earlier than v3.4, or a File System Interface (FSI).

Data deduplication is not an input/output (I/O) protocol. However, it does require an I/O protocol for data transfer, such as Fibre Channel Protocol (FCP), Common Internet File System (CIFS), Network File System (NFS), or an API.

Figure 2-1 illustrates how ProtecTIER data deduplication stores repeated instances of identical data in a single instance. This process saves storage capacity and bandwidth. Data deduplication can provide greater data reduction than previous technologies, such as Lempel-Ziv (LZ) compression and differencing, which is used for differential backups.

Using data deduplication does not always make sense because not all types of data can be deduplicated with identical efficiency. Data deduplication might interfere with other technologies, such as compression, encryption, or data security requirements. Data deduplication is not apparent to users and to applications.




Figure 2-1 Simplified data deduplication process

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

The effectiveness of data deduplication depends on many variables, which have a major effect on the effectiveness of data deduplication:

- ▶ The data change rate
- ▶ The number of backups
- ▶ The amount of repetitive or similar data in your backups
- ▶ The data retention period

For example, if you back up the exact same uncompressible data once a week for six months, you store the first copy and do not store the next 24, which would provide a 25:1 data deduplication ratio. If you back up an uncompressible file on week one, back up the exact same file again on week two and never back it up again, you have a 2:1 deduplication ratio.

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

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

Data deduplication can provide storage savings, but the benefit that you derive is determined by your data and your backup policies. Workloads with a high database content generally have the highest deduplication ratios. However, product functions, such as IBM Tivoli Storage Manager incremental forever backup, Oracle Recovery Manager (RMAN), or LiteSpeed for SQL Server, can affect the deduplication ratio.

Compressed, encrypted, or otherwise scrambled workloads typically do not benefit from deduplication, because the potential deduplication candidates are no longer similar. For more information, see 2.6, “Data types” on page 35.

2.1.1 HyperFactor, deduplication, and bandwidth savings

The cornerstone of ProtecTIER is HyperFactor, the IBM technology that deduplicates data inline as it is received from the backup application. ProtecTIER bandwidth-efficient replication, inline performance, and scalability directly stem from the technological breakthroughs inherent to HyperFactor.

HyperFactor is based on a series of algorithms that identify and filter out the elements of a data stream that was stored by ProtecTIER. Over time, HyperFactor can increase the usable capacity of an amount of physical storage by 25 times or more.

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

When a change occurs in the source site of a replication grid, the differences between the original and the copy elements are calculated and only the changes are sent over the replication link. This search is quick, and uses a small and efficient memory-resident index. After similar data elements are found, HyperFactor can compare the new data to the similar data to identify and store only the byte-level changes.

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

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

HyperFactor data deduplication uses a fixed-size 4 gigabyte (GB) Memory Resident Index to track similarities for up to 1 petabyte (PB) of physical disk in a single repository. Depending on the data deduplication ratio for your data, you could store much more than 1 PB of data on your disk array.

For example, with a ratio of 12:1, you can store 12 PB of data on 1 PB of a disk array. With the Memory Resident Index, HyperFactor can identify potentially duplicate data quickly for large amounts of data, and it does this action on data load or inline, reducing the amount of processing required for your data.

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

2.2 ProtecTIER HyperFactor deduplication processing

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

2.3 Components of a ProtecTIER system

The ProtecTIER data deduplication system consists of three main components:

- ▶ ProtecTIER server
- ▶ HyperFactor deduplication software
- ▶ Disk storage subsystem

Two of these components, the ProtecTIER server and the HyperFactor deduplication software, are always bundled together for your convenience. Depending on the model of ProtecTIER that you look at, you might have a bundled disk storage subsystem. Do not share the storage subsystem assigned to ProtecTIER with other applications. For an overview of ProtecTIER models, see 1.3, “ProtecTIER models” on page 11.

The components shown in Figure 2-2 are described in the next three sections.

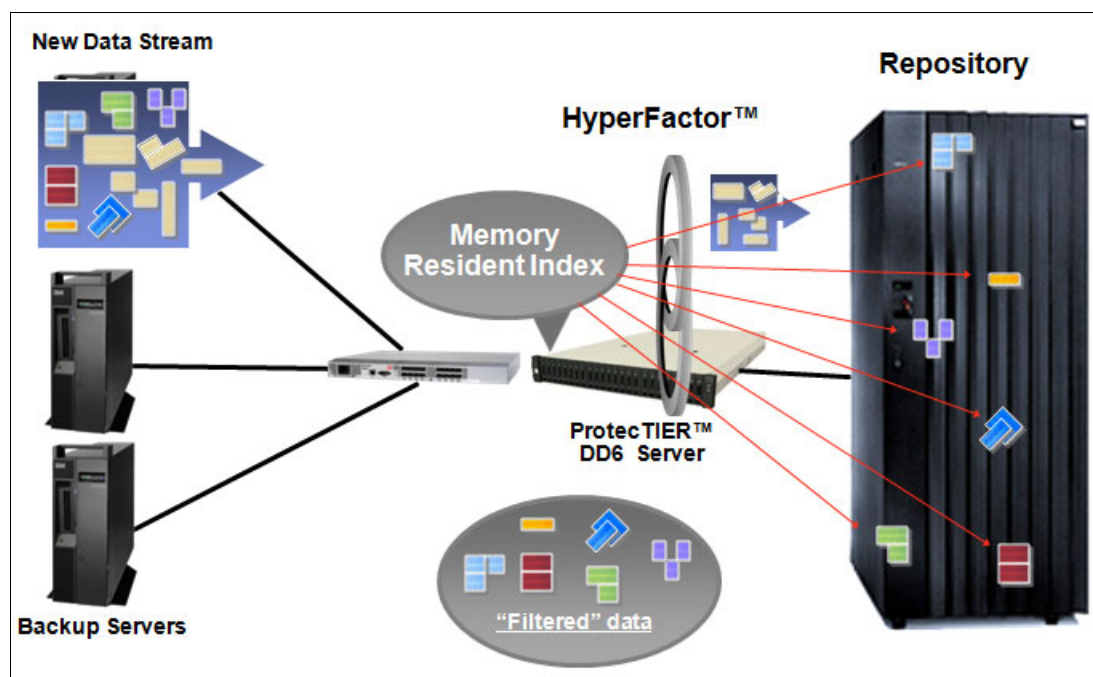


Figure 2-2 ProtecTIER components

2.3.1 ProtecTIER server

Every ProtecTIER deduplication system uses a server with the Red Hat Enterprise Linux (RHEL) operating system on which the HyperFactor software runs.

The IBM System Storage TS7620 ProtecTIER Deduplication Appliance Express (TS7620) comes as a bundle (3959 SM2 server) and LSI MegaRAID system storage) that enables up to 300 megabytes per second (MBps) performance, depending on the ProtecTIER interface and the number of storage expansion drawers that are connected.

Note: The TS7650 Appliance (3958 AP1) and TS7610 Express (3959 SM1) are withdrawn from marketing. For details about support, see the IBM Software Support Lifecycle website and search for ProtecTIER.

The TS7650G server (3958 DD5) is a high-performance configuration. It can support up to 1600 MBps and more, depending on the configuration of the installation environment. A ProtecTIER high availability (HA), active-active, two-server cluster configuration (in a cluster the servers are also known as *nodes*) is available with the TS7650G solution. You can have up to 2500 MBps and more depending on the configuration of the installation environment with a TS7650G two-server cluster.

ProtecTIER version 3.4 introduces the new hardware platform the TS7650G server (3958 DD6) which is also a high performance configuration. The performance characteristics are similar to those of the 3958 DD5 server, however the footprint of the installed system (as shown in Figure 2-2) reduces up to 80% because even the full HA configuration with two nodes occupies only two units in the rack.

2.3.2 HyperFactor deduplication algorithm

The HyperFactor data deduplication solution can be ordered in three interface styles:

- ▶ The VTL interface
- ▶ The OST API (Versions of ProtecTIER earlier than V3.4)
- ▶ The FSI with CIFS and NFS support (PGA v3.4)

Because the interface methods cannot be mixed, you must choose one or deploy multiple ProtecTIER models simultaneously.

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

2.3.3 Disk storage subsystem

Data that is processed by the ProtecTIER HyperFactor data deduplication software is stored on disk. The ProtecTIER appliance system, the TS7620 Appliance Express is pre-bundled with disk storage included in a ready to run configuration.

The ProtecTIER TS7650G Gateway server attaches to a wide variety of disk storage subsystems that separately must be made available. For a list of supported disk systems, see TS7650/TS7650G independent software vendor (ISV) and Interoperability Matrix at the IBM System Storage Interoperation Center (SSIC) web page:

<http://www.ibm.com/systems/support/storage/ssic/interoperability.wss>

Compression: ProtecTIER performs compression after the deduplication process completes unless you decide to turn off compression, which is not recommended. Allocation of disk arrays for ProtecTIER that run their own compression is not recommended, because no additional benefit is gained.

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

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

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

2.4 Benefits of ProtecTIER HyperFactor

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

2.4.1 Flexibility

ProtecTIER is independent from the backup application. You can combine multiple backup applications from different vendors (included in the ProtecTIER compatibility matrix) to work with one single ProtecTIER solution. All attached backup solutions directly benefit from the whole ProtecTIER deduplication potential; sharing the repository with multiple backup applications is possible.

2.4.2 High availability

ProtecTIER offers true highly available, active-active, dual-node clustering for VTL and OST (for versions earlier than V3.4) models. Mimicking the behavior of a physical tape library, when the correct setup is made, you can use the ProtecTIER solution to access your data even if a node is unavailable. The initial configuration of the FSI model is available as only a single node.

2.4.3 High performance, low storage requirements, and lower environmental costs

Data deduplication can reduce the amount of disk storage that is required to store data and keep it online. Performing restores from disk can be faster than restoring from tape, and having the data online for longer periods reduces the possibility that the required data might be sent offsite.

Inline deduplication has no need for more post-processing space, and therefore further reduces space requirements. If data deduplication reduces your disk storage requirements, then the environmental costs for running and cooling the disk storage are also reduced.

2.5 General ProtecTIER deduplication considerations

The following considerations and best practices can help you better understand what to do or not do regarding ProtecTIER deduplication.

2.5.1 Rethinking your overall backup strategy

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

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

Keeping the number of used cartridges low to save money no longer applies in the virtual world of ProtecTIER. Using as many cartridges in parallel as possible, to some extent, is a good idea. The maximum number of cartridges in a VTL is greater than 65,000. You do not need to use all of them, but you should plan on using more virtual resources than you would use physical resources. This guideline is true for virtual tape drives and virtual tape cartridges. This general approach is also true for FSI deployments, and also OST for versions earlier than 3.4.

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

If you perform incremental backups, especially for your databases, you might also want to rethink this process for critical applications. Multiple full backups, especially on a high frequency schedule, might appear to be a waste of space, but this situation is where you can benefit the most from ProtecTIER deduplication.

A ProtecTIER server has the best deduplication, the highest backup speed, and the highest restore speed if you write multiple full backups of the same objects to it. Your infrastructure should be up to the challenge because resources tend to sit idle during non-backup hours. So why not increase the usage of your already available resources?

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

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

2.5.2 Data reduction technologies should not be combined

ProtecTIER data deduplication is a data reduction technology. Compression is another data reduction technology. Tivoli Storage Manager (part of the IBM Spectrum Protect™ family) is an example of an application that provides its own brand of compression and deduplication. Tivoli Storage Manager also offers incremental forever backup, with which only changed data is backed up, so it can be thought of as a data reduction technology. There are many other potential data reduction technologies.

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

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

Some of the scenarios that enable the combination of data reduction technologies are described in this section.

Tivoli Storage Manager can combine both compression and deduplication. Details are explained in the Introduction to IBM Tivoli Storage Manager deduplication chapter of *Implementing IBM Storage Data Deduplication Solutions*, SG24-7888.

IBM DB2 database software can handle data in a way such that it can be compressed in DB2, but still achieves high deduplication ratios. For more information about using DB2 compression with a ProtecTIER repository, see 20.4.1, “Combining DB2 compression and ProtecTIER deduplication” on page 309.

2.5.3 Data streams must be in order

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

Encryption also results in shuffled data. A small change in an encrypted file produces a file that, to a deduplication solution, appears different. Potential deduplication candidates cannot be identified, because the patterns do not match anymore. Analyze your environment for other potential data shuffling causes, and aim to eliminate them.

2.5.4 Data organization in your ProtecTIER repository

The ProtecTIER repository is the place where your deduplicated data is stored. You can define one or many VTLs with multiple slots and cartridges. You can define one or many storage units for the OST API, or you can have multiple file shares for the FSI. No matter what type of ProtecTIER repository you use, logically segment your data and group similar backup types together.

This setup enables detailed deduplication analysis that is based on cartridge granularity that is done by your IBM Service Support Representative (SSR). If you can supply a list of virtual cartridges, or a virtual cartridge range that contains one special type of backed up data for detailed analysis, this setup provides valuable data that you can use to improve your data protection environment.

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

2.5.5 The dynamics of the ProtecTIER repository

In addition to the data that you write into your ProtecTIER repository, there are two other major effects of the ProtecTIER repository that must be understood. First, the repository dynamically reacts to the quality of your data.

If the data that you back up to the ProtecTIER repository suddenly changes and enables a higher deduplication ratio, the repository adapts and can store more data. If the quality of your data changes and enables only a reduced deduplication ratio, the ProtecTIER repository also reacts to this change, and less data can be stored.

Repository size: The ProtecTIER nominal repository size is calculated by using the following formula:

$$\text{Physical Repository Size} \times \text{HyperFactor Ratio} = \text{Available Free Space for you to write to (Nominal Free Space)}$$

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

A ProtecTIER repository is not directly aware of your data retention requirements. A ProtecTIER repository stores all data unless informed otherwise. An especially important point is for VTL emulations to specify whether or not you still need the data.

As an example, IBM Spectrum Control™ (formerly Tivoli Storage Manager) uses the RelabelScratch option of its library definition to communicate to the ProtecTIER repository that the space of a virtual cartridge can be freed up. Other backup applications might rely on housekeeping scripts to initiate a label sequence or write sequence from the beginning of tape, which has the same effect. Make sure to regularly free up unused space.

After you release the unused space, it becomes marked as *Pending* in the ProtecTIER repository. The ProtecTIER repository then automatically uses internal processes to optimize the available space for future backups:

- Internal Deleter processes reduce the Pending space and, in the process, create Fragmented space.
- Internal Defragger processes then reduce the Fragmented space.

In the Capacity section of the ProtecTIER GUI (Figure 2-3), the pie chart at the right side shows the nominal data, and you can see the pending space. In this example, the pending space is 16.9 terabytes (TB).




Figure 2-3 ProtecTIER repository with pending nominal space

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

2.5.6 ProtecTIER repository usage

From a technical standpoint, there is no problem with having a ProtecTIER repository that is 100% used (although from a usability standpoint it is not a good practice to permit a repository to get 100% full). After you reach a steady state where daily housekeeping frees up enough space to enable daily backups, a high usage is possible.

In reality, data tends to grow, so sooner or later you can face changed requirements for your ProtecTIER repository size. You should configure an automated message that informs you when the repository usage crosses a specified threshold value. Depending on the time you need to prepare a capacity upgrade of the ProtecTIER back-end storage, values greater than 80% or 90% can be selected to provide ample time for preparation.

To configure an automated message triggered by a threshold, follow these steps:

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

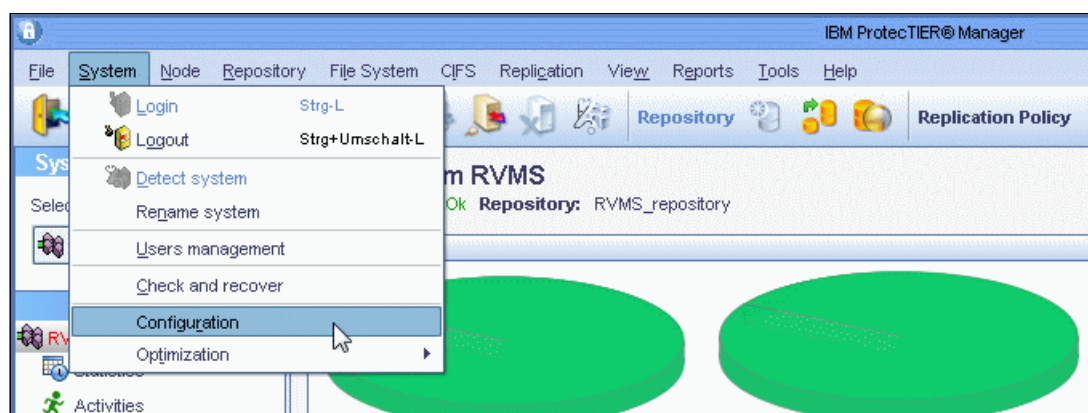


Figure 2-4 ProtecTIER Manager GUI

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

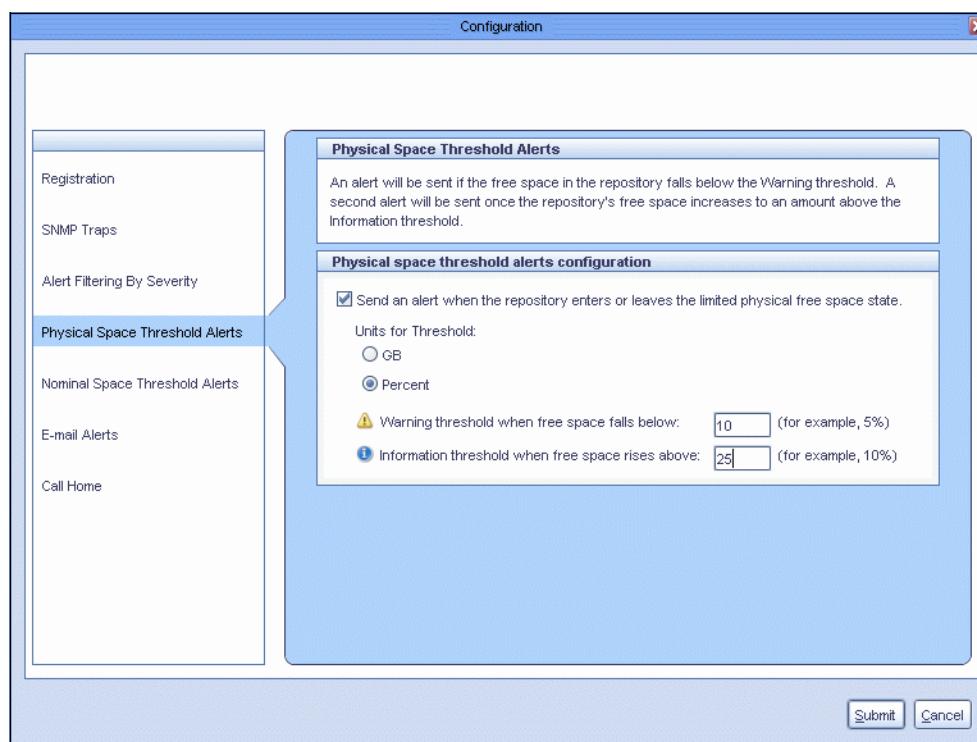


Figure 2-5 Physical Space Threshold Alerts

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

2.5.7 Compression

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

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

2.5.8 Encryption

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

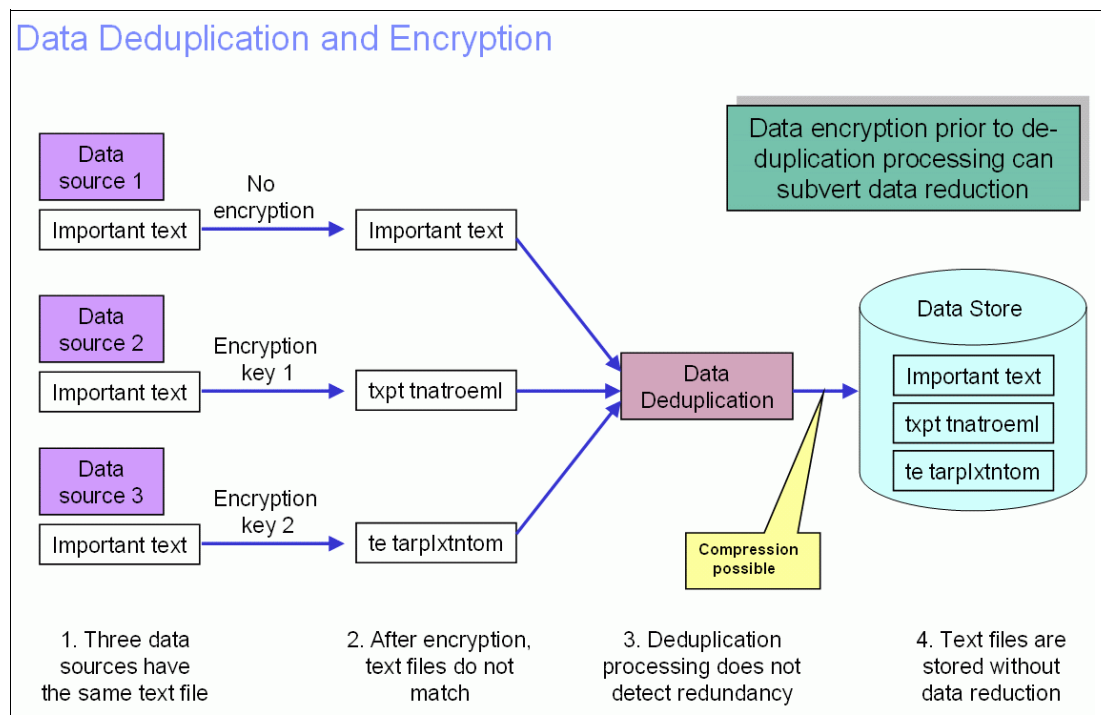


Figure 2-6 Challenges combining encryption with data deduplication

Important:

- If you prefer to run your environment with encryption, consider enabling disk storage-based encryption, for example, IBM System Storage DS8870, which features Full Disk Encryption (FDE).

If you prefer to have client-side encryption enabled, consider using a ProtecTIER solution without deduplication, as described in 2.5.2, “Data reduction technologies should not be combined” on page 29.

- Using an encryption switch between the ProtecTIER server and the storage system has been implemented in the field, but this is only supported under a request for product quotation (RPQ).

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

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

Backing up database logs: You can back up database logs to a ProtecTIER repository without issue, but be aware that it has an effect on your deduplication ratio.

2.5.10 Multiplexing

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

2.5.11 Tape block size

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

2.5.12 File size

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

2.6 Data types

Deduplication is primarily influenced by the type of data that you have. Depending on whether the data is structured to a high degree or unstructured, and possibly already compressed, deduplication yields a higher or lower ratio. For more information about data types, see Chapter 20, “Application considerations and data types” on page 295.

2.6.1 Candidates for a high factoring ratio

Potential candidates for a high deduplication ratio are all kinds of structured data. For example, databases are perfect candidates for deduplication, as is email. Most applications that deal with structured data, such as databases and email, offer some compression to reduce the amount of storage the application data needs.

Because these types of data are good candidates for data reduction in general, many application vendors already have implemented some compression, compaction, or defragmentation. Turning off these application-internal data reduction technologies, or ensuring that they do not affect the backup data stream, enables high deduplication ratios.

For an example of effectively using DB2 compression with a ProtecTIER repository, see 20.4.1, “Combining DB2 compression and ProtecTIER deduplication” on page 309.

2.6.2 Candidates for a low factoring ratio

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

The same situation generally applies to voice or audio data. Formats, such as *.mp3, *.aac, *.ogg, *.wma, or *.m4a, are also compressed. Backing up image files, video files, or audio files to a ProtecTIER repository results in a combination of data reduction technologies. This situation produces low deduplication ratios, because already reduced data cannot be reduced again (for more information, see 2.5.2, “Data reduction technologies should not be combined” on page 29).

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

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



Networking essentials

This chapter describes principal networking considerations, terminology, and concepts for the non-Fibre Channel (FC) network interface for front-end connectivity of the IBM System Storage TS7600 ProtecTIER family. The chapter discusses the File System Interface (FSI) with the options of Network File System (NFS) and Common Internet File System (CIFS) in a non-Virtual Tape Library (VTL) environment. The ProtecTIER server connects to the network with multiple Ethernet adapters, and supports a total throughput of hundreds of megabytes per second (MBps) over the network, per node. The chapter also describes network configuration preferred practices to support Internet Protocol (IP) configurations. It provides the scope and objectives, and main acronyms that are used in the subsequent sections.

Note: ProtecTIER GA Version 3.4 was released with only the Virtual Tape Library (VTL) interface support. File System Interface (FSI) support was added to ProtecTIER PGA 3.4 Version. For details, see the announcement letter:

http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/6/872/ENUSAP16-0076/index.html&lang=en&request_locale=en

This chapter describes the networking technologies that are used to set up the ProtecTIER front-end interfaces to the network. In addition, this section describes *teaming* or *bonding* in the network servers, 802.3ad link aggregation, and similar technologies on the local area network (LAN) switches, in addition to network topologies for standalone and redundant architectures. Some of these networking concepts also apply for ProtecTIER network replication environments.

This chapter describes the following topics:

- ▶ Network terminology
- ▶ General configuration considerations
- ▶ Bonding and teaming
- ▶ Preferred ProtecTIER bonding configuration

Notes:

- ▶ **Important:** Using the ProtecTIER FSI as a network-attached storage (NAS), in terms of storing primary data directly to it, is not supported. ProtecTIER FSI must be used with an application that is supported by ProtecTIER.
- ▶ **Terminology:** This chapter uses the term *bonding* when describing aggregating multiple network links to a single logical interface. You might have heard terms such as *network aggregation*, *link aggregation*, *port trunking*, *trunking*, *link bundling*, *Ethernet bonding*, *network bonding*, *network interface controller (NIC) bonding*, *802.3ad*, or *NIC teaming*, which also describe the same concept.
- ▶ **Cabling:** Multiple members of a single bond must be connected to the same network switch. If you use multiple physically independent networks, similar to a dual-fabric approach in FC terms, you are required to have at least two bonds. Each bond must be cabled to only one network switch.

3.1 Network terminology

The following terms are relevant to network configurations in general:

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

3.2 General configuration considerations

The following list describes several general preferred practices for FSI:

- ▶ Because ProtecTIER backups with FSI are done on a file share level, you should create a dedicated file share for each backup server that you use with ProtecTIER FSI. This file share makes more sophisticated load balancing scenarios possible.
- ▶ Make sure that the backup application runs in the context of the user who mounted the file system with write permission. If you experience access rights issues, this configuration most probably is the reason for them.
- ▶ There must be at least two different network subnets to separate the ProtecTIER management IP interface from the ProtecTIER file system IP interfaces. Otherwise, you cannot ensure that your backup traffic is using the FSI interface rather than the management interface.
- ▶ For FSI workloads, dedicated infrastructures might not exist for incoming and outgoing traffic of backup servers. To be sure that the environment does not suffer from infrastructure congestion, involve the engineers in charge of managing the network.
- ▶ If possible, implement bonding on all involved network devices, whether the devices are the ProtecTIER server, the backup server, or even the network switches. Enabling bonding only on the ProtecTIER server might not be enough to achieve the best results.

3.3 Bonding and teaming

To achieve high availability, load balancing, and increased throughput, you can use a network technology that is known by many names. You might have heard of terms such as network aggregation, link aggregation, port trunking, link bundling, Ethernet bonding, network bonding, NIC bonding, 802.3ad, or NIC teaming. All of these terms describe solutions that you can use to achieve high availability, load balancing, or increased throughput by combining multiple network interfaces and using them as one logical link.

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

Table 3-1 on page 42 lists the available ProtecTIER link aggregation modes, bonding options, and features for each link aggregation mode.

3.3.1 The three bonding modes of ProtecTIER

The ProtecTIER product supports three principal modes of bonding:

- ▶ High availability
- ▶ Round robin
- ▶ Link aggregation (L2, L2L3, L3L4); uses IEEE 802.3ad

For summaries of the bonding methodologies that are supported by ProtecTIER, see Table 3-1 on page 42.

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

Mode one: High availability

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

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

Mode two: Round robin

The Round Robin load balancing method uses a *balance-rr (round-robin) policy*. Outgoing network traffic is distributed across all members of the bond in sequential order. Incoming network traffic is still limited to one single network port (the primary). If one of the network interfaces fails, the other bond members take over. Outgoing traffic is distributed across the remaining bond members.

If the primary adapter for incoming traffic fails, the bond automatically selects a new primary adapter, and incoming traffic is handled from that one single network adapter. With this mode, you can achieve fault tolerance and load balancing. A potential unidirectional bandwidth increase for outgoing traffic is possible if you have multiple backup servers in your environment.

Bandwidth: With the round robin mode, you do not see any bandwidth increase if you have only two communication partners, for example, the ProtecTIER server and one backup server. Even enabling this mode on both of these machines does not enable you to use more than the bandwidth of one single interface, because the incoming traffic of both machines is dealt with only one single network interface.

If you have enough communication partners, for example, multiple backup servers and one ProtecTIER server, the increased bandwidth is used only during parallel restore of multiple backup servers.

Mode three: Link Aggregation (L2, L2L3, L3L4)

The third load balancing method is the most advanced method. If you set it up correctly, you can use the combined throughput of all of the involved network interfaces for incoming and outgoing traffic, balance the load across all available interfaces, and have a fault tolerant system at the same time. This load balancing method uses an IEEE 802.3ad Link Aggregation policy. ProtecTIER offers this mode with a variant in the transmit hash policy, the options are labeled as L2, L2L3, and L3L4.

To use this method, you must have a network infrastructure that fully supports it end-to-end. Since this bonding method relies on the IEEE 802.3ad dynamic link aggregation standard, it is also known as the 802.3ad mode. The servers and the switches they are connected to must support the 802.3ad standard, and load distribution is performed according to this standard.

You can use it to aggregate a group of interfaces of the same speed and duplex setting. Table 3-1 on page 42 summarizes the bonding methods and descriptions of the associated network layers. The 802.3ad standard does not mandate any particular distribution algorithms. However, no distribution algorithm ensures that the following actions do not occur:

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

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

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

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

Notes: T

- ▶ The 802.3ad dynamic link aggregation method is suitable to increase your throughput when you use 1 Gb network interfaces in combination with a single backup server.
- ▶ Combining all 1 Gb network interfaces into one single link aggregation group is not a preferred practice. Use multiple groups of two interfaces or four interfaces instead.

Summary of modes

Table 3-1 shows bonding modes, and options and features for each link aggregation mode implemented on ProtecTIER. The table uses HA (high availability) and RR (round robin) abbreviations.

Table 3-1 Bonding methods and available ProtecTIER link aggregations modes

ProtecTIER GUI and CLI	CLI	Bonding modes and options	Features
HA	HA	miimon=100 ^a mode=1	Fault tolerance
RR	RR	miimon=100 mode=0	<ul style="list-style-type: none"> ► Load balancing ► Fault tolerance ► Unidirectional bandwidth increase
L2	L2	miimon=100 mode=4 xmit_hash_policy=layer2 <ul style="list-style-type: none"> ► Based on MAC address ► 802.3ad switch support needed 	<ul style="list-style-type: none"> ► Load balancing ► Fault tolerance ► Bidirectional bandwidth increase^b
L2L3	L2L3	miimon=100 mode=4 xmit_hash_policy=layer2+3 <ul style="list-style-type: none"> ► Based on MAC address and IP addresses ► 802.3ad switch support needed 	<ul style="list-style-type: none"> ► Load balancing ► Fault tolerance ► Bidirectional bandwidth increase^c
L3L4	L3L4	miimon=100 mode=4 xmit_hash_policy=layer3+4 <ul style="list-style-type: none"> ► Based on IP addresses and ports ► 802.3ad switch support needed 	<ul style="list-style-type: none"> ► Load balancing ► Fault tolerance ► Bidirectional bandwidth increase possible^d

a. Miimon is a parameter to the bonding module. It specifies the MII link monitoring frequency in milliseconds. This determines how often the link state is inspected for link failures. A value of zero disables MII link monitoring. A value of 100 is a good starting point.

b. Outgoing traffic is spread by using a default transmit hash policy of Layer 2. The formula should be read as (source MAC XOR destination MAC) MODULO N (number of subordinates).

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

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

3.4 Preferred ProtecTIER bonding configuration

This section describes the suggested ProtecTIER bonding configuration to use with FSI. When your performance requirement for ProtecTIER FSI is below 500 MBps, you can use the default ProtecTIER bonding configuration of having all interfaces in one single aggregate. If you need more than 500 MBps, configure at least two shares that are exported across two individual IPs.

This configuration enables you to distribute the load across all of the available resources in the ProtecTIER. This configuration is only viable if your environment enables that setup to be efficiently integrated. If you, for example, have two or more backup servers, each of them should use its own FSI file system and file share. With this setup, you have no problems using multiple file shares or IPs on ProtecTIER at the same time.

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

The Layer 3 and Layer 2+3 configurations are suggested only if single bond is not applicable.

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

Table 3-2 clarifies the minimum number of interfaces, subnets, bonds, and IP addresses for different environments and performance goals. Using more of these items might be viable.

Table 3-2 Minimum numbers interfaces, subnets, bonds, and IP addresses for FSI

FSI interface	ProtecTIER performance goal ^a	Minimum # of subnets	Minimum # of bonds and IP addresses ^b
4x 1 Gb	< 500 MBps	1	1
4x 1 Gb	> 500 MBps ^c	2	2
2x 10 Gb	< 500 MBps	1	1
2x 10 Gb	> 500 MBps	2	2

a. The maximum performance that you can reach with your ProtecTIER setup is determined by the ProtecTIER sizing.

b. This column stresses the importance of assigning only one IP address per node (whether it is ProtecTIER or your backup server) to a subnet.

c. Assuming the maximum speed of a single 1 Gb network link is 110 MBps, you need a working 802.3ad setup or multiple subnets to reach these numbers.

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

Layer 2

In today's environments, this mode is the least likely mode to be used for optimization. It can be useful if the system is transmitting a large volume of non-IP traffic on the same VLAN. This mode can be used if IP traffic is being used but in a large single subnet that contains both the sending system and the target systems on the same VLAN. The traffic will be distributed across the interfaces based on the variety of MAC addresses.

Do not use this mode when the destination IP clients are on remote subnets. In this case there will be only the MAC of the sending device and the MAC of the default gateway router which will result in a single interface being used.

Layer 2 + Layer 3

This mode is best to use when IP is used and the destination clients have IP addresses on remote IP subnets. This mode is common in most network environments today. The traffic will be distributed across the interfaces based on the variety of IP addresses.

Layer 3 + Layer 4

This mode is the best to use when the transmitting device is sending traffic to single or small number of destination IP address but using a large number of TCP sessions. The traffic will be distributed across the interfaces based on the variety of TCP port numbers. This mode can also be used in the environment described in L2L3, but it will not necessarily provide additional improvements.

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

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

The default setting for the application interfaces is one application virtual interface that is assigned several physical network ports, depending on the model. This interface is configured in a bond, mode 0 (round robin). You should change the mode to L3L4 (IEEE 802.3ad) when a supporting switch is available.

Switch configuration: The switch must be configured for the L3L4 (IEEE 802.3ad) mode as well. If there is no switch support for IEEE 802.3ad, or the hosts are directly connected to the ProtecTIER server, the default mode should not be changed.

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

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

The main problem with this setup is related to performance if the ProtecTIER server is using the 1 GbE ports (versus using the 10 GbE configurations). Although the ports are configured to share the load on the ProtecTIER server side, the hosts (even if they are part of a bond or team) do not always know to load balance the activity to get the full throughput from the ports.

Load balancing mainly depends on the network cards that are installed on the hosts, and their implementation of teaming. Therefore, you should perform the following activities:

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

3.4.1 VLANs

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

As shown in Figure 3-1, a single switch topology, the ProtecTIER servers, and the hosts are connected to the same VLAN, with the same IP subnet, on the same physical switch.




Figure 3-1 Single switch configuration

When you connect the hosts and the ProtecTIER servers on multiple LAN switches, the connectivity between the switches must be able to transfer the data rate that is required for the backup. For best results, use 10 GbE connectivity between the switches. Another option is to define another link aggregation between the switches so that they can transfer the required bandwidth (Figure 3-2).

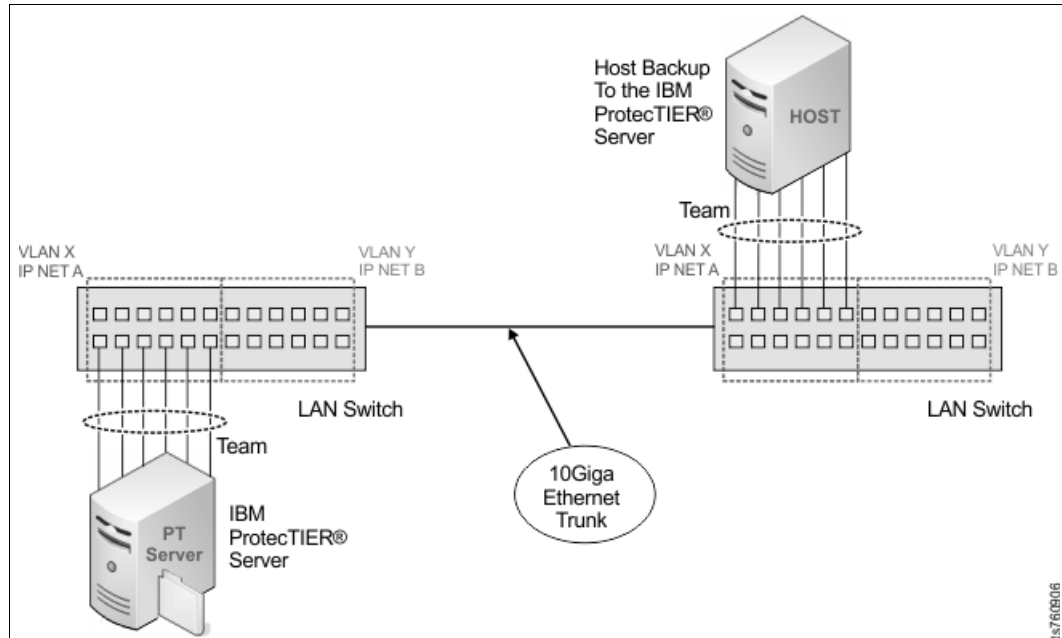


Figure 3-2 Multiple LAN switches configuration

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

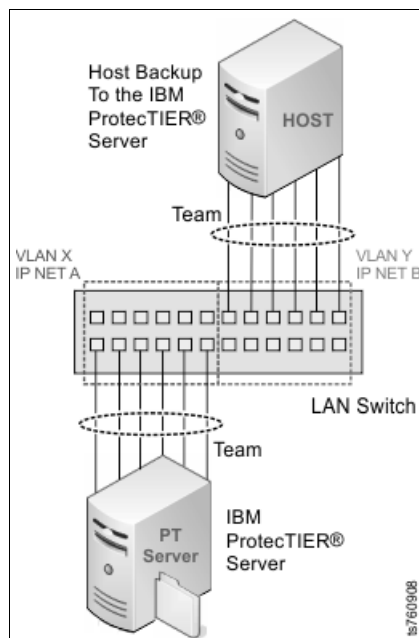


Figure 3-3 Separate VLANs and subnets configuration

3.4.2 IP addresses

You must configure unique IP addresses on the hosts and on the ProtecTIER servers if bonds are configured. If you are configuring bonds, each bond (or team) must be assigned a single IP address. Otherwise, each physical interface must be assigned a unique IP address.

On each system, host or ProtecTIER, each IP address that is configured must be on a different subnet. *Additional hosts and ProtecTIER servers can share the subnet.* For example, on the first ProtecTIER server, you can configure the following IP addresses (and additional sequential addresses):

- ▶ 192.168.151.1/24
- ▶ 192.168.152.1/24
- ▶ 192.168.153.1/24

In this case, the second ProtecTIER node can use the following addresses (and additional sequential addresses):

- ▶ 192.168.151.2/24
- ▶ 192.168.152.2/24
- ▶ 192.168.153.2/24

In this example, the first network is 192.168.151.0, and you can define 255 subnet addresses. Therefore, the first ProtecTIER server is using an address in this subnet (192.168.151.1), and the second ProtecTIER server can use a different address on the same subnet (192.168.151.2).

3.4.3 Routing the IP traffic

Static routes are a simple and effective way of instructing the host IP stack how to route IP traffic that is destined for specific subnets. This configuration is necessary whenever traffic to any specific subnet must be sent through a different gateway and possibly a different network interface than the default gateway definition would otherwise dictate.

If required, configure your static routes so that each port on the host can reach one virtual port on each ProtecTIER server to which it is connected. If possible, configure all IP addresses on the media servers on the same subnets that you defined on the ProtecTIER servers.

For details of how to configure static routes on the ProtecTIER, see the topic about configuring static routes in the *User's Guide for VTL Systems for IBM System Storage TS7600 with ProtecTIER V3.4*, GA32-0922.



Virtual Tape Library guidelines

This chapter describes general preferred practices for optimizing the IBM System Storage TS7600 ProtecTIER family Virtual Tape Library (VTL). It describes VTL concepts, methods, and system components.

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

This chapter describes the following topics:

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

4.1 ProtecTIER Virtual Tape Library introduction

The ProtecTIER VTL service emulates traditional tape libraries. Currently four types of libraries are supported: The TS3500 is the most common one; the VTL service can also emulate P3000, VTF, and V-TS3500 libraries. By emulating tape libraries, you can use ProtecTIER VTL to switch to disk backup without having to replace your entire backup environment. Your existing backup application can access virtual robots to move virtual cartridges between virtual slots and drives. The backup application perceives that the data is being stored on cartridges, although the ProtecTIER product stores data on a deduplicated disk repository.

4.2 General preferred practices for the virtual tape library

A ProtecTIER VTL can be optimized by the following these simple rules:

- ▶ Create more slots than are needed for future growth in the number or cartridges. Adding cartridges when there are free slots available is an online procedure, but adding more slots to the library is an offline procedure.
- ▶ Create cartridges with a fixed maximum cartridge size. You can accomplish this task by selecting **Max Cartridge Growth** in the tape cartridges creation menu.

Important: All virtual cartridges get an equal amount of space up to either the limit size that is configured when the virtual cartridge was created (maximum cartridge size), or the calculated amount of Nominal Space/Number of Cartridges, whatever comes first.

One possibility is to configure a virtual tape library to use cartridges with a fixed size. When this feature is enabled, the dynamic calculation of the virtual cartridges size is overridden. This is used sometimes when the virtual cartridges are migrated to physical cartridges using a backup application. This is a controlled procedure *because it is not a preferred configuration*. It must be requested to Level3 support through a PMR.

- ▶ Create small cartridge sizes to support as much parallelism as possible. The VTL emulates a physical tape library, so virtual cartridges behave in the same manner as physical cartridges with sequential access.

If you have many drives and large tape sizes, you might encounter a situation where the backup or restore is waiting for a large sized tape that is being used by another backup session. If you have small tapes with a capacity of 100 GB, for example, you decrease the probability of wait times for backup and restore operations.
- ▶ If you have a physical tape library that is connected to the same backup application that is using ProtecTIER, make sure that the libraries use different bar code ranges. This action avoids tape access conflicts and facilitates the identification of which tape is physical and which tape is virtual.
- ▶ Create only the number of cartridges that your repository can handle, maybe even fewer to control the repository allocation of different VTLs. You can estimate the nominal size of a repository by multiplying the physical size of the repository by the expected HyperFactor ratio. Then, divide it by the tape size you want and determine the optimized number of tapes.

4.3 Setting up the virtual library and cartridges

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

4.3.1 Creating libraries

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

1. Log on to the system on which you want to add a library.
2. From the menu bar, click **VT** → **VT Library** → **Create new library**. The Create new library wizard Welcome window opens (Figure 4-1).

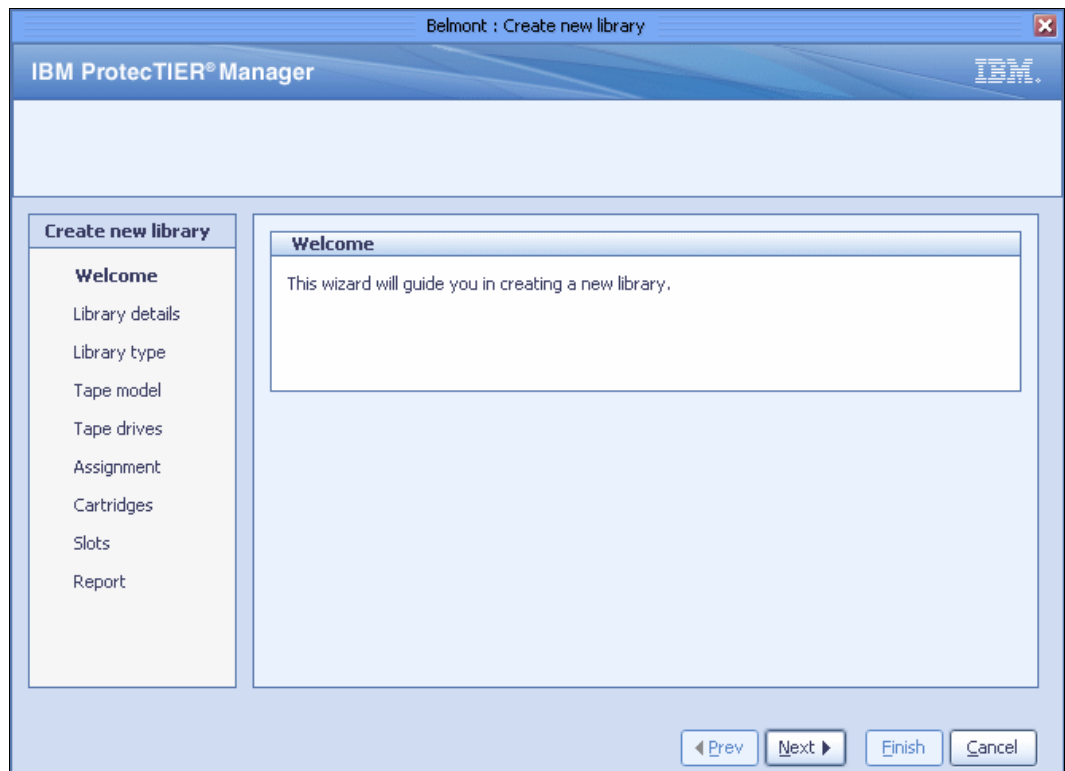


Figure 4-1 Create new library: Welcome window

3. Click **Next**. The Library details window opens (Figure 4-2).

vtlv7k : Create new library

IBM ProtecTIER® Manager

Create new library

- ✓ Welcome
- Library details**
- Library type
- Tape model
- Tape drives
- Assignment
- Cartridges
- Slots
- Report

Library name

Choose a name for this library.

VT name:

◀ Prev Next ▶ Finish Cancel

Figure 4-2 Create new library: Library details window

4. In the ProtecTIER VT name field, enter a name for the library, and click **Next**.

Note: A good practice is to use a unique naming convention for the VTL library, for example *VTL01*, to help differentiate VTL libraries from other Tape libraries.

5. The Library Type window opens (Figure 4-3).

By default, IBM TS3500 is selected.

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

- ▶ Old versions of Symantec NetBackup software require you to select **V-TS3500**. The latest versions of this backup application work fine with the library type **TS3500**.
- ▶ If you are using IBM Spectrum Protect which delivers the function of Tivoli Storage Manager, or any other backup application, you should select **TS3500**.

Click **Next**.

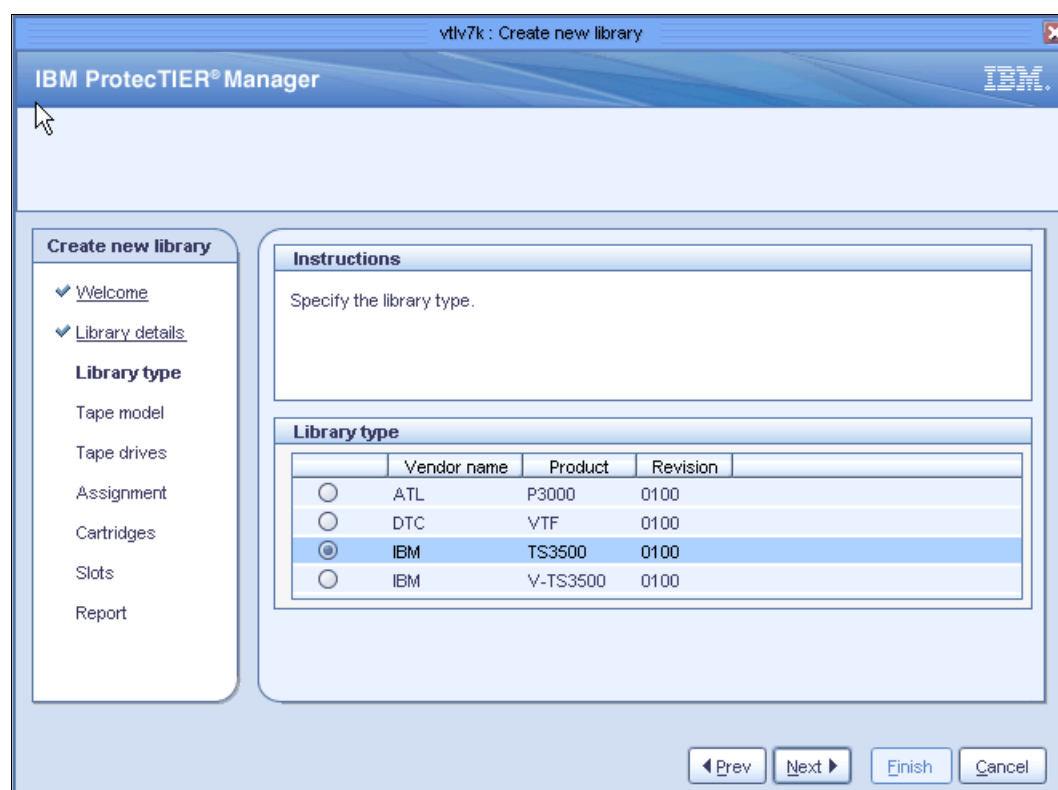
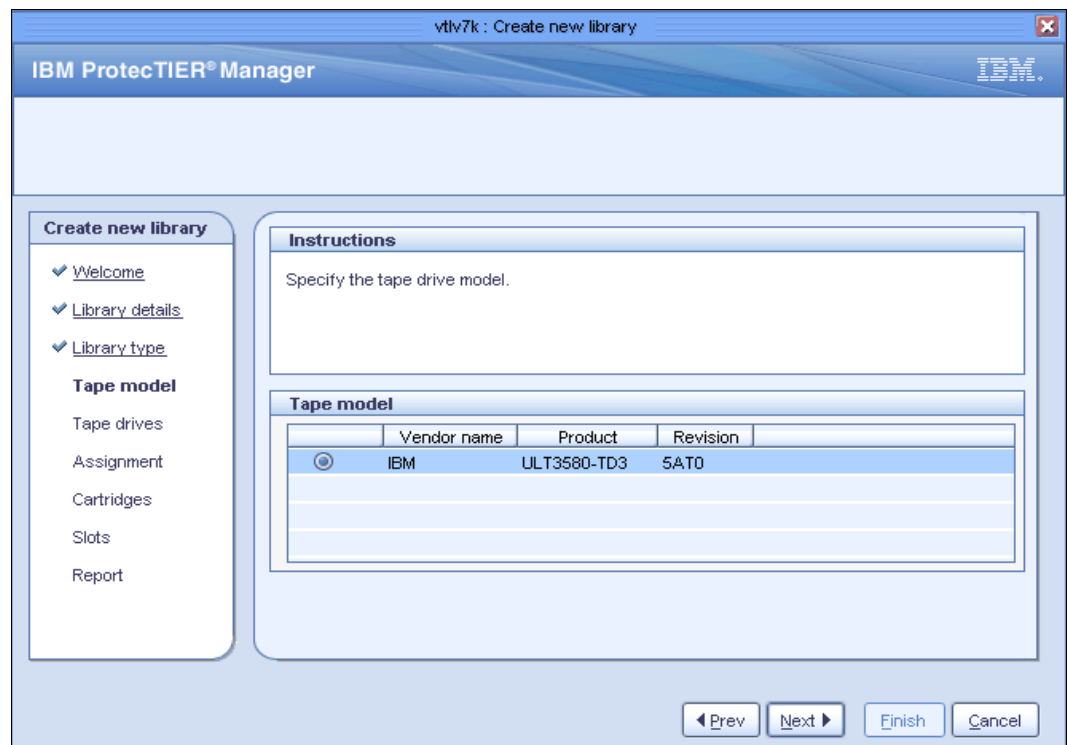


Figure 4-3 Create new library: Library type window

6. The Tape model window opens (Figure 4-4). Select the tape drive model that you want to use for your virtual library. The type of tape drive model that is displayed depends upon the type of library that was selected in the previous step. Click **Next**.



The screenshot shows the 'Create new library' window in IBM ProtecTIER Manager. The window title is 'vltv7k : Create new library'. The main header is 'IBM ProtecTIER® Manager' with the IBM logo on the right. On the left is a sidebar with a tree view under 'Create new library' containing: 'Welcome', 'Library details', 'Library type', 'Tape model' (selected), 'Tape drives', 'Assignment', 'Cartridges', 'Slots', and 'Report'. The main area has an 'Instructions' box with the text 'Specify the tape drive model.' Below it is a 'Tape model' section containing a table with columns: 'Vendor name', 'Product', and 'Revision'. The first row is selected and contains 'IBM', 'ULT3580-TD3', and '5AT0'. At the bottom right are four buttons: 'Prev', 'Next', 'Finish', and 'Cancel'.

	Vendor name	Product	Revision
<input checked="" type="radio"/>	IBM	ULT3580-TD3	5AT0
<input type="radio"/>			
<input type="radio"/>			

Figure 4-4 Create new library: Tape model window

7. The Tape drives window opens (Figure 4-5). In the Number of tape drives field, enter the number of tape drives to assign to the library.

vtlv7k : Create new library

IBM ProtecTIER® Manager

Create new library

- ✓ Welcome
- ✓ Library details
- ✓ Library type
- ✓ Tape model
- Tape drives**
- Assignment
- Cartridges
- Slots
- Report

Instructions

Specify the number of tape drives for all nodes in system vtlv7k. Maximum number of tape drives per node is limited by the license and memory.

Tape drives

Maximum number of tape drives that can be created on new library is 256.

Number of tape drives to create: 64

Number of tape drives on lbsdudup1a (Max 256)

◀ Prev Next ▶ Finish Cancel

Figure 4-5 Create new library: Tape drives window

The maximum number of virtual tape drives depends on the model of ProtecTIER. A ProtecTIER TS7620 Appliance Express supports a maximum of 128 virtual tape drives, a ProtecTIER TS7650 Gateway node supports a maximum of 256 virtual tape drives, and on a dual node cluster it supports a maximum of 512 virtual tape drives.

Figure 4-5 shows 64 drives created on the *lbsdudup1a* node.

Number of drives: Check with your backup application administrator to verify the number of drives and cartridges that is supported by your application.

8. Click **Next**. The Port assignment window opens, as highlighted in Figure 4-6.

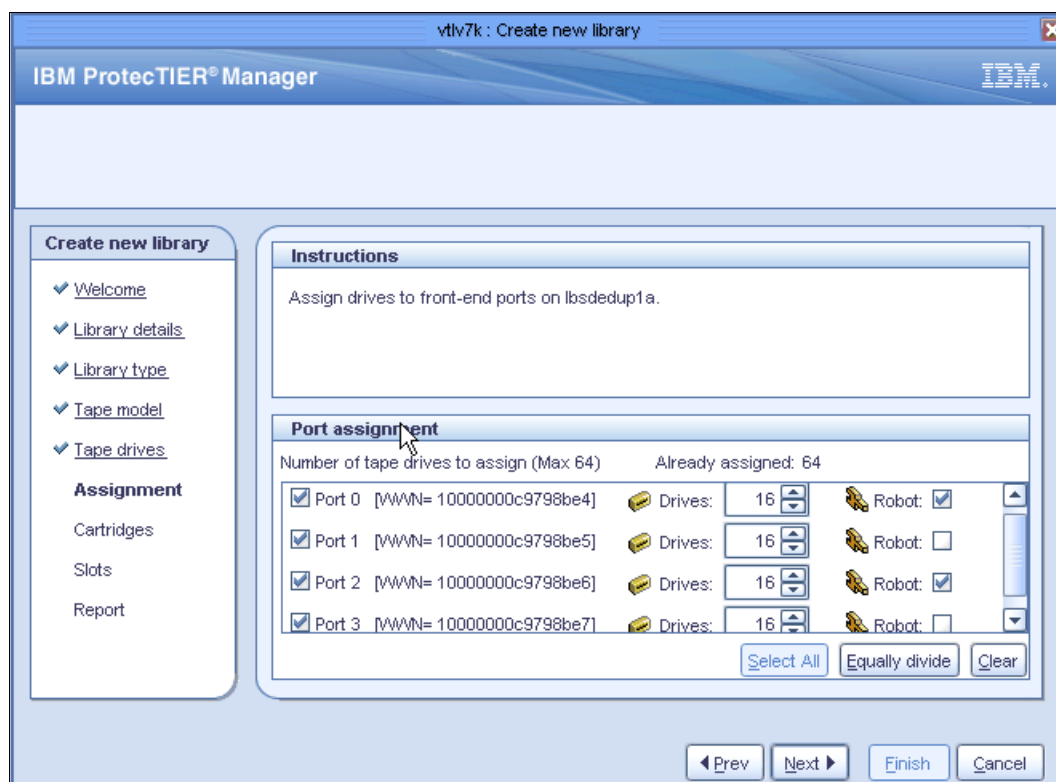


Figure 4-6 Create new library: Port assignment window

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

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

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

10. Click **Next**. If a second node exists in your cluster, the Assignment (2) window opens. Complete the same steps as you did in Figure 4-6 for your first cluster.

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

On a dual node cluster, the robot should be assigned on both nodes so that if one node is offline, the backup application still has a connection to the surviving node and a robot device on it. Remember that if the robot is not accessible, the backup application cannot move any virtual tapes in or out of the virtual drive.

11. Click **Next**. The Cartridges window opens (Figure 4-7). In the **No. of cartridges** field, enter the number of cartridges that you want to have in the library.

By default, ProtecTIER calculates a virtual cartridge size, but it does not use only the current nominal allocable space. Instead, ProtecTIER makes an internal projection for the allocable space that takes into account the current factoring ratio, the current percentage of used disk space, the expected factoring ratio, and the current percentage of free disk space.

Note: Adjusted Nominal Allocable Space = available physical space * ((current factoring ratio * current percentage of used disk space) + (expected factoring ratio * current percentage of free disk space))

The Virtual cartridge size shown on the Cartridges window (Figure 4-7) is calculated by dividing the Adjusted Nominal Allocable Space by the total number of cartridges that will exist on the repository. That is, all existent cartridges plus the ones that will be created for this library.

12. Limit the maximum size of the cartridges, especially if these will be migrated to physical cartridges later. Complete this task by selecting the **Max Cartridge Growth** check box and entering a cartridge size in gigabytes.

Although there is no “magic” number for cartridge size, the suggested guideline for cartridge maximum size is 100, 200, or even 400 GB, depending on the workload.

Figure 4-7 Create new library: Cartridges window

13. In the **Barcode seed** field (Figure 4-7), enter a value for the bar code seed. The bar code seed is the bar code that is assigned to the first cartridge created. Every cartridge that is added after the first cartridge is incrementally assigned a bar code that follows the previous one.

Tips:

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

14. Click **Next**. The Slots window opens (Figure 4-8).

The screenshot shows the 'Slots' window in the IBM ProtecTIER Manager. The window title is 'vtiv7k : Create new library'. The left sidebar shows a list of steps: Welcome, Library details, Library type, Tape model, Tape drives, Assignment, Cartridges, Slots (selected), and Report. The main area has an 'Instructions' section with a notice about replication targets and a 'Slots' section with two input fields: 'Number of slots (Max 63488)' set to 2,000 and 'Number of import/export slots (Max 1022)' set to 16. At the bottom are 'Prev', 'Next', 'Finish', and 'Cancel' buttons.

Figure 4-8 Creating new library: Slots window

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

Important:

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

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

Tips:

- ▶ The number of import/export slots will affect how many cartridges can be moved in and out of the library at a time in many operations, for example to move cartridges between the library and the shelf.
- ▶ The number of import/export slots should be considered in ratio to the total number of cartridges that are planned to exist in the repository. For example, for a total of 100 cartridges, 10 import/export slots is fine. However, for a total of 25,000 cartridges, 10 import/export slots is too low, because it would take 2500 iterations of a move to move all the tapes.
- ▶ The suggested guideline for the number of import/export slots is 500 - 1000, unless there is a limitation that overrules this number, such as a limitation imposed by the backup application.

17. Click **Next**. The Create new library wizard closes and a summary report opens. Click **Finish**.

Figure 4-9 and Figure 4-10 on page 60 show two parts to an example of the summary report output.

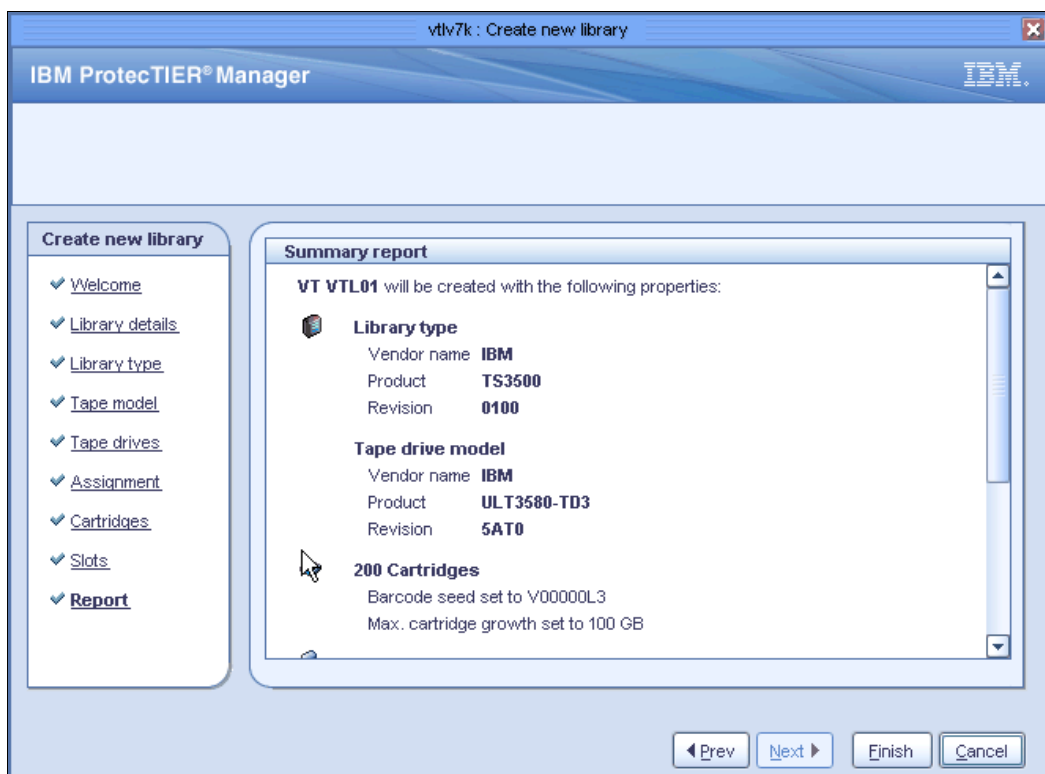


Figure 4-9 Summary report (part 1 of 2)

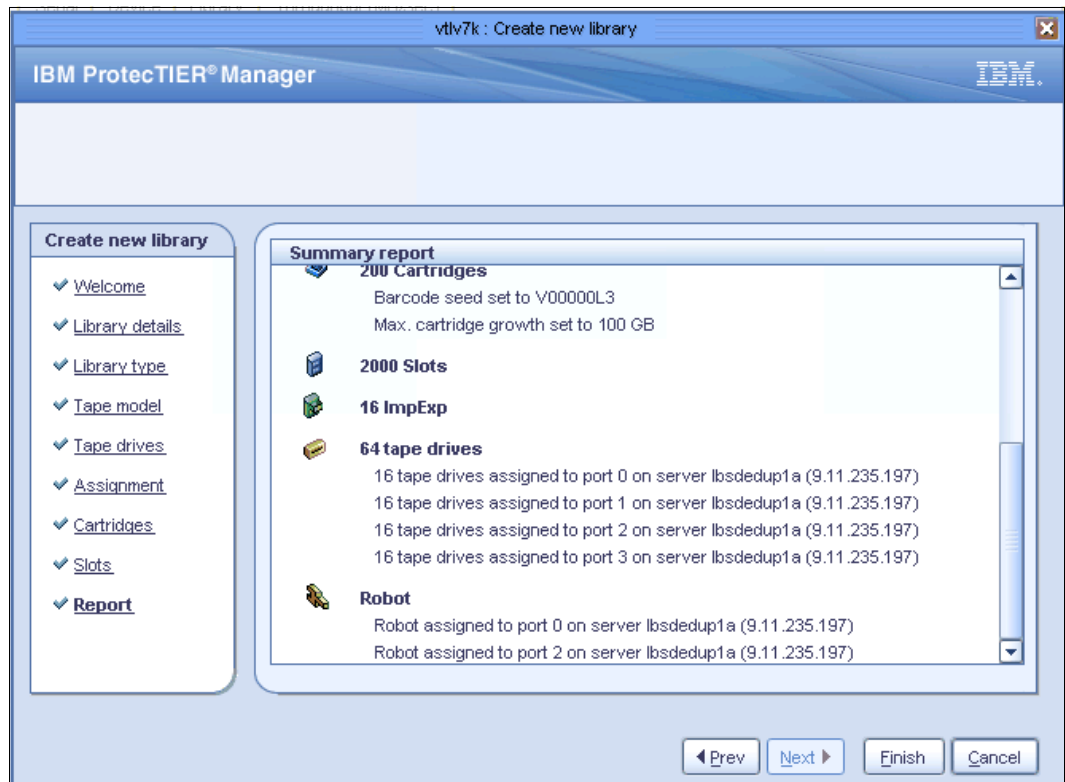


Figure 4-10 Summary report (part 2 of 2)

18. The Confirm operation window opens (Figure 4-11). Click **Yes**. The ProtecTIER system temporarily goes offline to create the library.

Important: All libraries go offline when the system creates a new library.

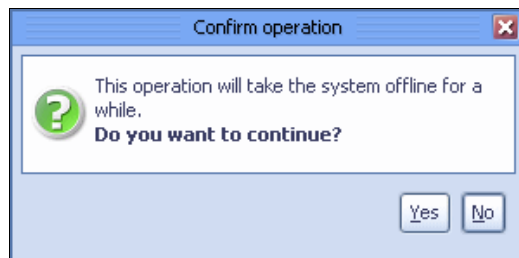


Figure 4-11 Create new library: Confirm operation window

The newly created library is displayed in the left navigation pane of the ProtecTIER Manager graphical user interface (GUI). The right context pane shows the details of the virtual library that you just created (Figure 4-12).

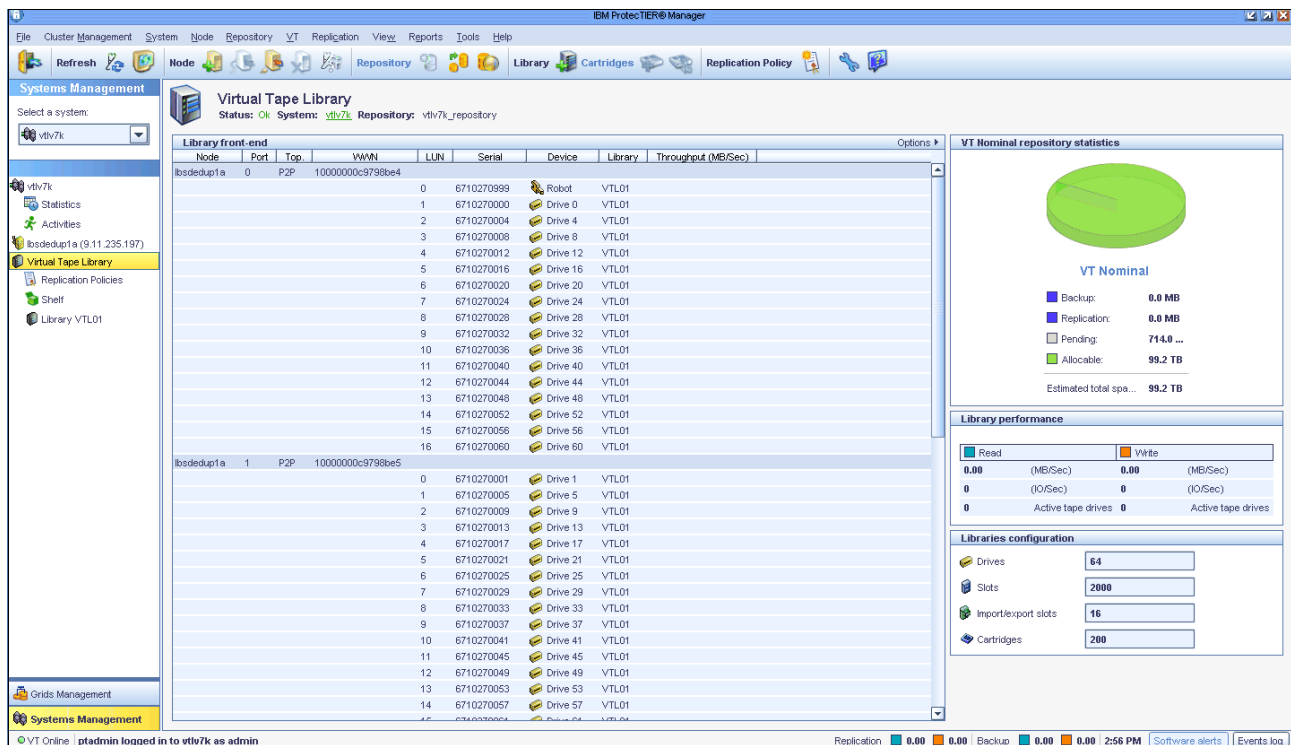
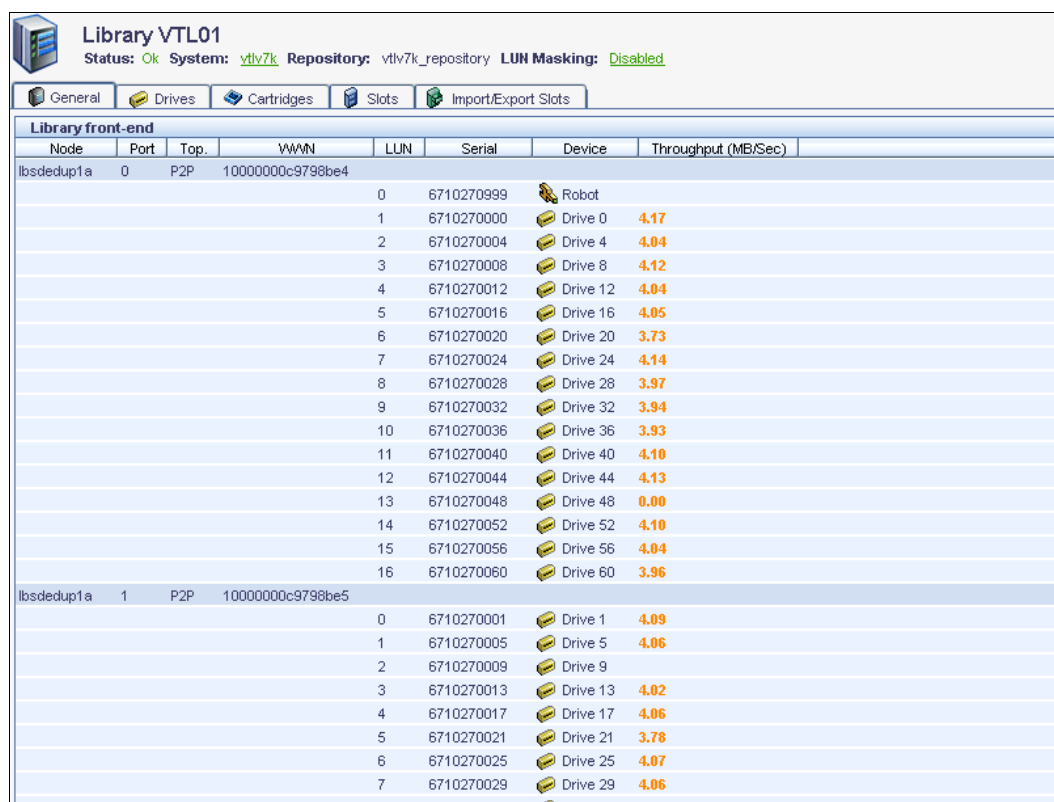


Figure 4-12 Newly created Virtual Tape Library

19. Select the new library (**Library VTL01**) in the Navigation to show the library details. The General tab of the logical library is shown in Figure 4-13.



The screenshot shows the 'Library VTL01' window with the 'General' tab selected. The window title bar includes 'Library VTL01', 'Status: Ok', 'System: vtlv7k', 'Repository: vtlv7k_repository', and 'LUN Masking: Disabled'. Below the title bar are tabs for 'General', 'Drives', 'Cartridges', 'Slots', and 'Import/Export Slots'. The main content area is titled 'Library front-end' and contains a table with columns: Node, Port, Top., WWN, LUN, Serial, Device, and Throughput (MB/Sec).

Node	Port	Top.	WWN	LUN	Serial	Device	Throughput (MB/Sec)
lbsdcdup1a	0	P2P	10000000c9798be4	0	6710270999	Robot	
				1	6710270000	Drive 0	4.17
				2	6710270004	Drive 4	4.04
				3	6710270008	Drive 8	4.12
				4	6710270012	Drive 12	4.04
				5	6710270016	Drive 16	4.05
				6	6710270020	Drive 20	3.73
				7	6710270024	Drive 24	4.14
				8	6710270028	Drive 28	3.97
				9	6710270032	Drive 32	3.94
				10	6710270036	Drive 36	3.93
				11	6710270040	Drive 40	4.10
				12	6710270044	Drive 44	4.13
				13	6710270048	Drive 48	0.00
				14	6710270052	Drive 52	4.10
				15	6710270056	Drive 56	4.04
				16	6710270060	Drive 60	3.96
lbsdcdup1a	1	P2P	10000000c9798be5	0	6710270001	Drive 1	4.09
				1	6710270005	Drive 5	4.06
				2	6710270009	Drive 9	
				3	6710270013	Drive 13	4.02
				4	6710270017	Drive 17	4.06
				5	6710270021	Drive 21	3.78
				6	6710270025	Drive 25	4.07
				7	6710270029	Drive 29	4.06
				8	6710270033	Drive 33	4.06

Figure 4-13 Logical library details

20. In the Logical library details window (Figure 4-13), you can click the tabs for details about each component, and also take the corresponding actions for the component. The components are as follows:

- **Drives.** Shows mounted tapes and the throughput of each logical tape drive. From this tab, you can, for example, reset a drive.
- **Cartridges.** Shows the cartridge inventory, including the tape size, its usage, and whether it is full. From the Cartridge tab, you can also see the cartridge replication synchronization status. For more information, see 22.3.7, “Gauging the replication completion status” on page 400.
- **Slots.** Represents the logical slots where the virtual tapes are stored. For example, if a virtual tape is mounted, the slot where the tape was located is empty while the tape is mounted, the same as a real library slot.

Slots tab: From the Slots tab, you can move tapes to a drive, to another slot, to the shelf, or to import/export slots. Although performing a manual activity is possible, you should move tapes by using the backup application to avoid mismatched library inventory in the backup application and the virtual tape library.

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

21.To the right of the General tab is a representation of the drives, online throughput statistics, and the dimensions of the selected VTL (Figure 4-14).

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

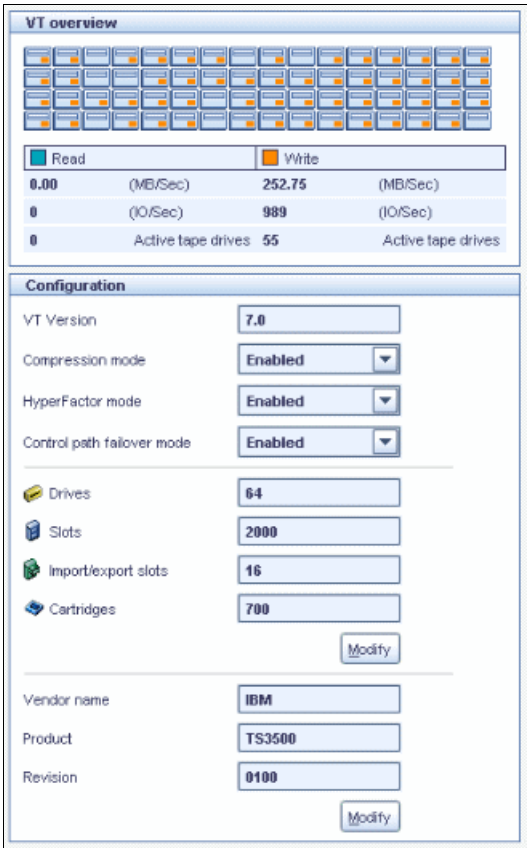


Figure 4-14 Logical tape library details and dimensions



ProtecTIER File System Interface: General introduction

This chapter describes network configuration considerations for setting up your network with the IBM System Storage TS7600 ProtecTIER family File System Interface (FSI). For the general configuration process, review the relevant user's guide for IBM TS7600 with ProtecTIER, for example *User's Guide for FSI Systems*, GA32-2235. This chapter describes how to create Network File System (NFS) exports and Common Internet File System (CIFS) shares on ProtecTIER FSI, and how to connect them to your backup host.

This chapter describes the following topics:

- ▶ ProtecTIER FSI network overview
- ▶ File System Interface guidelines for NFS
- ▶ File System Interface guidelines for CIFS
- ▶ FSI file system scalability

Important: ProtecTIER GA Version 3.4 was released with only the Virtual Tape Library (VTL) interface support. File System Interface (FSI) support was added to ProtecTIER PGA 3.4 Version. For details, see the announcement letter:

http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/6/872/ENUSA P16-0076/index.html&lang=en&request_locale=en

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

ProtecTIER FSI supports various backup applications, such as IBM Spectrum Protect, Symantec NetBackup, and EMC NetWorker. A list of supported backup applications is in Appendix B, "ProtecTIER compatibility" on page 457.

Starting with ProtecTIER Version 3.2, support is provided for Windows based servers through the CIFS protocol, and starting with ProtecTIER Version 3.3, support is provided for UNIX

clients through the NFS protocol. ProtecTIER emulates a UNIX or Windows file system behavior, and presents a virtualized hierarchy of file systems, directories, and files to UNIX NFS or Windows CIFS clients. These clients can perform file system operations on the emulated file system content.

You can use the ProtecTIER FSI to create multiple user file systems in a single ProtecTIER repository. When you create a user file system, the maximum size of the user file system is dynamically calculated by determining the total free nominal space in the repository and comparing it to the overall maximum user file system size of 256 TB.

The size of all file systems shrinks proportionally if the deduplication ratio goes lower than expected. If the deduplication ratio goes beyond the expected size, extending the file system size up to the 256 TB limit is possible by using the ProtecTIER Manager.

The FSI interface of ProtecTIER for UNIX and Windows clients is supported on a single node. Dual node cluster support is currently not available. However, a single node can serve multiple CIFS and NFS exports in the same repository.

On the current release, exporting a single FSI share through CIFS or NFS protocol is mutually exclusive. To change the export type from NFS to CIFS, you must delete the NFS share definition before you export it through CIFS, and vice versa. Disabling the share definition alone is not sufficient.

Important: ProtecTIER FSI support is intended for storing backup images that are produced by backup applications, and not for primary storage deduplication. ProtecTIER performs best when sequential streams are delivered to ProtecTIER rather than random input/output (I/O).

5.1 ProtecTIER FSI network overview

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

5.1.1 ProtecTIER network

ProtecTIER servers have several physical network ports. The number of ports varies based on the ProtecTIER model. Ports are used for management, replication, or file system-related operations from the hosts. Each port is assigned to one of these uses.

This configuration is achieved by assigning the physical ports to a virtual interface on the ProtecTIER server. The set of virtual interfaces on the ProtecTIER product includes external, replic1, replic2, fsi1, fsi2, on to fsi_n, as shown in Figure 5-4 on page 69. Each one of the virtual interfaces can have assigned one or more physical network ports.

The default setup of the ProtecTIER product assigns all of the FSI physical ports to a single virtual interface by using round robin load balancing. This setup can be changed as needed.

If more than one physical port is assigned to a virtual interface, be sure to configure the bonding methodology in this interface to align with the network environment to fulfill the wanted behavior in terms of performance and redundancy. For more information about the bonding methods that are available with the ProtecTIER product, see Chapter 3, “Networking essentials” on page 37.

5.1.2 Network configuration considerations

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

- ▶ Make sure that the backup application runs in the context of a user that has read and write permissions on the FSI share/export.
- ▶ You must have *at least* two different network subnets to separate the ProtecTIER management IP interface from the ProtecTIER file system (application) interfaces.
- ▶ For FSI workloads, you must have sufficient Transmission Control Protocol/Internet Protocol (TCP/IP) infrastructure for the incoming and outgoing traffic of backup servers. Ensure that you do not suffer from network bandwidth congestion.
- ▶ If bonding of network adapters is implemented, it must be implemented on all involved network devices, that is, the ProtecTIER server, the backup server, and the network switches. Enabling bonding only on the ProtecTIER server might not be enough to achieve the best results.

5.1.3 Connecting a ProtecTIER server to the network

To better understand the requirements that arise from using the FSI model of ProtecTIER in your network environment, look at the potential connections that you will deal with during the initial deployment. The diagrams displayed in this section reference the newest TS7650G DD6 model.

As shown in Figure 5-1, this example uses the connection (labeled **1**) to attach the ProtecTIER server to the customer network. Through this connection, you use the ProtecTIER Manager graphical user interface (GUI) and connect to the ProtecTIER server for management and configuration purposes.

The ProtecTIER IP replication feature shows two Ethernet connections (labeled **21** and **22**). By default, the replication workload is balanced across both ports.



Figure 5-1 ProtecTIER network interfaces for customer and replication network on Gateway model

To use the FSI feature on a ProtecTIER Gateway, you must prepare at least one (and as many as four when using 1 GbE cards) dedicated subnets for the backup server traffic to the ProtecTIER server. The data that is transferred from the backup server to the FSI interface must not use the customer network IP interface or the replication interfaces. For details about the cabling for other ProtecTIER models, review the chapter about hardware planning for the IBM TS7600 ProtecTIER system in *IBM System Storage TS7600 with ProtecTIER Version 3.3*, SG24-7968.

Figure 5-2 shows the 1 GbE interface model of the ProtecTIER server. The interfaces that are labeled **13**, **14**, **15**, and **16** are available to use the FSI traffic between the backup servers and the ProtecTIER server.




Figure 5-2 ProtecTIER network interfaces for FSI traffic from backup servers on a Gateway model for FSI 1 GbE

Now that you see all of the important interfaces for potential network traffic, you can review the configuration through the ProtecTIER Manager GUI.

To configure all networking-related aspects of the ProtecTIER server open the ProtecTIER Manager GUI and click **Node** → **Network configuration** (Figure 5-3).

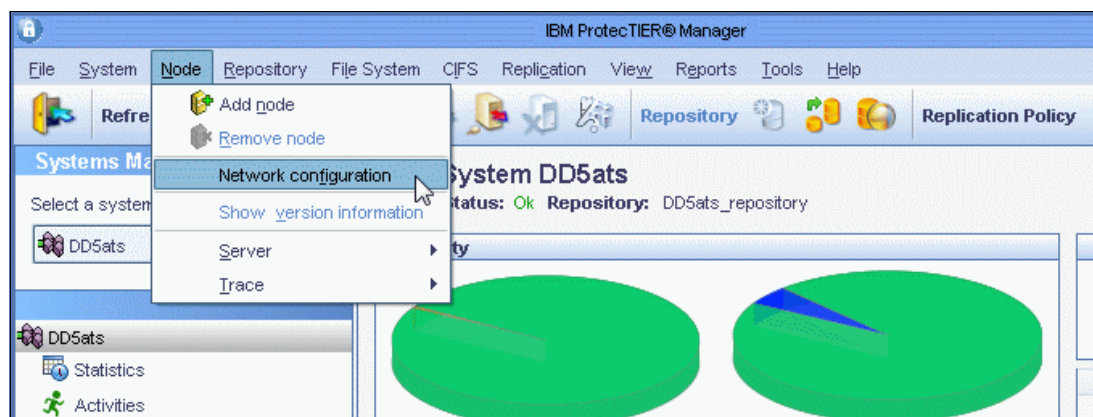








Figure 5-3 ProtecTIER Manager GUI network configuration window

The Network configuration window (Figure 5-4 on page 69) provides options to change the networking parameters.

DD5ats : Network configuration

Customize the network configuration by configuring the network addresses, load balancing methods, and assigning Ethernet ports to the devices.

 An external device has been modified. Applying these changes may disconnect communication with the server.

Network Configuration Table						
Name	L/B Method	IP Address	Netmask	Gateway	Assigned Ethernet	
external	High Availability	9.123.123.123	255.255.255.0	9.123.123.1	Assign ...	 eth0
replic1	High Availability	192.168.100.6	255.255.255.0		Assign ...	 eth2
replic2	High Availability	192.168.200.6	255.255.255.0		Assign ...	 eth9
fsi1	Round Robin	10.1.114.197	255.255.255.0		Assign ...	 eth4
fsi2	Round Robin	10.2.114.197	255.255.255.0		Assign ...	 eth6
fsi3	L3L4	192.168.152.1	255.255.255.0		Assign ...	
fsi4	L2L3	192.168.153.1	255.255.255.0		Assign ...	

Unassigned Ethernet Ports



Unassign ...  eth5  eth7

Figure 5-4 ProtecTIER network configuration window

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

Tip: Optionally, you can also configure your network from the ProtecTIER Service menu directly in the ProtecTIER server.

Multiple ways exist to set up your networking to ensure that you have an HA configuration, and that you distribute the load across all available resources. The default setup is a single virtual interface, *fsi1*, which consists of all four physical 1 Gb ports (Table 5-1).

Table 5-1 DD5 1 GbE Ethernet default port assignments - Gateway FSI 1 GbE

Network types	Virtual interfaces			Assigned physical ports			
	Network IP	LB	Subnet	Name	Speed	Slot	Port
External	External IP	RR	1	Eth2	1 GbE		Onboard
Application	fsi1	RR	2	Eth3	1 GbE	1	13
				Eth4	1 GbE	1	14
				Eth5	1 GbE	1	15
				Eth6	1 GbE	1	16
Replication	replic1	N/A	3	Eth0	1 GbE		Onboard
	replic2	N/A	4	Eth1	1 GbE		Onboard

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

If you use the configuration that is shown in Table 5-1 on page 69, all of your backup servers connect to the IP of the ProtecTIER application virtual interface fsi1. The default load-balancing (LB) method of round robin (RR) mode 1 works without special network infrastructure hardware requirements. This LB mode permits, depending on your network infrastructure, a unidirectional bandwidth increase.

This configuration means that, from the perspective of a single data stream that flows outbound from a ProtecTIER server, you can potentially benefit from up to 4 Gb of bandwidth, which is essentially the combined throughput of all four aggregated interfaces. It also means that restoring data from your ProtecTIER server to your backup server could be fast. For further details about port aggregation, see Chapter 3, “Networking essentials” on page 37.

Backing up your data creates a data stream that is directed toward the ProtecTIER server. Single data streams directed toward a ProtecTIER server do not benefit from the potential bandwidth increase when you use the round-robin LB method in this example.

To fully use the ProtecTIER server resources in this configuration, you must use multiple backup servers that back up to their respective file systems on the ProtecTIER server. To further optimize the potential throughput of single backup server environments, you must understand the link aggregation methods that can be used for load balancing and increasing throughput, as listed in Table 3-1 on page 42.

5.1.4 Replication

You can use ProtecTIER to define replication policies to replicate a file system's directories and all the objects that are contained in these directories recursively to remote ProtecTIER repositories without any disruption to the operation of the file system as a target for backup. It is possible to define up to 64 source directories per one replication policy, and to define up to three remote ProtecTIER destinations.

The replicated data in the remote destination can be easily used to restore data in the case of a Disaster Recovery (DR), or in the case of a DR test (without any interruption to the backup and replication procedures).

An important task is to enable the ProtecTIER system to supervise all the changes that are made to a directory, or to a set of directories, that is constantly defined in a replication policy. Therefore, you should not disable a replication policy unless this policy is no longer considered relevant.

If maintenance is scheduled for the network that is used for replication, a possibility (although not mandatory) is to suspend the replication to a specific destination. Suspending replication enables the ProtecTIER system to continue supervising all of the changes, but it does not attempt to send the replication data through the network for the time that is defined by the suspend operation. The suspend operation is limited in time, with a maximum suspend time of 72 hours.

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

5.1.5 Disaster recovery: Test

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

5.1.6 Disaster recovery: Event

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

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

For more information about this procedure, see the chapter about native replication and disaster recovery in *IBM System Storage TS7600 with ProtecTIER Version 3.3*, SG24-7968.

5.1.7 General FSI suggestions

Disable any encryption features in the backup server when you use ProtecTIER as the backup target, as shown in Table 5-2.

Table 5-2 Suggested settings

Parameter	Value in backup application
Compression	Disable
Deduplication	Disable
Encryption	Disable
Multiplexing	Disable

5.2 File System Interface guidelines for NFS

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

Important: ProtecTIER GA Version 3.4 was released with only the Virtual Tape Library (VTL) interface support. File System Interface (FSI) support was added to ProtecTIER PGA 3.4 Version. For details, see the announcement letter:

http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/6/872/ENUSA P16-0076/index.html&lang=en&request_locale=en

5.2.1 ProtecTIER NFS authentication and security management

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

- ▶ Port security
- ▶ Root squash/no root squash

Port security

To configure port security, go to the Properties tab of the NFS export wizard. The Port Security option is under the Details section (Figure 5-5 on page 73).

In the Details section, select whether you want to enable NFS clients to connect to ports higher than 1023. The port numbers 0 - 1023 are the well-known ports, also known as *system ports*. These TCP/IP port numbers permit only root users and services to run servers on these ports. Port 1024 and higher are also known as *user ports*.

Keep the default setting and leave the check box selected, as shown in Figure 5-5.

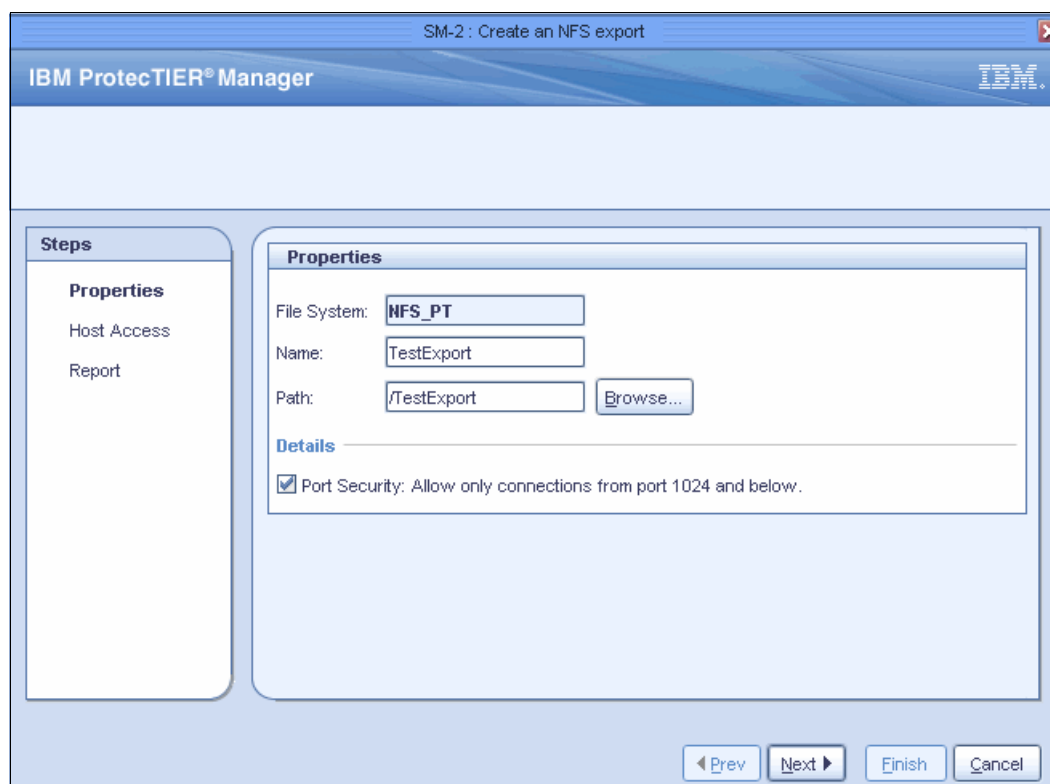


Figure 5-5 Create an NFS export wizard: Port Security definition

Root squash

Enabling this option prevents root users, from the NFS client systems, from having root privileges on the NFS export that are provided by ProtecTIER. If you do not enable root squash, any root user of a remote system might delete any user data on the mounted NFS export because root can delete any data of foreign users.

To prevent this action, the root squash option maps the root user ID 0 (UID) and group ID 0 (GID) to a customized UID. By doing this task, the remote root user cannot delete or modify any other data than the one that is created with the customized UID. Typically, the root squash function by default maps the UID to nfsnobody, but in the ProtecTIER implementation, the UID of that user is higher than the value that you are allowed to enter in the wizard's field.

Alternatively, the no_root_squash option turns off root squashing.

To select the User ID Mapping for root squash or no root squash, use the Add Host Access window shown in the Create an NFS export wizard (Figure 5-6).

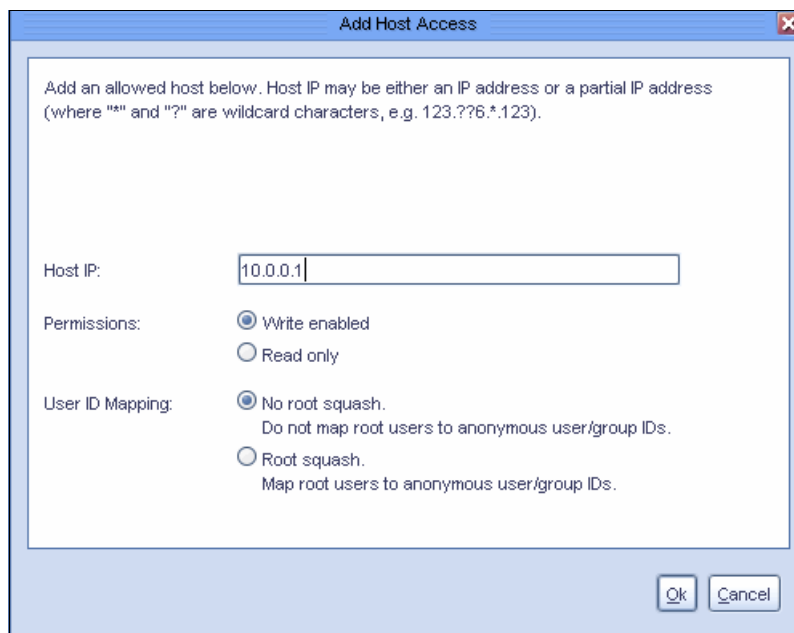


Figure 5-6 Create NFS Export wizard - User ID Mapping definition

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

Example 5-1 shows how to use the **grep** command to determine the UID and the GID of the user *nobody*. This user exists in the ProtecTIER system. You must log on to the ProtecTIER CLI using SSH to query the user account information.

Example 5-1 Determine the user ID and group ID of user nobody on the ProtecTIER server

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

The output of the commands in Example 5-1 shows that the numeric value for the user and group are both 99. You can use this number to configure the root user ID mapping, or to create a customized user account and a dedicated group to map one or more root accounts of the remote NFS clients.

If you decide not to use the existing *nobody* account, you can create your own customized group and several users, as shown in Example 5-2.

Example 5-2 Create a customized group and user

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

5.2.2 Understanding root squash

The basics of NFS root squash and no root squash are explained in “Root squash” on page 73. The following section demonstrates the effects of turning root squash on or off.

Example 5-3 shows a directory listing of a ProtecTIER NFS export. The `file1` file was created by a root user. Usually, the user ID of root is 0, but because root squash was turned on when the NFS export was defined, the root user ID is mapped to a defined UID (in this example, they are user ID 65536 and group ID 65536). The `file2` file was created by the `tsminst1`, which belongs to the `tsmsrvrs` group.

Example 5-3 Directory listing on an NFS share

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

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

Example 5-4 demonstrates that the root user ID is not allowed to delete `file2`, which belongs to `tsminst1`. The `delete` command fails with an error message `Permission denied`.

Example 5-4 Deleting files with root squash enabled in the NFS export definition

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

To demonstrate the power of the root user without the root squash function enabled, we modified the NFS export definition and disabled root squash. In comparison to Example 5-4 the root user can delete `file2` even if the file is owned by `tsminst1`. The result of the delete operation is shown in Example 5-5. The `file2` was deleted without any error.

Example 5-5 Deleting files with root squash disabled in the NFS export definition

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

5.3 File System Interface guidelines for CIFS

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

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

Important: ProtecTIER GA Version 3.4 was released with only the Virtual Tape Library (VTL) interface support. File System Interface (FSI) support was added to ProtecTIER PGA 3.4 Version. For details, see the announcement letter:

http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/6/872/ENUSA P16-0076/index.html&lang=en&request_locale=en

5.3.1 Mounting the NFS export in a UNIX system

When working with FSI, properly mounting the export in your host system is important. Example 5-6 shows the parameters to set when using a UNIX system.

Example 5-6 Suggested parameters for using a UNIX system

```
Linux: mount -o rw,soft,intr,nolock,timeo=3000,nfsvers=3,proto=tcp
<PTServerIPAdd>:/<ExportPath> /<mountpoint>
Solaris 10: mount -o rw,soft,intr,llock,timeo=3000,vers=3,proto=tcp
<PTServerIPAddress>:/<ExportPath> /<mountpoint>
Solaris 11: mount -o rw,soft,intr,llock,timeo=3000,vers=3,proto=tcp
<PTServerIPAddress>:/<ExportPath> /<mountpoint>
AIX: mount -o
rw,soft,intr,llock,timeo=3000,vers=3,proto=tcp,rsz=262144,wsz=262144
<PTServerIPAddress>:/<Exportpath> /<path>

[root@flash ~]# mount -o rw,soft,intr,nolock,timeo=3000,nfsvers=3,proto=tcp
10.0.25.129:/flash /mnt/flash
[root@flash ~]# mount
/dev/mapper/vg_flash-lv_root on / type ext4 (rw)
proc on /proc type proc (rw)
sysfs on /sys type sysfs (rw)
devpts on /dev/pts type devpts (rw,gid=5,mode=620)
tmpfs on /dev/shm type tmpfs (rw,rootcontext="system_u:object_r:tmpfs_t:s0")
/dev/sda2 on /boot type ext4 (rw)
/dev/sda1 on /boot/efi type vfat (rw,umask=0077,shortname=winnt)
/dev/mapper/vg_flash-lv_home on /home type ext4 (rw)
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)
10.0.35.129:/flash on /mnt/flash type nfs
(rw,soft,intr,nolock,timeo=3000,nfsvers=3,proto=tcp,addr=10.0.35.129)
backup_utility on /mnt/fuse type fuse.backup_utility (rw,nosuid,nodev)
```

After it is mounted, the export can be used with your backup application.

5.3.2 ProtecTIER authentication and user management

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

- ▶ Active Directory
- ▶ Workgroup

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

Active Directory and user IDs

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

Active Directory realm

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

Some helpful commands can be used to define the realm:

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

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

- ▶ From the ProtecTIER server, run the following command:

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

Example 5-7 shows output for the **net ads lookup** command.

Example 5-7 Output of the net ads lookup command

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

5.4 FSI file system scalability

An FSI file system can scale to the following values:

- ▶ Maximum virtual file systems per repository: 128; 48 on SM2
- ▶ Maximum nominal virtual file system size: 256 TB
- ▶ Maximum files per virtual file system: 1 million
- ▶ Maximum files per repository: 16 million; up to 3 million on SM2
- ▶ Maximum “open files” per replication (streams): 192; 64 on SM2



Host attachment considerations for Virtual Tape Library

This chapter describes the preferred practices for connecting hosts to the IBM System Storage TS7600 ProtecTIER family Virtual Tape Library (VTL), including device driver specifications for various operating system platforms, such as IBM AIX®, UNIX, Linux, and Solaris. This chapter also describes the suggested settings for LUN masking, persistent device name binding, and considerations about control path failover (CPF).

Note: This chapter applies only when you use the VTL emulation feature of ProtecTIER. It does not apply to the File System Interface (FSI).

This chapter describes the following topics:

- ▶ General suggestions (to connect a backup application host to a ProtecTIER VTL system)
- ▶ Device driver specifications
- ▶ LUN masking for VTL systems

For further details about the topics in this chapter and managing your tape environment, see the *IBM Tape Device Drivers Installation and User's Guide*, GC27-2130:

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

6.1 General suggestions

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

- ▶ Ensure that the operating system (OS) and version of your host, and your backup server version, are listed as supported in the IBM Interoperability Matrix and in the Backup Application independent software vendor (ISV) Support Matrix. For more information, see 6.2, “Device driver specifications” on page 80.
- ▶ Install the suggested device driver on the host, as specified in the IBM Interoperability Matrix. For more information, see 6.2, “Device driver specifications” on page 80.
- ▶ When possible, configure CPF to enable redundancy to access virtual robot devices. For more information, see 6.2, “Device driver specifications” on page 80.
- ▶ Set up persistent device naming to avoid changes on devices that are recognized by the OS after a system restart. For example, persistent naming can be configured under Linux by using the udev device manager. For more information, see 6.2, “Device driver specifications” on page 80.

When you set up persistent naming, do not use SAN Discovery in Tivoli Storage Manager (because version 7.1.3 was rebranded to IBM Spectrum Protect). The Tivoli Storage Manager SAN Discovery function discovers IBM devices that are based on the original OS device name, not based on customized devices names as they are created, for example, with udev.

- ▶ When you share a VTL across several backup hosts, enable the LUN masking feature and configure LUN masking groups, as described in 6.3, “LUN masking for VTL systems” on page 87.

6.2 Device driver specifications

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

To access the IBM SSIC and ProtecTIER ISV Support Matrix. More details are in Appendix B, “ProtecTIER compatibility” on page 457.

To ensure that your host hardware is supported, and that the firmware versions are at the minimum required levels. Review the Notes section of the ProtecTIER ISV Support Matrix. This section specifies which device driver must be installed on the host to work with the ProtecTIER VTL.

Table 6-1 summarizes which device driver to choose (either IBM Tape Device Driver, native OS driver, or ISV device driver) for each application. To confirm detailed information about version and specific configurations, see the latest release of the ProtecTIER ISV Support Matrix.

Table 6-1 Summary of recommended device drivers by each backup application

Backup application	IBM Tape Device Drivers	Native OS driver or ISV driver
IBM Tivoli Storage Manager (Rebranded to IBM Spectrum Protect as of v 7.1.3)	All platforms	Not applicable (NA)
Symantec Veritas NetBackup (NetBackup)	AIX with NetBackup version 6.5.2 and later. When an IBM tape driver is used, its multipath function must be disabled.	<ul style="list-style-type: none"> ▶ AIX with NetBackup older than version 6.5.2, requires the Symantec ovpass driver. ▶ Solaris: Requires the Symantec sg driver for drives and solaris st driver for the robot/changer devices. ▶ All other platforms.
EMC NetWorker	<ul style="list-style-type: none"> ▶ Windows ▶ AIX ▶ Solaris: Requires solaris sg driver for the robot changer devices and IBM tape driver for tape drives. 	<ul style="list-style-type: none"> ▶ Linux ▶ HP-UX ▶ Solaris
Commvault	<ul style="list-style-type: none"> ▶ Windows ▶ AIX 	All other platforms: Native OS drivers
HP Data Protector	Windows	All other platforms: Native OS drivers

Although the following sections list the specifications grouped by OS, always see the latest release of the ISV Support Matrix for current information. To access the latest IBM SSIC and ProtecTIER ISV Support Matrix, see Appendix A, “ProtecTIER compatibility” on page 465.

6.2.1 AIX specifications to work with VTL

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

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

6.2.2 Solaris specifications to work with VTL

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

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

6.2.3 Linux specifications to work with VTL

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

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

6.2.4 Windows specifications to work with VTL

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

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

6.2.5 IBM Tape Device Driver

For the IBM Tape Device Driver, an installation and user guide contain detailed steps to install, upgrade, or uninstall the device driver for all supported OS platforms. See the *IBM Tape Device Drivers Installation and User's Guide*, GC27-2130:

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

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

To download the IBM Tape Device Driver for your platform, complete the following steps:

1. Go to the IBM Fix Central website:
<http://www.ibm.com/support/fixcentral>
2. Click the **Select product** tab and complete the following steps:
 - a. From the Product Group drop-down menu, select **System Storage**.

- b. From the Select from System Storage drop-down menu, select **Tape systems**.
- c. From the next Select from Tape systems drop-down menu, select **Tape drivers and software**.
- d. From the Select from Tape drivers and software drop-down menu, select **Tape device drivers**.
- e. From the Platform drop-down menu, select your operating system. You can select the generic form of the platform (Linux) and *all* device drivers for that platform are listed.
- f. Click **Continue**. In the window that opens, select the download that you need.

6.2.6 Control path failover and data path failover

The path failover features ensure the use of a redundant path in the event that communication over the primary path fails. These path failover features are built in to the IBM Tape Device Drivers, and are enabled by default for ProtecTIER.

ProtecTIER offers path failover capabilities that enable the IBM Tape Device Driver to resend a command to an alternate path. The IBM Tape Device Driver initiates error recovery and continues the operation on the alternate path without interrupting the application.

Two types of path failover capabilities exist:

- ▶ **Control Path Failover (CPF)**. Control refers to the command set that controls the library (the SCSI Medium Changer command set).
- ▶ **Data Path Failover (DPF)**. Data refers to the command set that carries the customer data to and from the tape drives.

Path failover means the same in both: that is, where there is redundancy in the path from the application to the intended target (the library accessory or the drive mechanism, respectively), the IBM Tape Device Driver transparently fails over to another path in response to a break in the active path.

Note: CPF and DPF are activated when errors or events occur on the physical or transmission layer, such as cable failures, HBA hardware errors (and some software HBA errors), and tape drive hardware errors.

Both types of failover include host-side failover when configured with multiple HBA ports connected into a switch, but CPF includes target-side failover through the control paths that are enabled on more than one tape drive. DPF includes target-side failover for the dual-ported tape drives, but ProtecTIER does not virtualize dual-ported tape drives. Table 6-2 summarizes CPF and DPF support offered on ProtecTIER.

Table 6-2 CPF and DPF support enabled on ProtecTIER

Failover type	Host side	Target side
Control Path Failover	With multiple HBA ports connected into a switch	With robot accessible through multiple ports
Data Path Failover	With multiple HBA ports connected into a switch	Not supported

For more details about CPF and DPF see the Common extended features topic in *IBM Tape Device Drivers Installation and User's Guide*, GC27-2130.

Because CPF and DPF require use of the IBM Tape Device Driver, the backups applications connected to ProtecTIER that will support path failover features are those that use the IBM Tape Device Driver.

Note: CPF and DPF are not features owned by IBM Spectrum Protect (known as Tivoli Storage Manager prior to version 7.1) and are not related to the IBM Spectrum Control SAN Discovery.

In the ProtecTIER Manager, you can verify that CPF is enabled, which is the default, by checking the properties of a defined library in the Configuration window (Figure 6-1).

Tip: To use CPF, in addition to it being enabled for the library, make sure to have more than one robot enabled and available in the library.




Figure 6-1 Control path failover mode enabled at ProtecTIER Manager

Enabling control path failover in IBM Spectrum Protect (formerly Tivoli Storage Manager)

To enable CPF/DPF in an AIX system with Tivoli Storage Manager, enable path failover support on each SCSI medium changer by running the **chdev** command in AIX:

- ▶ `chdev -l smc0 -aalt_pathing=yes`
- ▶ `chdev -l smc1 -aalt_pathing=yes`

Primary and alternative paths

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

On AIX systems, on SCSI attached devices, **-P** is appended to the location field. On Fibre attached devices, **-PRI** is appended to the location field of the device (Example 6-1 on page 85). When a second logical device is configured with path failover support enabled for the same physical device, it configures as an alternative path.

On SCSI attached devices, **-A** is appended to the location field. On Fibre attached devices, **-ALT** is appended to the location field of the device (Example 6-1 on page 85). A third logical device is also configured as an alternative path with either **-A** or **-ALT** appended, and so on. The device driver supports up to 16 physical paths for a single device.

If `smc0` is configured first, and then `smc1` is configured, the `lsdev -Cc tape` command output is similar to Example 6-1 on page 85.

Example 6-1 Primary and alternative path example for Fibre attached devices on AIX

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

Configuring CPF: Detailed procedures of how to configure CPF for AIX and other platforms are in the topic about installing and configuring OS device drivers in *IBM System Storage TS7600 with ProtecTIER Version 3.3*, SG24-7968, and in the *IBM Tape Device Drivers Installation and User's Guide*, GC27-2130.

Redundant robots with Symantec NetBackup V6.5.2

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

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

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

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

If other paths to the library robot exist, you can configure them as alternative paths by enabling multiple path support in NetBackup. Use the NetBackup **robtest** utility to enable and manage multiple path support for library robots.

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

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

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

6.2.7 Persistent device naming

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

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

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

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

An example where persistent naming is useful is a specific host that always assigns the same device name to the first tape library and drives that it finds (Figure 6-2).¹

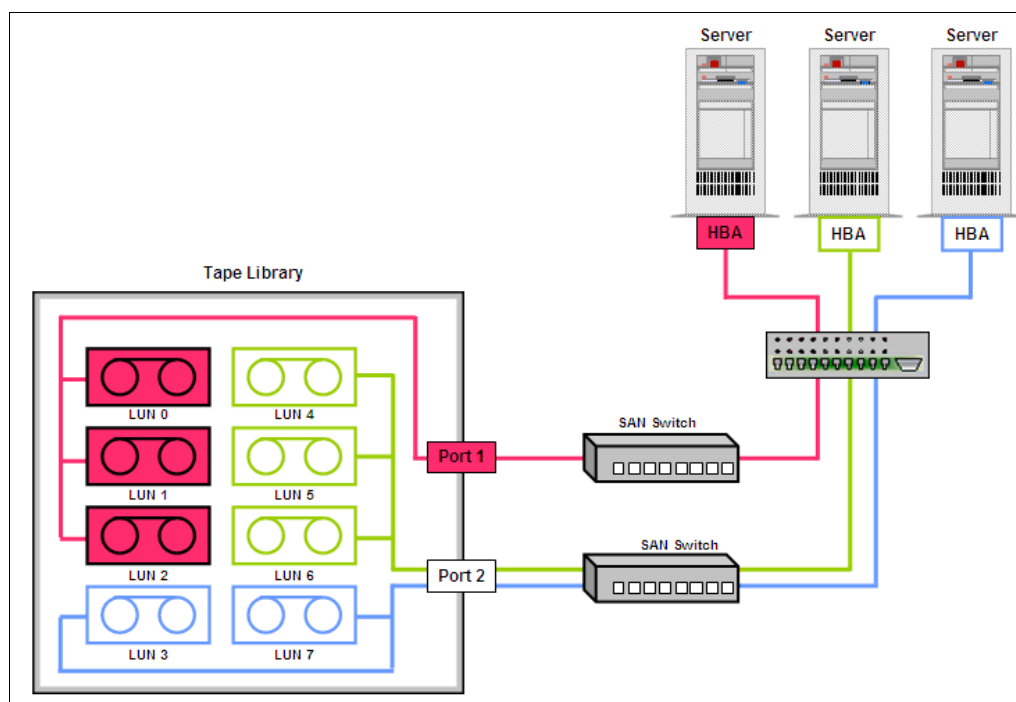


Figure 6-2 Persistent device name binding

Why persistent device naming matters

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

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

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

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

¹ Reference: <http://www.storagesearch.com/dataLink-art1.html>

Suppose that the tape devices are deleted from the system (by running `rmdev -d1 rmt1` and `rmdev -d1 rmt2`) before you restart the machine. On the next restart, if the existing rmt1 target 1 device is powered off, or not connected, the HBA initially configures two devices as Available with the following logical names:

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

If the previous rmt1, target 1 device is powered on after restart and the `cfgmgr` command is run; the HBA configures the device as rmt2 rather than rmt1:

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

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

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

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

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

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

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

Detailed procedures of how to configure persistent device naming for AIX and other platforms are in the *IBM Tape Device Drivers Installation and User's Guide*:

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

6.3 LUN masking for VTL systems

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

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

The LUN masking setup can be monitored and modified at any time during system operation. Every modification to LUN masking in a ProtecTIER server that might affect the host configuration requires rescanning by the host systems. By default, LUN masking is disabled.

Without LUN masking, all of the devices in the environment are visible to all of the FC attached hosts in the fabric if SAN zoning is set up accordingly. When you enable LUN masking, no LUNs are assigned to any backup host, and the user must create LUN masking groups and associate them with the backup hosts.

Figure 6-3 shows the management of a ProtecTIER environment. The ProtecTIER system includes several devices, such as tape drives and robots. Each device is assigned a LUN ID. The administrator manages two hosts, and each host has two HBA ports, where each HBA port has a unique worldwide name (WWN).

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

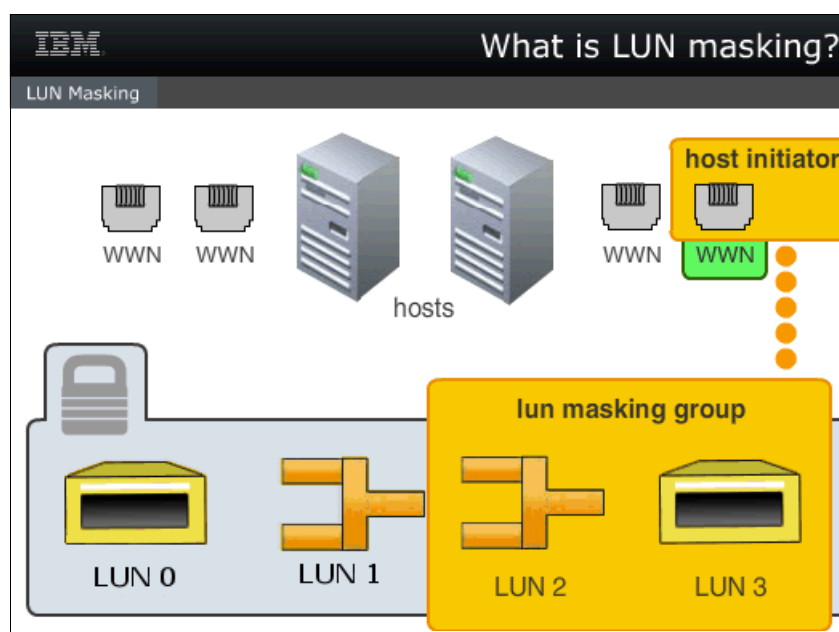


Figure 6-3 LUN masking scenario

6.3.1 LUN masking methods and preferred practices

Use LUN masking to manage device visibility. LUN masking conceals specific devices (tape drives or robots) from the view of host initiators while enabling a selected host initiator group to view them. Several preferred practices for LUN masking are as follows:

- ▶ Define host aliases to identify the host ports. When you define backup host aliases, use a practical naming scheme as in the following example:
 - hostname-FE0 (for front-end port 0)
 - hostname-P0 (for port 0)
- ▶ With more than two backup hosts, use LUN masking to load balance ProtecTIER performance across multiple front-end ports.

- ▶ Regardless of LUN masking, virtual drives are physically assigned to one front-end port, so backup hosts must be attached to that single port. For load balancing purposes, distribute drives across multiple front-end ports. If possible, distribute drives across all four front-end ports.
- ▶ Use LUN masking to establish two or more front-end paths to a backup server for redundancy. For example, the following configurations are relevant:
 - In environments with up to four backup servers, you can dedicate a single front-end port to each backup server rather than using LUN masking, but with the disadvantage of missing load balancing across multiple front-end ports and missing redundancy.
 - In environments where front-end ports are shared, and you want to prevent backup hosts from sharing, use LUN masking to isolate each backup host.

6.3.2 LUN masking configuration steps

Figure 6-4 shows the steps for a LUN masking configuration.

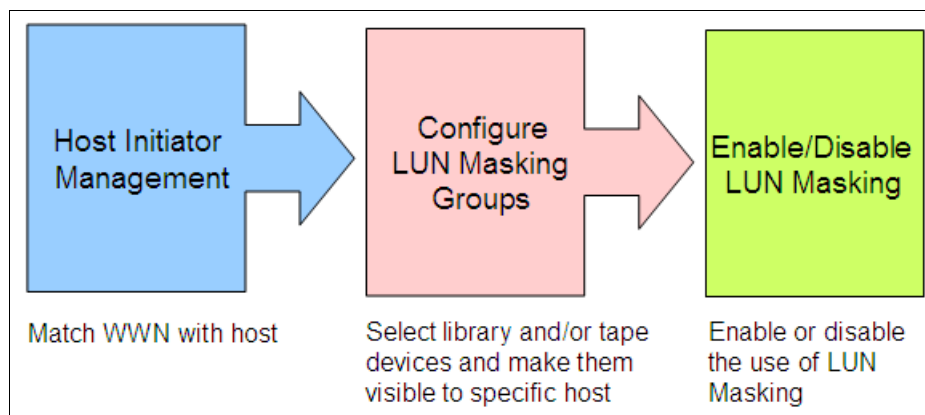


Figure 6-4 LUN masking configuration steps

Host initiator management

First, perform host initiator management by completing the following steps:

1. From ProtecTIER Manager, click **VT** → **Host Initiator Management**. A list of available host initiators is displayed (Figure 6-5).
2. Select one or more host initiators from the list, or manually add the host initiator by entering the appropriate WWN, as shown in Figure 6-5.

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

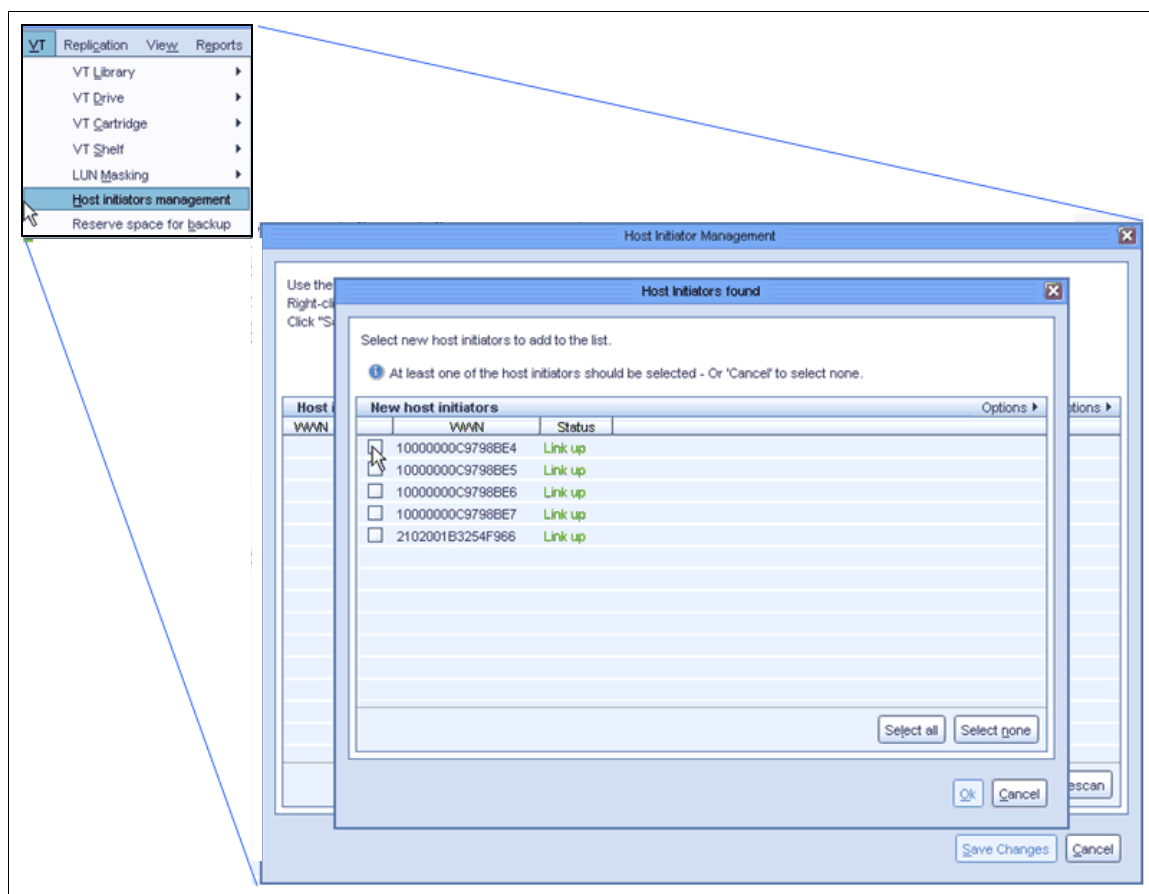


Figure 6-5 Host initiator management

3. You can also assign an alias to the WWN by clicking **Modify** (Figure 6-6). Aliases help you more easily identify which host is related to the WWN. The ProtecTIER worldwide port names (WWNs) are found in the Host Initiator Management window.

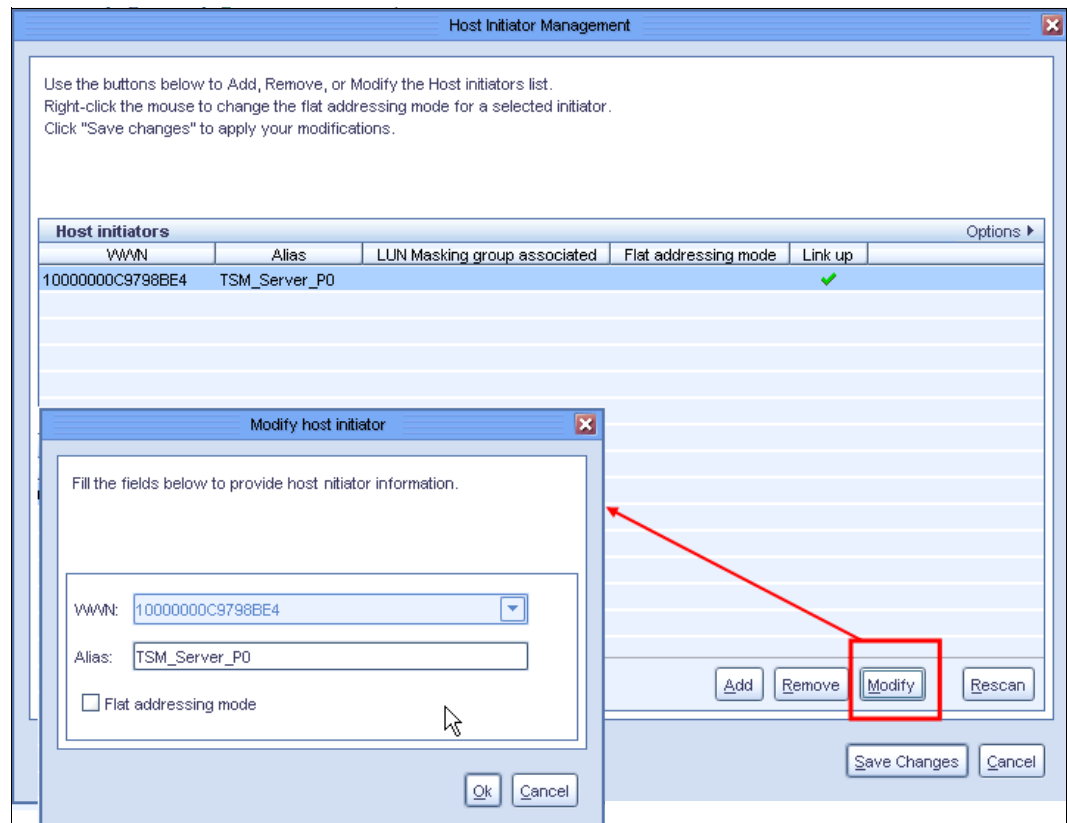


Figure 6-6 Modifying a host initiator alias

LUN masking groups

Now that you defined a host initiator, you can create a LUN masking group. Figure 6-7 shows the LUN Masking Group window.

Use this dialog to Create, Edit, and/or delete LUN Masking group(s). The updated configurations are only saved on the ProtecTIER repository once you click 'Save changes'. Until then, you can continue modifying the LUN masking groups. The changes are applied to the tape emulation devices when LUN masking is enabled.
Note that the changes made in the LUN Masking configuration require scanning by the host and by the back up application.

LUN Masking groups Options

Name
TSM_group

Group name: TSM_group

Selected Host Initiators

WWN	Alias
10000000c9798be4	TSM_Server_P0

Library Mappings

Library	Devices
vail	1 Robot 2 Drives

Selected Devices

Library	Device	Serial	LEA	WWN	Port	Node	LUN
vail	Robot	1255640999	0	10000000c9798be6	2	lbsdedup1a	0
vail	Drive 0	1255640000	0	10000000c9798be4	0	lbsdedup1a	0
vail	Drive 1	1255640001	1	10000000c9798be6	2	lbsdedup1a	1

Save Changes Cancel

Figure 6-7 LUN Masking Group window

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

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

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

Important:

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

Reassigning LUNs

After you modify a LUN masking group, unwanted gaps might occur in the LUN numbering sequence.

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

To reassign a LUN, complete the following steps:

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

3. The system displays the Reassign LUNs window, which has the following message (as shown in Figure 6-8):

You are about to renumber all the LUN values of the available devices in the group and all host connected must be rescanned

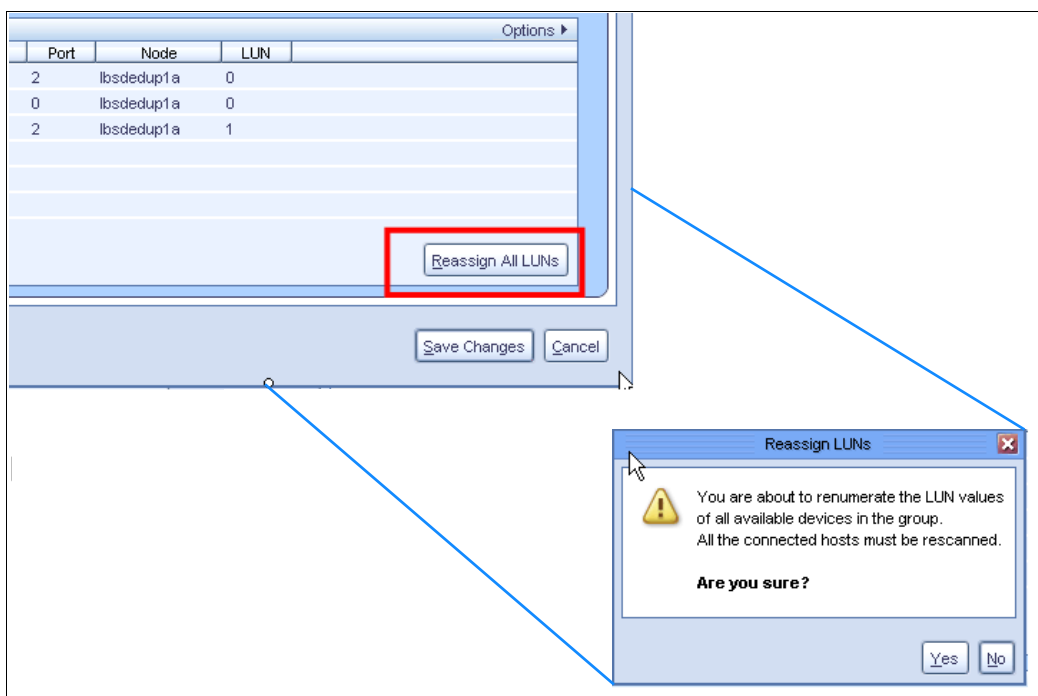


Figure 6-8 Reassigning LUNs

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

Enabling or disabling LUN masking

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

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

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

3. When the Enable/Disable LUN masking window opens, select **Enable LUN masking**, and click **OK**, as shown in Figure 6-9.

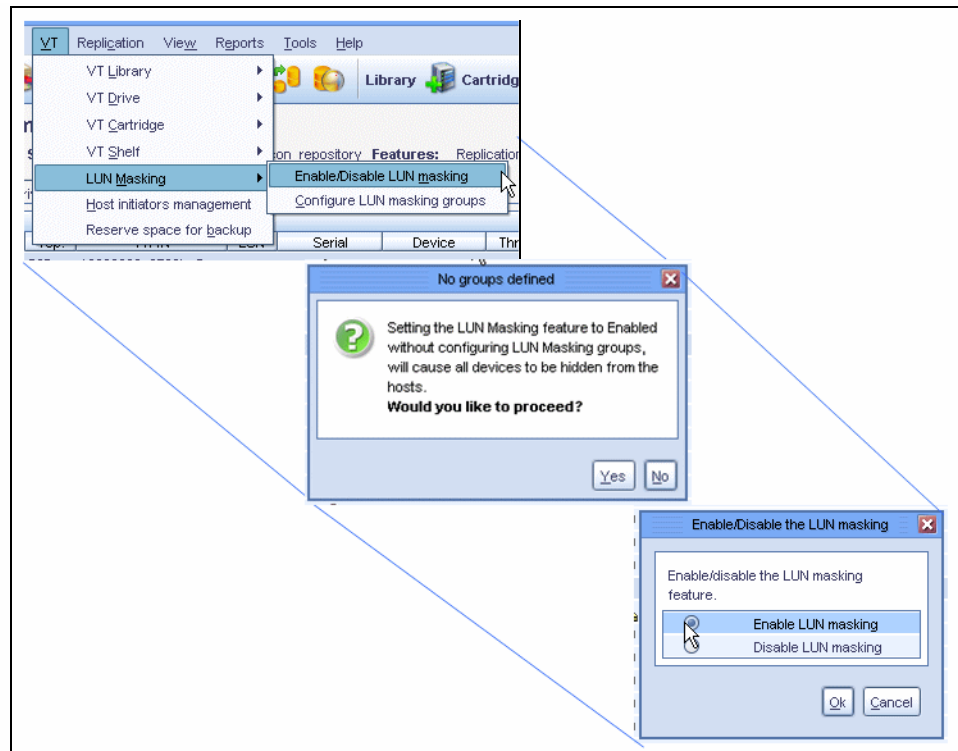


Figure 6-9 Enabling/Disabling LUN masking

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

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



Part 2

Back-end storage subsystems

This part describes IBM System Storage TS7600 ProtecTIER family preferred practices and configuration guidelines for specific back-end storage subsystems. It also describes principal rules to follow for storage resources that are used in a ProtecTIER environment, and contains the following chapters:

- ▶ Chapter 7, “Back-end storage overview” on page 99
- ▶ Chapter 8, “IBM FlashSystem 900 with System Storage TS7650G ProtecTIER Deduplication Gateway” on page 111
- ▶ Chapter 9, “IBM Storwize Family and SAN Volume Controller” on page 125
- ▶ Chapter 10, “IBM XIV Storage System” on page 157
- ▶ Chapter 11, “IBM System Storage DS8000” on page 167
- ▶ Chapter 12, “Backup management introduction” on page 179
- ▶ Chapter 13, “IBM Spectrum Protect” on page 187
- ▶ Chapter 14, “Symantec NetBackup and BackupExec” on page 207
- ▶ Chapter 15, “EMC NetWorker” on page 221
- ▶ Chapter 16, “HP Data Protector” on page 229
- ▶ Chapter 17, “IBM i and Backup, Recovery, and Media Services” on page 241
- ▶ Chapter 18, “Commvault” on page 257
- ▶ Chapter 19, “Veeam FSI-CIFS” on page 283

Supported storage systems: For a list of IBM and third-party disk storage subsystems that are supported by the ProtecTIER deduplication servers, see the TS7650/TS7650G ISV and Interoperability Matrix. For more information, see Appendix B, “ProtecTIER compatibility” on page 457.



Back-end storage overview

Back-end storage is one of the key components of the IBM System Storage TS7650G ProtecTIER Deduplication Gateway (TS7650G) implementation. It consists of the storage subsystems that attach to the ProtecTIER Gateway nodes.

In contrast to the ProtecTIER Appliance models, which are preconfigured with disks, the ProtecTIER Gateway attaches to back-end storage subsystems that then contain the repository. This chapter lists the key factors and common configuration requirements for back-end storage that is supported by the TS7650G.

This chapter describes the following topics:

- ▶ Overview
- ▶ Dependencies from a back-end storage subsystem view
- ▶ Dependencies from a ProtecTIER view
- ▶ Smart storage subsystems
- ▶ Key rules for a ProtecTIER server
- ▶ Storage arrays configuration
- ▶ Storage area network fabric

7.1 Overview

ProtectTIER back-end storage is a critical hardware component that holds the ProtectTIER repository. Three types of file systems make up the ProtectTIER repository:

Cluster database	Single 1 gigabyte logical unit number (1 GB LUN) that holds the cluster configuration information. It is in a metadata Redundant Array of Independent Disks (RAID) group, and is seen as a metadata file system. This item is also called a <i>quorum</i> .
Metadata	Stores all the aspects about the data that is backed up and indexed, but not the actual data. The references for restoring are kept in the metadata file systems. These file systems must be on RAID 10 arrays.
User data	Stores the actual data that is backed up, and referenced by new generations of backups. These file systems must be on RAID 6 or RAID 5 arrays.

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

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

Figure 7-1 on page 101 shows an example of the ProtectTIER Performance Planner. This tool is primarily used by the IBM Pre-sales and IBM Business Partners. Based on the planner, to achieve 500 megabytes per second (MBps) and accommodate a 120 terabyte (TB) repository with a 10:1 HyperFactor ratio, the storage array should be configured as follows:

► Metadata

- 24 drives each of 300 GB 15,000 revolutions per minute (RPM) of Fibre Channel (FC)
- 3 x 4+4 RAID 10 groups

The number of RAID groups given in this example does not account for the file system to store the cluster database, which is 1 GB in size and does not require a separated RAID group. Therefore, you can build two LUNs of 1200 GB each, one LUN of 1199 GB, and one LUN of 1 GB to present to the ProtectTIER to use for metadata. Or you can build the three LUNs of 1200 GB, and do the arrangements for the Cluster database at a disk partition level on the ProtectTIER server.

► User data

- 176 drives each of 1 TB 7200 RPM of Serial Advanced Technology Attachment (SATA)
- 22 x 6+2 RAID 6 groups

You can build 22 LUNs of 6 TB each, and present them to ProtectTIER to use for user data.

Figure 7-2 shows an example of the ProtecTIER Meta Data Planner. Based on information that is shown in Figure 7-2, the requirements for the metadata are as follows:

- ▶ Four metadata file systems (including the 1 GB file system for the Cluster database).
- ▶ The minimum sizes required for file systems are: 1 GB, 1042 GB, 1049 GB, and 650 GB.

Figure 7-2 Example of ProtecTIER Meta Data Planner

Important: Always assign the full RAID array capacity to the LUNs that are used for ProtecTIER metadata and user data. The only exception to this rule is the *Cluster database* file system of 1 GB, which *does not* require a separate RAID group.

7.1.1 ProtecTIER Planner tool

Availability: The ProtecTIER Planner tool is an IBM internal tool that is available to trained ProtecTIER specialists.

The core component for any capacity sizing and subsequent configuration effort is the ProtecTIER Planner. The primary methodologies to accurately size the required capacity and configure the disk and file system infrastructure depend on the correct usage of this tool.

The primary function of the ProtecTIER Planner enables field engineers to perform key activities when they size and configure the physical disk systems that are connected as back-end storage of the ProtecTIER server. The process starts at a high level, where a general capacity sizing is performed, based on key variables in the organization's environment.

Another aspect of the sizing process is to understand how many metadata and user data file systems are required based on disk technologies and RAID configurations to ensure correct performance. The ProtecTIER Performance Planner aids in this process.

The ProtecTIER Metadata Planner enables the field engineer to understand how many metadata file systems are required to support a repository of a certain size and the size of the metadata file systems.

There are other utilities in the ProtecTIER Planner, such as upgrading from a previous version (with the import of historical user data), and customizing any disk performance information based on unique user scenarios when planning performance.

For more information about capacity planning for a TS7650 Appliance environment, see *IBM System Storage TS7600 with ProtecTIER Version 3.3*, SG24-7968.

The ProtecTIER Planner should be used for capacity and performance planning while manually adapting it to a many-to-many environment at an early stage before the ProtecTIER deployment. The parameters to be considered are the same ones that are used in any replication deployment:

- ▶ System workload. Both local backup and replication, including incoming and outgoing.
- ▶ Network bandwidth. Available and required between the PT nodes on the grid.
- ▶ Time frames. Backup and replication windows or concurrent operation.

7.2 Dependencies from a back-end storage subsystem view

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

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

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

All IBM Storage Subsystems are based on standard hardware. For example, the IBM System Storage DS8000® is based on the IBM System p architecture, and the IBM Storwize® Family and SAN Volume Controller are based on the IBM System x architecture. With these architectures, the controller nodes of the storage subsystems tend to have multiple processors driving the input/output (I/O), and multiple cores per processor.

To use all of these resources, you must have multiple LUNs per controller at the same time. For example, with the IBM Storwize Family and SAN Volume Controller, you must have at least four arrays per controller to ensure the optimal usage of all computing cores.

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

7.3 Dependencies from a ProtecTIER view

ProtecTIER uses a special method to write data to disk. ProtecTIER writes data to disk by creating data in or appending data to the repository structure. These are the two principal modes of operation. Data that can be expired from a ProtecTIER perspective is not overwritten. It is deleted by the maintenance processes and is reshaped by the defragger practices before it is used again by write or append operations.

This concept of data management directly benefits clients by having multiple file systems that ProtecTIER can work with at the same time. If your ProtecTIER repository enables backup speeds of 2500 MBps or more, you need at least 32 file systems in your back-end storage for ProtecTIER to work with.

If you aim for a medium performance solution, you can use fewer than 32 file systems. Again, these assumptions are more general ones, and your individual system will be fine if these assumptions are implemented according to your individual sizing and planning.

7.4 Smart storage subsystems

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

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

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

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

These features provide the following primary enhancements:

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

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

Because of the design method that ProtecTIER uses to write data to disk, ProtecTIER prevents creating a higher load on a single array only, while others stay idle. ProtecTIER does not benefit from automatic hot spot avoidance.

ProtecTIER needs multiple LUNs to work with. Some of these LUNs are used to store the metadata, which must be RAID 10 to drive the heavily random write I/O. Some of these LUNs are used to store the user data, which must be RAID 6 or RAID 5 to drive the random read I/O. ProtecTIER evenly distributes the load across the storage system, and across all of the available LUNs.

ProtecTIER also uses a special method to extend the back-end storage during the ProtecTIER repository expansion process. Depending on the amount of storage that you want to add, you can add or extend metadata, add user data, or both. Depending on how the initial deployment of ProtecTIER was done, increasing the metadata might be done by adding

storage to existing metadata file systems (extend file system), or it might be done by adding some new metadata file systems and leave the existing ones untouched.

For the user data, because of the padding process that allocates all of the file system's storage immediately during implementation, the repository growth includes adding new file systems. Expanding the ProtecTIER user data does not involve growing individual LUNs or file systems. To grow the ProtecTIER user data, more user data LUNs and file systems must be added.

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

It is, however, theoretically possible that the usage of storage pool striping or rotate extents might collocate multiple and popular deduplication candidates in the repository on a reduced number of spindles.

7.5 Key rules for a ProtecTIER server

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

- ▶ Disk-based replication is supported only by request for product quotation (RPQ). Use the ProtecTIER native IP replication feature that is available in Version 2.5 and later. For more information about replication, see Part 4, "Replication and disaster recovery" on page 359.
- ▶ If you use a SAN switch to connect the TS7650G to the disk array, create dedicated zoning for ProtecTIER back-end ports. Do not mix the back-end ports with the front-end ports or any other SAN devices in the same zone.
- ▶ Dedicate the whole storage array to the TS7650G. If this configuration is not possible, make sure that the I/O requirements of ProtecTIER can be ensured and are not affected by other applications. Make sure that there is no congestion or oversubscription of resources because of other applications that might affect ProtecTIER arrays. Use zoning to isolate the TS7650G traffic from other applications. The TS7650G can never share RAID groups/arrays or LUNs with other applications.
- ▶ ProtecTIER creates a heavily random-read I/O pattern. About 80 - 90% of all I/O requests in a typical TS7650G environment are random reads. Because of the binary differential comparison, ProtecTIER creates this pattern even during backup traffic to ProtecTIER. The I/O pattern resembles the one for a database application. Therefore, implement suitable performance optimizations and tuning as suggest by the disk vendor for database I/O profiles.
- ▶ Because of the increased flexibility and the robustness to protect against cabling errors, use worldwide port name (WWPN)-based zoning (soft zoning). Direct attachment of back-end storage subsystems is supported for the Storwize and the DS8000 product families.
- ▶ RAID 10 must be used for metadata. Use RAID 6 for user data arrays with disk sizes larger then 900 GB. You can use RAID 5 for user data arrays with disks sizes of 900 GB or less.

7.6 Storage arrays configuration

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

7.6.1 General requirements

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

Storage array firmware

The storage array firmware level should be equal to or greater than the firmware version listed in the IBM SSIC and ProtecTIER ISV Support Matrix. More details are in Appendix B, “ProtecTIER compatibility” on page 457.

Storage cache

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

Host mapping

The ProtecTIER server host type (host connectivity settings) must be tuned for a Red Hat Enterprise Linux (RHEL) device-mapper-multipath client. For example, you should map metadata LUNs and user data LUNs as LNXCLUSTER (which has Automatic Value Transfer (AVT) disabled) for host mapping.

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

Storage arrays with only active-passive support

In storage arrays with only active-passive support, where LUNs can be accessed only by one disk controller at a time, LUN mapping must be interleaved between controllers to establish effective load balancing. Select the preferred path of the LUNs such that the following conditions are true:

- ▶ LUNs with even numbers are assigned to Controller A.
- ▶ LUNs with odd numbers are assigned to Controller B.

7.6.2 RAID considerations

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

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

Important: To maximize reliability of the ProtecTIER system, always use RAID 6 for solutions with larger disk sizes. Use RAID 5 only with disk sizes of 900 GB max . Today's large disk drive capacities cause long rebuild times in case of drive failure. In RAID 5 design, this implies a significant threat for *double disk failure* during RAID rebuild.

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

Metadata

The following list includes the RAID considerations regarding metadata:

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

Important: Because the average SATA disk provides only a limited number of I/O operations per second (IOPS) in comparison to a SAS or FC spindle, the use of SATA drives for metadata has a negative effect on performance.

User data

The following list includes the RAID considerations regarding user data:

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

7.6.3 LUNs

Create only one LUN per RAID group. The only exception to this rule is the cluster database of 1 GB, which can be co-located with any metadata LUN together on the same array.

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

Metadata

The following list includes guidelines for LUN creation regarding metadata:

- ▶ Metadata LUNs must be created on a different RAID group than user data LUNs.
- ▶ Create a 1 GB LUN for a cluster database on any RAID with a metadata LUN.
- ▶ The number and size of metadata LUNs are determined during the pre-installation discussion with a trained ProtecTIER specialist with the Metadata Planner (see Figure 7-2 on page 102).

User data

The following list includes guidelines for LUN creation regarding user data:

- ▶ As with metadata, the number of user data LUNs is determined during the pre-installation discussion with a trained ProtecTIER specialist. For optimal performance, create at least 24 LUNs.
- ▶ Having a repository with 24 or more LUNs is optimal for the best performance.
- ▶ The size of user data LUNs must be consistent.

Important: A high number of LUNs attached to ProtecTIER increases the length of boot time. Starting with ProtecTIER Version 3.2, the management of LUNs greater than 8 TB is improved. When ProtecTIER V3.3 works with LUNs greater than 8 TB, it splits them into logical volumes of a smaller size. Therefore, you can work with LUNs greater than 8 TB, but there is no benefit in performance in completing this action.

7.6.4 Expanding the repository

When you expand the repository, use the same spindle type and quality of RAID groups for metadata and user data. For example, if existing metadata LUNs were built on 4+4 RAID groups, then new metadata RAID groups must be at least 4+4. In this example, if 2+2 or 3+3 RAID groups are used, it degrades overall system performance because of an IOPS bottleneck.

7.7 Storage area network fabric

Directly connecting ProtecTIER nodes to hosts (backup servers) and storage arrays is possible. You also can connect the components into a SAN fabric. The connection between ProtecTIER nodes and hosts (backup servers) is referred to as a front-end connection, and the connection between ProtecTIER nodes and storage arrays is referred to as a back-end connection. For the updated list of supported SAN switches of ProtecTIER, see the IBM SSIC and ProtecTIER ISV Support Matrix. For more details, see Appendix B, “ProtecTIER compatibility” on page 457.

Figure 7-3 illustrates an example of SAN fabric and zoning.

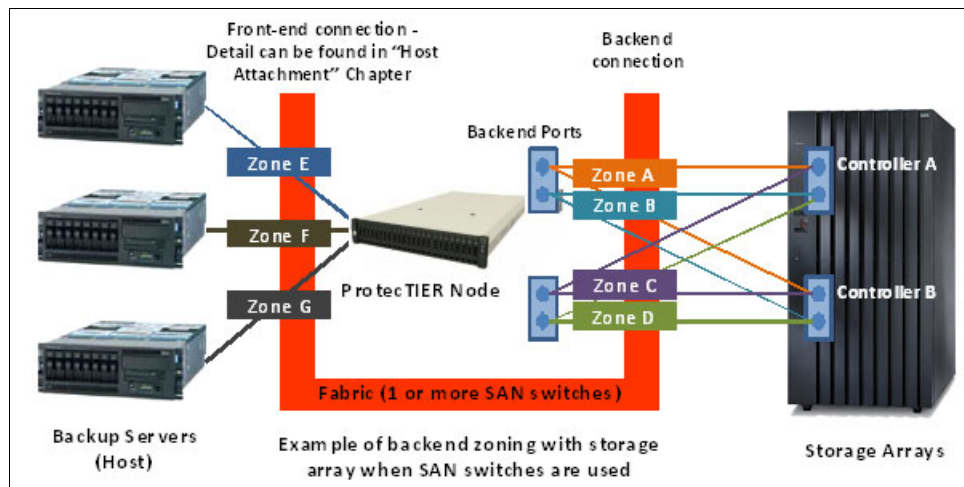


Figure 7-3 Example of SAN fabric and zoning

7.7.1 Two Fibre Channel paths to each storage controller

Each ProtecTIER node has four FC ports on the back end for storage connectivity. You should use two paths per storage controller.

7.7.2 Dedicated zones

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

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

7.7.3 Front-end zones

Create a dedicated zone for ProtecTIER front-end ports to the host (backup application). Do *not* mix this zone with other SAN devices. It should have only one initiator per zone (single-initiator zoning). For example, one port of your IBM Spectrum Protect (formerly Tivoli Storage Manager) server is zoned to all of the available ProtecTIER front-end ports. For more information about front-end connectivity to a host, see “Chapter 6, “Host attachment considerations for Virtual Tape Library” on page 79.

7.7.4 Back-end zones

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

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

The zoning topology that is shown in Figure 7-3 on page 109 is one example of front-end and back-end connectivity of a ProtecTIER server. For more information about front-end connectivity requirements, see Chapter 6, “Host attachment considerations for Virtual Tape Library” on page 79. For more information about back-end zoning of different storage systems, see Part 5, “Appendixes” on page 433.

7.7.5 SAN paths

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



IBM FlashSystem 900 with System Storage TS7650G ProtecTIER Deduplication Gateway

This chapter introduces the IBM FlashSystem storage product, its core value, benefits, technological advantages, and configuration steps to use it with the IBM System Storage TS7650G ProtecTIER Deduplication Gateway (TS7650G).

This chapter contains the following topics:

- ▶ Introduction to flash technology
- ▶ IBM FlashSystem overview
- ▶ General FlashSystem considerations
- ▶ Deploying the IBM FlashSystem 900 with IBM ProtecTIER
- ▶ General preferred practices integrating FlashSystem 900 with ProtecTIER

8.1 Introduction to flash technology

Flash technology in the data center is too relevant to be ignored for a few simple reasons:

- ▶ Since its introduction, flash memory has improved across all metrics, offering higher performance, density, and reliability, all of which translate to improved business efficiency.
- ▶ Flash cost per capacity and cost per transaction relative to hard disk-based storage make it very attractive to businesses that are attempting to maintain pace in a 24x7 competitive marketplace.
- ▶ Flash is easily integrated into existing data center environments, and provides an instant boost to the mission critical applications.

While flash in storage is pervasive in the data center, its implementation varies considerably among competitors and technologies. Some use it as a simple cache accelerator, and others implement it as yet another permanent data tier. The reality is that flash only matters when two conditions in the data center are met:


- ▶ Flash eliminates input/output (I/O) bottlenecks and generates higher levels of application efficiency (improved performance).
- ▶ Storage economics are improved by its use. It provides lower total cost of ownership (TCO), reducing costs, and faster return on investment (ROI) to the existing environment, enabling new business opportunities.

IBM FlashSystem™ storage delivers high performance, efficiency, and reliability for shared enterprise storage environments. It helps clients address performance issues with their most important applications and infrastructure. FlashSystem storage provides a hardware-only data path that unleashes all of the potential of flash memory. These systems are different from traditional storage systems, both in the technology and usage.

Figure 8-1 shows more details about FlashSystem 900 integrated with a TS7650G model DD5.

IBM FlashSystem with IBM TS7650G ProtecTIER Gateway

Extreme Performance



IBM FlashSystem 900

- Extreme performance with IBM MicroLatency™
- From 2 TB to 57 TB of usable capacity (using RAID 5)
- AES XTS 256 data-at-rest encryption for applications needing to safeguard valuable data.
- IBM-enhanced Micron MLC technology for higher storage density and improved endurance.
- Energy-efficient 2U form factor

with

Data Deduplication

IBM ProtecTIER TS7650G

- Improve backup and recovery and simplify disaster-recovery operations.
- Lower operational costs and energy usage
- Manage more data with less infrastructure




Figure 8-1 IBM FlashSystem with IBM TS7650G ProtecTIER Gateway

Note: Ensure that your host hardware is listed as supported on the IBM SSIC and ProtecTIER ISV Support Matrix. For more details, see Appendix B, “ProtecTIER compatibility” on page 457.

8.2 IBM FlashSystem overview

Flash technology has fundamentally changed the paradigm for IT systems, enabling new use cases and unlocking the scale of enterprise applications. Flash technology enhances performance, efficiency, reliability, and design of essential enterprise applications and solutions by addressing the bottleneck in the IT process (data storage), enabling truly optimized information infrastructure.

The FlashSystem shared flash memory systems offer affordable, high-density, ultra low-latency, high-reliability and scalable performance in a storage device that is both space-efficient and power-efficient. IBM Flash products, which can either augment or replace traditional hard disk drive storage systems in enterprise environments, empower applications to work faster and scale further.

In addition to optimizing performance, the FlashSystem family brings enterprise reliability and macro efficiency to the most demanding data centers, allowing businesses to accomplish the following tasks:

- ▶ Reduce customer complaints by improving application response time
- ▶ Service more users with less hardware
- ▶ Reduce I/O wait and response times of critical applications
- ▶ Simplify solutions
- ▶ Reduce power and floor space requirements
- ▶ Speed up applications, thereby enhancing the pace of business
- ▶ Improve usage of existing infrastructure
- ▶ Compliment existing infrastructure
- ▶ Mitigate risk

When used with the ProtecTIER System, FlashSystem can provide additional improvements and benefits:

- ▶ With the addition of data deduplication, the amount of disk storage (FlashSystem storage) that is required to store data and keep it online can be reduced.
- ▶ Performing restore operations from FlashSystem products can be faster than any currently existing storage device (disks or physical tapes).
- ▶ Running an environment with ProtecTIER and FlashSystem adds the ability to implement a robust disaster recovery process using the ProtecTIER Replication functionality.
- ▶ Running an environment with ProtecTIER and FlashSystem as back-end storage adds the ability to use other applications, such as Virtual Tape Library (VTL) or File System Interface (FSI).

From the client business perspective, FlashSystem provides focused benefits and value on four essential areas:

Extreme Performance	Enable business to unleash the power of performance, scale, and insight to drive services and products to market faster.
IBM MicroLatency®	Achieve competitive advantage through applications that enable faster decision making due to microsecond response times.
Macro Efficiency	Decrease costs by getting more from efficient use of information technology (IT) staff, IT applications and IT equipment due to the efficiencies flash brings to the data center.
Enterprise Reliability	Durable and reliable designs that use enterprise class flash and patented data protection technology.

The IBM FlashCore™ technology, used in IBM FlashSystem 900, employs several new and patented mechanisms to achieve greater capacity and throughput, at a lower cost than the previous IBM FlashSystem 840. Figure 8-2 shows the three major areas within IBM FlashCore technology and the unique IBM attributes of each one.

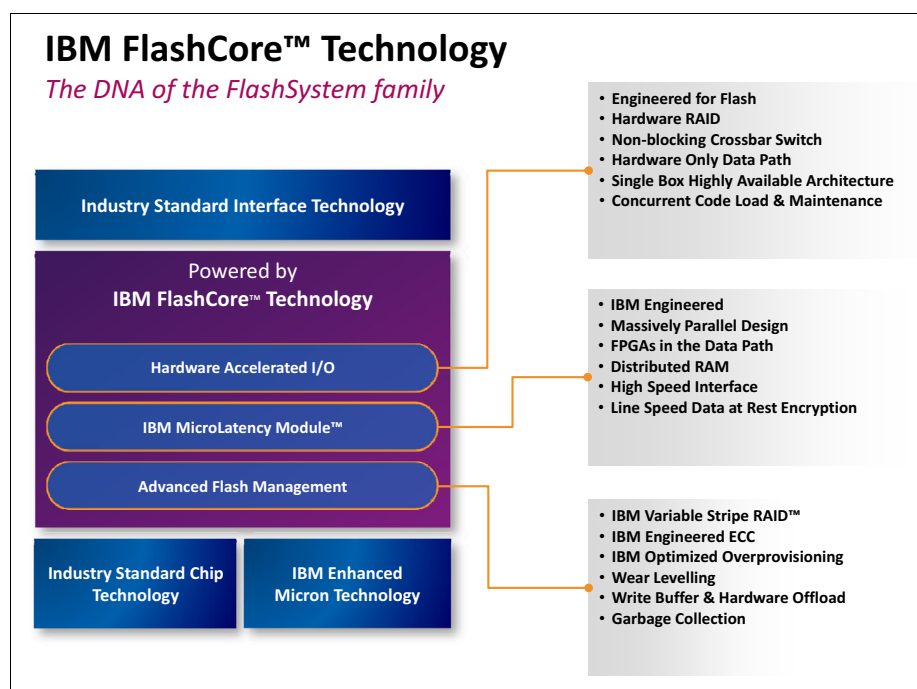


Figure 8-2 IBM FlashCore technology

8.3 General FlashSystem considerations

To make their backup environments faster and more efficient, IBM clients can have IBM FlashSystem integrated to the IBM TS7650 ProtecTIER Gateway. In this topic, the FlashSystem works as the back-end storage for the TS7650G ProtecTIER Deduplication Gateway, adding all its benefits to the backup environment that ProtecTIER provides.

IBM TS7650G ProtecTIER also provides disaster recovery capabilities that can be used by FlashSystem clients by replicating backup data to a different location. The ProtecTIER Internet Protocol (IP) replication function provides a powerful tool that can be used to design robust disaster recovery architectures. You electronically place backup data into vaults with much less network bandwidth, thus changing the paradigm of how data is taken off-site for safekeeping.

For VTL clients, the ProtecTIER IP replication feature can eliminate some of the expensive and labor-intensive handling, transport, and securing of the real tapes for disaster recovery (DR) purposes. For FSI clients, they can implement enhanced functionality only possible with disk, and avoid the limitation of tape emulation.

The IBM TS7650G ProtecTIER backend storage array is configured with two types of logical unit numbers (LUNs): *Metadata* and *user data*. Metadata LUNs are used to record where data is kept, and user data LUNs are used to store the actual data. Metadata LUN performance is critical, and 15,000 revolutions per minute (RPM) hard disk drive (HDD) spindles in a Redundant Array of Independent Disks (RAID10) 4+4 configuration are commonly used.

IBM FlashSystem 900 can be used by the IBM TS7650G ProtecTIER Gateway as the back-end storage device for metadata and user data LUNs.

A common use case for IBM FlashSystem 900 is for the IBM TS7650G ProtecTIER metadata LUNs. Compared to the cost of dozens of 15,000 or 10 kilobits (Kb) HDD spindles, using FlashSystem products for the IBM TS7650G ProtecTIER metadata LUNs can be more cost-effective. Also a more cost-effective approach might be to use FlashSystem for the entire IBM TS7650G ProtecTIER repository if high performance, but small capacity is needed.

Note: If FlashSystem is used day one for metadata, then all metadata expansions must be with FlashSystem storage. Plan for growth from the beginning.

For LUNs, plan for maximum metadata (2500 megabytes per second, or MBps) with fastest 15,000 8+8 configuration in the ProtecTIER Planner Tool. For more information about the ProtecTIER Planner Tool, see 7.1.1, “ProtecTIER Planner tool” on page 102. FlashSystem 900 is configurable with 2 - 48 terabytes (TB) of usable capacity for increased flexibility, which should cover most metadata configurations.

8.4 Deploying the IBM FlashSystem 900 with IBM ProtecTIER

The process to set up the environment to integrate the IBM ProtecTIER system and the FlashSystem 900 consists of the following steps:

1. Create the user data and metadata Volumes in the FlashSystem.
2. Map the volumes to the ProtecTIER System.
3. Create the File Systems and build the ProtecTIER repository.

For details about the initial setup for the FlashSystem 900, such as how to create the volumes and host mapping in the FlashSystem 900, see *Implementing IBM FlashSystem 900*, SG24-8271.

8.4.1 Creating user data and metadata volumes

Before creating the volumes that ProtecTIER will use for metadata and user data, consider that the suggested number of metadata RAID groups (managed disks, or MDisks) is determined by the Capacity Planning Tool during the pre-sales process. This number can be in the range of 2 - 10, based on repository size, factoring ratio, and performance needs.

The following steps describe the process of creating volumes and defining host access for the FlashSystem 900. Create each one of the metadata or user data volumes as follows:

1. Log in to the FlashSystem graphical user interface (GUI).
2. Select the **Volumes** option from the left navigation menu (Figure 8-3 on page 116).



Figure 8-3 Opening Volumes menu on the IBM FlashSystem 900

3. Select the **Create Volume** option.
4. Provide the details for the new volume: Name, Capacity, and Quantity (Figure 8-4).
5. Click **OK**.




Figure 8-4 Creating volumes on the IBM FlashSystem 900

6. Complete the previous setup parameters and repeat the same process until you have created all volumes for the MDisks that ProtecTIER will use for metadata or user data according to the ProtecTIER Capacity Planning Tool.

8.4.2 Mapping the FlashSystem Volumes to the ProtecTIER System

To map the FlashSystem volumes created in the previous section to the ProtecTIER system, perform the following steps.

Note: Before continuing, consider that you will need to update your storage area network (SAN) zoning configuration, and make sure that the IBM FlashSystem 900 is part of the same zone as the ProtecTIER system. In this environment, the host is the ProtecTIER system, and we assume in this section that you already created the host that will be attached to the required volumes in the FlashSystem 900.

Follow these steps:

1. Log in to the Flash System GUI web interface,
2. Select the **Hosts** → **Volumes by Host** (Figure 8-5).




Figure 8-5 Opening the Volumes by Host menu in the IBM FlashSystem 900

3. Expand the **Unmapped Volumes** section.

4. Right-click the volume that you want to map and select **Map to Host** (Figure 8-6).

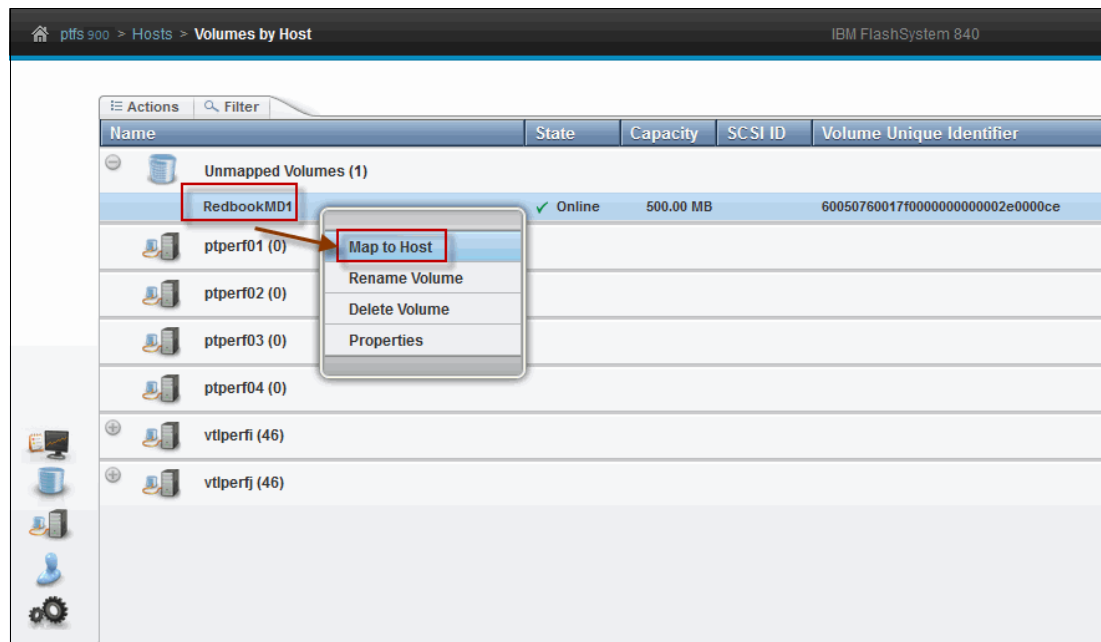


Figure 8-6 Open the Map to Host menu in the IBM FlashSystem 900

5. In the **Select the Host** list, select the host that you want to map and then click **Map** (Figure 8-7).




Figure 8-7 Select host in the IBM FlashSystem 900

8.4.3 Creating file systems and building the ProtecTIER repository

To create the file systems with ProtecTIER, complete the following steps:

1. Verify that the ProtecTIER node recognizes the volumes that are presented by the IBM FlashSystem. The **multipath -ll** command shows the LUNs and the paths that are connected to the storage subsystem. If the LUNs do not appear to be mapped to the storage, you must run a rescan on the Fibre Channel (FC) adapters, or restart the node.

Note: You can scan the FC adapters by running the following command:

```
echo "- -" > /sys/class/scsi_host/<host??>/scan
```

In this command, <host??> should be replaced by each FC adapter port.

2. You can also filter to see only part of the data, as shown in Example 8-1. The first command filters only the volumes of type FlashSystem. The second command shows the details of one of the devices.

Example 8-1 multipath command output

```
[root@lbsdedup1a ~]# multipath -ll | grep FlashSystem
[root@tepic ~]# multipath -ll | grep FlashSystem
mpath2 (20020c24008123fff) dm-2 IBM,FlashSystem
mpath1 (20020c24007123fff) dm-1 IBM,FlashSystem
mpath0 (20020c24006123fff) dm-0 IBM,FlashSystem
mpath3 (20020c24009123fff) dm-3 IBM,FlashSystem

[root@tepic ~]# multipath -ll | grep -A11 mpath0
mpath0 (20020c24006123fff) dm-0 IBM,FlashSystem
[size=2.0T][features=0][hwhandler=0][rw]
\_ round-robin 0 [prio=1][active]
  \_ 19:0:0:6 sdb 8:16 [active][ready]
  \_ round-robin 0 [prio=1][enabled]
    \_ 21:0:0:6 sdf 8:80 [active][ready]
.....
```

3. In the ProtecTIER Service menu, select the options 1) ProtecTIER Configuration (...) and then 6) File Systems Management (...), as shown in Example 8-2.

Example 8-2 ProtecTIER configuration menu

```
[root@lbsdedupla ~]# su - ptconfig
+-----+
| ProtecTIER Service Menu running on vela |
+-----+
| 1) ProtecTIER Configuration (...)      |
| 2) Manage ProtecTIER services (...)    |
| 3) Health Monitoring (...)            |
| 4) Problem Alerting (...)             |
| 5) Version Information (...)           |
| 6) Generate a service report           |
| 7) Generate a system view             |
| 8) Update ProtecTIER code              |
| 9) ProtecTIER Analysis (...)           |
|                                       |
| E) Exit                               |
+-----+
Your choice? 1

+-----+
| ProtecTIER Service Menu running on vela |
| ProtecTIER Configuration (...)          |
+-----+
| 1) Configure ProtecTIER node            |
| 2) Recover Configuration for a replaced server |
| 3) Configure machine serial number for a replaced server |
| 4) Configure RAS                       |
| 5) Update Time, Date, Timezone and Timeserver(s) |
| 6) Scan storage interconnections        |
| 7) File Systems Management (...)        |
| 8) Configure replication (...)          |
| 9) IP Network configuration (...)       |
| 10) Update Firmware (...)              |
| 11) Update the System's name            |
| 12) Validate configuration              |
| 13) Single Node - code upgrade (For Support Use ONLY) |
| 14) OS images Management (...)          |
| 15) Replace SAS Drive                  |
|                                       |
| B) Back                               |
| E) Exit                               |
+-----+
Your choice? 6
```

4. In the File Systems Management menu, you have options to display information about the devices or file systems configuration as shown in Example 8-3

You can configure the file systems on all available devices (Example 8-3), or for a single device.

Example 8-3 File Systems Management menu

```

+-----+
| ProtecTIER Service Menu running on lbsdedup1a |
| ProtecTIER Configuration (...)                |
| File Systems Management (...)                |
+-----+
| 1) Configure file systems on all available devices |
| 2) Create file system(s) on a single unused device |
| 3) Extend a file system with a new unused device  |
| 4) Update /etc/fstab                             |
| 5) Display configured devices                     |
| 6) Display unused devices                         |
| 7) Display GFS repository file systems            |
| 8) Display unused GFS file systems                |
| 9) Increase capacity completion (applicable for a second cluster node) |
|                                                    |
| B) Back                                           |
| E) Exit                                           |
+-----+
Your choice? 1
Begin Processing Procedure
Devices to be configured
...
Please confirm:? (yes|no) y
Creating physical volume           [ Done ]
Creating volume group              [ Done ]
Creating logical volume             [ Done ]
Creating file system                [ Done ]
...
End Processing Procedure Successfully

```

5. After you create the file system for all of the devices that are needed, you can go to the ProtecTIER Manager, add the node to it, and then select the **Repository** menu to create the repository.

For more information about how to create the repository from the ProtecTIER Manager GUI, see the topic about creating a repository for TS7650G Gateway in *IBM System Storage TS7600 with ProtecTIER Version 3.3*, SG24-7968.

8.5 General preferred practices integrating FlashSystem 900 with ProtecTIER

This section describes some preferred practices that you can use when working with ProtecTIER and IBM FlashSystem 900:

- ▶ Make sure that you use ProtecTIER v3.3.5.2 or later.
- ▶ Always use RAID 5 Configuration with FlashSystem 900.
- ▶ The suggested zoning is to use four paths between each FlashSystem volume and each ProtecTIER Server.

Deploy FlashSystem according to your environment needs:

- ▶ ProtecTIER metadata only:
 - Use any number of supported FlashSystem 900 modules in RAID 5 configuration according to the required repository metadata capacity.
 - Determine metadata capacity and layout based on ProtecTIER Metadata Planner tool.
- ▶ High I/O per second (IOPS) host structures and ProtecTIER metadata.
- ▶ Entire ProtecTIER repository (user data and metadata):
 - Use 12-module FlashSystem 900 in a RAID 5 configuration.
 - Create at least 40 user data volumes of equal size. User data volumes map to repository file systems:
 - Max 170 file systems in a ProtecTIER deduplication repository.
 - Ensure file systems take into account future growth considerations.
 - Create metadata volumes according to the ProtecTIER Metadata Planner output.

For more details about IBM FlashSystem 900, see the following publications:

- ▶ *FlashSystem 900 Product Guide*, TIPS1261
- ▶ *Implementing IBM FlashSystem 900*, SG24-8271

8.5.1 ProtecTIER Metadata Planner preferred practices for FlashSystem 900

The following preferred practices are for FlashSystem and the ProtecTIER Metadata Planner:

- ▶ Use 4+4 15,000 RPM FC drives as input for the Metadata Planner.
- ▶ Provide binary repository capacity input to the Metadata Planner.
- ▶ Determine and configure the FlashSystem 900 volumes according to the Metadata Planner output.

Example: Considering a 750 terabytes (TiB) repository, and 10:1 deduplication ratio, the FlashSystem 900 would require 16,204.8 gibibytes (GiB; 8 x 4 TB module). See Figure 8-8 for more details.

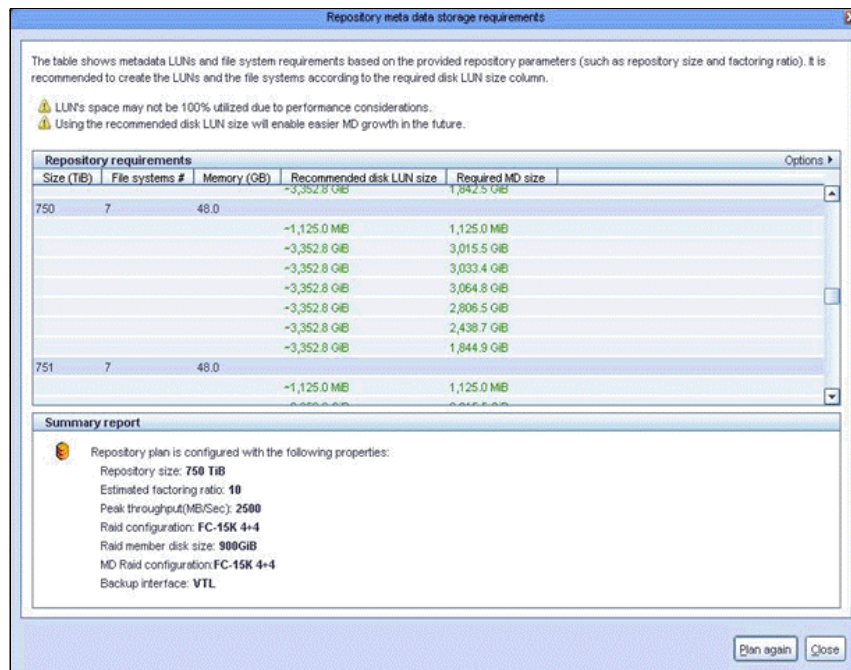


Figure 8-8 ProtecTIER Metadata Planner



IBM Storwize Family and SAN Volume Controller

This chapter describes the IBM Storwize V3700 (V3700), the IBM Storwize V5000 (V5000), the IBM Storwize V7000 and V7000 Unified, and the IBM SAN Volume Controller. It describes how these storage virtualization products can be connected to the IBM System Storage TS7600 ProtecTIER family system as back-end storage. This chapter also describes the suggested topology, volume (logical unit numbers, or LUNs) configuration, and settings.

The following topics are described:

- ▶ IBM Storwize V3700 and V5000 overview
- ▶ General considerations: V3700, V5000, and V7000
- ▶ IBM System Storage SAN Volume Controller, IBM Storwize V7000, and V7000 Unified Storage
- ▶ General notes
- ▶ Firmware level
- ▶ Fibre Channel connection topology
- ▶ User data and metadata pool: General suggestions
- ▶ Configuration steps

9.1 IBM Storwize V3700 and V5000 overview

The V3700 and V5000 are affordable, easy to use, and self-optimizing storage virtualization solutions with advanced functionality and reliability that is usually found only in more expensive enterprise systems.

The V3700 and V5000 are the newest members of the Storwize family, delivering flexible entry to midrange storage, using demonstrated IBM Storwize V7000 (V7000) and IBM System Storage SAN Volume Controller functions, management tools, and interoperability.

The V3700 and V5000 can operate in Internet Small Computer System Interface (iSCSI) and Fibre Channel (FC) environments. The V3700 and V5000 can be extended by small form factor (SFF) enclosures with 24 x 2.5 in. drives, and large form factor (LFF) enclosures with 12 x 3.5 in. drives. High-performance disk drives, high-capacity nearline serial-attached SCSI (NL-SAS) disk drives, and solid-state drives (SSDs) are supported.

The V3700 and V5000 use the proven technology of IBM Storwize V7000 to provide great performance and advanced features, such as virtualization, thin provisioning, copy services, and nondisruptive migration.

Table 9-1 shows the features of the Storwize family.

Table 9-1 Storwize Family features

Storwize Family hardware engines		
V3700	V5000	V7000
Up to 240 drives	Up to 480 drives per system and 960 drives in two-way clusters	Up to 504 drives per enclosure and 1056 per clustered system
Small <ul style="list-style-type: none">▶ Internal virtualization▶ Thin provisioning▶ Data migration▶ Advanced GUI▶ Multi-host support▶ VMware and Microsoft integration▶ Flash Copy (2040 images)▶ IBM Easy Tier®▶ Remote Mirror▶ Turbo performance	Medium V3700 plus these: <ul style="list-style-type: none">▶ Clustering (2)▶ Higher scalability▶ Higher performance▶ External virtualization▶ Thin provisioning▶ One-way data migration	Large V5000 plus these: <ul style="list-style-type: none">▶ Clustering (4)▶ Highest scalability▶ Highest performance▶ Compression▶ Encryption▶ Unified

For more information about the storage concepts to use and the general suggestions for ProtecTIER in combination with V3700 and V5000, see 9.2, “General considerations: V3700, V5000, and V7000” on page 127 up to 9.7, “User data and metadata pool: General suggestions” on page 154.

9.2 General considerations: V3700, V5000, and V7000

To tailor your V3700 or V5000 storage to ProtecTIER, you must use the command-line interface (CLI) to configure parts of the system. For information about how to use the graphical user interface (GUI) to support your configuration, see 9.3.1, “Storage virtualization introduction” on page 145 through 9.7, “User data and metadata pool: General suggestions” on page 154.

9.2.1 Configuration steps: ProtecTIER repository

Figure 9-1 shows the steps to configure the ProtecTIER repository on a Storwize V3700, V5000, and V7000.




Figure 9-1 Configuration steps

The steps are as follows:

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

9.2.2 Creating empty user data and metadata storage pools

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

Using the GUI

To create these items, complete the following steps in the GUI:

1. Hover the cursor over the Pools icon, and the system displays the Pools menu (Figure 9-2 on page 128). Click **MDisks by Pools**.




Figure 9-2 Pool menu

2. For each pool, click **New Pool** to create an empty pool.
3. Insert a name for the pool, such as MD_pool1 or UD_pool1 (Figure 9-3). At this time you might want to change the icon of the pool using the **Choose Icon** option. Click **Next** to continue with the pool creation process.

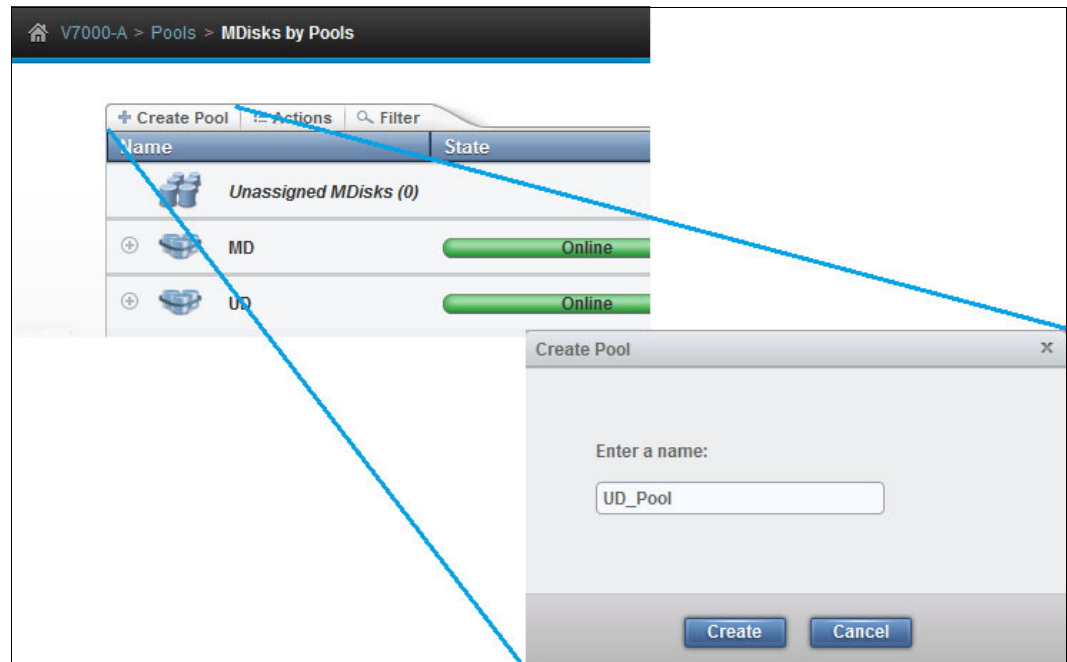


Figure 9-3 Creating pools

Tip: Set the extent size settings to 256 megabytes (MB). This size is the default for version 6. In version 7, the default is 1 gigabyte (GB). The 256 MB extent size is needed to support large repository dimensions of 1 petabyte (PB) or more of virtualization in the storage. Change the default warning threshold of 80% to 100% in order not to receive alerts about pool usage. The pool is fully allocated by the ProtecTIER LUNs. To do this task, use the CLI to create the pools, as shown in Example 9-1 on page 129.

4. At this point, the pool is empty. To add MDisks to the pool, see 9.2.3, “Creating the MDisk arrays or discovering unmanaged MDisks” on page 129. Repeat this procedure for the next pool. You need two pools: one for metadata and one for user data.

Using the CLI

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

Example 9-1 Creating a pool using the CLI

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

9.2.3 Creating the MDisk arrays or discovering unmanaged MDisk

When you use the Storwize V5000 or V7000 server, you can use internal or external disks to create MDisk.

Note: Storwize V3700 allows only internal disks; the SAN Volume Controller allows only external disks.

Using the GUI

To accomplish this task in the GUI, complete the following steps:

1. Hover the cursor over the **Pools** icon. The Pools menu opens (Figure 9-4).

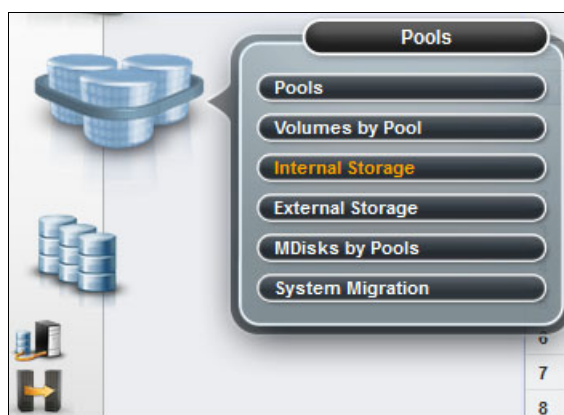


Figure 9-4 Pools menu

2. To create an MDisk with internal disks, click **Internal Storage** to display the list of candidate internal disks. Select the drive class that you want to configure, for example, SAS disks for metadata and NL_SAS (or Serial Advanced Technology Attachment, also known as SATA) for user data. Then, click **Configure Storage** (Figure 9-5 on page 130).

Note: If you want to determine the drives that will be spare, you can manually select them before this step and mark them **spare** in the **Use** column. For that, right-click the drive and select **Mark As** followed by the **spare** option. A good practice is to select one spare drive for each drawer and type of drive.



Figure 9-5 Configure storage

Important: Also remember that when you physically install the drives, if drives are of different sizes, they need to be placed in drawer 0 and drawer 1 equally, or equally spaced across equal number of drawers on each string.

3. If the MDisk is created from external storage, click **External Storage** to display the list of unmanaged MDisks. If no MDisks are displayed, click **Detect MDisks**. When the list of available MDisks opens, select the appropriate disks, right-click, and click **Add to pool**.
4. The Configure Internal Storage window is displayed. The default settings are shown in the window. Proceed as shown in Figure 9-6 on page 131:
 - a. To use different settings, click **Select a different configuration**.
 - b. Select the Redundant Array of Independent Disks (RAID) type, for example, **Balanced RAID 10** for a metadata array.
 - c. Click **Optimized for Performance** to force the GUI into creating RAID arrays of equal dimensions.
 - d. In the Number of drives to provision field, enter the number of disks to configure in your RAID arrays.
 - e. Ensure that the automatically calculated RAID layout resembles the RAID arrays that you want to create. For example, having eight disks in a Balanced RAID 10 creates a 4+4 array. Having eight disks in a RAID 6 creates a 6+2 array. Having five disks in a RAID 5 creates a 4+1 array.

Important: Closely monitor the output of the GUI in step e. Especially when you try to create multiple arrays concurrently, change the number of disks that you want to configure in a single step to avoid unwanted array sizes.

- f. Click **Next**.

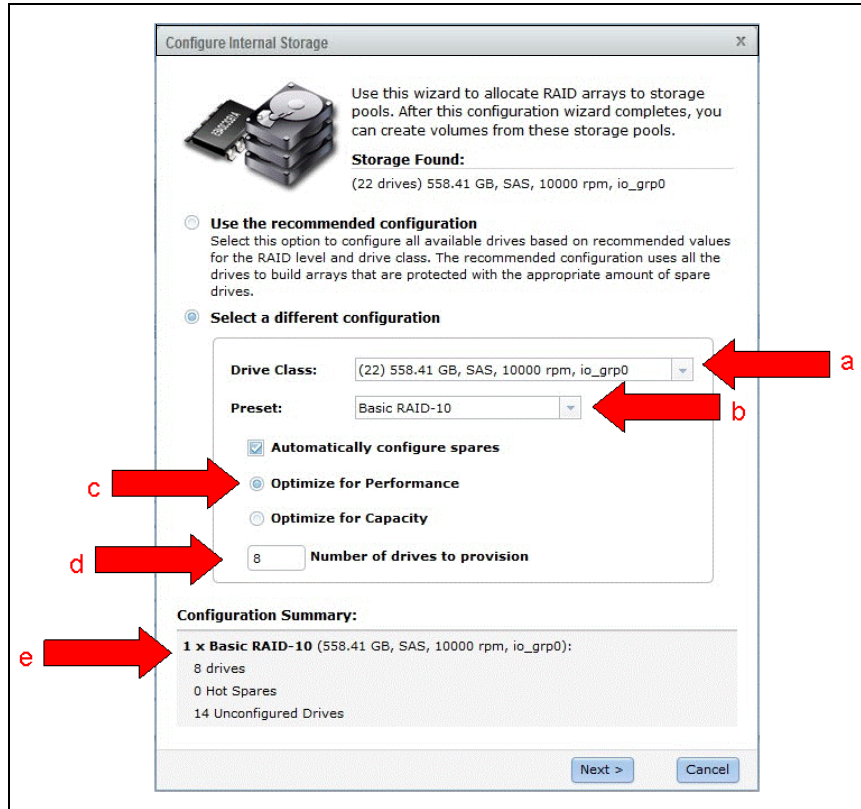


Figure 9-6 Configure Internal Storage window

Using the CLI

These steps can also be done by using the CLI, as shown in Example 9-2.

Example 9-2 MDisk array creation command

```
# svctask mkarray -drive 28:27:26:25 -level raid5 -sparegoal 0 -strip 128 UD_pool
# svctask mkarray -drive 8:6:5:10 -level raid10 -sparegoal 0 -strip 128 MD_pool
```

The numbers after the **-drive** flag represents the physical disk's ID, which is part of the array.

Important:

- ▶ The **-sparegoal 0** parameter permits array creation regardless of whether hot spare drives are available. With this setting, you must manually ensure that enough hot spare resources are available in your storage subsystem to satisfy your requirements.
- ▶ To manually create a balanced RAID 10 array, select the array members in a specific order. The **-drive** parameter uses the drives in the following order: Primary:Mirror:Primary:Mirror:Primary:Mirror, and so on. Be sure that each Primary disk is attached through one SAS chain, and each Mirror disk is attached through the other SAS chain.
- ▶ If you want to change the array attributes, use the **chararray** command as described in the product manuals. Also see the following IBM Knowledge Center web pages for IBM Storwize V3700 and V5000:
 - http://www.ibm.com/support/knowledgecenter/STLM5A_7.6.0/com.ibm.storwize.v3700.760.doc/svc_chararray.html?lang=en
 - http://www.ibm.com/support/knowledgecenter/STHGJ_7.6.0/com.ibm.storwize.v5000.760.doc/svc_chararray.html?lang=en

9.2.4 Creating volumes with a sequential virtualization type

The ProtecTIER architecture stripes the repository among all file systems and uses all the disks simultaneously. With this architecture, create the volume to accommodate the full size of each MDisk. Additionally, because of the nature of the ProtecTIER file system architecture and workload, *sequential* virtualization has better performance than striped virtualization.

To create virtual volumes (VDisks) with the sequential virtualization type, you must use the CLI with the root user ID. The GUI enables only the default configuration, which is striped. To create sequential volumes, complete the following steps:

1. Identify your MDisks by listing them. Run the **svcinfo lsmdisk** command (Example 9-3) to list them.

Example 9-3 The svcinfo lsmdisk output

```

IBM_2072:demo:webguest>lsmdisk
id name      status mode  mdisk_grp_id mdisk_grp_name capacity ctrl_LUN_#
controller_name UID tier
0 mdisk1 online array 2      MD_pool      2.2TB
generic_hdd
1 mdisk0 online array 2      MD_pool      2.2TB
generic_hdd
2 mdisk3 online array 1      UD_pool      2.2TB
generic_hdd
IBM_2072:demo:webguest>lsmdisk -unit mb -filtervalue
"mdisk_grp_name=MD_pool"
id name      status mode  mdisk_grp_id mdisk_grp_name capacity ctrl_LUN_#
controller_name UID tier
0 mdisk1 online array 2      MD_pool      2.2TB
generic_hdd
1 mdisk0 online array 2      MD_pool      2.2TB
generic_hdd
IBM_2072:demo:webguest>lsmdisk -unit mb -filtervalue
"mdisk_grp_name=UD_pool"

```


id	name	status	mode	mdisk_grp_id	mdisk_grp_name	capacity	ctrl_LUN_#
	controller_name	UID	tier				
2	mdisk3	online	array 1	UD_pool	2.2TB		
	generic_hdd						

- As explained in 9.3.1, “Storage virtualization introduction” on page 145, MDisk are composed of extents. The default extent size is 256 MB for version 6, and 1 GB for version 7. The number of extents from one MDisk to another MDisk can vary according to the SAN Volume Controller/Storwize quorum information. To use all the extents in the MDisk, you must verify the number of the free extents by running the **lsfreeextents** command (Example 9-4).

Example 9-4 lsfreeextents output

```
IBM_2072:demo:webguest>lsfreeextents mdisk1
id 0
number_of_extents 8929
IBM_2072:demo:webguest>lsfreeextents mdisk0
id 1
number_of_extents 8934
IBM_2072:demo:webguest>lsfreeextents mdisk3
id 2
number_of_extents 8837
```

- Take the number of free extents that are available in each MDisk and multiply them by the size of the extent, which is 256 MB. For example:
 - For mdisk1, the volume size in MB is 8929 (number_of_extents) x 256 MB (extent size) = 2285824 MB.
 - For mdisk0, the volume size in MB is 8934 (number_of_extents) x 256 MB (extent size) = 2287104 MB.
 - For mdisk3, the volume size in MB is 8837 (number_of_extents) x 256 MB (extent size) = 2262272 MB.

Important:

- If you created the MDisk using the default attributes for version 7, use 1024 MB in the previous calculations. However, a suggestion is to always use 256 MB when creating the MDisk.
- You must complete these calculations because the V3700/V5000 CLI requires you to specify the **-size** parameter if you use **-vtype seq** flag.

- Create the volume (VDisk) by using the **-vtype seq** flag, which means sequential type, using the value that was discovered in step 3 (Example 9-5).

Example 9-5 Sequential volume creation

```
IBM_2072:demo:webguest>mkvdisk -name MD_Quorum -mdiskgrp MD_pool -iogrp io_grp0
-mdisk mdisk0 -size 1024 -unit=mb -vtype seq
Virtual Disk, id [0], successfully created
IBM_2072:demo:webguest>mkvdisk -name MD_vol01 -mdiskgrp MD_pool -iogrp io_grp0
-mdisk mdisk0 -size 2284800 -unit=mb -vtype seq
Virtual Disk, id [1], successfully created
```

Hint: The **-size** value of 2284800 for MD_vol01 is derived from having the quorum on the same array. The number of free extents must be recalculated after quorum creation. The mdisk0 has 8934 extents of 256 MB in size. After you use four of those extents to create the quorum (4 x 256 MB = 1024 MB), the number of free extents is 8930. Multiplying 8930 (number of free extents) with 256 (extents size) leads to 2286080 MB as the value of **-size** for the metadata LUN that collocates with the quorum.

- Run the **svctask mkvdisk** command to create all user data and metadata volumes. Depending on the code version (starting with V7.5), a format will run, by default. Either specify the **-nofmtdisk** flag or increase the sync rate for the format duration.

For more details, see the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/ST3FR7_7.6.1/com.ibm.storwize.v7000.761.doc/svc_createsequentialorstripedmodevdisk_21kd92.html

Tip: The 1 GB metadata quorum volume can be created on any metadata MDisk.

- Run the **svcinfo lsvdisk** command to list all of the created volumes. In Example 9-6, you can see the VDisk user ID (UID), which identifies the LUN in the operating system (OS). You can also see the volume type as seq (sequential) and the volume size.

Example 9-6 svcinfo lsvdisk output

```
IBM_2072:demo:webguest>lsvdisk
id name      IO_group_id IO_group_name status mdisk_grp_id mdisk_grp_name capacity type
FC_id FC_name RC_id RC_name vdisk_UID          fc_map_count copy_count
fast_write_state se_copy_count RC_change compressed_copy_count
0 MD_Quorum 0      io_grp0      online 1          MD_pool      1.00GB seq
60050760008189D18800000000000047 0      1          empty      0
no      0
1 MD_vol01 0      io_grp0      online 1          MD_pool      2.18TB seq
60050760008189D18800000000000048 0      1          empty      0
no      0
IBM_2072:demo:webguest>lsmdisk
```

- You can also see the volumes that are created by using the GUI and going to the Volumes menu (Figure 9-7).




Figure 9-7 Listing volumes

You can customize columns by right-clicking above the column bar and selecting the columns you want (Figure 9-8).

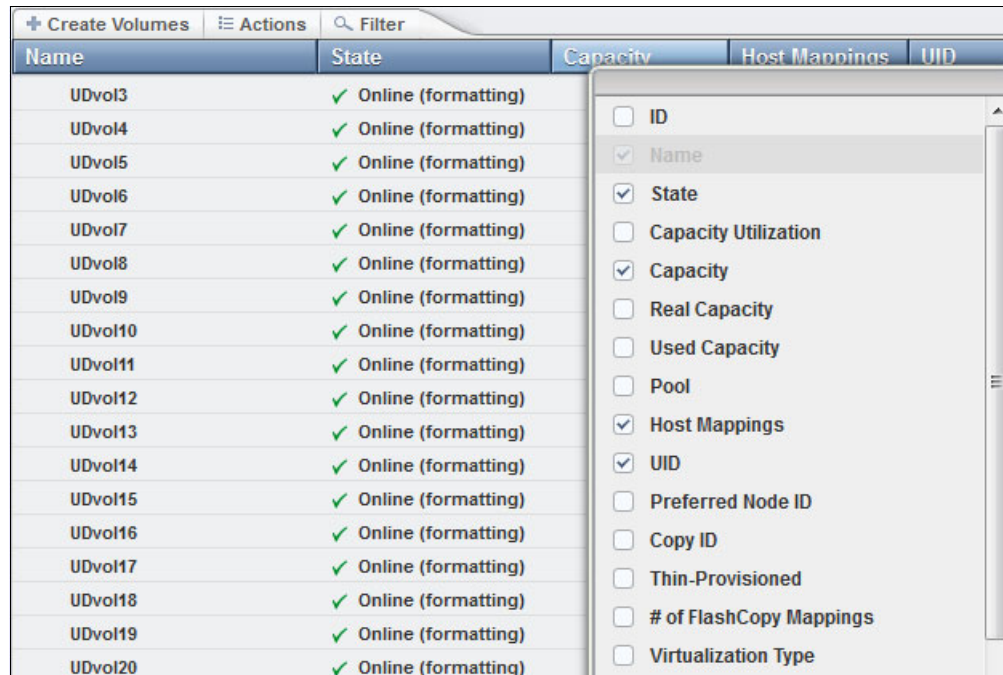


Figure 9-8 Columns options

9.2.5 Creating a host connection for the ProtecTIER nodes by using the GUI

To set up the host in the Storwize V3700/V5000 GUI, you must first know the worldwide port names (WWPNs) of the ProtecTIER nodes. To accomplish this task, complete these steps:

1. View the label that is attached to the FC adapter or download the `system_view.html` file by selecting option **7** (Generate a system view) from the **ptconfig** menu of your system (Example 9-7).

Example 9-7 Generating a system view report

```
+-----+
| ProtecTIER Service Menu running on vela |
+-----+
| 1) ProtecTIER Configuration (...)      |
| 2) Manage ProtecTIER services (...)   |
| 3) Health Monitoring (...)            |
| 4) Problem Alerting (...)             |
| 5) Version Information (...)          |
| 6) Generate a service report          |
| 7) Generate a system view            |
| 8) Update ProtecTIER code            |
| 9) ProtecTIER Analysis (...)          |
|                                       |
| E) Exit                              |
+-----+
Your choice?
Your choice? 7
Begin Processing Procedure

SystemView version 1.2
Generating Service Report                      [ Done ]
Systemview saved to /pt_work/systemview.html

End Processing Procedure Successfully
```

2. Open the `systemview.html` file with a browser. A list of selectable items is displayed. Click **QLogic HBAs** (Figure 9-9).

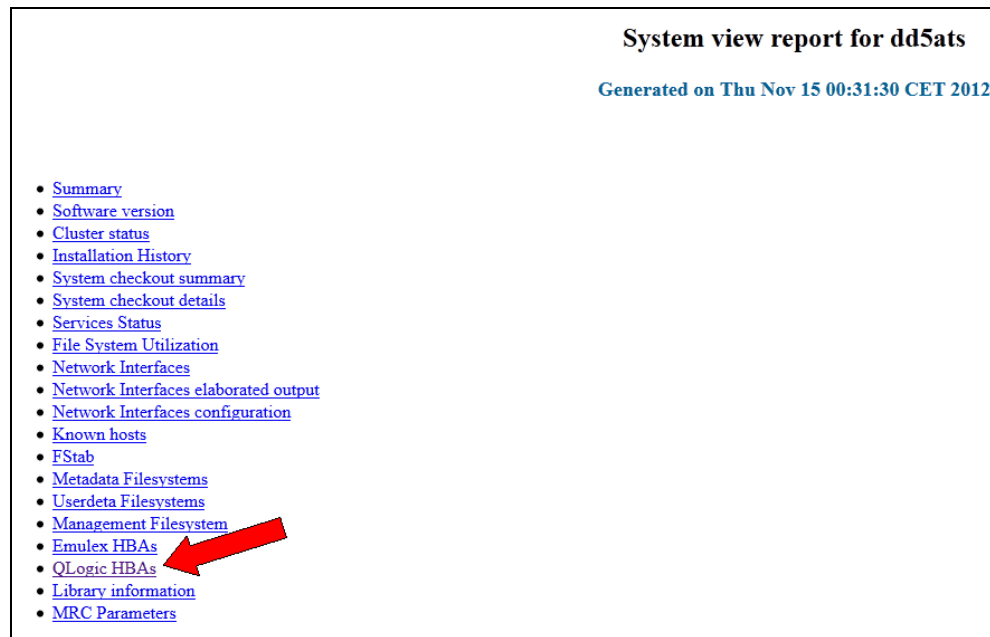


Figure 9-9 Browser showing the `systemview.html` file

Note: The QLogic HBA example shown in Figure 9-9 and Figure 9-10 applies only to DD4 and DD5 servers. The latest DD6 server uses Emulex adapters for both the front-end HBAs that connect from ProtecTIER to the backup server, and for the back-end HBAs that connect from ProtecTIER to the attached disk storage subsystem that acts as a repository.

3. The WWPN information of the QLogic back-end host bus adapter (HBA) is displayed, (Figure 9-10). Be sure to use the port names as marked in the figure. Do not use the node names.

- QLogic HBAs					
PCI ID	State	Speed	Port Name	Node Name	
0000:13:00.0	Online	8 Gbit	0x21000024ff3ae4d8	0x20000024ff3ae4d8	
0000:13:00.1	Online	unknown	0x21000024ff3ae4d9	0x20000024ff3ae4d9	
0000:18:00.0	Online	8 Gbit	0x21000024ff3ae4ba	0x20000024ff3ae4ba	
0000:18:00.1	Online	unknown	0x21000024ff3ae4bb	0x20000024ff3ae4bb	

Figure 9-10 QLogic HBA section of the `systemview.html` file

To create the host connection on your Storwize V3700/V5000 server, complete these steps:

1. Hover the cursor over the host type icon. The Hosts menu opens (Figure 9-11). Click **Hosts**.

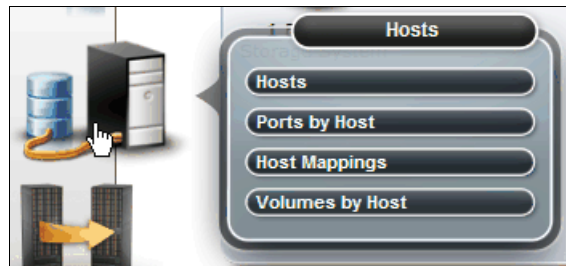


Figure 9-11 Hosts menu

2. In the next window, click **Add Host** to open the Add Host window (Figure 9-12).

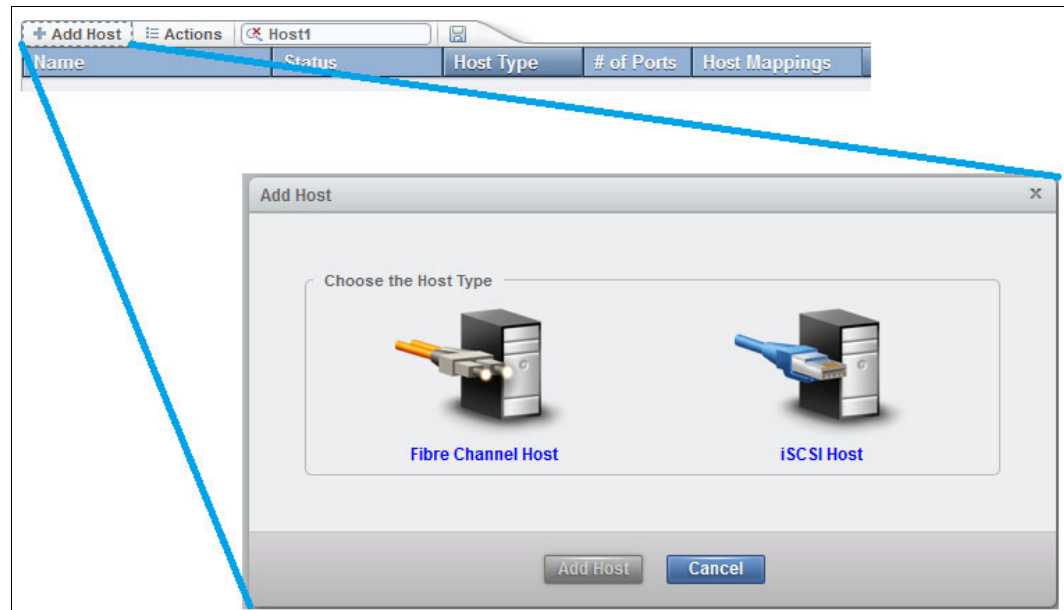
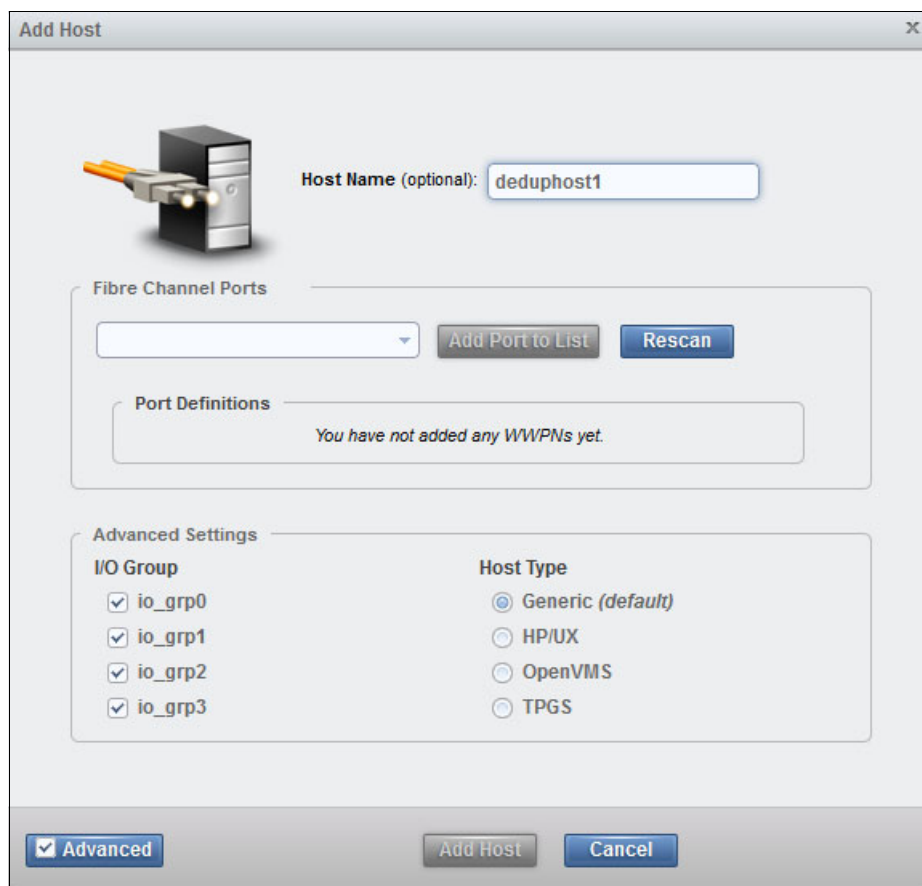


Figure 9-12 Add Host window

3. Click **Fibre Channel Host** and click **Add Host**.

4. In the next window (Figure 9-13), complete these steps:
 - a. Provide the name of your host. The host name is a name that is already specified.
 - b. Add your WWPNs to the Fibre Channel Ports selection box.
 - c. Keep the **Generic (default)** radio button selected as the Host Type.
 - d. Click **Add Host** to return to the main menu.



The image shows a software window titled "Add Host". At the top left is an icon of a server with a fibre optic cable. To its right is a text field labeled "Host Name (optional):" containing the text "deduphost1". Below this is a section titled "Fibre Channel Ports" containing a dropdown menu, an "Add Port to List" button, and a "Rescan" button. Underneath is a "Port Definitions" section with the message "You have not added any WWPNs yet." Below that is an "Advanced Settings" section. It has two columns: "I/O Group" with four checked checkboxes labeled "io_grp0", "io_grp1", "io_grp2", and "io_grp3"; and "Host Type" with four radio buttons, the first of which is selected and labeled "Generic (default)". The other radio buttons are labeled "HP/UX", "OpenVMS", and "TPGS". At the bottom of the window are three buttons: "Advanced" (with a checked checkbox), "Add Host", and "Cancel".

Figure 9-13 Expanded Add Host window

Host groups and clusters: The Storwize V3700 and V5000 storage subsystems do not use the concept of host group or cluster, which is used to map a volume to more than one host at the same time.

If you have a ProtecTIER dual-node server, you can one of two options:

- ▶ Create one host for each ProtecTIER node and, when you perform the volume mapping, use the same LUN ID for both nodes.
- ▶ Create a single host, but add the ports from both ProtecTIER nodes. This method is simpler, but if there are issues with one FC port, it is more difficult to identify which node the FC port belongs to.

9.2.6 Mapping volumes to a host

To make the volumes available to the ProtecTIER node, you must map the volumes to the ProtecTIER host by completing the following steps:

1. Hover the cursor over the Pools icon to open the Pools menu (Figure 9-14). Click **Volumes by Pool**.

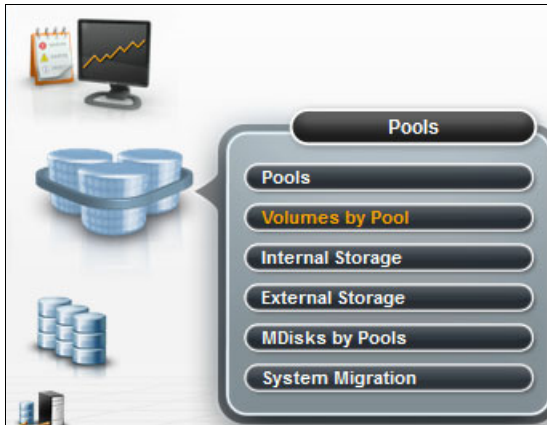


Figure 9-14 Pools menu

2. The list of volumes separated by pool opens (Figure 9-15). Mark the volumes, right-click, and click **Map to Host**.

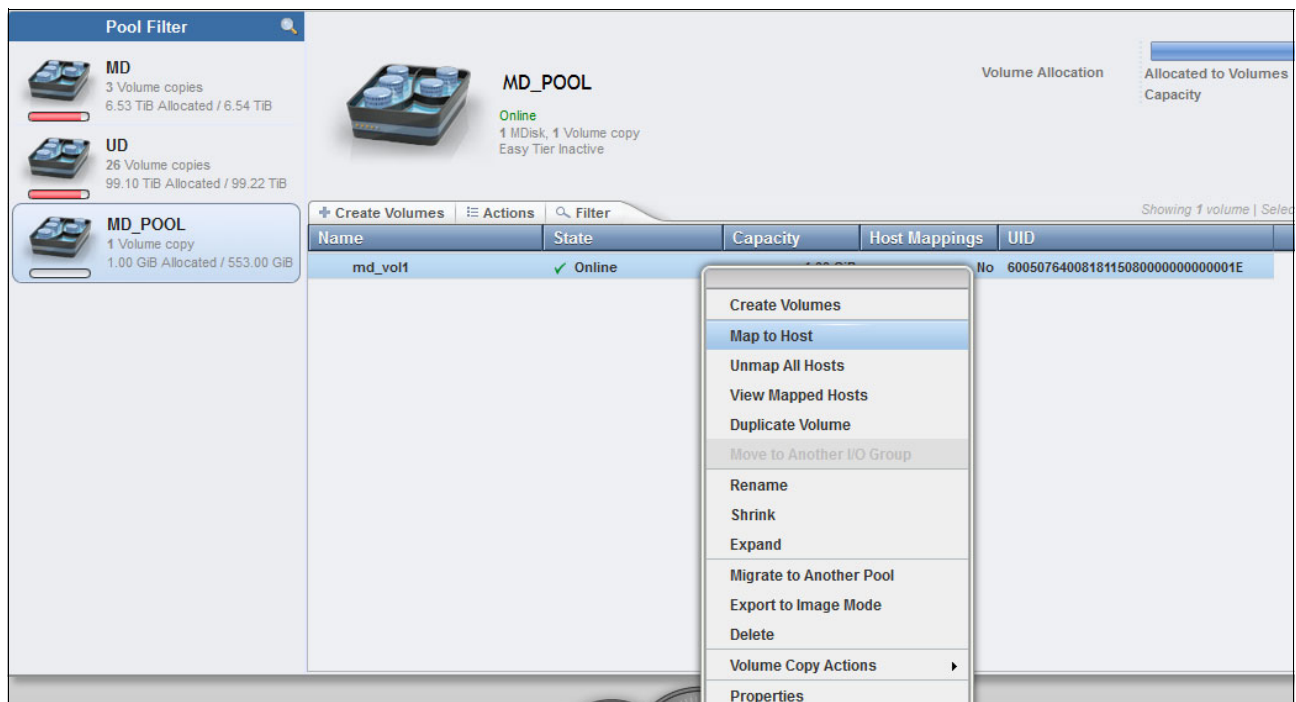


Figure 9-15 Volumes by pool

3. The **Modify Host Mappings** window opens (Figure 9-16). Click the **Host** drop-down menu and select the name of the host to which you want to map volumes. If the volumes are not highlighted in yellow, select volumes from the Unmapped Volumes pane and click the double arrow (>>) to move the volumes to the Volumes Mapped to the Host pane. Click **Map Volumes** or **Apply**.

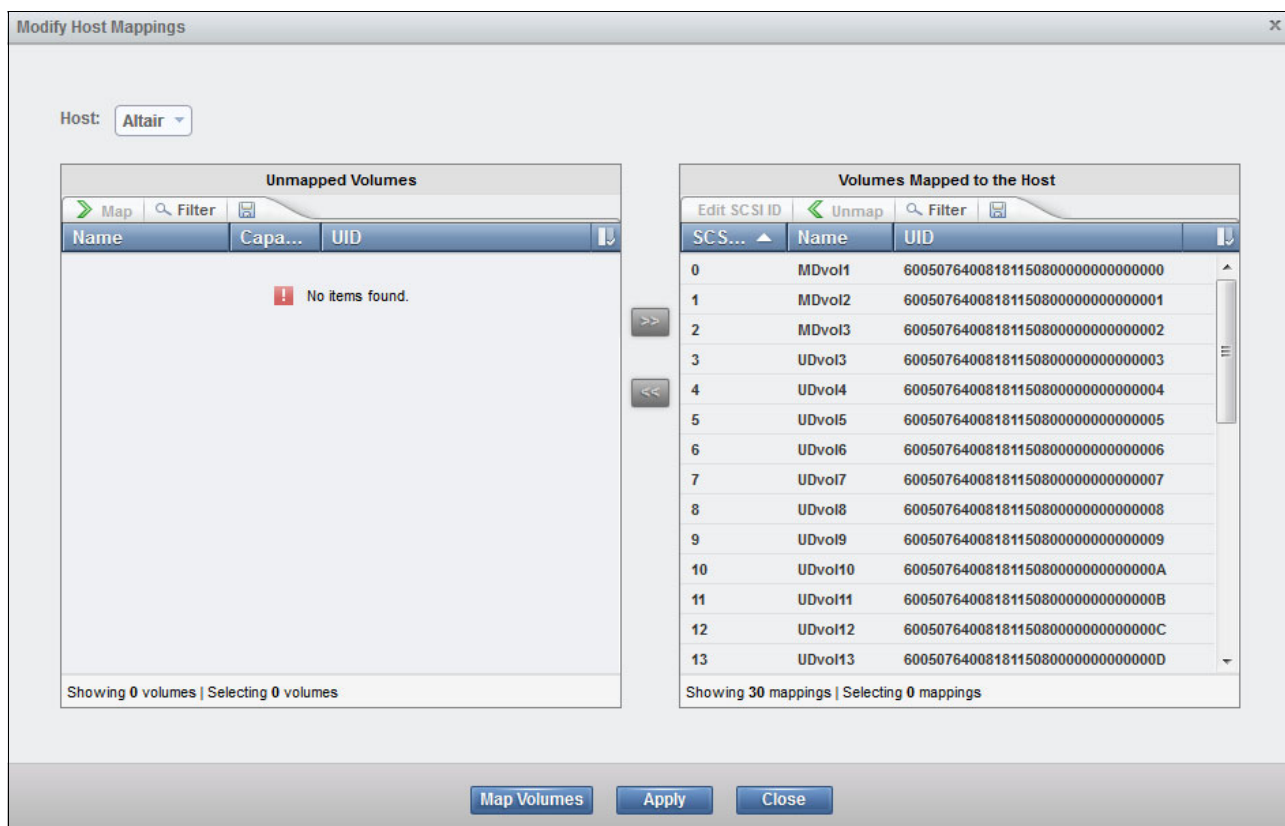


Figure 9-16 Modify mapping

Dual-node server

If you have a ProtecTIER dual-node server, after the repository is built on the second node, you must select the other node and proceed with the volume mapping again. You might receive a warning that states that the specific volume is already mapped (Figure 9-17). Verify that the mapping is correct and continue.

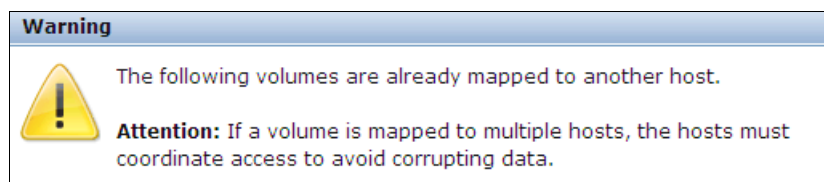


Figure 9-17 Volume mapping warning

9.2.7 Creating file systems and building the ProtecTIER repository

To create the file systems using ProtecTIER V3.4, complete the following steps:

1. Verify that the ProtecTIER node recognizes the volumes that are presented by the SAN Volume Controller and Storwize V3700 /V5000 server. The **multipath -ll** command shows the LUNs and the paths that are connected to the storage subsystem. If the LUNs do not appear to be mapped to the storage, you must run a rescan on the FC adapters or reboot the node.

Note: You can scan the FC adapters by running the following command:

```
echo "- - -" > /sys/class/scsi_host/<host??>/scan
```

In this command, <host??> should be replaced by each FC adapter port.

You can also filter to see only part of the data, as shown in Example 9-8. The first command filters only the volumes of type 2145, which is the Storwize V3700/V5000 type. The second command shows the details of one of the devices.

Example 9-8 multipath command output

```
[root@lbsdudup1a ~]# multipath -ll | grep 2145
mpath9 (360050768028580f8c800000000000e1) dm-8 IBM,2145
mpath8 (360050768028580f8c800000000000e0) dm-7 IBM,2145
mpath14 (360050768028580f8c800000000000e7) dm-27 IBM,2145
mpath7 (360050768028580f8c800000000000df) dm-6 IBM,2145
mpath6 (360050768028580f8c800000000000de) dm-5 IBM,2145
mpath12 (360050768028580f8c800000000000e4) dm-11 IBM,2145
mpath11 (360050768028580f8c800000000000e3) dm-10 IBM,2145
mpath10 (360050768028580f8c800000000000e2) dm-9 IBM,2145

[root@lbsdudup1a ~]# multipath -ll | grep -A11 mpath9
mpath9 (360050768028580f8c800000000000e1) dm-4 IBM,2145
[size=1.2T][features=1 queue_if_no_path][hw_handler=0][rw]
\_ round-robin 0 [prio=200][active]
  \_ 2:0:1:3 sdak 66:64 [active][ready]
  \_ 3:0:2:3 sda1 66:80 [active][ready]
  \_ 1:0:1:3 sdam 66:96 [active][ready]
  \_ 4:0:1:3 sdb5 68:96 [active][ready]
\_ round-robin 0 [prio=40][enabled]
  \_ 4:0:0:3 sdbk 67:224 [active][ready]
  \_ 1:0:0:3 sde 8:64 [active][ready]
  \_ 3:0:1:3 sdk 8:160 [active][ready]
  \_ 2:0:0:3 sdq 65:0 [active][ready]
```

2. In the ProtecTIER Service Menu panel (Example 9-9), select option **1** (ProtecTIER Configuration), and then select option **7** (File Systems Management).

Example 9-9 ProtecTIER configuration menu

```
+-----+
| ProtecTIER Service Menu running on vela |
+-----+
| 1) ProtecTIER Configuration (...)      |
| 2) Manage ProtecTIER services (...)   |
| 3) Health Monitoring (...)            |
| 4) Problem Alerting (...)             |
| 5) Version Information (...)          |
| 6) Generate a service report          |
| 7) Generate a system view            |
| 8) Update ProtecTIER code             |
| 9) ProtecTIER Analysis (...)          |
|                                       |
| E) Exit                              |
+-----+
```

Your choice? 1

```
+-----+
| ProtecTIER Service Menu running on vela |
| ProtecTIER Configuration (...)          |
+-----+
| 1) Configure ProtecTIER node           |
| 2) Recover Configuration for a replaced server |
| 3) Configure machine serial number for a replaced server |
| 4) Configure RAS                      |
| 5) Update Time, Date, Timezone and Timeserver(s) |
| 6) Scan storage interconnections       |
| 7) File Systems Management (...)       |
| 8) Configure replication (...)         |
| 9) IP Network configuration (...)       |
| 10) Update Firmware (...)             |
| 11) Update the System's name           |
| 12) Validate configuration             |
| 13) Single Node - code upgrade (For Support Use ONLY) |
| 14) OS images Management (...)         |
| 15) Replace SAS Drive                 |
|                                       |
| B) Back                              |
| E) Exit                              |
+-----+
```

Your choice?

Your choice? 7

- The File Systems Management menu opens. You have options to display information about the devices or file systems configuration. You have choices to configure the file systems on all available devices or for a single device. Example 9-10 shows the creation of a file system on a single unused device.

Example 9-10 File System Management menu

```

-----+
| ProtecTIER Service Menu running on lbsdedup1a
| ProtectTIER Configuration (...)
| File Systems Management (...)
|-----+
| 1) Configure file systems on all available devices
| 2) Create file system(s) on a single unused device
| 3) Extend a file system with a new unused device
| 4) Update /etc/fstab
| 5) Display configured devices
| 6) Display unused devices
| 7) Display GFS repository file systems
| 8) Display unused GFS file systems
| 9) Increase capacity completion (applicable for a second cluster node)
|
| B) Back
| E) Exit
|-----+
Your choice? 2
Begin Processing Procedure
Display of the available devices
Device:          Size:          Status
1. mpath0        1024.00M          Unused
2. mpath10       1286144.00M       Unused
3. mpath14       568832.00M        Unused
4. mpath11       1286144.00M       Unused
5. mpath6        1284864.00M       Unused
6. mpath7        1286144.00M       Unused
7. mpath8        1286144.00M       Unused
8. mpath9        1286144.00M       Unused

Please select device: 3
Please enter number of file systems[1-4]:1
Last file system - using the remaining free space
The file system to be created will have 568832.00 MiB size
Please confirm:? (yes|no) yes
Creating physical volume          [ Done ]
Creating volume group             [ Done ]
Creating logical volume           [ Done ]
Creating file system              [ Done ]

End Processing Procedure Successfully

```

- After you create the file system for all the devices that are needed, you can go to the ProtecTIER Manager, add the node to it, and then select the Repository menu to create the repository. For more information, see *IBM System Storage TS7650 ProtecTIER Deduplication Appliance Installation Roadmap Guide*, GA32-0921.

9.2.8 Expanding the repository

When you expand the repository, use the same spindle type and quantity of RAID groups for metadata or user data. For example, if the original two metadata LUNs were built on RAID 4+4 groups, then the added metadata RAID groups must be at least 4+4 to maintain the same level of performance. Using storage from 2+2 or 4+1 RAID groups, for example, for the expansion might degrade performance because of an input/output operations per second (IOPS) bottleneck.

Important: Upon installation, a critical step is that metadata be sized and configured, with future performance factored into the equation. Performance gain greatly depends on the initial metadata layout. Without the initial sizing done correctly, increasing the performance of ProtecTIER will require storage migration services. If migration services are required, and to avoid any issues, the migration should be performed by IBM experts or certified IBM Business Partners. You can increase the ProtecTIER user data capacity anytime by adding new metadata file systems or increasing existing metadata file systems if necessary.

The total number of volumes for both the ProtecTIER repository for metadata (MD) and user data (UD) should not exceed 170. To comply with the 1 PB ProtecTIER repository size limit, each volume size should not exceed 6 terabytes (TB).

Starting with ProtecTIER V3.2, the 8 TB restriction is removed. When the individual disk size is large (for example, 2 or 3 TB), use RAID 6 with 6+2 disk members.

9.3 IBM System Storage SAN Volume Controller, IBM Storwize V7000, and V7000 Unified Storage

Storage virtualization technology complements a virtual server environment. The IBM System Storage SAN Volume Controller (SAN Volume Controller), IBM Storwize V7000, and IBM Storwize V7000 Unified are IBM storage virtualization products that address high capacity, high throughput, and NAS connectivity needs with enhanced clustering capabilities and more processing power.

Note: The primary interface methods that are used to work with Storwize V7000 Unified are NAS protocols that enable file level I/O. ProtecTIER requires block level I/O for its potential back-end storage subsystems. Storwize V7000 Unified offers both file level I/O and block level I/O capabilities.

9.3.1 Storage virtualization introduction

Storage virtualization brings intelligence to the storage environment by implementing a virtualization layer between storage back-end disks and hosts. System administrators can view and access a common pool of storage on the storage area network (SAN). This functionality helps administrators use storage resources more efficiently and provides a common base for advanced functions, such as copy services, thin-provisioned volumes, and easy tiering.

Using storage virtualization offers several key benefits:

- ▶ Improves and optimizes capacity usage.
- ▶ Enables data migration of virtualized data with minimum application disruption.

- ▶ Facilitates a common platform for data replication, even if the back-end storage does not have this functionality.
- ▶ Provides a centralized point of control for storage provisioning across heterogeneous SAN environments.
- ▶ Reduces license costs for advanced functions in the virtualized storage systems.
- ▶ Increases storage administrator productivity.

The following IBM products provide this technology:

- ▶ SAN Volume Controller
- ▶ Storwize V7000
- ▶ Storwize V7000 Unified
- ▶ Storwize V3700 and Storwize V5000 (For details, see 9.2, “General considerations: V3700, V5000, and V7000” on page 127.)

SAN Volume Controller can virtualize multi-vendor storage disks to perform LUN provisioning in a high-end environment. In addition to performing external storage virtualization, Storwize V7000 has internal disks that constitute its storage capacity for a mid-range environment. Storwize V7000 Unified consolidates block (SAN) virtualization and file workloads (network-attached storage, or NAS) into a single storage system.

Supported vendor storage: A list of supported vendor storage disks for SAN Volume Controller is at the following web page:

<http://www.ibm.com/support/docview.wss?uid=ssg1S1003658>

Also see the IBM System Storage Interoperation Center (SSIC) website:

<http://www.ibm.com/systems/support/storage/ssic/interoperability.wss>

9.3.2 Terminology

This section describes the terminology used in this chapter that relates to storage virtualization. The following list includes the terms:

Storage pool	Also known as managed disk group (MDG). An MDG is a collection of storage capacity (MDisks) that provides the capacity requirements (extents) for volumes that are based on MDisks characteristics, such as disk speed and type.
Managed disk (MDisk)	An MDisk is composed of a Small Computer System Interface (SCSI) logical unit (LU) that a Redundant Array of Independent Disks (RAID) controller provides and that a clustered system manages. The MDisk is not visible to host systems on the SAN.
Volume/(VDisk)	Also known as virtual disk (VDisk). A volume is the representation of a SCSI disk that is presented to a host system.
Extents	A unit of data that manages the mapping of data between managed disks (MDisks) and volumes. By default, its size is 256 MB.

Figure 9-18 summarizes storage virtualization concepts. The disk RAID array from an external storage system (or from internal disks, when you use Storwize V7000) is presented to a SAN Volume Controller or Storwize V7000 as an MDisk. A set of MDisks forms an array, from which extents are taken to create the volumes (LUNs).

The volumes, now in virtualized mode, are presented to the hosts. In this sense, the hosts no longer see the back-end disks directly, and the SAN Volume Controller, as well as the Storwize V7000 family of products, behaves like a controller provisioning LUNs to the hosts.

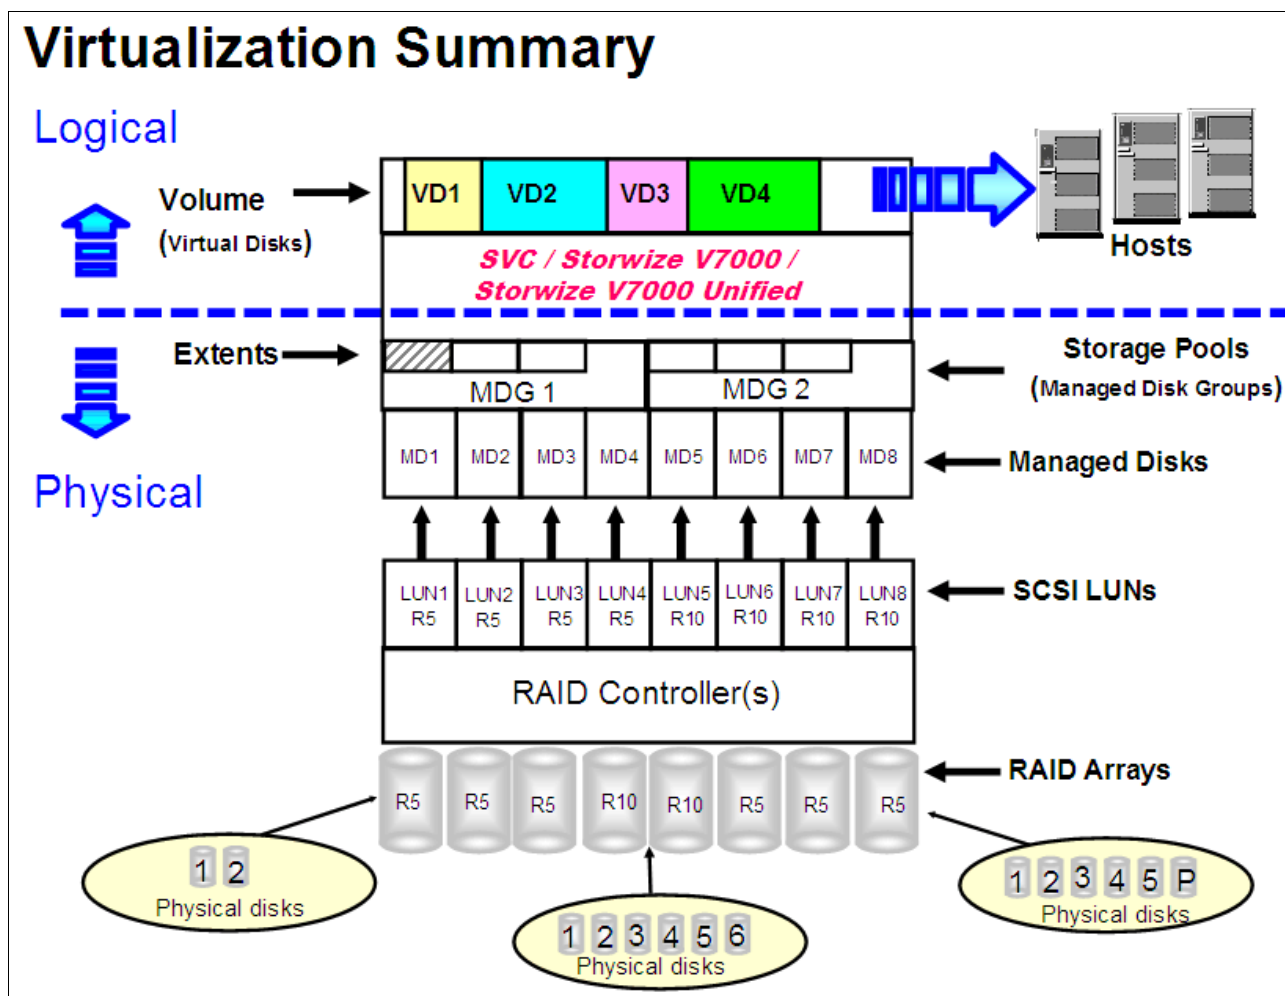


Figure 9-18 Storage virtualization concepts summary

Figure 9-19 shows the three types of virtualization for volumes:

- **Striped (default).** Takes an extent in turn from each managed disk that has free extents (or a subset of managed disks, known as a striped set) in the pool (MDG).
- **Sequential.** Allocates extents from one managed disk to create the volume (VDisk) if enough consecutive free extents are available on the chosen managed disk.
- **Image.** Creates a one-to-one direct mapping between a virtual disk and the managed disk that contains existing data, for example, an existing LUN from an external virtualized storage.

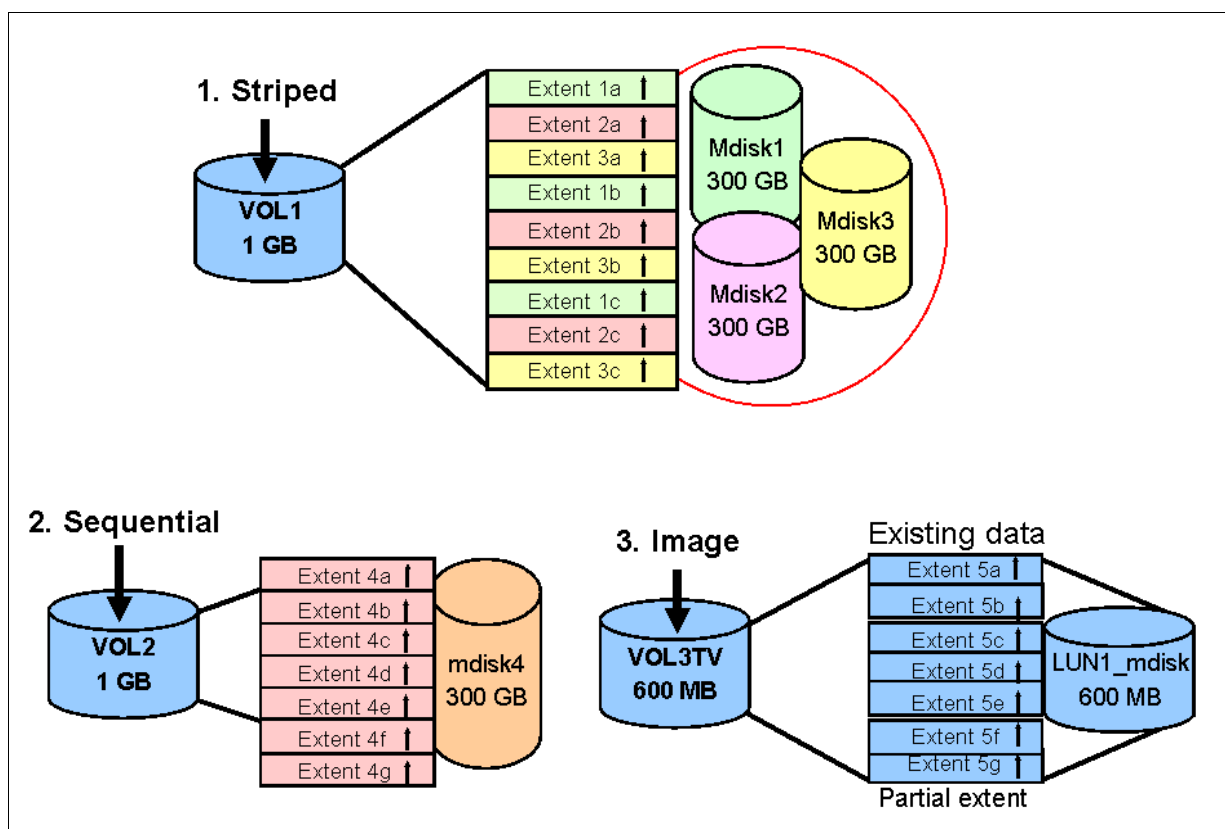


Figure 9-19 Three virtualization types for volumes

More information:

- ▶ SAN Volume Controller in the IBM Knowledge Center:
<http://www.ibm.com/support/knowledgecenter/STPVGU/welcome>
- ▶ SAN Volume Controller support information, configuration limits, and restrictions:
<http://www.ibm.com/support/docview.wss?uid=ssg1S1003658>
- ▶ Storwize V7000 in the IBM Knowledge Center:
<http://www.ibm.com/support/knowledgecenter/ST3FR7/welcome>
- ▶ Storwize V7000 support information, configuration limits, and restrictions:
<http://www.ibm.com/support/docview.wss?uid=ssg1S1003741>
- ▶ Storwize V7000 Unified in the IBM Knowledge Center:
http://www.ibm.com/support/knowledgecenter/ST5Q4U/landing/v7000_unified_welcome.htm
- ▶ Storwize V7000 Unified support information, configuration limits, and restrictions:
<http://www.ibm.com/support/docview.wss?uid=ssg1S1004680>

9.4 General notes

The following general notes apply for SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified storage subsystems:

- ▶ Because of the nature of the ProtecTIER product, storage resources are heavily used, so the SAN Volume Controller or Storwize V7000 storage subsystem should be dedicated solely to ProtecTIER activity.
- ▶ If dedicating the array to ProtecTIER is not possible, use zoning and LUN masking to isolate the IBM System Storage TS7650G ProtecTIER Deduplication Gateway (TS7650G) from other applications.

Important: The TS7650G must not share pools, MDisks, or volumes with other applications.

- ▶ When you use SAN peer-to-peer (P2P) topology to connect the TS7650G to the disk array, create a dedicated zone for the ProtecTIER back-end ports. Do not mix the back-end ports with the front-end ProtecTIER ports or any other SAN devices in the same zone.
- ▶ Ensure that you have redundant connections to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified nodes or node canisters.
- ▶ The Storwize V7000 Cluster is supported and can be used with your ProtecTIER solution.
- ▶ Create *only* two storage pools (also known as managed disk groups), where one is used for metadata and the other is used for user data. Creating more pools causes the storage system cache to be segmented among all pools. By having only two pools, you optimize the storage subsystem cache allocation.

- Use the *sequential* virtualization type for each MDisk rather than striped virtualization. Because of the nature of ProtecTIER file system architecture and workload, sequential virtualization shows better performance than striped.

When you expand the ProtecTIER repository, if you use sequential virtualization, the new sequential volumes maintain the same performance characteristics as the existing volumes. If you use the striped virtualization, you must manually check the extents distribution among the MDisk to ensure that they are balanced.

- Do not use thin-provisioned volumes. All of the storage that is presented to the ProtecTIER system is run through a “padding” process as part of the repository creation. This padding process immediately fills the thinly provisioned LUNS to 99% capacity.
- The number of paths to each volume *must not exceed eight paths* between the host and the SAN Volume Controller, Storwize V7000, or Storwize V7000 Unified environments.
- Each Storwize V7000/SAN Volume Controller volume has a preferred Storwize V7000 controller node to be bound to. Balance the volumes between the SAN Volume Controller, Storwize V7000, or Storwize V7000 Unified controller nodes. The ProtecTIER server has connections and SAN zoning for both nodes. By default, the allocation should equally divide all volumes among the controller nodes; do not interfere with this behavior.
- ProtecTIER is a *random-read* application. A range of 80 - 90% of I/O on a typical TS7650G environment is random read at block size of 64 kilobytes (KB). Implement suitable performance optimizations and tuning as suggested by the disk vendor for this I/O profile.
- Using IBM Easy Tier, set to “ON,” on the metadata or user data pool provides no benefit. The workload pattern of ProtecTIER does not produce enough hot extents to justify the usage of solid-state drives (SSDs).

Important: Always use RAID 6 in combination with Serial Advanced Technology Attachment (SATA) or nearline serial-attached SCSI (NL-SAS) drives or when no Native Replication is set up to maintain a copy of the data in a DR site.

LUN management: Starting with ProtecTIER Version 3.2, the management of LUNs greater than 8 TB is improved. When ProtecTIER V3.2 and later works with LUNs greater than 8 TB, it splits them in to logical volumes of smaller size. Therefore, you can work with LUNs greater than 8 TB, but there is no benefit in performance in completing this action.

9.5 Firmware level

The array firmware level must be equal to or greater than the firmware version that is listed in the IBM SSIC and ProtecTIER ISV Support Matrix. Refer Appendix B, “ProtecTIER compatibility” on page 457 for more details.

9.6 Fibre Channel connection topology

SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified all support redundant connections to clustered ProtecTIER nodes. To ensure full protection against the loss of any Fibre Channel paths from the ProtecTIER nodes to the Storwize V7000, use redundant host connections.

To meet the business requirements for high availability (HA), use SAN design practices to build a dual fabric network, two independent fabrics, or SANs. To increase the total number of switch ports in a fabric, switches can be connected together with one or more interswitch links (ISLs). Multiple ISLs between two switches do not increase the total number of paths to the volumes.

Direct attachment: ProtecTIER supports direct attachment to Storwize V7000. Follow Storwize V7000 preferred practices for direct connect attach. However, given HA preferred practices, a suggestion is to use a Fibre Channel (FC) switch with deployments of two or more arrays.

Connect each host to the appropriate single-ported host channels on the SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified controllers, also known as nodes or node canisters. These configurations have host and drive path failover protection. To ensure redundancy and stability, every ProtecTIER node must be configured to obtain two paths to each controller.

One port of each ProtecTIER server host adapter should be connected to one SAN fabric. For a volume (VDisk or LUN), which is owned by an I/O group, the number of paths from the SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified nodes to a host must not exceed eight. Use SAN zoning to limit the number of paths among ports across the SAN. This setting reduces the number of instances that the same volume (VDisk or LUN) is reported to a host operating system.

No more than four HBA ports of the ProtecTIER servers in total should be connected to the SAN. Each port should be zoned to a different set of the SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified ports to maximize performance and redundancy.

Figure 9-20 a sample redundant SAN fabric FC configuration with a ProtecTIER *single-node* configuration. Figure 9-21 on page 153 shows a redundant SAN fabric FC configuration with a ProtecTIER *dual-node* configuration.




Figure 9-20 Example redundant SAN fabric Fibre Channel configuration with a ProtecTIER single-node configuration

Figure 9-21 shows that each HBA port is zoned with two ports from each SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified node, providing a total of eight paths.




Figure 9-21 Example redundant SAN fabric Fibre Channel configuration with a ProtecTIER dual-node configuration

Multipath settings: You do not need to manually configure the multipath-specific settings (by running the **multipath** command). Configuration is done automatically by an installation script of the ProtecTIER product called **autorun**.

Example 9-11 shows the output after running **multipath -ll**, depending on the physical cabling and the SAN zoning configuration.

Example 9-11 Output from the **multipath -ll** command

```
mpath2 (360050768028109ed8800000000000009) dm-2 IBM,2145
[size=3.8T][features=1 queue_if_no_path][hwhandler=0][rw]
\_ round-robin 0 [prio=200][active]
\_ 6:0:1:8 sddb 70:144 [active][ready]
\_ 5:0:1:8 sddc 70:160 [active][ready]
\_ 4:0:1:8 sddd 70:176 [active][ready]
\_ 3:0:1:8 sdde 70:192 [active][ready]
\_ round-robin 0 [prio=40][enabled]
\_ 6:0:0:8 sdah 66:16 [active][ready]
\_ 3:0:0:8 sdai 66:32 [active][ready]
\_ 5:0:0:8 sdaj 66:48 [active][ready]
\_ 4:0:0:8 sdak 66:64 [active][ready]
```

The following line shows the format of the first four numbers in each line of Example 9-11 on page 153:

HOST (=ProtectTIER ports) : CHANNEL : ID (=different disk subsystem ports per Host port) : LUN (=Volume)

The following list includes the details of Example 9-11 on page 153 for SCSI multipath information for each path:

- ▶ The mpath2 device alias represents the LUN (volume) with ID 8 (unique ID in brackets shows a combination of the system WWN and LUN ID).
- ▶ Eight paths are available to reach the mpath2 device.
- ▶ All four ProtecTIER back-end ports (3, 4, 5, and 6) can reach this LUN ID 8.
- ▶ Each ProtecTIER port can reach two separate disk subsystem ports (for example, ProtecTIER port 6 sees disk subsystem ports 0 and 1).
- ▶ This is an active-active capable disk subsystem (there is no ghost status, such as the IBM DS4000® active-passive system can show).
- ▶ Only one PT to disk subsystem port path group is actively used (active) at one time; the other one is available for failover scenarios (enabled).
- ▶ I/O is load balanced among the [active] path group (round robin).

9.7 User data and metadata pool: General suggestions

This section provides general suggestions for setting up metadata and user data pools for optimum performance in your ProtecTIER environment. It also provides general guidelines for expanding the ProtecTIER repository.

9.7.1 Metadata pool

Consider the following items regarding the metadata pool:

- ▶ Use balanced RAID 10 groups for metadata MDisks (use a layout according to your planning requirements) with at least 4+4 members. The suggested number of metadata RAID groups (MDisks) is determined by the Capacity Planning tool during the pre-sales process. This number can be 2 - 10, based on repository size, factoring ratio, and performance needs.

Terminology:

- ▶ **Balanced RAID 10:** A RAID group creation in Storwize V7000 and Storwize V7000 Unified that uses drives to form the array from different enclosures. This setup protects against drive and enclosure failures, considering the SAS chain and the disk spare protection.
 - ▶ **SAS chain:** A cabling scheme for a string of expansion enclosures that provides redundant access to the drives inside the enclosures. This setup is achieved by having both node canisters in the control enclosure in the Storwize V7000.
- ▶ Do not share the metadata pool with the user data pool.
 - ▶ Create a single pool for metadata that contains all the metadata MDisks.

- ▶ Specify an extent size default of 256 MB for each metadata volume created.
- ▶ For each metadata MDisk, create a single metadata (MD) volume to accommodate the full size of each metadata MDisk, by using the *sequential* virtualization type.

Quorum volume: The 1 GB MD quorum volume can be created on any metadata MDisk.

9.7.2 User data pool

Consider the following items regarding user data pools:

- ▶ Use RAID for data protection and performance. The RAID type depends on the disk category, but is usually RAID 5 or RAID 6.
- ▶ Do not share the user data pool with the metadata pool.
- ▶ Create a single pool for user data that contains all the user data MDisk.
- ▶ Specify an extent size default of 256 MB for each user data volume that is created.
- ▶ For each user data MDisk, create a single user data (UD) VDisk/volume to accommodate the full size of each user data MDisk by using the *sequential* virtualization type.
- ▶ The size of UD volumes must be consistent. All UD volumes should be the same size.
- ▶ Comply with the suggestions that are related to the number of ProtecTIER file systems, as described in 7.3, “Dependencies from a ProtecTIER view” on page 103.

9.8 Configuration steps

The user interface and the commands of the Storwize V7000 are the same as the V3700 and V5000; therefore, see 9.2.1, “Configuration steps: ProtecTIER repository” on page 127 to complete the configuration in the V7000.



IBM XIV Storage System

This chapter addresses specific considerations for using the IBM XIV® Storage System as storage for IBM ProtecTIER servers.

This chapter describes the following topics:

- ▶ XIV Storage System hardware overview
- ▶ Fibre Channel switch cabling and zoning and zoning configuration for maximum performance with an XIV Storage System
- ▶ Configuring XIV Storage System for ProtecTIER server

10.1 XIV Storage System hardware

Figure 10-1 shows an example of the IBM XIV Storage System hardware and supported modules.




Figure 10-1 XIV Storage System hardware

XIV Storage System supports configurations of 6, 9, 10, 11, 12, 13, 14, or 15 modules (Table 10-1):

- ▶ Modules 1 - 3 and 10 - 15 are disks only and are called data modules.
- ▶ Modules 4 - 9 have disks and host interfaces and are called interface modules.

Table 10-1 Configurations of modules

Number of modules	6	9	10	11	12	13	14	15
Interface Module 9 state	Empty	Disabled	Disabled	Enabled	Enabled	Enabled	Enabled	Enabled
Interface Module 8 state	Empty	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
Interface Module 7 state	Empty	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
Interface Module 6 state	Disabled	Disabled	Disabled	Disabled	Enabled	Enabled	Enabled	Enabled
Interface Module 5 state	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled

Number of modules	6	9	10	11	12	13	14	15
Interface Module 4 state	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
FC ports	8	16	16	20	20	24	24	24
Usable capacity (1 / 2 / 3 / 4 / 6 TB)	28 TB 55 TB 84 TB 112 TB 169 TB	44 TB 88 TB 132 TB 177 TB 267 TB	51 TB 102 TB 154 TB 207 TB 311 TB	56 TB 111 TB 168 TB 225 TB 338 TB	63 TB 125 TB 190 TB 254 TB 382 TB	67TB 134 TB 203 TB 272 TB 409 TB	75 TB 149 TB 225 TB 301 TB 453 TB	81 TB 161 TB 243 TB 325 TB 489 TB

10.2 Fibre Channel switch cabling and zoning

For maximum performance with an XIV Storage System, connect all available XIV Storage System Interface Modules and use all of the back-end ProtecTier ports. For redundancy, connect Fibre Channel cables from the ProtecTIER server to two Fibre Channel switched fabrics.

If a single XIV Storage System is being connected, each Fibre Channel switched fabric must have six available ports for Fibre Channel cable attachment to the XIV Storage System. Generally, two connections are used for each interface module in XIV Storage System. Typically, XIV Storage System Interface Module port 1 is used for Fibre Channel switch 1, and port 3 is used for switch 2 (Figure 10-2).




Figure 10-2 Cable diagram for connecting a TS7650G to XIV Storage System

When you use a partially configured XIV Storage System rack, see Figure 10-2 to locate the available FC ports.

10.2.1 Zoning configuration

For each ProtecTIER disk attachment port, multiple XIV Storage System host ports are configured into separate isolated zone pairing in a 1:1 manner:

- ▶ All XIV Storage System Interface Modules on port 1 are zoned to the ProtecTIER host bus adapter (HBA) in slot 6, port 1 and HBA in slot 7, port 1.
- ▶ All XIV Storage System Interface Modules in port 3 are zoned to the ProtecTIER HBA in slot 6, port 2 and HBA in slot 7, port 2.

Information: The suggested way is to connect port 1 and port 3 of the XIV Storage System Interface Modules because they are predefined for host I/O. Other ports might be predefined for XIV Storage System replication. Also, ports 1 and 3 are distributed across the XIV internal dual-port adapters. So, using ports 1 and 3 minimizes administrative effect and ensures protection against an adapter failure of XIV I/O modules.

Each Interface Module in the XIV Storage System has a connection with both ProtecTIER HBAs. A typical ProtecTIER configuration uses 1:1 zoning (one initiator and one target in each zone) to create zones. These zones connect a single ProtecTIER server with a 15 module XIV Storage System with all six Interface Modules (Example 10-1).

Example 10-1 Zoning example for an XIV Storage System attachment

Switch 1:

Zone 01: PT_S6P1, XIV_Module4Port1
Zone 02: PT_S6P1, XIV_Module6Port1
Zone 03: PT_S6P1, XIV_Module8Port1
Zone 04: PT_S7P1, XIV_Module5Port1
Zone 05: PT_S7P1, XIV_Module7Port1
Zone 06: PT_S7P1, XIV_Module9Port1

Switch 02:

Zone 01: PT_S6P2, XIV_Module4Port3
Zone 02: PT_S6P2, XIV_Module6Port3
Zone 03: PT_S6P2, XIV_Module8Port3
Zone 04: PT_S7P2, XIV_Module5Port3
Zone 05: PT_S7P2, XIV_Module7Port3
Zone 06: PT_S7P2, XIV_Module9Port3

This example has the following characteristics:

- ▶ Each ProtecTIER back-end host bus adapter (HBA) port sees three XIV Storage System Interface Modules.
- ▶ Each XIV Storage System Interface Module is connected redundantly to two separate ProtecTIER back-end HBA ports.
- ▶ There are 12 paths (4 x 3) to one volume from a single ProtecTIER server.

10.3 Configuring XIV Storage System for ProtecTIER server

An IBM Service Support Representative (SSR) uses the ProtecTIER Capacity Planning tool to size the ProtecTIER repository metadata and user data. Capacity planning varies, because it depends heavily on your type of data and expected deduplication ratio.

The planning tool output includes the detailed information about all volume sizes and capacities for your specific ProtecTIER installation. If you do not have this information, contact your IBM SSR to get it.

Tip: When you calculate the required size of the repository for ProtecTIER, always account for three factors:

- ▶ Factoring ratio (can be estimated using the IBM ProtecTIER Performance Calculator)
- ▶ Throughput
- ▶ Size of the repository

The factoring ratio and the size of the repository directly impact the size of the metadata volumes; the bigger these values are, the bigger the required metadata volumes will be, while higher throughput will require a larger number of metadata volumes. These three factors must be considered for the initial installation, always accounting for the future growth of the business.

You must configure the XIV Storage System before the ProtecTIER system is installed by an SSR. To configure the system, complete the following steps:

1. Configure an XIV Storage System for the ProtecTIER system. Set the snapshot space to zero because creating snapshots on XIV Storage System is not supported by a ProtecTIER server.
2. Set up volumes in the XIV Storage System with the ProtecTIER Capacity Planning tool and the Create repository planning wizard output. Starting with ProtecTIER V3.2.0, you can select XIV MD Device 8 +8 for the MD RAID Configuration from the ProtecTIER Manager when you are using the Create repository planning wizard.

The ProtecTIER Capacity Planning tool output gives you the metadata volume size and the size of the 32 user data volumes. Also configure a Quorum volume with a minimum of 1 GB, or 17 GB for XIV Storage System because that is the smallest volume size that can be created, in case the solution needs more ProtecTIER servers in the future.

Important: Your IBM SSR can use the Capacity Magic tool to calculate the size of the XIV volumes accurately to ensure that you get the expected volume sizes.

3. Map the volumes to the ProtecTIER server or ProtecTIER cluster.

Example of configuring an XIV Storage System

If you want to set up a ProtecTIER environment with a 79 TB XIV Storage System and a deduplication factoring ratio of 12, use the following volumes sizes:

- ▶ 2 x 1571 GB volumes for metadata: Make these volumes equal to each other, and nearest to the XIV Storage System allocation size, in this case, 1583 (see Figure 10-4 on page 163).
- ▶ 1 x 17 GB volume for Quorum (see Figure 10-5 on page 163): It must be 17 GB because that is the XIV Storage System minimum size.
- ▶ 32 x *<Remaining Pool Space available>*, which is 75440: Dividing 75440 by 32 means that user data LUNs on the XIV Storage System should be 2357 GB each (see Figure 10-6 on page 164).

Memory totals: When you have an XIV Storage System Gen 3 full rack, you can have up to 243 TB of available space. With that configuration, and also if you have more than one XIV Storage System connected to the ProtecTIER server, you might need more than 32 LUNS. For the best performance, do not exceed the LUN size of 6 TB.

For example, two full racks equal 486 TB. Dividing this number by 6 (as in the 6 TB suggested LUN size), you need roughly 81 LUNs. Create the necessary metadata LUNs of the suggested size and the 6 TB LUNs for user data.

Always upgrade the firmware of your XIV Storage System to the current supported level. For more information, see Appendix B, “ProtecTIER compatibility” on page 457.

For XIV Storage System series, capacity, and connectivity details, see this web page:

<http://www.ibm.com/systems/storage/disk/xiv/specifications.html>

As illustrated in the following examples, the XIV Storage System client GUI makes this calculation easy for you. Enter the number of volumes to create, then drag the slider to the right to fill the entire pool. The GUI automatically calculates the appropriate equivalent amount.

Create a pool size for the capacity that you want to use for the ProtecTIER Deduplication Gateway with the XIV Storage System GUI, as shown in Figure 10-3. Normally, this would be 100% of the available space on the XIV. Create one single pool only; there is no benefit to having multiple pools.

Information: On older firmware levels of XIV, you might be required to have multiple pools to allocate all the space that you want to use with ProtecTIER. Besides a slightly increased administrative effect, there is no drawback to completing this action.

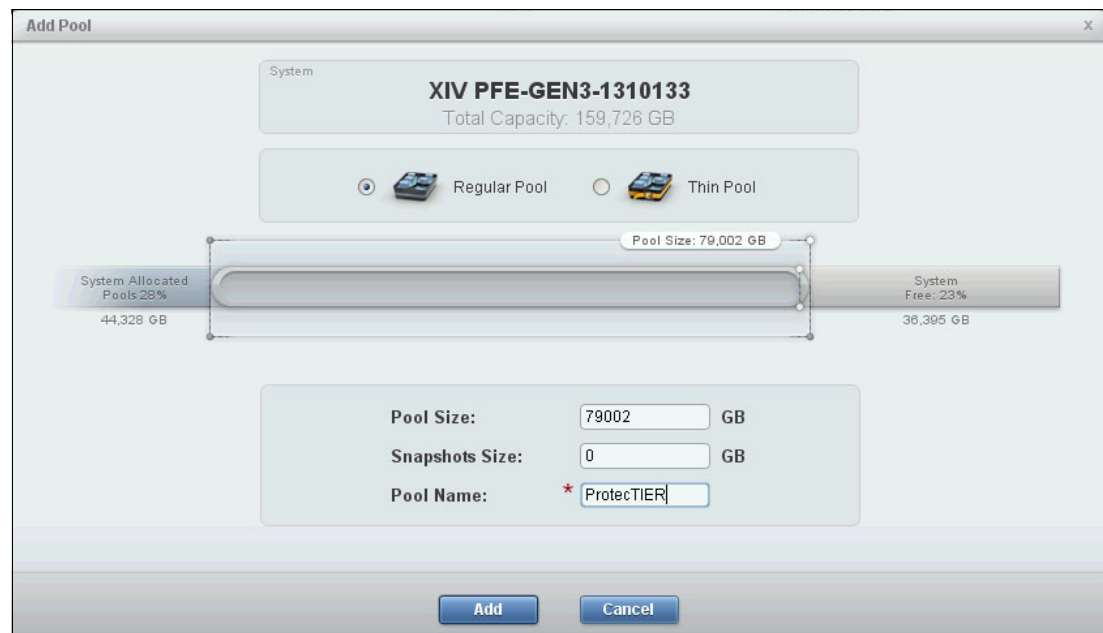


Figure 10-3 Creating a pool

Tip: Use a regular pool and zero the snapshot reserve space. Snapshots and thin provisioning are not supported when XIV Storage System is used with a ProtecTIER server.

Figure 10-4 shows the creation of the metadata volumes.

The screenshot shows the 'Create Volumes' dialog box. At the top, there is a 'Select Pool' dropdown menu set to 'ProtecTIER', with 'Total Size: 79,002 GB' displayed below it. A horizontal slider bar is visible, with a small rectangular selection box positioned on the left side. Below the slider, there are three input fields: 'Number of Volumes' set to '2', 'Volume Size' set to '1583' with a 'GB' unit dropdown, and 'Volume Name' set to 'MetaData' with a red asterisk indicating a required field. To the right of the 'Volume Name' field are two small buttons labeled '1' and '2'. At the bottom of the dialog are 'Create' and 'Cancel' buttons.

Figure 10-4 Creating metadata volumes

Figure 10-5 shows the creation of the quorum volume.

The screenshot shows the 'Create Volumes' dialog box. At the top, there is a 'Select Pool' dropdown menu set to 'ProtecTIER', with 'Total Size: 79,002 GB' displayed below it. A horizontal slider bar is visible, with a small rectangular selection box positioned on the left side. Below the slider, there are three input fields: 'Number of Volumes' set to '1', 'Volume Size' set to '17' with a 'GB' unit dropdown, and 'Volume Name' set to 'Quorum' with a red asterisk indicating a required field. At the bottom left of the dialog, there is a red error message that reads 'Volume size cannot be zero'. At the bottom of the dialog are 'Create' and 'Cancel' buttons.

Figure 10-5 Creating a quorum volume

Figure 10-6 shows the creation of volumes for user data. The arrows show dragging the slider to use all of the pool. This action automatically calculates the appropriate size for all volumes.




Figure 10-6 Creating user data volumes

If you have a ProtecTIER cluster (two ProtecTIER servers in a high availability solution), complete the following steps:

1. Create a cluster group (Figure 10-7).





Figure 10-7 Creating a cluster group

2. Add a host that is defined for each node to that cluster group.

3. Create a cluster definition for the highly available ProtecTIER cluster (Figure 10-8).



The 'Add Cluster' dialog box is shown. It has a title bar with 'Add Cluster' and a close button. The main area contains two fields: 'Name:' with a red asterisk and a text box containing 'ProtecTIER', and 'Type' with a dropdown menu showing 'default'. At the bottom are 'Add' and 'Cancel' buttons.

Figure 10-8 Adding a cluster definition

4. Right-click the cluster and select **Add Host** (Figure 10-9).

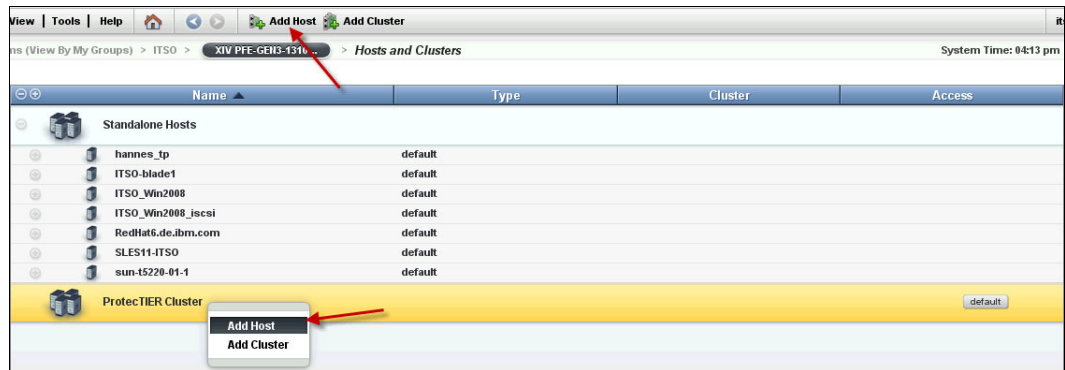
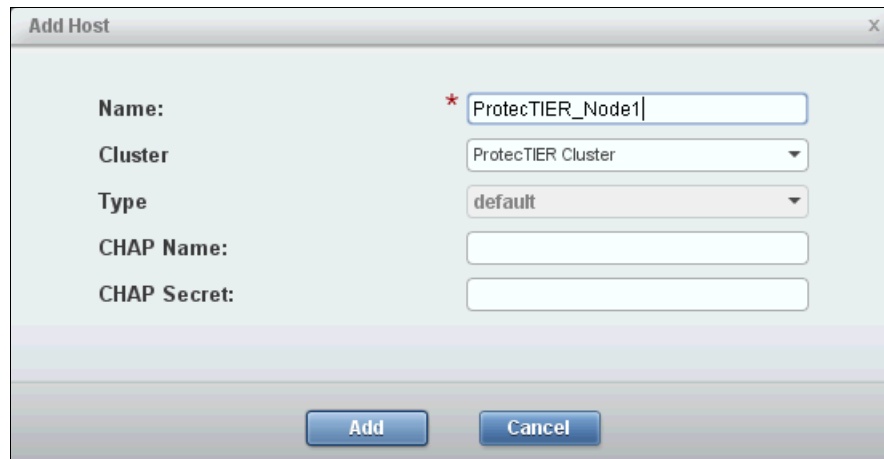


Figure 10-9 Adding a host to the cluster

5. Enter the information for the new ProtecTIER host and click **Add** (Figure 10-10).

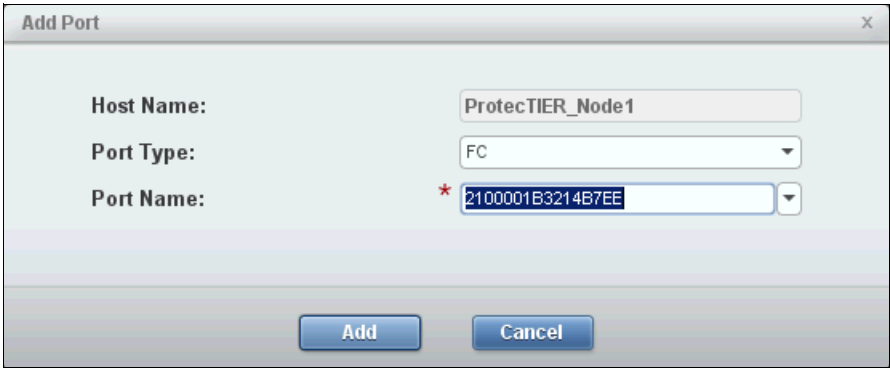


The 'Add Host' dialog box is shown. It has a title bar with 'Add Host' and a close button. The main area contains five fields: 'Name:' with a red asterisk and a text box containing 'ProtecTIER_Node1'; 'Cluster' with a dropdown menu showing 'ProtecTIER Cluster'; 'Type' with a dropdown menu showing 'default'; 'CHAP Name:' with an empty text box; and 'CHAP Secret:' with an empty text box. At the bottom are 'Add' and 'Cancel' buttons.

Figure 10-10 Adding a ProtecTIER host to a cluster

6. Find the worldwide port names (WWPNs) of the ProtecTIER servers. WWPNs can be found in the name server of the Fibre Channel switch. If zoning is in place, they are selectable from the menu. Alternatively, they can also be found in the basic input/output system (BIOS) of the HBA cards and then entered manually in to the XIV Storage System graphical user interface (GUI).

7. Add the WWPNs to the ProtecTIER servers (Figure 10-11).



The 'Add Port' dialog box is shown. It has three fields: 'Host Name' with the value 'ProtecTIER_Node1', 'Port Type' with a dropdown menu showing 'FC', and 'Port Name' with a dropdown menu showing '2100001B3214B7EE'. There is a red asterisk next to the Port Name field. At the bottom are 'Add' and 'Cancel' buttons.

Figure 10-11 Adding the WWPN of ProtecTIER server 1 to the cluster

Figure 10-12 shows the WWPNs that are added to the hosts.

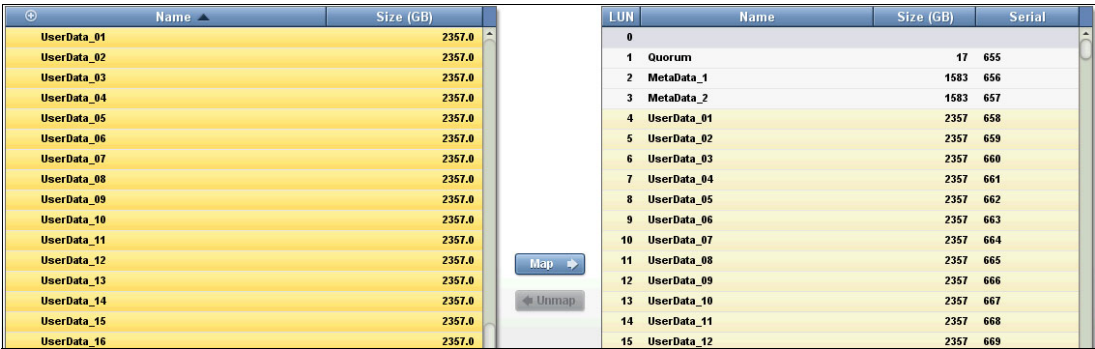


Name	Type	Cluster	Access
ProtectTIER Cluster			
ProtecTIER_Node1	default	ProtectTIER Cluster	default
2100001B321468ED	FC		
2100001B3214B7EE	FC		
ProtecTIER_Node2	default	ProtectTIER Cluster	
2100001B32145CEC	FC		
2100001B321494EE	FC		

Figure 10-12 ProtecTIER WWPNs added to host and cluster definitions

8. Map the volumes to the ProtecTIER cluster. In the XIV Storage System GUI, right-click the cluster name or on the host if you have only one ProtecTIER server, and select **Modify LUN Mapping**. Figure 10-13 shows you what the mapping view looks like.

Tip: If you have only one ProtecTIER server, map the volumes directly to the ProtecTIER server.



Name	Size (GB)
UserData_01	2357.0
UserData_02	2357.0
UserData_03	2357.0
UserData_04	2357.0
UserData_05	2357.0
UserData_06	2357.0
UserData_07	2357.0
UserData_08	2357.0
UserData_09	2357.0
UserData_10	2357.0
UserData_11	2357.0
UserData_12	2357.0
UserData_13	2357.0
UserData_14	2357.0
UserData_15	2357.0
UserData_16	2357.0

Map

Unmap

LUN	Name	Size (GB)	Serial
0			
1	Quorum	17	655
2	MetaData_1	1583	656
3	MetaData_2	1583	657
4	UserData_01	2357	658
5	UserData_02	2357	659
6	UserData_03	2357	660
7	UserData_04	2357	661
8	UserData_05	2357	662
9	UserData_06	2357	663
10	UserData_07	2357	664
11	UserData_08	2357	665
12	UserData_09	2357	666
13	UserData_10	2357	667
14	UserData_11	2357	668
15	UserData_12	2357	669

Figure 10-13 Mapping LUNs to the ProtecTIER cluster



IBM System Storage DS8000

This chapter addresses specific considerations for using the IBM System Storage DS8000 as a storage system for IBM ProtecTIER servers.

This chapter describes the following topics:

- ▶ DS8000 series overview and suggested RAID levels
- ▶ General considerations for planning tools, metadata, user data, firmware levels, replication
- ▶ Rotate extents: Striping and when to use it

Note: The DS8000 series models DS8100, DS8300, DS8700, and DS8800 (types 2107 and 242x, and models 931, 932, 9B2, 941, and 951) have been withdrawn from marketing. The currently available DS8000 products have been replaced by the IBM System Storage DS8870 (type 242x model 961).

These predecessor products are still all supported as back-end storage for ProtecTIER, and will be covered in this chapter.

For a list of supported disk storage systems, see the IBM System Storage Interoperation Center (SSIC) and ProtecTIER ISV Support Matrix. More details are in Appendix B, “ProtecTIER compatibility” on page 457.

11.1 DS8000 series overview

The DS8000 family is a high-performance, high-capacity, and resilient series of disk storage systems. It offers high availability (HA), multiplatform support, and simplified management tools to provide a cost-effective path to an on-demand configuration.

Table 11-1 describes the DS8000 products (DS8300, DS8700, DS8800, DS8870).

Table 11-1 DS8000 hardware comparison

	DS8300	DS8700	DS8800	DS8870
Processor	P5+ 2.2 GHz 4-core	P6 4.7 GHz 2 or 4-core	P6+ 5.0 GHz 2 or 4-core	P7 4.228 GHz 2, 4, 8, 16-core
Processor memory	32 - 256 GB	32 - 384 GB	16 - 384 GB	16 - 1024 GB
Drive count	16 - 1024	16 - 1024	16 - 1536	16 - 1536
Enterprise drive options	FC - 73, 146, 300, 450 GB	FC - 300, 450, 600 GB	SAS2 - 146, 300, 450, 600, 900 GB	SAS2 - 146, 300, 600, 900 GB
Solid-state drive (SSD) options	73, 146 GB	600 GB	300, 400 GB	200, 400, 800 GB 1.6 TB
Nearline drive options	1 TB	2 TB	3 TB	4 TB
Drive enclosure	Megapack		High-density, high-efficiency Gigapack	
Max physical capacity	1024 TB	2048 TB	2304 TB	2304 TB
Power supply	Bulk	Bulk	Bulk	DC-UPS
Rack space for SSD Ultra Drawer	No	No	No	Yes
Redundant Array of Independent Disks (RAID) options	RAID 5, 6, 10	RAID 5, 6, 10	RAID 5, 6, 10	RAID 5, 6, 10
Internal fabric	Remote Input/Output (RIO)-G	Peripheral Component Interconnect Express (PCIe)		
Maximum logical unit number (LUN)/count key data (CKD) volumes	64,000 total	64,000 total	64,000 total	65,280 total
Maximum LUN size	2 TB	16 TB	16 TB	16 TB
Host adapters	IBM ESCON x2 ports 4 Gb FC x4 ports	4 Gb FC x4 ports 8 Gb FC x4 ports	8 Gb FC x4 or x8 ports per adapter	8 Gb FC x4 or x8 ports per adapter 16 Gb FC x4 ports per adapter
Host adapter slots	32	32	16	16
Maximum host adapter ports	128	128	128	128
Drive interface	2 Gbps FC-AL	2 Gbps FC-AL	6 Gbps SAS2	6 Gbps SAS2
Device adapter slots	16	16	16	16
Cabinet design	Top exhaust		Front-to-back	

Table 11-2 compares key features of the DS8000 products.

Table 11-2 DS8000 key feature comparison.

	DS8300	DS8700	DS8800	DS8870
Self-encrypting drives	Optional	Optional	Optional	Standard
Point-in-time copies	IBM FlashCopy®, FlashCopy SE	Same plus Spectrum Control Protect Snapshot, Remote Pair FlashCopy	Same	Same
Smart Drive Rebuild	No	Yes	Yes	Yes
Remote Mirroring	Advanced mirroring	Same plus Global Mirror Multi Session, Open IBM HyperSwap® for AIX	Same	Same
Automated drive tiering	No	Easy Tier Gen 1, 2, 3	Easy Tier Gen 1, 2, 3, 4	Same
Thin provisioning	Yes	Yes	Yes	Yes
Storage pool striping	Yes	Yes	Yes	Yes
I/O Priority Manager	No	Yes	Yes	Yes
Graphical user interface (GUI)	DS Storage Manager	Same plus enhancements	New XIV-like GUI	Same
Dynamic provisioning	Add/Delete	Same plus depopulate rank	Same	Same

For more information about the DS8000 family products, see the following web page:

<http://www.ibm.com/systems/storage/disk/ds8000/overview.html>

11.1.1 Disk drives

This section describes the available drives for the DS8000 products.

Solid-state drives (SSD)

SSDs are the best choice for I/O-intensive workloads. They provide up to 100 times the throughput and 10 times lower response time than 15,000 revolutions per minute (RPM) spinning disks. They also use less power than traditional spinning disks.

SAS and Fibre Channel disk drives

Serial-attached Small Computer System Interface (SAS) enterprise drives rotate at 15,000 or 10,000 RPM. If an application requires high performance data throughput and continuous, intensive I/O operations, enterprise drives are the best price-performance option.

Serial Advanced Technology Attachment (SATA) and NL-SAS

The 4 TB nearline SAS (NL-SAS) drives are both the largest and slowest of the drives available for the DS8000 family. Nearline drives are a cost-efficient storage option for lower intensity storage workloads, and are available since the DS8800. Because of the lower usage and the potential for drive protection throttling, these drives are not the optimal choice for high performance or I/O-intensive applications.

11.1.2 RAID levels

The DS8000 series offers RAID 5, RAID 6, and RAID 10 levels. There are some limitations:

- ▶ RAID 10 for SSD is not standard and is only supported through a request for product quotation (RPQ).
- ▶ SSDs cannot be configured in RAID 6.
- ▶ Nearline disks cannot be configured in RAID 5 and RAID 10.

RAID 5

Normally, RAID 5 is used because it provides good performance for random and sequential workloads and it does not need much more storage for redundancy (one parity drive). The DS8000 series can detect sequential workload. When a complete stripe is in cache for destaging, the DS8000 series switches to a RAID 3-like algorithm. Because a complete stripe must be destaged, the old data and parity do not need to be read.

Instead, the new parity is calculated across the stripe, and the data and parity are destaged to disk. This action provides good sequential performance. A random write causes a cache hit, but the I/O is not complete until a copy of the write data is put in non-volatile storage (NVS). When data is destaged to disk, a write in RAID 5 causes four disk operations, the so-called write penalty:

- ▶ Old data and the old parity information must be read.
- ▶ New parity is calculated in the device adapter.
- ▶ Data and parity are written to disk.

Most of this activity is hidden to the server or host because the I/O is complete when data enters the cache and non-volatile storage (NVS).

RAID 6

RAID 6 is an option that increases data fault tolerance. It allows additional failure, compared to RAID 5, by using a second independent distributed parity scheme (dual parity). RAID 6 provides a read performance that is similar to RAID 5, but has more write penalty than RAID 5 because it must write a second parity stripe.

RAID 6 should be considered in situations where you would consider RAID 5, but there is a demand for increased reliability. RAID 6 is designed for protection during longer rebuild times on larger capacity drives to cope with the risk of having a second drive failure in a rank while the failed drive is being rebuilt. It has the following characteristics:

- ▶ Sequential read of about 99% x RAID 5 rate.
- ▶ Sequential write of about 65% x RAID 5 rate.
- ▶ Random 4 K 70% R/30% W IOPS of about 55% x RAID 5 rate.
- ▶ The performance is degraded with two failing disks.

Note: If enough disks are available and capacity is not an issue at an installation, then using a RAID 6 array for best possible protection of the data is always a better approach.

RAID 10

A workload that is dominated by random writes benefits from RAID 10. In this case, data is striped across several disks and concurrently mirrored to another set of disks. A write causes only two disk operations compared to the four operations of RAID 5. However, you need nearly twice as many disk drives for the same capacity compared to RAID 5.

Therefore, for twice the number of drives (and probably cost), you can perform four times more random writes, so considering the use of RAID 10 for high-performance random-write workloads is worthwhile.

11.2 General considerations

This section describes general considerations for the ProtecTIER Capacity Planning tool and some guidelines for the usage and setup of metadata and user data in your DS8000. For an example of the ProtecTIER Capacity Planning tool, see 7.1, “Overview” on page 100.

11.2.1 Planning tools

An IBM System Client Technical Specialist (CTS) expert uses the ProtecTIER Capacity Planning tool to size the ProtecTIER repository metadata and user data. Capacity planning always differs because it depends heavily on your type of data and the expected deduplication ratio. The planning tool output includes detailed information about all volume sizes and capacities for your specific ProtecTIER installation. If you do not have this information, contact your IBM Sales Representative to get it.

Tip: When calculating the total amount of physical storage required to build the repository for ProtecTIER, always account for three factors:

- ▶ Factoring ratio (estimate by using the IBM ProtecTIER Performance Calculator)
- ▶ Throughput
- ▶ Size of the repository user data

The factoring ratio and the size of the repository directly impact the size of the metadata volumes: the bigger these two values are, the bigger the required metadata volumes size will be requested. The throughput directly impacts the number of metadata values: the higher the desired throughput, a larger number of metadata volumes will be requested.

Consider the three factors for the initial installation, always accounting for the future growth of the business.

11.2.2 Metadata

Consider the following items about metadata:

- ▶ Use the ProtecTIER Capacity Planning tool and the Create repository planning wizard output to determine the metadata requirements for your environment.
- ▶ You must use RAID 10 for metadata. Use high-performance and high-reliability enterprise class disks for metadata RAID 10 arrays.
- ▶ When possible, do not use SATA disks, because RAID 10 is not supported by SATA disks and because ProtecTIER metadata has a heavily random read I/O characteristic. If you require a large physical repository and have only SATA drives available in the storage system, you must use the **rotateextents** feature for all LUNs to ensure the equal

distribution of ProtecTIER workload across all available resources. For more information about the usage of the **rotateextents** feature, see 11.3, “Rotate extents: Striping and when to use it” on page 173.

11.2.3 User data

Consider the following items about user data:

- ▶ ProtecTIER is a random-read application. 80 - 90% of I/O in a typical ProtecTIER environment is random read. Implement suitable performance optimizations and tuning as suggested for this I/O profile.
- ▶ For SATA drives or large capacity disk drives, use RAID 6 with 6 + 2 disk members for increased availability and faster recovery from disk failure.
- ▶ Do not intermix arrays with different disk types in the metadata (MD) and the user data (UD) because smaller disk types hold back the performance of larger disk types and degrade the overall system throughput.
- ▶ For smaller capacity FC or SAS drives, use RAID 5 with at least five disk members per group.
- ▶ Create an even number of LUNs in each pool.

Important: With ProtecTIER and DS8000, create LUNs that are all the same size to avoid performance degradation.

Starting with ProtecTIER Version 3.2, the management of LUNs greater than 8 TB is improved. When ProtecTIER uses LUNs greater than 8 TB, it splits them into logical volumes of smaller size. Therefore, you can work with LUNs greater than 8 TB, but there is no benefit in performance in completing this action.

Always use RAID 6 for SATA or NL-SAS drives for the user data LUNs. With SAS drives, only RAID 6 is supported.

- ▶ Do not use thin provisioning with your ProtecTIER. Thin provisioning technology enables you to assign storage to a LUN on demand. The storage can present to a host a 10 TB LUN, but allocate only 2 TB of physical storage at the beginning. As data is being written to that LUN and the 2 TB are overpassed, the storage will assign more physical storage to that LUN, up until the LUN size is reached.

When the ProtecTIER repository is built, ProtecTIER reserves all of the space in the user data file systems, by writing zeros (padding) in the entire file system. This padding process voids the reason for using thin provisioning.

11.2.4 Firmware levels

Ensure that you are using supported firmware levels. When possible, use the current supported level. For compatibility information, see the IBM SSIC and ProtecTIER ISV Support Matrix. More details are in Appendix B, “ProtecTIER compatibility” on page 457.

11.2.5 Replication

Do not use disk-based replication, because disk-based replication features are not supported by the ProtecTIER product. Rather than using the replication feature of the DS8000, use the

ProtectTIER native replication. For more information about replication, see Part 4, “Replication and disaster recovery” on page 359.

11.3 Rotate extents: Striping and when to use it

This section describes the **rotateexts** feature, and when to use it, or not use it, in your ProtectTIER environment. The rotate extents (**rotateexts**) feature is also referred to as *Storage Pool Striping* (SPS). In addition to the rotate volumes extent allocation method, which remains the default, the rotate extents algorithm is an extra option of the **mkfbvol** command.

The rotate extents algorithm evenly distributes the extents of a single volume across all the ranks in a multirank extent pool. This algorithm provides the maximum granularity that is available on the DS8000 (that is, on the extent level that is equal to 1 GB for fixed-block architecture (FB) volumes), spreading each single volume across multiple ranks, and evenly balancing the workload in an extent pool.

Depending on the type and size of disks that you use in your DS8000 server, and your planned array size to create your ProtectTIER repository, you can consider using **rotateexts**. Because the ProtectTIER product already does a good job at equally distributing the load to the back-end disks, there are some potential scenarios where you should not use **rotateexts** (as described in 11.3.1, “When not to use rotate extents” on page 174).

Important: For ProtectTIER performance, the most critical item is the number of spinning disks in the back end. The spindle count has a direct effect on the ProtectTIER performance. Sharing disk arrays between ProtectTIER and some other workload is *not* recommended. This situation directly impacts your ability to reach your wanted performance. Because you do not share disks between ProtectTIER and other workloads, assigning the full array capacity to the ProtectTIER server is suggested.

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

- ▶ 4+4 RAID 10 or 3+3+2S RAID 10
- ▶ 7+1 RAID 5 or 6+1+S RAID 5
- ▶ 6+2 RAID 6 or 5+2+S RAID 6

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

If you use 3 TB SATA disks to create your arrays, you can have these array dimensions:

- ▶ Creating a 6+2 RAID 6 with a 3 TB disk results in a potential LUN size of 18 TB.
- ▶ Creating a 5+1+S RAID 5 with a 3 TB disk results in a potential LUN size of 15 TB.

In this case, use **rotateexts** to equally distribute the ProtectTIER load to the DS8000 across all available resources. The rotate extents feature helps you create smaller LUNs when ProtectTIER code older than Version 3.2 is installed.

Important: ProtecTIER metadata that is on the RAID 10 arrays has a heavily random write I/O characteristic. ProtecTIER user data that is on RAID 5 or RAID 6 arrays has a heavily random read I/O characteristic. You should use high-performance and high-reliability enterprise-class disk for your metadata RAID 10 arrays.

11.3.1 When not to use rotate extents

Rotate extents (**rotateexts**) is a useful DS8000 feature that can be used to achieve great flexibility and performance with minimal effort. ProtecTIER includes special requirements where using the **rotateexts** feature does not always make sense.

The ProtecTIER product does a great job of equally distributing its load to its back-end disks and directly benefits all available resources, even without **rotateexts**. The typical ProtecTIER write pattern does not create hot spots on the back-end disk, so **rotateexts** does not contribute to better I/O performance.

If the repository needs to be grown, the addition of more disks to already existing extent pools, or the addition of another extent pool with all new disks, creates storage that has different performance capabilities than the already existing ones. Adding dedicated arrays with their specific performance characteristics enables the ProtecTIER server to equally distribute all data across all LUNs. So, all back-end LUNs have the same performance characteristics and therefore behave as expected.

Consider the following example. You want to use 300 GB 15,000 RPM FC drives for metadata and user data in your DS8000. To reach the wanted performance, you need four 4+4 RAID 10 arrays for metadata. Because you use FC drives, go with RAID 5 arrays and configure all user data file systems with 6+1+S RAID 5 or 7+1 RAID 5. With this approach, you do not create RAID 10 arrays with ranks that have a hot spare requirement.

As shown in Figure 11-1, the following example needs some work in order to be aligned with preferred practices.

A16	Normal	RAID 5 (6+P+S)	Assigned	R16	4	300	15 Enterprise
A17	Normal	RAID 5 (6+P+S)	Assigned	R17	4	300	15 Enterprise
A18	Normal	RAID 5 (6+P+S)	Assigned	R18	4	300	15 Enterprise
A19	Normal	RAID 5 (6+P+S)	Assigned	R19	4	300	15 Enterprise
A20	Normal	RAID 5 (7+P)	Assigned	R20	4	300	15 Enterprise
A21	Normal	RAID 5 (7+P)	Assigned	R21	4	300	15 Enterprise
A22	Normal	RAID 5 (7+P)	Assigned	R22	4	300	15 Enterprise
A23	Normal	RAID 5 (7+P)	Assigned	R23	4	300	15 Enterprise
A50	Normal	RAID 10 (3*2+2S)	Unassigned		3	300	15 Enterprise
A51	Normal	RAID 10 (4*2)	Unassigned		3	300	15 Enterprise
A52	Normal	RAID 10 (3*2+2S)	Unassigned		3	300	15 Enterprise
A53	Normal	RAID 10 (4*2)	Unassigned		3	300	15 Enterprise
A54	Normal	RAID 5 (7+P)	Unassigned		0	300	15 Enterprise
A55	Normal	RAID 5 (7+P)	Unassigned		0	300	15 Enterprise
A56	Normal	RAID 5 (6+P+S)	Unassigned		1	300	15 Enterprise
A57	Normal	RAID 5 (6+P+S)	Unassigned		1	300	15 Enterprise
A58	Normal	RAID 5 (6+P+S)	Unassigned		1	300	15 Enterprise
A59	Normal	RAID 5 (6+P+S)	Unassigned		1	300	15 Enterprise
A60	Normal	RAID 5 (7+P)	Unassigned		1	300	15 Enterprise
A61	Normal	RAID 5 (7+P)	Unassigned		1	300	15 Enterprise
A62	Normal	RAID 5 (7+P)	Unassigned		1	300	15 Enterprise
A63	Normal	RAID 5 (7+P)	Unassigned		1	300	15 Enterprise

Figure 11-1 DS8000 layout example with bad RAID 10 arrays

Figure 11-2 shows the **dsc1i** output of the **lsextpool** command and the names that are assigned to extent pools.

```
dsc1i> lsextpool
```

Name	ID	stgtype	rankgrp	status	availstor (2^30B)	%allocated	available	reserved
numvols								
=====								
=====								
TS_RAID10_0	P18	fb	0	below	1054	0	1054	1054
TS_RAID10_1	P19	fb	1	below	1054	0	1054	1054
TS_RAID5_0	P22	fb	0	below	14494	0	14494	0
TS_RAID5_1	P23	fb	1	below	15024	0	15024	0

Figure 11-2 lsextpool output

Look more closely at the extpools p18, p19, and p22 in Figure 11-3, and extent pool p23 in Figure 11-4 on page 176.

```
dsccli> lsrank -extpool p18
```

ID	Group	State	datastate	Array	RAIDtype	extpoolID	stgtype
=====							
R20	0	Normal	Normal	A20	10 P18		fb
R60	0	Normal	Normal	A60	10 P18		fb

```
dsccli> lsrank -extpool p19
```

ID	Group	State	datastate	Array	RAIDtype	extpoolID	stgtype
=====							
R21	1	Normal	Normal	A21	10 P19		fb
R61	1	Normal	Normal	A61	10 P19		fb

```
dsccli> lsrank -extpool p22
```

ID	Group	State	datastate	Array	RAIDtype	extpoolID	stgtype
=====							
R16	0	Normal	Normal	A16	5 P22		fb
R18	0	Normal	Normal	A18	5 P22		fb
R22	0	Normal	Normal	A22	5 P22		fb
R50	0	Normal	Normal	A50	5 P22		fb
R52	0	Normal	Normal	A52	5 P22		fb
R54	0	Normal	Normal	A54	5 P22		fb
R56	0	Normal	Normal	A56	5 P22		fb
R58	0	Normal	Normal	A58	5 P22		fb
R62	0	Normal	Normal	A62	5 P22		fb

Figure 11-3 Extent pool attributes for p18, p19, and p22

```
dsccli> lsrank -extpool p23
```

ID	Group	State	datastate	Array	RAIDtype	extpoolID	stgtype
=====							
R17	1	Normal	Normal	A17	5	P23	fb
R19	1	Normal	Normal	A19	5	P23	fb
R23	1	Normal	Normal	A23	5	P23	fb
R51	1	Normal	Normal	A51	5	P23	fb
R53	1	Normal	Normal	A53	5	P23	fb
R55	1	Normal	Normal	A55	5	P23	fb
R57	1	Normal	Normal	A57	5	P23	fb
R59	1	Normal	Normal	A59	5	P23	fb
R63	1	Normal	Normal	A63	5	P23	fb

Figure 11-4 Extent pool tabulates for p23

Align the respective ranks to dedicated DS8000 cluster nodes by grouping odd and even numbers of resources together in extent pools.

To ensure that you do not use **rotateexts** but keep specific repository LUNs to stick to dedicated 4+4 arrays, use the **chrank** command to reserve ranks and make them unavailable during fixed block volume creation by completing the following steps:

1. Reserve the rank 61 in extent pool p19 to make it unavailable during volume creation (Example 11-1).

Example 11-1 Reserve rank r61 with extent pool p19

```
dsccli> chrank -reserve r61
CMUC00008I chrank: Rank R61 successfully modified.
```

2. Verify the successful execution of the command by running the **lsrank** command (Example 11-2).

Example 11-2 lsrank command

```
dsccli> lsrank -l
```

ID	Group	State	datastate	Array	RAIDtype	extpoolID	extpoolnam	stgtype	exts	usedexts
=====										
R21	1	Normal	Normal	A21	10	P19	TS_RAID10_1	fb	1054	0
R61	1	Reserved	Normal	A61	10	P19	TS_RAID10_1	fb	1054	0

3. After verification, create your first of two metadata LUNs in this extent pool (Example 11-3).

Example 11-3 Create the first of two metadata LUNs

```
dsccli> mkfbvol -extpool P19 -cap 1054 -name ProtMETA_#d -volgrp V2 1900
CMUC00025I mkfbvol: FB volume 1900 successfully created.
```

- After volume creation, verify that the allocated 1054 extents for the newly created fixed block volume 1900 are all placed into rank R21 (Example 11-4).

Example 11-4 Verify allocated extents for new volume in rank r21

```
dsccli> lsrank -l
```

ID	Group	State	datastate	Array	RAIDtype	extpoolID	extpoolnam	stgtype	exts	usedexts
R21	1	Normal	Normal	A21	10 P19	TS_RAID10_1	fb		1054	1054
R61	1	Reserved	Normal	A61	10 P19	TS_RAID10_1	fb		1054	0

- Now, you can release the second rank in your extent pool to enable volume creation on it (Example 11-5).

Example 11-5 Release rank r61

```
dsccli> chrack -release r61
```

CMUC00008I chrack: Rank R61 successfully modified.

- Create the fixed block volume that is used as the metadata LUN (Example 11-6).

Example 11-6 Create metadata LUN

```
dsccli> mkfbvol -extpool P19 -cap 1054 -name ProtMETA_#d -volgrp V2 1901
```

CMUC00025I mkfbvol: FB volume 1901 successfully created.

- After volume creation, verify that the newly allocated extents are all placed in the second rank R61 (Example 11-7).

Example 11-7 Verify that new extents are placed in the second rank r61

```
dsccli> lsrank -l
```

ID	Group	State	datastate	Array	RAIDtype	extpoolID	extpoolnam	stgtype	exts	usedexts
R21	1	Normal	Normal	A21	10 P19	TS_RAID10_1	fb		1054	1054
R61	1	Normal	Normal	A61	10 P19	TS_RAID10_1	fb		1054	1054



Backup management introduction

This chapter describes the suggested settings that are common to all backup servers, and includes information about interoperability, software compatibility, zoning, and more. The subsequent chapters provide information about IBM Spectrum Protect, starting at version 7.1.3 (formerly known as IBM Tivoli Storage Manager), Symantec NetBackup, and Commvault, and provide an overview and procedural information about VMware and other backup applications.

The following chapters focus on preferred practices for specific backup applications in the IBM ProtecTIER Virtual Tape Library (VTL) and the File System Interface (FSI) for Common Internet File System (CIFS) and Network File System (NFS) environments.

This chapter contains the following topics:

- ▶ Introduction
- ▶ General suggestions
- ▶ General advice for backups
- ▶ ProtecTIER integration with backup applications
- ▶ Backup application terminology cross-reference
- ▶ Backup application catalog
- ▶ Remote cloning of virtual tapes

Note: ProtecTIER GA Version 3.4 was released with only the Virtual Tape Library (VTL) interface support. File System Interface (FSI) support was added to ProtecTIER PGA 3.4 Version. For details, see the announcement letter:

http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/6/872/ENUSA P16-0076/index.html&lang=en&request_locale=en

12.1 Introduction

Many backup servers have features and settings that are used to optimize performance when writing data to real tape cartridges. In the case of a VTL environment, the ProtecTIER repository presents a VTL with virtual drives and cartridges to the backup server, and some settings that are optimized for real tape are not required.

This might have a detrimental effect on the ProtecTIER deduplication factor and performance. Check the current settings of your backup server, and adjust those settings that do not follow the preferred practices presented in the following sections.

12.2 General suggestions

This section provides an overview of general suggestions regarding interoperability, software, backup applications, tape, and storage area network (SAN) zoning that are common to all backup servers. This section also describes compression, encryption, multiplexing, tape block sizes, and types of data that are targeted for backup.

12.2.1 Interoperability

Check the IBM Interoperability Matrix to ensure that the version of your backup server and operating system are supported by the ProtecTIER product. Also, ensure that your server host bus adapter (HBA) is compatible with the ProtecTIER product. More details are in Appendix B, “ProtecTIER compatibility” on page 457.

12.2.2 Software compatibility

Ensure that your backup server version platform and operating system version are listed in the supported hardware and software list for the ProtecTIER product. More details are in Appendix B, “ProtecTIER compatibility” on page 457.

12.2.3 Software, backup application, and operating system

Ensure that the backup application software is updated to the latest version supported by ProtecTIER. This action can affect the overall factoring performance. Also, ensure that the OS of the backup server is updated to the most recent patch or maintenance level. This action can affect overall IBM HyperFactor performance. More details are in Appendix B, “ProtecTIER compatibility” on page 457.

12.2.4 Tape library zoning

The backup server must have a dedicated HBA port or ports for the ProtecTIER VTL. This port or ports can be shared with a physical tape library. However, the physical tape library must not be in the same SAN zone as the VTL. When it is not possible to dedicate HBA ports for VTL and physical tape library, have different zones to separate the traffic.

Port sharing: Although sharing a Fibre Channel (FC) port between physical and virtual tape is possible when they are in different SAN zones, never share a port with disk attachments. Tape and disk devices require incompatible HBA port settings for reliable operation and optimal performance characteristics. Under stress conditions (high I/O rates for tape, disk, or both), where disk and tape subsystems share a common HBA port, stability problems have been observed.

Other SAN zoning suggestions

The following list provides suggestions that are common to all backup servers:

- ▶ Use zones that are based on the worldwide port name (WWPN).
- ▶ Use two-member zones, that is, one initiator port and one target port per zone. If this configuration is not possible, consider at least having one single initiator on each zone.
- ▶ Do not mix the front-end zoning (zoning between the ProtecTIER server and backup application server) with the back-end zoning (zoning between the ProtecTIER server and the ProtecTIER back-end storage). Including ProtecTIER back-end ports and front-end ports in the same SAN zone causes problems.
- ▶ Do not include more than one single ProtecTIER front-end port in a SAN zone. To avoid problems, a ProtecTIER front-end port must not see other ProtecTIER front-end ports.
- ▶ For each backup server, create a separate zone for each HBA that accesses ProtecTIER virtual resources.
- ▶ Before you create worldwide name (WWN) zones on a SAN switch, you must obtain the WWPN of each port for both your ProtecTIER server and your host computer.
- ▶ If you plan to use control path failover (CPF), you can zone your host to all ProtecTIER ports. You have more than one instance of the same robot that is recognized in the OS, but it is ready when the CPF redirects the traffic to the other port without requiring zoning changes.
- ▶ If you are not planning to use CPF, zone the host to some ports of the VTL, to balance the backup traffic among the HBAs. Spread drives across the HBAs.

Table 12-1 on page 181 shows an example of a zone from a system where one backup application server has two HBAs. The workload distribution between the two HBAs is needed and there is no intention to use CPF. Each IBM Spectrum Protect (formerly Tivoli Storage Manager) HBA discovers two front-end ports of the VTL. The tape devices in the OS appear only once, and the load is distributed.

Table 12-1 Example of a zone where load distribution is wanted

Initiator	Target
TSM_HBA_Port0	VTL_FrontEnd_port0
TSM_HBA_Port0	VTL_FrontEnd_port2
TSM_HBA_Port1	VTL_FrontEnd_port1
TSM_HBA_Port1	VTL_FrontEnd_port3

12.2.5 Compression

Compression scrambles the data that is sent to the ProtecTIER server, which makes pattern matching difficult. This data scrambling affects data matching rates, even if the same data is sent each time. The ProtecTIER product compresses the data that it sends to the back-end storage disks after the virtual tape drives receive and deduplicate the data. Disable any compression features for the ProtecTIER that are defined on the backup server and clients. Any type of server and client compression, deduplication, or encryption negatively affects the deduplication ratio on the ProtecTIER system.

12.2.6 Encryption

Encryption makes each piece of data that is sent to the ProtecTIER server unique. Encryption affects the data matching rates and the factoring performance. Even if the same data is sent each time, it appears as different data to the deduplication engine. Disable any encryption features in your backup server and client application.

12.2.7 Multiplexing

Do not use the multiplexing feature of any backup application with the ProtecTIER product. Although the ProtecTIER product works with these features, the benefits (disk savings) of the HyperFactor algorithm and compression is reduced. Disable any multiplexing features on the backup server and clients.

12.2.8 Tape block sizes

Unless otherwise indicated, to optimize the backup server performance, set the block size for data that is sent to the (virtual) tape drives to *256 kilobytes (KB)* or greater.

12.2.9 Type of data that is backed up

Another factor that affects performance in a ProtecTIER environment is the *type* of data that is targeted for backup. Some data is well-suited for data deduplication and other data is not. For example, small files (less than 32 KB in size) commonly found in operating systems do not factor well, although the built-in compression might reduce their stored size. For more information about data types, see Chapter 20, “Application considerations and data types” on page 295.

Reevaluate your current backup workloads. Decide which backups are not good candidates for ProtecTIER deduplication.

12.3 General advice for backups

Generally, the preferred method of operation for using the ProtecTIER product is to imitate the procedure that is used with physical cartridges. Implement the time frame mode of operation so that, for every 24-hour cycle, there is a backup window and then a replication window. The user must ensure that there is enough bandwidth (Transmission Control Protocol/Internet Protocol (TCP/IP) and SAN) and time allotted so that there is no overlap and no replication backlog.

The following steps present 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/database (DB) backup for disaster recovery purposes.
3. Set up the system so that replication starts and completes before the next backup cycle starts.
4. The user must have a complete and easily recoverable set of their latest daily backups, including the backup application catalog image.
5. If a disaster occurs, the user can revert to the last completed set of backups. So the recovery point objective (RPO) is in the 24-hour window that is typical for the service level agreement (SLA).

12.4 ProtecTIER integration with backup applications

The ProtecTIER repository interfaces with backup applications in two ways:

- ▶ VTL
- ▶ FSI (FSI-CIFS and FSI-NFS)

Table 12-2 summarizes the backup application support for each type of ProtecTIER interface as of Version 3.4. In the table, “Y” indicates “yes.”

Table 12-2 ProtecTIER interfaces and backup applications

Backup application	VTL	FSI-CIFS	FSI-NFS
IBM Spectrum Protect starting at V7.1.3 (formerly known as Tivoli Storage Manager)	Y	Y	Y
Symantec Veritas NetBackup (NetBackup)	Y	Y	Y
EMC NetWorker (Legato)	Y	Y	Y
Commvault	Y	Y	Y ^a
HP Data Protector	Y	Y	Y
Symantec BackupExec	Y	Y	Y

a. Supported as of Version 3.4

12.5 Backup application terminology cross-reference

Each backup application has its own terminology. Table 12-3 describes the terms that are used by each backup application.

Table 12-3 Backup application vocabulary cross-reference

Term definition	IBM Spectrum Protect (formerly known as Tivoli Storage Manager)	EMC NetWorker	NetBackup	Commvault
The object that is saved in to the backup application, for example, a file or a database table.	Backup	Save set	Image	Backup set
The physical box that connects to the backup devices, such as a tape device.	Tivoli Storage Manager Server or Library Manager	Storage node	Media server	Media agent
The location where the master database of the backups is stored.	Tivoli Storage Manager Server	NetWorker server	-	CommServe
Repository where the information that enables access and decision-making for backups is stored.	Tivoli Storage Manager database	-	Catalog	Database
A system that has data to be backed up, but has no metadata information about the backed-up data.	Client or node	Client node	-	Client
A system that has data to be backed up, and has direct attached storage, typically tape drives.	LAN-free client	Storage node (remote)	-	LAN-free Media Agent
Application that runs on a client to send data to be backed up.	Tivoli Storage Manager Client	Save	-	Data agent

12.6 Backup application catalog

The backup application catalog/database (DB) is a list of the cartridges or storage devices that are used for backup and includes the following information:

- ▶ The date when the backup was performed
- ▶ A list of files that are associated with the backup
- ▶ Retention period
- ▶ Other backup application-specific information

The backup application supports one catalog or DB per backup server instance. In many cases, the primary and disaster recover (DR) sites have two separate backup servers, each with its own DB or catalog. To efficiently read replicated cartridges at the DR site, the DR site backup server must have access to either the catalog or DB of the primary backup server or an exact copy of it.

Two main backup environment topologies are as follows:

- Single domain backup environment

A single domain backup environment shares a catalog across the primary and DR sites. Its entries are visible for both servers always. In a single domain environment, the backup application is fully aware of the whereabouts and status of the cloned cartridges.

Note: The single domain backup environment works well with Symantec NetBackup, but does not work with Tivoli Storage Manager.

- Multiple domain backup environment

A multiple domain environment requires the user to recover the DR backup server by using a copy of the catalog or DB that matches the replicated repository cartridges that are set before restoring their replicated data from the ProtecTIER system at the DR site.

12.7 Remote cloning of virtual tapes

Remote cloning is the process of using a secondary (DR) site to clone cartridges. You can use ProtecTIER replication to offload tape cloning to your secondary site. Many users replicate their data from the primary site to the secondary (DR) site, and then move it from the disk-based repository on to physical tape cartridges for long-term retention.

One of the advantages of this practice at the secondary site is that it shifts the burden of cloning to physical tape from the production environment to the DR site location. The DR site cloning operation uses the cartridge replicas at the ProtecTIER VTL shelf of the destination. The process imitates the commonly used physical process for the transportation of physical cartridges from the primary site to a DR site.

This feature is effective in single domain backup deployments because in these environments the backup application servers at both sites share the catalog and can be concurrently connected to the ProtecTIER systems. The replication visibility switch control feature is used in these environments. The cartridges to be cloned are moved from the primary repository to the secondary repository and then cloned to physical tapes.



IBM Spectrum Protect

This chapter suggests settings for IBM Spectrum Protect (formerly Tivoli Storage Manager), and includes information about interoperability, software compatibility, zoning, and more.

The IBM ProtecTIER product can be deployed as a Virtual Tape Library (VTL) or File System Interface (FSI) to IBM Spectrum Protect. This chapter describes IBM Spectrum Protect with VTL and with FSI.

Rebranding: Beginning with version 7.1.3, IBM Tivoli Storage Manager was rebranded to the name IBM Spectrum Protect. Versions prior to 7.1.3 deployed with IBM ProtecTIER are referred to as Tivoli Storage Manager.

- ▶ For more information about the branding transition, see this web page:
<http://www.ibm.com/support/docview.wss?uid=swg21963634>
- ▶ For an overview of IBM Spectrum Storage, see this web page:
<http://www.ibm.com/systems/storage/spectrum/>
- ▶ For general VTL considerations for IBM Spectrum Protect servers, see this web page:
<http://www.ibm.com/support/docview.wss?uid=swg21425849>
- ▶ For preferred practices and configuration of IBM Spectrum Protect in your ProtecTIER FSI environment, see 13.3, “IBM Spectrum Protect: FSI” on page 200.

This chapter contains the following topics:

- ▶ IBM Spectrum Protect VTL
- ▶ IBM Spectrum Protect: Preferred options
- ▶ IBM Spectrum Protect: FSI

13.1 IBM Spectrum Protect VTL

Combining the advanced capabilities and features of IBM Spectrum Protect with the powerful performance-enhancing and cost reducing capabilities of the ProtecTIER product provide IT organizations with a cost-effective way to improve the performance, reliability, and scalability of data protection.

Important: The IBM Tivoli Storage Manager parser does not recognize the backed-up file as a Tivoli Storage Manager stream when you use random access mode. Do not use random access mode with IBM Spectrum Protect, or Tivoli Storage Manager.

For more information about planning for the IBM Tivoli Storage Manager parser, and estimating the benefits of the Tivoli Storage Manager parser by using the ProcessCSV tool, see Appendix C, “ProtecTIER parsers” on page 461.

13.2 IBM Spectrum Protect: Preferred options

Review the following IBM Spectrum Protect server and client options. If necessary, change the options to enable optimum performance of the ProtecTIER server.

- ▶ IBM Spectrum Protect uses a 256 KB block size to write to tape. This block size is tied to the tape drive and the HBA. To optimize this block size with your operative system, configure an adequate memory page size that efficiently handles the IBM Spectrum Protect block size.

The most well known example is AIX, where 64 KB is the page size that adjusts better for resource utilization in terms of memory paging.

Note: Allowed page sizes in AIX are 4 KB, 64 KB and 16 MB. Using a 4 KB page size requires 64 cycles of acquiring the segment control block (SCB) lock; using a 64 KB page size requires 4 cycles of acquiring the SCB lock; and using a 16 MB page size will require only 1 cycle, but it would result in inefficient memory usage

- ▶ If IBM Spectrum Protect is installed under a NAS network, the Network Data Management Protocol (NDMP) feature can be configured to set the block size to 256 KB. See the following web pages:
 - NDMP functionality in IBM Spectrum Protect and NetApp filers:
<http://www.ibm.com/support/docview.wss?uid=swg27046965>
 - Tape libraries and drives for NDMP operations:
http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.5/srv.admin/c_ndmp_ops_libs_drives.html
 - NDMP requirements:
http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.5/srv.admin/r_ndmp_requirements.html
- ▶ Disable client compression. Keep the default parameter as **COMPRESSION NO** in the IBM Spectrum Protect Backup Client option file, or update the IBM Spectrum Protect client node definition in the IBM Spectrum Protect server with the **update node <node_name> compression=no** parameter.
- ▶ Set the server option **MOVEBATCHSIZE** to 1000 (the default value).

- ▶ Set the server option **MOVESIZETHRESHOLD** to 2048 (the default value).
- ▶ When you define the library in the ProtecTIER server, select the TS3500 as the library to be emulated by ProtecTIER (more information about setting up the virtual library is in 4.3.1, “Creating libraries” on page 51).
- ▶ When you use Windows-based IBM Spectrum Protect servers, use IBM Tape Driver, not the IBM Spectrum Protect included drivers.
- ▶ With the ProtecTIER product, servers can share one virtual library, or you can create a separate virtual library for each IBM Spectrum Protect server.
- ▶ Set the IBM Spectrum Protect device class to represent the Ultrium LTO3 tape without compression by using the **FORMAT=ULTRIUM3** parameter.
- ▶ Configure the estimated capacity size, in the IBM Spectrum Protect device class, to represent the virtual tape size that is defined in the VTL, by using the **ESTCAPacity** parameter in the device class definition.
- ▶ When you define the library in the IBM Spectrum Protect Server, if the IBM Spectrum Protect Server version is Version 6.3 or higher, use the **LIBType=VTL** parameter. When using this library type, there can be no mixed media in the library and a path must be defined for all the drives to all of the servers or storage agents that use the library.
- ▶ For best results, create as many drives as necessary for your environment, taking into consideration multi-streamed and parallel backups to improve performance. If more drives are required, you can logically partition the VTL into multiple libraries and assign drives to each library. Operating system and storage area network (SAN) hardware configurations might impose limitations on the number of devices that can be used in the VTL library.
- ▶ Use **RELABELSCRatch=yes** in the IBM Spectrum Protect library definition to specify that the server relabels volumes that are deleted and returned to scratch to free up the space in the VTL repository. Without relabeling scratch tapes, the allocated space in the ProtecTIER repository is not released.

IBM Spectrum Protect command syntax: To verify the command syntax and other topics related to IBM Spectrum Protect, see the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/SSGSG7/landing/welcome_ssgsg7.html

13.2.1 LAN-free backups with the ProtecTIER product

Local area network (LAN)-free backups are simpler with the ProtecTIER product because there are increased tape resources and fewer hardware restrictions. ProtecTIER configured as a VTL has the advantage of presenting greatly increased tape resources to the backup server. So, you are able to perform LAN-free backups to the ProtecTIER server without considering the limitations that are normally applied to these backups, such as tape drive availability.

If you have many LAN-free clients, then a possibility is that your LAN-free backup windows are dictated not entirely by business needs but also by hardware availability. With the ProtecTIER product and its maximum of 256 virtual tape drives per ProtecTIER node, you can virtually eliminate any previous hardware restrictions, and schedule your backups as and when they are required by your business needs. The IBM System Storage TS7620 ProtecTIER Deduplication Appliance Express (TS7620) can support only up to 64 virtual drives per node.

LUN masking: Enable logical unit number (LUN) masking with LAN-free clients. LUN masking reduces the administration of path creation at the IBM Spectrum Protect server. For more information, see 6.3.1, “LUN masking methods and preferred practices” on page 88.

13.2.2 Data streams

You might be able to reduce your current backup window by taking full advantage of the throughput performance capabilities of the ProtecTIER product. If tape drive availability is limited for concurrent backup operations on your information technology (IT) storage management (IBM Spectrum Protect) server, you can define a greater number of virtual drives. Reschedule backups to run at the same time to maximize the number of allowable parallel tape operations on ProtecTIER servers.

You can also increase the number of parallel streams for the IBM Spectrum Protect database backup using the **NUMSTREAMS** parameter (maximum of 4).

Important:

- ▶ If you choose to implement this strategy, you might need to increase the value of the **MAXSESSIONS** option on your IBM Spectrum Protect server to specify the maximum number of simultaneous client sessions that can connect to the server. For more information, see the IBM Knowledge Center:
http://www.ibm.com/support/knowledgecenter/SSGS67_7.1.5/srv.reference/r_opt_server_maxsessions.html?lang=en
- ▶ You might also need to update the maximum number of mount points (**MAXNUMMP**) in the IBM Spectrum Protect node registration to specify the maximum number of mount points a node can use on the server or storage agent only for backup operations.
- ▶ In the IBM Spectrum Protect client option file, you can set the resource utilization (**RESOURCEUTILIZATION**) parameter to specify the level of resources that the IBM Spectrum Protect server and client can use during processing. For more information, see the IBM Knowledge Center:
http://www.ibm.com/support/knowledgecenter/SSGS67_7.1.4/client/r_opt_resourceutilization.html?lang=en
- ▶ See *Harnessing the Power of ProtecTIER and Tivoli Storage Manager*, SG24-8209.

13.2.3 Reclamation

Continue to reclaim virtual storage pools that are on the ProtecTIER server. The thresholds for reclamation might need some adjustment until the system reaches a *steady state*. In a steady state, the fluctuating size of the virtual cartridges stabilizes and you can decide what the fixed reclamation limit ought to be.

When you decide how many virtual cartridges to define, consider the current storage pool **reusedelay** value. This value is equal to the number of days that your IBM Spectrum Protect database backups are retained. The same delay period applies to your storage pools that store data on ProtecTIER virtual cartridges. You might need to increase the number of pools to ensure that you always have scratch cartridges available for backup.

Note: A fix pack exists that must be applied to certain IBM Spectrum Protect versions so that **REUSEDelay** and **RELABELSCRATCH** work correctly; otherwise, you receive this error:

IC78915: RELABELSCRATCH PARAMETER DOES NOT WORK ON VTL WITH REUSEDelay
PARAMETER GREATER THAN 0

For more information, go to the following web page:

<http://www.ibm.com/support/docview.wss?uid=swg1IC78915>

13.2.4 Collocation

When you use a virtual library, consider implementing collocation for your primary storage pools.

Collocation means that all of your data for a node or node group is contained on the same set of virtual cartridges. You can also collocate by file system or group of file systems. Because you do not have any of the limitations of physical cartridges that are normally associated with this feature (such as media and slot consumption), you can enable the option.

Collocating data with similar expiration characteristics

As much as possible, collocate data with similar expiration characteristics, and then let that data expire. This collocation practice minimizes reclamation and helps reduce the IBM Spectrum Protect workload. It also reduces the risk of replicated cartridges being out of synchronization because of the timing of the reclamation activity.

13.2.5 Migration

When using ProtecTier to migrate data from Disk Storage Pools you can use as many IBM Spectrum Protect migration processes (**MigProcess**) as possible in the definition of the Storage Pool and accelerate your administration cycle with this process.

This is even more highly suggested with IBM Spectrum Protect version 7 and the Multi-threaded Storage Pool Migration feature that improves the server's efficiency when running migration from stgpools with DISK devclass for nodes with a large number of file spaces. It applies to migration to storage pools that do not have collocation specified, or that have collocation by file space.

13.2.6 Backing up your IBM Spectrum Protect database

Starting with Tivoli Storage Manager V6.3, and also IBM Spectrum Protect starting at version 7.1.3, in addition to using the **backup db** command along with using up to four streams (**NUMStreams=4**), another integrated deduplication efficiency option (**DEDUPDEVICE=YES**) is available for the IBM DB2-based IBM Spectrum Protect database. The same command is suggested when using DB2 9.7 Fix Pack 4 or later. It specifies that a target storage device supports data deduplication. When set to YES, the format for backup images is optimized for data deduplication devices, making backup operations more efficient.

Note: Allowed page sizes in AIX are 4 KB, 64 KB and 16 MB. Using a 4 KB page size requires 64 cycles of acquiring the segment control block (SCB) lock, using a 64 KB page size requires four cycles of acquiring the SCB lock, using a 16 MB page size will require only one cycle, but it can result in inefficient memory usage.

13.2.7 Physical tape

Depending on your data protection requirements, a necessary step might be to copy the deduplicated data to physical tape. You can complete this task by using standard IBM Spectrum Protect copy storage pools that have device classes that direct data to physical libraries and drives.

Tip: Estimate the number of drives that can be used to move data to physical tapes and consider a single stream performance. Remember that when data is sent to physical tapes, it is rehydrated. *Rehydrate* is a term for reading all parts of the data and assembling it to fully match the original backed up data.

Migrating data

Do not expect an effective deduplication when you migrate your existing data from physical tape to the ProtecTIER repository if the data was originally backed up without preferred practices in place. Use the most current version of the ProtecTIER product so that you implement the appropriate IBM Tivoli Storage Manager parser, which maximizes your overall deduplication factoring ratio.

13.2.8 Avoiding mount conflicts

To avoid a mount conflict, increase the number of drives (according to your needs) up to 512 per dual-node cluster (256 per node). The TS7620 Appliance Express supports up to 64 virtual drives per node. Depending on your IBM Spectrum Protect version or operating system, these maximum values might change.

Cartridge size: As a rule, an optimal cartridge size is 100 gigabyte (GB) to reduce the reclamation load and to enable concurrent replication, but this can vary depending on the kind of data being backed up.

13.2.9 Multiple streams from the client with the resourceutilization parameter

When possible, use multiple streams for the client backup. Try using four or more concurrent streams when you need maximum performance. You can set up multiple streams by modifying the `dsm.opt` (Windows) or `dsm.sys` (UNIX) file on the client and specify the **resourceutilization** parameter. For more information, see the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.4/client/r_opt_resourceutilization.html?lang=en

The **resourceutilization** option increases or decreases the ability of the IBM Spectrum Protect client to create multiple sessions. For backup or archive, the value of **resourceutilization** does not directly specify the number of sessions that are created by the client.

However, this setting specifies the level of resources that the IBM Spectrum Protect server and client can use during backup or archive processing. The higher the value, the more sessions that the client can start if it deems necessary. The range for the parameter is 1 - 10.

When the option is not set, which is the default, then only two sessions are created on the server. The default **resourceutilization** level is 0, and it enables up to two sessions running on the server, one for querying the server and one for sending file data.

The value **resourceutilization=5** permits up to four sessions (two for queries and two for sending data), and **resourceutilization=10** permits up to eight sessions (four for queries and four for sending data) with the server. The relationship between **resourceutilization** and the maximum number of sessions that is created is part of an internalized algorithm and, as such, is subject to change.

The following values are the preferred settings:

- ▶ For workstations: resourceutilization 1
- ▶ For a small server: resourceutilization 5
- ▶ For a large server: resourceutilization 10

Table 13-1 lists the relationships between **resourceutilization** values and the maximum sessions that are created. Producer sessions scan the client system for eligible files. The remaining sessions are consumer threads and are used for data transfer. Subtract the producer sessions that are listed in the table from the maximum number of sessions to determine the number of consumer threads.

Table 13-1 Relationship between the RESOURCEUTILIZATION value and maximum sessions created

RESOURCEUTILIZATION value	Maximum number of sessions	Unique number of producer sessions
1	1	0
2	2	1
3	3	1
4	3	1
5	4	2
6	4	2
7	5	2
8	6	2
9	7	3
10	8	4
Default (0)	2	1

Note: When you use IBM Spectrum Protect for Virtual Environments - Data Protection for VMware to back up virtual machines in a VMware vSphere environment, there is one mount point used per VM backup (for example, one virtual tape drive). Consider backing up multiple virtual machines using the **vmmaxparallel** option to take advantage of the ProtecTier multi-stream feature.

For more information about how to use IBM Spectrum Protect for Virtual Environments - Data Protection for VMware with the ProtecTIER product, see the following web page:

<http://www.ibm.com/support/docview.wss?uid=swg27021081>

13.2.10 Accommodating increased sessions

Ensure that the **MAXSESSIONS** setting on the IBM Spectrum Protect server can accommodate the increased sessions. The default value for **MAXSESSIONS** is 25. Set this parameter in the IBM Spectrum Protect server options file (IBM Spectrum Protect must be halted and then restarted) or run the **setopt** command, as shown in Example 13-1.

Example 13-1 Set MAXSESSIONS parameter

```
tsm: SERVER1>setopt MaxSessions 100
```

```
Do you wish to proceed? (Yes (Y)/No (N)) y
```

```
ANR2119I The MAXSESSIONS option has been changed in the options file.
```

Also, update the **NODE** definition on the IBM Spectrum Protect server to enable more than one mount point (**MAXNUMMP**). For more information, see the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_node_update.html?lang=en

13.2.11 IBM Spectrum Protect storage pool selection

When you select storage pools to restore or to retrieve data, the server evaluates the number of volumes that are required for the operation and selects the storage pool with the fewest volumes.

Usually, a VTL that is set up with small logical volumes often has data that is spread over more volumes than the data in a physical tape library. As a result, the server selects the physical tape storage pool, which has fewer volumes, rather than the faster VTL storage pool.

To force the server to ignore the number of volumes when you select a storage pool to restore or to retrieve data, use the **IGNORENUMVOLSCHECK** IBM Spectrum Protect server option.

Storage pool selection: For more information about the storage pool selection, see the following web page:

<http://www.ibm.com/support/docview.wss?uid=swg21417248>

13.2.12 Technical overview

Figure 13-1 illustrates a typical IBM Spectrum Protect environment that uses the ProtecTIER product. The IBM Spectrum Protect environment is straightforward. The IBM Spectrum Protect servers are connected to storage devices (disk, real tape, or virtual tape), which are used to store data that is backed up from the clients.

Every action and backup set that is processed by IBM Spectrum Protect is recorded in the IBM Spectrum Protect database. Without a copy of the IBM Spectrum Protect database, a IBM Spectrum Protect server cannot restore any of the data from the storage devices.

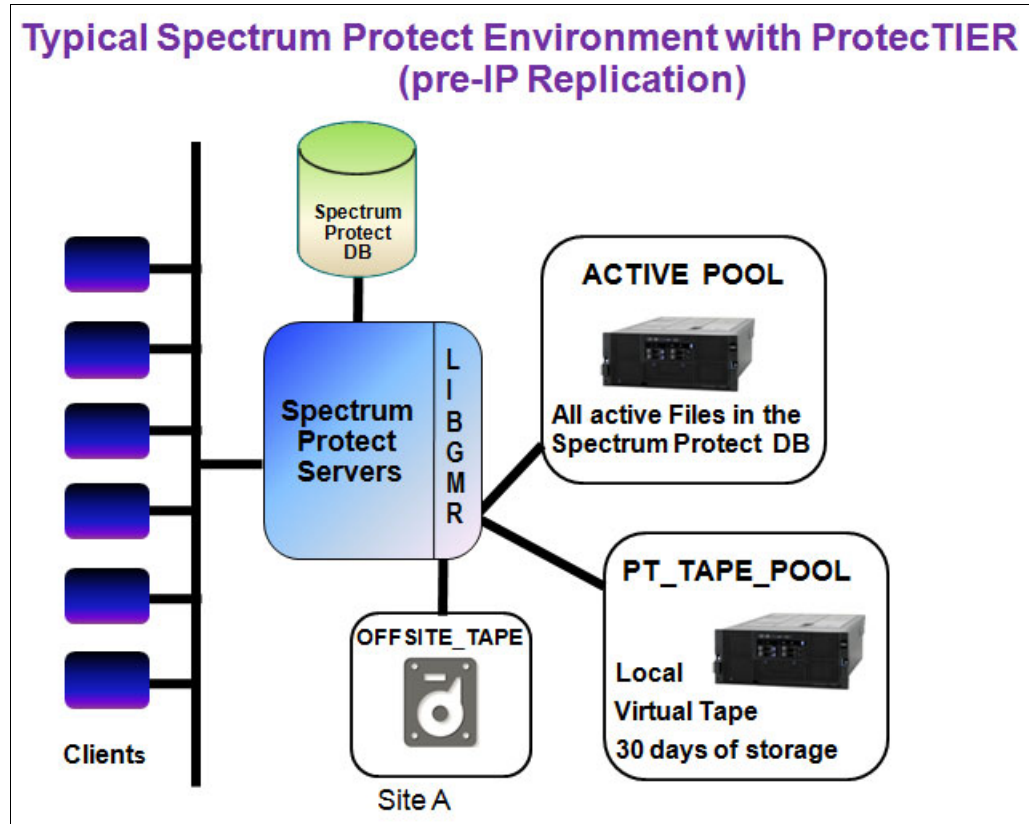


Figure 13-1 Typical IBM Spectrum Protect Environment with ProtecTIER (pre-IP replication)

The ProtecTIER product provides a virtual tape interface with the IBM Spectrum Protect servers and supports the creation of two storage pools:

- ▶ The ACTIVE IBM Spectrum Protect pool
- ▶ The ONSITE TAPE pool (called PT_TAPE_POOL)

Using the configuration that is shown in Figure 13-1, the user creates a storage pool to create real physical tapes to take offsite (called OFFSITE_TAPE). The user sizes the PT_TAPE_POOL (ProtecTIER system) to store all active client files plus about 30 days worth of inactive client files on virtual tape.

The user creates an ACTIVE POOL, which is also hosted on the ProtecTIER system. The ACTIVE POOL pool contains the most recent (active) files that are backed up from all client servers. Client restoration information comes from the ACTIVE Pool.

13.2.13 Advantages of an IBM Spectrum Protect environment with ProtecTIER

The configuration that is shown in Figure 13-1 on page 195 eliminates the usage of physical tape in the data center and enables faster restoration, because the information is coming from the ProtecTIER disk-based virtual tape versus real tape.

Volume definition: When you predefine the volume at IBM Spectrum Protect for backup, this volume is not parsed.

13.2.14 IBM Spectrum Protect configuration with VTL

This section describes the preferred practices when you configure IBM Spectrum Protect to work with the ProtecTIER VTL product. Although the operations mentioned in this section are described using the command-line interface (CLI), be aware that this operation can also be executed by using the Operations Center. To enable a ProtecTIER server to work with IBM Spectrum Protect, complete the following steps:

1. Ensure that devices (robot and tapes) are recognized by the operating system.
Example 13-2 shows how the devices appear in an AIX server.

Output: For the sake of brevity, some of the output in these examples is shortened.

Example 13-2 Robot (smc0, smc1) and tapes (rmt0 - rmt63) in the AIX

```
lbsserv38> lsdev -Cc tape
rmt0 Available 02-08-02 IBM 3580 Ultrium Tape Drive (FCP)
rmt1 Available 02-08-02 IBM 3580 Ultrium Tape Drive (FCP)
rmt2 Available 02-08-02 IBM 3580 Ultrium Tape Drive (FCP)
...
rmt61 Available 03-08-02 IBM 3580 Ultrium Tape Drive (FCP)
rmt62 Available 03-08-02 IBM 3580 Ultrium Tape Drive (FCP)
rmt63 Available 03-08-02 IBM 3580 Ultrium Tape Drive (FCP)
smc0 Available 02-08-02 IBM 3584 Library Medium Changer (FCP)
smc1 Available 03-08-02 IBM 3584 Library Medium Changer (FCP)
```

2. In the administrative console (**dsmadm**), you can define the library by using the VTL library type and **relabelscratch** parameters (Example 13-3).

Example 13-3 Defining a library in the IBM Spectrum Protect server

```
TSM:SERVER1> define library ptlibrary libtype=vtl relabelscratch=yes shared=yes
ANR8400I Library PTLIBRARY defined.
```

Library tape VTL: The library type VTL was introduced in Tivoli Storage Manager V6.3 to improve communication between Tivoli Storage Manager and the ProtecTIER VTL.

If you are using a version of Tivoli Storage Manager older than V6.3, you can use **libtype=scsi** instead.

- If you are using Tivoli Storage Manager V6.3 or later, you can create the path to the library and its subordinate drives by running the **perform libaction** command. For more information, see 13.2.16, “Defining and deleting IBM Spectrum Protect libraries with many drives” on page 199. Example 13-4 shows the **perform libaction** command.

Example 13-4 The perform libaction command

```
tsm: SERVER1>perform libaction ptlibrary action=define device=/dev/smc0
prefix=vtldr
ANR1720I A path from SERVER1 to PTLIBRARY has been defined.
ANR8955I Drive DR00 in library PTLIBRARY with serial number  is updated with
the newly discovered serial number 4641632000.
ANR1720I A path from SERVER1 to PTLIBRARY DR00 has been defined.
```

Tip: To run the **perform libaction** command, the **SANDISCOVERY** parameter must be set to **on**. The SAN discovery function relies on the operating system device names, such as IBMtape0 or IBMchanger0. If you configured customized device names, for example, on Linux with udev, the SAN discovery function does not use those device names. To enable SAN discovery, run the following command:

```
tsm: SERVER1> setopt SANDISCOVERY on
Do you wish to proceed? (Yes (Y)/No (N)) y
ANR2119I The SANDISCOVERY option has been changed in the options file.
```

The IBM Spectrum Protect server **SANDISCOVERY** setting must be disabled for IBM Spectrum Protect CPF/DPF functionality. The ProtecTIER product exports multiple tape drives with the same WWPN, and the SAN discovery feature does not work as expected, so it must be turned off.

The **SANDISCOVERY** setting can be turned on temporarily so that IBM Spectrum Protect can perform the **libaction** command. It can then be turned off when you use control path failover (CPF)/data path failover (DPF) in IBM Spectrum Protect.

- To display the SAN devices, run the **query san** command (Example 13-5).

Example 13-5 query san command output

```
TSM:SERVER1>query san
```

Devicetype	Vendor	Product	Serial Number	Device
-----	-----	-----	-----	-----
LIBRARY	IBM	03584L32	0046416329990402	/dev/smc0
DRIVE	IBM	ULT3580-TD3	4641632000	/dev/rmt0
DRIVE	IBM	ULT3580-TD3	4641632001	/dev/rmt1

- If you are using a version of Tivoli Storage Manager earlier than Version 6.3, you must manually define the path of the robot and all of its subordinate drives, as shown in Example 13-6.

Drives definition: Define the names for the drive in Tivoli Storage Manager that represent the VTL.

Example 13-6 Manually defining the virtual tapes drives

```
tsm: SERVER1>def path SERVER1 PTLIBRARY srct=SERVER destt=library dev=/dev/smc0
ANR1720I A path from SERVER1 to PTLIBRARY has been defined.
```

```
tsm: SERVER1>q path
```

Source Name	Source Type	Destination Name	Destination Type	On-Line
SERVER1	SERVER	PTLIBRARY	LIBRARY	Yes

```
tsm: SERVER1>def drive PTLIBRARY drive1
ANR8404I Drive DRIVE1 defined in library PTLIBRARY.
```

```
tsm: SERVER1>def path SERVER1 DRIVE1 srct=SERVER destt=drive library=PTLIBRARY
device=/dev/rmt0
ANR1720I A path from SERVER1 to PTLIBRARY DRIVE1 has been defined.
```

Important: The HBA wrapper files that are included with the IBM Spectrum Protect server package (except on AIX) provide communication with the virtual library. If AIX:/usr/lib/libhbaapi.a (provided by AIX with the HBAAPI installation) is not correctly configured, the following error might occur:

```
ANR1803W SAN discovery module /opt/tivoli/tsm/server/bin/dsmqsan is not
installed correctly.
ANR1791W HBAAPI wrapper library libHBAAPI.a(shr_64.o) failed to load or is
missing.
ANR1792W HBAAPI vendor library failed to load or is missing.
ANR8396E PERFORM LIBACTION: Library PTLIBRARY is not capable of discovering
the drives that it owns.
```

To resolve this error, ensure that the SAN discovery module can run, has the setuid bit turned on, and is owned by root. The SAN discovery module is called dsmqsan, and must be in the server or storage agent executable directory, as shown here:

```
chown root:system /opt/tivoli/tsm/server/bin/dsmqsan
chmod 4755 /opt/tivoli/tsm/server/bin/dsmqsan
```

- Label the virtual tapes. Example 13-7 shows the `label libvol` command that is required for label creation.

Example 13-7 The label libvol command

```
tsm: SERVER1>label libvol PTLIBRARY checkin=scratch search=yes
labelsource=barcode
ANS8003I Process number 8 started.
```

Tip: If the `AUTOLabel = yes` parameter is defined in the IBM Spectrum Protect library definition, you can run the `checkin labelsource=barcode` command.

7. Define the device class in the IBM Spectrum Protect Server for the library (Example 13-8).

Example 13-8 Device class configuration

```
tsm: SERVER1>define devclass pvtldevclass library=ptlibrary devtype=lto
estcapacity=100000M format=ultrium3
ANR2203I Device class PTCLASS
defined.
```

8. The remaining tasks vary for each client environment. You must create storage pools by using the device class that is configured for the VTL. Then, you must update the management classes and backup/archive copy groups. After these steps are complete, you can explore the advantages of the ProtecTIER product.

13.2.15 Updating to a VTL library type

You can update an existing small computer system interface (SCSI) library to a VTL library type by running the following command:

update library <libname> LIBTYPE=VTL

Updating the library to a VTL library type enables ProtecTIER to make more accurate assumptions and skip unnecessary SCSI validations. Depending on the operating system, you might run into limitations, such as 300 - 500 maximum drives.

Setting the **LIBTYPE=VTL** also eliminates the restriction of defining only 120 tape drives when you use IBM Spectrum Protect and the ProtecTIER product. This feature is available only with Tivoli Storage Manager V6.3 or later. If you have a previous version of Tivoli Storage Manager, you might want to review APAR IC66116 (Large number of tape drives can cause volume mounts to perform slowly) at the following website:

<http://www.ibm.com/support/docview.wss?uid=swg1IC66116>

Updating the library: The **UPDATE LIBRARY** function does not support mixed media, such as drives with different device types or device generations in the same library (LTO2 and LTO3). Also, this function requires that online paths be defined for servers and storage agents to all drives in the library. If paths are missing or offline, the performance levels degrade to the level of the SCSI library.

13.2.16 Defining and deleting IBM Spectrum Protect libraries with many drives

Version 6.3 introduced the **PERFORM LIBACTION** command, which automatically defines the library/path/drive structure in a single command. This action is helpful when you use the ProtecTIER product because it usually contains many drives. When you run the **PERFORM LIBACTION** command, complete the following steps:

1. Create a path to the library.
2. Scan for all drives that belong to the library and define all drives that are found in the library.
3. Define paths to the drives found in the library, respecting the virtual tape drive numbering.

Use IBM Spectrum Protect to define and delete IBM Spectrum Protect libraries (VTL and SCSI library types only) that contain many drives by running the following command:

PERFORM LIBAction <library_name> device=xxx action=define

Example 13-9 shows a partial output of this command.

Example 13-9 Partial output of a perform libaction command

```
tsm: SERVER1>perform libaction ptlibrary action=define
device=/dev/smc0 prefix=vtldr
ANR1720I A path from SERVER1 to PTLIBRARY has been defined.
ANR2017I Administrator ADMIN issued command: PERFORM LIBACTION
PTLIBRARY action=define device=/dev/smc0 prefix=dr
ANR2017I Administrator ADMIN issued command: DEFINE DRIVE
PTLIBRARY DR00
ANR8404I Drive DR00 defined in library PTLIBRARY.
ANR2017I Administrator ADMIN issued command: DEFINE PATH SERVER1
DR00
SRCTYPE=SERVER DESTTYPE=DRIVE LIBRARY=PTLIBRARY
...
ANR2017I Administrator ADMIN issued command: DEFINE PATH SERVER1
DR00
SRCTYPE=SERVER DESTTYPE=DRIVE LIBRARY=PTLIBRARY DEVICE=/dev/rmt32
ANR8955I Drive DR00 in library PTLIBRARY with serial number is
updated with
the newly discovered serial number 4641632001.
```

Run the **PERFORM LIBACTION** command to set up a single SCSI or virtual tape library (VTL).

If you are setting up or modifying your hardware environment and must create or change many drive definitions, the **PERFORM LIBACTION** command can make this task much simpler. You can define a library and then define all drives and paths to the drives. If you have an existing library that you want to delete, you can delete all existing drives and their paths.

The **PERFORM LIBACTION** command can be used only for SCSI and VTL libraries. If you are defining drives and paths for a library, the **SANDISCOVERY** option must be supported and enabled.

Note: For more information, see the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.srv.doc/t_vtl_libaction_perform.html?cp=SSGSG7_7.1.1%2F3-10-2-1-2-2

13.3 IBM Spectrum Protect: FSI

This section provides the following information:

- ▶ Steps and preferred practices for configuring and setting up IBM Spectrum Protect for backup and restore.
- ▶ IBM Spectrum Protect parameters and settings for best performance with ProtecTIER FSI.
- ▶ Configuration steps and parameters to configure IBM Spectrum Protect sequential-access disk (FILE) device classes for usage with ProtecTIER FSI.

The required steps and preferred practices for IBM Spectrum Protect with ProtecTIER FSI-CIFS, or IBM Spectrum Protect on a UNIX system with a ProtecTIER FSI-NFS, are identical unless otherwise noted.

Note: ProtecTIER GA Version 3.4 was released with only the Virtual Tape Library (VTL) interface support. File System Interface (FSI) support was added to ProtecTIER PGA 3.4 Version. For details, see the announcement letter:

http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/6/872/ENUSA P16-0076/index.html&lang=en&request_locale=en

13.3.1 Setting up backup and restore on IBM Spectrum Protect

Before you set up the new IBM Spectrum Protect device classes for the FSI export, disable the server option directly. On a Linux system, configure the value for the **DIRECTIO** parameter, as shown in Example 13-10, in the `dsmserv.opt` file. IBM Spectrum Protect does try to use direct I/O to files on NFS shares when it is using `dataformat=native` (the default) on the storage pools.

Example 13-10 Disable direct IO for storage pools on NFS exports

```
DIRECTIO NO
```

To set backup and restore on IBM Spectrum Protect, start the IBM Spectrum Protect server by running the following command on the UNIX host:

```
/opt/tivoli/tsm/server/bin/dsmserv
```

Note: This operation can be executed by using the Operations Center.

The `dsmserv` command starts the IBM Spectrum Protect server. After the server is running, complete the following steps from the server in preparation for performing an IBM Spectrum Protect backup:

1. With ProtecTIER FSI, create sequential I/O device classes when you store the IBM Spectrum Protect volumes in an FSI share. A random file device class is not supported. The definition of the device class is the only step that differs between a Windows and UNIX IBM Spectrum Protect configuration, because it contains the operating system-specific file system notations.

Therefore, two dedicated examples are shown for each operating system type. To create the device class, run the command that is shown in Example 13-11 from the administrative command line of the IBM Spectrum Protect server.

Example 13-11 Define the sequential device class for UNIX with FSI-NFS

```
DEFINE DEVclass PT1_fsi_devclass DEVType=FILE MOUNTLimit=192 MAXCAPacity=8G  
DIRectory=/mnt/puck_tsm2_nfs1,/mnt/puck_tsm2_nfs2
```

Alternatively if you configure a IBM Spectrum Protect server on Windows with a ProtecTIER CIFS share, run the command that is shown in Example 13-12.

Example 13-12 Define a sequential device class for Windows with FSI-CIFS

```
DEFINE DEVclass PT1_fsi_devclass DEVType=FILE MOUNTLimit=192 MAXCAPacity=8G  
DIRectory=\\FSI_IP_alias\\share1
```

Important: When you specify a file device class in IBM Spectrum Protect on Windows, do not specify the IP address when you configure the directory parameter. The full path of the IBM Spectrum Protect volumes is stored in the IBM Spectrum Protect database, including the network path with the Internet Protocol (IP) address.

Whenever the IP address of the FSI share changes or when you want to switch to a replicated file system on a remote repository, this action is not possible. Therefore, specify a host alias of the FSI interface IP in your domain server or create an entry in the operating system hosts file and specify the alias rather than any IP address.

Note: The MAXCAPacity=8GB suggestion is no longer relevant when using ProtecTIER Post General Availability V3.3.3 or later. From this code level and later, there is no suggestion to limit the backup file size and in fact the bigger the better

The parameter for the **define devclass** command is described with the default values. The suggested values for ProtecTIER FSI, along with additional explanations, are described in Table 13-2.

Table 13-2 Define devclass command parameters, default values, and suggested values

Parameter	Description	Suggested or maximum values	Default value
MOUNTLimit	Number of parallel streams. For VTL, this is the number of parallel mounted virtual cartridges. For FSI, this specifies the number of IBM Spectrum Protect volumes that can be opened simultaneously for read and write operations.	16, with a maximum of 192 parallel streams for IBM System Storage TS7650G ProtecTIER Deduplication Gateway (TS7650G) and a maximum of 64 streams for TS7620 SM2	Default is 20.
MAXCAPacity	Maximum file size a IBM Spectrum Protect volume can reach. Based on test results, we suggest a maximum volume size of 8 GB for deviceType=file. ^a	8 GB ^a	Default is 2 GB.
DIRectory	Location of the IBM Spectrum Protect volumes. When you define a file device class, the location is either a CIFS or NFS share. If the server needs to allocate a scratch volume, it creates a file in one of these directories. Furthermore, the distribution of incoming data streams is distributed across the available directory definitions, based on round robin. If one path is unavailable, then IBM Spectrum Protect selects the next available path. If a path is unavailable for read/restore operations, the operation might fail. Therefore, it is the administrators responsibility to implement failover mechanisms, such as bonding.	Specify at least two file shares, as shown in Example 13-11. Furthermore, the administrator needs to implement fault tolerance mechanisms for the network shares, such as bonding, to prevent path failures.	Default is the current working directory of the server when the command is issued.

a. The MAXCAPacity=8 GB suggestion is no longer relevant when using ProtecTIER PGA V3.3.3 or later. From this code level and later, there is no suggestion to limit the backup file size, and in fact the bigger the better.

2. To create your ProtecTIER storage pool, run the command that is shown in Example 13-13 from the IBM Spectrum Protect server.

Example 13-13 Defining a storage pool

```
DEfINE STGpool PT1_stgpool PT1_fsi_devclass P0oltype=Primary Description="stgpool on
ProtecTIER NFS share" REClaim=90 RECLAIMProcess=16 MAXSCRatch=200 DEDUPlicate=no
DataFormat=NATive
```

Important: When you use the **define stgpool** command, see Table 13-3 on page 203, for the parameter value descriptions.

Table 13-3 The define stgpool command parameters, default values, and suggested values

Parameter	Description	Suggested value	Default value
REClaim	(Optional) Specifies when the server reclaims a volume. This action is based on the percentage of reclaimable space on a volume. Reclamation makes the fragmented space on the volumes usable again by moving any remaining unexpired files from one volume to another volume, therefore making the original volume available for reuse.	90	60
RECLAIMProcess	(Optional) Specifies the number of parallel processes to use for reclaiming the volumes in this storage pool.	TS7650: 3 TS7620: 3	1
MAXSCRatch	Specifies the maximum number of scratch volumes that the server can request for this storage pool. To identify the number of scratch volumes, determine the ProtecTIER file system size that you configured by looking at the size of a scratch volume and reduce the number by 2 to avoid exceeding the file system size: $MAXSCRatch = (PT\ FS\ size\ in\ gigabytes / 8\ GB)^a - 2$ For example: File system size=10 TB $MAXCAPacity = 8\ GB^a$ $MAXSCRatch = (10 * 1024 / 8) - 2 = 1278$	The number of scratch volumes is based on the file system size that you configured on the ProtecTIER system.	
REUsedelay	Specifies the number of days that must elapse after all files are deleted from a volume before the volume can be rewritten or returned to the scratch pool.		Default value is 0, which means that a volume can be rewritten or returned to the scratch pool when all the files are deleted from the volume.
MIGProcess	Specifies the number of parallel processes to use for migrating files from the volumes in this storage pool.	TS7650/TS7620: 3	Default value is 1.

Parameter	Description	Suggested value	Default value
DATAFormat	Specifies the data format to use to back up files to this storage pool and restore files from this storage pool.	Native. Important: Do not use the following option that can adversely affect performance and deduplication ratio: DATAFormat=nonblock	Native
DEDuplicate	If you use a version earlier than Tivoli Storage Manager V6.2, the DEDuplicate option might not exist.	No	No

a. The MAXCAPacity=8 GB suggestion is not relevant anymore when using ProtecTIER PGA V3.3.3 or later. From this code level and higher there is no suggestion to limit the backup file size, and in fact the bigger the better.

3. To create your policy domain, run the following command:

```
DEfIne D0main domain1 BACKRETention=30 ARCHRETention=365
```

4. To create a policy set, run the following command:

```
DEfIne P0licyset domain1 policyset1
```

5. To create a management class, run the following command:

```
DEfIne mgmtclass domain1 policyset1 mgmtclass1
```

6. To create a copy group, run the following commands:

```
DEfIne C0pygroup domain1 policyset1 mgmtclass1 STANDARD Type=Backup FREQuency=0
DESTination=PT1 VERExists=NOLimit VERDeleted=5 RETExtra=30 RETOnly=60
DEfIne C0pygroup domain1 policyset1 mgmtclass1 STANDARD Type=Archive DESTination=PT1
RETVer=365 RETInit=CREATion
```

7. To set the default management class in the policy, run the following command:

```
ASsign DEFMGmtclass domain1 policyset1 mgmtclass1
```

8. To activate the policy, run the following command:

```
ACTivate P0licyset domain1 policyset1
```

9. To register a client, run the following command:

```
REGister Node <node name> <node password> PASSExp=0 D0main=domain1 COMPression=no
ARCHDElete=yes BACKDElete=yes Type=Client MAXNUMMP=80 DEDUPlication=SERVEROnly
```

DEDuplicate option: If you use a version earlier than Tivoli Storage Manager Server V6.1 or Tivoli Storage Manager Client V6.2, the **DEDuplicate** option does not exist. Do not add it to the command. Server-side data deduplication was introduced with Tivoli Storage Manager V6.1. Client-side data deduplication was introduced with Version 6.2.

10. To set the client option file, open the directory where the IBM Spectrum Protect client is installed and open the dsm.opt file. The file options should be set (Example 13-14).

Example 13-14 Client option file settings

```
NODENAME <the node name given at the register node command>
TCPSEVERADDRESS <IBM Protect server ip>
resourceutilization 10
compression no
deduplication no
```

Encryption: The encryption option should not be activated, otherwise it will negatively impact the ProtecTIER's Hyperfactor capability.

13.3.2 Parameters for best performance with ProtecTIER FSI

Table 13-4 shows the IBM Spectrum Protect parameters and settings for best performance with ProtecTIER FSI.

Table 13-4 IBM Spectrum Protect parameter settings for best performance with ProtecTIER FSI

Value	Parameter	Component	Server or client
File	DEVType	Devclass	IBM Spectrum Protect Server
8 GB ^a	MAXCAPacity	Dev class	
16	MOUNTLimit		
Native	DATAFormat	Stg pool	
<File system size GB>/<8> -2	MAXSCRatch		
90	REClaim		
Do not set this option.	DEDUPlicate		
0	REUsedelay		
No	COMPression	REGister Node	
16	MAXNUMMP		
No	compression	Option file	IBM Spectrum Protect Client
No	deduplication		
Do not set this option.	encryptiontype		
10	resource utilization		

- a. The MAXCAPacity=8 GB suggestion is no longer relevant when using ProtecTIER PGA V3.3.3 or later. From this code level and higher there is no suggestion to limit the backup file size, and in fact the bigger the better.

Tips:

- ▶ You should not use the **define volume** command. Use the IBM Spectrum Protect server to handle the file creation by setting the **MAXSCRatch** parameter of **define stgpool**.
- ▶ Running the **define volume** command might affect ProtecTIER performance.
- ▶ When you run the **define stgpool** command, you should use the **DATAFormat=native** parameter. Do not use **nonblock**, because it might adversely affect the performance and the deduplication ratio.
- ▶ When you work with the ProtecTIER product, you should use the **devclass** file, but not the **devclass** disk.



Symantec NetBackup and BackupExec

This chapter describes the suggested settings and procedural information to integrate the IBM ProtecTIER product with Symantec NetBackup (NetBackup) environments. The suggested configurations and results can vary, so be sure that you review the configuration with a NetBackup specialist to determine the best configuration for your environment. This chapter also briefly describes Symantec BackupExec in a File System Interface (FSI) environment.

The ProtecTIER product can be deployed as a Virtual Tape Library (VTL) or FSI to NetBackup. This chapter describes NetBackup with VTL and FSI.

This chapter contains the following topics:

- ▶ NetBackup overview
- ▶ General suggestions for NetBackup
- ▶ NetBackup in a VTL environment
- ▶ NetBackup in an FSI environment
- ▶ Symantec BackupExec in an FSI environment

Note: ProtecTIER GA Version 3.4 was released with only the Virtual Tape Library (VTL) interface support. File System Interface (FSI) support was added to ProtecTIER PGA 3.4 Version. For details, see the announcement letter:

http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/6/872/ENUSA P16-0076/index.html&lang=en&request_locale=en

14.1 NetBackup overview

Symantec NetBackup is an Open Systems Enterprise backup software solution. Its architecture has the following three main building blocks:

Clients	The systems with the data that require backing up.
Media servers	The systems that are connected to the clients used for the backup devices. They can be considered as the data mover in the NetBackup environment, and require specific bandwidth considerations. Media servers can also increase performance because they distribute the network load.
Master server	The system that manages and controls the NetBackup environment: the backups, archives, restores, catalog, and so on. Typically, the catalog is also stored on the master server.

Note: The *catalog* is the database that contains information about backup images and configuration information.

In all configurations, at least one media server is required. The media server has access to and manages the storage unit. The master and media servers can be installed on the same hardware. Several media servers can be installed. Each media server controls and manages its own data. NetBackup clients write the backup data to a media server over local area network (LAN) and Internet Protocol (IP), but the client and media server can be installed on the same hardware.

Note: A suggestion is for the master server to be installed on a dedicated system. However, in small environments it can be installed to act as both master and media server.

In general, a media server uses its own storage unit. A storage unit can be either a disk staging device or a tape storage unit. If a tape storage unit is intended to be shared over several media servers, then an additional license, Shared Storage Option (SSO), is required.

The ProtecTIER product can eliminate or reduce the usage of SSO because it can emulate many virtual tape drives, so sharing might no longer be required.

For detailed product information, see the NetBackup web page:

<https://www.veritas.com/product/backup-and-recovery/netbackup>

14.2 General suggestions for NetBackup

The ProtecTIER product can be deployed as a VTL or FSI application when it integrates with NetBackup. Follow the suggestions that are described in this section to ensure optimal deduplication and performance.

Examine the following configuration options in NetBackup and, if necessary, change them to assist with the optimal performance of any ProtecTIER deployment:

- Ensure that you have adequate NetBackup licenses before you do the implementation.
- Verify the ProtecTIER compatibility with NetBackup server hardware and software, operating system, and storage area network (SAN) switches.

- ▶ Disable multiplexing, compression, and encryption.
- ▶ Select **Allow Multiple Data Streams within policies** to enable a backup job to run in simultaneous streams. It is suggested that the number of streams be set to 16 or less.
- ▶ Ensure that your host properties Global Attributes are set appropriately for your configuration.
- ▶ Do not mix disks and tapes on a single Small Computer System Interface (SCSI) bus.
- ▶ For each backup server, create a separate zone for each host bus adapter (HBA) that can access ProtecTIER virtual resources. Use zones that are based on a worldwide port name (WWPN).
- ▶ The Number of Data Buffers is a minimum of 16, and it is suggested to have 32, where the value of the NUMBER_DATA_BUFFER buffer is 32.
- ▶ Although the value of the SIZE_DATA_BUFFER buffer should be at least 262144 (256 KB), 524288 (512 KB) is suggested. On IBM AIX and Linux systems, these buffers can be configured by creating the files on the NetBackup media server (Example 14-1).

Example 14-1 Create the SZ_DATA_BUFFER file

```
-- /usr/opensv/netbackup/db/config/SIZE_DATA_BUFFERS --
--/usr/opensv/netbackup/db/config/NUMBER_DATA_BUFFERS
```

Note: Occasionally the configured parameters in NUMBER_DATA_BUFFER and SIZE_DATA_BUFFER are not supported by the HBA cards, and you might experience failed jobs during backup operations. One solution is to be sure that the HBA front-end cards have the latest driver and firmware versions installed. Check the card's vendor information to perform the proper upgrade procedure for your hardware.

- ▶ The procedure to configure the buffers when using Windows servers varies from the procedure used with UNIX based servers. To configure the NUMBER_DATA_BUFFER and the SIZE_DATA_BUFFER, follow the instructions at the following web page.

https://www.veritas.com/support/en_US/article.TECH18422

Preferred practice

Customers in Linux and UNIX environments that use NetBackup and the PDISC-PRLI (registration request) loop can greatly reduce the amount of noise in the SAN by setting the AVR_SCAN_DELAY entry in the vm.conf file to a high value.

This is a parameter that reflects the number of seconds between normal scan cycles. The minimum for number_of_seconds is 1. The maximum is 180. A value of zero converts to one second. The default value is 15 seconds. The default low value is used to minimize tape mount time but introduces noise to the SAN. It should also be noted that the benefit of the low value helps in physical tape environments where tape mounts take a long time, but irrelevant in VTL systems, where tape operations are instant.

For more information, see the *Symantec NetBackup Administrator's Guide for UNIX and Linux, Release 7.5*:

https://www.veritas.com/support/en_US/article.DOC5157

14.3 NetBackup in a VTL environment

Examine the following configuration options in NetBackup and if necessary, change them to assist with the optimal performance of ProtecTIER VTL deployments:

- ▶ Backup and replication activities can run concurrently; however, a preferred practice is to use separate time windows for the backup and replication operations.
- ▶ Ensure that tape encryption for images that go to the ProtecTIER repository is disabled.
- ▶ Use a block size of 512 kilobytes (KB) for the best deduplication ratio and performance balance.
- ▶ When you create the virtual IBM System Storage TS7650G ProtecTIER Deduplication Gateway (TS7650G) library, select either the data transfer controller (DTC) emulation (for creating a p3000 library) or V-TS3500 (identical to TS3500 emulation). This setting is a requirement from Symantec for NetBackup support.
- ▶ When you use Windows Master or Media servers, consider using NetBackup device drivers rather than native windows device drivers, per Symantec suggestions.
- ▶ For each backup server, create a separate zone for each HBA that can access ProtecTIER virtual resources. Use zones that are based on a worldwide port name (WWPN).
- ▶ If you use more than one HBA on the server to access virtual tape libraries and drives, you might see duplicate virtual tape drive definitions or robotics definitions on your server. To resolve this issue, you can enable zoning, use persistent binding, or ignore the duplicate devices.

14.4 NetBackup in an FSI environment

This section provides steps and preferred practices for configuring and setting up Symantec NetBackup (NetBackup) for backup and restore in an FSI environment. This section also provides NetBackup parameters and settings for best performance with ProtecTIER FSI-Common Internet File System (CIFS) and FSI-Network File System (NFS).

14.4.1 NetBackup in an FSI-CIFS environment

This section provides preferred practices for configuring and setting up NetBackup in an FSI-CIFS environment for backup and restore. This section also provides NetBackup parameters and settings for best performance.

Setting up for backup and restore

The following sections describe the steps for configuring NetBackup for backup and restore for NetBackup 7.0.1 and higher versions in an FSI-CIFS environment. These sections also describe the steps for configuring backup and restore for a version of NetBackup before Version 7.0.1.

Configuring AdvancedDisk sharing with CIFS for NetBackup V7.0.1 and later versions

You can configure AdvancedDisk sharing on Windows by using CIFS. Support for AdvancedDisk sharing on Windows with CIFS is introduced with Symantec NetBackup

V7.0.1. For details about the specific additional steps before the disk pool can be configured, see the following web page:

https://www.veritas.com/support/en_US/article.TECH158427

Also see the topic about implementing the ProtecTIER File System Interface, which is in *IBM System Storage TS7600 with ProtecTIER Version 3.3*, SG24-7968.

Mount point and drive letter: The mount point is set as \\FSI_IP\CIFS_name on ProtecTIER FSI by running the following command:

```
nbdevconfig.exe -createdv -stype AdvancedDisk -storage_server windows-ubmu9k3  
-dv \\9.11.109.130\ForNBU
```

To use AdvancedDisk, mount a network drive and assign it a letter. A drive letter must be assigned in order for AdvancedDisk to work.

Creating a disk pool using a shared volume

This section describes the steps to create a disk pool using a shared volume. When you create a disk pool using a shared volume, ensure that you select only the shared volumes.

Complete the following steps:

1. In the NetBackup Administration Console, click **Media and Device Management**.
2. From the list of wizards in the Details pane, click **Configure Disk Pool** and follow the wizard's instructions.
3. Select the type of Disk Pool you are creating. In this example, the type is AdvancedDisk. Click **Next**.
4. Select the storage server for which you want to create the disk pool and click **Next**. The Disk Pool Configuration wizard shows the list of mount points that can be used.
5. Select the shared volumes that you want to use. Click **Next**.

Important: Ensure that you select from the list only the shared volumes that you created previously to work with ProtecTIER FSI.

6. In the Disk Pool Properties window, set **Limit I/O streams**. By setting the Maximum I/O streams option on the disk pool, you can limit the total number of jobs that access the disk pool concurrently, regardless of the job type.

Storage unit configurations

Storage units that are used with disk type AdvancedDisk are based on the disk pool rather than individual disks. To determine whether storage units exist for the disk pool, open the Administration Console and click **NetBackup Management** → **Storage** → **Storage Units**.

Creating policy and choosing the newly created storage unit

To create a policy, complete the following steps:

1. In the NetBackup Administration Console, expand **NetBackup Management** → **Policies**.
2. Right-click **Policies** and click **New Policy**.

3. In the Change Policy window (Figure 14-1), select the newly created storage unit from the **Policy storage** drop-down menu.

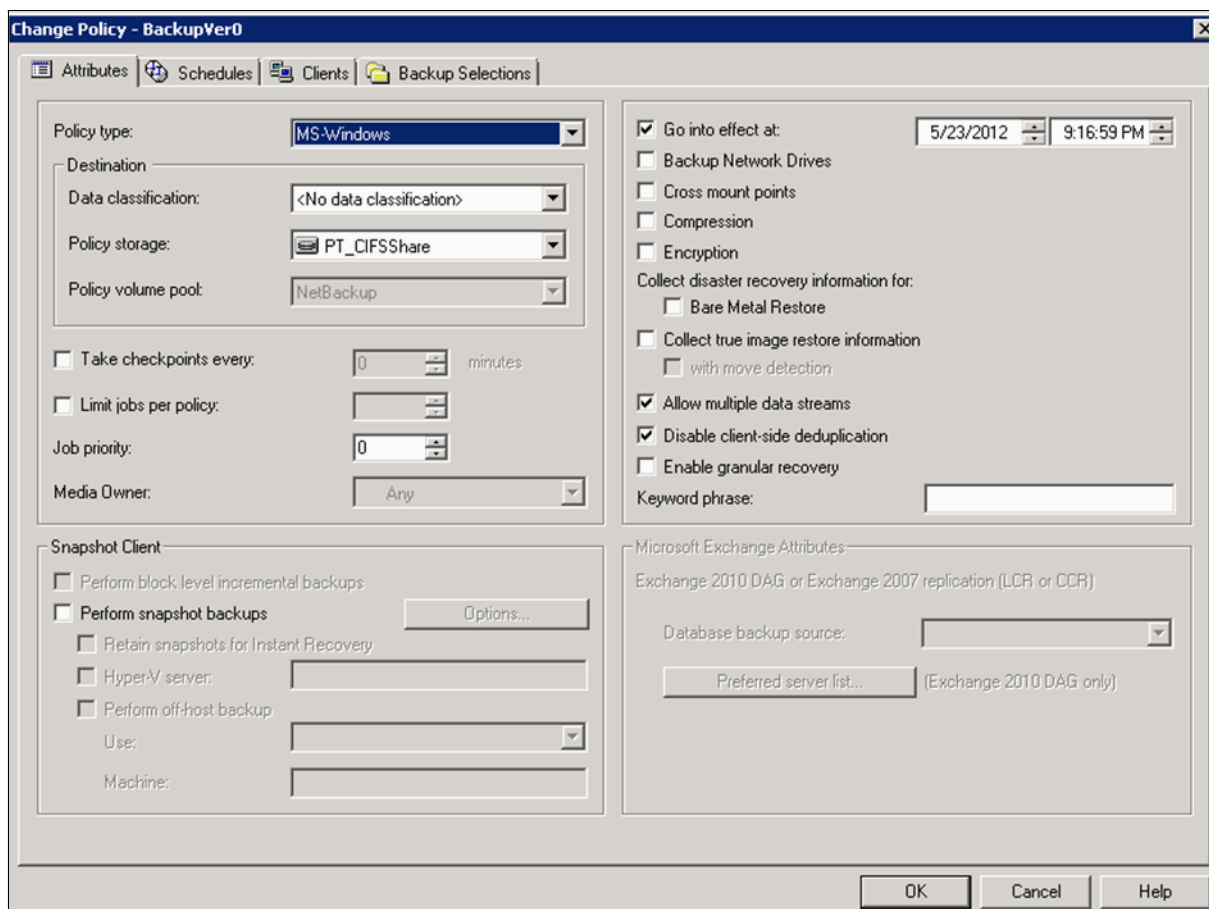


Figure 14-1 Change policy - disable encryption and compression¹

4. Specify the date and time that your policy goes into effect and select **Go into effect at**; and then clear **Compression** and **Encryption** check boxes.

Note: Compression, Encryption, and client-side deduplication should be disabled, as shown in the figure.

5. Also use the **Schedules**, **Clients**, and **Backup Selections** tabs to select more options.

Performing automatic backups or manual backups

This section provides the steps that are needed to perform automatic or manual backups by using NetBackup in your ProtectTIER FSI-CIFS environment. Complete the following steps:

1. In the NetBackup Administration Console, expand **NetBackup Management** → **Policies**.
2. Right-click the policy name in the left pane.
3. To do a manual backup, click **Manual Backup**.
4. To do a manual backup *and* activate the policy, select the **Go into effect at** check box in the **Policy Attributes** tab. The policy must be active for NetBackup to use the policy.

¹ Symantec, Reprinted by permission.

Important: The **Go into effect at** attribute specifies when the policy can begin to schedule backups. For example, if today is Monday and you enter Wednesday at 12:01 AM, the policy does not run until that time or later. Use the **Go into effect at** attribute to configure a series of policies in advance of when the policies need to become active.

Performing restore operations

To perform a restore operation by using NetBackup in your ProtecTIER FSI environment, complete the following steps:

1. Start the Backup, Archive, and Restore client interface. Click **Select for Restore**, and select the **System State** check box.
2. In the Actions menu, click **Start Restore of Marked Files**.
3. In the Restore Marked Files dialog box, click **Restore everything to its original location and Overwrite the existing file**.

Attention: Do not redirect the System State restore to a different host. System State is computer-specific. Trying to restore it to a different computer can result in an unusable system.

4. Click **Start Restore**.

Parameters for best performance with ProtecTIER FSI-CIFS

This section provides guidelines and parameters for best performance with ProtecTIER FSI-CIFS. It describes network file considerations, and the usage of maximum I/O streams with disk pools for optimum performance. You can adjust the Network Buffer Size, Data Buffer Size, and Number of Data Buffers for performance enhancements.

Multiplexing: For optimum performance in your Protectier FSI-CIFS environment, do not enable multiplexing.

Setting compression, encryption, and client-side deduplication

Compression, encryption, and client-side deduplication should be disabled, as described in “Creating policy and choosing the newly created storage unit” on page 211.

Network file system considerations

The AdvancedDisk storage implementation presents all mounted file systems as disk volumes to NetBackup, including network file systems (such as NFS and CIFS). For more information about using network file systems, see the this web page:

https://www.veritas.com/support/en_US/article.TECH158427

Using Maximum I/O streams with disk pools

The Maximum Concurrent Jobs setting on the Storage Units menu limits the number of backup or write jobs using each storage unit, but does not limit the number of restore or read activities that might be going on concurrently with the write activity. This situation can cause unexpected I/O contention on the disk pool. For more information, see the this web page:

https://www.veritas.com/support/en_US/article.TECH158427

Changing the disk pool I/O streams option

To update the default number of I/O streams per volume, complete the following steps:

1. In the NetBackup Administration Console, expand **Media and Device Management** → **Devices** → **Disk Pools**.
2. Select the disk pool that you want to change in the Details pane.
3. Click **Edit** → **Change**, and the Change Disk Pool window opens, where you can change the value in the Maximum I/O Streams pane.

Setting the maximum jobs per client

To set or change the maximum jobs per client, from the NetBackup Administration Console on the master server, complete the following steps:

1. Expand **NetBackup Management** → **Host Properties** → **Master Server**.
2. Open the host properties of the master server.
3. Click **Global Attributes**.

The Global Attributes properties apply to the currently selected master servers. The Global Attributes properties affect all operations for all policies and clients. The default values are adequate for most installations, but can be changed.

4. Set **Maximum jobs per client**. The Maximum jobs per client property applies to all clients in all policies.

This property specifies the maximum number of backup and archive jobs that NetBackup clients can perform concurrently.

14.4.2 NetBackup in an FSI-NFS environment

This section provides preferred practices for configuring and setting up NetBackup in an FSI-NFS environment for backup and restore. It also suggests NetBackup parameters and settings for best performance.

Creating a disk pool using a shared volume

This section describes the steps to create a disk pool using a shared volume. When you create a disk pool using a shared volume, be sure to select only the shared volumes that you configured previously to work with ProtecTIER. Complete the following steps:

1. In the NetBackup Administration Console, click **Media and Device Management**.
2. From the list of wizards in the Details pane, click **Configure Disk Storage Servers**.
3. Select the type of disk storage you are configuring. The selected disk storage type should be AdvancedDisk. Click **Next**.
4. Select the storage server and media server for which you want to create the disk pool. Click **Next**.
5. The wizard guides you through the steps that are required to create a disk pool and a storage unit that uses the newly created disk pool. Click **Next**.
6. Select the type of Disk Pool you are creating. The selected disk pool type should be AdvancedDisk. Click **Next**.
7. Select the storage server for which you want to create the disk pool. Click **Next**.
8. The Disk Pool Configuration wizard shows the list of mount points that can be used. Select the FSI-NFS mounted point to add to the disk pool. Click **Next**.

Note: Select *only* the ProtecTIER mounted folder. Preferably, one mount should be defined per disk pool.

9. In the Disk Pool Properties window, set the **Limit I/O streams** parameter to 16. This parameter limits the total number of I/O streams that access the disk pool concurrently, regardless of the job type. The ProtecTIER suggestion is for a total of 16 I/O streams. Click **Next**.
10. The disk pool configuration is completed. Click **Next**.
11. Select **Create a storage unit that uses the disk pool** to configure a storage unit.

Storage unit configurations

Storage units that are used with disk type AdvancedDisk are based on the disk pool rather than individual disks. The Disk Pool Configuration wizard enables you to create a storage unit as part of the creation of the disk pool. To determine whether storage units exist for the disk pool, open the Administration Console and click **NetBackup Management** → **Storage** → **Storage Units**.

To create a storage unit, complete the following steps:

1. In the NetBackup Administration Console, select **NetBackup Management** → **Storage** → **Storage Units**.
2. Click **Actions** → **New** → **Storage Unit**.
3. Complete the fields in the New Storage Unit dialog box.
4. Set the maximum concurrent jobs to 16 for a better performance.

Note: When using a ProtecTIER version older than V3.3.3, set the maximum fragment size parameter to 8192 MB (8 GB). When using ProtecTIER PGA V3.3.3 or later, the maximum fragment size suggestion is no longer relevant. From this code level and later there is no suggestion to limit the backup file size, and in fact a bigger size is better.

Creating a policy and choosing the newly created storage unit

To create a policy, complete the following steps:

1. In the NetBackup Administration Console, expand **NetBackup Management Policies**.
2. Right-click **Policies** and click **New Policy**.
3. From the Change Policy window, select the newly created storage unit from the **Policy storage** drop-down menu. **Schedules**, **Clients**, and **Backup Selections** can be selected from the corresponding tabs.
4. *Do not select:* Under the Attributes tab, ensure that the following check boxes are *not* selected:
 - Compress
 - Encrypt
5. *Do select:* Under the Attributes tab, ensure that the following check boxes *are* selected:
 - Disable client-side deduplication
 - Allow multiple data streams
6. Under the Schedules tab, double-click every relevant schedule and ensure that the value of **Media Multiplexing** in Attributes tab is set to 1 (one).

Creating configuration files on the media server

These files adjust the responding rate between a media server and ProtecTIER FSI-NFS, and are needed for the ProtecTIER FSI-NFS to work correctly.

To create the configuration files, complete the following steps:

1. Run the following commands:

```
touch /usr/opensv/netbackup/db/config/DPS_PROXYNOEXPIRE
echo "1800" > /usr/opensv/netbackup/db/config/DPS_PROXYDEFAULTSENDTMO
echo "1800" > /usr/opensv/netbackup/db/config/DPS_PROXYDEFAULTRECVTMO
```

Note: The 1800 value should be the only value in these two files. In general, allowed values are in the range of 10 - 3600. For more information about this value, see this web page:

https://www.veritas.com/support/en_US/article.TECH156490

2. Do one *either* of the following operations:

- Restart **nbrmms** (NetBackup Remote Manager and Monitor Service) on the media server by running these commands:

```
kill nbrmms
/usr/opensv/netbackup/bin/nbrmms
```

- Stop and restart all services on the MSDP media server by running these commands:

```
/usr/opensv/netbackup/bin/goodies/netbackup stop
/usr/opensv/netbackup/bin/goodies/netbackup start
```

Setting the maximum jobs per client

To set or change the maximum jobs per client, from the NetBackup Administration Console on the master server, complete the following steps:

1. Expand **NetBackup Management** → **Host Properties** → **Master Server**.
2. Open the host properties of the master server.
3. Click **Global Attributes**.

Note: The Global Attributes properties apply to the currently selected master servers and affect all operations for all policies and clients. The default values are adequate for most installations, but can be changed.

4. Set **Maximum jobs per client**. The Maximum jobs per client property applies to all clients in all policies. This property specifies the maximum number of backup and archive jobs that NetBackup clients can perform concurrently.

Tip: For best performance, a suggestion is to set the Maximum jobs per client to 16.

Performing a backup

This section provides the steps that are needed to perform automatic or manual backups by using NetBackup in your ProtecTIER FSI-NFS environment. Complete the following steps:

1. In the NetBackup Administration Console, expand **NetBackup Management** → **Policies**, and right-click the policy name in the left pane.
2. To do a manual backup, click **Manual Backup**.
3. To do a manual backup *and* to activate the policy, select the **Go into effect at** check box in the **Policy Attributes** tab. The policy must be active for NetBackup to start the policy.

Important: The **Go into effect at** attribute specifies when the policy can schedule backups. For example, if today is Monday and you enter Wednesday at 12:01 AM, the policy does not run until that time or later. Use the **Go into effect at** attribute to configure a series of policies in advance of when the policies need to become active.

Performing a restore

This section provides the steps that are needed to perform a restore by using NetBackup in your ProtecTIER FSI-NFS environment. Complete the following steps:

1. Start the Backup, Archive, and Restore client interface. Click **Select for Restore**, and then select the **System State** check box.
2. From the Actions menu, click **Start Restore of Marked Files**.
3. From the Restore Marked Files dialog box, select the option that suits your needs for example: **Restore everything to its original location** and **Overwrite the existing files**.

Attention: Do not redirect System State restore to a different host. System State is computer-specific. Restoring it to another computer can result in an unusable system.

4. Click **Start Restore**.

Parameters for best performance with ProtecTIER FSI-NFS

Table 14-1 summarizes the required and suggested settings for ProtecTIER preferred practices with NetBackup.

Table 14-1 Summary of settings for ProtecTIER with NetBackup

Component	Parameter	Value
Disk pool definition	Disk pool type	AdvancedDisk
	Limit I/O streams	16 ^a
Storage unit definition	Maximum fragment size	8192 MB (8 GB) ^b
	Maximum concurrent jobs	16 ^a
Policy definition	Compress	Disabled (Not Checked)
	Encrypt	Disabled (Not Checked)
	Allow multiple data streams	Enabled (Checked)
	Disable client-side deduplication	Enabled (Checked)
	Media Multiplexing	1
Global Attributes	Maximum jobs per client	16 ^a

a. To learn about the maximum number of streams that are supported, see 14.2, “General suggestions for NetBackup” on page 208.

b. When using a ProtecTIER version older than V3.3.3 set the maximum fragment size parameter to 8192 MB (8 GB). When using ProtecTIER PGA V3.3.3 or later the maximum fragment size suggestion is no longer relevant. From this code level and later there is no suggestion to limit the backup file size, and in fact the bigger the better.

Network file system considerations

The AdvancedDisk storage implementation presents all mounted file systems as disk volumes to NetBackup, including network file systems (such as NFS). For more information about using network file systems, go to the following web page:

https://www.veritas.com/support/en_US/article.TECH158427

14.5 Symantec BackupExec in an FSI environment

This section provides steps and preferred practices for configuring and setting up Symantec BackupExec (BackupExec) for backup and restore. It also provides BackupExec parameters and settings for best performance with ProtecTIER FSI.

Use the steps in this section to create a backup-to-disk folder and also create a backup job and a restore job.

Creating a backup-to-disk folder

To create a backup-to-disk folder, complete the following steps:

1. Open the BackupExec GUI and go to **Devices**.
2. Start the Configure Devices Assistant.
3. Create a backup-to-disk folder.
4. Configure the parameters of the folder. Figure 14-2 shows the preferred parameters (full path name, and maximum concurrent jobs).




Figure 14-2 Create a backup-to-disk folder

Note: When using a ProtecTIER version older than V3.3.3, set the maximum fragment size parameter to 8192 MB (8 GB). When using ProtecTIER PGA V3.3.3 or later the maximum fragment size suggestion is no longer relevant. From this code level and later there is no suggestion to limit the backup file size, and in fact the bigger the better.

Creating the backup job to perform backups

To create a backup job to perform backups, complete the following steps:

1. Open the Job Monitor window.
2. Start the backup wizard by clicking the **Backup** icon.
3. Use the Backup wizard suggested parameters. The Backup wizard prompts you with **What do you want to back up?** Click **This computer** and then click **Next**.
4. Click **Full backup job** and then click **Next**.
5. Choose the time that you want to run this backup at (for example, for a one-time backup, click **Run the full backup now** and then click **Next**).
6. Select **Backup-to-disk folder** and, in the drop-down list, select the folder that you configured (Figure 14-3). Then, click **Next**.

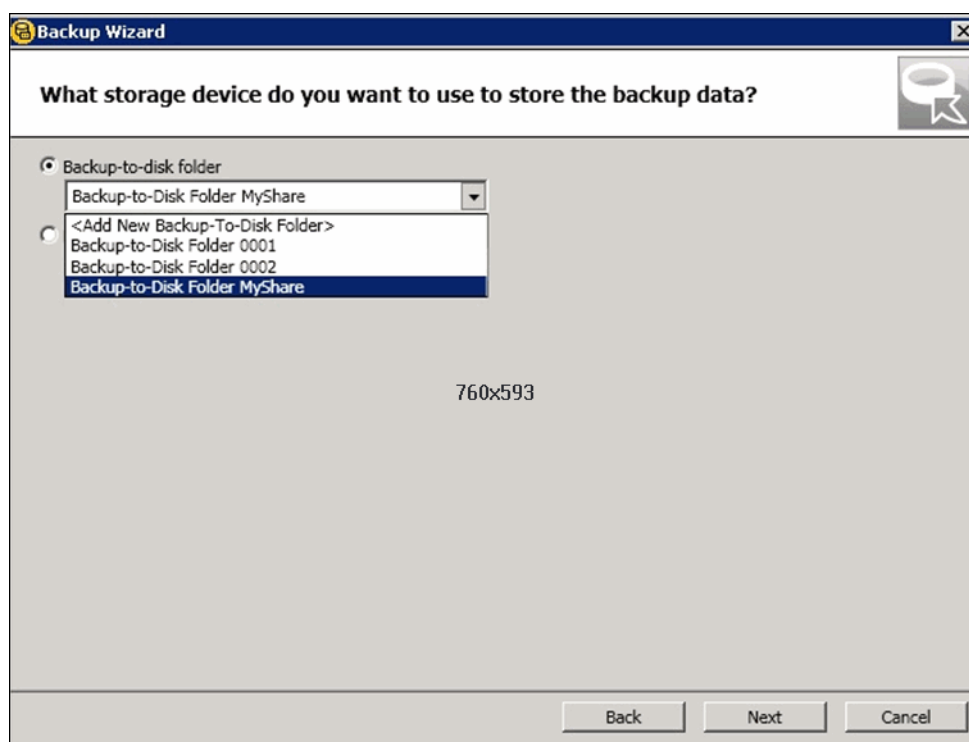


Figure 14-3 Choose backup-to-disk folder from the Backup Wizard window

7. Choose a retention period for keeping backups and click **Next**.
8. Give the backup job a meaningful name and click **Submit**.

Jobs: You can see both the job running (in the Job Monitor view) and the files being written (in the share).

Creating the restore job to perform a restore

To create a restore job to perform a restore, complete the following steps:

1. Open the Job Monitor window.
2. Start the restore wizard by clicking the **Restore** icon.
3. Follow the Restore wizard instructions to create the restore job to perform a restore.



EMC NetWorker

EMC NetWorker (NetWorker; formerly known as Legato NetWorker) is an enterprise backup application that provides central backup management for various applications on various operating systems, with various backup methods to various types of storage media.

The IBM ProtecTIER product can be deployed as a Virtual Tape Library (VTL) or a File System Interface-Common Internet File System (FSI-CIFS) share and FSI-Network File System (NFS) export to NetWorker to enhance its data protection ability.

This chapter describes suggested settings and procedures to integrate ProtecTIER VTL and FSI in a NetWorker environment in order to achieve optimal backup throughput and the factoring ratio of the ProtecTIER system. The suggested configurations and results might vary in different environments. Review the configuration with your NetWorker specialist for the best configuration that fits into your environment.

This chapter contains the following topics:

- Overview
- EMC NetWorker in a VTL environment
- EMC NetWorker in an FSI environment

Note: ProtecTIER GA Version 3.4 was released with only the Virtual Tape Library (VTL) interface support. File System Interface (FSI) support was added to ProtecTIER PGA 3.4 Version. For details, see the announcement letter:

http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/6/872/ENUSA P16-0076/index.html&lang=en&request_locale=en

15.1 Overview

EMC NetWorker is a centralized and automated backup and recovery product for heterogeneous enterprise data. The NetWorker Server hosts the configuration information and NetWorker databases that track the backups and volumes. It runs on all major operating systems, such as AIX, Linux, Windows, Oracle Solaris, and HP-UX. Apart from server software, the NetWorker server always has the NetWorker Storage Node and NetWorker Client installed.

The NetWorker Storage Node is the host that has direct access to tape or disk media. It uses this access to read and write data to storage devices. It sends tracking information only to the NetWorker Server. The NetWorker Client is installed on the customer's servers to generate save sets, and sends them to or retrieves them from the storage node. Different clients are available for the integration of special applications, such as NetWorker for IBM DB2.

A NetWorker *datazone* consists of one NetWorker Server, several NetWorker Storage Nodes, and several NetWorker Clients. Figure 15-1 illustrates the integration of NetWorker components into a datazone.

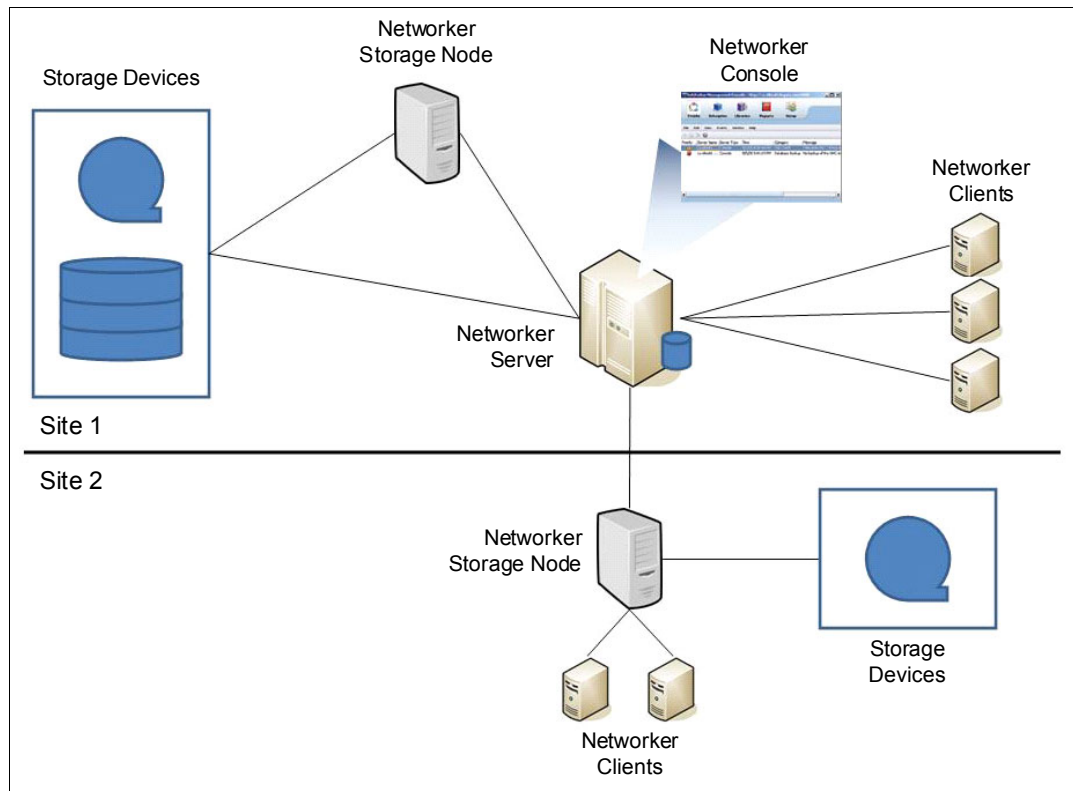


Figure 15-1 NetWorker components in a NetWorker datazone

NetWorker writes data to volumes and identifies the volumes with specified volume IDs. A volume is a physical storage media, such as magnetic tape, an optical disk, or a file system. The volume ID can be changed only when a volume is relabeled. NetWorker organizes the volumes into media pools. A ProtecTIER virtual cartridge or FSI share is seen as a volume in NetWorker and is sorted into specific media pools.

15.2 EMC NetWorker in a VTL environment

The ProtecTIER product can be deployed as a VTL when you integrate it with NetWorker. In traditional tape backups, tape drive sharing can be costly with NetWorker because of the need for an additional license for dynamic drive sharing. With VTL, many virtual tape drives can be created and dedicated to different storage nodes without wasting tape resources.

15.2.1 General suggestions

Follow these general suggestions to achieve optimum performance and factoring ratio of the ProtecTIER product with NetWorker:

- ▶ Ensure that you have adequate NetWorker licenses before you do the implementation.
- ▶ Check the compatibility of ProtecTIER with the NetWorker Server hardware and software, operating system, and storage area network (SAN) switches.
- ▶ Review the configuration with your NetWorker specialist to avoid an adverse performance effect to your environment.
- ▶ Disable client compression and encryption.
- ▶ Whenever possible, use the same Network Time Protocol (NTP) or time server for the ProtecTIER and NetWorker servers to ease maintenance and troubleshooting tasks.

15.2.2 Suggestion if a ProtecTIER server is used as a VTL

Follow these suggested configurations to achieve optimum performance and the optimum factoring ratio of the ProtecTIER server in your VTL environment:

- ▶ Ensure that you have the VTL option enabled or an appropriate number of autochanger and library slot licenses for NetWorker.

Attention: Two licenses are required for EMC NetWorker to function: the *Autochanger Module* (which must be compatible with the number of slots that are created in the ProtecTIER server) and *Virtual Tape Library Capacity* (for the number of terabytes that exist in the ProtecTIER server). Without these two licenses, NetWorker cannot function correctly and you might receive the following error:

The VTL is not properly licensed. Make sure the VTL is properly enabled and authorized (VTL frame enabler needed). See your EMC representative.

- ▶ Consider creating dedicated tape drives for each storage node.
- ▶ Use ProtecTIER logical unit number (LUN) masking if multiple storage nodes share virtual tape drives. For more information, see 6.3, “LUN masking for VTL systems” on page 87.
- ▶ Add virtual tape drives gradually.
- ▶ Follow the NetWorker device driver guidelines in Chapter 6, “Host attachment considerations for Virtual Tape Library” on page 79.
- ▶ Use persistent binding for virtual tape drives. For more information, see , “When the multipath feature is enabled, it defaults to running in automatic mode. The automatic mode automatically scans for all paths for each library robot at each tldcd daemon start, requiring no additional setup.” on page 85.
- ▶ Disable common device interface (CDI) on all virtual tape drives.
- ▶ Use 512 kilobytes (KB) I/O for the virtual tape drives to ensure a good factoring ratio.

- ▶ Disable multiplex by setting parallelism to 1 (one) on all virtual drives. You can do so by editing the properties of each virtual tape drive to set the target sessions and maximum sessions value to 1.
- ▶ Enable the Auto Media Management for virtual tape libraries to support space reclamation. For details about Auto Media Management, see the Media Management section in the *EMC NetWorker Administrator Guide*, which is available at the NetWorker Information Hub website:
<http://nsrd.info/docs.html>
- ▶ Increase the NetWorker media multiplexor daemon (**nsrmmmd**) polling interval and **nsrmmmd restart** interval if you see many **nsrmmmd** restart activities during the backup session. Increase the value gradually, such as **nsrmmmd** polling interval=6, restart interval=4, followed by polling interval=9, restart interval=6, and so on.

15.3 EMC NetWorker in an FSI environment

This section provides steps and preferred practices for configuring and setting up EMC NetWorker (NetWorker) for backup and restore. It also provides NetWorker parameters and settings and suggestions for best performance with ProtecTIER FSI. The suggestions apply to FSI-CIFS and FSI-NFS unless stated otherwise.

15.3.1 Creating a Windows user for EMC NetWorker

To configure EMC NetWorker (NetWorker) for Windows with ProtecTIER FSI-CIFS, you must create a Windows User for NetWorker. Follow the steps in this section to configure NetWorker and run a simple backup job.

Changing NetWorker Services logon properties to the Windows user

Complete the following steps:

1. In the Services window, right-click **NetWorker Remote Exec Service** and select **Properties**.
2. The NetWorker Remote Exec Service Properties window opens. Click the **Log on** tab, click **This account**, enter the Windows user account that you created. Click **OK**.
3. You receive a message that the changes do not take effect until the service stops and restarts. Click **OK**.
4. Repeat these steps with the NetWorker Backup and Recover Server service.
5. Right-click **NetWorker Remote Exec Service** and click **Restart** to restart the service. This action also restarts the NetWorker Backup and Recover Server service.

Creating a user in NetWorker

Create a user for NetWorker to access your CIFS share:

1. Open the NetWorker Management Console.
2. From the Console window, click **Setup**.
3. In the left pane, right-click **Users** and click **New**. The Create User dialog box opens.
4. Enter the user name, the appropriate role information (for example, Console Application Administrator, Console User, and Console Security Administrator), and Password to create a user, and then click **OK**.

15.3.2 Setting up for backup and restore

The following section details the steps for setting up your ProtecTIER FSI-CIFS and FSI-NFS environment to perform backups and restores by using EMC NetWorker.

15.3.3 General configuration suggestions

The following general configuration suggestions can help you configure a NetWorker server for use with ProtecTIER FSI shares:

- ▶ Consider a dedicated network for backup. To review the FSI Guidelines for network preparation, see Chapter 5, “ProtecTIER File System Interface: General introduction” on page 65.
- ▶ Ensure that you have the NetWorker DiskBackup option enabled.
- ▶ The NetWorker storage node supports only advanced file type devices (AFTD) for FSI shares.
- ▶ Use the Universal Naming Convention (UNC) path for device configuration for CIFS shares and references to the FSI mount point on Linux.
- ▶ Whenever possible, create different ProtecTIER file systems for different storage nodes.
- ▶ NetWorker does not span ongoing saves across multiple AFTD devices. NetWorker suspends all saves being written to AFTDs when the device is full until more space is made available on the device. Plan the AFTDs correctly to have sufficient space that is based on the backup size and retention period.
- ▶ Do not share a single FSI share across multiple storage nodes, and create one AFTD in each FSI share on the storage node.
- ▶ You can configure NetWorker to split a single save set into multiple sessions in one AFTD. In the AFTD device, each save set is stored as a separate file with a unique saveset ID regardless of its parallelism. Therefore, the effect to the deduplication should be minimal.
- ▶ When you install the NetWorker server on Windows, create a dedicated user (Workgroup) or add the user (Active Directory) to the Windows Admin and Backup Operator Group of the NetWorker Windows Server. Change the NetWorker Remote Exec server and NetWorker Backup/Restore Server services to be started by the newly created user with administrator rights. The same user must be added as a user of the CIFS shares in the ProtecTIER server with write enabled permission.
- ▶ When you install the NetWorker server or storage node on Linux, the read/write permission to the NFS export is granted based on host Internet Protocol (IP) addresses and not on user-level granularity.

For information about AFTD, see Device Operations in *EMC NetWorker Administrator Guide*.¹

¹ NetWorker documentation is at the NetWorker Information Hub: <http://nsrd.info/docs.html>

Creating the device

To create a device in the NetWorker Administration window, use the Device Configuration wizard and complete the following steps:

1. In the NetWorker Administration window, right-click **Devices** and click **New Device wizard**.
2. In the Device Configuration wizard, click **Advanced File Type Device (AFTD)** and then click **Next**.
3. The Select Storage node window opens. Click **Enter device path**, and in the box that opens, enter the address to the FSI share. Example 15-1 shows the configuration for a FSI-CIFS share by using the Universal Naming Convention (UNC) path.

Example 15-1 Create a device on an FSI-CIFS share

\\10.0.20.134\networker

For Linux systems, the path definition is similar to Example 15-2.

Example 15-2 Create a device on an FSI-NFS export

/mnt/fsi_shares/networker

Note: Before creating the AFTD, you must mount the export on the Linux file system.

4. In the next window, click **Backup pool type** and click **Next**.
5. In the next window, set the Number of Target Sessions (number of client's connections in parallel) and click **Next**.
6. The next window should display a summary of the information. Click **Configure**.
7. After the configuration completes, you should see a confirmation window. Click **Finish**.

15.3.4 Setting the information to be backed up

Now that you have set up your device, you must define which information you are backing up to this device. To set up the information to be backed up, complete the following steps:

1. In the NetWorker Administration window, click **Configuration** and then click **Groups → Default**.
2. Right-click the existing group and select **Client Backup Configuration → Modify**.
3. In the first window that opens, do not make any changes. Click **Next**.
4. The networker data window opens, and shows the operating system, version, and so on. Verify this information and click **Next**.
5. The next window opens and shows all the directories for all the units that are available. Choose the information that you want to back up and click **Next**.
6. The next window opens and shows the times for the Browse and Retention policy. Make your required changes and click **Next**.
7. The next window opens. Although you can change the Group for this policy, for this example, do not change any data. Click **Next**.

8. Although in the next window, you can choose another Storage Node for Backup and Restore, for this example, do not change anything. Click **Next**.
9. A summary window opens. If all the information is correct, click **Modify**.
10. A confirmation window opens. Click **Finish**.

15.3.5 Setting the time for the backup

Either set the time for your backup or start the backup now by completing these steps:

1. In the NetWorker Administration window, click **Configuration**.
2. Click **Groups**, right-click **Default Group**, and then click **Properties**.
3. In the Properties dialog box, you can change the time in the Setup pane. In the drop-down box menu, you can either choose whether the backup starts automatically or at a later time. Choose your option and click **OK**.

15.3.6 Performing a restore

To perform a restore, complete the following steps:

1. Open the NetWorker User.
2. Select the server that backed up the information, and click the **Recover** icon.
3. Select the information that you want to restore, and click the **Recover** icon to start the restore operation.

15.3.7 Parameters for best performance with ProtecTIER FSI

This section describes how to get in to diagnostics mode to run diagnostic tests on devices. This section also describes how to turn off compression and encryption in NetWorker.

Setting compression and encryption to none

To set compression and encryption to none, complete the following steps:

1. In the NetWorker Administration window, click **Devices**.
2. Click **View** → **Diagnostic Mode** on the toolbar. The Devices window opens.
3. Click **Devices** in the navigation tree (on the left). The Devices detail table is displayed.
4. Double-click the mounted device from the detail table. The Device Properties window opens.
5. In the Device Properties window, click the **Cloud** tab (Figure 15-2 on page 228).

The Compression and Encryption settings are under the Options pane.

6. Set Compression and Encryption to **none**.

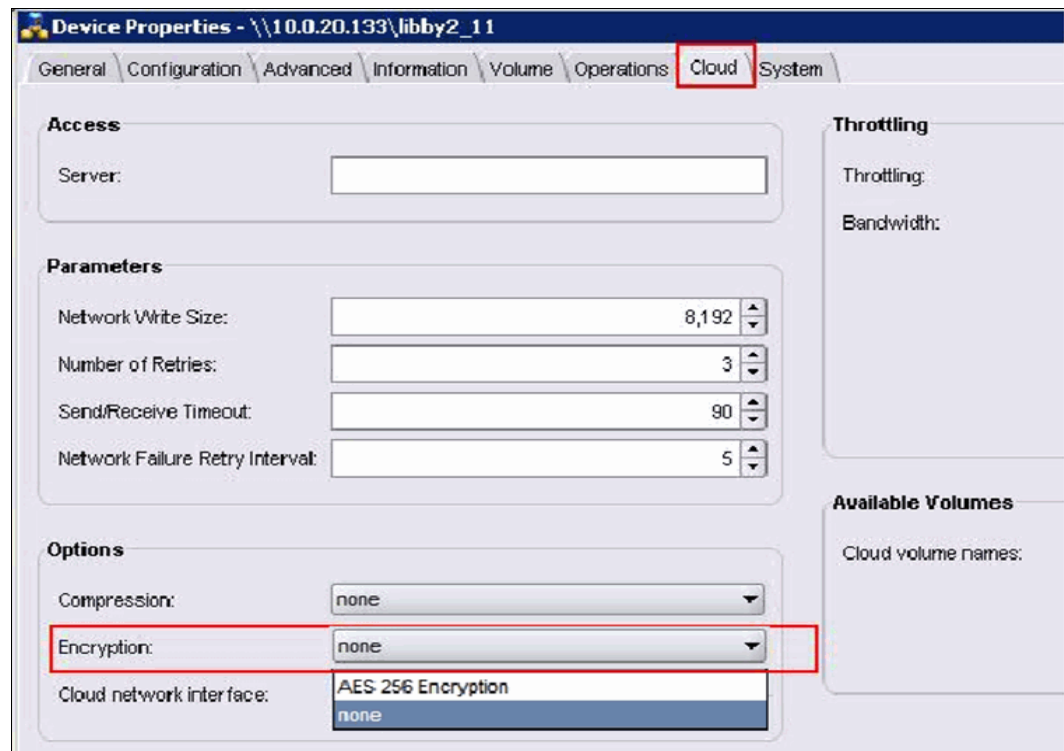


Figure 15-2 Device Properties window: disable encryption



HP Data Protector

This chapter describes the settings and parameters to modify in HP Data Protector (HPDP) in IBM ProtecTIER environments so that maximum performance can be enabled.

Most of the changes and settings are related to the Virtual Tape Library (VTL) mode of the ProtecTIER product, although the File System Interface-Common Internet File System (FSI-CIFS) mode includes straightforward steps to enable deduplication in your HPDP storage that is provisioned by ProtecTIER.

This chapter contains the following topics:

- ▶ HP Data Protector with ProtecTIER
- ▶ HP Data Protector in a VTL environment

Note: ProtecTIER GA Version 3.4 was released with only the Virtual Tape Library (VTL) interface support. File System Interface (FSI) support was added to ProtecTIER PGA 3.4 Version. For details, see the announcement letter:

http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/6/872/ENUSA/P16-0076/index.html&lang=en&request_locale=en

16.1 HP Data Protector with ProtecTIER

The HPDP supports a wide range of applications, disk and tape devices, data protection tools, and any hypervisor, typically from one centralized console. The HPDP also provides protection for Microsoft Exchange, Microsoft SharePoint, Microsoft SQL Server, Systems Applications and Products (SAP), Oracle, and many more business-critical applications in enterprise environments.¹

The IBM ProtecTIER deduplication solution is compatible with HPDP installed as a backup application. The deployment of ProtecTIER with HPDP is supported by IBM as a VTL or FSI-CIFS.

This section provides principal configuration steps and parameters that are suggested to be enabled in HPDP to achieve good deduplication results, optimal backup and restore performance, and system stability.

16.1.1 HP Data Protector architecture with ProtecTIER

The HPDP infrastructure is based on a *cell* layout, as shown in Figure 16-1. The cell is a network environment, typically in the data center or at a remote branch office, that consists of a *Cell Manager*, *client systems*, and *devices*. The Cell Manager is a central management point where HPDP is installed and operated.

The Cell Manager incorporates the HPDP *Internal Database* (IDB) that maintains the references about all backup files from clients. After the Cell Manager is installed and configured, you might want to add systems to be backed up (protected). These systems become client systems.

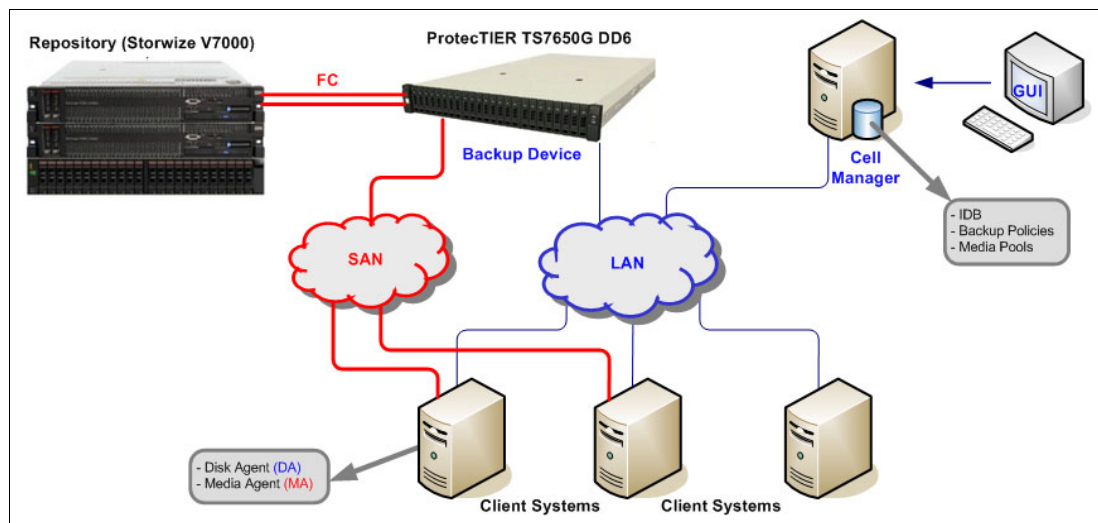


Figure 16-1 The cell diagram of HP Data Protector with ProtecTIER as a VTL or FSI

When client systems send backup files to the HPDP, it stores them on *media* in backup *devices*. In the environments with IBM ProtecTIER deduplication solutions, these media are represented by virtual tape cartridges. Devices are depicted by virtual tape libraries that are provisioned by a ProtecTIER server (VTL mode of the ProtecTIER product). In the case of a FSI-CIFS installation of the ProtecTIER product, the HPDP accesses the CIFS shares from the ProtecTIER server as a media repository for backups.

¹ Source: <http://www.hp.com>

16.2 HP Data Protector in a VTL environment

The ProtecTIER product can be configured either as VTL or FSI, but not both modes concurrently. This section describes the setup and configuration in a VTL environment for backup and restore.

The Cell Manager is the key system in the cell. It is responsible for the following tasks:

- ▶ Operates the cell from a single control point.
- ▶ Maintains the IDB records with references to back up data, media IDs, and backup sessions.
- ▶ Runs the core HPDP software.
- ▶ Runs Session Manager, which triggers backup and restore jobs and records the results in IDB.

Two core components are on the client systems:

- ▶ Disk Agent, also called Backup Agent, which performs a backup of static files in the operating system. An additional utility, called Application Agent, can back up databases online without disruption to the database services.
- ▶ Media Agent offers an interface to the directly attached or storage area network (SAN)-zoned media devices, including virtual tape drives allocated from your ProtecTIER server. To grant control over the attached Small Computer System Interface (SCSI) tape library to the source systems for direct backup, use General Media Agent.

The following sections provide the key HPDP parameters and configuration steps to improve your backup and restore performance with the ProtecTIER server as a VTL.

The HPDP offers three types of disk-based backup devices:

- ▶ The stand-alone file device is the simplest disk-based backup device. It contains a single slot, to which data can be backed up. The maximum capacity is 2 terabytes (TB).
- ▶ The file jukebox device consists of multiple slots to which you can back up data. Each slot offers a maximum capacity of 2 TB. These volumes must always be created manually.
- ▶ The file library device is similar in concept to the file device class in IBM Tivoli Storage Manager. It has multiple slots that are called file depots, which are created and deleted automatically as required. The maximum capacity of each slot is also 2 TB.

For the FSI type of ProtecTIER operation, the File Library device should be used. Select a slot (file depot) capacity 100 - 300 gigabytes (GB) as the preferred practice for optimal performance.

Tip: When an FSI-enabled ProtecTIER server is used together with HPDP, define the maximum slot (file depot) capacity as 100 - 300 GB. The maximum slot capacity size is based on the complexity of your backup infrastructure and an average daily backup volume. In typical medium to large environments, set the file depot size to 200 GB.

After the capacity is configured, do not change it; doing so leads to inefficient space management in the ProtecTIER/HPDP file library.

16.2.1 Enabling the robotic bar code reader

The IBM ProtecTIER Deduplication Gateway in a VTL configuration emulates an IBM System Storage TS3500 Tape Library with Ultrium LTO3 tape drives. It also emulates the numbering of virtual LTO3 tape cartridges that use virtual bar code labels. This feature enables backup applications to perform an inventory of the tape cartridges in the virtual library the same way a real physical tape library with a bar code reader performs an inventory.

To reduce the time that is needed to perform media inventory in the HPDP Backup Device (in this case, VTL), enable robotic bar code reader support. Use these virtual bar codes as a medium label.

Figure 16-2 shows how to enable a robotic bar code reader by selecting devices from the **Environment** folder in the Devices and Media window. Select the **bar code reader support** and the **Use bar code as medium label on initialization** check boxes.

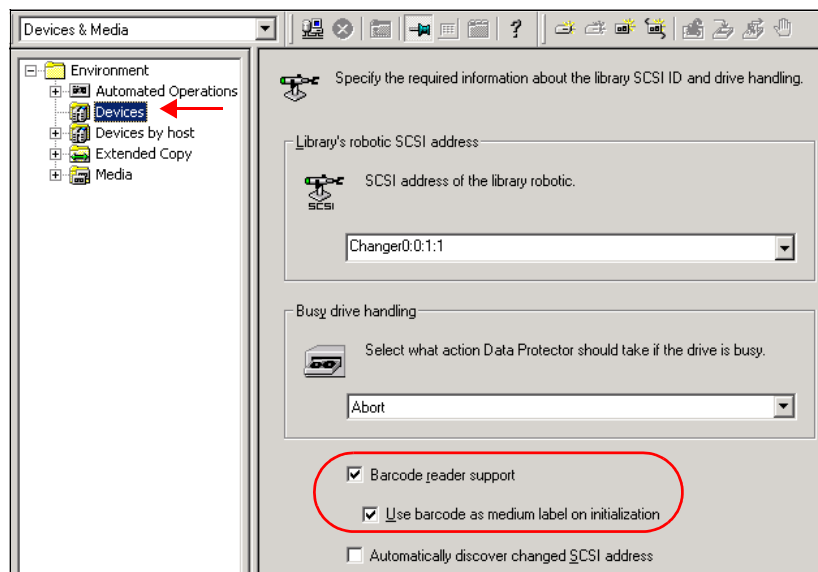


Figure 16-2 Enable robotic bar code reader

On the same window, at the bottom, also enable SCSI reserve/release robotic control to reserve the robotic control only for HPDP operations. This action restricts other systems from working with SCSI reservations on the same target SCSI device, if these systems are zoned to different servers. A typical example is an installation process of a new server, where LAN-free backups are planned and tape zoning is prepared in advance.

16.2.2 Increasing the tape block size

Increase the block size of the tape device from the default value of 64 KB to a minimum of 256 KB and a maximum of 1 MB. The ProtecTIER product supports all values up to a block size of 1 MB; however, a block size of 512 KB provides the best deduplication ratio and performance balance. Increasing the block size reduces the number of virtual tape headers.

Figure 16-3 shows how to increase the tape block size to 512 KB. In this example, select the VTL_LTO3_01 drive under the Drives folder. Then, click the **Settings** tab; your Default Media Pool, VTL_TAPE_ONSITE, is displayed. Click **Advanced**; the Advanced Options window opens. Select the **Sizes** tab, and set a value of 512 KB in the **Block size (kB)** field.

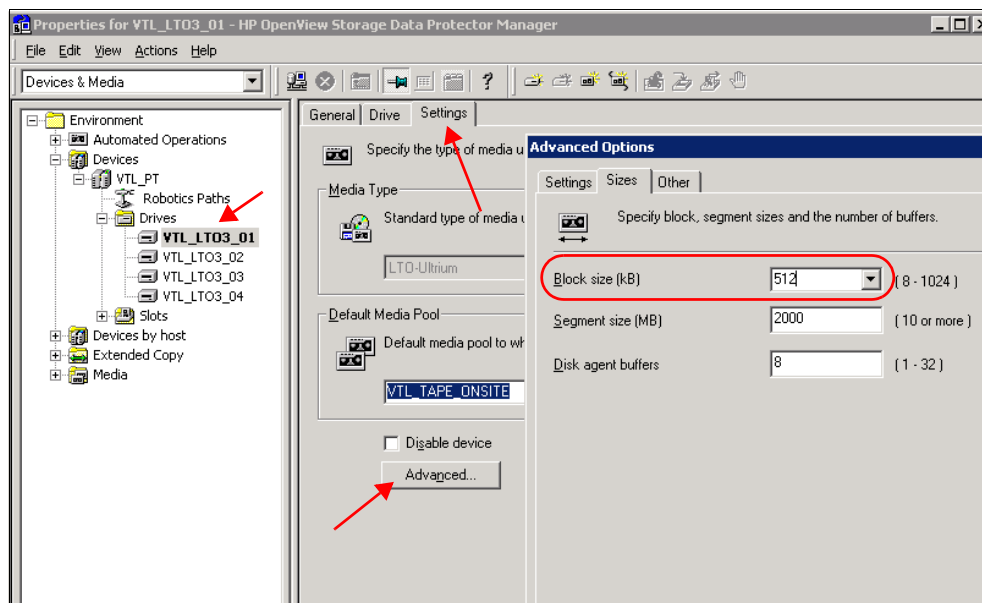


Figure 16-3 Increasing the tape block size

If your Media Agent is running on a Windows server, you must modify the registry.

Important: Before you make any changes in the registry, make a consistent backup.

To modify the registry, complete the following steps:

1. Open the registry editor by running **regedit.exe** and navigating to the following path:
 \HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ql2300\Parameters\Device
 Figure 16-4 shows the registry key.

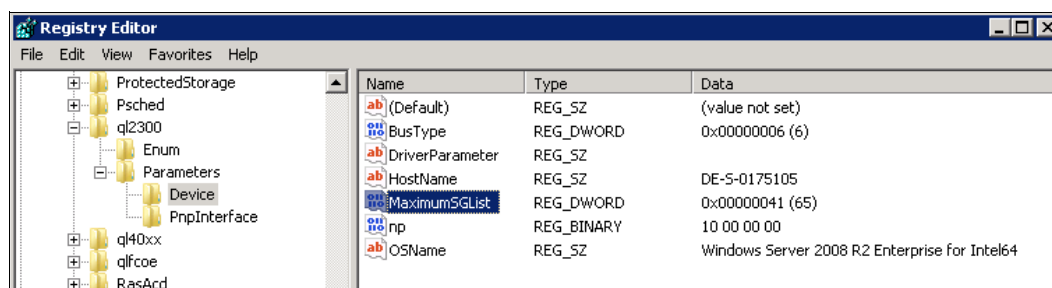


Figure 16-4 Windows registry records of QLogic HBA

Note: The registry value **MaximumSGList** exists only in Windows 2008 by default. In Windows 2003, you must create it by adding a DWORD of the same name.

2. Modify the value of the tape block size to 512 KB (524288 bytes). Enter this value as a decimal number, not hexadecimal (Figure 16-5). Calculate the "Value data" as follows:

"Value data"=(wanted Block Size in bytes)/4096+1,

In this case, the value data has this result:

"Value data"=524288/4096+1=129




Figure 16-5 Setting a greater tape block size

3. Restart your server to make the changes effective.

Important: The **MaximumSGList** parameter is a value in the range of 0 - 255. Use this adjustment sparingly, because scatter/gather lists use non-paged memory. Avoid setting it arbitrarily to a high value. For more information, see the following web page:

<http://support.microsoft.com/kb/280793>

To increase the block size above 256 KB for HP-UX, IBM AIX, Linux, or Solaris with the HPDP, add the following variable to the /opt/omni/.omnirc file:

OB2LIMITBLKSIZE=0.

To support a block size of 1024 KB, set the **st_large_recs** parameter to a nonzero value.

16.2.3 Enabling the lock name

The virtual robotic arm in your VTL is responsible for maintaining records of where all virtual cartridges are. It tracks which tapes are in each storage slot, entry/exit ports, the robotics gripper, and tape drives. These robotic arms are available to all Media Agents where direct backup to tape is needed.

Enable the lock name to prevent a collision when HPDP tries to use the same physical device in several backup sessions at the same time. Complete these steps as shown in Figure 16-6:

1. In this example, select the VTL_LTO3_01 drive from the Drives folder.
2. Click the **Settings** tab; you see your Default Media Pool, VTL_TAPE_ONSITE.
3. Click the **Advanced** tab; the Advanced Options window opens.
4. Click the **Other** tab and select the **Use Lock Name** check box.

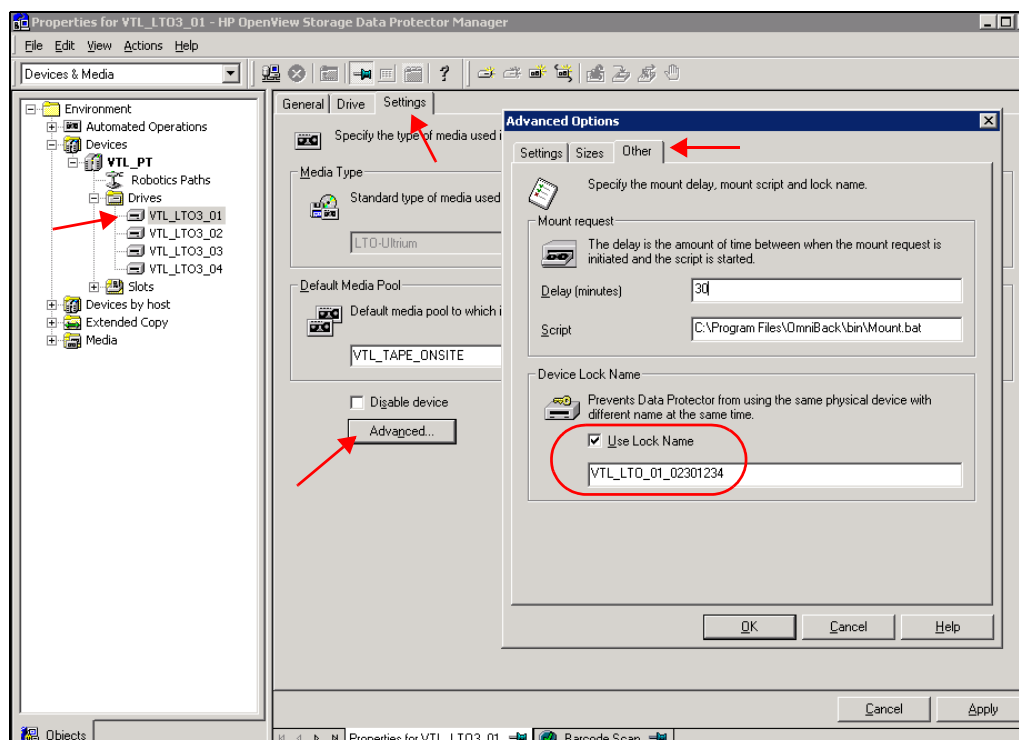


Figure 16-6 Enable lock name

16.2.4 Disabling compression, encryption, and CRC chksum

Compression, encryption, hash-based application deduplication, cyclic redundancy check (CRC) of data integrity, and other processes, affect the ProtecTIER deduplication factoring ratio. Avoid using these techniques when the IBM ProtecTIER solution is implemented as a component of your backup environment.

Although HPDP offers these features, ensure that they are disabled for virtual tape devices that provisioned by the ProtecTIER product. Figure 16-7 on page 236 shows how to deactivate those options by clearing them:

1. Disable compression, encryption, and CRC by selecting a file system, in this example FS_Daily_Incr.
2. Then, click the **Options** tab, and the Filesystem Options window opens.
3. Clear the **Software compression** and **Encode** check boxes.

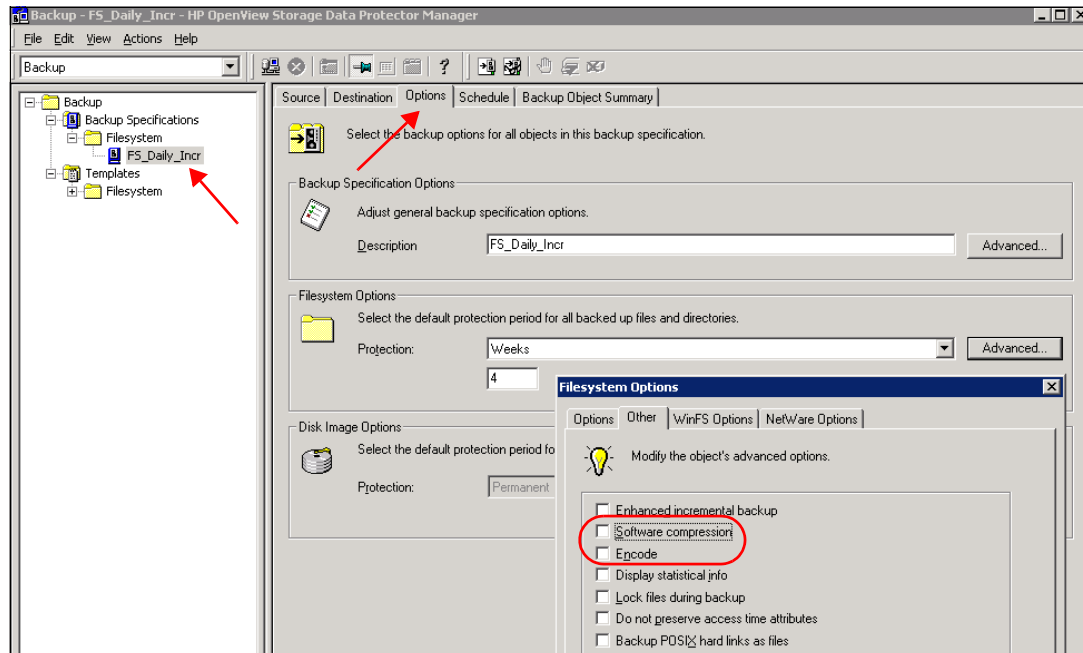


Figure 16-7 Disable compression, encryption, and CRC

You can use HPDP to use multiple Media Devices for different types of backups, or for different client systems (Disk Agents). For certain devices, such as physical tape drives or external disk arrays, it is beneficial to have these storage efficiency techniques enabled. Activate them individually for each backup process to these Media Devices, and ensure that the ProtecTIER virtual tape repository is not affected. This action is exceptionally important when deduplication is enabled.

16.2.5 Hosts multipath support

The IBM System Storage TS7650G ProtecTIER Deduplication Gateway (TS7650G) includes two Emulex dual-port LPe12002 8 gigabits per second (Gbps) Fibre Channel (FC) adapters for the front-end connections to the backup server and, optionally, to the LAN-free client systems (by using a SAN infrastructure).

For high availability (HA) or performance reasons, you might use more than one host bus adapter (HBA) on the server to access VTLs and drives. However, if you implement this configuration, you might see duplicate virtual tape drive definitions or robotics on your server, each accessible through a different HBA.

HPDP supports this multipathing function, although certain settings must be defined. Before you make such changes, ensure that the data path failover (DPF) and control path failover (CPF) are supported and enabled by the IBM Tape Device Driver that is installed on a relevant platform.

Hint: Use a dedicated IBM Tape Device Driver for ProtecTIER virtual tape devices. Ensure that multipathing is enabled in the operating system before you enable multipathing in the HPDP Devices and Media Manager.

For examples of how to enable multipathing on a tape device (either robotics or tape drive) in HPDP, see Figure 16-8 and Figure 16-9.

Figure 16-8 shows how to enable multipathing for the robotic path for a tape library:

1. From the Devices & Media drop-down menu, select the VTL_PT library.
2. Select the **MultiPath device** check box.

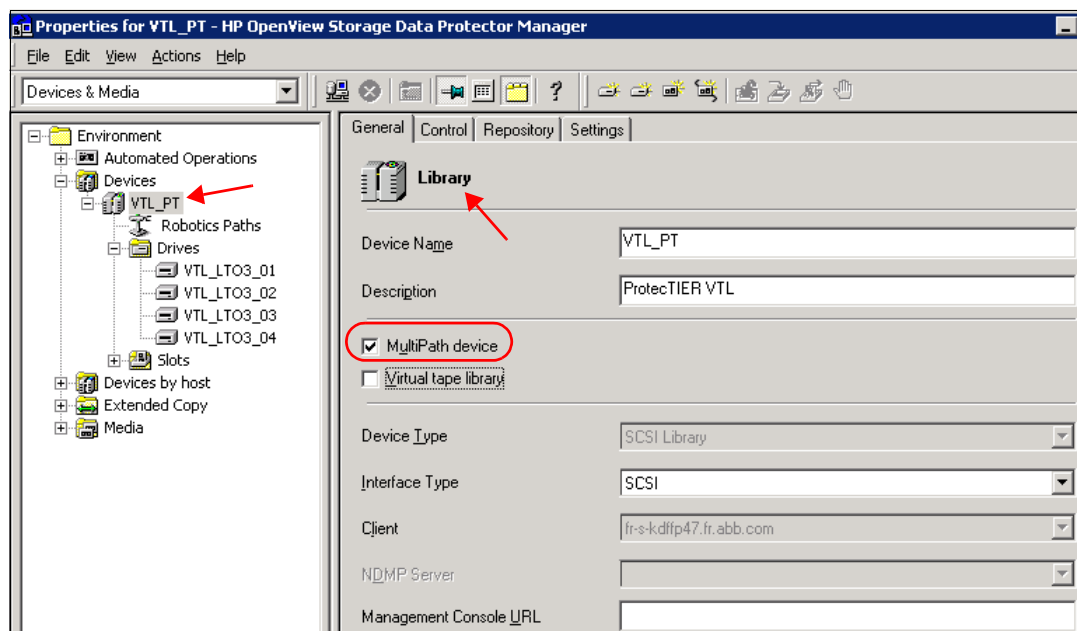


Figure 16-8 Multipath device on library robotics device

Figure 16-9 shows how to enable multipathing for the robotic path for a drive in a tape library:

1. From the Devices & Media drop-down menu, select the drive (in this case, VTL-LTO3_01).
2. Select the **MultiPath device** check box.

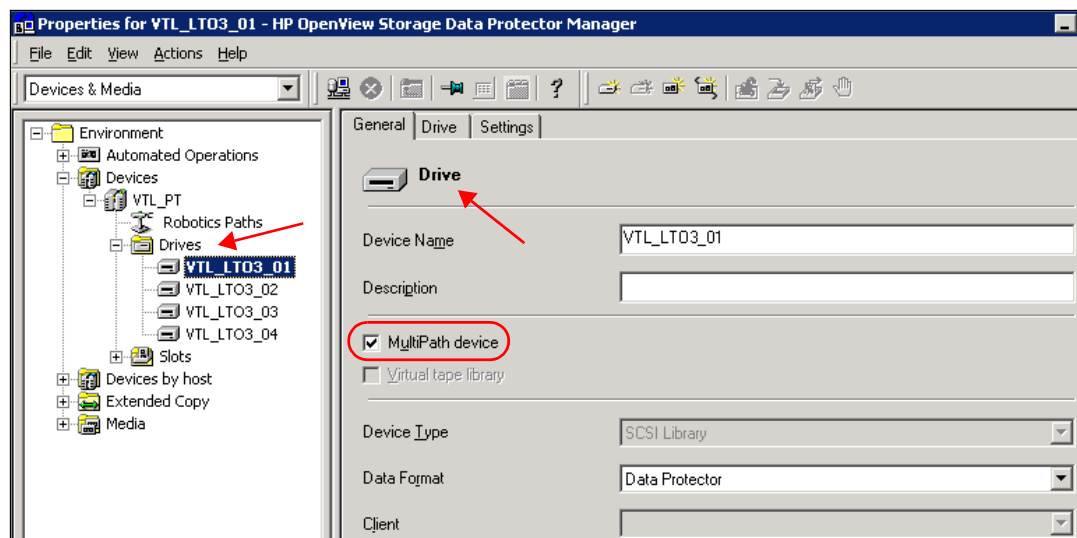


Figure 16-9 Multipath device settings for tape drives

Although the activation of multipathing in the operating system (by using a tape device driver) requires a reboot of the server, the setting of a specific device with multipathing in HPDP is a dynamic process and does not require a reboot. This change is effective immediately.

16.2.6 Load balancing

The load balancing feature specifies how many tape drives are available (minimally) at the start of the backup and how many in total are reserved at the start of this backup to the client system. The Data Protector uses the optimum amount of resources that are required for the tape devices workload. Using inappropriate settings of both parameters can make reserved tape drives unavailable for other backup sessions until the current backup job finishes.

Figure 16-10 provides an example of a setup for load balancing during a backup operation:

1. Select **Backup** from the drop-down menu.
2. For the backup specifications, select the FS_Daily_incr file system.
3. Click the **Destination** tab. Five tape drives are displayed in the Backup pane. Select the VTL_LTO3_04 and VTL_LTO3_05 tape drives, and set the minimum number of tape drives that are available to run backups to 2. Set the maximum number of tape drives to reserve to 5.

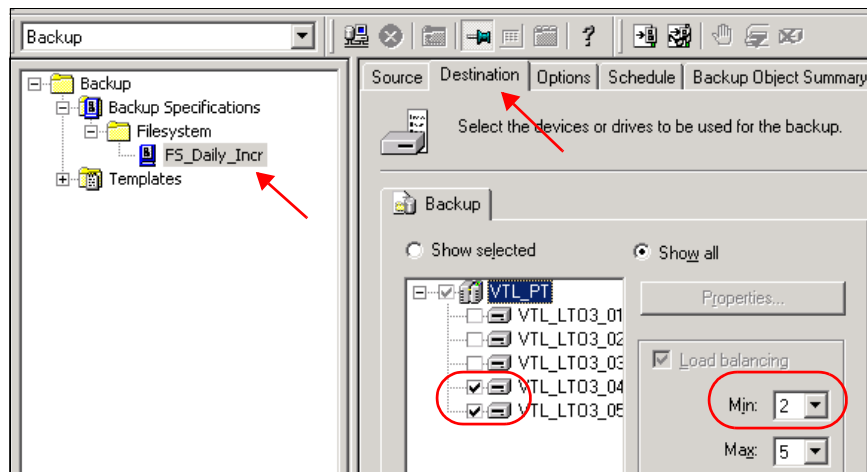


Figure 16-10 Load balancing

In this scenario, HPDP checks whether there is a minimum of two tape drives available for backup; if yes, the backup starts in two parallel streams. During backup, the HPDP can reserve up to five tape drives, making those drives unavailable to other client systems that might start backup later that day. HPDP selects any available drive that you assigned in the backup specification for this load balancing.

Tip: Do not use load balancing when you back up a few large objects, typically databases. In this scenario, HPDP is often not able to effectively balance the load among such devices. Define a dedicated backup policy with disabled load balancing for such objects.

16.2.7 Using a mirroring functionality

You can use the HPDP object mirror devices functionality to write the same data to several media simultaneously during a backup session. In some cases, this feature can replace vaulting or migrating the data between libraries and can decrease the usage of resources and increase performance of backup jobs.

Modify the backup policy settings, as shown in Figure 16-11. The figure shows the usage of a mirror. To add the object mirror functionality:

1. Select the file system (in this example, select FS_Daily_Incr).
2. Click the **Destination** tab and select the relevant drives. As illustrated in Figure 16-11, from the Mirror 1 tab, select the VTL_LTO3_01 and VTL_LTO3_02 drives.
3. Click the **Add Mirror** tab.

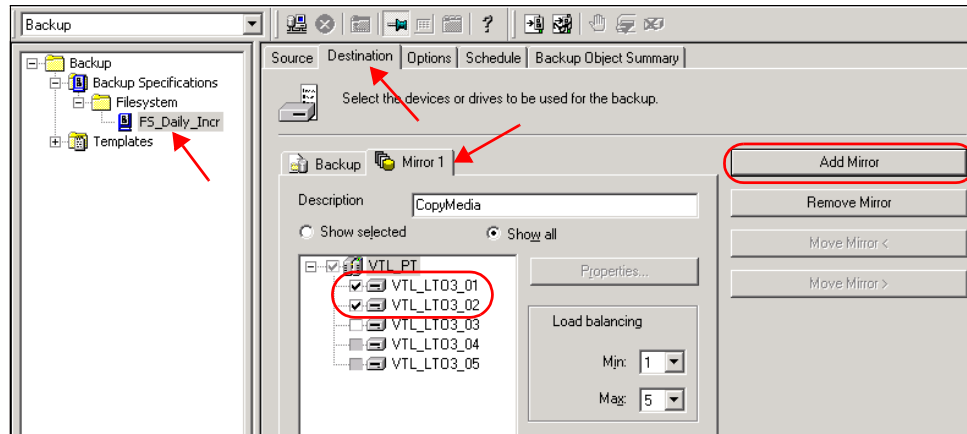


Figure 16-11 Usage of a mirror

16.2.8 Troubleshooting logs

The following information might be helpful with preliminary analysis troubleshooting:

- ▶ CLI output commands are from the Data Protector host server:
 - #devbre -dev
 - #sanconf -list_drivers
- ▶ Data Protector log files are in the following directories:
 - For Windows systems:
 - <Data_Protector_home>\log
 - For HP-UX:
 - /var/opt/omni/log and /var/opt/omni/server/log
 - For other UNIX systems:
 - /usr/omni/log

For more information about troubleshooting of HPDP and log analysis, see this web page:

http://h20564.www2.hp.com/hpsc/doc/public/display?docId=emr_na-c02029318

The general information and documents library for HPDP is available at this web page:

<http://www8.hp.com/h20195/v2/GetHTML.aspx?docname=c04109270>



IBM i and Backup, Recovery, and Media Services

The IBM ProtecTIER product extends the disk backup solution for IBM i with minimal disk usage by using the IBM HyperFactor deduplication feature.

The ProtecTIER product is deployed as a Virtual Tape Library (VTL) that emulates tape devices for IBM i. IBM i stores its backup data in ProtecTIER virtual cartridges as though they were physical cartridges. Users have the flexibility to create virtual tape drives and virtual cartridges that are based on their needs. IT centers that require a copy of backup data to be offsite can use ProtecTIER replication to replicate the virtual cartridges to a secondary site.

Because the ProtecTIER product emulates tape devices, you can always share it between IBM i and Open Systems. It is highly suggested to conduct a correct sizing based on your environment before the implementation.

This chapter focuses on using Backup, Recovery, and Media Service (BRMS) to integrate the ProtecTIER VTL product in to an IBM i environment.

This chapter describes the following topics:

- ▶ IBM i overview
- ▶ Integration of IBM i and ProtecTIER in a VTL environment
- ▶ Configuration of BRMS for ProtecTIER
- ▶ Deploying ProtecTIER with BRMS for disaster recovery (DR)
- ▶ Use DUPMEDBRM to duplicate media from ProtecTIER to physical media
- ▶ Device configuration
- ▶ BRMS movement of replicated volumes: considerations

17.1 IBM i overview

This section describes some IBM i-specific features, which you should be familiar with to understand how IBM i backups are performed.

Today, when you purchase IBM i, you purchase IBM POWER® hardware that can run IBM AIX, Linux on Power, or IBM i. You load your operating system into each logical partition (LPAR). In this example, load IBM i.

IBM i (formerly known as IBM AS/400, IBM eServer™ iSeries, and IBM System i®) is an integrated system that includes hardware, software, security, a database, and other components. IBM i has a flexible architecture where software is independent from hardware, so changing one has little effect on the other one.

IBM i and ProtecTIER: For information about IBM i resources for the ProtecTIER product, see the following web pages:

- ▶ Configuring tape libraries, in the IBM Knowledge Center:
http://www.ibm.com/support/knowledgecenter/ssw_ibm_i_72/rzam4/rzam4config.htm
- ▶ BRMS, in IBM developerWorks:
<http://www.ibm.com/developerworks/ibmi/brms>
- ▶ IBM Removable Media on IBM i, in IBM developerWorks:
<http://www.ibm.com/developerworks/ibmi/media>

17.1.1 Integrated file system

IBM i contains an integrated file system (IFS) that provides a common interface with other file systems. This file system enables applications that are written on other file systems, such as UNIX and Windows, to access data that is stored in IBM i.

17.1.2 Integrated database

IBM DB2 is the integrated relational database in IBM i. DB2 is tightly integrated into the entire system, making it efficient and easy to use by IBM i applications. DB2 is used by various applications, from traditional host-based applications to client/server applications to business intelligence applications.

17.1.3 Object-based architecture

IBM i is an object-based operating system. Unlike most other operating systems where everything is seen as a file, IBM i sees everything as an object. These objects include database files, user profiles, job queues, compiled programs, word-processing documents, and so on. Objects are categorized by type, which enables the users to specify what type of objects are required for a task. The IBM i operating system provides an interface to define operations that can be performed on objects and to provide instructions about the usage of the encapsulated data.

17.1.4 Libraries

IBM i groups objects into libraries. A library is a directory of objects, and it is an object that is used to find other objects in the directory. An object can exist in only one library. A library cannot reference other libraries except for the library that is called QSYS, which contains all the other libraries. Libraries can be associated or referenced by user profiles or applications.

17.1.5 Backup considerations in IBM i

IBM i offers a wide range of backup recovery options. They are intended to help you accomplish the following tasks:

- ▶ Make your *save* operations faster and more efficient.
- ▶ Keep your system available for your users.
- ▶ Plan and manage your backup and recovery.

Before you implement a backup solution in IBM i, consider performing the following actions:

- ▶ Determine the save objects and how often to save them.
 - Consider performing daily saves of the libraries and objects that regularly change, such as application libraries, user profiles, configuration objects, and parts of IFS. The objects that regularly change have a higher possibility of restoration in a shorter period as compared to objects that do not change regularly.
 - Consider performing weekly full system saves. The full system saves provide a baseline copy of your system that can be used for restoring all system objects in the event of a disaster. Alternatively, save all user libraries (*ALLUSR) every week.
- ▶ Determine the save window that is based on the following items:
 - Affordable downtime of save objects.
 - Affordable IBM i system downtime for a full system backup.
- ▶ Determine the recovery time and availability options.
- ▶ Test the developed backup and recovery strategy.

17.2 Integration of IBM i and ProtecTIER in a VTL environment

This section describes the considerations and preferred practices for configuring IBM i in a ProtecTIER environment using the VTL interface.

17.2.1 Backup considerations with ProtecTIER

Using VTL is not necessarily faster than physical tape backup. IBM tape products have been tested and work efficiently with IBM i. IBM i is able to achieve 90% - 100% of tape drive speed in an environment with fewer tape drives. You often require multiple streams in a VTL to achieve the same performance throughput as physical tapes. In this scenario, Backup, Recovery, and Media Service (BRMS) is useful in managing the tape media resources for parallel saves.

BRMS enables users to better manage resources. Although BRMS does not support the sharing of resources, it does support the sharing of media and saved history.

BRMS tracks what you saved, when you saved it, and where it is saved. When you need to perform a recovery, BRMS ensures that the correct information is restored from the correct tapes in the correct sequence.

17.2.2 Suggested ProtecTIER and IBM i configuration

Follow these guidelines when you configure ProtecTIER in your IBM i environment:

- ▶ Configure parallel or concurrent saves to reduce save windows.
- ▶ Use BRMS for automation. Ensure that you are licensed for BRMS.
- ▶ If using BRMS, you will also require a license for Media Storage Extensions, 57XX SS1, Option 18.
- ▶ The ProtecTIER product supports IBM i from Version 5 Release 4 onward, and both the IOP and IOPIess Fibre Channel (FC) adapters are supported. For more information, see the IBM SSIC and ProtecTIER ISV Support Matrix. More details are in Appendix B, “ProtecTIER compatibility” on page 457.
- ▶ Create a separate VTL if you are sharing one ProtecTIER system for IBM i and other platforms.
- ▶ Do not mix disk and tape on one FC adapter for performance reasons, because a tape workload is always large block size but a disk workload is typically small block. The ProtecTIER solution is VTL; it is considered a tape workload when integrated with IBM i.
- ▶ When you create virtual tape devices in ProtecTIER Manager, choose IBM TS3500 for your library type and LTO3 (ULT3580-TD3) for the drive type.
- ▶ Carefully design the drive attachment and robot attachment on each port (see Figure 17-1 for port assignment). Usually, you would spread the drives across all ports for optimal performance, but in IBM i, there are certain rules for using tape that you might want to consider. Some of the rules are listed in the next bullet.

Port	WWN	Drives	Robot
<input checked="" type="checkbox"/> Port 0	[WWN= 10000000c97c62e2]	6	<input checked="" type="checkbox"/>
<input type="checkbox"/> Port 1	[WWN= 10000000c97c62e1]	0	<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/> Port 2	[WWN= 10000000c97c5f4a]	6	<input type="checkbox"/>
<input type="checkbox"/> Port 3	[WWN= 10000000c97c5f49]	0	<input checked="" type="checkbox"/>

Buttons: Select All, Equally divide, Clear

Figure 17-1 Port assignment

- ▶ IBM i cannot split the backup evenly across the FC adapters if more than one FC tape adapter is used in one LPAR. It is advised for you to run multiple saves simultaneously for a ProtecTIER deployment, the tape library, tape drive layout, and SAN zoning must be designed carefully. For more information, see SAN design for IBM i Tape and ProtecTIER: <http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS2997>

Several rules for designing your ProtecTIER environment with IBM i are as follows:

- Multipath is not supported in IBM i for the FC tape adapter. If required, create two storage area network (SAN) zones to the tape devices, one active and one inactive. You must manually activate the inactive SAN zone when the active zone fails.

See the following link for more details:

<http://ibm.co/11BW5sI>

- A maximum number of logical unit numbers (LUNs), or addresses, are supported by an FC tape adapter in IBM i. The LUN count includes both drives and control paths:
 - 16 LUNs on an IOP FC adapter
 - 64 LUNs per port on an IOPIess FC adapter

Table 17-1 is an example of how to count the number of LUNs per FC tape adapter.

Table 17-1 Example of LUN counts per FC tape adapter

Number of tape libraries	Number of drives per port	Number of connected ProtecTIER ports (robot enabled at all connected ports)	LUN (device) count
1	8	1	9 (1 control path and 8 data paths)
1	8	4	12 (4 control paths and 8 data paths)
16	1	1	32 (16 control paths and 16 data paths)

- ▶ When you create the virtual tape cartridges, note the number of tape cartridges, because each cartridge requires a slot.
- ▶ Create extra slots for your VTL in case you need some slots for extra tape cartridges, because the library is placed offline when you want to change any of the logical library dimensions.
- ▶ Create enough import/export slots to hold all tape cartridges that you eject (replicate) every day.
- ▶ Enabling LUN masking on the ProtecTIER server is suggested.

For details about LUN masking in a ProtecTIER VTL environment, see 6.3, “LUN masking for VTL systems” on page 87.

- ▶ IBM i has a limitation of a maximum of 15,000 elements (places where tape cartridges are stored, such as slots, drives, picker, import/export slots, and so on). So, the total number of tape cartridges slots, extra slots, and import/export slots of each tape library must not exceed 15,000 elements.

Elements: The maximum number of elements was increased to from 5000 to 15,000 elements with the following program temporary fixes (PTFs):

- ▶ R611: MF50093, MF55406, MF55407, and MF55408
- ▶ R7.1: MF55396, MF55397, MF55409, and MF55393(MF55600)

- ▶ To allow more than 250 virtual tape libraries, be sure the following PTFs are applied:
 - V6R1M1: MF56115
 - IBM i 7.1: MF56114
 - IBM i 7.2: in base
- ▶ Do not use compression or compaction when you save to virtual cartridges on the ProtecTIER server.

17.3 Configuration of BRMS for ProtecTIER

This section describes the fundamental structure of the BRMS application for IBM i systems. Also provided is BRMS terminology, and suggested configurations of BRMS.

17.3.1 BRMS overview

You can plan, control, and automate the backup, recovery, and media management services for your IBM i systems with BRMS. Figure 17-2 shows how BRMS integrates the tape devices, which can be ProtecTIER virtual tape devices.

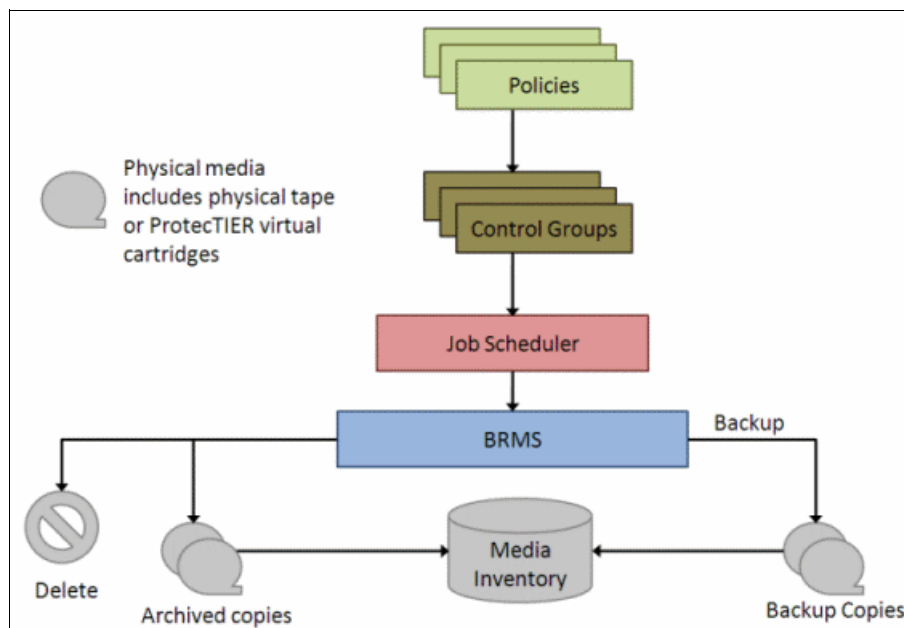


Figure 17-2 Fundamental structure diagram of BRMS

BRMS terminology

Several common BRMS terms are as follows:

Media	A tape cartridge (volume) that holds the saved data.
Media identifier	A name that is given to a media.
Media class	A logical grouping of media with similar physical, logical, or both characteristics.
Control group	A group of libraries, objects, folders, spool files, and other types of data that share common characteristics or a group of items that you want to group it together for backup purposes. It determines which data is to be processed.

Policies	Defines how BRMS operations are processed. There are different types of policies, which include system policy, backup policy, recovery policy, media policy, and move policy. A policy determines how data is to be processed.
QBRM and QUSRBRM	These IBM i libraries contain BRMS-related objects and management information.

17.3.2 Suggested configurations of BRMS

VTL versus physical tape drives: Configuration of BRMS for ProtecTIER virtual tape devices is no different from configuration for physical tape devices. Relevant BRMS version guidelines and steps to configure BRMS tape devices are available in the IBM BRMS product information section at the following web page:

<http://www.ibm.com/systems/i/support/brms/prodinfo.html>

This section focuses on the suggested BRMS configurations that are relevant to ProtecTIER virtual tape devices. The following list includes some guidelines:

- ▶ Use **STRSST** to reset the tape adapter that is attached to the ProtecTIER node to recognize newly created virtual tape devices in ProtecTIER system.
- ▶ The ProtecTIER VTL takes the next available TAPMLBxx name by default. Rename it to match the name that you used in ProtecTIER Manager to keep your environment orderly.
- ▶ The ProtecTIER product emulates LTO3 tape devices. When you define the BRMS media class in a ProtecTIER environment, create a separate media class for virtual LTO3 volumes from the media class for physical LTO3 volumes. Use a *ULTRIUM3 density.

Tip: You can make the virtual LTO3 cartridges any size you want, but not the LTO3 size. The cartridges' performance matches whatever your ProtecTIER configuration supports, and is not tied to LTO3 speed.

- ▶ Add the virtual cartridges in to BRMS by running **ADDMLMBRM** command and be sure to initialize them to reflect the virtual cartridge's actual label (bar code). Use the actual command rather than navigating there from the **WRKMLMBRM** menu.
- ▶ To choose your replication destination with ProtecTIER Manager, complete the following steps:
 - If you choose a remote VTL as a destination, configure the move policy for the selected media so that the media can be moved out from the local VTL to the shelf after a save operation. The replica cartridges can then be moved to a remote VTL and be accessible in remote IBM i.
 - Do not configure a move policy if you want the media to stay in the library after save operations. You can perform a manual move when you want to access it from a remote site.

Configuring a parallel save with BRMS

You can configure a parallel save by specifying the number of parallel device resources in the BRMS control backup control group attributes. An example of the needed specified attributes is shown in Figure 17-3.

```
Change Backup Control Group Attributes

Group . . . . . : SAMPLE

Type information, press Enter.

Media policy for:
  Full backups . . . . . *BKUPCY      Name, F4 for list
  Incremental backups . . . . . *BKUPCY  Name, F4 for list
  Backup devices . . . . . *BKUPCY      Name, F4 for list

Parallel device resources:
  Minimum resources . . . . . 2          1-32, *NONE, *AVAIL
  Maximum resources . . . . . 4          1-32, *AVAIL, *MIN
  Sign off interactive users . . . . . *BKUPCY *YES, *NO, *BKUPCY
  Sign off limit . . . . . *BKUPCY      0-999 minutes, *BKUPCY
  Default weekly activity . . . . . *BKUPCY SMTWTF(S/F/I), *BKUPCY
  Incremental type . . . . . *BKUPCY      *CUMUL, *INCR, *BKUPCY
  Force full backup days . . . . . *BKUPCY 0-365, *NOMAX, *BKUPCY

F3=Exit  F4=Prompt  F12=Cancel
```

Figure 17-3 Example of configuring parallel saves with BRMS

17.4 Deploying ProtecTIER with BRMS for disaster recovery

This section describes various scenarios for replication and performing DR by using ProtecTIER with BRMS. The first scenario describes how to recover an IBM i production system with BRMS, where BRMS is installed on both the production site and the DR server in one BRMS network. The second scenario describes how to recover an IBM i production system with BRMS installed on the production system, and there is no BRMS installed at the DR site.

17.4.1 BRMS available at the production site and DR site

In this scenario, both the production and DR sites have an IBM i server with a ProtecTIER server connected. BRMS is installed on both the production and DR server in one BRMS network, and the information about the media, devices, and so on, is synchronized between two BRMS systems.

The production IBM i performs backups to a locally attached ProtecTIER system with local BRMS. Replication is established from the production ProtecTIER system to the DR ProtecTIER system.

Replication setup

To set up replication, complete the following steps:

1. Configure the ProtecTIER replication policy with the destination as the VTL of the DR ProtecTIER server and with *Visibility Change* enabled.
2. Configure a move policy for these virtual cartridges so that the virtual cartridges can be moved to the shelf after they are ejected from production. Then, the replica cartridges are moved from the shelf to the I/O station (import/export slots) of the selected VTL at the DR site when BRMS movement is run by running **MOVMEDBRM** or **STRMNTBRM** **MOVME** (*YES) commands.
3. As soon as the cartridge is ejected on the source side, or moved using BRMS, it is placed on the shelf at the source, and moved into the I/O station in the target VTL. This happens automatically only if *Visibility Change* is enabled in the replication policy, otherwise the cartridge will end up on the shelf of both the source and the destination of the replication pair.

Notes:

- ▶ Replication only occurs if the target volume is located on the shelf in the Destination repository.
- ▶ If for some reason, the target volume is moved to a logical library at the Destination, then the replication for this particular volume will be suspended until it is returned to the shelf by ejecting it from the virtual library.

Important:

- ▶ Use BRMS to move media within a virtual library or from a virtual library to the shelf.
- ▶ Avoid using the ProtecTIER Manager GUI to perform this type of tasks as it can cause the backup application to lose track of the location of a particular volume in the libraries.

Disaster at the production site

If a disaster occurs at the production site, complete the following steps:

1. Use different volumes for backups when running BRMS.
2. Generate a ProtecTIER replication statistic (.csv report) that includes statistics for all replica cartridges, including sync time.
3. Review the .csv report to determine whether the cartridges you want to restore have consistent data. If they do not contain consistent data, consider restoring from an earlier set of backups, as described in “Assessing the cartridges’ status and synchronizing with the catalog” on page 410 and “Recovering the data” on page 410.
4. Restore the IBM i system from the consistent replicas of cartridges in the DR ProtecTIER server.
5. You can perform daily saves to the DR ProtecTIER server during the outage of the production system.

Failback

After the production system is running and connected to the production ProtecTIER server, complete the following steps:

1. Create a failback policy to replicate all the cartridges or just the ones that were written during the outage to the production ProtecTIER server.
2. When all the needed virtual cartridges are replicated to the production ProtecTIER server, stop the failback replication by leaving DR mode from the DR ProtecTIER Manager.

For details regarding the failback process and other operations related to disaster recovery, see Chapter 22, “Disaster recovery deployment with backup applications” on page 395.

You can now resume the daily saves from the production system.

17.4.2 No BRMS at the DR site

In this scenario, the production IBM i performs backups to a locally attached ProtecTIER system with BRMS. Replication is established from the production ProtecTIER system to the DR ProtecTIER system, and there is no BRMS installed at the DR site.

Disaster at the production site

If there is a disaster at the production site, restore the entire system to a server at a DR site by completing the following steps:

1. Enter DR mode from the DR ProtecTIER server.
2. Generate a ProtecTIER replication statistic (.csv report) that includes statistics for all replica cartridges, including sync time.
3. Review the .csv report to determine whether the cartridges you want to restore have consistent data. If they do not contain consistent data, consider restoring from an earlier set of backups, as described in “Assessing the cartridges’ status and synchronizing with the catalog” on page 410 and “Recovering the data” on page 410.
4. After you obtain the list of consistent media, restore the IBM i system with the BRMS recovery report. Completing the following steps might also be necessary:
 - a. Restore System i Licensed Internal Code.
 - b. Restore the operating system.
 - c. Restore the BRMS product and associated libraries QBRM and QUSRBRM.
 - d. Restore BRMS-related media information.
 - e. Restore user profiles.
 - f. Restore System libraries QGPL, QUSRSYS, and QSYS2.
 - g. Restore configuration data.
 - h. Restore IBM product libraries.
 - i. Restore user libraries.
 - j. Restore document library.
 - k. Restore IFS objects and directories.
 - l. Restore spooled files, journal changes, and authorization information.
5. You can perform the daily saves on the recovered IBM i system at the DR site to the DR ProtecTIER server during the outage.

Alternatively, if there is an FC connection from the production system to the DR ProtecTIER server, and depending on the type of failure on the production site, you can establish a connection for the DR ProtecTIER server to the production system and restore the required objects from the DR ProtecTIER server.

Examples of BRMS policies and control groups

In this example, BRMS on a production IBM i and on a DR IBM i are in the BRMS network. Create the source VTL (TS7650SRC1) on the production ProtecTIER system, and create the target VTL (TS7650TGT1) on the DR ProtecTIER system. Both VTLs are known to the BRMS network.

Complete the following steps:

1. After the virtual cartridges are added to BRMS and initialized, set up the BRMS move policy to move the cartridges from the location TS7650SRC1 to TS7650TGT1. The cartridge stays in the TS7650TGT1 until it expires. Give the move policy a name, TS7650 (Figure 17-4).

The Media policy in Figure 17-4 specifies TS7650.

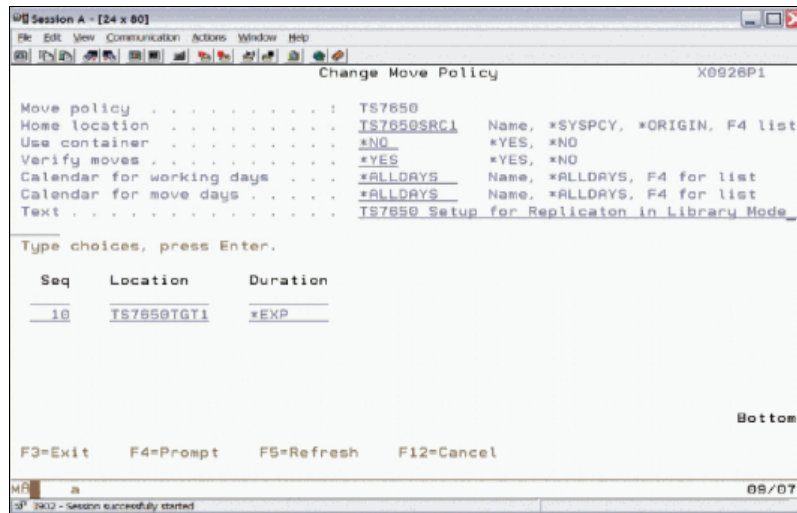


Figure 17-4 Example of a media policy

In Figure 17-8 on page 253, the Home Location specifies the source VTL. Under the Location column, type the name of the target location.

2. Create a BRMS media policy to use this move policy (Figure 17-5).




Figure 17-5 Example of change control group attributes

3. Create a BRMS backup control group to use the media policy with the VTL on the production ProtecTIER system (Figure 17-6).

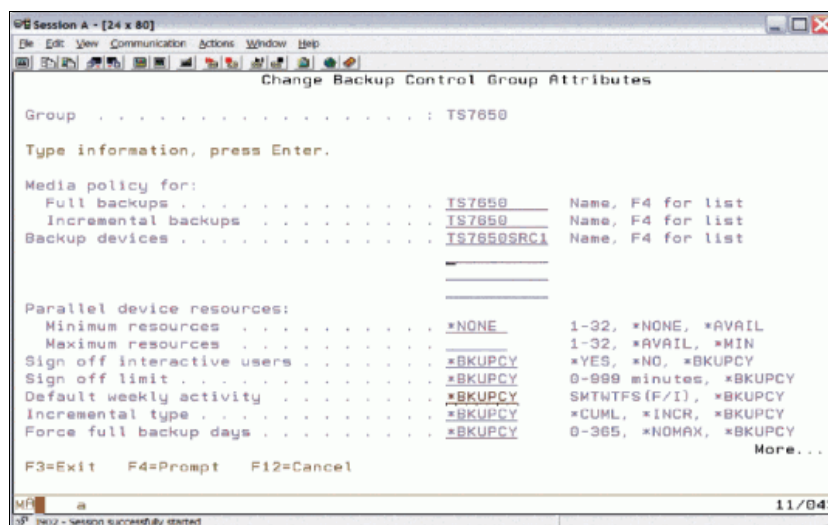


Figure 17-6 Example of verify media move

In Figure 17-6, the following field values are specified:

- The Full backups field specifies TS7650.
 - The Incremental backups field specifies TS7650.
 - The Backup devices field specifies TS7650SRC1.
4. On ProtecTIER Manager, configure the replication policy with a destination in the target VTL (TS7650TGT1). The moved media has an Inserted status in the DR IBM i system.
 5. Verify the media move on the DR IBM i system, which makes the cartridge available (Figure 17-7). The Location column specifies the name of the target VTL.




Figure 17-7 Displaying media status

- The cartridge is now available at DR site (Figure 17-8). The Status column for the moved cartridges specifies Available.




Figure 17-8 Example of change move policy

17.5 Use DUPMEDBRM to duplicate media from ProtecTIER to physical media

When duplicating media from the ProtecTIER to physical media, the BRMS **DUPMEDBRM** command copies the contents of a single volume, a single volume in a media set, a media set, or a set of marked saved items to other volumes. The following sections describe the steps and configuration to use the **DUPMEDBRM** command.

Note: Throughout the following sections, using the **DUPMEDBRM** command, the **COMPACT** parameter on the command line needs to be changed to ***YES** or ***DEV**.

17.5.1 Enable append to multiple parallel volumes during DUPMEDBRM

BRMS can be configured to enable multiple parallel volumes to be appended to the same volume when running the command **DUPMEDBRM TOSEQNBR(*END)**. The first volume will start on a scratch volume and all other volumes in the parallel set will be appended to that volume until it is full:

```
CRTDTAARA QUSRBRM/Q1AALWMFDP TYPE(*CHAR) LEN(1)
```

The following notes apply when performing restores from duplicated media:

- ▶ When performing restores from the duplicated single volume, parallel resources need to be set to a minimum and a maximum of 1 **-PRLRSC(1 1)**.
- ▶ Ensure that only expired media is available in the media class specified in the **MEDPCY** parameter on the **DUPMEDBRM** command. If there are any volumes in that media class, they will be appended too.
- ▶ Restores from the duplicate media might take longer.
- ▶ To change back to not duplicate to a single volume, delete the data area:
DLTDTAARA QUSRBRM/Q1AALWMFDP

17.5.2 Support multiple batch jobs for the DUPMEDBRM command

In Version 5R4M0 and later, the **Duplicate Media using BRMS (DUPMEDBRM)** command supports using multiple batch jobs to asynchronously duplicate media sets. The batch jobs used for duplication will be submitted using the **Job description** and **Job queue** fields in the **BRMS submitted jobs** section of the **BRMS System Policy**.

The job log from the **DUPMEDBRM** command will contain messages indicating which batch jobs are performing the duplication. Each of these duplication jobs will send completion or error messages to the BRMS log, so the BRMS log should be monitored to verify that each of the duplication jobs completed successfully.

To enable **DUPMEDBRM** batch job options for the current job, run the following command:

```
CALL PGM(QBRM/Q1AOLD) PARM('DUPBATCH' '*SET' 'nn')
```

Note: In this command, *nn* is the number of batch jobs to use. This value must be greater than or equal to 00 and less than or equal to the number of device resources available to be used during the duplication. 00 indicates to use the default behavior.

To display the current job's **DUPMEDBRM** batch job options, run the following command:

```
CALL PGM(QBRM/Q1AOLD) PARM('DUPBATCH' '*DISPLAY')
```

To remove the current job's **DUPMEDBRM** batch job options, run the following command:

```
CALL PGM(QBRM/Q1AOLD) PARM('DUPBATCH' '*REMOVE')
```

Restrictions

The following restrictions apply when running the **DUPMEDBRM** command:

- ▶ The call to **DUPBATCH** only applies to the current job.
- ▶ The number of values specified for the **Input volume list (FROMVOL)** parameter on the **DUPMEDBRM** command must include more than one volume identifier to be duplicated.
- ▶ If you need to submit the **DUPMEDBRM** command to batch, a program must be created that has the call to **DUPBATCH** and the **DUPMEDBRM** command in it. This program can then be submitted to batch.
- ▶ Batch jobs will only be used when a volume name, ***LIST** or ***SEARCH** is specified for the **From volume identifier (VOL)** parameter. Batch jobs will not be used when ***SCHST** or ***RESUME** is specified for the **From volume identifier (VOL)** parameter.
- ▶ If the media being duplicated was created during a parallel save, then restores from the duplicated media will require the same or a lesser number of resources used during the **DUPMEDBRM** command.

- ▶ The **Output volume list(TOVOL)** parameter value must be *MOUNTED. Specifying output volumes is not supported.
- ▶ The DUPMEDBRM command is not supported when using the autodup function that can be set up using media policies.

Note: The following PTFs are required for BRMS releases 7.1 and earlier:

- ▶ 7.1 S147935
- ▶ 6.1 S147934
- ▶ V5R4M0 S147933

17.6 Device configuration

When configuring devices on the ProtecTIER, ensure that LUN masking is enabled. Because users can now create more virtual tape drives, it is suggested that tape resources not be shared, and that each system have its own devices.

17.7 BRMS movement of replicated volumes: considerations

Most customers are going to have two sites with a ProtecTIER product at each site, and then systems at each of those sites. It is unlikely that a system will be physically attached to both ProtecTIER products. Example 17-1 shows a two-site type of scenario where all systems are in a BRMS network.

Example 17-1 Two site scenario with a ProtecTIER at each site, then systems at each site

SOURCEPT with **SYSA**, **SYSB** and **SYSC** attached
TARGETPT with **SYSD**, **SYSE** and **SYSF** attached

In the previous Example 17-1, there is a *save* on SYSA, SYSB, or SYSC to **SOURCEPT** and the media is replicated to **TARGETPT**.

BRMS movement runs on SYSA and the volume is ejected on **SOURCEPT**, moved to the shelf and on **TARGETPT**. It is then moved from the shelf into the IO station.

Until the volumes expire, now when they must move back to **SOURCEPT**, if movement is run on SYSA, SYSB, or SYSC, it will logically change the location of the volume, but it will not eject the volume from **TARGETPT**.

To successfully perform this scenario, SYSA, SYSB, or SYSC must be physically connected to that ProtecTIER. The Remove Tape Cartridge (**RMVTAPCTG**) command requires physical connection to the tape library. For this to work, you must run movement as depicted in Example 17-2:

Example 17-2 SYSA, SYSB, or SYSC physically connected to that ProtecTIER

On **SYSA**, **SYSB** or **SYSC** - **MOVMEDBRM LOC(SOURCEPT)** and
 On **SYSD**, **SYSE** or **SYSF** - **MOVMEDBRM LOC(TARGETPT)**



Commvault

This chapter describes all of the necessary steps for a successful integration of Commvault Simpana 9 Backup and Recovery software with the IBM ProtecTIER deduplication solutions to get the best factoring ratio and optimal performance.

ProtecTIER can be deployed as Virtual Tape Library (VTL) or File System Interface (FSI) to Commvault. This chapter describes Commvault with VTL, Commvault with FSI-Common Internet File System (CIFS) for Windows based servers, and FSI-Network File System (NFS) for UNIX clients.

Important: For preferred practices and configuration of Commvault in your ProtecTIER FSI environment, see 18.3, “Commvault FSI” on page 271.

The other versions of Commvault Hitachi Data Protection Suite 7 and 8 are also fully supported by the ProtecTIER product.

This chapter describes the following topics:

- ▶ Commvault introduction
- ▶ Commvault with ProtecTIER VTL
- ▶ Commvault FSI

Note: ProtecTIER GA Version 3.4 was released with only the Virtual Tape Library (VTL) interface support. File System Interface (FSI) support was added to ProtecTIER PGA 3.4 Version. For details, see the announcement letter:

http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/6/872/ENUSA P16-0076/index.html&lang=en&request_locale=en

18.1 Commvault introduction

Commvault Simpana is an enterprise backup and recovery software, which consists of fully integrated modules for backup and recovery, archiving, replication, search, and eDiscovery, all managed from a single user interface (UI).¹

These are the key capabilities and benefits of Commvault:

Common Technology Platform	Incorporates full compatibility across disk and tape products.
Virtual Server Protection	Can protect virtual machines with a block-based backup approach.
Migration Tools	Can migrate from different backup software, such as Symantec NetBackup.
Central Management	Helps operate the Simpana backup and recovery software from a central management interface.
Gen3 Deduplication	Incorporates an integrated and embedded hash-based deduplication solution.

Important: Never enable Commvault integrated deduplication if the IBM ProtecTIER solution is in place. Doing so severely degrades the benefit of IBM HyperFactor deduplication, which offers greater space savings than Commvault hash-based techniques.

Global Reporting	Quickly identify the status of your data protection environment, including the backup results, storage occupancy, and more.
Capacity Licensing	Pay for the amount of protected data independently of the complexity of the backup environment.

The Commvault Simpana Backup and Recovery software is a robust suite of data management capabilities that are built on a unified platform. By using it, you simplify data protection operations.

18.1.1 Commvault components

This section introduces the key components of Commvault Simpana Backup and Recovery software.

CommCell

The CommCell feature provides a set of storage management tools that you can use to move and manage your critical data. You can use these tools to store and retrieve data that is associated with computer systems in your enterprise. The system consists of integrated software modules that can be grouped in a CommCell configuration.

¹ Source: <http://www.commvault.com>

Each CommCell configuration consists of the following key components:

- ▶ One or more of the following *Client Agents*:
 - DataAgents that perform backup and recovery operation
 - Archive Management Agents
 - Quick Recovery agents that create Quick Recovery volumes
 - ContinuousDataReplicator to perform data replication from source to destination clients
- ▶ The *Common Technology Engine*, consisting of the following components:
 - One CommServe server
 - One or more MediaAgents
- ▶ *Storage Resource Manager* for analyzing and reporting of stored information
- ▶ *CommCell Console* for central management and operation of CommCell
- ▶ *Content Indexing and Search* engine for easy and fast data discovery

Common Technology Engine

The Common Technology Engine consists of software modules that provide the necessary tools to manage and administer the Client Agents and also manage the storage media that are associated with the CommCell configuration.

CommServe

The CommServe server ties the CommCell components together; it is the coordinator and administrator of the CommCell components. The CommServe server communicates with all agents in the CommCell to initiate data protection, management, and recovery operations. Similarly, it communicates with MediaAgents when the media subsystem requires management. It maintains a database that contains all the information that relates to the CommCell configuration. In addition, it provides several tools to administer and manage the CommCell components.

MediaAgent

The MediaAgent transfers data between the client computers and the storage media. Each MediaAgent communicates locally or remotely to one or more storage devices, which contain the storage media. The system provides support for various storage devices, including the IBM ProtecTIER virtual tape libraries and virtual tape drives.

CommCell Console

The CommCell Console is the graphical user interface (GUI) that enables you to control and manage the CommCell element. The CommCell Console can be run in two ways:

- ▶ As a stand-alone application, which can be installed directly on any computer that can communicate with the CommServe storage manager.
- ▶ As a remote web-based application using Java Web Start, which enables you to remotely access the CommCell Console by using the web browser.

Content Indexing and Search

You can use Content Indexing and Search to search and perform data discovery operations in your CommCell group. Use this powerful component to search both online and stored data. Administrators, Compliance Officers, and users can use it to search and restore data from several applications, such as FSI, Microsoft Exchange, Microsoft SharePoint, and IBM Lotus Notes, in the CommCell configuration. The search and restore operations can be performed by using either the CommCell Console or the web-based Search Console, which are controlled by a security module.

18.2 Commvault with ProtecTIER VTL

Commvault supports the installation of IBM ProtecTIER deduplication gateways in either a VTL or a FSI-CIFS configuration. However, each of these options requires specific planning and configuration steps that are further explained in the next sections. The general concept of the ProtecTIER product and Commvault integration is presented in Figure 18-1.

Important: For preferred practices and configuration of Commvault in your ProtecTIER FSI environment, see 18.3, “Commvault FSI” on page 271.

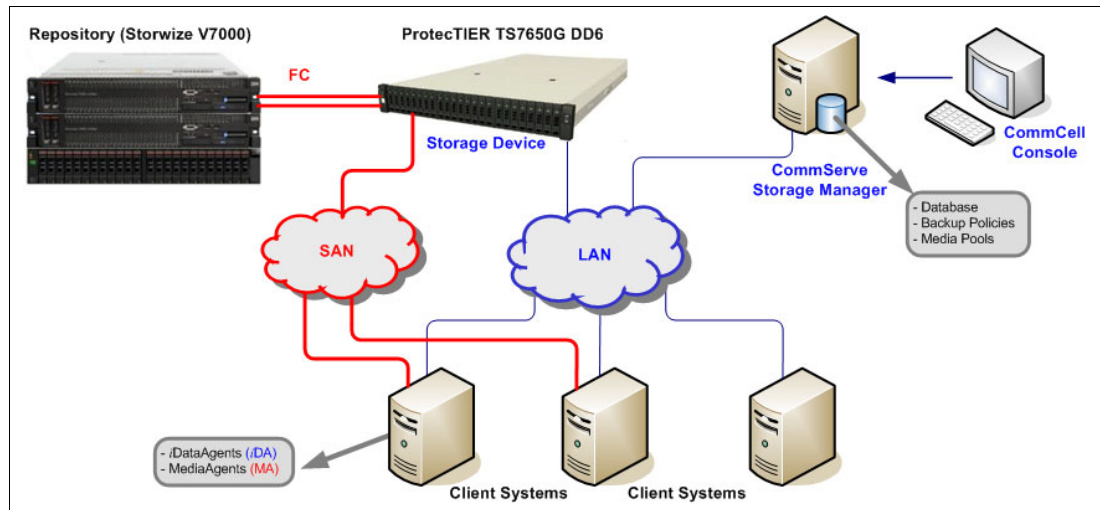


Figure 18-1 Conceptual overview of CommCell with ProtecTIER

The CommServe server is the core component in your CommCell implementation that manages all the data that is backed up from the clients. It has either ProtecTIER virtual tape drives, or file system shares that are configured as a storage repository for data deduplication.

In the case of the VTL installation, with the required storage area network (SAN)-based backup, the ProtecTIER virtual tape drives are also zoned to the client systems and configured by using MediaAgents. The iDataAgents then transfer the backup data to the CommServe repository, represented by the ProtecTIER deduplication solution.

Important: When you configure VTL on CommServe or client systems, be sure that the latest version of the IBM Tape Device Driver is correctly installed. Data path and control path should be enabled, and enabling the Persistent Naming feature is highly suggested. For guidelines about persistent device name binding, control path failover (CPF), and data path failover (DPF), see Chapter 6, “Host attachment considerations for Virtual Tape Library” on page 79.

18.2.1 Commvault configuration

This section guides you through the initial configuration of the VTL IBM System Storage TS3500 Tape Library (TS3500) with Ultrium LTO3 virtual tape drives, which are emulated by the IBM ProtecTIER deduplication gateway, using Commvault Simpana 9. Furthermore, it provides you with the mandatory parameters and the preferred practices that must be set to achieve the best factoring ratio from your ProtecTIER server.

Initial configuration

Initially, you must set up the physical cabling or SAN zoning, install your Commvault software, and configure your ProtectTIER server with the storage repository. After these tasks are completed, you must perform the initial configuration of the VTL in the CommServe server and on all relevant Clients' MediaAgents.

If you assigned tape drives to more than one ProtectTIER front-end port, a special procedure is needed to scan the VTL correctly because of a Commvault feature that is called *Exhaustive Detection* and its behavior. To work around Exhaustive Detection, complete the following steps for an existing, automatically detected tape library:

1. Stop all input/output (I/O) operations and ensure that a backup is not running. On the Commvault Service Control Manager, ensure that all Commvault services are running.
2. Log on to the CommCell Console by using your administrator account and password, click the CommCell Control Panel icon, then click **Library and Drive Configuration** (Figure 18-2).

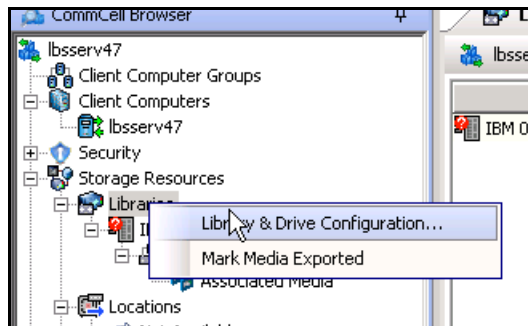


Figure 18-2 Library and drives configuration menu

3. Under the Available MediaAgents pane, double-click the server name. This action moves the server to the pane at the right (Figure 18-3). Click **OK**.

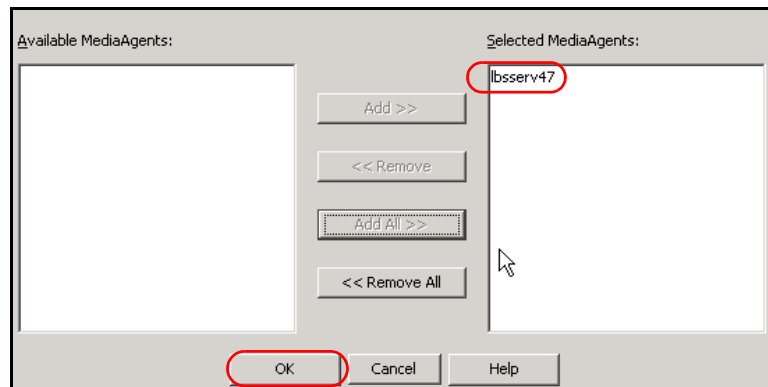


Figure 18-3 Select the MediaAgent

4. In the General tab, under Libraries, right-click the robot, click **Deconfigure** (Figure 18-4), and confirm the operation. You might be requested to confirm twice.




Figure 18-4 Unconfigure the existing tape library

5. Because it is not configured, right-click the robot again, click **Delete**, and confirm the operation. The library should now be deleted. Click **Start** and then **Exit**.
6. Now, you must define the library and drives again with the appropriate settings. In the toolbar icons, click the **Control Panel** and double-click **Library and Drive Configuration**. From the available MediaAgents, double-click the server (lbsserv47) to confirm your selection, and click **OK**.
7. In the General tab, right-click **Libraries** and click **Detect/Config Devices**. The Detect Library window opens (Figure 18-5). Ensure all the check boxes are clear.




Figure 18-5 Detect new library

8. Verify the correct detection of all devices in the generated log file window (Figure 18-6) and return to the General tab.



Figure 18-6 Detection log window

9. In the General tab, open the current node and navigate to the **Libraries** overview (Figure 18-7). Right-click the unconfigured tape library tag and select **Configure**.

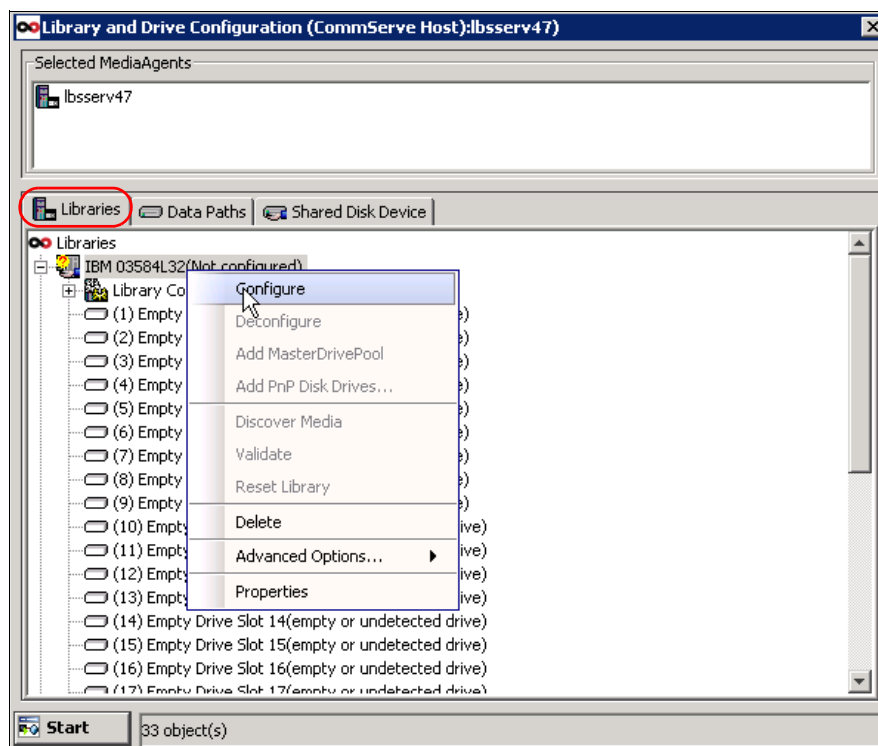


Figure 18-7 Configure the library robotics

10. The Configuration window opens, where you can select **Library only** or **Library and all Drives** to be configured. Click **Library only** (Figure 18-8). In most cases, the possibility of tape drive configuration is not available yet. Click **Yes** to confirm that the library has a bar code reader.

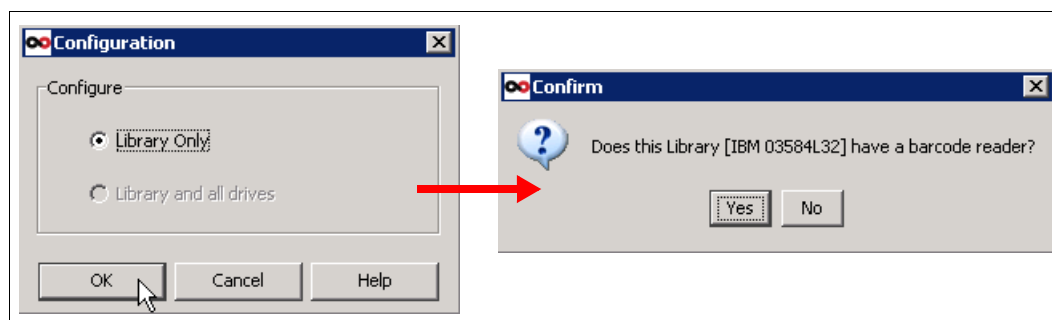


Figure 18-8 Confirm the library robotics configurations

11. Perform a full scan of the detected tape library by using the Full Scan option from the library menu (Figure 18-9).

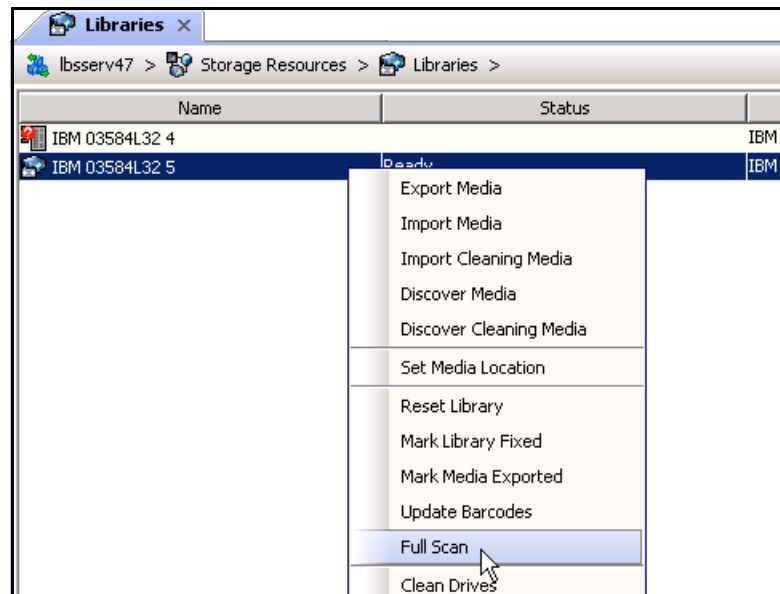


Figure 18-9 Perform a full scan of the library

12. Configure the first detected tape drive from the Libraries tab (Figure 18-10). This action enables you to configure all other tape drives in the future in one step.

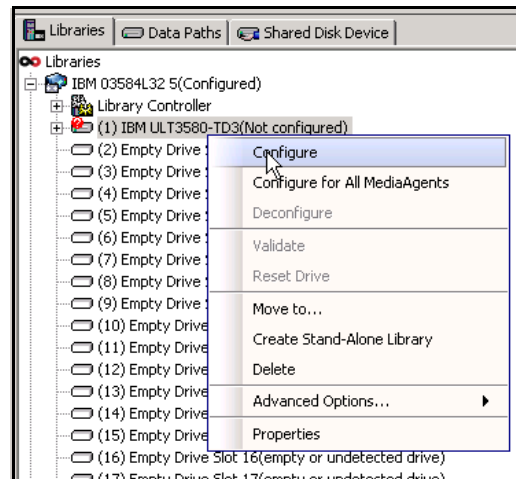


Figure 18-10 Configure first tape drive

13.Ensure that Exhaustive Detection is not used (Figure 18-11).




Figure 18-11 Do not use Exhaustive Detection

14.After the first tape drive is correctly configured, configure the remaining tape drives by using Exhaustive Detection. Using Exhaustive Detection saves time in configuring each remaining tape drive (Figure 18-12).




Figure 18-12 Configuration of remaining tape drives

15. When prompted, select the **Ultrium V3** type of media for a Discovery Media Option and verify that all the remaining tape drives are correctly detected and configured (Figure 18-13).




Figure 18-13 Verify the discovered tape drives

Important: If a library is present before the procedure, you must update the existing storage policies to the newly configured tape library. This action is not needed in the case of an initial configuration of the Commvault environment with the ProtecTIER server.

18.2.2 Data multiplexing

Data multiplexing is a Commvault licensed feature that enables multiple backup receivers to combine their backup data streams into one data writer. The data of multiple backup streams is then written to the single media. Figure 18-14 is an overview of how the data multiplexing feature works.

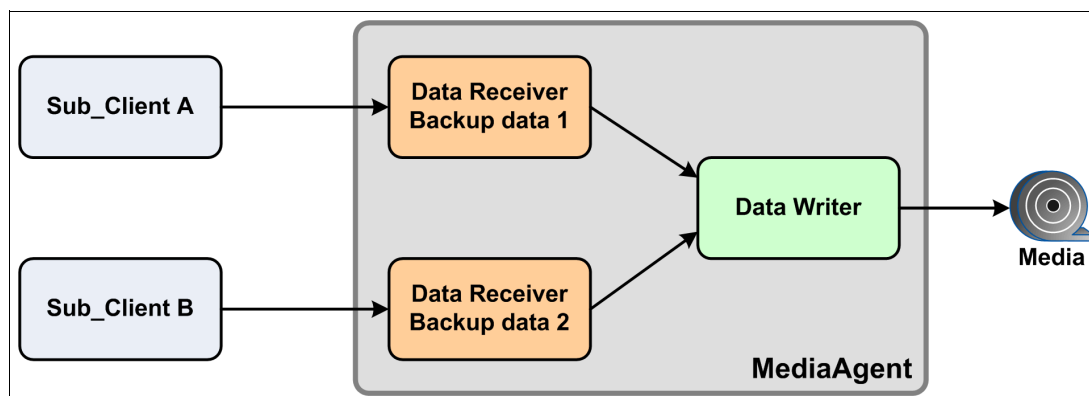


Figure 18-14 The concept of data multiplexing

Data multiplexing negatively impacts the ProtecTIER factoring ratio, so data multiplexing must be disabled. To prevent your subclients from using this feature, data multiplexing must be disabled from the Media tab of Copy Properties dialog box of the primary copy. In the CommCell browser window, navigate to **Storage Policies**, right-click the policy that you need to update, and select **Properties** (Figure 18-15).

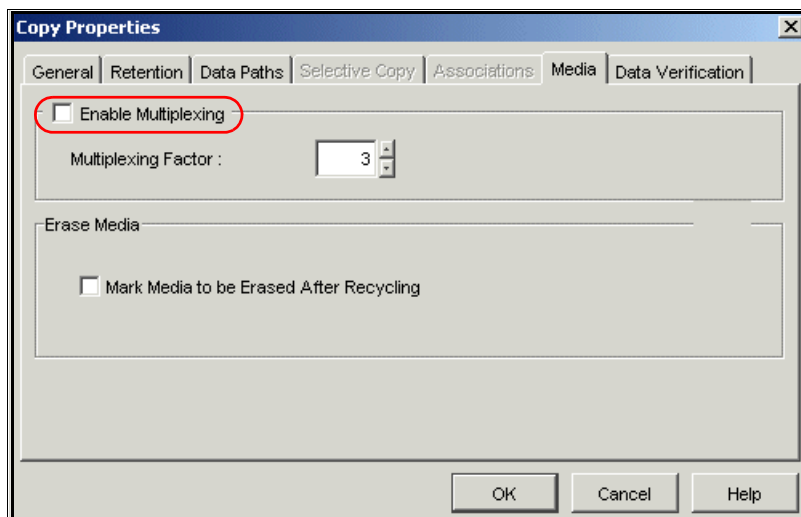


Figure 18-15 Disable data multiplexing

Multistreaming with Oracle jobs

The Oracle iDataAgent applies multiplexing rules for multiple jobs. Also, when you have multiplexing enabled for an Oracle job with multiple streams, all the streams of the job use the available drives sequentially (when one drive is full, it continues to the next). This task is enabled by setting the **JMEnableMultiplexingForOracleAgents** parameter in the CommServe database as follows:

```
Insert into GXGlobalParam values ('JMEnableMultiplexingForOracleAgents','1')
```

This parameter can be used only for Oracle jobs from the CommCell Console or when you initiate the job by using the **qoperation** backup command. In on-demand Oracle jobs, data multiplexing is enabled by default. Disable this feature by using the **QB_NO_MULTIPLEX_STREAM** option.

18.2.3 Hardware compression

You can use the IBM TS3500 Tape Libraries, which are emulated by the ProtecTIER server, to use hardware compression on virtual LTO3 tape drives. This hardware compression affects the factoring ratio, and therefore must be disabled.

To disable hardware compression, follow these steps:

1. In the CommCell browser window, navigate to **Storage Policies**.
2. Right-click the policy that you need to update, and select **Properties**.
3. Select the **Data Paths** tab of the Copy Properties dialog box of the primary copy.
4. Ensure that the **Hardware Compression** check box is clear (Figure 18-16).

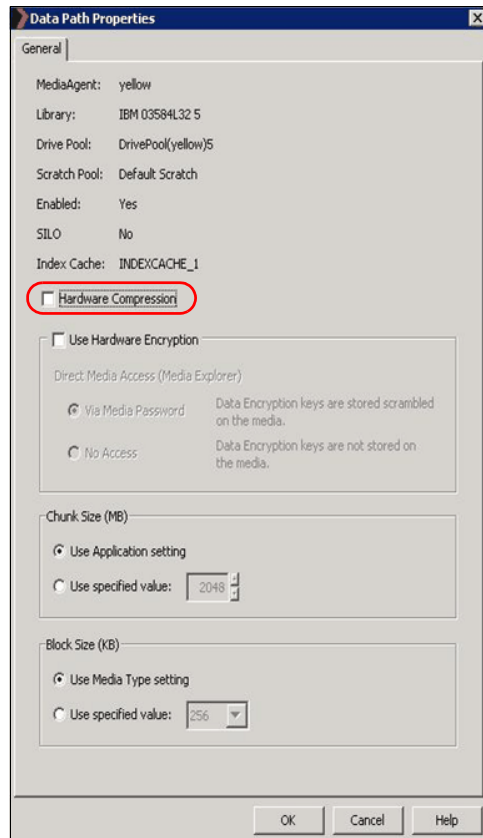


Figure 18-16 Disable hardware compression

18.2.4 Data encryption

The Commvault software supports three types of data encryption:

- ▶ Client level encryption of backup data
- ▶ Auxiliary copy level of data that is stored in CommServe media
- ▶ Hardware encryption on tape drives

None of these types of data encryption should be used with the IBM ProtecTIER solution, or the factoring ratio will be compromised or suppressed. This section briefly describes how to avoid using client encryption. Hardware encryption is not offered by the ProtecTIER emulated VTL, and an auxiliary copy of backup data is typically not stored in the ProtecTIER repository. The auxiliary copy is typically used for physical tapes, which are sent to the offsite location or local vault regularly.

Commvault is able to use the following type of encryption algorithms (Cipher) with different block sizes:

- ▶ Blowfish: 64-bit block size, 128- or 256-bit key length
- ▶ Advanced Encryption Standard (AES): 128-bit block size, 128- or 256-bit key length
- ▶ Serpent: 128-bit block size, 128- or 256-bit key length
- ▶ Twofish: 128-bit block size, 128- or 256-bit key length
- ▶ 3-DES (Triple Data Encryption Standard): 64-bit block size, 192-bit key length

To disable client-level encryption feature, right-click the client in the CommServe Console and select its properties. In the Encryption tab, ensure that the **Encrypt Data** check box is clear. This action automatically disables the data encryption for all relevant subclients that belong to the same Client's system (iDataAgents).

18.2.5 Alternative data paths

A data path is a licensed feature that integrates the MediaAgent, Library, Drive Pool, and Scratch Pool features that are used by the storage policy copy to provide backup operations. Each storage policy copy has a unique, single data path by default. For high availability purposes, you can define alternative data paths for each storage policy copy.

The Alternate Data Paths (ADP, also known as GridStor) feature provides the following benefits:

- ▶ Automatically redirects a backup stream to an alternative data path, if one of the components in the default data path is not available.
- ▶ Alternative data paths can be used to minimize media usage by routing data protection operations from several subclients to the same storage policy and the same media, rather than creating several storage policies by using a different media for each subclient.
- ▶ Load balancing (round robin) between alternative data paths provides the mechanism to evenly distribute backup operations between available resources.

If a storage policy is created during the library configuration process, a default data path is created for the primary copy. The default data path is created by using the MediaAgent, Library, Drive Pool, and default Scratch Pool combination for drive pools that are configured in the library. If you create a storage policy, you must specify a Library, MediaAgent, Drive, and Scratch Pool combination for the primary copy.

To update the properties, follow these steps:

1. In the CommCell browser window, navigate to **Storage Policies**.
2. Right-click the policy that you need to update.
3. Select **Properties**. Additional data paths for the primary copy can be defined on the Data Paths tab of the Copy Properties dialog box (Figure 18-17).²

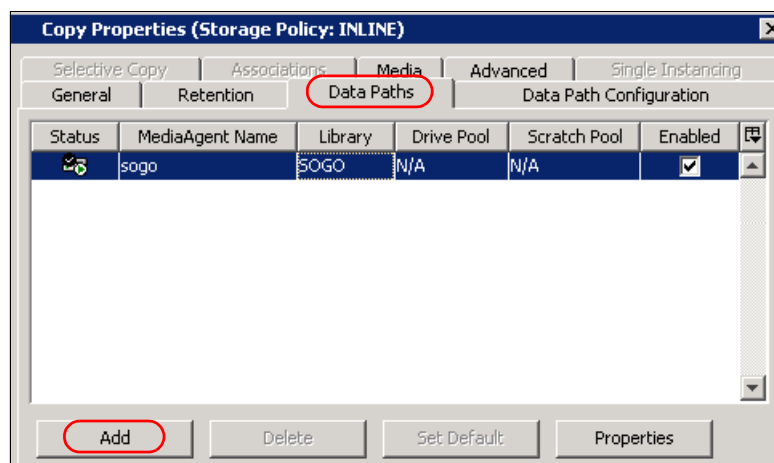


Figure 18-17 Adding new data paths²

18.3 Commvault FSI

This section provides steps and preferred practices for configuring and setting up Commvault for backup and restore. It also provides Commvault parameters and settings for best performance with ProtecTIER FSI (versions earlier than 3.4) for Windows based servers through the CIFS protocol (FSI-CIFS) and for UNIX clients through the NFS protocol (FSI-NFS).

Note: ProtecTIER GA Version 3.4 was released with only the Virtual Tape Library (VTL) interface support. File System Interface (FSI) support was added to ProtecTIER PGA 3.4 Version. For details, see the announcement letter:

http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/6/872/ENUSA P16-0076/index.html&lang=en&request_locale=en

18.3.1 Setting up backup and restore in a CIFS environment

This section provides steps and preferred practices for configuring and setting up Commvault for backup and restore in an FSI-CIFS environment.

Adding a disk library

To add a disk library, complete the following steps:

1. From the Control Panel, double-click **Library and Drive Configuration**.
2. Select the MediaAgents whose devices you want to detect or display from the Available MediaAgents pane.
3. Click **Add >>** to move the MediaAgents to the Selected MediaAgents pane.

² Source: <http://www.commvault.com>

4. Click **OK**.
5. From the Library and Drive Configuration window, click the **Start** menu, select **Add**, and then click **Disk Library** from the menu.

Adding a shared mount path

To add a shared mount path, complete the following steps:

1. In the Add Mount Path dialog box, click **Network Path**. Complete the Connect As, Password, and Verify Password fields with the information of the CIFS share user account. Write permission is defined when you set CIFS authentication to either Workgroup or Active Directory.
2. Click **Network Path** → **Folder** and enter \\FSI_IP\CIFS_name. This example uses \\9.11.109.130\bpgc.

Note: When you specify the network path for the mount point, use host name aliases rather than static IP addresses. By specifying the host name (that is, from /etc/hosts or c:\windows\system32\System32\drivers\etc), you have flexibility if the IP address of the ProtecTIER FSI changes or you want to switch from the primary copy to a copy of the file system on a remote ProtecTIER system.

Defining the BackupSet and associating it with the storage policy

To define a BackupSet and associate it with a storage policy, complete the following steps:

1. Collapse the **Client Computers** menu in the CommCell Browser
2. Choose the BackupSet you want to use, highlight the subclient, and click **Properties**.
 - a. Edit the content, and add the paths that you want to back up.
 - b. Select the Storage Policy name, which is automatically created you add a disk library, from the **Storage Policy** drop-down list in the **Data Storage Policy** tab.

Performing a backup

To perform a backup, complete the following steps:

1. Choose the BackupSet, right-click the subclient that you want to back up, and choose **Backup**.
2. Click **Full and Immediate** backup, and then click **OK**.

Performing a restore

To perform a restore, complete the following steps:

1. From the CommCell Browser, navigate to **Client Computers** → **Client** → **File System** → **defaultBackupSet**.
2. Right-click the default subclient and then click **Browse Backup Data**.
3. Expand **defaultBackupSet** and right-click the folder that you want to restore.
4. Click **Restore Current Selected** or **Restore All Selected**.

If you want to restore to the same folder or another CIFS share folder, complete the User Name, Password, and Confirm Password fields with the information of the CIFS share user account with write permission. Write permission is defined when you set CIFS authentication to either Workgroup or Active Directory.

Complete the following steps:

1. From the Restore Options for All Selected Items dialog box, click **Advanced**.
2. The Advanced Restore Options window opens. Click **Advanced Restore Options**, click the **General** tab, and click the **Impersonate User** option.
3. In the User Name and Password boxes, enter a user name and password that has access privileges. In the Confirm Password box, type the password again. Click **OK**.

18.3.2 Parameters for best performance with ProtecTIER FSI-CIFS

This section suggests several parameters for the best performance with ProtecTIER FSI.

Clearing hardware compression/using hardware encryption

To disable hardware encryption, complete the following steps:

1. Right-click the storage policy copy and click **Properties**.
2. Click the **Data Paths** tab.
3. Click the data path for which you want to change the hardware compression and then click **Properties**.
4. From the Data Path Properties dialog box, clear the **Hardware Compression** check box to disable hardware compression. Clear the **Use Hardware Encryption** check box to disable hardware encryption (Figure 18-18).

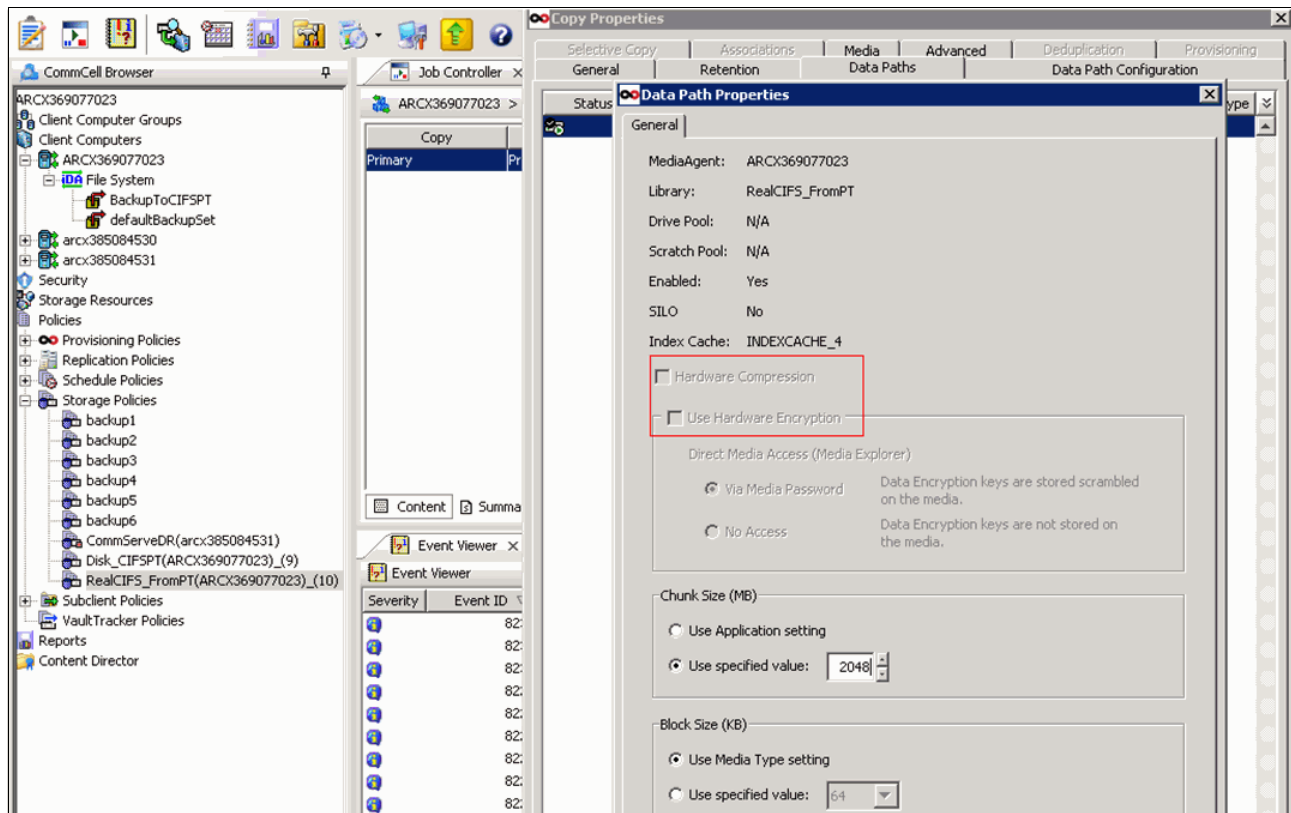


Figure 18-18 Clearing/enabling hardware compression

5. Click **OK** to save your changes.

Disabling data multiplexing

To disable data multiplexing, complete the following steps:

1. In the right pane of the CommCell Browser, right-click the storage policy copy for which you want to configure multiplexing, and select **Properties**.
2. From the Copy Properties (Media) window, clear the **Enable Multiplexing** check box (Figure 18-19).

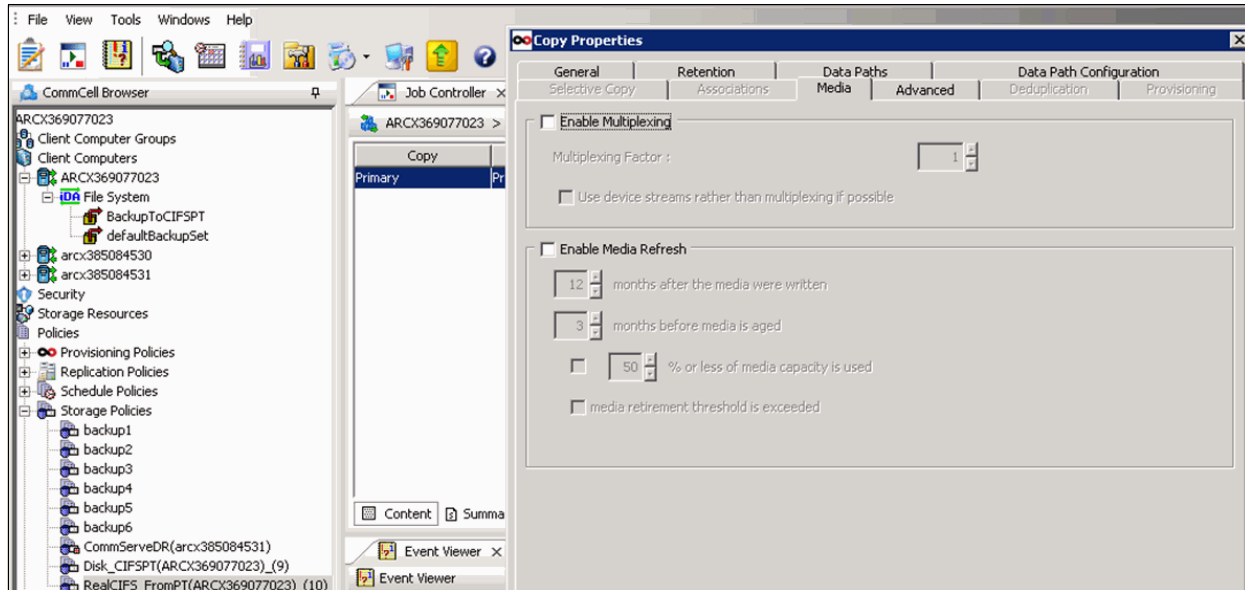


Figure 18-19 Clear Enable Multiplexing

3. Click **OK** to save your changes.

Disabling software compression

To disable software compression, complete the following steps:

1. From the CommCell Browser, right-click the subclient for which you want to disable software compression and then click **Properties**.
2. Click the **Storage Device** tab and under the Data Storage Policy tab, select the storage policy from the Storage Policy list.

3. Click the **Data Transfer Option** tab and select **Off** for the Software Compression option for this subclient (Figure 18-20).
4. Click **OK** to save your changes.

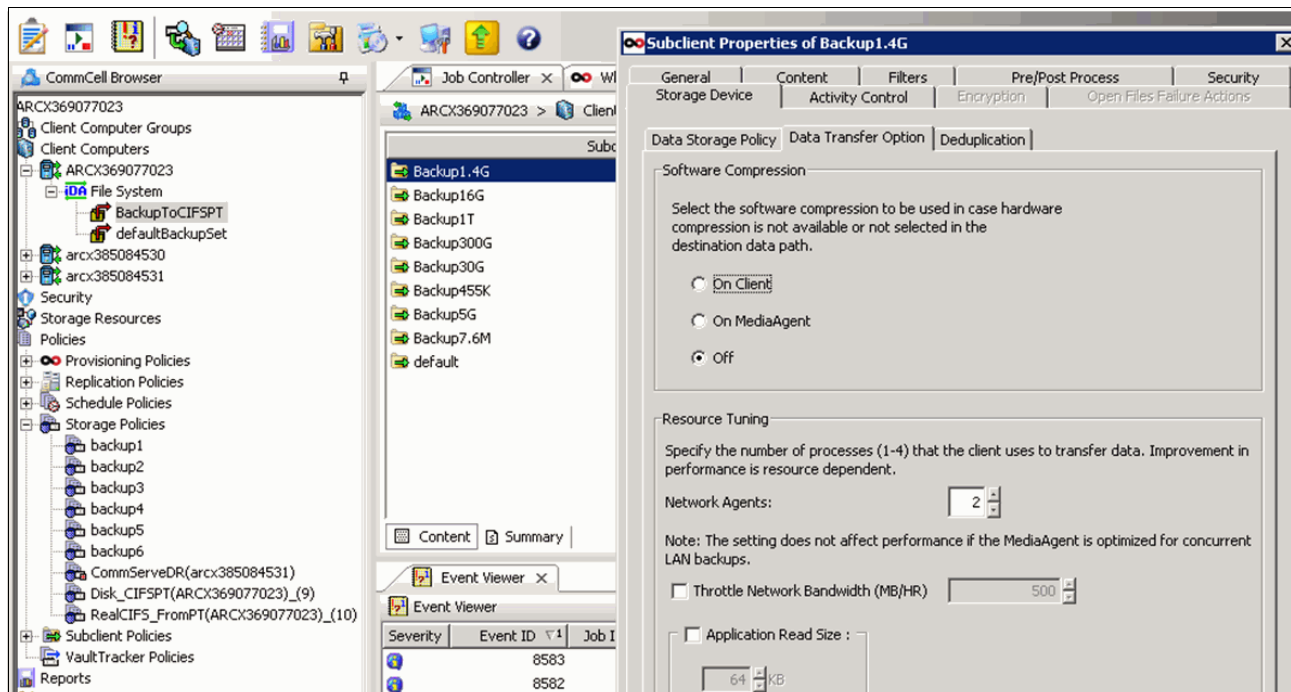


Figure 18-20 Set software compression off

Disabling deduplication

To disable deduplication, complete the following steps:

1. From the CommCell browser, right-click the subclient for which you want to disable software compression and then click **Properties**.
2. Click the **Storage Device** tab and, from the Data Storage Policy tab, select the storage policy from the **Storage Policy** list.
3. Click the **Deduplication** tab and clear the **Enable Deduplication** check box for this subclient (Figure 18-21).
4. Click **OK** to save your changes.

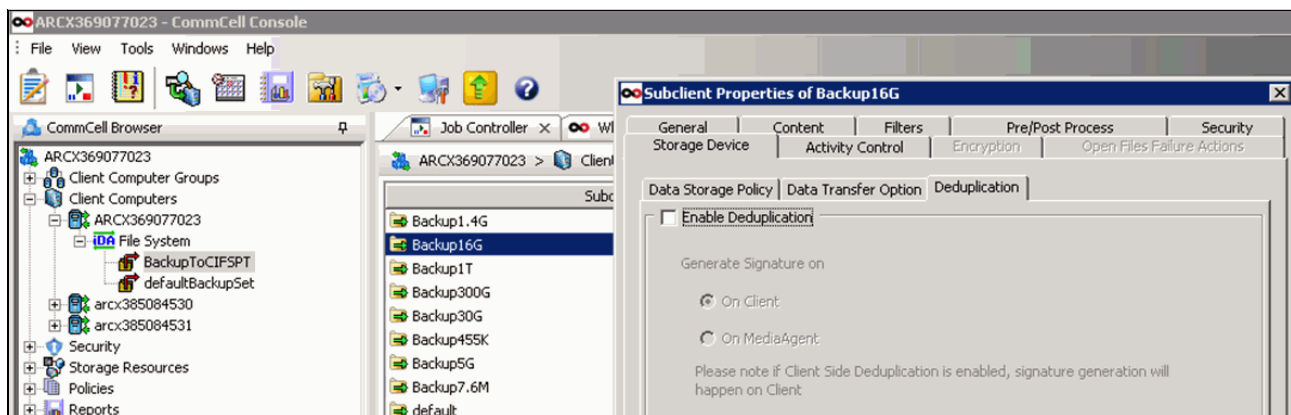


Figure 18-21 Disable deduplication

Configuring multiple streams for backups and changing the maximum number of data streams

You do not have to enable multistreaming if you have multiple mount points that point to the same physical drive. For detailed steps regarding configuring multistreaming and changing the maximum number of data streams, see the following web page:

http://documentation.commvault.com/commvault/release_9_0_0/books_online_1/english_us/prod_info/windows.htm

You can configure the automatic definition of alternative data paths if multiple MediaAgents share the ProtecTIER VTL:

1. In the CommCell browser window, navigate to **Storage Policies**, right-click the policy that you need to update, and select **Properties**.
2. Navigate to the **Data Path Configuration** tab of your Storage Policy **Copy Properties** window (Figure 18-22).

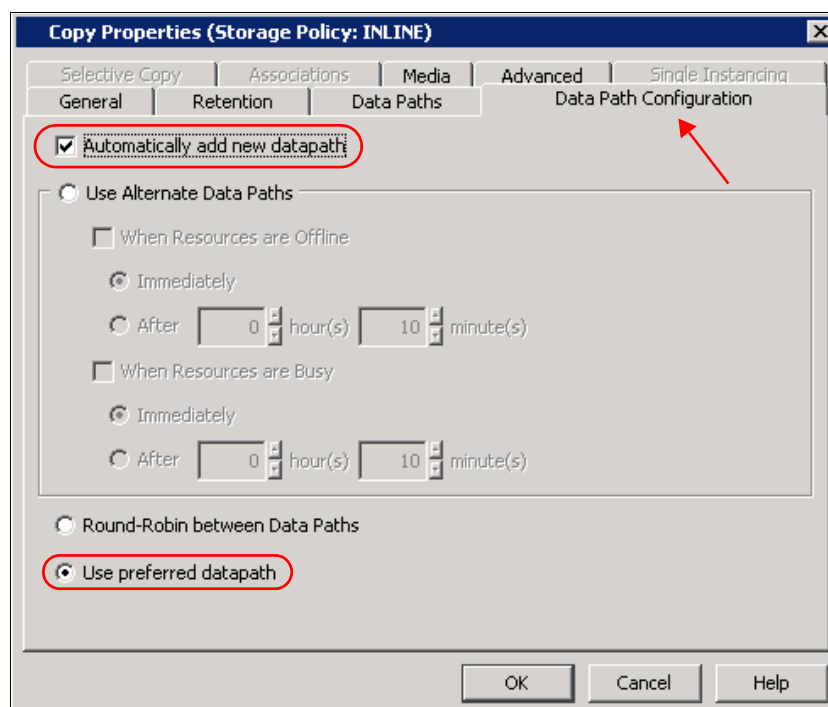


Figure 18-22 Automatically add new data paths

Always use the **Use preferred data path** option when you define new alternative data paths. Avoid using round-robin load balancing between the resources, because it has a negative effect on the factoring ratio. This negative effect occurs because the load balancing is done on the tape drive pool level, not on the MediaAgent level.

For more information about configuring alternative data paths, see the Commvault documentation:

http://documentation.commvault.com/commvault/release_9_0_0/books_online_1/

18.3.3 Setting up backup and restore in an NFS environment

This section provides steps and best practices for configuring and setting up Commvault for backup and restore in an FSI-NFS environment. It also provides Commvault parameters and settings for best performance with ProtectTIER FSI-NFS.

Adding a disk library

To add a disk library, complete the following steps:

1. From the Control Panel, double-click the **Library and Drive Configuration**.
2. Select the MediaAgents whose devices you want to detect or display from the Available MediaAgents pane.
3. Click **Add >>** to move the MediaAgents to the Selected MediaAgents pane and click **OK**.
4. From the Library and Drive Configuration window, click the **Start** menu, select **Add**, and then click **Disk Library** from the menu.
5. Write an alias for the new Disk Library and check the **Automatically create storage policy for the new data paths** option.
6. Click **OK**. The Add Mount Path dialog box opens for the Adding a shared mount path option.
7. In the Add Mount Path dialog box, leave the **Disk Device** and **Base Folder** options as they are. Confirm that in the MediaAgent box the correct client is selected.
8. Click Local Path Folder and select the */<mountpoint>* of the NFS export in the new window that appears.
9. Click **OK**.

Defining the BackupSet and associating it with the storage policy

To define a BackupSet and associate it with a storage policy, complete the following steps:

1. From the CommCell Browser, navigate to **Client Computers** → **Client** → **FileSystem**. Right-click **FileSystem** and select **All Tasks** → **Create new Backup Set**.
2. Input the name of the new BackupSet.
3. Select the Storage Policy name, which is automatically created when you add a disk library, from the **Storage Policy** drop-down list, and click **OK**.
4. Choose whether to create a backup schedule. For example, choose **Do not schedule**, and click **OK**.
5. Right-click the **BackupSet** created and select **All Tasks** → **Create new Backup Subclient**.
6. Input the name of the new subclient under the General tab of the subclient properties dialog box.
7. Add the file path that you want to back up under the Content tab of the subclient properties dialog box.
8. Under the **Storage Device** → **Data Storage Policy** tab of the subclient properties dialog box, select the Storage Policy name, which is automatically created when you add a disk library, from the **Storage Policy** drop-down list.
9. Click **OK**.

Preferred practices for BackupSets:

- ▶ Have a total of 16 streams that are used by all BackupSets (one or more) with ProtecTIER FSI-NFS.
- ▶ The total number of streams per one BackupSet is the sum of all of the streams started by every subclient that is defined in this BackupSet.
- ▶ By default, each subclient uses one stream. To change this default and configure the number of streams per subclient, you can perform the following steps:
 - a. Click the **General** tab under the subclient properties dialog box.
 - b. Select the **Allow multiple data readers in a drive or mount point** option.
 - c. Input the number of streams that you want into the Number of Data Readers option.
- ▶ Make sure to configure the correct number of streams per individual subclient, the correct number of subclients per individual BackupSet, and then the correct number of BackupSets to align with the suggestion of the total number of streams being 16.
- ▶ If you need to run more than 16 streams per BackupSet, be aware that the maximum number of streams that are supported by each BackupSet is 16. This is a Simpana limitation.

Disabling hardware and software compression and encryption

To disable hardware compression and encryption, complete the following steps:

1. In the CommCell Browser's right pane, click **Policies** → **Storage Policies**. Right-click the storage policy that was created before and click **Properties**.
2. Click the **Data Paths** tab.
3. Click the data path for which you want to change the hardware compression, and then click **Properties**.
4. From the Data Path Properties dialog box, clear the **Hardware Compression** check box to disable hardware compression. Clear the **Use Hardware Encryption** check box to disable hardware encryption.
5. Click **OK** to save your changes.

To disable software compression, complete the following steps:

1. From the CommCell Browser, right-click the subclient for which you want to disable software compression and then click **Properties**.
2. Click the **Storage Device** tab and, under the Data Storage Policy tab, select the storage policy from the Storage Policy list.
3. Click the **Data Transfer Option** tab and select **Off** for the Software Compression option for this subclient.
4. Click **OK** to save your changes.

Configuring the chunk size

To configure the chunk size, complete the following steps:

1. In the CommCell Browser's right pane, click **Policies** → **Storage Policies**. Right-click the storage policy that was created before and click **Properties**.
2. Click the **Data Paths** tab.
3. Click the data path that you want to change and then click **Properties**.

4. From the Data Path Properties dialog box, select **Use specified value** option in **Chunk Size** area and set it to 8192 MB.

Note: The Chunk size suggestion is no longer relevant when using ProtecTIER V3.3.3 or later. From this code level and later, there is no suggestion to limit the backup file size, and in fact the bigger the better.

5. Click **OK** to save your changes.

Disabling data multiplexing

To disable data multiplexing, complete the following steps:

1. In the right pane of the CommCell Browser, right-click the storage policy copy for which you want to configure multiplexing, and select **Properties**.
2. From the Copy Properties window, clear the **Enable Multiplexing** check box.
3. Click **OK** to save your changes.

Disabling de-duplication

To disable deduplication, complete the following steps:

1. From the CommCell browser, right-click the subclient for which you want to disable software compression and then click **Properties**.
2. Click the **Storage Device** tab and, from the Data Storage Policy tab, select the storage policy from the Storage Policy list.
3. Click the **Deduplication** tab and clear the **Enable Deduplication** check box for this subclient.
4. Click **OK** to save your changes.

Configuring device streams

To configure device streams, complete the following steps:

1. In the CommCell Browser's right pane, select **Policies Storage Policies**. Right-click the storage policy that was created before and click **Properties**.
2. Click the **General** tab and set **Device Streams** to 16.
3. Click **OK** to save your changes.

Notes:

- Commvault requires that the number of streams that are configured in the Storage Policy should be equal to or greater than the specified number of data readers that are defined per subclient. If you need more than 16 streams in your subclients, this setting should be adjusted.
- You do not have to enable multistreaming if you have multiple mount points that point to the same physical drive. For detailed steps about configuring multistreaming and about changing the maximum number of data streams, contact Commvault support or review the documentation available at the Commvault web page:

<http://bit.ly/25aJEIT>

Configuring semaphores of media agent OS

For detailed steps about configuring the number of Linux Semaphores, contact Commvault support or review the documentation available at the Commvault web page:

<http://bit.ly/1oYjY97>

Note: This setting is required only on a Linux Media Agent system.

Performing a backup

To perform a backup, complete the following steps:

1. From the CommCell Browser, click **Client Computers** → **Client** → **FileSystem**.
2. Choose the **BackupSet**, right-click the subclient that you want to back up, and select **Backup**.
3. Click **Full and Immediate backup**, and then click **OK**.

Performing restore

To perform a restore, complete the following steps:

1. From the CommCell Browser, click **Client Computers** → **Client** → **FileSystem** → **defaultBackupSet**.
2. Right-click the default subclient, and then click **Browse Backup Data**.
3. Select **Browsing time** or **Browse the Latest Data backedup**. Click **OK**.
4. Expand **Default Backup Set** and right-click the folder that you want to restore.
5. Click **Restore Current Selected** or **Restore All Selected**.
6. In the window that opens, choose whether you want to overwrite files and the restore destination, either to the same folder or a different destination path.
7. Click **OK**.

18.3.4 Parameters for best performance with ProtecTIER FSI-NFS

Table 18-1 summarizes the required and suggested settings for ProtecTIER best practices with Commvault.

Table 18-1 Summary of settings for ProtecTIER with Commvault

Component	Parameter	Value
Data path properties	Hardware Compression	Disabled (Not checked)
	Use Hardware Compression	Disabled (Not checked)
	Chunk Size use specified value	Enabled (Checked) 8192 MB ^a (8 GB)
Copy properties	Enable Multiplexing	Disabled (Not checked)
Data Transfer option	Software Compression	Off
Deduplication tab	Enable Deduplication	Disabled (Not Checked)
General tab of the Storage Policy Properties	Device Streams	16 ^{b c}
Media Agent OS Properties	Number of Linux Semaphores	Required only for Linux Media Agent system ^d

a. The Chunk size suggestion is no longer relevant when using ProtecTIER pGA V3.3.3 or later. From this code level and later there is no suggestion to limit the backup file size, and in fact the bigger the better.

b. The number of streams that are configured in the Storage Policy should be equal to or greater than the specified number of data readers.

c. You do not have to enable multistreaming if you have multiple mount points that point to the same physical drive. For detailed steps about configuring multistreaming and about changing the maximum number of data streams, contact Commvault support or review the documentation available on the Commvault website: <http://bit.ly/1ucJhpo>

d. This setting is only required for Linux Media Agent system. For details about configuring the number of Linux Semaphores, contact Commvault support, or review the documentation available on the Commvault website: <http://bit.ly/1oYjY97>



Veeam FSI-CIFS

This chapter describes steps and preferred practices for configuring and setting up Veeam Backup & Replication for backup and restore. The chapter also provides parameters and settings for best performance with ProtecTIER File System Interface-Common Internet File System (FSI-CIFS).

Before continuing, be sure that your ProtecTIER System for FSI-CIFS is configured and that you completed the following tasks:

- ▶ A file system is created.
- ▶ CIFS authentication is configured.
- ▶ CIFS Share is created.

For general information about how to configure the ProtecTIER System for FSI, review Chapter 5, “ProtecTIER File System Interface: General introduction” on page 65.

Information about Veeam Backup & Replication is at the following web page:

<https://www.veeam.com/vm-backup-recovery-replication-software.html>

Note: ProtecTIER GA Version 3.4 was released with only the Virtual Tape Library (VTL) interface support. File System Interface (FSI) support was added to ProtecTIER PGA 3.4 Version. For details, see the announcement letter:

http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/6/872/ENUSA P16-0076/index.html&lang=en&request_locale=en

This chapter describes the following topics:

- ▶ Setting up backup infrastructure
- ▶ Setting up backup and restore
- ▶ Parameters for best performance with ProtecTIER-CIFS
- ▶ Advanced backup repository settings
- ▶ Summary of parameters for best performance

19.1 Setting up backup infrastructure

Before creating backup and restore jobs, you must plan and set up a backup infrastructure. A typical Veeam backup infrastructure includes the following components:

- ▶ Backup repository
- ▶ Backup proxy
- ▶ Virtual servers (ESX(i))

This section provides steps to set up these components for the Veeam backup infrastructure.

19.1.1 Creating backup repository

A backup repository is a location for storing backup data on the target side. This section describes how to configure a Windows-based backup repository that is built on a ProtecTIER FSI-CIFS share.

To create the backup repository, complete the following steps:

1. Start the **Veeam Backup & Replication** console.
2. In the left pane, open the **Backup Infrastructure** view.
3. Right-click the **Backup Repositories** node and select **Add Backup Repository**.
4. In the wizard, specify a name for the repository and a description for future reference. Click **Next**.
5. Choose **Shared folder** as the type of backup repository you want to create. Click **Next**.
6. In the **Shared folder** field (Figure 19-1), specify the Universal Naming Convention (UNC) path to the share folder and select the **This share requires access credentials** check box. Enter an account with administrative privileges on the share. Click **Next**.

The screenshot shows the 'New Backup Repository' wizard in the Veeam Backup & Replication console. The 'Share' step is active, showing a 'Shared folder' field with the UNC path '\\9.11.109.144\\Veeam'. The 'This share requires access credentials' checkbox is checked. The 'Username' field contains 'PTFSI\\administrator' and the 'Password' field contains '[To change the saved password, click here]'. The 'Write data to this share' section has 'Directly from backup proxy server' selected. Navigation buttons at the bottom include '< Previous', 'Next >', 'Finish', and 'Cancel'.

Figure 19-1 Creating a backup repository

7. For the Repository location (Figure 19-2), click **Populate** to calculate the capacity and free space on the selected share folder.

Note: A suggestion is to set the **Limit maximum concurrent jobs count** parameter to 16.

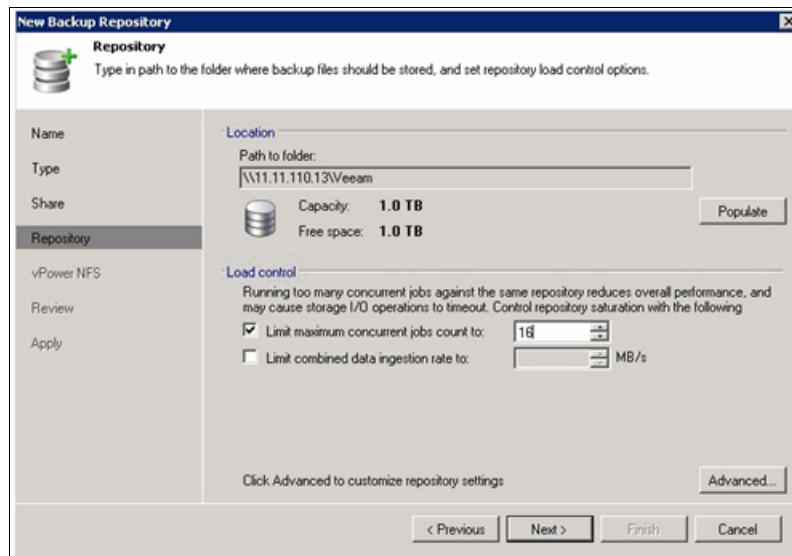


Figure 19-2 Setting maximum concurrent jobs count

8. Click **Advanced** to select the following two parameters:
 - Align backup file data blocks.
 - Decompress backup data blocks before storing.

Note: Review Table 19-1 on page 291 for suggested settings of these parameters.

9. Click **Finish**.

19.1.2 Adding VMware backup proxy

To add a VMware backup proxy, complete the following steps:

1. Start the **Veeam Backup & Replication** console.
2. In the left pane, open the **Backup Infrastructure** view.
3. Right-click the **Backup Proxies** node and select **Add VMware backup proxy**.
4. From the **Choose server** list, select **any other server** as a new backup proxy.

To select another server as a backup proxy, you can click **Add new** to open the **New Windows Server** wizard to add a new server.

5. Click **Finish**.

Tip: For best performance, a suggestion is for the server that is installed with Veeam proxy to have a direct storage area network (SAN) connection to the virtual machine disk.

19.1.3 Adding a VMware Server

To add a VMware Server, complete the following steps:

1. Start the **Veeam Backup & Replication** console.
2. In the left pane, open the **backup Infrastructure** view.
3. Right-click the **Managed servers** node and select **Add server**.
4. In the **Add server** window, select **VMware vSphere**.
5. Enter the DNS name or IP address of the VMware vCenter server in the **DNS name or IP address** field and provide a description. For example: 10.0.0.1. Click **Next**.
6. Enter a user account with local administrator privileges on the server you are adding, as shown in Figure 19-3.




Figure 19-3 Adding VMware server

7. Click **Finish**.

19.2 Setting up backup and restore

This section provides steps for configuring Veeam Backup & Replication V6.5 to perform backup and restore in a ProtecTIER FSI-CIFS environment.

19.2.1 Creating a new backup job for backup

To create a new backup job, complete the following steps:

1. Start the **Veeam Backup & Replication** console.
2. In the left pane, open the **Backup & Replication** view.
3. Right-click the **Job** node and select **Backup** to start a new backup job wizard.
4. Enter a name of new backup job in the **Name** field and provide a description. Click **Next**.

- Click **Add** to select VMs that should be backed up and click add. Use the **Recalculate** button to refresh the total size value after you add new VMs to the job (note that it is suggested to have one job per VM). Click **Next**.
- From the **Backup repository** list, select a repository that you previously created to store the created backup files (Figure 19-4).




Figure 19-4 Creating backup job

- Click **Finish**

19.2.2 Performing backup

To perform a backup, complete the following steps:

- Start the **Veeam Backup & Replication** console.
- In the left pane, open the **Backup & Replication** view.
- Select the node **Jobs** → **Backup**, right-click a backup job in the list, and select **Start** to manually run the backup.

Note: The first time that you run the backup in this way, it is a full backup; after that, it will be an incremental backup. If you want to run a full backup again, right-click a backup job and select **Active Full**.

19.2.3 Performing restore

To perform a restore, complete the following steps:

- Start the **Veeam Backup & Replication** console.
- On the **Home** tab, click **Restore** and select **VMware**.
- In the Restore from backup section, choose **Entire VM**. Click **Next**.
- Click **Add VM** → **from backup** to select the wanted VM in the list of available backup jobs.
- Click **Point** to select the wanted restore point for the VM. Click **Next**.

6. Select **Restore to a new location, or with different settings**. Click **Next**.
7. In the **Host** field, specify a host on which the VM should run.
8. In the **Resource pool** list, select a resource pool to which the VM should be recovered.
9. In the **Datastore** field, select which disk and disk type should be used to restore the VM.
10. In the **Folder** field, change a new name under which the VM should be restored and registered. Click **Next**.
11. Select the **Power on VM after restoring** check box to start the VM immediately after recovery.
12. Click **Finish** to start restore.

19.3 Parameters for best performance with ProtecTIER-CIFS

For best performance with ProtecTIER-CIFS, use the suggested parameters that are described in this section.

19.3.1 Disabling data de-duplication and compression

By default, the **Enable inline data deduplication** check box is selected and **Compression** is optimal. These settings have a negative effect on the ProtecTIER deduplication factoring ratio. Therefore, disable both **Enable inline data deduplication** and **Compression**.

To disable those settings, complete the following steps:

1. Start the **Veeam Backup & Replication** console.
1. In the left pane, open **Backup & Replication** view and select **Backup** node.
2. Right-click a backup job and select **Edit**.
3. In the Storage field, click **Advanced**.
4. Click the **Storage** tab (Figure 19-5) and make the following changes:
 - a. Clear the **Enable inline data deduplication** check box to disable data deduplication.
 - b. Select **None** for the compression Level to disable data compression.

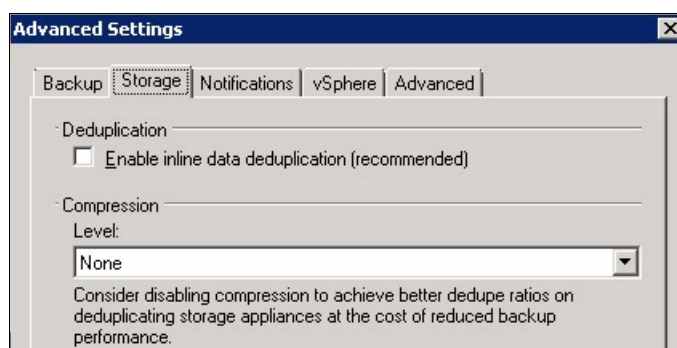


Figure 19-5 Disabling data deduplication and compression

19.3.2 Setting data block size for backup to SAN or local storage

To set the data block size for backup to SAN or local storage, complete the following steps:

1. Start the **Veeam Backup & Replication** console.
2. In the left pane, open **Backup & Replication** view and select **Backup** node.
3. Right-click a backup job and select **Edit** to start the Edit Backup Job wizard.
4. In the Storage field, click **Advanced**.
5. Under Storage optimizations (Optimize for), select **Local target** (Figure 19-6) for backup to SAN, direct-attached storage (DAS), or local storage. Click **OK**.




Figure 19-6 Setting data block size

19.3.3 Setting Incremental backup mode

Veeam Backup & Replication provides two methods for storing backup files:

- ▶ Reversed incremental backup
- ▶ Forward incremental backup

Reversed incremental backup

Reversed incremental backup will rebuild a full backup file for each incremental backup. During incremental backup, Veeam Backup & Replication copies only the data blocks that have changed since the last job, and then “injects” these changes into the previous full backup file (.vbk) to rebuild it to the most recent state of a VM. Meanwhile, it creates a reversed incremental file (.vrb) that contains the block data replaced when rebuilding the full backup file.

Forward incremental backup

Forward incremental backup as a normal incremental backup contains only changed block data since the last job (whether full or incremental) and creates backup files (.vib) until the next full backup.

Generally, the reversed incremental method is more convenient for restore. However, the forward incremental method will have better performance results than the reversed incremental method for backups. Therefore a suggestion is to use the forward incremental method with ProtecTIER.

To use the forward incremental method complete the following steps:

1. Start the **Veeam Backup & Replication** console.
2. In the left pane, open the **Backup & Replication** view and select **Backup** node.
3. In the backup list, right-click a backup job and select **Edit**.
4. Select the **Storage** tab and click **Advanced**.

5. Select the **Incremental** and **Perform active full backups periodically** check boxes, as shown in Figure 19-7.




Figure 19-7 Setting incremental backup mode

Tip: For best performance, disable *synthetic* backup, which performs incremental backups after the first full backup, and instead perform active full backup weekly because completing a synthetic backup takes more time.

19.4 Advanced backup repository settings

This section describes how to set up data blocks alignment and decompression methods.

19.4.1 Setting data blocks alignment

For storage systems using fixed block size, select the **Align backup file data blocks** check box. Veeam Backup & Replication will align VM data saved to a backup file to a 4 kilobits (Kb) block boundary, providing better deduplication across backup files.

19.4.2 Setting data blocks decompression

If backup data is compressed at the source side before writing to the target, select the **Decompression backup data blocks before storing** check box. Veeam Backup & Replication will compress VM data, transmit it over the local area network (LAN), decompress data on the target side, and write raw VM data to the storage device to achieve a higher deduplication ratio.

For more information about setting these parameters, see 19.1.1, “Creating backup repository” on page 284.

Figure 19-8 shows how to configure the advanced backup repository settings.

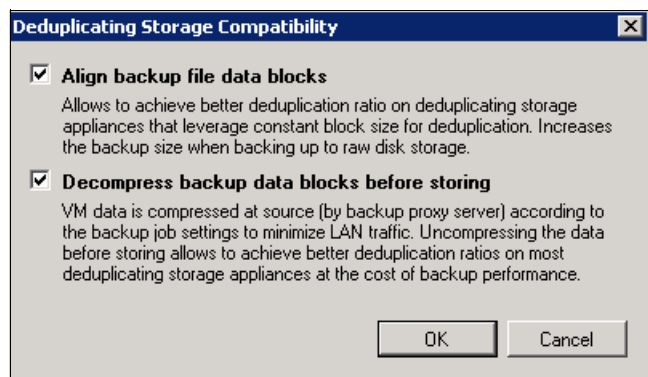


Figure 19-8 Setting advanced backup repository

19.5 Summary of parameters for best performance

Table 19-1 summarizes the required and suggested settings for ProtecTIER preferred practices with Veeam Backup & Replication.

Table 19-1 Summary of parameters for best performance

Component	Parameter	Value
Backup job: Storage definition	Deduplication	Disabled (Not selected)
	Compression level	None
	Storage optimizations	Local target
Backup job: Backup definition	Backup Mode	Forward Incremental
	Active Full backup	Enable (Selected)
Backup Repository	Limit Maximum concurrent job count	16
	Align backup file data blocks	Enable (Selected)
	Decompress backup data blocks	Enable (Selected)



Part 3

Application considerations

This part describes settings and parameters that are modified for optimum deduplication ratios and performance when you work with specific data types. The applications that it focuses on are RMAN Oracle, IBM Domino, Microsoft Exchange, Microsoft SQL Server, IBM DB2, and VMware.

This part contains the following chapter:

- Chapter 20, “Application considerations and data types” on page 295.



Application considerations and data types

This chapter described guidelines for the settings and parameters that are modified in your applications, and specific data types for optimal performance and deduplication factoring ratios.

This chapter describes the following topics:

- ▶ IBM Domino
- ▶ Microsoft Exchange
- ▶ Microsoft SQL Server
- ▶ DB2
- ▶ Oracle
- ▶ SAP
- ▶ VMware

Readers of the relevant sections should be familiar with the backup and restore concept of the managed application or data type. Therefore, this chapter does not provide steps for configuring backup applications.

Notes:

- ▶ ProtecTIER GA Version 3.4 was released with only the Virtual Tape Library (VTL) interface support. File System Interface (FSI) support was added to ProtecTIER PGA 3.4 Version. For details, see the announcement letter:
http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/6/872/ENUSAP16-0076/index.html&lang=en&request_locale=en
- ▶ Beginning with ProtecTIER V3.4, the OpenStorage (OST) interface is not supported.
- ▶ Beginning with version 7.1.3, Tivoli Storage Manager was rebranded to IBM Spectrum Control. The scenarios in this chapter were conducted with a version prior to 7.1.3, so this chapter uses the name Tivoli Storage Manager for legacy and reference purposes.

20.1 IBM Domino

This section describes the settings and parameters that should be modified in IBM Domino environments to enable the optimal factoring for the ProtecTIER product.

20.1.1 Common server

IBM Domino employs a common email or application server for several members of the company or network. Clients usually run backup policies from the common server (the Domino server) that stores all the data on the physical or virtual tape.

Domino servers in enterprise or secured environments are configured in *application clusters*. In contrast, the server-based clusters and the shared storage resources that are assigned to the application are available on both (or more) cluster nodes simultaneously. Access to the them is fully controlled by the clustered applications, in this case, Domino servers.

A Domino mail environment typically has a configuration of active-active clusters, where each Domino server is always active only on a dedicated node of the dual-node cluster, never fails over, and both Domino applications (that is, the applications running on the server //cluster) control the same storage resources. However, only the dedicated portion of the databases (set of mail files) is served at one time by the single node.

The common understanding of the application failover to the standby cluster node does not apply in Domino environments. If there is a node failure, the application that is running on a live node takes full management control over all portions (all mail files) of storage resources instantaneously.

From the Domino functional perspective, the following categories of Domino server installations exist:

- ▶ An email server that supports IBM Lotus Notes, IMAP, POP3, SMTP, and WebMail access (IBM iNotes®).
- ▶ An application server where the Lotus Notes client provides the application run time.
- ▶ A database server that offers Notes Storage Facility.
- ▶ A web server that enables Lotus Notes clients to access the data through a web browser.
- ▶ A directory server for authentication services (hub/gateway).
- ▶ Instant messaging and web conferencing, also known as IBM Sametime®.

This section focuses on email, application, and database servers, which usually hold the most amount of data in Domino server repositories. The other listed features are highly transactional services with small amounts of data, and are therefore not optimal candidates for ProtecTIER deduplication.

20.1.2 Existing backup and disk space usage

Running the existing backup commands and using the general suggested methods cause a low factoring ratio of the data even if the change is low. These actions reduce the benefit of using the ProtecTIER solution and disk space usage.

The root cause is the inherent compaction of the database, which reshuffles Notes Storage Format (NSF) files inside the database. Although this function reduces space usage from the perspective of Domino, it also changes the layout and data pattern of every NSF.

The next time that the ProtecTIER server receives blocks from these databases, they all look unique, so the factoring ratio is low. Ratios of 1:2 - 1:3 are possible in environments with Domino Version 7. However, running compaction that is based on the Domino **DELETE** operation is a preferred practice for Domino, so disabling it is not a solution.

Simply stated, the compaction saves the primary expensive storage on Domino server and increases the performance of mailbox operation, especially in clustered environments, so there is no single client that wants to disable it.

regarding deduplication efficiency, another factor to consider is the method that is used to store email attachments. Working documents are compressed and archived by one of the widely available compression tools, and are converted to various file formats (.zip, .rar, .tar, .7z, and others), which do not factor optimally. The same is true for media files in email attachments, such as compressed pictures (.jpg and .gif), movies (.mpg, .mov, and .avi), or music (.mp3).

20.1.3 Domino attachments and object service

Deduplication can be improved by using a feature in the Domino environment: *Domino Attachment and Object Service* (DAOS). DAOS removes all the email or application attachments from the Notes Storage Format (NSF) files and stores them separately in the server file system in a *single occurrence* as a Notes Large Object (NLO). This feature has been available since Domino Version 8.5.

See examples of a dedicated storage repository for NSF and NLO objects:

- ▶ Figure 20-1 shows the location of mail files in Domino storage.
- ▶ Figure 20-2 on page 299 shows the folder structure of DAOS attachments.

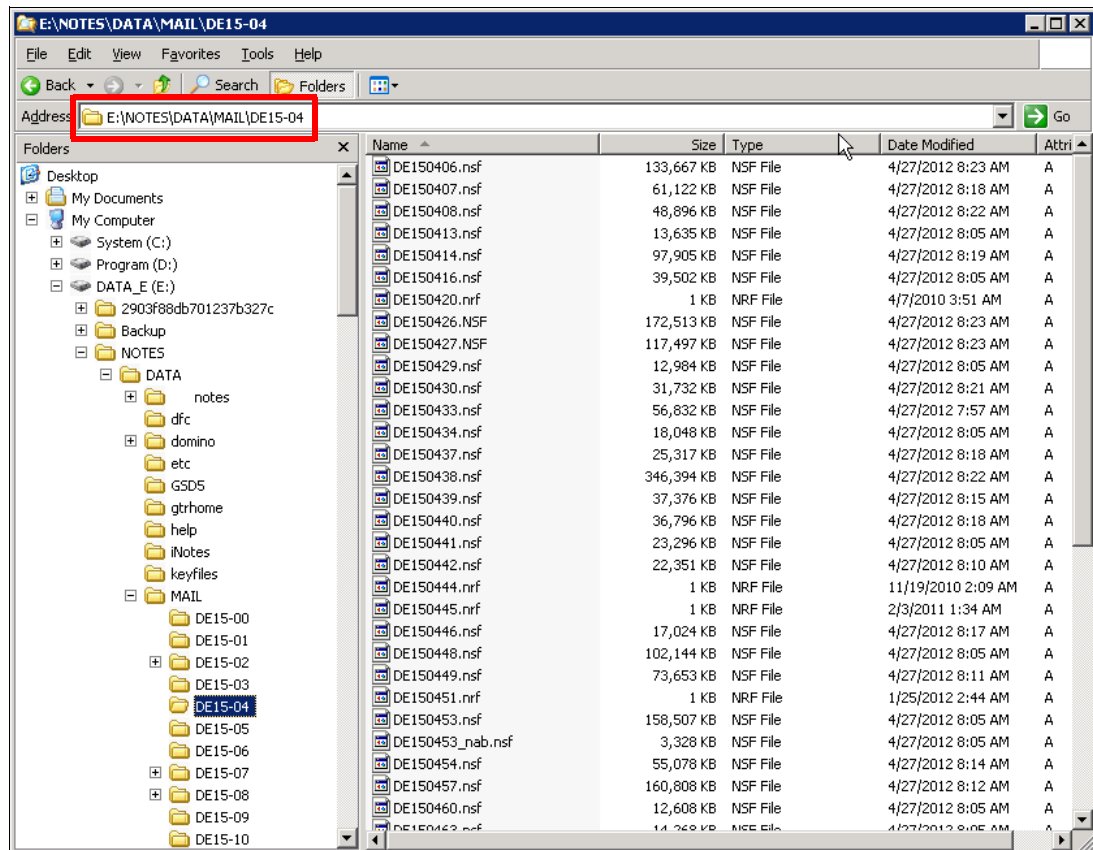


Figure 20-1 The location of mail files in Domino storage

Figure 20-2 shows another example of a dedicated storage repository for NSF and NLO objects.




Figure 20-2 The folder structure of DAOS attachments

DAOS divides the database objects into two categories of items:

- ▶ Database items in .nsf files
- ▶ Attachment items in .nlo files (attachment items)

The NLO holds only one instance of each attachment and the multiple NSF files that contain relevant metadata links, and reuses it. DAOS reduces the effect of using the **DELETE** option because the DAOS layout does not hold each attachment multiple times. This arrangement mitigates the compaction effect and the NLO change is marginal.

Backing up the NLO files in the DAOS repository can be done either while the Domino server is down, or when it is up and running. The backup does not require the usage of any Domino API-based utilities. After the NLO files are initially written, Domino never modifies their contents, so the backup mechanism does not work around file-write activity. NLO files can be backed up as with any other generic files in the file system. Only the NLO files that are complete and not in the process of being written to or renamed must be backed up.

Any files that are busy can be skipped until the next backup job runs. Most backup applications automatically skip files that they cannot read because of other activity.

Important: If Domino is running during a backup process (online backup), an important step is to first back up all NSF files before you proceed with the NLO backup because the metadata references in the NSFs are related to the newly detached NLO files.

In typical Domino environments, dedicated and independent processes are used to back up NSF and NLO files. The NLO file is just a flat file on the disk, but the NSF file is considered to be in database format, which means different tools are used to back up each of the files (for example, Tivoli Data Protection for Mail, in contrast with the Tivoli Storage Manager Backup/Archive client). Obviously, operating system flat file backups are not retained in the backup server for the same period as online backups of Domino NSF files.

Important: Ensure that the retention period of NLO file backups is at least the same or longer than the longest retention period used for online backups of NSF files (monthly, quarterly, or yearly backups).

Domino incorporates the feature to keep NLO files on disk for a period after the link to them becomes invalid and the NLO files are not needed anymore. If this DAOS retention period is longer than the backup retention of NSF files in a backup server, the previous statement does not apply. This case is the only exception where backup retention on NLO files does not play a role.

A backup copy is never used for the restoration of mail files with detached data, and the relevant NLO file is still retained on the disk by DAOS. However, the minimal backup retention is still needed to protect data against a disaster or file system corruption.

The disk footprint savings with DAOS apply to the backup processing as well. The NLO files represent the static data that used to be in the NSF, and was backed up every cycle even though it had not changed. In a typical mail environment, a large reduction in the NSF footprint, plus a small amount of NLO data, translates almost directly into a reduction in the backup footprint. In addition to the duplicate data being eliminated, the mail file data is also separated into static and dynamic components.

By applying an incremental backup regimen to the static NLO data, only the NLO files that were created since the last backup cycle need to be processed. Those files typically represent a small amount of data compared to the entire set of NLO files.

In the incremental backup case, duplicate NLOs are not backed up again. Therefore, the space savings from DAOS are directly proportional to the number of duplicate NLOs seen in the environment, and the backup time savings is the product of the space that is saved and the backup throughput.

The ProtectTIER server greatly benefits from this whole behavior. We have seen a factoring ratio of 3 - 5 times higher than before the DAOS is enabled.

20.1.4 Applying the DAOS solution

to assess the benefit of using DAOS, run the DAOS Estimator tool by issuing the **DAOSest** command in the Domino server console. Perform this action outside of your office hours, because the estimation procedure impacts the server performance, especially mail servers with hundreds to thousands of users.

The histogram result of the estimation job is shown in Example 20-1 on page 301:

- ▶ The top line shows the number of files.
- ▶ The bottom line show the specific size of the detachment threshold.
- ▶ The middle line shows what the result represents in percentage of all attachments.

Example 20-1 Relative and percentage numbers of attachments of different size

=====											
66362	19744	33629	35311	161416	90946	18550	3458	426	22	0	
=====											
0.0%	0.1%	0.3%	0.6%	7.5%	20.4%	31.5%	25.1%	11.4%	3.2%	0.0%	
=====											
4k	8k	16k	32k	64k	1MB	5MB	20MB	100MB	1GB	>1GB	
=====											

The summary of the estimation process is also shown in a different form in Example 20-2 on page 303.

Before anything is done with DAOS, there are some prerequisites that must be addressed. These prerequisites might not all apply in your situation, but an important task is to verify them to ensure that any changes that are needed can be accommodated. These items are all requirements for enabling DAOS, and are not optional.

Consultation: Before you implement DAOS, consult with your Domino representative.

The prerequisites are as follows:

- ▶ Disable SCOS Shared mail

The Single Copy Object Store (SCOS) is an older approach to attachment consolidation. This feature is not compatible with DAOS and must be disabled before you enable DAOS.

- ▶ Disable NSFDB2

The NSFDB2 is a feature that you can use to store NSF data in DB2 running either on the same or a different server. This feature is also not compatible with DAOS and must be disabled on every NSF application that participates in DAOS.

- ▶ Upgrade Domino server

Although DAOS was introduced in Domino V8.5.0, many important stability and performance improvements were made in subsequent releases. Therefore, all new DAOS deployments should use Domino 8.5.3 or later.

- ▶ Enable transaction logging

The DAOS depends on transaction logging for correct operation. Because DAOS must update several locations simultaneously, it is important that all those updates succeed or fail (and are later rolled back) as a unit.

- ▶ Adjust backup/restore processes

You must have reliable backup and restore procedures in a production environment to avoid the possibility of data loss. DAOS adds some complexity to the backup and restore process, so you must have a well-established backup and restore foundation for DAOS. Transaction logging introduces additional features that provide even better recovery options.

- ▶ Upgrade Names.nsf design

The design of the Names.nsf file was changed to accommodate DAOS, and the Server document has a tab that covers the DAOS settings. Names.nsf must use the new pubnames.ntf template on all Domino servers that are enabled for DAOS.

20.1.5 ProtecTIER considerations

In contrast to the general suggestions for DAOS deployment on Domino servers, this section summarizes preferred practices or limitations that apply when the ProtecTIER deduplication solution for backup and recovery is in place:

- Disable NLO compression.

The mail attachments that are represented by NLO files use a certain amount of disk space on the Domino server. Domino administrators tend to enable one of the available compression techniques on the attachments in NSF files. If no attachment compression is enabled on the NSF files, or if Huffman compression is being used, then enabling LZ1 compression can save a significant amount of disk space. Run **compact -ZU** to enable LZ1 compression.

Tip: To achieve the best factoring ratio, avoid using the **-ZU** flag during compaction.

- Disable design and data document compression.

Another Domino space-saving feature is design and data document compression. Enabling these compression forms can also save disk space, but they have a negative effect on deduplication results in the ProtecTIER server. The savings from these features are independent from DAOS and do not achieve the level of savings that you can make with the ProtecTIER solution.

Tip: Do not compress Design and Data documents.

- Consider the compacting frequency.

Compacting less frequently is not popular with Domino administrators. However, it does *not* have a significant effect on the performance of a Domino server, mailboxes, or the storage capacity that is used by the Domino server. Complicating factors are backups with retention periods and the database file unique identifier, also called *database instance identifier* (DBIID).

When the DBIID is changed by running a compact job (which is the default action, unless the **-b** parameter is not specified), a Domino server always considers this database as eligible for full backup, regardless whether the next backup job is scheduled as an incremental only.

With a backup schedule that consists of weekly full backups (during the weekend) and daily incremental backup of databases with a changed DBIID (during weekdays), you should perform compact jobs on a weekly basis before the full backup occurs. This setup has a positive effect on the ProtecTIER factoring ratio.

Tip: Schedule compact jobs less frequently and ensure that they always complete before the next full (or selective) backup of NSF databases. Incremental backup does not back up Domino NSF files, unless the DBIID has changed.

- Compact only selected databases.

Not all Domino databases must be compacted regularly. If the percentage of *white space* (unused space) in the database is, for example, less than 10% of the mailbox size, consider excluding this database from the compaction. The space savings of such a compact job is negligible, but your factoring ratio decreases. Use the **-S 10** option to direct the compact task to databases only with 10% or more of the white space. The database DBIID still changes, unless the **-b** option is not used.

Tip: Do not compact databases that use storage space efficiently.

- Disable encryption of attachments.

When access to server resources is restricted to responsible personnel only and there is minimal or no risk of data exposure, Domino administrators should disable encryption on detached email attachments (NLO files). The enabled encryption has a negative effect on the ProtectTIER factoring ratio, because each block of data that is sent to the ProtectTIER server behaves as a unique block of data. Although the encryption is enabled by default, you can disable it by adding the following parameter to the `Notes.ini` file:

```
DAOS_ENCRYPT_NLO=0
```

The setting cannot be changed retroactively, and the only way to remove encryption from an existing DAOS installation is to completely disable DAOS.

Encryption: Encryption has a negative effect on factoring ratios. Disable it if possible.

- Define the appropriate thresholds of the DAOS process.

If the attachment is larger than the minimum participation size, it is stored in the DAOS repository. If it is smaller, it is still stored in the NSF, as it would be without the DAOS feature enabled.

Choosing a size that is too large results in too few attachments being stored in DAOS (low yield), which reduces the savings that DAOS can offer and the ProtectTIER product can benefit from. Conversely, choosing too small of a size can result in a high yield, resulting in an unmanageable number of files in the DAOS repository. The statistics in Example 20-2 show the DAOS minimum size versus the number of NLOs and disk space that is required.

Example 20-2 DAOS minimum size versus the number of NLOs and disk space

0.0 KB will result in	429864 .nlo files using	180.7 GB
4.0 KB will result in	363502 .nlo files using	136.6 GB
8.0 KB will result in	343758 .nlo files using	130.5 GB
16.0 KB will result in	310129 .nlo files using	128.2 GB
32.0 KB will result in	274818 .nlo files using	119.4 GB
64.0 KB will result in	113402 .nlo files using	110.4 GB
1.0 MB will result in	22456 .nlo files using	85.8 GB
5.0 MB will result in	3906 .nlo files using	47.9 GB
20.0 MB will result in	448 .nlo files using	17.6 GB
100.0 MB will result in	22 .nlo files using	3.8 GB

Look for a value that yields about 80 - 90% of the theoretical maximum of the DAOS repository size. Although that value might sound low, it is generally the best trade-off between the DAOS benefits and the resulting number of files.

Hint: Determine the appropriate DAOS size when the attachment is offloaded from the NSF to the NLO file.

20.1.6 Preparing Domino databases for DAOS

The task of preparing Domino databases for DAOS should be performed only once. To accomplish this task, complete the following steps:

1. Depending on your operating system, choose the most appropriate procedure:
 - If the Domino server is running Windows, click **Start** → **Programs** → **Lotus Applications** → **Lotus Domino Server**.
 - If the Domino server is running UNIX, enter the following command in the command-line interface (CLI):

```
/opt/lotus/bin/server
```
2. Double-click **n1notes.exe**, and go to the workspace window.
3. Browse for the **names.nsf** file. The location is usually **E:\Lotus** → **Domino** → **Data**.
4. Click **Configuration** → **Servers** → **All Server Documents**. Then, open the document that is related to your Domino server.
5. Double-click the page to change it to edit mode. Select the **Transactional Logging** tab.
6. Set the following parameters:
 - Log path: **logdir**
 - Logging style: **Circular**
 - Maximum log space: **512 MB**
7. Save your parameters and close the window.
8. Shut down the Domino server by entering the following command at the CLI:

```
exit <password>
```
9. Add the following line to the **notes.ini** file:

```
CREATE_R85_DATABASE=1
```
10. Start the Domino server again and use the password.

Starting time: For the initial startup sequence after you make these changes, it might take several minutes for the start sequence to run.

11. Complete steps 5 - 7 to edit the server document again. Open the **DAOS** tab. You might need to scroll to the right to see the tab.
12. Update the following parameters:
 - Store Attachments in DAOS: **ENABLED**
 - Minimum size: **4096**
 - DAOS base path: **daos**
 - Defer deletion: **30 days**
13. Restart the Domino Server by entering the following command at the CLI:

```
restart server [password]
```

In the next compaction, you see the DAOS directory that is created in the Domino data directory. That directory contains the following entry:

```
0001/<really_long_name>.n1o
```

20.2 Microsoft Exchange

This section describes the suggested settings for the Microsoft Exchange (Exchange) environment to improve the backup throughput and the factoring ratio of the ProtecTIER server. The examples that are used in this section are based on IBM Tivoli Storage Manager, but the suggested settings apply to most enterprise backup applications. Some of the settings might not be available in other backup applications. Contact the backup application provider for additional information.

20.2.1 Defragmentation

Defragmentation is commonly used in the Microsoft Exchange environment to recover the disk efficiency of fragmented disks. The defragmentation process rearranges the data that is stored on the disk and creates continuous storage space. There are two types of defragmentation processes: online defragmentation and offline defragmentation.

Online defragmentation

The online defragmentation process removes objects that are no longer being used while Exchange databases remain online. Before Microsoft Exchange 2010, the online defragmentation ran as part of daily Mailbox database maintenance, although this Mailbox database maintenance can be scheduled to run at different times.

As of Exchange 2010, online defragmentation is separated from the Mailbox database maintenance process and it runs continuously in the background. Additional details about database defragmentation are available in at the following web page:

<http://technet.microsoft.com/en-us/library/bb125040.aspx#NewESE>

Offline defragmentation

Offline defragmentation is a manual process that creates a database file and copies database records without the white space from the original database file to the newly created database file. When the defragmentation process is complete, the original database is removed and the new database file is renamed as the original.

Offline defragmentation is not part of regular Mailbox database maintenance. It can be done only when the Mailbox database is in an offline state, and this action requires much storage space, because both the original database file and newly created database file must coexist on the disk during the defragmentation process.

20.2.2 Suggestions for Microsoft Exchange

The following list includes the suggested processes and settings of the backup applications to optimize the performance and factoring ratio of the ProtecTIER server:

- ▶ Perform a daily full backup rather than daily incremental backups where only transaction logs are backed up.
- ▶ Create one backup job for each database (or for each storage group) without multistreaming to keep similar data blocks in the same stream.

- ▶ Create concurrent backup jobs if there is more than one database (or more than one storage group) in the Exchange servers to improve overall backup throughput. Example 20-3 shows how to create different backup jobs for different databases in Tivoli Storage Manager. Remember to increase the number of mount points for the client node if multiple databases are housed in one Exchange server.

Example 20-3 Create multiple backup jobs for different databases with Tivoli Storage Manager

```

TDPEXCC BACKup <Storage Group 01/ Mailbox Database 01> full
TDPEXCC BACKup <Storage Group 02/ Mailbox Database 02> full
TDPEXCC BACKup <Storage Group 03/ Mailbox Database 03> full

```

- ▶ Disable compression and encryption in Exchange databases and backup applications.
- ▶ If personal archive files (.pst) are backed up, do not enable compression and encryption whenever possible.
- ▶ Configure a longer interval for an online defragmentation process, for example, reschedule the daily database maintenance to be on a weekly basis.
- ▶ Consider a LAN-free backup that enables data to be sent directly to storage devices.

20.2.3 Microsoft Exchange 2010

Because the online defragmentation is moved out from the daily maintenance process and it runs continuously in background, there is no option to disable or schedule online defragmentation. This situation impacts the factoring ratio. However, we do expect deduplication for Exchange 2010 because Single Instance Storage (SIS) is no longer supported in Exchange 2010. You can find more information about SIS in the Microsoft Exchange team blog at the following web page:

<http://msexchange.com/archive/2010/02/22/454051.aspx>

20.3 Microsoft SQL Server

This section describes the suggested settings for the Microsoft SQL environment to improve the backup throughput and factoring ratio of the ProtecTIER server. The examples that are used in this section are based on IBM Tivoli Storage Manager, but the suggested settings apply to most enterprise backup applications. Some of the settings might not be available in other backup applications. For more information, contact the backup application provider.

20.3.1 Integrating the ProtecTIER server with Microsoft SQL Server backup

A ProtecTIER server can be integrated with traditional backup applications, or with the native SQL server backup utility to back up the Microsoft SQL server.

To back up a Microsoft SQL server with traditional backup applications, such as Tivoli Storage Manager (as of version 7.1.3, rebranded to IBM Spectrum Protect), the ProtecTIER product can be deployed as a Virtual Tape Library (VTL) and the File System Interface (FSI) to work in conjunction with the backup applications.

To back up the Microsoft SQL Server with the native SQL server backup utility, the ProtecTIER product can be used as CIFS shares through FSI deployment. This section describes the suggested ways to integrate the ProtecTIER server with different backup methods.

Using native SQL Server backup

You can back up the Microsoft SQL Server with the native SQL server backup and restore utility, where data files can be backed up directly to backup media without using third-party backup applications. The backup media can be disk or tape devices. Most administrators choose to back up to disk rather than to tape devices because the native SQL server backup does not have a tape media management capability similar to other backup applications.

The ProtecTIER product can be deployed as an FSI that provides CIFS shares to a Microsoft SQL server, and the native SQL server backup can use the CIFS share as the destination (Figure 20-3).

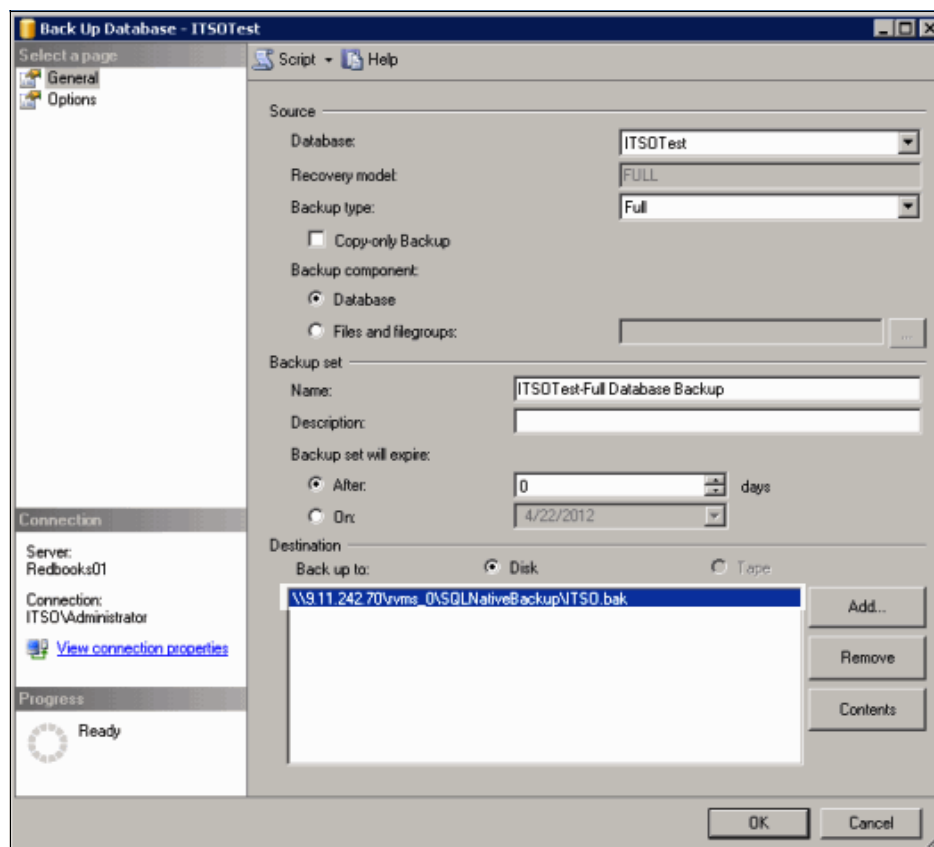


Figure 20-3 Use the ProtecTIER CIFS share as the destination of the SQL native backup

For more information about ProtecTIER FSI, see Chapter 5, “ProtecTIER File System Interface: General introduction” on page 65.

Using third-party backup applications

Most backup applications support Microsoft SQL backup through an SQL backup agent, for example, IBM Tivoli Storage Manager Data Protection for Microsoft SQL Server. Backup applications use the ProtecTIER server as tape devices, file devices, or disk storage as backup destinations during the backup server configuration. For details about how to integrate the ProtecTIER server with different types of back-end storage, see Part 2, “Back-end storage subsystems” on page 97.

20.3.2 Index defragmentation

The Microsoft SQL Server maintains indexes to track table updates, and these indexes can become fragmented over time. Heavily fragmented indexes might affect the database query performance, so Microsoft SQL Server uses a defragmentation feature to reorder the index row in continuous pages.

The defragmentation process rebuilds the indexes by compacting the index pages and reorganizing the index rows. This process results in the indexes being seen as new data blocks in backup streams, which can affect the deduplication process that identifies unique data at block level.

Index defragmentation can adversely affect database workload performance. You should perform index defragmentation only when it is necessary. Before you defragment, see the following web page:

<http://technet.microsoft.com/en-us/library/cc966523.aspx>

20.3.3 Suggestions for Microsoft SQL Server

The following list includes some suggestions for the Microsoft SQL Server to improve the backup throughput and deduplication ratio of the ProtecTIER server:

- ▶ Perform full backups whenever possible.
- ▶ When you use the ProtecTIER server as CIFS shares, always use the Universal Network Convention (UNC) path rather than the Windows mapped drive to ensure that the correct CIFS shares are used, and to avoid the Windows connection timeout issue.
- ▶ Do not schedule index defragmentation on a regular basis. Perform index defragmentation only when necessary.
- ▶ Disable compression and encryption in the Microsoft SQL server and backup applications.
- ▶ Limit the number of backup streams to one stream for one database backup, or one stream per physical volume, if one single large database is split into multiple physical volumes. Setting the stream to 1 gives the best factoring ratio, but it might affect overall backup performance. Set the number of the stream to the minimal number that does not inhibit the performance by using the following option:

STRIPes=1

- ▶ Use a larger buffer size for a better deduplication ratio. Increase the buffer size slowly from the default buffer size of backup application, but do not exceed the amount of buffer that can be handled by the system memory.
- ▶ Limit the number of input/output (I/O) buffers in a backup stream. Ideally, there should be two buffers per stream, with one buffer for reading data from an SQL Server and the other for sending data to the backup applications, as shown by the following settings:
 - BUFFER=2
 - BUFFERSize=1024
 - SQLBUFFER=0
 - SQLBUFFERSize=1024

20.3.4 LiteSpeed for SQL Server

LiteSpeed for SQL Server is a backup utility that compresses and encrypts the SQL database before the data is stored in backup devices. The factoring ratio is greatly impacted if the data is compressed and encrypted before it reaches the ProtecTIER repository. The ProtecTIER product offers little deduplication benefit if LiteSpeed is used for SQL server backup.

20.4 DB2

This section describes the settings and parameters that should be modified in DB2 environments to enable the maximum performance and optimum factoring for the ProtecTIER server. It also explains why combining DB2 compression with ProtecTIER deduplication is possible.

Updating DB2: Update your DB2 to Version 9.7 Fix Pack 4 or later and use the **DEDUP_DEVICE** option for backing up your database. This action results in the best deduplication ratio. DB2 compression types can work with deduplication.

20.4.1 Combining DB2 compression and ProtecTIER deduplication

DB2 offers multiple options to use compression in conjunction with database rows, database values, or both. Run the **select tabname,compression from SYSCAT.TABLES** command to verify the settings for your database.

Table 20-1 list the available compression types.

Table 20-1 DB2 compression types

SYSCAT.TABLES values	Compression type active
R	Row compression is activated if licensed. A row format that supports compression can be used.
V	Value compression is activated. A row format that supports compression is used.
B	Both value and row compression are activated.
N	No compression is activated. A row format that does not support compression is used.

With DB2 compression, data in the database is compressed on a table-row basis. These compressed rows are written to disk as DB2 pages with a default size of 4 K. Changes in a DB2 database with compression enabled affect only the data in these specific DB2 pages; the changes do not affect the entire database because of block based compression, which is different from traditional compression approaches. Also, after changes occur in the database, only the changed pages are recompressed.

Effectively, compressed DB2 pages are not apparent to HyperFactor, and a large sequence of compressed pages factors well if they are not changed. Effectively, there is no general penalty for using compression in DB2. The data change rate affects this deduplication ratio; regardless of whether you use compression, the behavior is the same.

Remember, even if DB2 compression does have a friendly synergy with ProtecTIER deduplication, the full deduplication potential can be reached only with all sorts of data reduction technology, such as disabled compression.

Important: Using another form of compression with DB2 database backups, for example, Tivoli Storage Manager (version 7.1.3 was rebranded to IBM Spectrum Protect) compression or the compression feature of another backup software, still impacts your achievable deduplication ratio.

20.4.2 Upgrading the DB2 database to improve deduplication

The most suggested method for improving deduplication is to upgrade the DB2 database to DB2 9.7 Fix Pack 4 or later to be able to use the **DEDUP_DEVICE** option. This special feature improves the DB2 backup process to make it *deduplication friendly*. Update to DB2 9.7 Fix Pack 4 or later. Only with this version or any later version can you experience the full benefit of the optimized DB2 data handling for deduplication devices.

You can download DB2 Fix Packs for DB2, for Linux, UNIX, and Windows, and IBM DB2 Connect™ products from the following web page:

<http://www.ibm.com/support/docview.wss?uid=swg27007053>

To fully understand the improvements of the **DEDUP_DEVICE** option, look at the default DB2 database backup behavior. When a DB2 backup operation begins, one or more buffer manipulator (db2bm) threads are started. These threads are responsible for accessing data in the database and streaming it to one or more backup buffers. Likewise, one or more media controller (db2med) threads are started and these threads are responsible for writing data in the backup buffers to files on the target backup device.

The number of db2bm threads that is used is controlled by the **PARALLELISM** option of the **BACKUP DATABASE** command. The number of db2med threads that is used is controlled by the **OPEN n SESSIONS** option. Finally, a DB2 agent (db2agent) thread is assigned the responsibility of directing communication between the buffer manipulator threads and the media controller threads.

This process is shown in Figure 20-4.

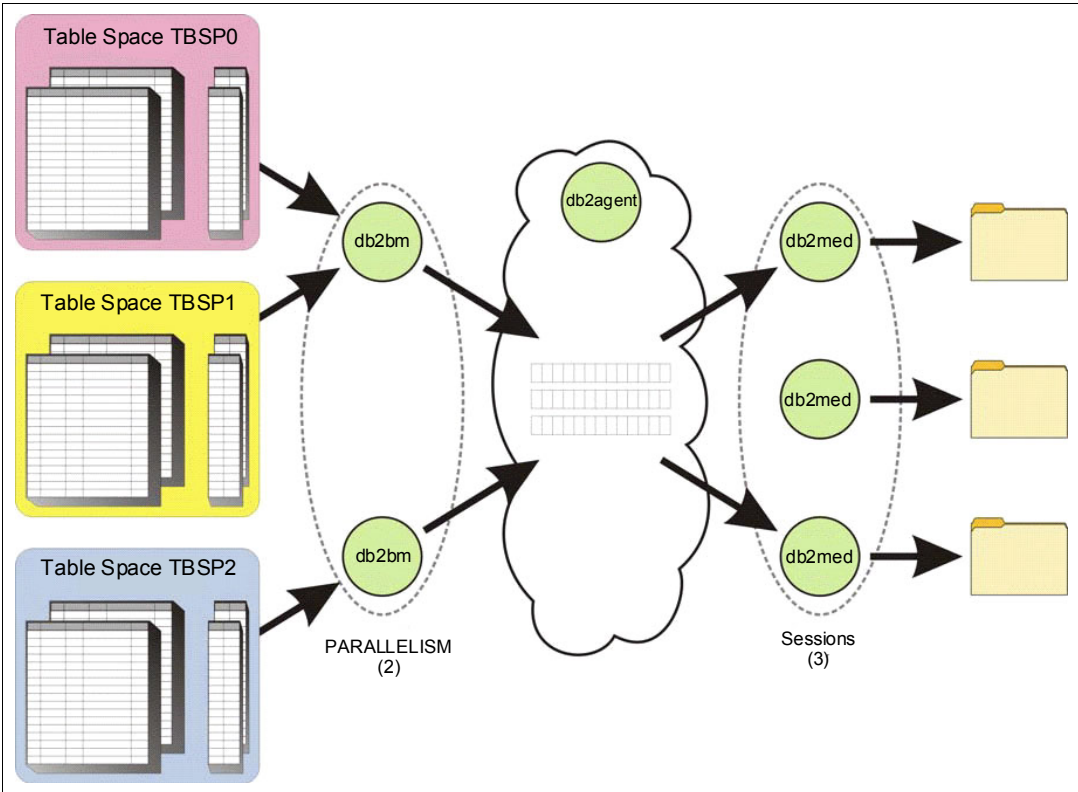


Figure 20-4 DB2 backup process model

Without the **DEDUP_DEVICE** option, data that is retrieved by buffer manipulator (db2bm) threads is read and multiplexed across all of the output streams that are being used by the media controller (db2med) thread. There is no deterministic pattern to the way in which data is placed in the output streams that are used (Figure 20-5). As a result, when the output streams are directed to a deduplication device, the device thrashes in an attempt to identify chunks of data that are already backed up.

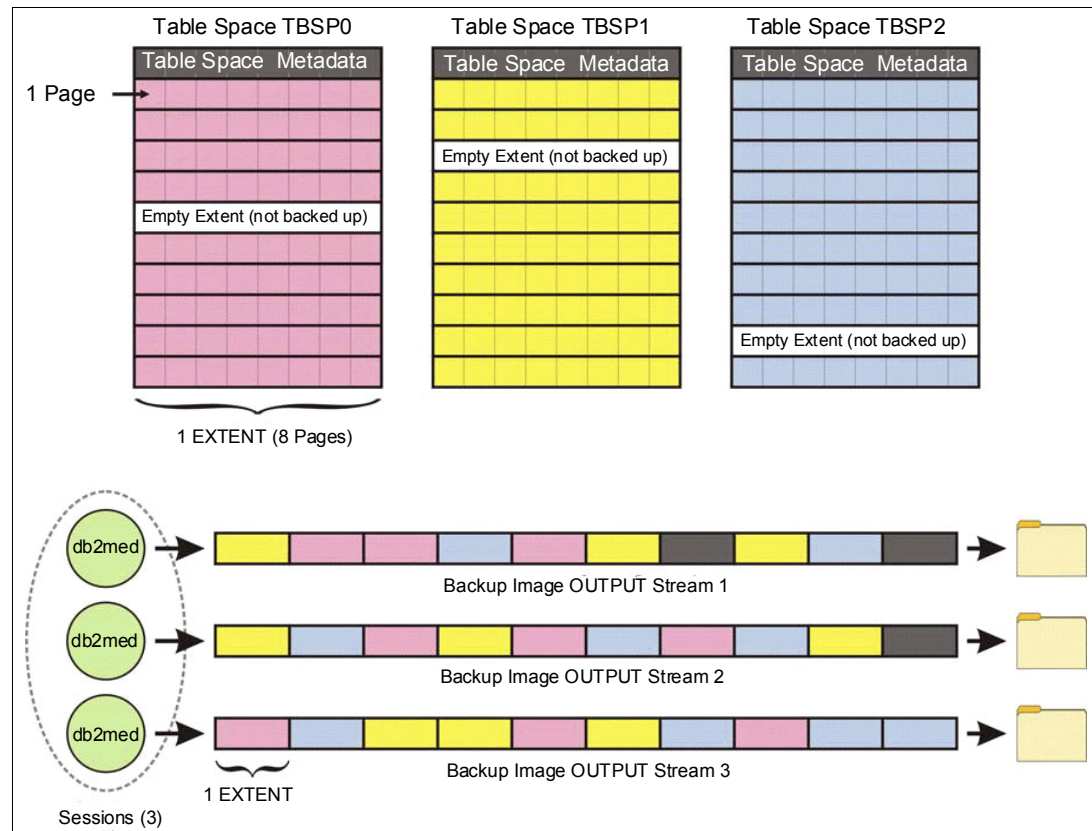


Figure 20-5 Default database backup behavior

20.4.3 DB2 DEDUP_DEVICE setting

When the **DEDUP_DEVICE** option is used with the **BACKUP DATABASE** command, data that is retrieved by buffer manipulator (db2bm) threads is no longer read and multiplexed across the output streams that are being used by the media controller (db2med) threads. Instead, as data is read from a particular table space, all of that table space's data is sent to only one output stream. Furthermore, data for a particular table space is always written in order, from lowest to highest page. As a result, a predictable and deterministic pattern of the data emerges in each output stream, making it easy for a deduplication device to identify chunks of data that are already backed up.

Figure 20-6 illustrates this change in backup behavior when the **DEDUP_DEVICE** option of the **BACKUP DATABASE** command is used.

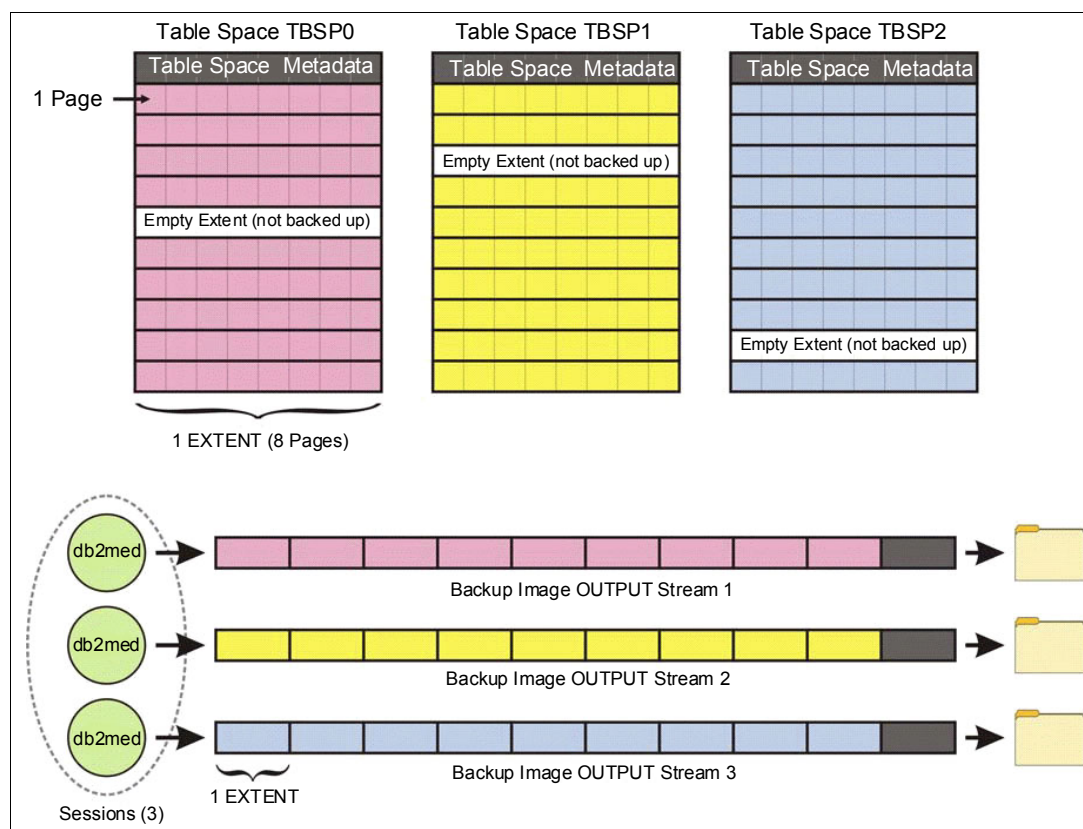


Figure 20-6 Database backup behavior with the **DEDUP_DEVICE** option

When you use the **DEDUP_DEVICE** option, each table space is backed up to a dedicated tape drive. Using a number of virtual tape drives that is equal to or greater than the number of table spaces you want to back up is suggested.

If the database contains table spaces larger than others (above 30% of the entire database size), it prolongs the backup. If this situation affects the backup window, consult your DB2 support to assist you in splitting the larger table spaces and making them smaller. Also, communicate this information to the DB2 database planning staff so that future deployments can directly benefit of the improved deduplication without any drawback.

20.4.4 Example of **DEDUP_DEVICE** setting

Example 20-4 uses 16 tape drives to back up your 16 table spaces that ideally are of equal size using Tivoli Storage Manager, in parallel, using the **DEDUP_DEVICE** option.

Example 20-4 Multistreamed backup of DB2

```
db2 backup db <database_name> use tsm open 16 sessions dedup_device exclude logs
```

This action results in the best possible deduplication ratio. With DB2 9.7 Fix Pack 4 and later, the DB2 self-tuning capability is able to support this backup command by choosing all tuning values automatically.

20.4.5 Excluding logs from the DB2 database backup

Use the **exclude logs** parameter to avoid backing up your database logs to the same destination as your database. Database logs tend to have a 100% change rate and therefore have a negative effect on your overall HyperFactor ratio. Instead, redirect the archive logs directly to storage with no further active data reduction technology. Using the **include logs** parameter with the DB2 backup command results in archive logs being automatically added to the backup images. This action causes different patterns in the backup streams and reduces deduplication efficiency.

20.4.6 DB2 suggested settings without DEDUP_DEVICE

Backing up to a deduplication device when the **DEDUP_DEVICE** option is not available can still be optimized by applying some rules. The DB2 settings in Table 20-2 provide the best deduplication efficiency for backing up without the **DEDUP_DEVICE** option.

Table 20-2 Suggested DB2 settings

DB2 parameter	Suggested value	Description
sessions/ OPEN n SESSIONS	Minimum ^a	Change the value to read the data at the required backup rate.
buffers/ WITH num-buff BUFFERS	Parallelism + sessions + 2	The numbers of buffers should be #sessions + #parallelism +2. Also, the following calculation must fit: (num-buffers * buffer-size) < UTIL_HEAP_SZ (UTIL_HEAP_SZ is the database utility heap size.).
buffer/ BUFFER buff-size	16384	This value requires much memory. If this value is too much for your environment, use the largest possible BUFFER value instead. The value of this parameter is specified in multiples of 4 KB pages.
parallelism/ PARALLELISM	Minimum ^a	Change the value to read the data at the required backup rate.

a. Select the minimum value to configure an acceptable backup window time frame. Although a value of 1 is the best for deduplication, it might increase backup times in large multi-table space databases.

Setting: The large **BUFFER** size of 16384 is the setting with the most effect on your HyperFactor deduplication. The bigger the **BUFFER** value is, the better your deduplication ratio is.

20.4.7 Example of DB2 command using sessions, buffers, and parallelism

Example 20-5 shows an example of a DB2 backup command using four sessions, eight buffers, a buffersize of 16384, and a parallelism of 2.

Example 20-5 Database backup command

```
db2 backup db <databasename> use tsm open 4 sessions with 8 buffers buffer 16384  
parallelism 2
```

Tip: Always use the same parameters for restore as you did for backup (number of sessions, buffers, buffer size, and parallelism) to ensure maximum restore performance.

20.5 Oracle

Oracle Recovery Manager (RMAN) is a backup and recovery utility for Oracle databases. The RMAN backs up Oracle databases directly to the disk or to other storage devices using third-party backup applications. Backup applications interface with RMAN to back up Oracle databases with various storage devices, such as tape, or file system.

The ProtecTIER server can be deployed as a VTL or FSI. For more details about how to set up VTL, and FSI, see Chapter 4, “Virtual Tape Library guidelines” on page 49 and Chapter 5, “ProtecTIER File System Interface: General introduction” on page 65.

This section describes the optimal settings and guidelines of RMAN to improve the backup throughput and factoring ratio of the ProtecTIER solution. The focus is on preferred practice parameters in Oracle RMAN environment when using ProtecTIER network-attached storage (NAS) as a disk target for RMAN backups.

20.5.1 Suggested RMAN settings

The following list describes several suggested settings for RMAN:

- ▶ Backup routine:
 - Perform daily backups whenever possible. Performing a full backup enables the simplest and fastest restoration.
 - If full backup is not feasible, consider using incremental backups (level 0 - level 1 backups).
 - Backup all database components including data files, control files, redo logs, spfile, and archive logs (where applicable).
- ▶ Archive log mode:
 - For online backups of Oracle databases, the database must have archive log mode enabled.
 - Make sure that multiple redo log groups are defined, and that their members' size is sufficient to complete.
 - Archive locations must be saved outside of ProtecTIER NAS.
 - Keep all archive logs generated between backups.
 - Enable **ARCHIVELOG** mode for your database. Run **ARCHIVELOG** as often as possible. Make sure that the RMAN backup scripts use separate commands when backing up data (Tablespaces, data files, control files) and a separate command for backing up archive logs. Separating the data files and archive logs to two separate backup streams will provide better deduplication because offline archive logs are unique and cannot be deduplicated. Run archive log backups as often as possible for better point-in-time recovery capabilities.
- ▶ Disable compression and encryption in Oracle databases and backup applications.
- ▶ Disable or minimize multiplexing. Multiplexing enables RMAN to combine data blocks from different files into a single backup set, which impacts the factoring ratio. RMAN multiplexing is affected by the following two parameters:
 - The **FILESERSET** parameter determines how many files should be included in each backup set. Set **FILESERSET=1** to send only one file per backup set in each channel (backup stream).

- The **MAXOPENFILES** parameter defines how many files RMAN can read from the Oracle source simultaneously. Set **MAXOPENFILES=1** so that RMAN does not read from more than one file at a time.

Example 20-6 shows a multiplexing calculation with various **FILESERSET** and **MAXOPENFILES** settings.

Example 20-6 Calculation of multiplexing in RMAN

Scenario 1: **FILESERSET=6**, **MAXOPENFILES=3**, number of data files=4
 Multiplex = 3 (Limiting by the **MAXOPENFILES** setting)

Scenario 2: **FILESERSET=2**, **MAXOPENFILES=3**, number of data files=4
 Multiplex = 2 (Limiting by the **FILESERSET** setting)

Scenario 3: **FILESERSET=8**, **MAXOPENFILES=4**, number of data files=2
 Multiplex = 2 (Limiting by the number of data files)

Scenario 4: **FILESERSET=1**, **MAXOPENFILES=1**, number of data files=4
 Multiplex = 1 (Limiting by the **FILESERSET** and **MAXOPENFILES** settings)

- Increase the number of parallel backup streams to improve backup throughput. Ensure that the number of ProtecTIER virtual tape drives that are available for Oracle backup matches the number of parallel streams that are configured in RMAN. For example, this value is enabled in the definition of the Tivoli Storage Manager client on the Tivoli Storage Manager server by using the **MAXNUMP=32** parameter. Set **PARALLELISM=32** (up to 64).

Figure 20-7 depicts a case study that shows the factoring ratio and a backup throughput result with different multiplexing and parallel channel settings. The result is taken from a case study of Oracle database backup with Tivoli Storage Manager, and a 30-day retention period on ProtecTIER virtual tape. A full backup that is performed on alternate days averages a 5% data change rate between the full backups.




Figure 20-7 Example of multiplexing and parallelism effect on the HyperFactor ratio

Environmental variations: The parameters that are used in this test are not absolute requirements. Different environments might produce different results, depending on the data change rate and backup practices. Fine-tune the RMAN settings in your environment gradually to get the settings that do not inhibit performance.

Be aware of the following usage guidelines:

- ▶ Using Direct NFS (dNFS, Oracle 11gR2 introduced dNFS). An optimized NFS client that provides faster and more scalable access to NFS storage located on NAS storage devices. dNFS was tested and evaluated but is currently not supported. Use standard client kernel NFS instead.
- ▶ Using Automatic Storage Manager (ASM). The use of ASM was tested using ASMLib with ProtecTIER and found to be fully functional. The use of ASM for storing the database provided better performance for backups using less concurrent channels.
- ▶ Using hidden underscore parameters. The use of Oracle hidden, undocumented, and unsupported parameters is strongly *not* recommended when using ProtecTIER. Some performance white papers suggest changing `_backup_disk_bufcnt` and `_backup_disk_bufsz` hidden parameters. Tests were performed for those setups and it is *not* recommended to change those default hidden parameters when using ProtecTIER.

20.5.2 Mounting NFS Oracle Server to ProtecTIER NAS

For RMAN backups to work with ProtecTIER NAS, the following mount parameters are suggested to be used from the Oracle Servers. Example 20-7 has ProtecTIER Server (proctier) exporting share /RMAN and Oracle server using mount point /mnt/RMAN for destination target for RMAN backups.

Example 20-7 Suggested parameters from Oracle Servers for RMAN backups with ProtecTIER NAS

```
mount -o rw,soft,intr,nolock,timeo=3000,nfsvers=3,proto=tcp protectier:/RMAN  
/mnt/RMAN
```

Attention: A suggestion is to use the soft mount attributes because using the hard mount option continues to retry I/O forever, which might lead to Oracle database crashes. If the ProtecTIER server goes down and the NFS client is using the hard mount, eventually the retry forever will succeed, or it could result in negative consequences on the Oracle DB server side.

The suggested timeout value to use with soft is five minutes. It is much quicker than ProtecTIER rebooting, then the Oracle NFS client will detect it and stop trying, only leading to a backup failure. When the Oracle DBA re-initiates the RMAN backup jobs for the failed data files, the data will be rewritten with no data loss.

Oracle requires the NFS mount point to be set to hard mount for NAS storage system. Oracle assumes that the mount point will be used to contain tablespaces and archive logs. Therefore, any down time with NAS storage will prevent Oracle from restarting or mounting the database. When an NFS mount point is being used by an Oracle instance, the instance checks for specific mount options.

If any of the option is *incorrect*, this Oracle instance will issue the following message and the operation will fail:

ORA-27054: NFS file system where the file is created or resides is not mounted with correct options

This behavior is described in document 359515.1 at the My Oracle Support website:

<https://support.oracle.com/>

Because the suggested soft mount option for ProtecTIER NAS and Oracle Server requires hard mount, the messages from Oracle that report incorrect mount options (shown in Example 20-8) are listed in the alert.log file.

Example 20-8 Messages indicate incorrect mount option during RMAN backups to ProtecTIER NAS

```
WARNING:NFS file system /mnt/ORC2_RMAN1 mounted with incorrect options (rw,vers=3,
rsize=1048576, soft, intr,
nolock,ptoto=tcp,timeo=3000,retrans=2,sec=sysaddr=10.1.2.135)
WARNING:Expected NFS mount options: rsize>=32768,wsiz>=32768,hard,
Wed Jul 10 03:10:31 2013
WARNING:NFS file system /mnt/ORC2_RMAN1 mounted with incorrect options (rw,vers=3,
rsize=1048576, soft, intr,
nolock,ptoto=tcp,timeo=3000,retrans=2,sec=sysaddr=10.1.2.135)
WARNING:Expected NFS mount options: rsize>=32768,wsiz>=32768,hard,
Wed Jul 10 03:10:31 2013
WARNING:NFS file system /mnt/ORC2_RMAN1 mounted with incorrect options (rw,vers=3,
rsize=1048576, soft, intr,
nolock,ptoto=tcp,timeo=3000,retrans=2,sec=sysaddr=10.1.2.135)
WARNING:Expected NFS mount options: rsize>=32768,wsiz>=32768,hard,
Wed Jul 10 03:10:31 2013
```

To work around the Oracle NFS checks, with error message ORA-27054, use the following commands:

```
-- Set event until database shutdown
alter system set events '10298 trace name context forever,level 32';

-- Set event after database restart
alter system set event="10298 trace name context forever, level 32" scope =
spfile;
```

These commands will prevent any NFS checks of the database which will enable the customer to use the suggested mount options, especially using the soft mount options during RMAN backups to ProtecTIER NAS.

To re-enable Oracle's NFS checks, use the following SQL command:

```
-- Set event until database shutdown
alter system set events '10298 trace name context off';
-- Set event after database restart
alter system set event="10298 trace name context off" scope = spfile;
```

Note: This setting will apply to all mount points used by the instance, and not for only those that are used by RMAN. If you are using NFS mount point for data file or archive log containers, you should use the suggested mount point settings mentioned in My Oracle Support node #359515.1 for those mount points.

20.5.3 Using ProtecTIER NAS to run RMAN incremental merge backups

Using Oracle RMAN incremental backups, you can decrease the data that needs to be backed up, including only datafile blocks that have changed since a specified previous backup. The goal is to back up only those data blocks that have changed since a previous backup.

This kind of backup is permitted on the ProtecTIER, but has no significant advantages in terms of deduplication. The backup space will be lower, but the time to recover from those backups will be longer. The fastest backup in these terms is to use image copy backups.

Because image copy backup might take much time and storage space, Oracle suggests using the *Incrementally Updated Backups* feature, which enables you to avoid the overhead of taking full image copy backups of datafiles, while providing the same recovery advantages as image copy backups.

At the beginning of a backup strategy, RMAN creates an image copy backup of the datafile. Then, at regular intervals, such as daily or weekly, level 1 incremental backups are taken and applied to the image copy backup, rolling it forward to the point in time when the level 1 incremental was created.

During restore and recovery of the database, RMAN can restore from this incrementally updated copy and then apply changes from the redo log, with the same results as restoring the database from a full backup taken at the System Change Number (SCN) of the most recently applied incremental level 1 backup.

A backup strategy based on incrementally updated backups can help minimize time required for media recovery of your database. For example, if you run scripts to implement this strategy daily, at recovery time you never have more than one day of redo to apply.

The down side of this backup is that only one valid full backup exists at all times, making a recovery to a point prior to the latest backup impossible.

Using the ProtecTIER Cloning feature, you can now duplicate the backed up directory to a new directory without any physical effect (100% deduplication) and applying the incremental changes to one of the copies. This method will ensure that any unchanged data in the datafiles will be deduplicated, and all new or changed information will be available. When using this method, multiple copies of a backup can be saved at the same time, and the restore of the database to any version will be available at all times.

Example 20-9 and Example 20-10 on page 320 show an instance of an implementation of RMAN incremental merge backup with ProtecTIER clone.

Example 20-9 Master script for running the backup and clone

```
#!/bin/bash
# Run rman using catalog to backup incremental
rman TARGET / CATALOG rman/rman@rman CMDFILE rman_inc.rcv

# Run clone to end of backup
NOW=$(date +%Y%m%d_%H%M)
loginInline="ptadmin,ptadmin"
sourceFS=ORC2_RMAN1
targetFS=ORC2_RMAN1
sourceDir="ORC2_RMAN1/RMAN_backups"
targetDir="ORC2_RMAN1/clone_${NOW}"
```

```
ssh ptadmin@protectier /opt/dtc/ptcli/ptcli CloneDirectory --loginInline
$loginInline --force --sourceFS $sourceFS --targetFS $targetFS --sourceDir
$sourceDir --targetDir $target-Dir
```

Example 20-10 RMAN Backup script (rman_inc.rcv) using 12 channels

```
run {
allocate channel c1 device type disk format '/mnt/ORC2_RMAN1/RMAN_backups/%U' ;
allocate channel c2 device type disk format '/mnt/ORC2_RMAN1/RMAN_backups/%U' ;
allocate channel c3 device type disk format '/mnt/ORC2_RMAN1/RMAN_backups/%U' ;
allocate channel c4 device type disk format '/mnt/ORC2_RMAN1/RMAN_backups/%U' ;
allocate channel c5 device type disk format '/mnt/ORC2_RMAN1/RMAN_backups/%U' ;
allocate channel c6 device type disk format '/mnt/ORC2_RMAN1/RMAN_backups/%U' ;
allocate channel c7 device type disk format '/mnt/ORC2_RMAN1/RMAN_backups/%U' ;
allocate channel c8 device type disk format '/mnt/ORC2_RMAN1/RMAN_backups/%U' ;
allocate channel c9 device type disk format '/mnt/ORC2_RMAN1/RMAN_backups/%U' ;
allocate channel c10 device type disk format '/mnt/ORC2_RMAN1/RMAN_backups/%U' ;
allocate channel c11 device type disk format '/mnt/ORC2_RMAN1/RMAN_backups/%U' ;
allocate channel c12 device type disk format '/mnt/ORC2_RMAN1/RMAN_backups/%U' ;

CROSSCHECK BACKUP;

RECOVER COPY OF DATABASE WITH TAG 'LVLO_MERGE_INCR';

BACKUP CHECK LOGICAL INCREMENTAL LEVEL 1 CUMULATIVE COPIES=1
FOR RECOVER OF COPY WITH TAG 'LVLO_MERGE_INCR' DATABASE;

sql 'ALTER SYSTEM ARCHIVE LOG CURRENT';

BACKUP CHECK LOGICAL AS COMPRESSED BACKUPSET FILESPERSET 10 ARCHIVELOG ALL DELETE
INPUT;

DELETE NOPROMPT OBSOLETE;
DELETE NOPROMPT EXPIRED BACKUP;

release channel c1;
release channel c2;
release channel c3;
release channel c4;
release channel c5;
release channel c6;
release channel c7;
release channel c8;
release channel c9;
release channel c10;
release channel c11;
release channel c12;
}
```

When using incremental backup, a strong suggestion is to enable Change Block Tracking for faster incremental backup:

```
ALTER DATABASE ENABLE BLOCK CHANGE TRACKING;
ALTER DATABASE ENABLE BLOCK CHANGE TRACKING USING FILE
'/mnt/RMAN_backup/rman_change_track.f' REUSE;
```

20.5.4 Using ProtecTIER NAS to store Oracle Data Pump exports

Data Pump Export is an ready-for-use Oracle utility for unloading data and metadata into a set of operating system files called a dump file set. The dump file set can be imported only by the Data Pump Import utility. The dump file set can be imported on the same system, or it can be moved to another system and loaded there.

The dump file set is made up of one or more disk files that contain table data, database object metadata, and control information. The files are written in a proprietary, binary format. During an import operation, the Data Pump Import utility uses these files to locate each database object in the dump file set and import it to the database.

Export and import are commonly used for creating logical backups of a full database, schema or a single object.

Oracle Data Pump Export and Import can use the ProtecTIER NAS location, when mounted as an NFS location using the mounting options and the event setting described previously.

Because the export dump files are binaries, the deduplication ratio can vary based on changes made to the exported objects in a single dump file. Compressing the dump files using the **COMPRESSION=DATA_ONLY** parameter will cause the deduplication to be less efficient, and is therefore *not* recommended. Using the parallelism option **PARALLEL=n** parameter will multiplex the export into files and will speed up the export operations.

Deduplication on export files is decent, and was tested to be as good as 1:10 on multiple exports to the same object.

20.5.5 Using ProtecTIER NAS to store Oracle database files and offline redo logs

Placing data files on the ProtecTIER storage will affect data base performance. This kind of usage might hinder the performance of the ProtecTIER as well. Because Oracle data files are always open and data is being written into it, putting the data files on ProtecTIER storage will cause the ProtecTIER to constantly evaluate the data for deduplication.

Attention: Even though the ProtecTIER NAS was not designed for it, internal tests have proven that it was possible to create table spaces on it, but it is highly *not* recommended.

Using ProtecTIER NAS as a single archive log destination will work, but again, it is *not* recommended. The purpose of using ProtecTIER NAS is to provide a target for backups. Using ProtecTIER NAS as a single archive log destination solution might affect the database performance, and in some cases might cause the database to hang.

The use of the ProtecTIER storage to reduce costs by storing read-only or offline table spaces was tested and found to be working and valid.

20.5.6 Other suggestions for RMAN

When working with RMAN and ProtecTIER, other suggestions are relevant:

- ▶ As a preferred practice, use client mount parameters for Oracle database servers:

```
mount -o rw,soft,intr,nolock,timeo=3000,nfsvers=3,proto=tcp protectier:/RMAN /mnt/RMAN
```
- ▶ Disable Oracle NFS check by using the following SQL command for each Oracle database instance:

```
SQL> alter system set event="10298 trace name context forever, level 32";
```
- ▶ To disable Oracle NFS checks during the lifecycle of the database instance, use the following command:

```
SQL> alter system set event="10298 trace name context forever, level 32" scope = spfile;
```
- ▶ To enable Oracle NFS checks (when no longer using RMAN backups to ProtecTIER NAS), use the following command:

```
SQL> alter system set events '10298 trace name context off';
```
- ▶ Do not use ProtecTIER NAS to store or migrate tablespaces, tables or datafiles. It is to be used only as a target for RMAN backups.
- ▶ Do not use ProtecTIER NAS as primary or alternate archive log destination.
- ▶ Because this scenario suggests disabling the Oracle NFS check, Oracle Servers that are using NAS from other vendors (Netapp, EMC, Hitachi, or IBM N-Series) to store Oracle tablespaces should not use ProtecTIER NAS for Oracle RMAN backups.
- ▶ Oracle Real Application Clusters (RAC) are currently not supported.
- ▶ Oracle Direct NFS (dNFS) is being tested and evaluated, but is currently not supported. Use standard client kernel NFS instead.

In summary, ProtecTIER can be deployed as either Serial Backup Tape (SBT) or Disk target for RMAN backups. Using ProtecTIER NAS as a target for RMAN disk backups greatly strengthen the benefits of Oracle database administrators (DBAs) gaining complete control of their RMAN backups, and at the same time using the benefits of deduplication and scalability enhancements.

Additionally, disabling hard mount options will ensure the Oracle DB that no issues should occur with the database being crashed during timeout between the ProtecTIER NAS unit and the Oracle DB server.

20.6 SAP

This section describes settings and parameters to be modified for optimum performance when you are working with specific data types, such as SAP integrated with Tivoli Storage Manager.

Note: Beginning with version 7.1.3, Tivoli Storage Manager was rebranded to IBM Spectrum Control. The scenarios in this topic were conducted with a version earlier than 7.1.3. For legacy and reference purposes, we continue to refer to Tivoli Storage Manager.

20.6.1 SAP introduction

SAP is an acronym for *Systems Applications and Products*. SAP provides a common centralized database for all the applications that are running in an organization. The database instance is a mandatory installation component for the installation of an SAP system.

SAP supports the following databases:

- ▶ Oracle
- ▶ MS SQL Server
- ▶ IBM DB2 Universal Database™ for UNIX and Windows
- ▶ SAP liveCache technology
- ▶ MaxDB
- ▶ IBM DB2 Universal Database for z/OS®
- ▶ IBM DB2 Universal Database for iSeries
- ▶ IBM Informix®

For more database and operating system support information, see the Product Availability Matrix (PAM) at the SAP Service Marketplace. Log in at the following address:

<http://service.sap.com/pam>

20.6.2 Data protection for SAP

Data protection for the SAP server involves steps to protect all of the software components that are needed by the SAP system to operate. The base components of the SAP server are the operating system, the SAP application server, the database instance, and the data files. Each component requires different data protection techniques.

The SAP system uses the relational database as main storage for all SAP data and meta information. This main storage is the basis for the tight integration of all SAP application modules and ensures consistent data storage. Data in the SAP database is unique for every company, and if the data is lost, it cannot be simply reinstalled in the same manner as an operating system is reinstalled. Therefore, be especially careful when you plan the protection of the data that is stored in the SAP database.

Protection of the SAP database

The protection of the SAP database has two parts: protecting the database binary files and configuration files, and protecting data that is stored in the data files.

Database binary files and configuration files are typically backed up as part of the operating system or file system backup. The backup of the database data files and other supporting structures that are associated with SAP data should be performed by a dedicated tool that is designed especially for the database backup. You can use database backup and restore tools to perform backup and restore data in a consistent state.

The backup tools can also perform an online backup of the database and backup of the redo log files just after the log files are archived. A backup of a database creates a copy of the database's data files, control files, and, optionally, log files. It then stores these files on backup media.

A *consistent backup*, also called an offline backup or cold backup, is a backup of all the data files in the database that is taken when all interim changes are physically written to the data files. With a consistent backup, partial changes from the log files that are not written to the data files are not backed up. If you restore a database from a consistent backup, the database is in a consistent state when the restore operation finishes.

Also note the following information:

- ▶ For an Oracle database, a consistent backup can be taken only when the database is shut down for the entire duration of the backup procedure.
- ▶ For a DB2 Universal Database (UDB), a database must be deactivated, or the instance must be stopped before the backup operation starts.

The database must stay inactive until the backup finishes. This action ensures that there are no data changes on the database at the time the backup is being taken. A consistent backup is always a backup of the entire database; it cannot be a partial or incremental backup.

You can take an offline backup by using either a dedicated database backup tool, such as Oracle Recovery Manager, BR*Tools, the **DB2 BACKUP** command, or a non-database backup tool, such as the Tivoli Storage Manager backup archive client. The dedicated database backup tools ensure that all the objects that are required for the successful database restore are included in the backup image.

The database backup tool also ensures that the location, time stamp, type, and other information about the backup copy is registered in the repository, such as BR*Tools logs or a DB2 database history file. Using the metadata in the repository, backup tools can perform an automatic restore that is based on the specified time stamp without prompting for the backup images to restore and their location.

IBM offers products for both data protection and data retention and reduction. For example, in the SAP environment, there are Tivoli Storage Manager for Enterprise Resource Planning (ERP) and Tivoli Storage Manager for Database for data protection. For data retention, you can use IBM DB2 CommonStore for SAP. Both solutions can use a Tivoli Storage Manager server for the media manager.

Note: Beginning with version 7.1.3, Tivoli Storage Manager was rebranded to IBM Spectrum Control. The scenarios in this topic were conducted with a version earlier than 7.1.3. For legacy and reference purposes, this book continues to use the name Tivoli Storage Manager.

Tivoli Storage Manager for ERP, formerly known as Tivoli Data Protection for SAP, is a component of Tivoli Storage Manager family that provides a complete backup solution for SAP databases. The current version supports Oracle and DB2 only.

The following features are available for Tivoli Storage Manager for ERP:

- ▶ Handles large amounts of data
- ▶ Optimized processor usage that reduces the overall time for backup and restore
- ▶ Optimized for an SAP environment
- ▶ Supports multiple management classes

These solutions are illustrated in Figure 20-8.

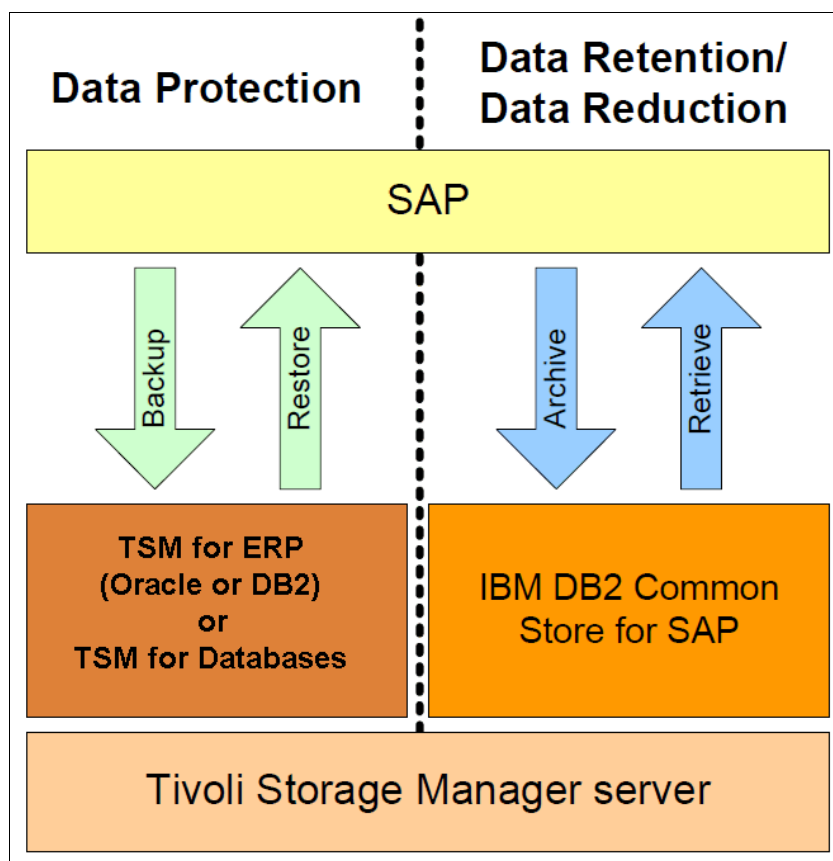


Figure 20-8 Backup and archival products

Additional information: For more information, see the following web page:

<http://www.ibm.com/software/tivoli/products/storage-mgr-erp/>

20.6.3 Integration of Tivoli Storage Manager for ERP with SAP

Tivoli Storage Manager for ERP is fully integrated in to the SAP environment. The communication between the backup and archive server is performed by an application programming interface (API) called *ProLE*. This API is shared with other Tivoli Data Protection products. ProLE runs as a background process and provides communication with the Tivoli Storage Manager server. Figure 20-9 on page 326 shows a sample architecture of Tivoli Storage Manager for ERP integrated with SAP.

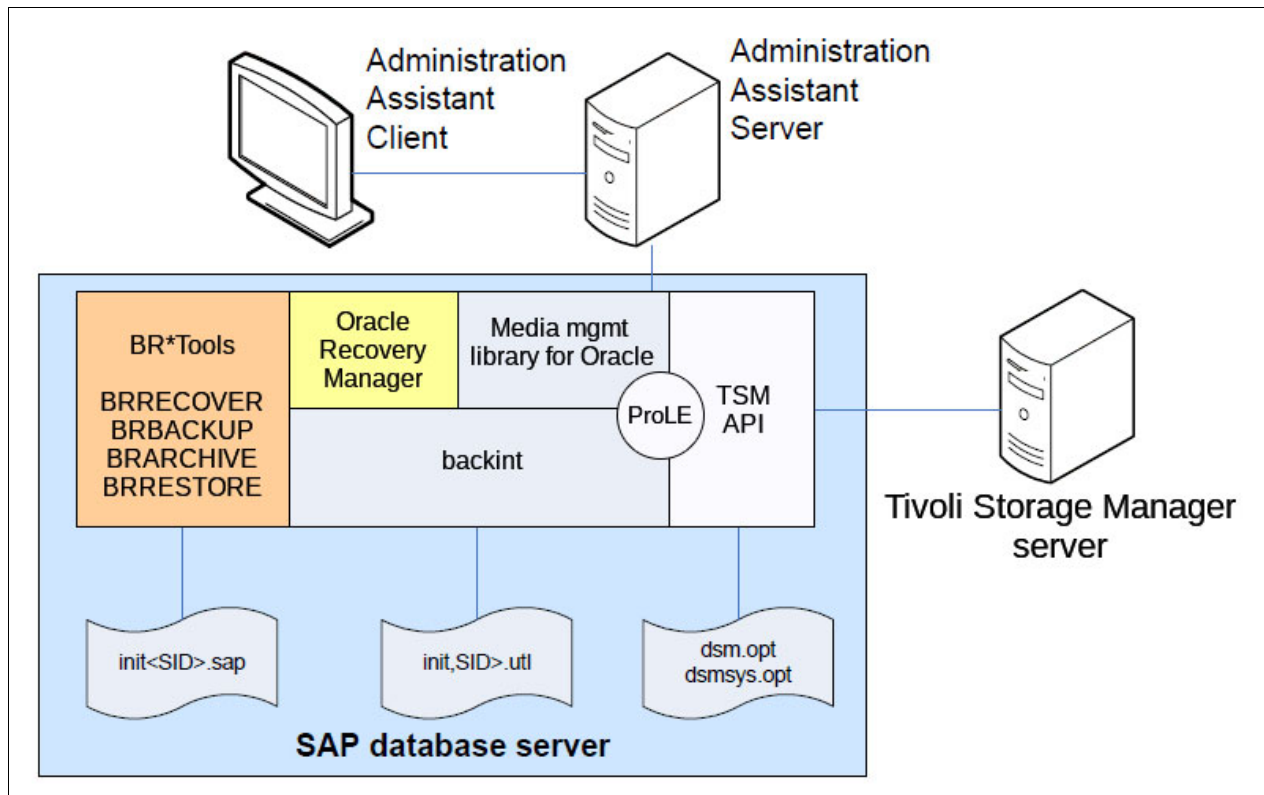


Figure 20-9 Tivoli Storage Manager for ERP sample scenario

Additional information: For more information, see the IBM Knowledge Center:

http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3/index.jsp?topic=%2Fcom.ibm.itism.nav.doc%2Ft_protect_dperp.html

20.6.4 Tivoli Storage Manager for ERP for Oracle database

Tivoli Storage Manager for ERP is a client and server program that manages backups and restores in conjunction with the Tivoli Storage Manager. With Tivoli Storage Manager for ERP, it is possible to handle SAP database backups, and it includes the ability to manage backup storage and processing independently from normal SAP operations.

Furthermore, Data Protection for SAP in combination with Tivoli Storage Manager provides reliable, high performance, and repeatable backup and restore processes to manage large volumes of data more efficiently.

For Oracle databases, two options exist to implement a backup using Tivoli Storage Manager:

- ▶ Tivoli Storage Manager for ERP using the **BACKINT** interface
- ▶ Tivoli Storage Manager for ERP using Oracle Recovery Manager (RMAN)

With the integration, it is possible to follow the ERP backup and restore procedures and to use the integrated SAP database utilities **BRBACKUP**, **BRARCHIVE**, **BRRESTORE**, and **SAPDBA** for backup and restore. Other SAP-related files (executable files) are backed up by using Tivoli Storage Manager standard techniques for file backup and restore, for example, incremental backup, file filtering, and point-in-time recovery.

Tivoli Storage Manager for ERP for Oracle using BACKINT

Using this feature, you can perform the traditional Oracle online backup with automation provided by **BACKINT**. Figure 20-10 shows the data interface between Oracle Databases and Tivoli Storage Manager for ERP for Oracle using the **BACKINT** interface.

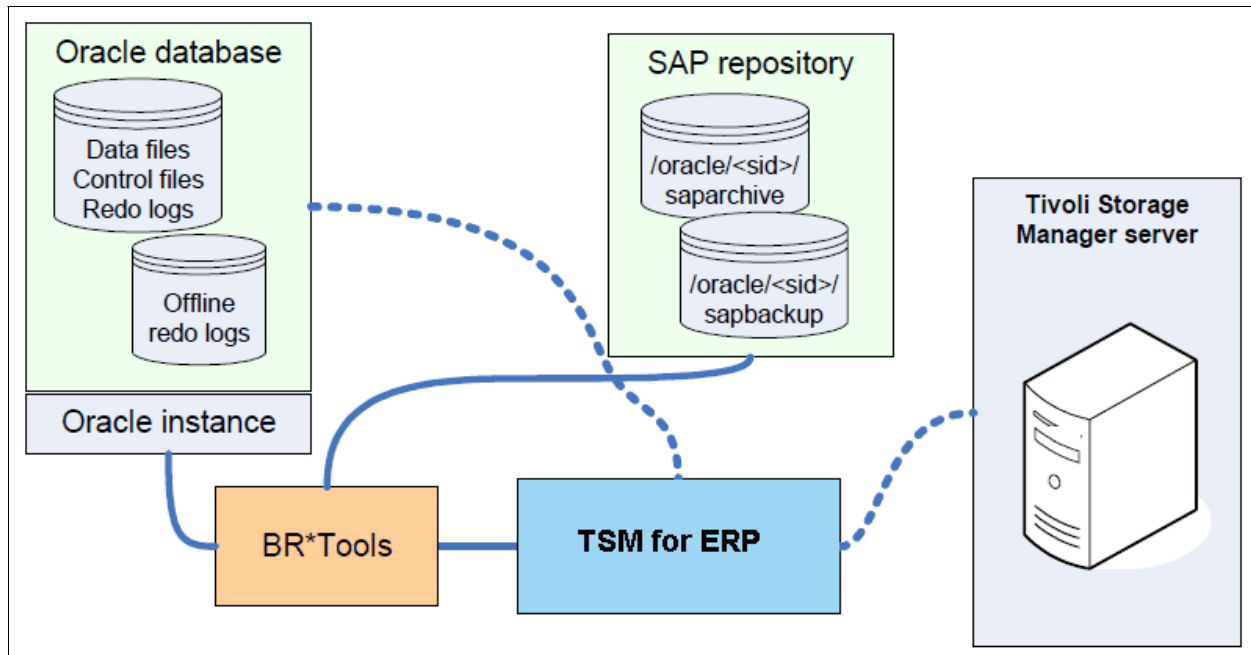


Figure 20-10 Tivoli Storage Manager for ERP for Oracle using BACKINT

The backup proceeds as follows:

1. BR*Tools takes control.
2. **BRBACKUP** calls the Tivoli Storage Manager for ERP by using **BACKINT**.
3. **BACKINT** changes the table spaces to backup mode with the following command:

```
alter tablespace <tablespace name> begin backup
```
4. **BACKINT** using Tivoli Storage Manager for ERP reads all the data files and saves them to Tivoli Storage Manager server.
5. BR*Tools updates the catalog with information about the backed up data file.

Logs: BR*Tools logs are stored in the `/oracle/<SID>/saparch` directory.

BRBACKUP automatically backs up the logs and profiles after every backup operation. In the case of bare metal restore or disaster recovery, logs and profiles must be restored to enable BR*Tools to restore data files. The process can be simplified if the logs and profiles are backed up by a Tivoli Storage Manager backup archive client during the file system backup.

Using this method, the chosen data files are sent to Tivoli Storage Manager one by one. No compression or block checking is performed at this level.

When a database is in backup mode, the amount of redo logs that are written to disk increases because Oracle writes the entire dirty block to the disk, not just the updated data. In some cases, when the backup routine fails for any reason, the data file remains in active backup mode, which can cause some performance effect and additional I/O to the disk.

Tivoli Storage Manager for ERP for Oracle using RMAN

Using this feature, you can take advantage of all the facilities that are provided by RMAN. In general, RMAN is able to perform a backup in less time compared to the traditional backup using **BACKINT** because RMAN sends only used data blocks (in an Oracle data file) to Tivoli Storage Manager. The other interesting feature is block checking, which discovers bad blocks as soon as they occur.

In addition, you can use the Oracle Recovery Manager (RMAN) utility to run some tasks that are not provided by BR*Tools, such as incremental backups, releasing backup versions, and catalog maintenance.

Note: Beginning with version 7.1.3, Tivoli Storage Manager was rebranded to IBM Spectrum Control. The scenarios in this topic were conducted with a version earlier than 7.1.3. For legacy and reference purposes, we continue to refer to Tivoli Storage Manager.

Figure 20-11 shows the data interface between Oracle Database and Oracle for SAP using RMAN.

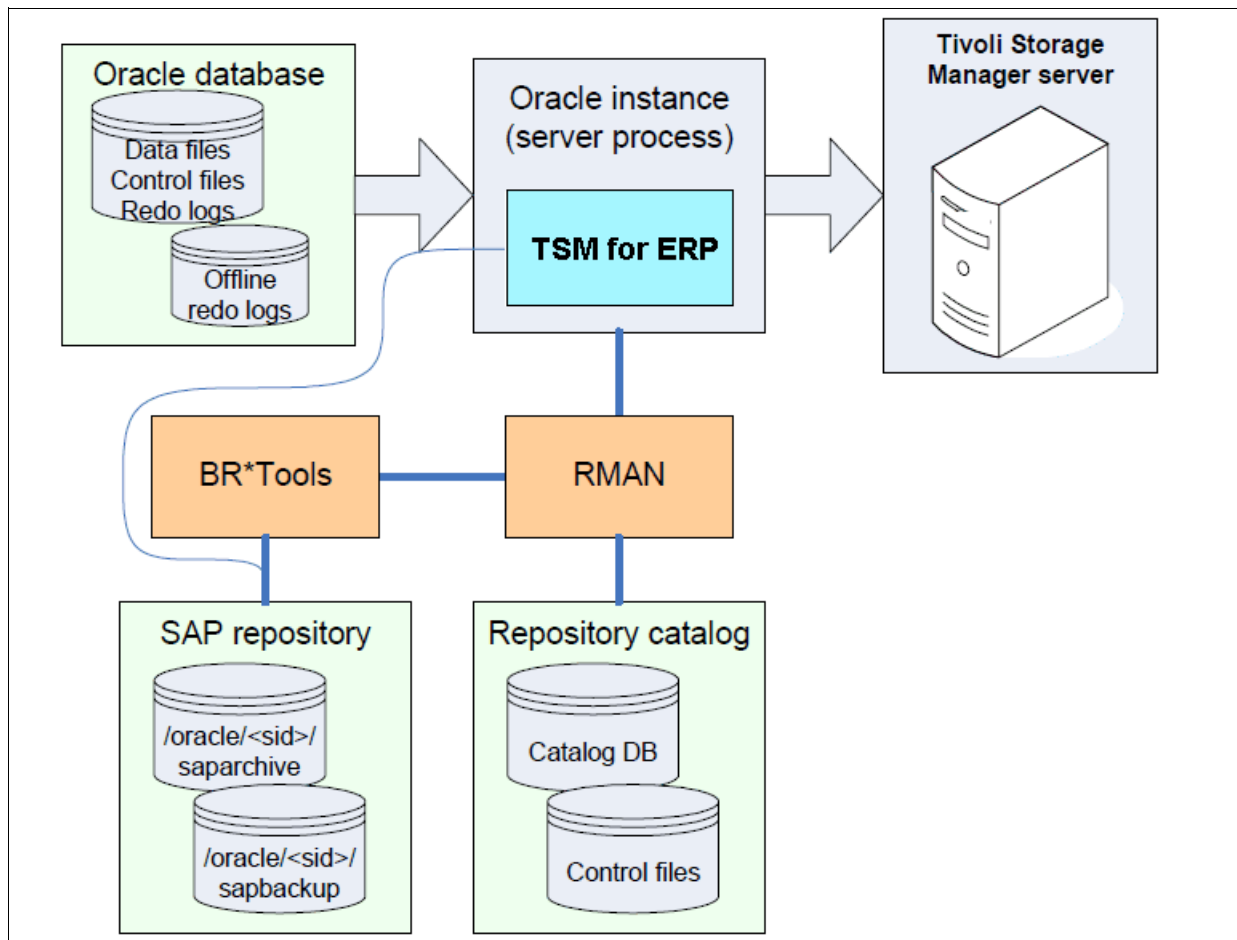


Figure 20-11 Tivoli Storage Manager for ERP for Oracle using RMAN

20.6.5 Tivoli Storage Manager for ERP for DB2

Tivoli Storage Manager for ERP for Oracle for DB2 was created to provide an intelligent interface to manage backup and restore by using Tivoli Storage Manager. It is fully integrated in to the SAP environment. The backup command **DB2 BACKUP DATABASE** and the restore command **DB2 RESTORE DATABASE** are run at the DB2 CLI, which calls the Tivoli Data Protection for SAP for DBA module.

The backup and restore of the DB2 log files is provided by the BR*Tools commands **BRARCHIVE** and **BRRESTORE**. In addition, you can use the Tivoli Storage Manager for ERP for DB2 Tools BackOM and the built-in Log Manager.

Figure 20-12 shows the data interface between DB2 Databases and Tivoli Storage Manager for ERP for DB2.

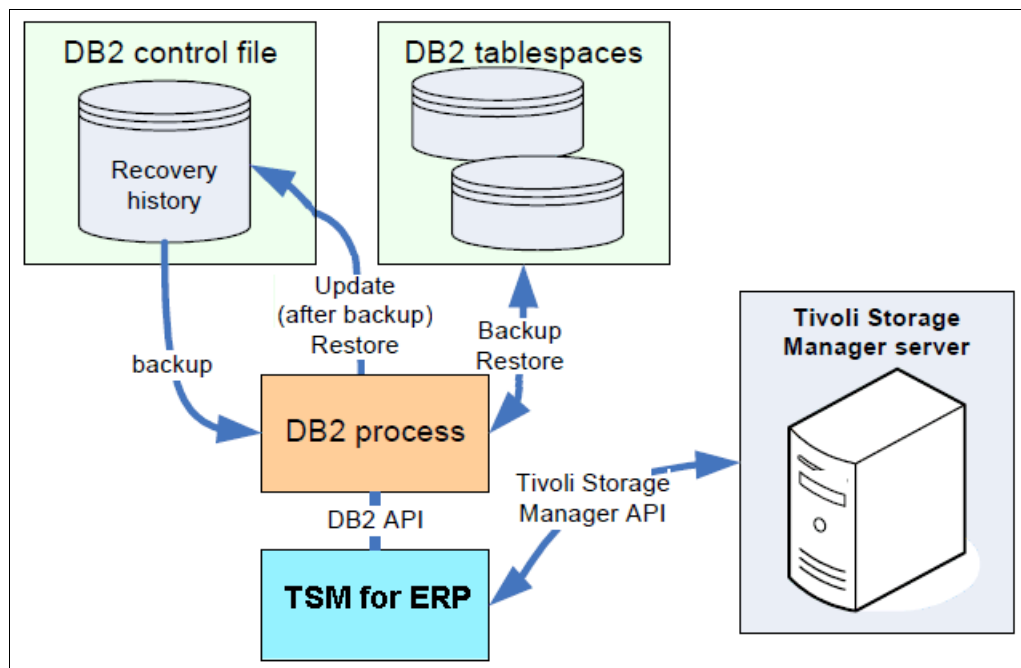


Figure 20-12 Tivoli Storage Manager for ERP for DB2

The archiving of DB2 offline log files is provided by the SAP tool **BRARCHIVE**. The retrieval of DB2 offline log files is provided by the SAP tool **BRRESTORE** and by the Tivoli Storage Manager for ERP tool BackOM. As of DB2 Version 9.X, offline log files can be archived and retrieved with the DB2 built-in Log Manager.

The DB2 command line processor (CLP) interprets commands for the DB2 database and passes control to a DB2 Server Process. In the case of Tivoli Storage Manager for ERP, the **LOAD <libraryname>** option causes DB2 to start the Tivoli Storage Manager for ERP shared library.

This process runs during the backup or restore, loads the library dynamically, and communicates with it through the Tivoli Storage Manager API. To start a backup or restore, the DB2 CLP communicates with the DB2 Server Process, providing the server process with the relevant information for processing the database.

Additional information: For more information about backup methodologies for SAP, see *SAP Backup using Tivoli Storage Manager*, SG24-7686.

All backup solutions that are described in this section can be integrated with advanced backup techniques, such as LAN-free backup, parallel transfer of backup data to and from Tivoli Storage Manager server, or multiplexing.

Reduction of processing time: Implementation of these techniques can reduce backup and restore times, and eliminate the effect of backup data transfers on LAN throughput.

20.6.6 SAP BR*Tools for Oracle using BACKINT

SAP BR*Tools for Oracle is a package of utilities, developed by SAP AG to protect and manage SAP data that is stored in Oracle databases. BR*Tools supports functions for online, offline, partial, or full backups of databases (**BRBACKUP**) and backups of archived redo logs (**BRARCHIVE**). It has functions for database restore and recovery (**BRRECOVER** and **BRRESTORE**).

In addition to using BR*Tools for database recoverability tasks, you can also have it serve as a tool for creating homogeneous database copies, and to assist with database migration to different platforms or database versions.

BRBACKUP is the BR*Tools utility that enables online or offline backup of database files (data files, control files, and online redo log files). **BRBACKUP** can be used to back up individual data files, table spaces, or the entire Oracle database. **BRBACKUP** also backs up the BR*Tools configuration profiles and logs that are required for the database's disaster recovery.

The smallest unit that can be saved with **BRBACKUP** is a file. You can use **BRBACKUP** for backing up both files in the database and non-database files and directories. Use the **backup_mode** command from the Initialization Profile `init<DBSID>.sap` file or the **brbackup -m|-mode** command option for this purpose.

Before the offline backup is taken, **BRBACKUP** automatically closes the database and opens it when the backup is accomplished. **BRBACKUP** can also change the status of the table space to be backed up to **BEGIN/END BACKUP**.

You can also instruct **BRBACKUP** to use software compression. The software compression client can enhance the backup, especially if the network is slow.

Compression: If you plan to send data to a ProtecTIER server, do not enable software compression; it might affect the overall deduplication ratio.

The most frequently used **BRBACKUP** function is a full database backup. Example 20-11 shows **BRBACKUP**.

Example 20-11 Online backup by using the databases BRBACKUP tool

```
$su - cptadm
$BRBACKUP -c -u / -t ONLINE_CONS -m FULL -p /oracle/CPT/102_64/dbs/initCPT.sap
```

You can perform a full backup by running **BRBACKUP** with the following options:

- ▶ The mode option (**-mode/-m**) is set to FULL or ALL.
- ▶ You can start a full backup either in online mode (**-type/-t online_cons**) or in offline mode (**-type offline**). In the case of the **online_cons** type, the offline redo log files that are generated during the full backup are also backed up to the same media.

- The backup storage media is defined by the BR*Tools profile file that is specified by the **BRBACKUP** parameter **-profile/-p**.
- The user name and password that is used by **BRBACKUP** to log on the Oracle database system is specified by the parameter **-user/-u**. If you are working as a DBA user that is authenticated to the database by the OS (\$OPSuser), you can use "/" as value of this parameter.

The parameter **-confirm/-c** stands for an unattended mode, which is mostly used in the backup scripts, so BR*Tools does not prompt you for confirmations.

Archived redo log backup functions

BRARCHIVE provides functions for offline redo log files backup in Oracle databases that run in archiving mode. If archiving is enabled, a database cannot overwrite an active log file until the content is archived. Whenever an active redo log is filled, the database performs a log switch and starts writing to another log file. The full redo log files are archived by Oracle background processes into the archive log directory.

The redo log is the most important database component for a recovery from a crash, media failure, or user failure. Therefore, at least the production databases should be configured in archiving mode. To prevent the archive log directory from filling up, **BRARCHIVE** should be run periodically to move the offline redo logs from the archive directory to the backup media.

BR*Tools and Tivoli Storage Manager for ERP with Oracle

BR*Tools interacts with Tivoli Storage Manager for ERP with Oracle through the **BACKINT** interface. The communication of BR*Tools and **BACKINT** occurs as follows:

1. The BR*Tools utility **BRBACKUP** informs Oracle of what data must be backed up and puts the database into the correct backup state (online or offline backup).
2. **BRBACKUP** calls Tivoli Data Protection for ERP using the **BACKINT** interface with a list of all files to be backed up.
3. Tivoli Data Protection for ERP reads all the requested files from the database and reports back to **BRBACKUP**. **BRBACKUP** adds these files to the repository that contains all processed backups.
4. **BACKINT** transfers the data to the Tivoli Storage Manager server by using the Tivoli Storage Manager Client API.
5. The BR*Tools updates the repository that contains information about the status of the files.

BR*Tools configuration

To configure BR*Tools, complete the following steps:

1. The BR*Tools configuration is stored in the `init<SID>.sap` initialization profile file. The configuration file contains parameters that affect the performance of backup and restore functions. The file is in the following default locations:
 - On UNIX: `<ORACLE_HOME>/dbs`
 - On Windows: `<ORACLE_HOME>\database`
2. Some parameters that are specified in the profile can be overridden if the BR*Tools programs are called with different command options. In the BR*Tools profile, you can specify the backup adapter that is used to transfer data (**cpio**, **BACKINT**, or **RMAN**).
3. If you set up BR*Tools to use the **BACKINT** adapter, you need to reference the appropriate **BACKINT** profile (*.utl file) in the BR*Tools profile. If you want to instruct BR*Tools to use Oracle **RMAN**, you must define the **RMAN** channel parameters in the BR*Tools profile.

4. The configuration profile of Tivoli Storage Manager for ERP is defined in the `init<SID>.utl` file, which is in the same directory as the `BR*Tools` profile (`init<SID>.sap`). The configuration parameters in the `init<SID>.utl` file include the Tivoli Storage Manager node name and the management classes to be used for data files backup and offline redo logs backup. If the backup retention is going to be controlled by Tivoli Storage Manager for ERP, you can set up the number of backup versions to be kept in this file.
5. The configuration file of the Tivoli Storage Manager API (`dsm.sys`) is, by default, stored in the Tivoli Storage Manager API installation directory (specified by the environmental variable `DSMI_DIR`). The configuration file of the Tivoli Storage Manager API Client defines the network settings (protocol and network address of the Tivoli Storage Manager server) to enable communication between the API client and the Tivoli Storage Manager server.
6. You also specify in this file the authentication type (**PASSWORDACCESS**) that the Tivoli Storage Manager API client uses to connect to the Tivoli Storage Manager server. Additionally, if the Storage Agent is operable on the local node in this file, you can instruct the Tivoli Storage Manager API client to use the LAN-free backup (by using the **LANFREE yes|no** option).
7. Instruct BR*Tools to use the **BACKINT** interface by setting the **backup_dev_type** parameter in the SAP initialization file (`init<SID>.sap`) as follows:

```
backup_dev_type = util_file
```
8. Instruct BR*Tools to use the `init<SID>.utl` file (created by the Tivoli Storage Manager for ERP installation wizard) by setting the **util_par_file** parameter in the SAP initialization file:

```
util_par_file=<full_path>/init<SID>.utl
```

Archiving functions: Tivoli Storage Manager for ERP uses the Tivoli Storage Manager archive functions to transfer data to Tivoli Storage Manager server and back. Therefore, the management classes that are assigned to Tivoli Storage Manager for ERP (in the `init<SID>.utl` file) must have an archive copy group defined.

For more information, see the IBM Knowledge Center:

http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3/index.jsp?topic=%2Fcom.ibm.itsm.nav.doc%2Ft_protect_dperp.html

20.6.7 SAP BR*Tools for Oracle using RMAN with Tivoli Storage Manager

Configure BR*Tools for use with the RMAN Tivoli Storage Manager channel as follows:

- On the Tivoli Storage Manager server, complete the following steps:
 - a. Define a policy domain with two management classes that are used to transfer data and logs. Define an archive management class in each of the management classes. If the retention control is performed at the Tivoli Storage Manager server, specify **RETVER=<days>** for each archive copy group. If the retention control is performed at Tivoli Storage Manager for ERP, specify **RETVER=no limit**.
 - b. Register the Tivoli Storage Manager node with the defined domain. Update the parameter **MAXNUMMP** for the Tivoli Storage Manager node to **MAXNUMMP=2** (based on the parallelism that is required).
- On the client node, complete the following steps:
 - a. Update or create the `DSM.OPT` and `DSM.SYS` files to configure the Tivoli Storage Manager API client. The **PASSWORDACCESS** parameter must be set to "PROMPT" in this configuration.
 - b. Set up the environment values **DSMI_DIR** and **DSMI_LOG** for the Oracle OS user.

- c. Install IBM Tivoli Storage Manager for ERP - Oracle on the Oracle server with SAP installed on it.
- d. Configure the client resources for Oracle server in the IBM Tivoli Storage Manager for ERP configuration file (<ORACLE_HOME>\dbs\init<SID>.utl).
- e. Check the defined Tivoli Storage Manager node name and Tivoli Storage Manager management classes to be used for the backup of offline redo log files and data files. Ensure that the **SERVER** parameter refers to an existing stanza in the DSM.SYS file.
- f. If the retention control is driven by Tivoli Storage Manager for ERP, set the **MAX_VERSIONS** parameter.
- g. Switch to the Oracle instance owner and update the Tivoli Storage Manager node password for Oracle by running the following command:

```
backint -p <ORACLE_HOME>\dbs\init<SID>.utl -f password
```
- h. Ensure that RMAN can access the Tivoli Storage Manager for ERP API. The following links must exist (be created):
 - `ln -s /usr/tivoli/tsm/tdp_r3/ora/libtdp_r3.<ext>`
 - `/usr/lib/libobk.<ext> ln -s /usr/lib/libobk.<ext>`
 - `$ORACLE_HOME/lib/libobk.<ext>`
- i. Instruct BR*Tools to use RMAN by setting the **backup_dev_type** and **rman_parms** options in the SAP initialization file (init<SID>.sap) as follows:
 - `backup_dev_type = rman_util`
 - `rman_parms="ENV=(XINT_PROFILE=<ORACLE_HOME>/dbs/init<SID>.utl,PROLE_PORT=<portnumber>,&BR_INFO)"`
- j. Instruct BR*Tools to use the init<SID>.utl file for Tivoli Storage Manager specific parameters by setting the **util_par_file** parameter in the SAP initialization file:

```
util_par_file=<path to Tivoli Storage Manager for ERP util file - init<SID>.utl>
```

20.6.8 SAP BR*Tools for Oracle: Using RMAN to configure DB2 to use Tivoli Storage Manager

Configure DB2 to use Tivoli Storage Manager for ERP as follows:

- On the Tivoli Storage Manager server, complete the following steps:
 - a. Define a policy domain with two management classes that are used to transfer data and logs. Define an archive copy group for both management classes. If the retention control is performed at the Tivoli Storage Manager server, specify **RETVER=<days>** for each archive copy group. If the retention control is performed at Tivoli Storage Manager for ERP level, specify **RETVER=no limit**.
 - b. Register Tivoli Storage Manager node with the defined domain. Update the parameter **MAXNUMMP** for Tivoli Storage Manager node to **MAXNUMMP=2** (based on the parallelism that is required).
- On the client node, complete the following steps:
 - a. Update or create the Tivoli Storage Manager API client option files DSM.OPT and DSM.SYS. The **PASSWORDACCESS=GENERATE** parameter must be set for this configuration.
 - b. Configure the environment values **DSMI_DIR**, **DSMI_CONFIG**, and **DSMI_LOG** in the DB2 instance owner user's profile. You must restart the DB2 instance to make the parameters effective for DB2.

- c. Install Tivoli Storage Manager (As of version 7.1.3 known as IBM Spectrum Protect) for ERP - DB2 on the DB2 UDB server, with SAP already installed. You can use the installation wizard to specify the name of the Tivoli Storage Manager server stanza (in DSM.SYS), the Tivoli Storage Manager node name, and the management classes to be used for the backup of data and archived logs.
- d. Check the client resource for the Tivoli Storage Manager server in the Tivoli Storage Manager for ERP configuration file /db2/<SID>/tdp_r3/init<SID>.utl. Verify that the following environment variables are set correctly in the DB2 owner user's profile:
 - XINT_PROFILE
 - DB2_VENDOR_LIB
 - TDP_DIR
- e. Switch to the DB2 instance owner and update the Tivoli Storage Manager client password for DB2 node by running the following command:


```
$/usr/tivoli/tsm/tdp_r3/db264/backom -c password
```
- f. Restart the DB2 instance.
- g. Optionally, you can set up DB2 automatic log management so that the archived logs are sent to Tivoli Storage Manager by using the Tivoli Storage Manager media management library that is provided by Tivoli Storage Manager for ERP. This task can be accomplished by setting the DB2 configuration parameters **LOGARCHMETH1** and **LOGARCHOPT1** as follows:
 - update db cfg for <SID> using LOGARCHMETH1
VENDOR:/usr/tivoli/tsm/tdp_r3/db264/libtdpdb264.a
 - update db cfg for <SID> using LOGARCHOPT1 /db2/<SID>/tdp_r3/vendor.env
- h. If you use the direct log backup method that is specified in step g, you should also specify the **FAILARCHPATH** db2 configuration parameter. **FAILARCHPATH** points to a directory that is used as a temporary storage for offline logs in case the Tivoli Storage Manager server is unavailable, which can prevent the DB2 from filling up the log directory. This is the command syntax:


```
update db cfg for <SID> using FAILARCHPATH <offline log path>
```

20.6.9 Preferred practices for Tivoli Storage Manager for ERP with ProtecTIER

The configuration profile of Tivoli Storage Manager for ERP is defined in the `init<SID>.utl` file, which is in the same directory as the `BR*Tools` profile (`init<SID>.sap`).

When the ProtecTIER VTL is defined for Tivoli Storage Manager, some settings must be done in Tivoli Storage Manager for ERP to optimize this integration.

Set the following information in the `init<SID>.utl` file:

- Disable multiplexing. The **MULTIPLEXING** parameter specifies how many files are read simultaneously and are multiplexed. If a file is multiplexed, it can affect the deduplication ratio. Set **MULTIPLEXING=1**.
- Use as many backup sessions in parallel as possible. The **MAX_SESSIONS** parameter defines the number of parallel sessions to be established. The valid range of **MAX_SESSIONS** is 1 - 32. Also define the **SESSIONS** parameter in each Tivoli Storage Manager stanza in the `.utl` file to specify the maximum number of sessions in that Tivoli Storage Manager server stanza.

Important: The **MAX_SESSIONS** parameter setting must not exceed the number of tape drives that are available simultaneously to the node in the Tivoli Storage Manager servers to be accessed. This maximum is established by the **MAXNUMMP** parameter settings in the Tivoli Storage Manager node definition.

- Disable compression by configuring **RL_COMPRESSION=NO**. The **RL_COMPRESSION** parameter specifies whether a null block compression of the data should be performed before transmission to Tivoli Storage Manager. Although **RL_COMPRESSION** introduces additional processor load to the SAP server, throughput can be improved when the network is the bottleneck, but it can affect the ProtecTIER deduplication ratio.

On the Tivoli Storage Manager server, complete the following steps:

1. Update the **MAXNUMMP** parameter for the Tivoli Storage Manager node to **MAXNUMMP=x**, where x should be the number of parallels required. This number should match the **MAXSESSION** parameter that is set in the `.utl` file. The **MAXNUMMP** parameter specifies the maximum number of mount points a node can use on the server only for operations, such as backup and archive.
2. Update the **COMPRESSION** parameter for the Tivoli Storage Manager node to **COMPRESSION=NO**. This setting specifies that the client node does not compress its files before it sends them to the server for backup and archive.

20.7 VMware

In addition to the now available vStorage APIs, the vStorage APIs for Data Protection (VADP) are also available. VADP replaces the VMware Consolidated Backup (VCB) framework, and offers multiple methods to improve your VMware backup. With the new VADP comes the option to use incremental virtual machine image backups by using the changed block tracking (CBT) feature.

In contrast to the full virtual machine image backup, CBT reduces the amount of backed up data because only the changed blocks that are compared to the last full backup are backed up. With CBT enabled, the backup operation backs up only the changed blocks, which results in a high data change rate for the ProtecTIER server, because only new data is backed up. For ProtecTIER deduplication to perform optimally, run at least one full backup per week.

Incremental backups: If you use incremental virtual machine image backups, run at least one full virtual machine image backup per week to optimize your deduplication ratio.

Follow the general preferred practices and the Tivoli Storage Manager (as of version 7.1.3, rebranded to IBM Spectrum Protect) practices that are described in Chapter 13, “IBM Spectrum Protect” on page 187.

20.7.1 Technical overview

VMware ESX is installed directly on the hardware and does not require any specific operating system. It is a virtualization platform that is used to create the virtual machines (VMs) as a set of configuration and disk files that perform all the functions of a physical machine.

vCenter Server

The vCenter server is a service that acts as a central administration point for ESX hosts that are connected to a network. This service directs actions on the virtual machines and the hosts. The vCenter server is the working core of the vCenter.

Multiple vCenter servers can be joined to a linked mode group, where you can log on to any single vCenter server to view and manage the inventories of all the vCenter server systems in the group.

With vCenter, an administrator can manage every component of a virtual environment. ESX servers, VMs, and extended functions, such as Distributed Resource Scheduler (DRS), vMotion, and VM backup, all access the vCenter server by using the vSphere Client GUI.

20.7.2 Settings and tuning for VMware and Tivoli Storage Manager

When you set up VMware for Tivoli Storage Manager (as of version 7.1.3, Tivoli Storage Manager has been rebranded to IBM Spectrum Protect) there are guidelines for using the vStorage API, *changed block tracking* (CBT), and format specification for formats for virtual disk files. This section provides a brief description of those guidelines.

vStorage API

The vStorage application programming interfaces (APIs) for data protection enable backup software to protect system, application, and user data in your virtual machines in a simple and scalable way. These APIs enable backup software to perform the following actions:

- ▶ Perform full, differential, and incremental image backup and restore of virtual machines
- ▶ Perform file-level backup of virtual machines using supported Windows and Linux operating systems
- ▶ Ensure data consistency by using Microsoft Volume Shadow Copy Services (VSS) for virtual machines that run supported Microsoft Windows operating systems

Changed block tracking (CBT)

Virtual machines that run on ESX/ESXi hosts can track disk sectors that change. This feature is called CBT. On many file systems, CBT identifies the disk sectors that are altered between two change set IDs. On Virtual Machine File System (VMFS) partitions, CBT can also identify all the disk sectors in use. CBT is useful when you set up incremental backups.

Virtual disk formats

When you perform certain virtual machine management operations (such as creating a virtual disk, cloning a virtual machine to a template, or migrating a virtual machine), you can specify a format for the virtual disk file. However, you cannot specify the disk format if the disk is on an NFS data store. The NFS server determines the allocation policy for the disk. The disk formats listed in this section are supported.

Thick format

This format is the default virtual disk format. The thick virtual disk does not change its size, and from the beginning occupies the entire data storage space that is provisioned to it. Thick format does not zero the blocks in the allocated space.

Conversion: Converting a thick disk format to a thin provisioned format is not possible.

Thin provisioned format

Use this format to save storage space. For the thin provisioned format, specify as much data storage space as the disk requires based on the value that you enter for the disk size. However, the thin disk starts small and, at first, uses only as much data storage space as the disk needs for its initial operations.

Thin disk considerations: If a virtual disk supports clustering solutions such as fault tolerance, you cannot make the disk thin. If the thin disk needs more space later, it can grow to its maximum capacity and occupy the entire data storage space that is provisioned to it. Also, you can manually convert the thin disk to thick disk.

20.7.3 Backup solutions

This section describes backup solutions and prerequisites for VMware backup by using the ProtecTIER server and Tivoli Storage Manager.

Full VM backup on ProtecTIER

You can use Tivoli Storage Manager Data Protection (DP) for VMware to back up and restore VM data through SAN-based data movement. The two data paths where data movement is possible as follows:

- ▶ The first data path is from the VMware data store to the vStorage server through SAN.
- ▶ The second data path is from the vStorage server to the ProtecTIER server. It could be SAN-based if Tivoli Storage Manager LAN-free is used.

The backup data path uses the *Tivoli Storage Manager for SAN* feature (LAN-free backup). In this book, LAN-free backup is used on VMware.

The DP for VMware stores virtual machine full backup images (full-VM) as a collection of control and data files. The data files contain the contents of virtual machine disk files, and the control files are small metadata files that are used during full VM restore operations and full VM incremental backups. In most cases, VMs are cloned according to a predetermined template. As a result, there is huge duplication of data. The ProtecTIER solution, in conjunction with Tivoli Storage Manager, deduplicates such data.

Prerequisites to VMware backup using ProtecTIER and Tivoli Storage Manager

Check that the following items are complete before you use VMware to back up your data with the ProtecTIER product and Tivoli Storage Manager (rebranded as IBM Spectrum Protect beginning with version 7.1.3):

- ☐ The ProtecTIER repository exists.
- ☐ The ProtecTIER deduplication function is enabled.
- ☐ The Tivoli Storage Manager server is installed with a license.
- ☐ The Tivoli Storage Manager storage agent is installed on a vStorage server, if LAN-free is used.
- ☐ The Tivoli Storage Manager Backup-Archive client is installed on the vStorage server.

Tip: For the best performance, the vStorage server must have separate HBA ports and each port must be connected to the ProtecTIER repository and the disk subsystem that stores the VMware data store.

Prerequisites for ESX/ESXi

Check that the following items are complete before you use VMware ESX/ESXi with the ProtecTIER repository and Tivoli Storage Manager:

- ☐ The host must be running ESX/ESXi Version 4.0 or later.
- ☐ The VMware vCenter Server must be Version 4.1.x or later.
- ☐ For incremental backup, the virtual machine that owns the disks to be tracked must be hardware Version 7 or later.
- ☐ CBT must be enabled for the virtual machine. (In the vSphere client, click **Edit** → **Settings** → **Options** → **Advanced/General** → **Configuration Parameters**.)
- ☐ The configuration of the virtual machine (.vmx) file must contain the following entry:
 ctkEnabled = "TRUE"
- ☐ For each virtual disk, the .vmx file must contain the following entry:
 scsiox.ctkEnabled = "TRUE"
- ☐ For each virtual disk and snapshot disk, a .ctk file must exist (Example 20-12).

Example 20-12 Both the virtual disk and snapshot disk have an associated .ctk file

```
vmname.vmdk  
vmname-flat.vmdk  
vmname-ctk.vmdk  
vmname-000001.vmdk  
vmname-000001-delta.vmdk  
vmname-000001-ctk.vmdk
```

VMware topology

Check that the topology shown in Figure 20-13 is in place before you use VMware with ProtecTIER and Tivoli Storage Manager.




Figure 20-13 Tivoli Storage Manager/VMware topology

The following conditions should be true:

- ▶ The Tivoli Storage Manager backup-archive client on the vStorage server must read GUEST OS data from Disk Subsystem by SAN.
- ▶ The Tivoli Storage Manager Storage agent on the vStorage server must write GUEST OS data to ProtecTIER by SAN.
- ▶ The Tivoli Storage Manager server writes control data to the ProtecTIER repository through the SAN.

20.7.4 Zoning

The tables in this section describe the required fabric zoning for the ProtecTIER repository, the hosts, and Tivoli Storage Manager:

- ▶ The HBA ports are described in Table 20-3 on page 340.
- ▶ A SAN zoning example is listed in Table 20-4 on page 340.

Table 20-3 HBA ports

Item	Port
ProtectTIER	PT front-end port_0 PT front-end port_1
Tivoli Storage Manager Storage Agent	Tivoli Storage Manager Storage Agent port_0 Tivoli Storage Manager Storage Agent port_1
Tivoli Storage Manager Server	Tivoli Storage Manager server port_0 Tivoli Storage Manager server port_1
ESX/ESi Server	ESX/ESXi server port_0 ESX/ESXi server port_1
XIV	XIV_Module4 port_0 XIV_Module5 port_0 XIV_Module6 port_0 XIV_Module4 port_1 XIV_Module5 port_1 XIV_Module6 port_1

Table 20-4 SAN zoning examples

Zone name	Zone members
Zone_XIV_TSM_StorageAgent_0	Tivoli Storage Manager Storage Agent port_0 XIV_Module4 port_0 XIV_Module5 port_0 XIV_Module6 port_0
Zone_XIV_TSM_StorageAgent_1	Tivoli Storage Manager Storage Agent port_1 XIV_Module4 port_1 XIV_Module5 port_1 XIV_Module6 port_1
Zone_PT_TSM_StorageAgent_0	Tivoli Storage Manager Storage Agent port_0 PT front-end port_0
Zone_PT_TSM_StorageAgent_1	Tivoli Storage Manager Storage Agent port_1 PT front-end port_1
Zone_PT_TSM_Server_0	Tivoli Storage Manager server port_0 PT front-end port_0
Zone_PT_TSM_Server_1	Tivoli Storage Manager server port_1 PT front-end port_1
Zone_ESX_XIV_0	Tivoli Storage Manager Storage Agent port_0 XIV Module4 port_0 XIV Module5 port_0 XIV Module6 port_0
Zone_ESX_XIV_1	Tivoli Storage Manager Storage Agent port_1 XIV Module4 port_1 XIV Module5 port_1 XIV Module6 port_1

20.7.5 Configuring the ProtecTIER server

This section describes all the steps that are required to create and to configure a new VTL in the ProtecTIER server. Also described is the optional procedure to enable and configure LUN masking for the Tivoli Storage Manager server and the Tivoli Storage Manager storage agent (vStorage server in a Tivoli Storage Manager environment).

To create the VTL, complete the following steps:

1. From the **VT** drop-down menu of ProtecTIER Manager, click **VT → VT Library → Create new library**. The Create new library window opens.
2. Input the name of the library in the VT name field, and press **Next**. The Library type window opens in the Create new library window.
3. Select **IBM TS3500** as the library type to simulate. The Tape model window opens.
4. Select **IBM ULT3580-TD3** as the tape model to simulate. The Port Assignment window opens.
5. Create the robot and the drives (such as one robot and 20 drives). The drives are assigned, crossing all front-end ports to ensure better performance. The Cartridges window opens.
6. Create cartridges that are based on the backup policy (for example, 100 cartridges). The Slots window opens.
7. Create 100 slots and 32 import/export slots by selecting **100** in the **No. of slots** selection box, and **32** in the **Number of import/exports** slots selection box. Click **Next**.

Important: The creation of the library takes the system offline for a few minutes.

8. (Optional) Enable and configure LUN masking for the Tivoli Storage Manager server and the Tivoli Storage Manager storage agent (vStorage server). If you have multiple backup servers that are connected to the ProtecTIER server, enabling the LUN masking feature is suggested. Enable and configure LUN masking by completing the following steps:
 - a. From the expanded list on the left side of the ProtecTIER Manager window, click **VT → LUN Masking → Enable/Disable LUN masking**. ProtecTIER Manager notifies you that you if you enable the LUN masking feature without configuring LUN masking groups, the devices are hidden from the hosts, and prompts you to confirm whether you want to proceed with this process.
 - b. When the Enable/Disable LUN masking dialog box opens, select **Enable LUN masking**, and click **OK**.
 - c. From the expanded list on the left side of the ProtecTIER Manager window, click **VT → LUN Masking → Configure LUN masking groups**. The LUN Masking window opens.
 - d. In the “Selected Host Initiators” frame, click **Add**. The Host Initiator Management window opens.
 - e. Create LUN masking groups for the Tivoli Storage Manager server and the Tivoli Storage Manager storage agent by adding the WWPNs of the Tivoli Storage Manager server and the Tivoli Storage Manager storage agent to the list.
 - f. Select the check box beside each of the added ports, and click **Save Changes**. The LUN Masking window opens.
 - g. In the Library Mappings frame, click **Add**, and add the library that is called “TSM_VMW” to the library mappings list.
 - h. Click **Save Changes**.

20.7.6 Installing the tape driver on the Tivoli Storage Manager server and the Tivoli Storage Manager storage agent

You can install the tape driver on the Tivoli Storage Manager and the Tivoli Storage Manager storage agent for Windows and Linux based systems.

Note: Beginning with version 7.1.3, Tivoli Storage Manager was rebranded to IBM Spectrum Control. The scenarios in this topic were conducted with a version earlier than 7.1.3. For legacy and reference purposes, we continue to refer to Tivoli Storage Manager.

The Windows Device Manager shows the tape changer as “Unknown Medium Changer” and the tape drives as “IBM ULTRium III 3580 TAPE DRIVE” (Figure 20-14).

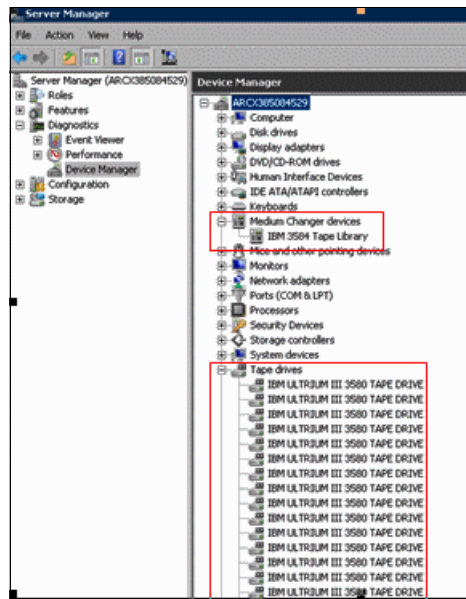
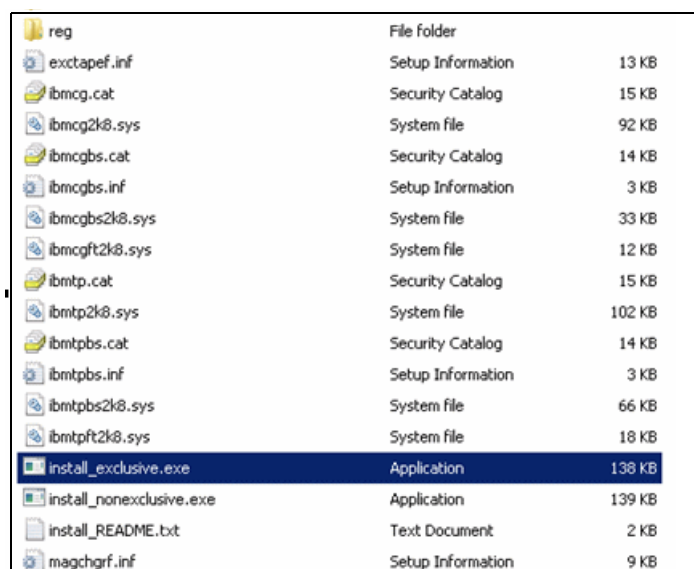


Figure 20-14 Server Manager window

Procedure for Windows systems

Complete the following steps:

1. Download the IBM Tape Device Driver from the following website:
ftp://ftp.software.ibm.com/storage/devdrv/FTPSITE_IS_SUNSET.html
2. Run **install_exclusive.exe** (Figure 20-15) to install the IBM Tape Driver for Tivoli Storage Manager. The installation application initiates, and when complete, displays a dialog box that notifies you that the installation was successful.



reg	File Folder	
exctapef.inf	Setup Information	13 KB
ibmcg.cat	Security Catalog	15 KB
ibmcg2k8.sys	System file	92 KB
ibmcgbs.cat	Security Catalog	14 KB
ibmcgbs.inf	Setup Information	3 KB
ibmcgbs2k8.sys	System file	33 KB
ibmcgft2k8.sys	System file	12 KB
ibmtp.cat	Security Catalog	15 KB
ibmtp2k8.sys	System file	102 KB
ibmtpbs.cat	Security Catalog	14 KB
ibmtpbs.inf	Setup Information	3 KB
ibmtpbs2k8.sys	System file	66 KB
ibmtpft2k8.sys	System file	18 KB
install_exclusive.exe	Application	138 KB
install_nonexclusive.exe	Application	139 KB
install_README.txt	Text Document	2 KB
magchgrf.inf	Setup Information	9 KB

Figure 20-15 Installation program

3. After installation is complete, Windows Device Manager shows the tape changer as “IBM 3584 Tape Library”, and the tape drives as “IBM ULTRIUM III 3580 TAPE DRIVE” (Figure 20-16).

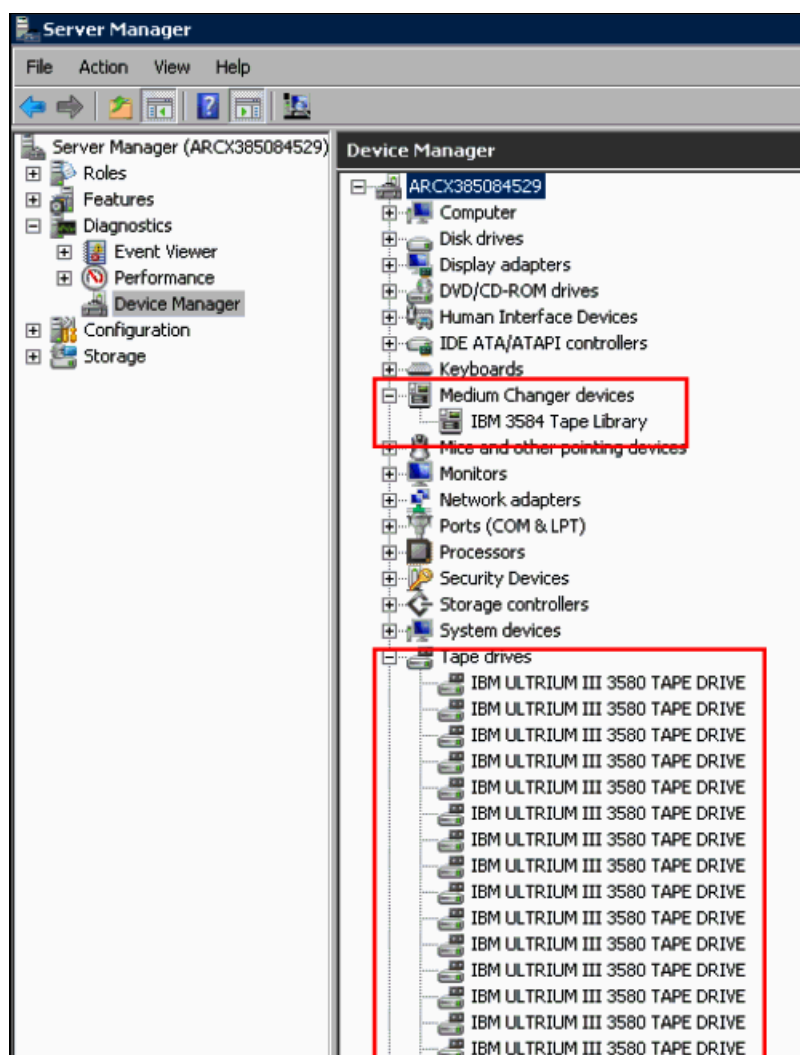


Figure 20-16 Server manager window - showing renamed changer and tape drive

Procedure for Linux systems

Complete the following steps:

1. Download the tape device driver from the following website:
ftp://ftp.software.ibm.com/storage/devdrv/FTPSITE_IS_SUNSET.html
2. Install the following RPM packages:
 - lin_tape-1.54.0-1
 - lin_taped-1.54.0-1

20.7.7 Tivoli Storage Manager storage agent configuration

This section describes the process of configuring the Tivoli Storage Manager storage agent to establish communication with the Tivoli Storage Manager server in a VMware environment.

To accomplish this task, complete the following steps:

1. To establish communication for Tivoli Storage Manager server and Tivoli Storage Manager storage agent, run the commands that are shown in Example 20-13 at the CLI of the Tivoli Storage Manager storage agent.

Example 20-13 Establish communication for Tivoli Storage Manager server and Tivoli Storage Manager storage agent

```
dsmsta.exe setstorageserver  
myname=ARCX385084529  
mypassword=<user_password>  
myhladdress=x.x.110.38  
servername=ARCX3650N1332  
serverpassword=open1sys  
hladdress=x.x.110.65  
lladdress=1500
```

2. Disable automatic mounting of volumes on the Tivoli Storage Manager storage agent host by running the following command at the CLI prompt of the Tivoli Storage Manager storage agent:

```
diskpart > automount disable > exit
```

Important: The **diskpart** command is necessary to keep the Tivoli Storage Manager storage agent from damaging the SAN volumes that are used for raw disk mapping (RDM) virtual disks.

3. To enable the Tivoli Storage Manager storage agent to access online GUEST OS data, click **Server manager** → **Storage** → **Disk Management** → **Online** (Figure 20-17).

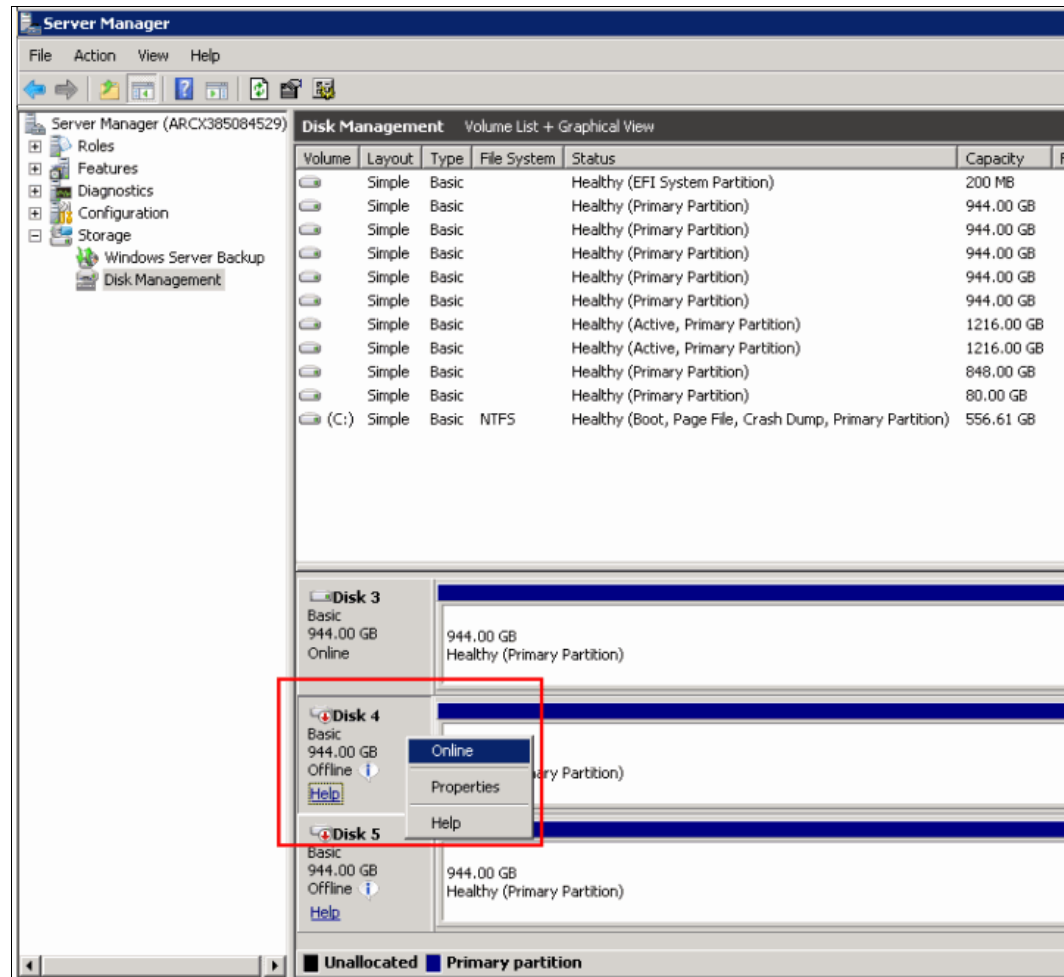
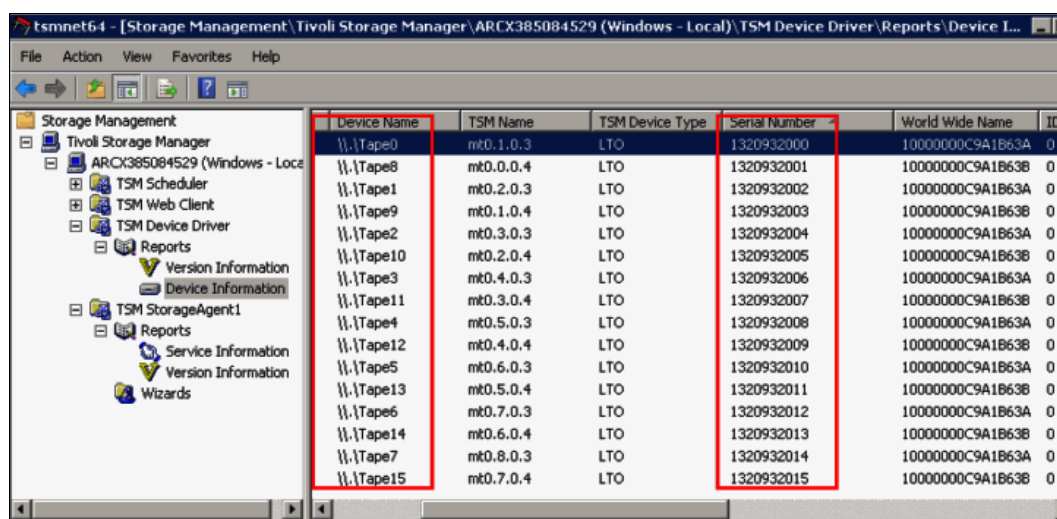


Figure 20-17 Server Manager window shows Tivoli Storage Manager storage agent set to Online

- Note the device name and serial number of the Tivoli Storage Manager storage agent (Figure 20-18). You need this information to define the path in the Tivoli Storage Manager server in a later step.



Device Name	TSM Name	TSM Device Type	Serial Number	World Wide Name	ID
\\.\Tape0	mt0.1.0.3	LTO	1320932000	10000000C9A1B63A	0
\\.\Tape8	mt0.0.0.4	LTO	1320932001	10000000C9A1B63B	0
\\.\Tape1	mt0.2.0.3	LTO	1320932002	10000000C9A1B63A	0
\\.\Tape9	mt0.1.0.4	LTO	1320932003	10000000C9A1B63B	0
\\.\Tape2	mt0.3.0.3	LTO	1320932004	10000000C9A1B63A	0
\\.\Tape10	mt0.2.0.4	LTO	1320932005	10000000C9A1B63B	0
\\.\Tape3	mt0.4.0.3	LTO	1320932006	10000000C9A1B63A	0
\\.\Tape11	mt0.3.0.4	LTO	1320932007	10000000C9A1B63B	0
\\.\Tape4	mt0.5.0.3	LTO	1320932008	10000000C9A1B63A	0
\\.\Tape12	mt0.4.0.4	LTO	1320932009	10000000C9A1B63B	0
\\.\Tape5	mt0.6.0.3	LTO	1320932010	10000000C9A1B63A	0
\\.\Tape13	mt0.5.0.4	LTO	1320932011	10000000C9A1B63B	0
\\.\Tape6	mt0.7.0.3	LTO	1320932012	10000000C9A1B63A	0
\\.\Tape14	mt0.6.0.4	LTO	1320932013	10000000C9A1B63B	0
\\.\Tape7	mt0.8.0.3	LTO	1320932014	10000000C9A1B63A	0
\\.\Tape15	mt0.7.0.4	LTO	1320932015	10000000C9A1B63B	0

Figure 20-18 Serial numbers of Tivoli Storage Manager storage agent

Persistent naming: The example in Figure 20-18 on page 347 does not show the usage of persistent naming. Follow the guidelines in , “When the multipath feature is enabled, it defaults to running in automatic mode. The automatic mode automatically scans for all paths for each library robot at each tldcd daemon start, requiring no additional setup.” on page 85.

If you do not use persistent naming, take a screen capture of the Tivoli Storage Manager management console so that you can have a readily accessible record of the tape device information in the Tivoli Storage Manager storage agent.

20.7.8 Tivoli Storage Manager server configuration

This section describes the procedure to define and to configure the Tivoli Storage Manager server in a VMware environment through the Tivoli Storage Manager server CLI. To accomplish this task, complete the following steps:

- Define the server for the storage agent by running the following command at the Tivoli Storage Manager server CLI:

```
define server ARCX385084529 SERVERPassword=admin HAddress=x.x.110.38
LLAddress=1500 COMMmethod=TCPIP
```
- Set the server name to Tivoli Storage Manager Server by running the following command:

```
set servername ARCX3650N1332
```
- Set the password by running the following command:

```
set serverpassword admin
```
- Set the Tivoli Storage Manager server IP address by running the following command:

```
set serverhladdress x.xx.xxx.xx
```

5. Set the Tivoli Storage Manager server port by running the following command:

```
set server11address 1502
```
 6. Create the library by running the following command:

```
define library VMW_LIB libtype=scsi autolabel=yes shared=yes RELABELSCRatch=yes
```
 7. Choose all devices that are related to tsminst1.
 8. Define a library path from the Tivoli Storage Manager server to the physical OS devices by running the following command:

```
DEFINE PATH ARCX3650N1332 VMW_LIB srctype=server desttype=library  
autodetect=yes device=/dev/IBMchanger0
```
 9. Define all the drives by running the following commands:

```
- define drive VMW_LIB drive0  
- define drive VMW_LIB drive1
```
 10. Define the drives path from Tivoli Storage Manager server to the physical OS devices by running the following commands:

```
- define path ARCX3650N1332 drive0 srctype=server desttype=drive  
library=VMW_LIB autodetect=yes device=/dev/IBMtape0  
- define path ARCX3650N1332 drive1 srctype=server desttype=drive  
library=VMW_LIB autodetect=yes device=/dev/IBMtape1
```
 11. Define the drives path from Tivoli Storage Manager storage agent to the physical OS devices by running the following commands:

```
- define path ARCX385084529 drive0 srctype=server desttype=drive  
library=VMW_LIB  
autodetect=yes device=\\.\Tape0  
- define path ARCX385084529 drive1 srctype=server desttype=drive  
library=VMW_LIB  
autodetect=yes device=\\.\Tape1
```
- Important:** Ensure that the device on the Tivoli Storage Manager server and the device on the Tivoli Storage Manager storage agent, which are mapped to the same drive path, have the same serial number. They are essentially the same device. (For the serial numbers in the Tivoli Storage Manager storage agent, see your notes from step 4 on page 347.)
12. Query the drive and verify that it has a status of online by running the following command:

```
query drive
```
 13. Check in and label the cartridges by running the following command:

```
label LIBVOL VMW_LIB search=yes labelsource=barcode CHECKIN=scratch  
overwrite=yes waitt=0
```
 14. Define a device class by running the following command:

```
define devclass LTOCLASS3 library=VMW_LIB devtype=lto format=ULTRIUM3C
```
 15. Define a storage pool by running the following command:

```
define stgpool VMW_POOL LTOCLASS3 pooltype=primary maxscratch=99999
```


16. Define a domain by running the following command:

```
DEFine D0main VMW_DOMAIN BACKRETention=60 ARCHRETention=365
```

17. Define a policy set by running the following command:

```
DEFine POLicyset VMW_DOMAIN VMW_POLICY
```

18. Define a management class by running the following command:

```
DEFine MGmtclass VMW_DOMAIN VMW_POLICY LTOCLASS3
```

19. Define a copy group by running the following command:

```
DEFine C0pygroup VMW_DOMAIN VMW_POLICY LTOCLASS3 DESTination=VMW_POOL
```

20. Define an archive copy group by running the following command:

```
DEFine C0pygroup VMW_DOMAIN VMW_POLICY LTOCLASS3 Type=Archive  
DESTination=VMW_POOL
```

21. Assign the default management class by running the following command:

```
ASSign DEFMGmtclass VMW_DOMAIN VMW_POLICY LTOCLASS3
```

22. Activate the policy set by running the following command:

```
ACTivate POLicyset VMW_DOMAIN VMW_POLICY
```

23. Register the node for Tivoli Storage Manager BAC by running the following command:

```
register node ARCX385084529 admin passexp=0 userid=admin  
domain=VMW_DOMAINcompression=no type=client DATAWritepath=lanfree  
DATAReadpath=lanfree hladdress=x.x.110.38 lladdress=1502 archdelete=yes  
backdelete=yes maxnummp=999
```

Compression value: Specify the value of compression as no.
For LAN-free backup, specify the value of **DATAWritepath** and **DATAReadpath** as lanfree.

24. Give the appropriate permissions to the administrator by running the following command:

```
GRant AUTHority admin CClasses=SYstem
```

20.7.9 Tivoli Storage Manager client installation

Install the Tivoli Storage Manager client in a VMware environment by using the IBM Tivoli Storage Management installation program:

1. In the wizard (Figure 20-19), select all the components to install, and click **Next**.

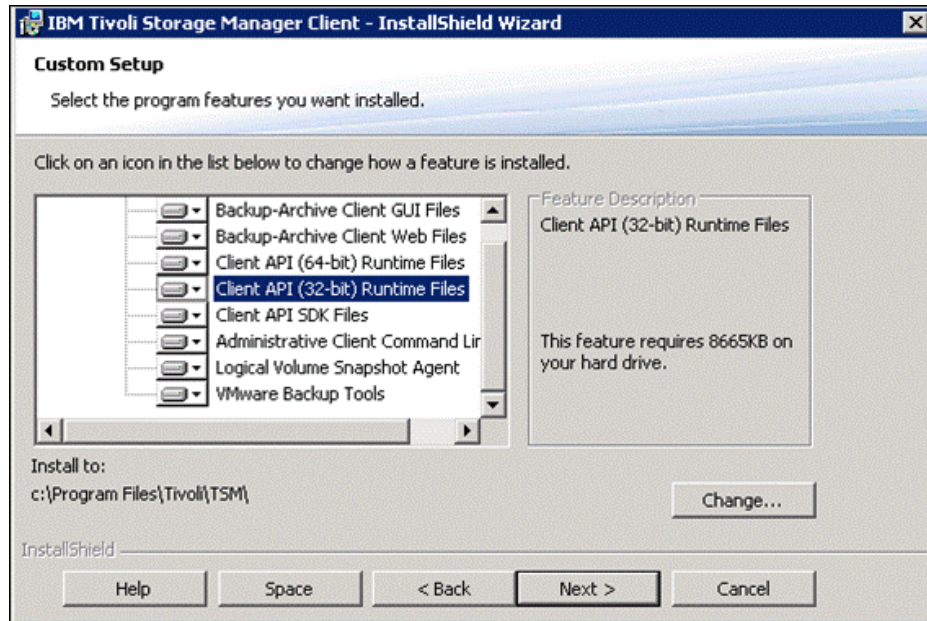


Figure 20-19 Tivoli Storage Manager installation program

Important: Ensure that VMware Backup Tools are installed.

2. When you are prompted to select the type of installation, click **Complete** and finish the installation.

20.7.10 Disabling compression and deduplication on Tivoli Storage Manager

Disable compression and deduplication in Tivoli Storage Manager by using the Tivoli Storage Manager GUI as follows:

1. From the Tivoli Storage Manager GUI drop-down menus, click **Edit** → **Client Preferences** (Figure 20-20). The Client Preferences window opens.

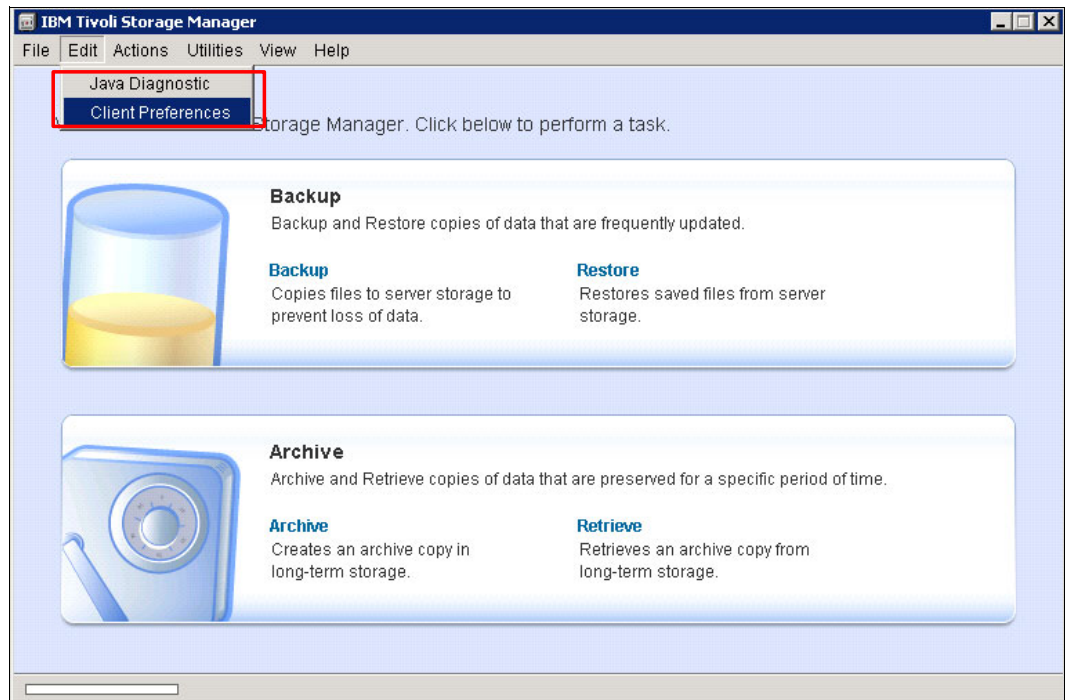


Figure 20-20 Tivoli Storage Manager GUI with Client Preferences selected

2. In the menu at the left, select **Deduplication** (Figure 20-21). The Deduplication Preferences page opens.

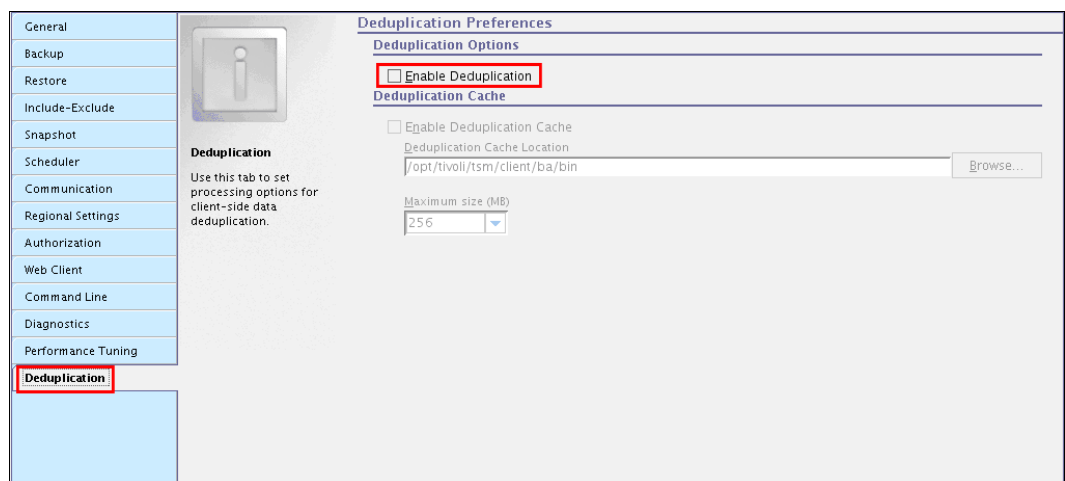


Figure 20-21 Deduplication Preferences window

3. Disable Tivoli Storage Manager deduplication by clearing the **Enable Deduplication** check box.

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

4. In the menu at the left, click **Backup** (Figure 20-22). The Backup Preferences page opens.

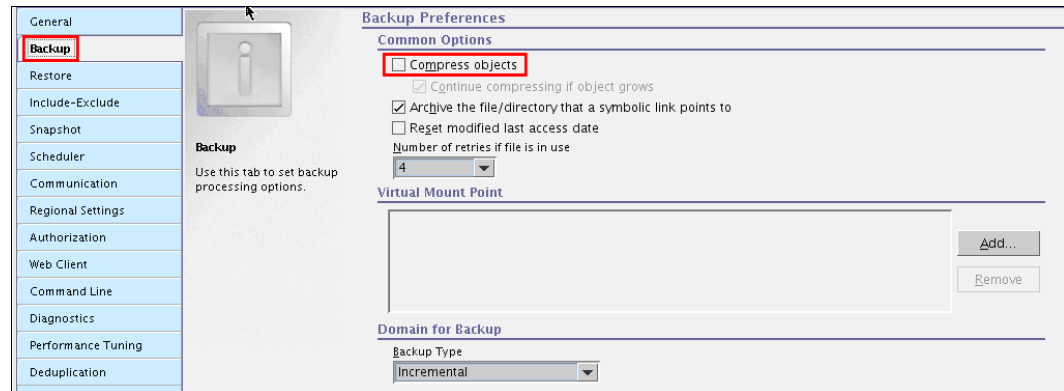


Figure 20-22 Backup Preferences window

5. Disable Tivoli Storage Manager compression by clearing the **Compress Objects** check box.

20.7.11 Configuring a full VM backup through the vStorage API

Configure a full VMware backup from the Tivoli Storage Manager GUI as follows:

1. From the Tivoli Storage Manager GUI drop-down menus, click **Edit** → **Client Preferences**. The Client Preferences window opens.
2. From the menu at the left, click **VM Backup** (Figure 20-23).

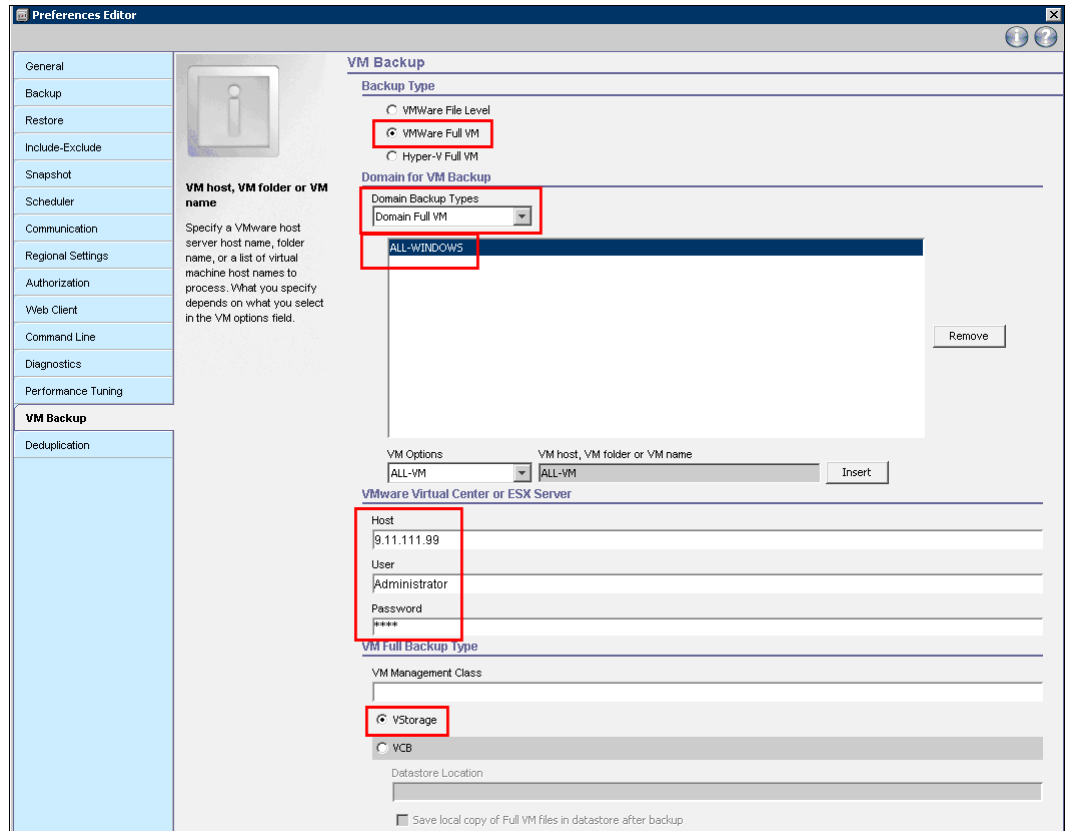


Figure 20-23 VM Backup window

3. Under Backup Type, select **VMware Full VM**.
4. In the Domain for VM Backup selection box, select **Domain Full VM**.
5. In the VM Options selection box, select **All-VM** (If you are using Windows Guest OS, select **ALL-WINDOWS**) and click **Insert**.
6. Enter the IP address, user name, and password of the vCenter.
7. Under VM Management Class, select **VStorage**.

20.7.12 VMware Guest OS backup to ProtecTIER

Use one of these methods to implement VMware Guest OS backup to the ProtecTIER server:

- ▶ Backing up VM by using the Tivoli Storage Manager GUI
- ▶ Full VM backup by using the CLI
- ▶ Incremental VM backup by using the CLI
- ▶ Restoring VM by using the Tivoli Storage Manager GUI

Backing up VM by using the Tivoli Storage Manager GUI

Complete the following steps:

1. From the Tivoli Storage Manager GUI drop-down menus, click **Actions** → **Backup VM** (Figure 20-24).

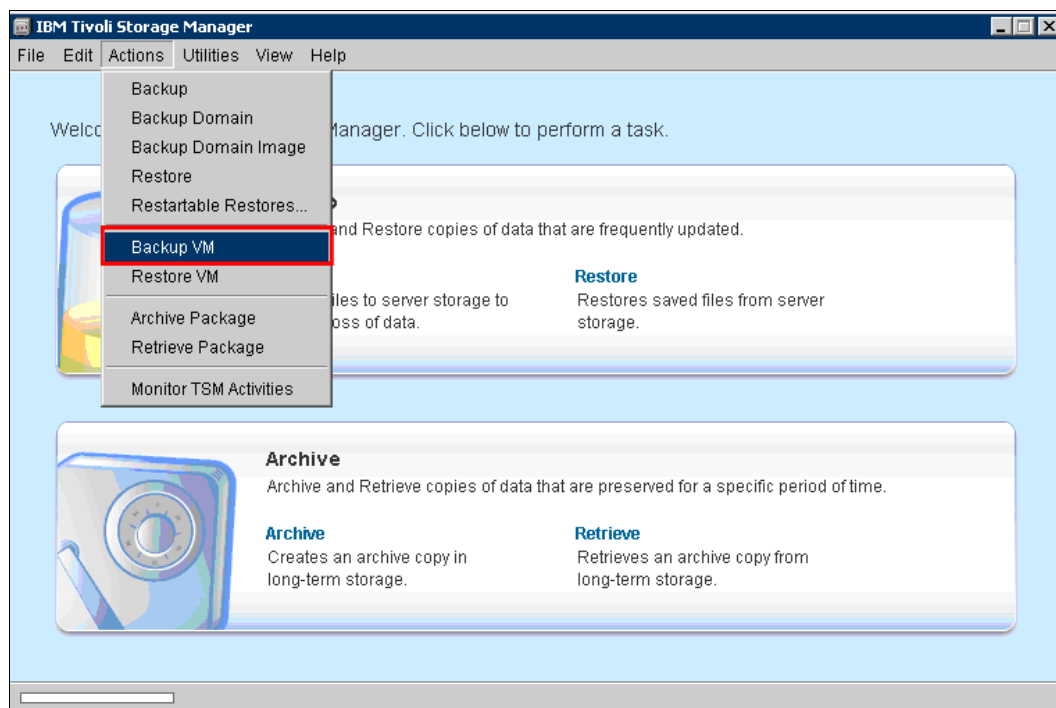


Figure 20-24 Tivoli Storage Manager GUI showing the Backup VM menu cascade

2. The Backup Virtual Machine window opens (Figure 20-25). In the Backup selection box, select the backup type: **VMware Full VM (vStorage)** or **VMware Full VM (incremental)**.

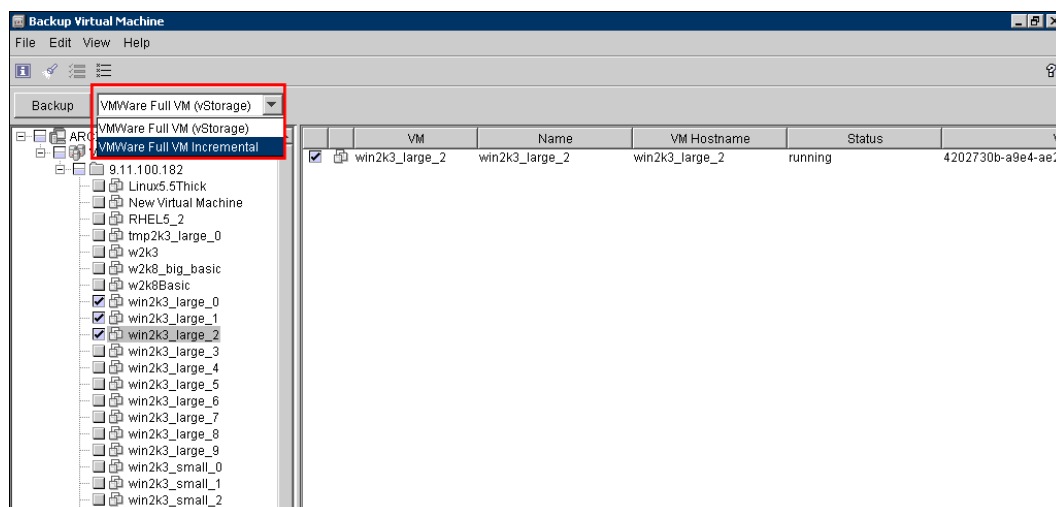


Figure 20-25 Choose the type of backup

Full VM backup by using the CLI

Start a full VM backup; use the **mode=full** backup parameter in the following command:

```
PS C:\Program Files\Tivoli\tsm\baclient> ./dsmc backup vm win2k3_large_3  
-mode=full -vmbackuptype=fullvm
```

Tivoli Storage Manager backs up all of the inspected data, and the value of the total destruction ratio is 0%. The system displays output similar to Example 20-14.

Example 20-14 Output log: full backup

```
Total number of objects inspected: 1  
Total number of objects backed up: 1  
Total number of objects updated: 0  
Total number of objects rebound: 0  
Total number of objects deleted: 0  
Total number of objects expired: 0  
Total number of objects failed: 0  
Total number of subfile objects: 0  
Total number of bytes inspected: 15.00 GB  
Total number of bytes transferred: 15.00 GB  
LanFree data bytes: 15.00 GB  
Data transfer time: 842.86 sec  
Network data transfer rate: 18,661.04 KB/sec  
Aggregate data transfer rate: 12,541.90 KB/sec  
Objects compressed by: 0%  
Total data reduction ratio: 0.00%  
Subfile objects reduced by: 0%  
Elapsed processing time: 00:20:54
```

Incremental VM backup by using the CLI

Start an incremental VM backup; use the **mode=incremental** backup parameter as follows:

```
C:\Program Files\Tivoli\TSM\baclient> ./dsmc backup vm win2k3_large_3  
-mode=incremental -vmbackuptype=fullvm
```

Tivoli Storage Manager backs up only the changed data that is found by VMware CBT, so the value of the data deduction ratio is 99.77%. The output is shown in Example 20-15.

Example 20-15 Output log - full backup

```
Total number of objects inspected: 1  
Total number of objects backed up: 1  
Total number of objects updated: 0  
Total number of objects rebound: 0  
Total number of objects deleted: 0  
Total number of objects expired: 0  
Total number of objects failed: 0  
Total number of subfile objects: 0  
Total number of bytes inspected: 15.00 GB  
Total number of bytes transferred: 38.63 GB  
LanFree data bytes: 38.63 GB  
Data transfer time: 2.77 sec  
Network data transfer rate: 13,461.05 KB/sec  
Aggregate data transfer rate: 2,689.30 KB/sec  
Objects compressed by: 0%  
Total data reduction ratio: 99.77%  
Subfile objects reduced by: 0%  
Elapsed processing time: 00:00:13
```

Restoring VM by using the Tivoli Storage Manager GUI

Complete the following steps:

1. In the Tivoli Storage Manager GUI, click **Actions** → **Restore VM** (Figure 20-26).

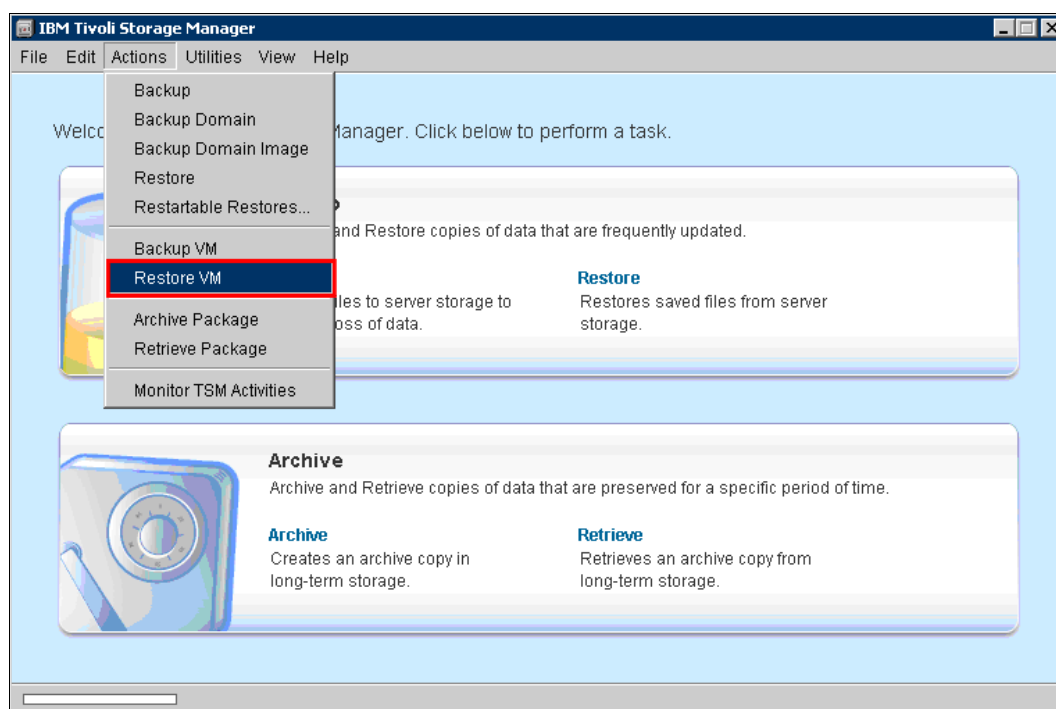


Figure 20-26 Tivoli Storage Manager GUI with Restore VM menu cascade

2. The Restore Virtual Machine window opens (Figure 20-27). Select the version to restored, either full backup (**FULL**) or incremental backup (**INCR**).

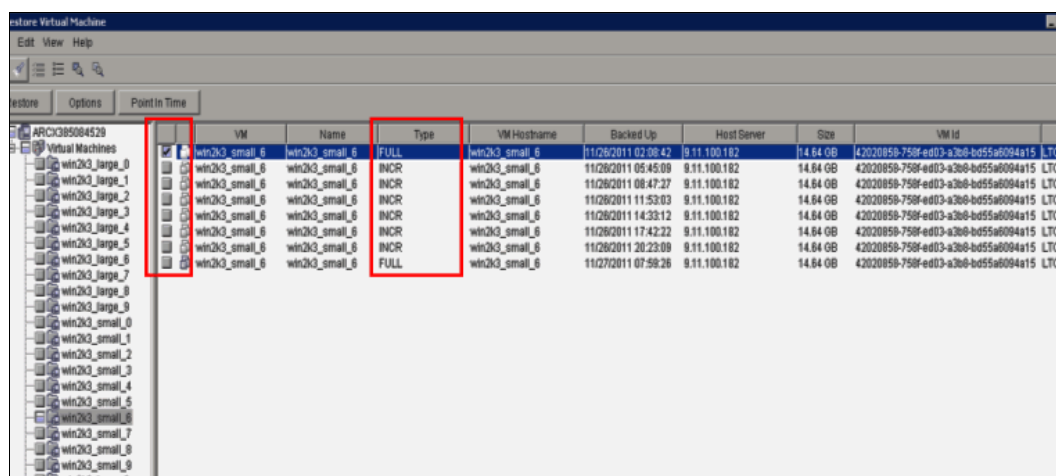


Figure 20-27 Restore Virtual Machine window

3. Tivoli Storage Manager prompts you to select whether to restore to the original location or to a new location (Figure 20-28). If you choose to restore to a new location, enter the following information, and click **Restore**:

- Name: The Guest OS name that is managed by vCenter.
- Datacenter: The name of the data center that stores the new Guest OS.
- Host: The IP of the ESX server that stores the new Guest OS.
- Datastore: The name of data store that stores the new Guest OS.

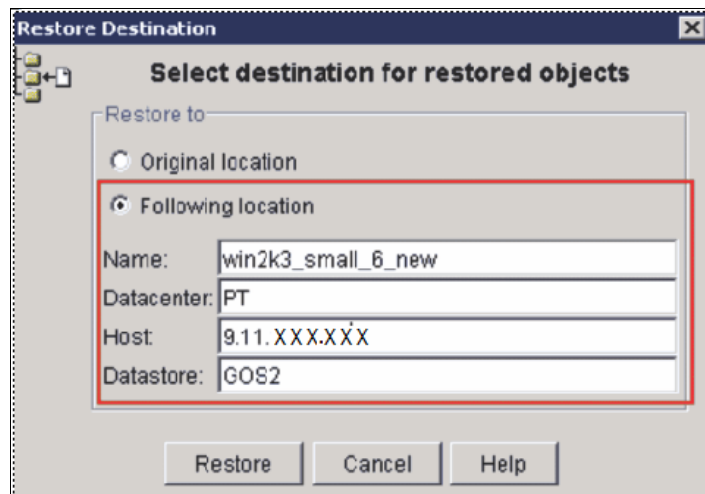


Figure 20-28 Restore Destination window

Replication and disaster recovery

ProtecTIER with replication enables virtual tape cartridges to be replicated from multiple primary sites (spokes) to a central secondary location (hub) for enhanced disaster recovery (DR) and business continuity (BC) capabilities. This part describes replication. Also included are concepts such as replication deployment, solution optimization, and the procedures to deploy replication to work with specific backup applications.

This part contains the following chapters:

- ▶ Chapter 21, “ProtecTIER replication” on page 361
- ▶ Chapter 22, “Disaster recovery deployment with backup applications” on page 395



ProtecTIER replication

Information technology (IT) organizations that use an IBM ProtecTIER system with replication can easily expand the coverage of that replication to all of the applications in their environment. You can create replication policies to set rules for replicating data objects across ProtecTIER repositories. This chapter describes the purpose of replication and the enhanced features of the latest ProtecTIER code.

This chapter describes the procedures for replication deployment, including preinstallation steps, creation of the replication grid, and synchronization of the primary and secondary repositories. It also describes upgrading the existing system and enabling replication.

In addition, this chapter provides the principal rules and guidance of replication deployment for the ProtecTIER product in environments with Virtual Tape Library (VTL), and File System Interface (FSI). It also describes the concepts, procedures, and considerations that are related to optimizing replication performance, including the procedures to automate and script the daily operations.

This chapter primarily focuses on preferred practices for planning, configuration, operation, and testing of ProtecTIER native replication. The concept, detailed planning, and implementation of native replication is described in *IBM TS7650 with ProtecTIER V3.4 User's Guide for VTL Systems*, GA32-0922. This chapter describes the following topics:

- ▶ ProtecTIER IP replication
- ▶ Native replication
- ▶ Replication policies
- ▶ Visibility switching
- ▶ Principality
- ▶ Replication Manager
- ▶ Initial synchronization
- ▶ Replication schedules
- ▶ Replication backlog
- ▶ Replication planning
- ▶ Bandwidth validation utility
- ▶ Planning ProtecTIER replication
- ▶ The backup application database backup

21.1 ProtecTIER IP replication

The ProtecTIER IP replication function (Figure 21-1) provides a powerful tool that you can use to design robust disaster recovery architectures. You electronically place backup data into vaults with much less network bandwidth, therefore changing the paradigm of how data is taken off-site for safe keeping. The ProtecTIER IP replication feature can eliminate some of the expensive and labor-intensive handling, transport, and securing of the real tapes for disaster recovery purposes.

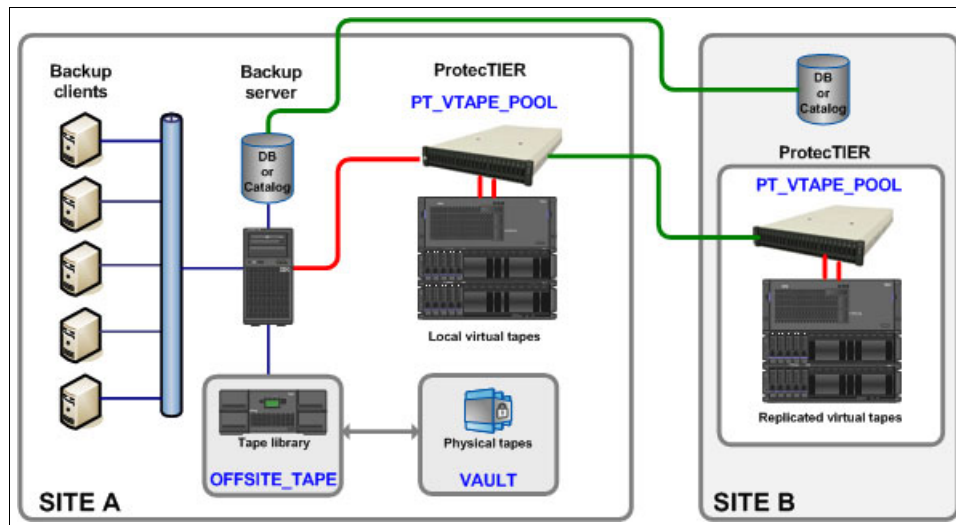


Figure 21-1 IP replication in a backup and recovery environment

Figure 21-1 illustrates how the ProtecTIER IP replication function can be used in a backup and recovery environment. This particular client uses this feature to replicate all of the virtual tapes in onsite PT_VTAPE_POOL off-site. It also backs up all backup application databases or catalogs to virtual tapes. These database backup virtual tapes are also replicated to Site B.

If a disaster occurs, this client can restore the backup server environment on Site B, which is connected to a ProtecTIER VTL. It contains the backup application database (catalog) together with all of the client backup files on virtual tapes in the PT_VTAPE_POOL pool.

21.2 Native replication

ProtectTIER native replication enables data replication capability across repositories, among ProtecTIER systems, which are connected to the wide area network (WAN). Because the ProtecTIER product deduplicates data before storing it, only the changes, or unique elements of data, are transferred to the DR site over the replication link. This feature can translate into substantial savings in the bandwidth that is needed for the replication TCP/IP link.

In early versions of ProtecTIER, the repository replication was handled by the disk array systems. Starting with Version 2.4, ProtecTIER introduced the function that is known as *native replication*, where the replication of deduplicated data became a function of ProtecTIER. Deduplicated data is replicated to a secondary ProtecTIER system through TCP/IP rather than relying on the back-end disk arrays and their associated infrastructure.

21.2.1 One-to-one replication

The initial replication design of ProtecTIER consisted of two ProtecTIER systems, with one system that is designated as the source and the other system that is designated as the target. The target system (or hub) was dedicated to receive incoming replicated data and was not eligible to take local backups.

21.2.2 Many-to-one replication

ProtecTIER Version 2.4 expanded the native replication functionality and introduced the many-to-one replication grid. Also known as *spoke and hub*, up to 12 source systems (spokes) can all replicate to a single target ProtecTIER system (hub) simultaneously. The hub system can provide DR functionality for one or *more* spokes concurrently, and the hub system can accept and deduplicate local backup data. The hub system cannot replicate outgoing data.

21.2.3 Many-to-many replication

ProtecTIER Version 3.1 built upon existing replication technology and introduced the many-to-many bidirectional replication grid. Up to four systems (all hubs) can be joined into a bidirectional replication grid and can accept and deduplicated local backup data, replicate data to up to three other ProtecTIER systems in the grid, and receive incoming replicated data from up to three other ProtecTIER systems in the grid.

21.2.4 VTL replication

VTL replication can be set up as one-to-one, many-to-one, and many-to-many. The ProtecTIER VTL service emulates traditional tape libraries. Your existing backup application can access virtual robots to move virtual cartridges between virtual slots and drives. The backup application perceives that the data is being stored on cartridges, although the ProtecTIER product stores data on a deduplicated disk repository. In a VTL replication scenario, data is replicated at the virtual tape cartridge level.

21.2.5 FSI replication

With FSI replication, up to eight ProtecTIER FSI systems can be included in the bidirectional replication group (Figure 21-2). Each FSI system can replicate deduplicated data to as many as three other remote ProtecTIER FSI systems. Data is replicated at the file system level.




Figure 21-2 FSI replication topology group

21.2.6 Replication grid

A replication grid is a logical set of repositories that can replicate from one repository to other repositories. A ProtecTIER system must be a member of a replication grid before the system is able to create replication policies.

All ProtecTIER systems are capable of replication. Different models of ProtecTIER systems, such as the TS7650G Gateway and TS7620 Appliance Express models, can be part of the same grid. You can have more than one replication topology group in the same grid. A grid can also contain various types of replication groups, such as groups of VTL and FSI. A single replication grid can include up to 24 ProtecTIER systems.

Important: A ProtecTIER system can be a member of only one replication grid in its lifetime. After a ProtecTIER system joins a grid, it is no longer eligible to join any other ProtecTIER replication grid.

21.2.7 Replication topology group

A replication topology group defines the relationship between ProtecTIER systems in a replication grid. A group includes many-to-one, many-to-many, and FSI groups. A replication grid can have multiple topology groups of various types, as shown in Figure 21-3.

Note: A ProtecTIER system can be a member of only one topology group at a time. A ProtecTIER system can move from one topology group to another in the same grid.

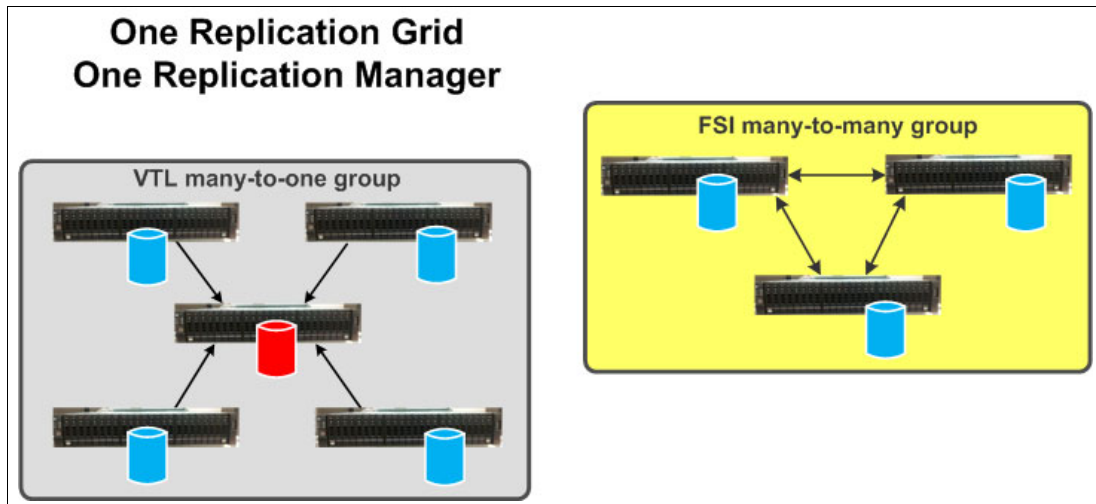


Figure 21-3 Multiple topology groups in a single replication grid.

21.3 Replication policies

The rules for replicating ProtecTIER data objects (VTL cartridges, and FSI file systems) are defined in *replication policies*. Replication policies for VTL and FSI data objects are defined on the ProtecTIER system.

When the backup application writes to a ProtecTIER data object (VTL cartridge or FSI file system) that is part of a replication policy, the ProtecTIER software conducts a check on the object and places it in the replication queue. For VTL cartridges the check includes determining the replication priority, for FSI file systems the replication policies do not have a priority option.

Data objects created in the primary site repository are read/write enabled so that the backup application at the primary site has full control of them and their content. Data objects replicated to the DR site are set in a read-only mode.

In VTL replication, only one cartridge *instance* can be in a library at the same time; all replicas are on the virtual shelf in the disaster recovery site repository.

Tip: At any time, you can override the default location of any VTL cartridge and manually move the replica from the virtual shelf to a library in the repository of the disaster recovery site.

Before replication occurs, the *dirty bit* technology system determines which data objects need to be replicated. Both the local and secondary sites hold synchronized data for each of their data objects. The destination site then references this synchronized data to determine which data (if any) should be transferred. The replication mechanism has two types of data to transfer:

Metadata Data that describes the actual data and carries all the information about it.

User data The actual backed up data.

A data object is marked as *synced* after its data finishes replicating from the primary to the secondary site. So, at the time of synchronization, the local objects and their remote replicas are identical. Before replication starts running, the system ensures that only unique new data is transferred over the TCP/IP link.

Attention: If you delete a data object in the source repository, then all the replicas are also deleted in the target repositories.

Network failure: If a network failure occurs during replication, the system continues to try, for up to seven consecutive days, to complete the replication tasks. After seven days, a replication error is logged and the replication tasks are no longer retried automatically by the system.

21.4 Visibility switching

Visibility switching is the automated process that transfers the visibility of a VTL cartridge from its master to its replica and vice versa. The visibility switching process is triggered by moving a cartridge (that is part of a replication policy) to the source library import/export slot. The cartridge then disappears from the import/export slot and appears at the destination library import/export slot, and at the source the cartridge is moved from the import/export slot to the shelf.

To move the cartridge back to the source library, the cartridge must be ejected to the shelf at the destination library. The cartridge then disappears from the destination library and reappears at the source import/export slot.

Important: By design, the source sends visibility changes to the destination when the replication trigger is completed. Consequently, there might be more replication triggers in queue for the same cartridge, and they will be run later. This behavior is correct by design.

Trying to make any tests on replicated carts, without closing replication backlog first, can cause failures.

21.5 Principality

Principality is the privilege to write to a cartridge (set it to read/write mode). The principality of each cartridge belongs to only one repository in the grid. By default, the principality belongs to the repository where the cartridge was created. The cartridge metadata information includes the principality repository ID field.

Principality can be changed from one repository to another, by which the writing permissions are transferred from one repository to another. Principality can be changed in the following cases:

- ▶ During normal replication operation, from the repository where the cartridge was created to the DR repository, if the cartridge is 100% synced.
- ▶ During failback process, if the principality belongs to one of the following repositories:
 - The DR repository
 - The original primary repository, and this site is the destination for the failback.
 - The original primary repository with the following exceptions:
 - The original primary repository is out of the replication grid.
 - The target for the failback is a repository that is defined as a replacement repository through the ProtecTIER repository replacement procedure.

21.6 Replication Manager

The ProtecTIER *Replication Manager*, also known as Grid Manager, is the part of the software that is used to remotely manage the replication configuration. From the Replication Manager, you can build and maintain the replication infrastructure and repository relationships.

The preferred practice is to designate the DR site system as the Replication Manager.

A possibility is to have a ProtecTIER Replication Manager (Grid Manager) installed in a node that is not one of the systems that the Grid Manager is managing. To have a dedicated host as a Grid Manager requires a request for product quotation (RPQ), which must be requested by your IBM marketing representative.

The ProtecTIER Replication Manager that is installed on a ProtecTIER node can manage only one grid with up to 24 repositories. If a dedicated server is chosen, *and approved by the RPQ process*, it can manage up to 64 grids with 256 repositories in each grid.

You must activate the Grid Manager function before you can add it to the list of known Grid Managers on the ProtecTIER Manager GUI, as shown in Figure 21-4 on page 368.

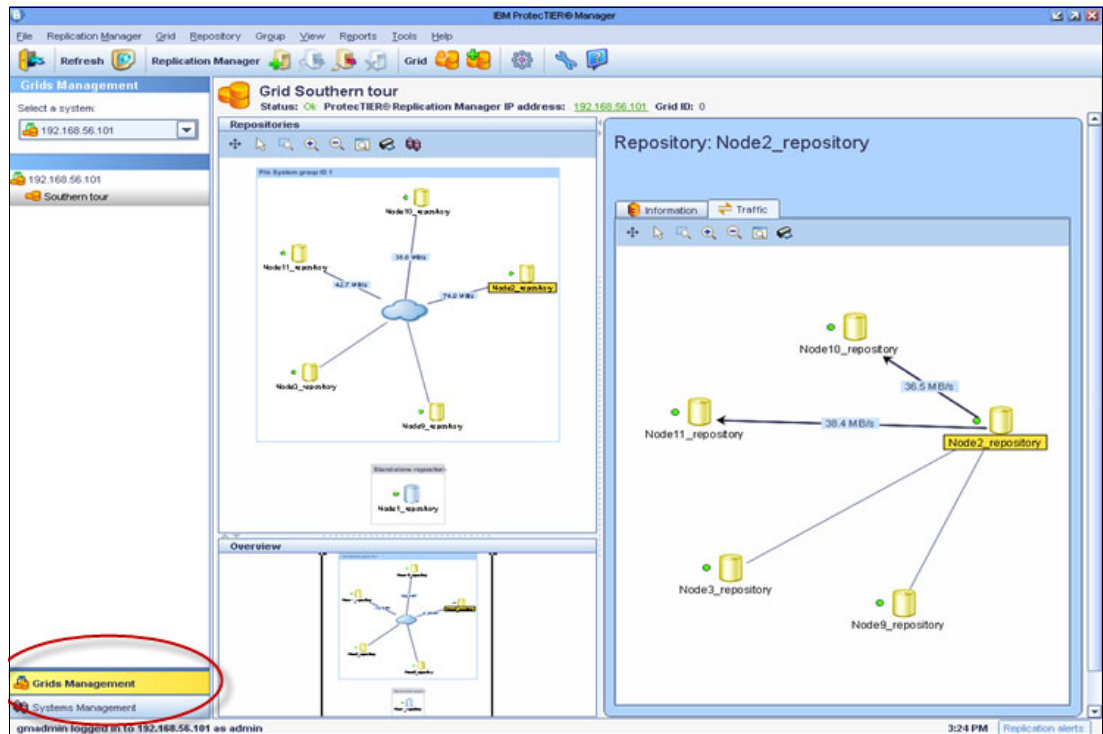


Figure 21-4 Designate ProtecTIER system as a ProtecTIER Replication Manager

To designate a ProtecTIER system as a Grid Manager, use the **menu** command (Figure 21-5):
 ProtecTIER Configuration > Configure replication

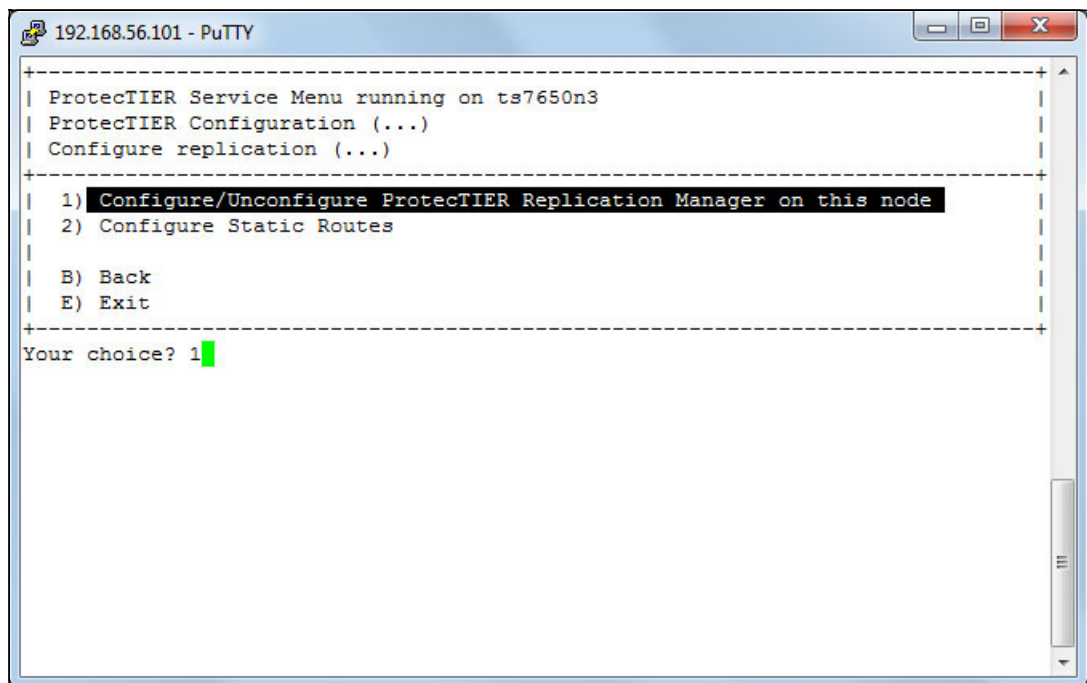


Figure 21-5 Enable Replication Manager function

21.7 Initial synchronization

When a new ProtecTIER system is configured as a replication target (secondary) for an already existing ProtecTIER system (primary), it is necessary to synchronize the primary system with the secondary system.

A deployment of a second ProtecTIER server at a secondary (DR) site has an effect on the planning cycle because the first replication jobs use more bandwidth than required after deduplication takes effect. So, when you prepare for replication deployment, bandwidth is an important consideration.

During the planning cycle, the planners and engineers must consider the amount of nominal data that will be replicated, and the amount of dedicated bandwidth. It might be necessary to implement a replication policy that enables the first replication job to complete before the next backup activity begins.

Note: For the initial replication, you must arrange enough network bandwidth to account for the full nominal size of the data to be replicated.

Two methods can be used; both methods focus on gradually adding workload to the replication policies over time.

Gradual management of policies over time

This is the preferred method, whether you are deploying a new system or adding replication to an existing system. In this method, you add new replication policies over time, and manually ensure that the total daily volume of replicated data remains in the bandwidth limit.

Replication policies with a gradual increase are preferred to stay in the available network bandwidth boundaries and in the time frame that is scheduled for replication activity.

Priming the DR repository at a common locality with the primary system

Priming the DR system at a primary site first and then moving it to its DR location has limited practical value, and is not the preferred choice. In a multisite deployment, this method is a poor choice:

- ▶ If you take this approach, you must manage the synchronization process again when the systems are placed in to their final location.
- ▶ If you are synchronizing a full, partial, or even a newly started repository, the system must have sufficient network bandwidth for primary and secondary systems to synchronize in the available time frame.

21.8 Replication schedules

The ProtecTIER product offers two modes of operation for the replication activity:

- ▶ *Scheduled replication* occurs during a predefined time frame.
- ▶ *Continuous replication* runs constantly.

The mode of operation can be set on both sending (spoke) and receiving (hub) ProtecTIER systems. All defined replication policies operate in one of these modes. In most cases, scheduled replication is the best approach. It enables administrators to accurately plan for performance, and to better ensure that service level agreements (SLAs) are met. The replication mode of operation is a system-wide option and it affects all polices in the system.

By default, data objects are continuously being replicated from the primary (local) site to the repository at the DR site. Optionally, a replication schedule can be defined to limit replication activity to specific time slots during the week.

21.8.1 Continuous replication

Continuous replication can run concurrently with the backup operation. Typically, it requires a larger system to enable concurrent operations at a high throughput rate. This option can affect backup performance because the “read” function is shared between the deduplication processes and the replication operation. Consider the following aspects when you plan continuous replication:

- ▶ Data automatically starts replicating to a DR site repository soon after it is written to the primary ProtecTIER system.
- ▶ Replication runs faster (up to 100% of available performance) if the primary system is not performing backup or restore activity.
- ▶ If it is running concurrently, replication is prioritized lower than backup or restore in the ProtecTIER system.

Continuous replication is available or suggested in the following situations:

- ▶ A system has consistently low bandwidth.
- ▶ The operation calls for few backup windows that are spread throughout the day.
- ▶ Deploying a multisite scenario, especially across multiple time zones.

21.8.2 Scheduled replication

The scheduled replication occurs during a predefined time frame, which is the suggested mode for most applications. This mode imitates the procedure that is used with physical tapes that are being transported to a DR site after backup is completed. This method enables users to keep complete sets of backup data together with a matching backup application catalog or database for every 24 hour period.

With this approach, consider the following information:

- ▶ Backups are allowed to finish without performance effect from replication.
- ▶ The user defines the start and end of the replication time frame.
- ▶ Replication activity begins at the predefined time.
- ▶ Replication stops at the end of the time window specified; each cartridge in transit stops at a consistent point at the end of the time window.
- ▶ Replication does not occur outside of the dedicated time window.

During a scheduled replication, replication activity has the same priority as backup and restore activity. If backup and restore activity takes place during the same time frame, they are equally weighted and processed in a first-in-first-out manner. Because the overall system throughput (backup and restore plus replication) can reach the maximum configured rate, the backup duration might vary.

Tip: Because backup, restore, and replication jobs access the same back-end disk repository, contention between these processes can slow them down. You must plan for and configure the ProtecTIER system resources to accommodate these types of activities to finish their tasks in the wanted time frames.

However, the ProtecTIER system remains available to the backup application throughout the time frame that is dedicated to replication. So if a backup or restore operation is necessary during the replication time frame, the operation can be performed.

A primary benefit of scheduled replication is the ability to strictly classify when the ProtecTIER server uses the network infrastructure, and accurately isolate the usage of the network. This mode of operation is aligned with the backup and DR activity, where users manage a specific backup time frame and schedule replicating jobs that follow the backup.

21.8.3 Centralized Replication Schedule Management

Each ProtecTIER system has the optional ability to schedule both incoming and outgoing replication activity by using a weekly schedule that is divided into one-half hour time slots. Only one schedule exists for a ProtecTIER system that governs all replication policies on that system. Figure 21-6 shows the Set Replication Timeframe window.

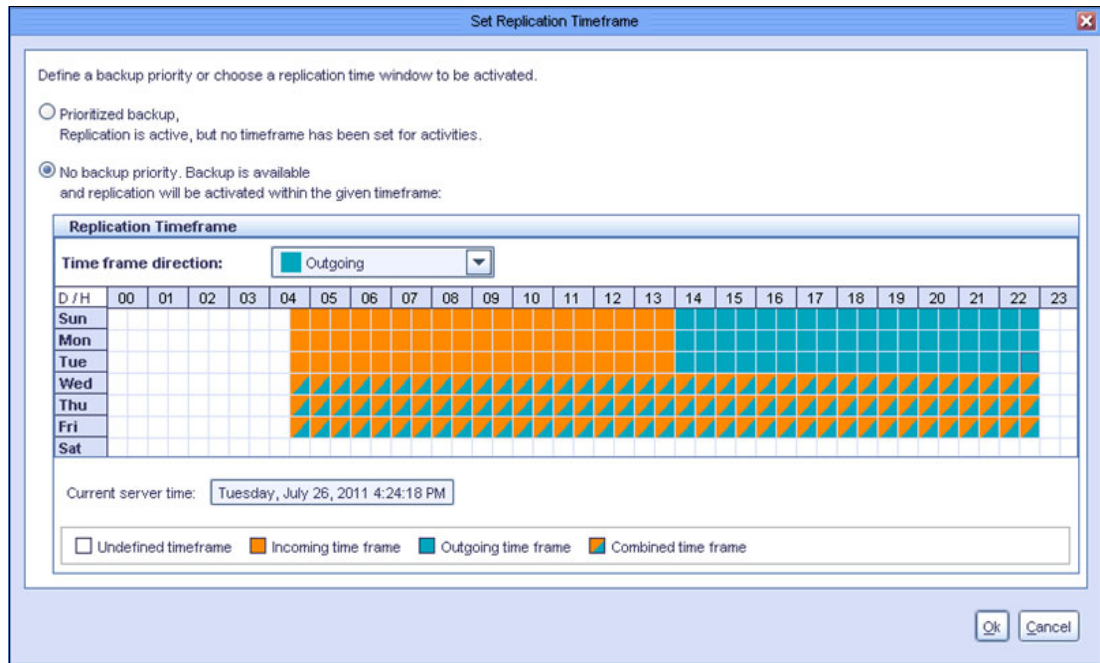


Figure 21-6 Replication schedule

Schedules can be set on both sending (spoke) and receiving (hub) ProtecTIER systems.

Important: Avoid time window conflicts when you define time frames at the hub and at the spokes:

- ▶ No synchronization mechanism exists to foresee misalignments, so if you set the hub and spokes to different time slots, replication never runs.
- ▶ Ensure that the hub has enough time frame slots to accommodate all of the spokes' combined time frames.

Starting with Version 3.1, ProtecTIER introduced the Centralized Replication Schedule Management function. Using this function, you can view and set replication schedules for all the nodes in a grid and visually check time frame alignment between nodes, as shown in Figure 21-7.

Note: Centralized Schedule Management is available in the Grid Management view of the ProtecTIER Manager GUI

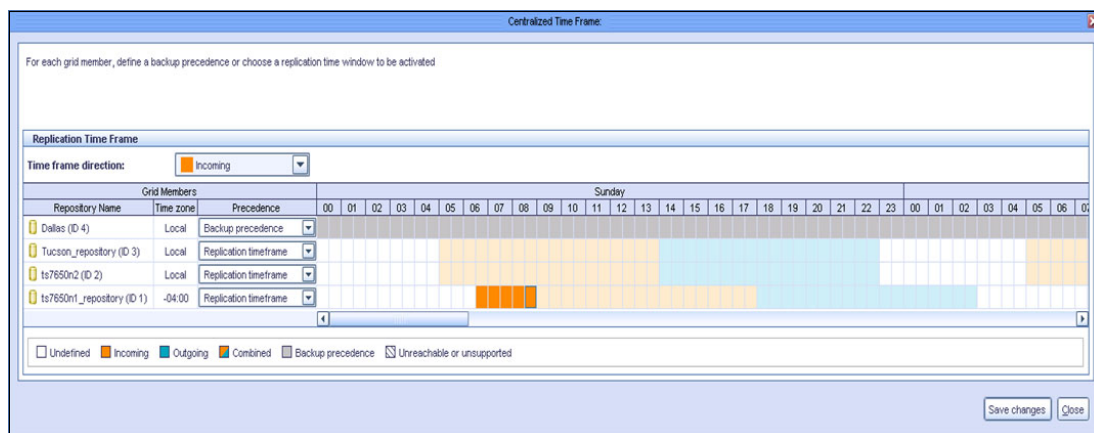


Figure 21-7 Centralized Replication Schedule Management

21.8.4 Replication rate control

There are enhanced system replication throttling and dynamic *system resource allocation* functions for incoming and outgoing replication. ProtecTIER replication offers the following enhanced features and benefits:

- ▶ Setting replication performance limits: It can be either a nominal or a physical limit. The nominal performance limit reflects the overall resource consumption of the system. The physical performance limit reflects the network transfer rate of the replication network.
- ▶ Enhancements to the replication rate control mechanism: Currently, the Replication Rate Control (RRC) is used when a user does not provide a time frame and the system replicates continuously. The *rate calculation* determines the maximum rate that is possible in both levels of system usage (IDLE and BUSY), and normalizes the rate.

- A GUI feature that provides an at-a-glance view of the proportion of the repository data, replication data, local backup data, and free space, as shown in Figure 21-8.




Figure 21-8 Repository usage by category

The nominal and physical throughput (data flow rate) can be limited by setting the RRC. The following information must be considered:

- Both the nominal and physical amounts of data that are being processed or transferred.
- The ability to send and receive new unique data between spokes and the hub.
- ProtecTIER validates all the new or updated objects at the target repository before it makes them available for the user. Setting the replication rate control enables the user to limit the nominal and physical throughput (data flow rate of replication). This feature can be used on spokes and on the hub for both sending and receiving. The values that are set in the physical throughput might, but do not necessarily, affect those values that are set in the nominal throughput, and vice versa. However, *when you use both methods, the physical settings override the nominal ones.*

Setting a nominal limit

When you set a nominal limit, you define the maximum ProtecTIER server system resources that can be used to process the replication data. The nominal throughput directly affects the replication data flow and the load on both the source and destination repositories.

On a ProtecTIER system that performs both backup and replication, setting a nominal replication limit enables you to select a replication throughput limit such that it does not compete with the backup operation for system resources. Setting the limit on a source repository ensures that the backup operation realizes the total possible throughput minus the nominal limit set.

For example, on a node with a performance capacity of 500 MBps that performs backup and replication concurrently, you might set the following limits:

- ▶ 300 MBps when replication is running on its own
- ▶ 100 MBps when replication is running concurrently with a backup

Setting a physical limit

When you set a physical limit, you limit replication network bandwidth consumption by the ProtecTIER server. This limit is intended to be used when the network is shared between the ProtecTIER server and other applications so that all applications can run concurrently. The physical throughput limit restrains the amount of resources that the replication processes can use. This limit reduces the total load on the replication networks that are used by the repository.

Although this limit can be set at either the spoke, the hub, or both, it is typically set at the spoke. Setting a limit at the hub results in *de facto* limitations on all spokes.

21.8.5 Setting replication rate limits

You can limit the replication rates and throughput (Figure 21-9) in the following situations:

- ▶ During a backup or restore operation
- ▶ When there is no backup or restore activity
- ▶ In a defined replication time frame

Replication Rate Limits

Select and define the maximum physical and nominal throughput (in MB/sec) options to set the replication rate limits. Leave the checkboxes empty for unlimited replication rates. The Default system settings values are automatically adjusted when the peak throughput is increased. With selecting the default system settings, the values will be auto adjusted when the peak throughput will be increased.

Choose the incoming and outgoing replication rate limits mode:

☐ Default system settings
☐ Combined settings
☒ Individual settings

Physical throughput (MB/Sec)

☐ Physical throughput limits

Incoming	Outgoing
Unlimited	Unlimited

Nominal throughput (MB/Sec)

☐ Limit when no backup or restore load

Incoming	Outgoing
Unlimited	Unlimited

☐ Limit during backup or restore load

Incoming	Outgoing
Unlimited	Unlimited

☐ Limit within a replication timeframe

Incoming	Outgoing
Unlimited	Unlimited

Ok Cancel

Figure 21-9 Setting replication rate limits

21.8.6 Limiting port bandwidth consumption

Bandwidth throttling (physical limit) controls the speed at which replication operates, where the user can specify a maximum limit for the network usage. By default, there is no configured bandwidth limit; the ProtecTIER server uses as much bandwidth as it can.

If the physical network layer consists of dark fiber or other high-speed network infrastructure, there is typically no reason to limit replication throughput. However, if the ProtecTIER server is running over a smaller network pipe that is shared by other applications, you can restrict the maximum physical throughput that is used by ProtecTIER replication.

This parameter is adjustable per Ethernet replication port on all nodes in the replication grid. It applies only to outgoing data. Set it at the source (sending) system. If the source system is composed of a dual-node cluster, be sure to consider all ports of both nodes when setting the limit.

For example, to hold ProtecTIER replication to a limit of 100 MBps, set each of the four available Ethernet replication ports to 25 MBps, as shown on Figure 21-10. 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 capacity. By default, the setting per port is *Unlimited*.

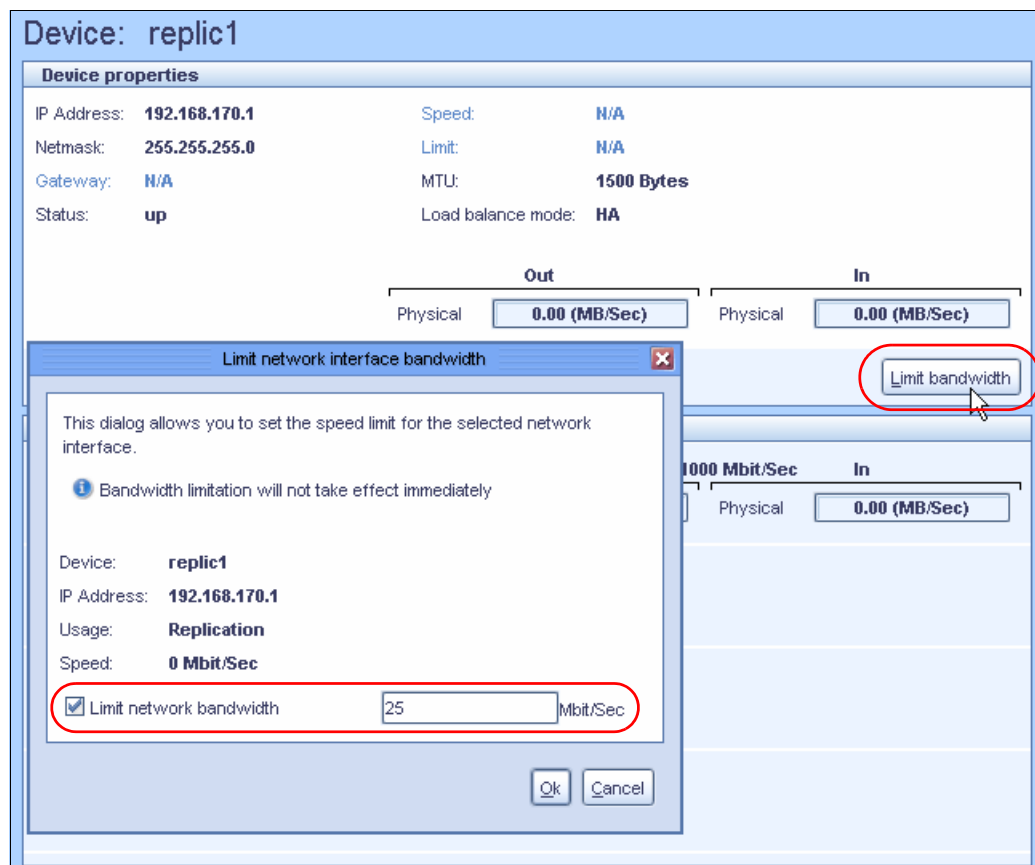


Figure 21-10 Potential modification of the replication interface limit

Changing the bandwidth: If the bandwidth limitation is changed during replication, the change does not take effect immediately. If replication begins after the bandwidth limitation change, the effect is immediate.

21.9 Replication backlog

When replication activity is started, the source system builds a list of new and changed data blocks and sends that list to the receiving system. The receiving system checks the list and determines which data blocks it must synchronize with the source system and then sends requests for the transferal of data blocks. Now, there is a *backlog* of data to replicate. The source system monitors and displays the amount of backlog replication data in the ProtecTIER Manager GUI Activities view.

Having a backlog of data to replicate is not inherently a problem. A potential problem is indicated when the amount of backlog replication data does not go down over time.

If there is an unscheduled long network or DR site outage, the replication backlog might become too large for the system to catch up. A prolonged replication backlog might be an indication of insufficient available bandwidth that is allocated for the replication operation. In an optimal situation, the replication backlog should follow the backup activities (Figure 21-11).

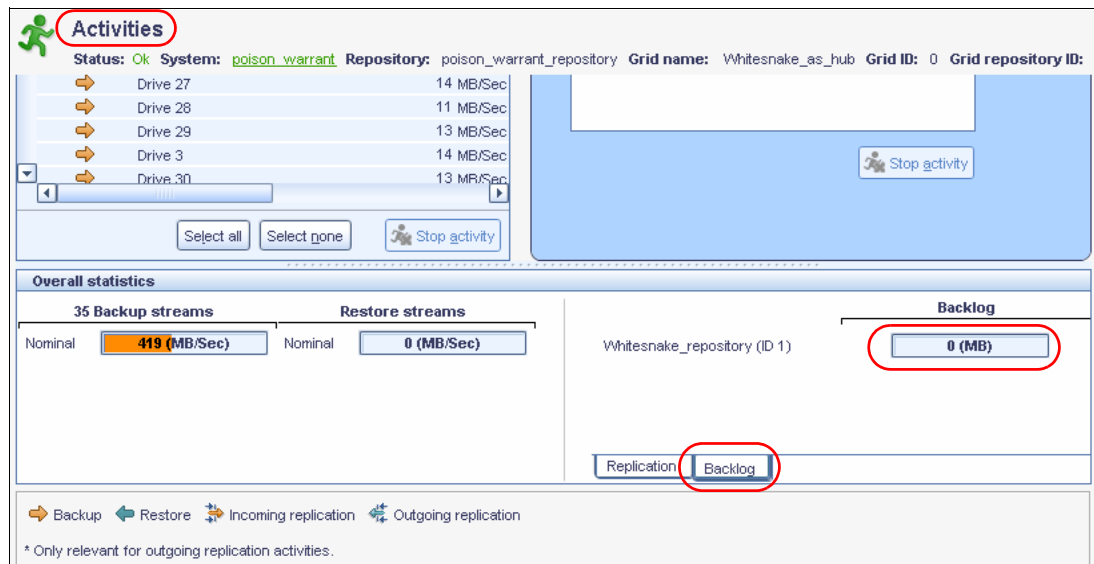


Figure 21-11 Backlog status in replication activities

Use either of these methods to dismiss the replication backlog:

- From the ProtecTIER Manager replication Policy view, select a specific policy and click **Stop activity**.
- From the ProtecTIER Manager replication Activities view, select a specific activity and click **Stop activity**.

These options cause the replication backlog to be discarded; if the replication activities are restarted, the replication backlog will be calculated at that time.

Stopping replication tasks

SLAs: For the system to support the organization's set of SLAs, enough bandwidth must be allotted for replication during the replication window so that all the policies are run in the allotted time.

Stopping replication tasks removes them from the list of pending and running tasks. These tasks are automatically returned to the replication queue if the specific cartridge is in one of the following states:

- ▶ Appended
- ▶ Ejected from the library
- ▶ Selected for manual execution

One way to prevent these replication tasks from rerunning is to mark those cartridges as *read-only* either on the ProtecTIER server or by the backup application. These cartridges are not used for further backups, and therefore do not replicate. New (scratch) sets of cartridges are used for subsequent backups, and do not contain backlog data that does not need to be replicated.

Tip: When backup activity is resumed, using a different set of bar codes can enable having the new data replicated, and skip replication of the data from the old cartridges.

21.9.1 SNMP alerts for replication backlog

ProtecTIER provides a Simple Network Management Protocol (SNMP) method for monitoring backlog data and notifying you if backlog data becomes greater than a user-defined threshold setting, as shown in Figure 21-12.

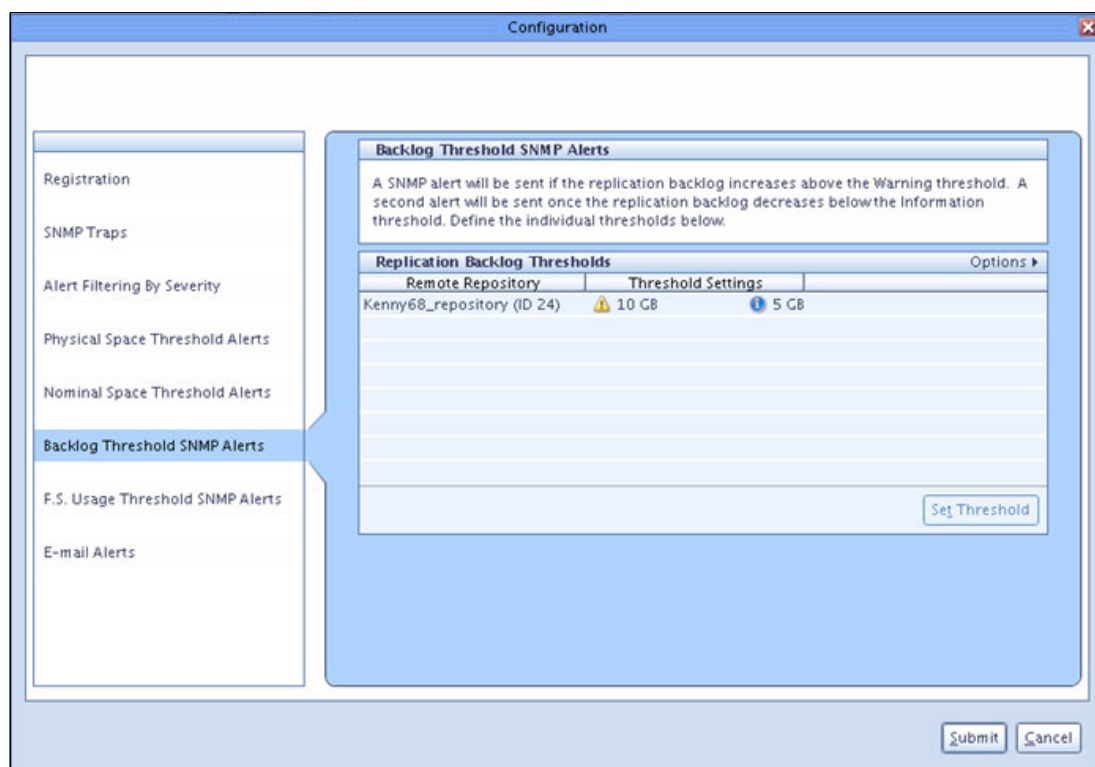


Figure 21-12 Replication backlog SNMP Alerts Reserving space for local backup data

ProtecTIER can reserve local-backup-only space for the hub repository. You can use this enhancement to exclusively assign a portion of a hub repository for local backups. This enhancement was added to ensure that capacity is reserved only for local backup. Replication cannot be written to this portion of the hub repository.

Error notifications display if the repository hub areas that are reserved for local backup or replication are reaching maximum capacity. Figure 21-13 shows the window for this enhancement.

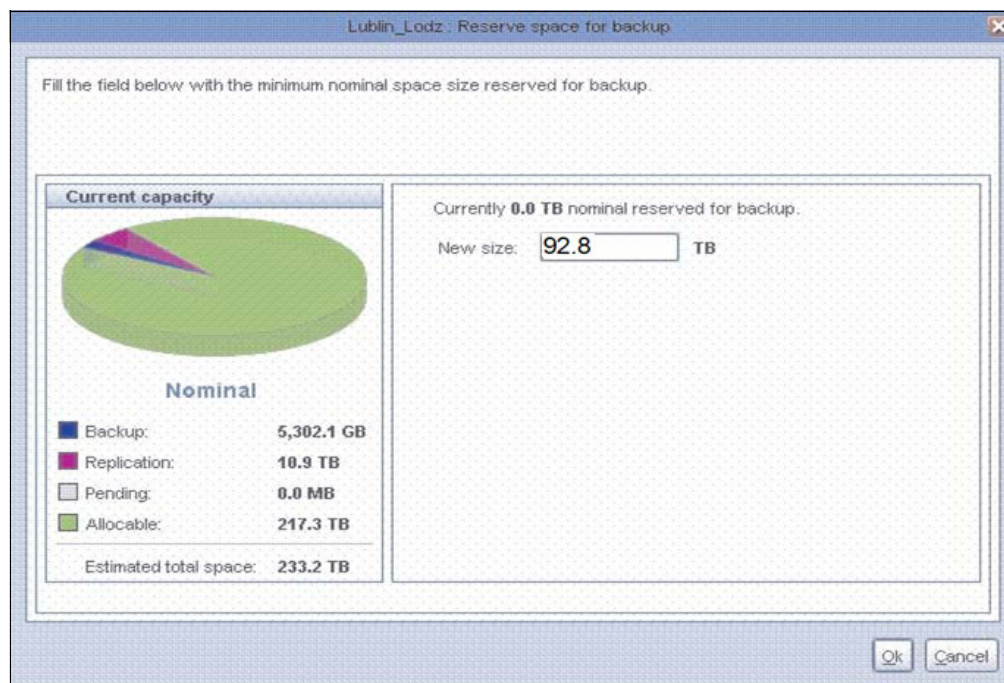


Figure 21-13 Current capacity

21.10 Replication planning

The planning process for ProtecTIER systems with replication deployed requires more input and considerations beyond the individual capacity and performance planning that is needed for a system that is used only as a local VTL, or FSI. When a multiple-site, many-to-one, or many-to-many replication strategy is deployed, the entire configuration, including all spokes and hubs, must be evaluated.

The planning of a many-to-one replication environment is similar to the planning of a one-to-one replication strategy. The only difference is that you must combine all replication loads (and potentially a local backup) for the hub. The ProtecTIER Planner tool should be used at the planning stages before any ProtecTIER deployment Bandwidth sizing and requirements. For more details, see 7.1.1, "ProtecTIER Planner tool" on page 102.

The ProtecTIER server replicates only the new or unique deduplicated data. Data that was deduplicated on the primary server is not sent to the DR site. However, the DR site (hub) must synchronize with all of the data on the primary server to ensure 100% data integrity.

For example, two cartridges are at the primary site, cartridge A and cartridge B, and each contains the same 1 GB of data:

- ▶ Replicating Cartridge A transfers 1 GB of physical (which equals nominal) data. The data is new to the DR site repository.
- ▶ Replicating Cartridge B transfers 0 GB of physical data to the DR site. Because the same data was transferred with the Cartridge A replication, all of the Cartridge B data exists at the DR site repository. Following the replication action, 1 GB of nominal data is indexed on Cartridge B at the DR site.

Throughput: The maximum replication throughput for any scenario depends on many factors, such as the data type and change rate.

Tip: When you configure a system with unbalanced replication network lines, the total throughput is reduced to the slowest line.

The preferred practice is for both networks in the replication network topology to be set at the same speed at the source. If need be, set the speed of the port, by using `ethtool`, of the faster network to the same speed as the slower network. This is an example of the command:

```
ethtool -s eth2 speed 100
```

21.10.1 Replication throughput barriers

The following types of replication data transfer throughput barriers have been identified: physical data-transfer barrier and nominal data barrier.

Physical data-transfer barrier

This barrier results from the Ethernet ports used for replication on the ProtecTIER (1 Gbps or 10 Gbps).

- ▶ 1 Gbps = 1000 Mbps = 125 MBps. If two 1 Gbps ports are used for replication, the maximum possible physical transfer rate is 250 MBps (125 MBps x 2).
- ▶ 10 Gbps = 10,000 Mbps = 1250 MBps. If two 10 Gbps ports are used for replication, the maximum possible physical transfer rate is 2500 MBps (1250 MBps x 2).

The speed of the Ethernet ports is a reference of a physical speed limit that cannot be overpassed. Nevertheless, these physical speed limits are not usually reached, due to many factors that can reduce the transfer rates:

- ▶ TCP's handshake phase. The three-way handshake imposes a certain latency penalty on every new TCP connection.
- ▶ Latency. Depending upon many factors along the network span, the latency in any WAN varies, but must never exceed 200 ms. If so, it might decrease the system replication throughput. For more information about this topic, contact your network administrator.
- ▶ Packet loss. *Packet loss across the network should be 0%.* Any other value indicates a major network problem that must be addressed before replication is deployed. For more information about this topic, contact your network administrator.

Nominal data barrier

The nominal data barrier results from the maximum processing capability of a given ProtecTIER system (3958-DD6):

- ▶ A single node system might support up to 1,660 MBps of nominal data backup ingest, replication, or a combination of these activities.
- ▶ A dual-node clustered system might support sustainable rates of up to 2,990 MBps of nominal data backup ingest, replication, or a combination of these activities.

Note: Maximum specifications are based on a TS7650G 3958-DD6 and a correctly configured back-end disk array. Typical restores are approximately 15 - 20% faster than backups.

21.10.2 Calculating the replication data transfer

Use the following formula to calculate the replication data transfer. The formula estimates the number of gigabytes of changed data to be sent across the network, and adds 0.5% for control data.

Replication data transfer = daily backup × (Change rate + 0.5%)

Example 21-1 shows this formula with values.

Example 21-1 Replication of a 6 TB daily backup with change rate of 10%

replication data transfer = 6,144 GB × (10% + 0.5%) = **645.12 GB**

In this scenario, 645.12 GB of physical data is replicated to the secondary site, rather than 6 TB of nominal data that would otherwise be transferred without deduplication.

21.10.3 Calculating replication bandwidth

Use this formula to calculate the required replication bandwidth:

Replication bandwidth = replication data transfer ÷ available replication hours

Example 21-2 shows this formula with values.

Example 21-2 For a replication window of 10 hours

replication bandwidth = 645.12 GB ÷ 10h = 64.51 GB per hour

The WAN bandwidth must be able to transfer an average 64.51 GB per hour, which represents the requirements for an 18.34 MBps link between spoke and hub.

Tip: Continuous replication operation (24 hour replication concurrent with a backup operation) is rarely the suggested mode of operation. Add 10% of the required bandwidth for a buffer in case of network outages or slowdown periods.

21.10.4 Ports for replication in firewalled environments

In a firewalled environment, you must open the following TCP ports so that IP replication can function properly:

- ▶ The replication manager uses TCP ports 6202, 3501, and 3503.
- ▶ The replication operation between any two repositories uses TCP ports 6520, 6530, 6540, 6550, 3501, and 3503.

ProtectTIER replication does not use any User Datagram Protocol (UDP) ports.

21.11 Bandwidth validation utility

The `pt_net_perf_util` network testing utility is included as part of the ProtectTIER software package. As a part of the installation process, the installer must ensure that the ProtectTIER nodes at both sites (sender and receiver) can run this utility concurrently.

The objective of the `pt_net_perf_util` utility is to test maximal replication performance between two ProtectTIER repositories. It does so by emulating the network usage patterns of the ProtectTIER native replication component. This utility does not predict replication performance, but it might discover performance bottlenecks.

Tip: It is not necessary to build a repository or configure the ProtectTIER back-end disk to run the `pt_net_perf_util` test tool.

The utility includes the following requirements:

- ▶ Red Hat Enterprise Linux Version 5.2 or later.
- ▶ Standard external utilities that are expected to be in the current path are as follows:
 - ping
 - netstat
 - getopt
 - echo

The `pt_net_perf_util` utility uses the `iperf` tool internally. Both tools are installed as part of the ProtectTIER software installation; `pt_net_perf_util` is installed under path `/opt/dtc/app/sbin`, and `iperf` is installed under path `/usr/local/bin`.

Important: Prior to ProtectTIER Version 3.2, the `pt_net_perf_util` utility had the option to use internally either the `iperf` or `nuttcp` tools, but the option to use `nuttcp` was removed.

The utility has two modes of operation: client and server. The client is the ProtectTIER system that transmits the test data. The server is the ProtectTIER system that receives the replication data (also known as the target server). Based on the data that is 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 that replication actually takes place. The replication target must be tested as the server, because the flow of data is to that server from the client. However, it is also important to test the reverse direction to measure the bandwidth performance during disaster recovery failback. Network bandwidth is not always the same in both directions.

21.11.1 Using the bandwidth validation utility to test the data flow

Consider the following generalities before you start the bandwidth validation process:

- ▶ Before you run the utility, the ProtecTIER services on both client and server need to be stopped.
- ▶ The server must be started and running before the client.
- ▶ Each test runs for 5 minutes (300 seconds). Because there are five tests, the process takes about 25 minutes.

Use these steps to test network performance between two ProtecTIER systems on a WAN:

1. Stop the services on both ProtecTIER systems that participate in the test.

Unless otherwise indicated, use the ProtecTIER Service menu to stop ProtecTIER services:

Manage ProtecTIER services > Stop ProtecTIER services only (including GFS)

2. Start the server mode of the **pt_net_perf_util** utility on the target server (the ProtecTIER system that receives the replication data).

Example 21-3 shows how to start the **pt_net_perf_util** utility in server mode. The **-s** flag indicates to the utility to start as server.

Example 21-3 Start pt_net_perf_util server mode

```
[root@tintan ~]# /opt/dtc/app/sbin/pt_net_perf_util -s
```

```
-----  
Server listening on TCP port 5001  
TCP window size: 85.3 KByte (default)  
-----
```

3. Start the client mode of the **pt_net_perf_util** utility on the client (the ProtecTIER system that sends the replication data).

Example 21-4 shows how to start the **pt_net_perf_util** utility in client mode.

The **-c <server>** flag indicates to the utility to start as a client, and to connect to the given server. The **-t** flag indicates the seconds to run each test. Without the **-t** flag, the utility will not run, and an error (ERROR: -t not specified) along with the utility usage will be displayed. Unless otherwise indicated, use 300 seconds to start the client.

Example 21-4 Start pt_net_perf_util client mode

```
[root@torito ~]# /opt/dtc/app/sbin/pt_net_perf_util -c 10.0.5.44 -t 300
```

```
*** Latency
```

4. The utility automatically performs the tests in sequence; wait until the client completes all tests. Example 21-5 shows the output of the client after all tests completed running.

Example 21-5 Output of the pt_net_perf_util client

```
[root@torito ~]# /opt/dtc/app/sbin/pt_net_perf_util -c 10.0.5.44 -t 300
```

***** Latency**

```
PING 10.0.5.44 (10.0.5.44) 56(84) bytes of data.
```

```
--- 10.0.5.44 ping statistics ---
```

```
300 packets transmitted, 300 received, 0% packet loss, time 298999ms
```

```
rtt min/avg/max/mdev = 0.066/0.118/0.633/0.039 ms
```

```
*** Throughput - Default TCP
[ 3] 0.0-300.0 sec 32.9 GBytes 942 Mbits/sec
```

```
*** Throughput - 1 TCP stream(s), 1024KB send buffer
[ 3] 0.0-300.0 sec 32.9 GBytes 942 Mbits/sec
```

```
*** Throughput - 16 TCP stream(s), 1024KB send buffer
[SUM] 0.0-300.4 sec 32.9 GBytes 942 Mbits/sec
```

```
*** Throughput - 127 TCP stream(s), 1024KB send buffer
[SUM] 0.0-303.7 sec 33.2 GBytes 940 Mbits/sec
```

Number of TCP segments sent: 4188852

Number of TCP retransmissions detected: 969530 (23.145%)

As the tests are run, the server prints output of each test. Example 21-6 shows excerpts of the output on the server.

Example 21-6 Output on the pt_net_perf_util server

```
[root@tinTAN ~]# /opt/dtc/app/sbin/pt_net_perf_util -s
-----
Server listening on TCP port 5001
TCP window size: 85.3 KByte (default)
-----
[ 4] local 10.0.5.44 port 5001 connected with 10.0.5.41 port 43216
[ ID] Interval      Transfer    Bandwidth
[ 4] 0.0-300.0 sec 32.9 GBytes 941 Mbits/sec
[ 5] local 10.0.5.44 port 5001 connected with 10.0.5.41 port 43227
[ ID] Interval      Transfer    Bandwidth
[ 5] 0.0-300.0 sec 32.9 GBytes 941 Mbits/sec
[ 4] local 10.0.5.44 port 5001 connected with 10.0.5.41 port 43238
[ 5] local 10.0.5.44 port 5001 connected with 10.0.5.41 port 43239
...
[ ID] Interval      Transfer    Bandwidth
[ 16] 0.0-300.4 sec 2.06 GBytes 58.8 Mbits/sec
[ ID] Interval      Transfer    Bandwidth
[ 18] 0.0-300.4 sec 2.05 GBytes 58.7 Mbits/sec
...
[ 20] local 10.0.5.44 port 5001 connected with 10.0.5.41 port 43264
[ 4] local 10.0.5.44 port 5001 connected with 10.0.5.41 port 43265
...
[ ID] Interval      Transfer    Bandwidth
[101] 0.0-303.9 sec 269 MBytes 7.43 Mbits/sec
[ ID] Interval      Transfer    Bandwidth
[ 88] 0.0-303.9 sec 224 MBytes 6.18 Mbits/sec
...
[ ID] Interval      Transfer    Bandwidth
[ 13] 0.0-306.9 sec 186 MBytes 5.08 Mbits/sec
[SUM] 0.0-306.9 sec 33.2 GBytes 930 Mbits/sec
```

- Restart the ProtecTIER services after finishing the tests. Unless otherwise indicated, use the ProtecTIER Service menu to stop ProtecTIER services:

Manage ProtecTIER services > Start all services

21.11.2 Interpreting the results

The utility performs five foreground tests and one background test (TCP) re-transmissions (versus total TCP segments sent on the five foreground tests).

Interpretations of the following tests is based on the example output that is shown in Example 21-5 on page 383. Results vary in each environment.

Test 1: Latency

This test checks the nominal network link latency and packet loss. The following results can be interpreted:

- ▶ There was 0% packet loss.
- ▶ The average round-trip-time (RTT) was 0.118 ms.

The latency in WAN topologies might vary, but must never exceed 200 ms. Contact your network administrator if latency reports more than 200 ms. Higher latency values cause a major deterioration in replication throughput. Packet loss must be 0%, because any other value implies a major network problem.

Test 2: Throughput - Default TCP

This test checks the maximum TCP throughput by using a single data stream with default TCP settings. The following results can be interpreted:

- ▶ The test ran for 300 seconds.
- ▶ 32.9 GB of data was transferred.
- ▶ The average throughput was 942 Mbps.

Remember: 1 MB = 1,048,576 bytes. 1 MBps = 1,000,000 Bps

Test 3: Throughput - 1 TCP stream(s), 1024 KB send buffer

This test checks the maximum TCP throughput by using a single data stream with a 1 MB send buffer. The following results can be interpreted:

- ▶ The test ran for 300 seconds.
- ▶ 32.9 GB of data was transferred.
- ▶ The average throughput was 942 Mbps.

Test 4: Throughput - 16 TCP stream(s), 1024 KB send buffer

This test checks the maximum TCP throughput by using 16 data stream with a 1 MB send buffer. The following results can be interpreted:

- ▶ The test ran for 300.4 seconds.
- ▶ 32.9 GB of data was transferred.
- ▶ The average throughput was 942 Mbps.

The megabits per second reported in this test is the maximum replication performance the system can achieve if the backup environment uses three or fewer cartridges in parallel.

Test 5: Throughput - 127 TCP stream(s), 1024 KB send buffer

This test checks the maximum TCP throughput by using 127 data streams with a 1 MB send buffer. The following results can be interpreted:

- ▶ The test ran for 303.7 seconds.
- ▶ 33.2 GB of data was transferred.
- ▶ The average throughput was 940 Mbps.

The throughput value that is 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. If this number is lower than anticipated, contact your network administrator.

Test 6: TCP retransmissions versus total TCP segments sent

This test compares the total TCP transmissions sent with the number of packets that are lost and retransmitted. The following results can be interpreted:

- ▶ A total of 4,188,852 TCP segments were sent during the five tests.
- ▶ 969,530 were lost and retransmitted.
- ▶ The retransmission rate is 23.145%.

The retransmission rate imposes a direct penalty on the throughput, because the retransmission of these packets take up bandwidth. The retransmission can be caused by the underlying network (for example, packet dropping by an overflowed router). It can also be caused by the TCP layer (for example, retransmission because of packet reordering). Segment loss can be caused by each of the network layers.

Important: TCP retransmission larger than 2% might cause performance degradation and unstable network connectivity. Contact your network administrator to resolve this issue.

21.11.3 Repository replacement

Use the repository replacement function when you want to fail back to a different repository or rebuild a repository. To accomplish this task, complete the following steps:

1. Cancel the pairing of the original repositories in the replication manager.
2. Take the original primary repository out of the replication grid.

Important: If a new repository replaces the original one, then the new repository must be installed and join the replication grid.

3. Run the ProtecTIER repository replacement wizard and specify the repository to be replaced and the replacement repository.

After the disaster recovery situation ends and the primary repository is restored or replaced, you can return to normal operation with the replacement repository on the production site as the primary site.

Figure 21-14 shows how to leave ProtecTIER DR mode by selecting **Replication** → **Replication Disaster Recovery** → **Leave DR mode**.

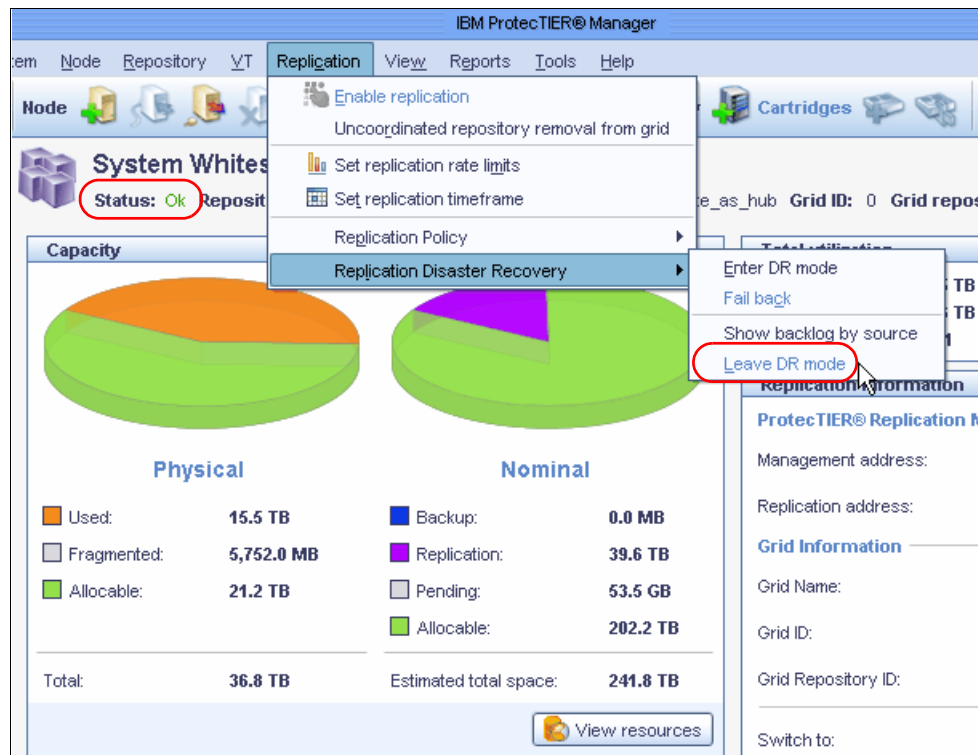


Figure 21-14 Leaving ProtecTIER DR mode

Important: Leaving DR mode should always be preceded by a failback action.

For more information, see *IBM TS7650 with ProtecTIER V3.4 User's Guide for VTL Systems*, GA32-0922-10.

Cartridge ownership takeover

Cartridge ownership takeover enables the local repository, or hub, to take control of cartridges that belong to a deleted repository. Taking ownership of the cartridges on a deleted repository enables the user to write on the cartridges that previously belonged to the replaced (deleted) repository. This process is also known as a change of *principality*.

Cartridge ownership: The repository can take ownership of a cartridge only if the repository is defined on the Replication Manager as the replacement of the deleted repository.

21.12 Planning ProtecTIER replication

This section provides case studies of planning and sizing ProtecTIER Replication. Both many-to-one (spoke and hub) replication environment and many-to-many bidirectional replication scenarios are described.

21.12.1 Deployment planning scenario: many-to-many

This section shows a deployment planning scenario for four sites, each with a dual-node gateway, building a maximum of four repositories, many-to-many VTL configuration, with various replication strategies.

VTL and FSI systems are configured in many-to-many replication groups, so this same sizing strategy applies, but the throughput numbers vary for each type.

At the time of the writing of this book, the maximum listed speed for a dual-node DD6 VTL gateway is 2500 MBps, so all calculations are based on this speed. As ProtecTIER technology improves, the rated performance numbers continue to increase. For the current published ratings, see the following web page:

<http://www.ibm.com/systems/storage/tape/ts7650g/index.html>

Assume that the following characteristics of the replication grid are present:

- ▶ All processes have the same maximum rate (2,500 MBps).
- ▶ All data exists at all sites.

Maximum backup with no replication

With no data replication, a maximum of 24 hours can be used to accept backup data. One 24-hour time slot translates to 216 TB per day for each system. *This is not a recommended configuration; it is included in this scenario for purposes of the example.*

Figure 21-15 shows an overview of this scenario.

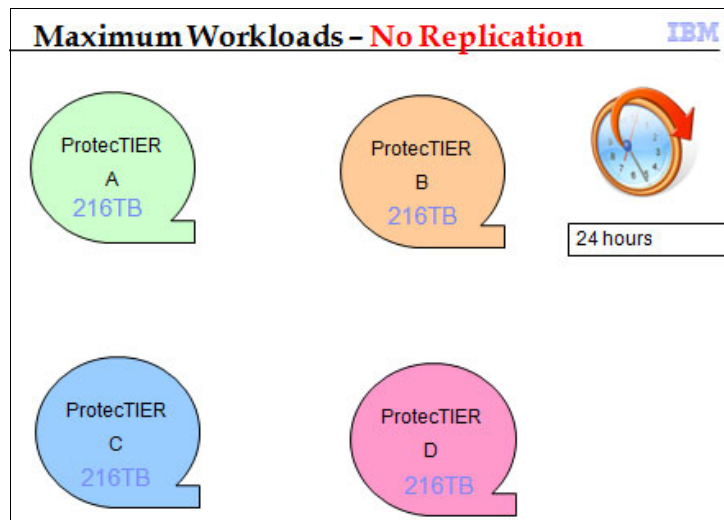


Figure 21-15 Maximum workload with no replication

Maximum workload with one replicated copy

For this example, all four ProtecTIER systems receive and replicate the same maximum amount of data that is possible in a 24 hour period. Because the workloads are equal, you can divide the 24 hour period into three equal time slots:

- ▶ One backup process (all four nodes accept backup at the same time.)
- ▶ One incoming replication processes
- ▶ One outgoing replication processes

With one data replication for each node, a maximum of 8 hours can be used to accept backup data. One 8 hour time slot translates to 72 TB per day for each system.

Figure 21-16 shows an overview of this scenario.

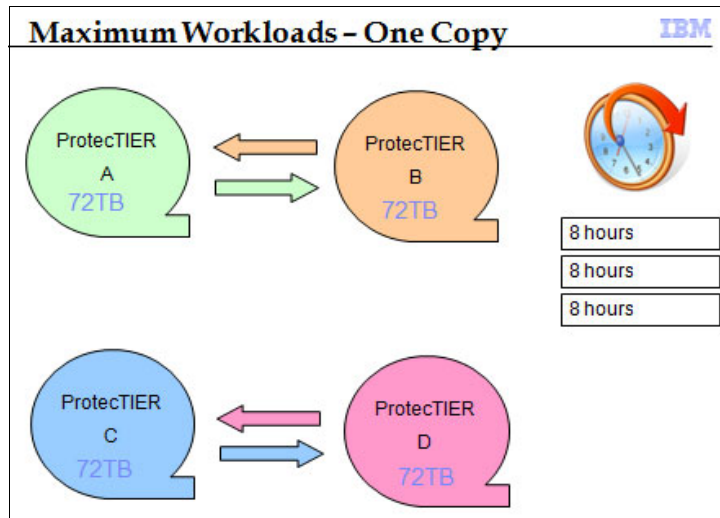


Figure 21-16 Maximum workload with one replicated copy

Two replicated copies

For this example, all four ProtecTIER systems receive and replicate the same maximum amount of data that is possible in a 24 hour period. Because the workloads are equal, you can divide the 24 hour period into five equal time slots:

- ▶ One backup process (all four nodes accept backup at the same time.)
- ▶ Two incoming replication processes
- ▶ Two outgoing replication processes

With two data replications for each node, a maximum of 4.8 hours can be used to accept backup data. One 4.8 hour time slot translates to 43 TB per day for each system.

Figure 21-18 on page 390 shows an overview of this scenario.

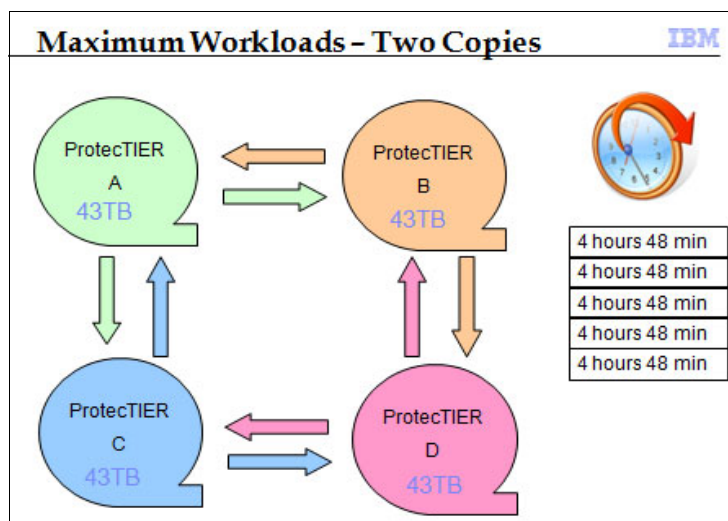


Figure 21-17 Maximum workload with two replicated copies

Three replicated copies

For this example, all four ProtecTIER systems receive and replicate the same maximum amount of data that is possible in a 24 hour period. Because the workloads are equal, you can divide the 24 hour period into seven equal time slots:

- ▶ One backup process (all four nodes accept backup at the same time)
- ▶ Three incoming replication processes
- ▶ Three outgoing replication processes

With three data replications for each node, a maximum of 3.4 hours can be used to accept backup data. One 3.4 hour time slot translates to 30.6 TB per day for each system.

Figure 21-18 shows an overview of this scenario.

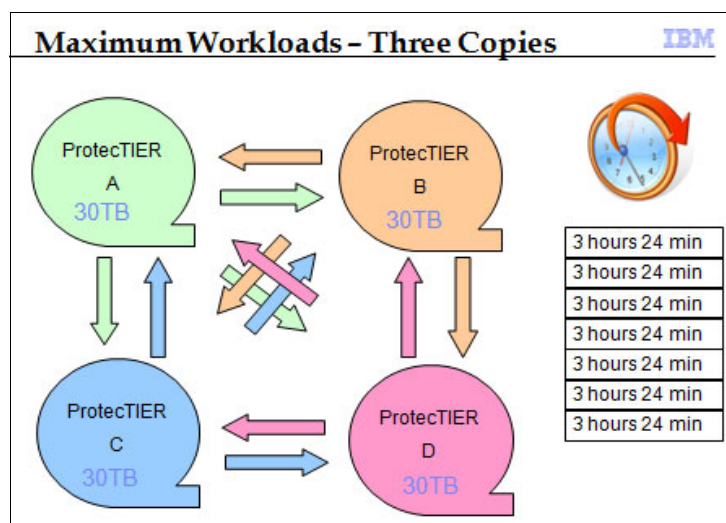


Figure 21-18 Maximum workload with three replicated copies

Table 21-1 depicts activity in different time frames in a 4-way many-to-many replication configuration.

Table 21-1 Example of a 4-way many-to-many replication configuration

Time frame	Activity
1	All systems process backups at 2500 MBps for 3.4 hours (30.6 TB).
2	System C replicates to B at 2500 MBps for 3.4 hours. System D replicates to A at 2500 MBps for 3.4 hours.
3	System C replicates to D at 2500 MBps for 3.4 hours. System A replicates to B at 2500 MBps for 3.4 hours.
4	System C replicates to A at 2500 MBps for 3.4 hours. System B replicates to D at 2500 MBps for 3.4 hours.
5	System B replicates to A at 2500 MBps for 3.4 hours. System D replicates to C at 2500 MBps for 3.4 hours.
6	System D replicates to B at 2500 MBps for 3.4 hours. System A replicates to C at 2500 MBps for 3.4 hours.
7	System B replicates to C at 2500 MBps for 3.4 hours. System A replicates to D at 2500 MBps for 3.4 hours.

21.12.2 Many-to-one replication

The specific performance numbers would vary depending of the ProtecTIER model used, but this same process can be followed for IBM System Storage TS7620 small and medium business (SMB) appliances and TS7650G Gateway systems when sizing and planning a replication scenario. Figure 21-19 shows a many-to-one replication example.

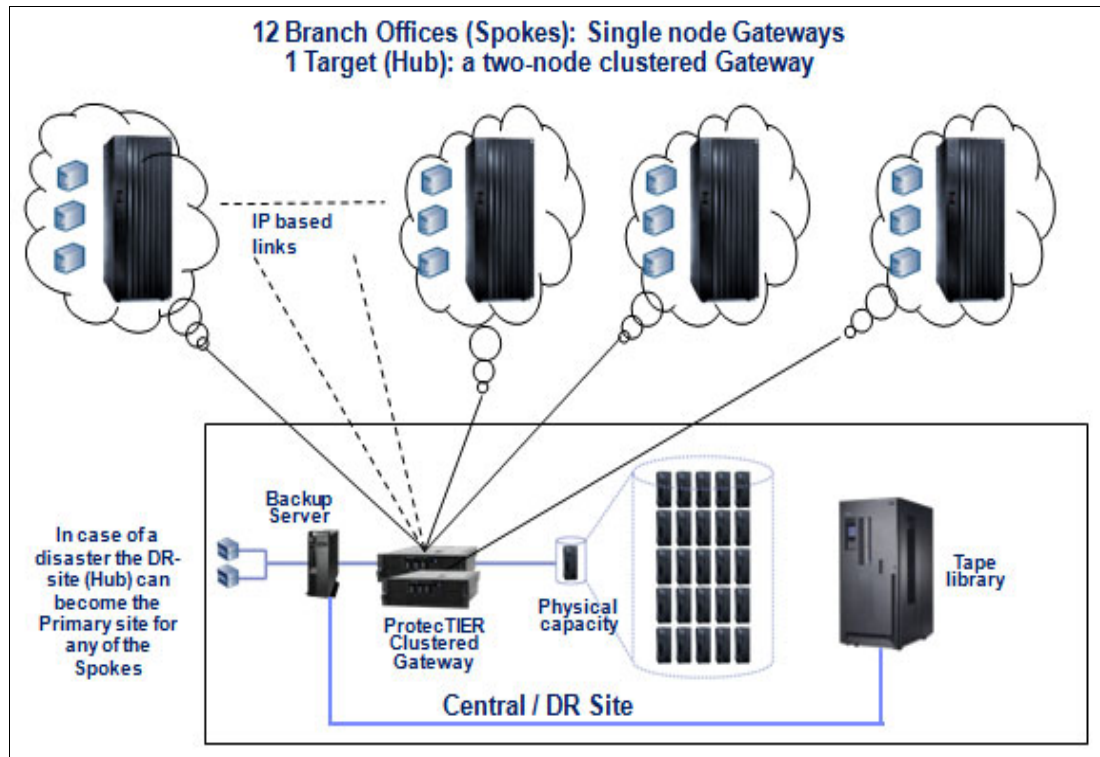


Figure 21-19 Many-to-one replication example

Assumptions

The case modeling is based on the following assumptions:

- ▶ A maximum environment of 12 Spoke systems and one hub system.
- ▶ Eight hour backup windows (hub and spokes).
- ▶ A 16 hour replication window.
- ▶ All windows are aligned, meaning that the eight hour backup window is the same actual time at all 13 ProtecTIER systems (spokes and hub).
- ▶ Adequate bandwidth between all spokes and hub.
- ▶ A 10:1 deduplication ratio throughout the system.
- ▶ Data change rate at spokes does not saturate physical reception capabilities at the hub.

Maximum workloads assumed

The values for the maximum workloads are as follows:

- ▶ Hub backup:
 - 8-hour backup window.
 - 9 TB per hour (2500 MBps).
 - 72 TB of nominal daily backup at the hub.
 - 16-hour replication window.
 - 9 TB per hour (2500 MBps) replication performance.
 - 144 TB of nominal data can be replicated from all spokes.
- ▶ Spoke backup:
 - 8-hour backup window.
 - 144 TB for all 12 spokes = 12 TB of daily backup data per spoke.
 - 12 TB/eight hours = 1.5 GB per hour or 1500 MBps sustained for eight hours.
 - A spoke can potentially back up 72 TB of nominal data, but can replicate only 12 TB because of configuration constraints.

Sizing the repositories for spokes and hub

This section provides examples for sizing repositories for spokes and hubs. It also provides examples for calculating local backup space and incoming replication space.

Example of spoke repository sizing

In this example, each spoke can process up to 12 TB per day of local backup data, with a 10:1 deduplication ratio.

To size the spoke repository in this example, complete the following steps:

1. Assuming a 10:1 deduplication ratio, approximately 1200 GB of new data must be replicated to the hub per backup. The total daily space for 27 incremental backups is calculated as follows:
 $1,200 \text{ GB} \times 27 \text{ incrementals} = 32,400 \text{ GB}$ (or ~32 TB) of physical space (for incrementals)
2. With a backup compression ratio of 2:1, add 6 TB for the first “full” backup (12 TB at 2:1 compression):
 $32 \text{ TB} + 6 \text{ TB} = 38 \text{ TB}$ of physical space for incrementals and full backup
3. Calculate the space that is necessary for spare capacity by multiplying the total physical space that is needed by 10%:
 $38 \text{ TB} \times 10\% = 3.8 \text{ TB}$ of spare capacity
4. Calculate the total physical repository for each spoke by adding the total physical space that is needed and the spare capacity:
 $38 \text{ TB} + 4 \text{ TB} = 42 \text{ TB}$

Example of hub repository sizing

The hub repository must be sized to handle 27 days of local backups and 27 days of incoming replication from all 12 spokes plus approximately 10% spare capacity.

Local backup space

In this example, the hub system can back up 72 TB in the 8-hour window. The first full backup at the hub requires 36 TB of physical space (72 TB @ 2:1 compression ration). With a 10:1 deduplication ratio, the hub accumulates 7.2 TB of new data for each of the next 27 days.

The following example is the calculation for the local backup space:

$$36 \text{ TB} + 194.4 \text{ TB} (7.2 \text{ TB} \times 27 \text{ days}) = 230.4 \text{ TB}$$

Incoming replication space

To calculate the incoming replication space in this example, complete the following steps:

1. Calculate the hub repository space for a full backup of all 12 spokes at 2:1 compression:
 $(12 \text{ TB} \times 12 \text{ spokes})/2 = 72 \text{ TB}$ of repository space
2. Assuming a 10:1 deduplication ratio, approximately, 1,200 GB (1.2 TB) of new data per spoke must be replicated to the hub per backup. Calculate the new data received daily at the hub from all spokes:
 $1200 \text{ GB} \times 12 \text{ spokes} = 14.4 \text{ TB}$ of new data
3. The total daily space for 27 incremental backups is calculated as follows:
 $14.4 \text{ TB} \times 27 \text{ incrementals} = 388 \text{ TB}$ of physical space
4. The total hub repository space that is necessary to accommodate the 27 incremental backups and one full backup is:
 $230.4 \text{ TB} + 388 \text{ TB} + 40 \text{ TB} (10\% \text{ spare capacity}) = 464 \text{ TB}$ for hub repository space

21.13 The backup application database backup

Figure 21-20 illustrates a typical backup and DR environment using the ProtecTIER product.

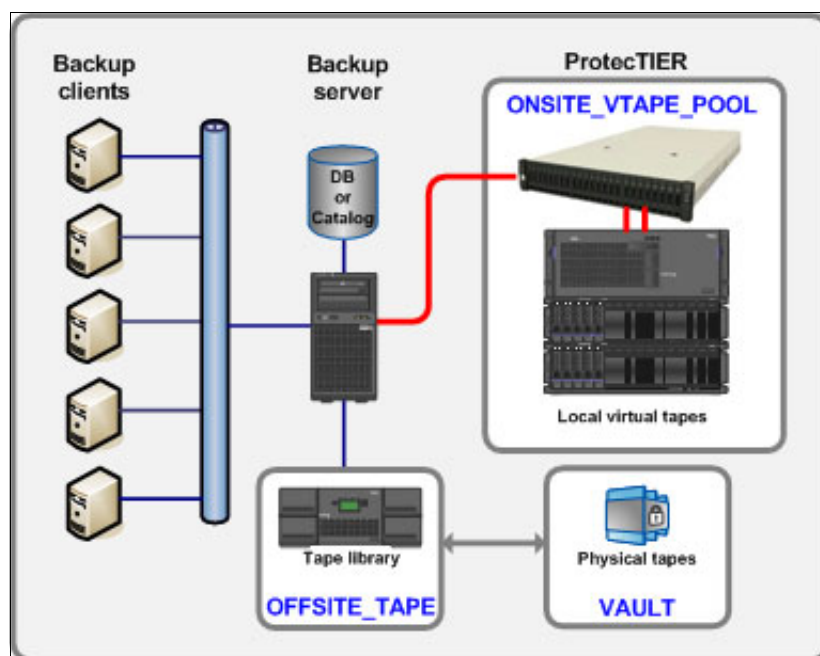



Figure 21-20 Typical backup and DR environment using the ProtecTIER product

The backup application environment is straightforward. The backup application servers are connected to storage devices (disk, real tape, or virtual tape). Every action and backup set that the backup servers process is recorded in the backup application database or catalog. The catalog is at the heart of any recovery operation. Without a valid copy of the database or catalog, restoration of data is difficult, and sometimes even impossible.

The ProtecTIER server provides a virtual tape interface to the backup application server, which you can use to create tapes, as represented by `ONSITE_VTAPE_POOL` in Figure 21-20 on page 393. The client can also maintain another tape library to create real tapes to take off-site, called `OFFSITE_TAPE` in Figure 21-20 on page 393.

`ONSITE_VTAPE_POOL` is where most client recoveries and restores come from. The key advantage of this architecture is that restoration occurs much faster because the data is coming from the ProtecTIER disk-based virtual tape rather than from real tape.



Disaster recovery deployment with backup applications

This chapter provides preferred practices, general rules, and setup for disaster recovery (DR) considerations for the following backup applications:

- ▶ IBM Spectrum Protect (as of version 7.1.3, Tivoli Storage Manager was rebranded to IBM Spectrum Protect)
- ▶ Symantec NetBackup
- ▶ EMC NetWorker
- ▶ Commvault

It also provides effective techniques to perform DR processes and the related failback process after the primary location recovers from the disaster situation.

This chapter contains the following topics:

- ▶ Disaster recovery operations
- ▶ ProtecTIER replication overview
- ▶ Disaster recovery operations with VTL
- ▶ Disaster recovery operations with FSI
- ▶ Entering ProtecTIER DR mode
- ▶ The backup application catalog
- ▶ Single domain and multiple domains
- ▶ Deploying replication with specific backup applications
- ▶ IBM Spectrum Protect
- ▶ Symantec NetBackup deployment with ProtecTIER replication
- ▶ EMC NetWorker deployment with ProtecTIER replication
- ▶ Commvault

Note: ProtecTIER GA Version 3.4 was released with only the Virtual Tape Library (VTL) interface support. File System Interface (FSI) support was added to ProtecTIER PGA 3.4 Version. For details, see the announcement letter:

http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/6/872/ENUSA P16-0076/index.html&lang=en&request_locale=en

22.1 Disaster recovery operations

Disaster recovery (DR) is the process of recovering production site data at a secondary (or tertiary) location. Disaster recovery is useful if a disaster occurs or a situation occurs where the production (or primary) site goes offline. The secondary site, hub, or DR site, can take the place of the production site until the primary site comes back online.

In ProtecTIER VTL environments, entering DR mode is done by using the ProtecTIER Manager graphical user interface (GUI) when you are logged on to the DR ProtecTIER system. When you enter ProtecTIER DR mode, all incoming replication and visibility switching activities from the failed production site to the DR site (hub) are blocked. The DR node remains eligible for local backup and restore activity.

When the primary site is rebuilt or replaced, you can then use the ProtecTIER Failback function to return the systems to their normal state and resume the backup and replication operations. When the primary site comes back online, previously replicated and newly created objects can be moved to the main production site by using the failback process. The primary site then resumes its roll as a production site.

Note: The Disaster Recovery and Failback menu options are available for VTL environments. For FSI environments, you can block the replication traffic by selecting **Replication** → **Control Replication traffic** in the ProtecTIER Manager GUI. This Replication menu also provides other operations for DR purposes, such as Cloning.

For details about the DR concepts for VTL and FSI, see the topic about native replication and disaster recovery in *IBM System Storage TS7600 with ProtecTIER Version 3.3*, SG24-7968.

22.2 ProtecTIER replication overview

ProtecTIER replication enables virtual tape cartridges (VTL mode) or file system volumes (FSI mode) to be replicated from the primary site to secondary or multiple sites. This capability is intended to enhance DR and business continuity (BC).

Users can select some or all of their data objects to be replicated to one or more sites. ProtecTIER replication provides DR capability and enhanced availability in accessing backed up data for restore operations across sites. Replication can run in parallel with backup and restore activities or in dedicated scheduled time frames.

Data transfer is started based on trigger points in either on-demand continuous replication or replication schedules. Data verification and validation is done at the DR site to ensure the integrity of the transferred data before you make the virtual cartridge or tape available.

In summary, replication works as follows:

- ▶ Replication is a process that transfers logical objects such as cartridges from one ProtecTIER repository to another repository.
- ▶ The replication function enables ProtecTIER deployment to be distributed across sites. Each site can have a single or clustered ProtecTIER environment.
- ▶ Each ProtecTIER environment has at least one ProtecTIER server.
- ▶ Each ProtecTIER server that is a part of a replication grid has two dedicated replication ports that are used for replication.

- ▶ Replication ports are connected to the customer's WAN. By default, each ProtecTIER server is configured on two replication subnets.
- ▶ The replication groups are configured and managed in the ProtecTIER Grid Manager.

22.3 Disaster recovery operations with VTL

This section describes DR operations in a VTL environment. It describes managing cartridges with replication by using basic disaster recovery and visibility switching, cartridge replication requirements, and importing/exporting slots allocation and searching in a VTL.

22.3.1 Replication data transfer

When the replication action is started either manually or based on a policy, the source (primary) ProtecTIER system carries out the following procedures:

- ▶ Initiates the sync-cartridge function between its own (source, that is, primary site) repository and the destination (remote, DR site) repository.
- ▶ Reads the unique replication data units on requests from the remote ProtecTIER system based on what it is missing.
- ▶ Sends the unique data elements, by using Transmission Control Protocol (TCP), over the wide area network (WAN) to the remote (DR) site.

At the same time, the destination ProtecTIER system performs the following handshake actions in this order:

1. Calculates the relevant cartridges' *sync* point from where the replication must start.
2. Receives many data units concurrently as part of the replication action.
3. Verifies cyclic redundancy check (CRC) for all replicated data before it becomes available as part of the data object.

After the CRC check is successful, the system moves each of the verified data elements in to scope and makes it available at the DR site.

22.3.2 Managing cartridges after replication

When the replication jobs complete, the handling of the replicated cartridges is determined by the replication policy. The cartridge is always replicated to the ProtecTIER shelf at the target site. When the source cartridge is ejected from the source system virtual library, one of two options happens: It is left off the shelf or it is automatically moved to the import/export slot of the target virtual library. The automated cartridge movement is known as *visibility switching*.

These two options are commonly referred to as follows:

- ▶ Basic disaster recovery (DR)
- ▶ Visibility switching

Basic disaster recovery (DR)

The replicated cartridge stays on the target shelf after replication completes. Basic DR is similar to disaster recovery with a physical Tape Library where cartridges are kept on a physical shelf at the DR site or at a remote storage facility. When the source site fails, physical cartridges can be rapidly transferred to the DR location and imported in to the standby library. The same concept exists with the ProtecTIER product when you use the basic DR mode.

Cartridges are on the virtual shelf of the target system and are ready to be imported to an existing or new virtual library.

Visibility switching

Visibility switching is similar to a warm backup site practice. Physical cartridges are shipped from the source location to the DR site and stored in the physical slots of the standby library. When a disaster is declared, the necessary cartridges are immediately available for recovery at the DR site.

Added value of visibility switching

In a virtual world, the difference in processor usage between the previously described modes is minimal. Importing cartridges from the virtual shelf to a library is fast and requires little effort. It can be done from anywhere (with network access to the system). Storing cartridges on the virtual shelf does not make the DR system less reactive for recovery. As a result, the recovery time objective (RTO) that a ProtecTIER replication-based DR solution offers represents an improvement over a physical tape-based solution.

The advantage of visibility switching is more versatile management for backup applications that support cross-site distributed tape library management of their catalogs. Backup applications that can manage multiple sites through a universal catalog or database can use *automated cartridge movement*. Automated cartridge movement easily moves cartridges from site to site without using any other interface other than the backup application.

With a replication policy that is configured to use visibility switching, when the backup application ejects a cartridge from a source library the cartridge then appears at an import and export slot of a designated DR site library (pending completion of replication). Likewise, cartridges can be moved back to the source site library by using the reverse process. Having control of the cartridge movement through the backup application simplifies the process of cloning cartridges to physical tape at the target site.

By eliminating the usage of *physical tape* from the replication process, a few more steps by the ProtecTIER server are required if visibility switching is not used. Because backup applications cannot handle the same bar code at multiple sites, cartridges are visible in only one library at a time. Therefore, with the ProtecTIER server, you cannot have a cartridge visible in two libraries even if the data exists in both locations. To create a physical copy at the DR site without visibility switching, the administrator must perform the following actions:

- ▶ Import the replica into the target library after replication completes.
- ▶ Import the source cartridge back into the primary library when the clone job completes.

Some well-known backup applications that support single domain are Symantec NetBackup, EMC NetWorker, and IBM System i BRMS.

For applications that do not support a single domain, the automated visibility switching mechanism is of no real value. Each library is managed by a separate entity with no shared knowledge of the replicated volume's content and whereabouts. In these environments, the local backup server must proceed through a recent backup catalog or database that describes the content of the associated data volumes. Every set of replicated data cartridges that is imported into a target library needs to be preceded by a recent catalog or database update.

After the backup server at the target site is updated, the creation of physical tapes requires the moving of the cartridges. The cartridges are moved from the shelf to the library and exported back to the shelf when you are finished. This part of the tape creation procedure is the same as it is in a single domain environment.

22.3.3 Cartridge replication requirements

The following list describes the main cartridge replication requirements:

- ▶ A ProtecTIER server does not replicate a cartridge if the target instance is in a library.
- ▶ Cartridges should be exported back to the shelf after cloning.
- ▶ A cartridge that remains in the target library after a cloning prevents further replication.
- ▶ Before you move cartridges into a source library, verify that a cartridge is not left in the target library.

22.3.4 Importing/exporting slots allocation in VTL

When you create a virtual library, the default is to create eight import/export slots. A low import/export slot count can affect your DR process. For this reason, allocate *more than eight* import/export slots for DR purposes (the maximum value is 1022 per library and 4096 per repository). Allocation of more import/export slots is also important if visibility switching is used and in cases of a DR strategy with a heavy cartridge ejection requirement is implemented.

Tip: Create enough import/export slots so that all of your DR cartridges are processed in one single step.

For example, if you need to move 32 cartridges for your initial DR operation, create at least 32 import/export slots in your virtual library. This configuration reduces the DR complexity.

22.3.5 Import/export slots searching

When you move replicated cartridges from the virtual shelf to a library, the backup application needs to scan and find the new cartridges in the import/export slots. Scanning for the new cartridges is a time-consuming action. Alternatively, you can create more empty slots (for example, when you create the library, choose X cartridges and X+100 slots). The slot locations are already known to the backup application, and therefore reduce the scan time that is required to scan for new cartridges.

22.3.6 Automation of daily operation

The automation process of moving cartridges between sites and performing clone-to-physical-tape operations at the secondary site is more suitable in single-domain backup application environments. Some of the major backup applications, such as NetBackup, EMC NetWorker, and Backup, Recovery, and Media Services (BRMS), support this type of environment.

The following elements are an example of a possible automation opportunity in a NetBackup backup application environment:

- ▶ Vault profile

You can create a vault profile for ejecting cartridges from a library by running the following command:

```
vltrun <vault profile name>
```

- ▶ You can inject a cartridge at the DR site to the library by running the following command:

```
vltinject <vault profile name>
```

- ▶ Bar codes

- You can eject cartridges from a library by running the following command:

```
vmchange -res -multi_eject -ml <barcodeA:barcodeB:...:barcodeZ>
```

- You can inject cartridges to a library by running the following command:

```
vmchange -res -multi_inject -ml <barcodeA:barcodeB:...:barcodeZ>
```

- ▶ Inventory command

This command scans the import/export slots and injects all available cartridges to the library:

```
vmupdate -rt <robot type>-m<robot #> -empty_map
```

Scripting the inject/eject commands

Vault and inject/eject commands can be scripted to run periodically on the backup application host. The script triggers an automatic cartridge movement from the import/export slot to the library whenever the relevant cartridge is in the import/export slot. This process ensures that there are free import/export slots. Example 22-1 shows this script.

Example 22-1 Example script of inject/eject commands

```
#!/bin/csh
while (1)
  vltinject myVault
  sleep 600
end
```

Scripting the inventory command

Scripting the inventory command is *not* recommended because it scans the robot and therefore might take a long time to complete in libraries with many cartridges.

22.3.7 Gauging the replication completion status

You can gauge the replication completion status through the ProtecTIER Manager GUI or the ptcli command-line tool.

To use the GUI, follow the steps in Figure 22-1: Click Shelf (1), select **Replica properties** from Shelf cartridges (2), and sort the display list by the In-Sync status (3).

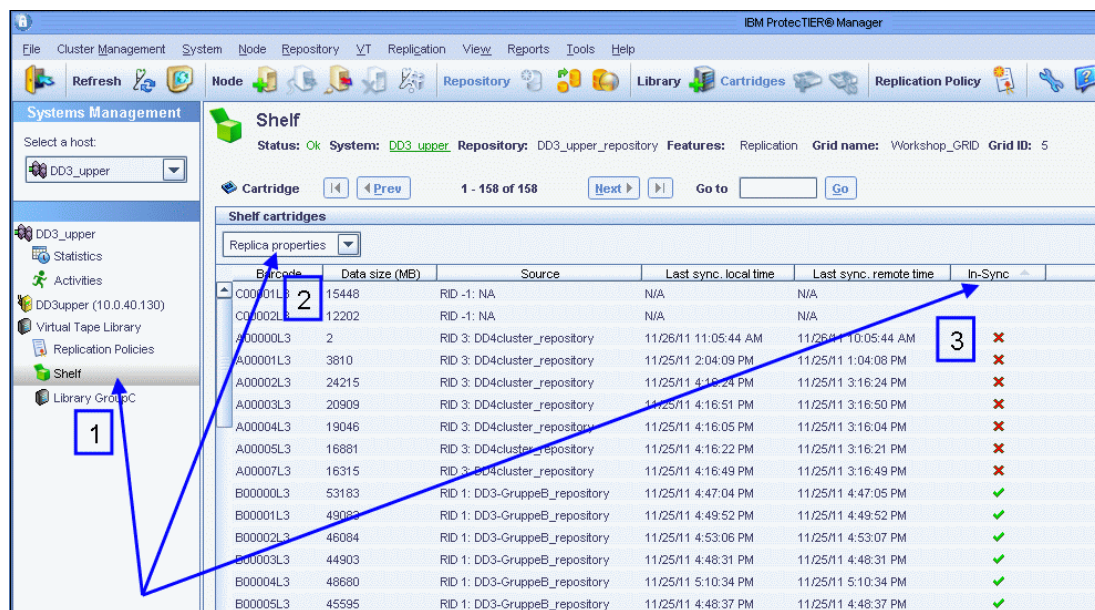


Figure 22-1 Gauging the replication status with the ProtecTIER Manager GUI

If you prefer to gauge the replication status with the IBM ProtecTIER command-line interface (ptcli) tool, complete the following steps. The ptcli tool is located in the /opt/dtc/ptcli directory on each ProtecTIER node. These commands must query the replication target machine (your hub machine) to successfully work.

1. To use the ptcli tool, first create a ptcli profile by running the following command:

```
./ptcli -p ptadmin
```

Using ptcli: The ptcli tool requires being an ptadmin user. As a ptadmin user you can run commands to gather information for gauging of the replication completion status. The ProtecTIER product still “obeys” the single-logon rule.

You cannot run these commands if an active ptadmin user is already logged on. You can use a **force** option for the ptcli command to end an existing active session. Be careful when you use the **force** option because issuing these commands in a script can affect an active admin session.

2. Generate a fresh copy of the inventory statistics by running the following command:

```
./ptcli InventoryRefresh --login ptadmin --ip 10.0.40.120 -force
```

3. Query for all out-of-sync cartridges by running the following command:

```
./ptcli InventoryFilter --login ptadmin --ip 10.0.40.120 --querytype replica --query in_sync=false --output /tmp/dirty_carts.csv
```

The output that is generated in the /tmp/dirty_carts.csv file lists all cartridges with pending replication tasks. Should the command in step 3 create no output, this situation indicates that the hub is synchronized with all spokes.

22.4 Disaster recovery operations with FSI

You can use the ProtecTIER product to configure replication policies to replicate file system's directories and all objects that are contained in these directories recursively to remote ProtecTIER repositories without any disruption to the operation of the file system as a target for backup.

Defining up to 64 source directories and up to three remote ProtecTIER destinations per replication policy is possible. The replicated data in the remote destination can be easily used to restore data in the case of a disaster recovery, or in the case of a disaster recovery test (without any interruption to the backup and replication procedures).

Be sure to enable the ProtecTIER product to supervise all the changes that are made to a directory, or to a set of directories, that are defined in a replication policy. Therefore, you should not disable a replication policy unless this policy is no longer considered relevant. If there is a scheduled maintenance of the network that is used for replication, it is possible (though not mandatory) to *suspend* the replication to a specific destination.

Suspend enables the ProtecTIER product to continue supervising all of the changes made; it does not attempt to send the replication data through the network for the time that is defined by the suspend operation (the suspend operation is limited in time, you can suspend replication activity on a policy for up to 72 hours).

In the case where a policy was disabled for some reason, a new replication destination directory (RDD) must be defined to re-enable the policy. The ProtecTIER product does not need to replicate all of the data from scratch if the old RDD was not deleted. It needs to create only the new structure and metadata in the new RDD. Therefore, you should not delete the old RDD until a new cycle of replication to the new RDD is complete.

The following list summarizes the information:

- ▶ File system replication is a major feature of the ProtecTIER FSI project. It provides an inherent DR ability for ProtecTIER FSI customers.
- ▶ As opposed to VTL replication, FSI must replicate file's data and metadata, and the directory tree structure.
- ▶ FSI replication enables directory replication capability across repositories.
- ▶ FSI replication provides high availability in accessing backed up data across sites.
- ▶ FSI structural changes are replicated one by one to maintain the order of changes.
- ▶ Data replication is separated from structural replication to prevent delaying the order of changes.
- ▶ Upon a write operation, a file closure creates two replication triggers, one for data and one for structure.
- ▶ Structural changes create a structural trigger to replicate the change.
- ▶ Renaming a directory that is included in a replication policy is not allowed and results in a client error.
- ▶ Manual replication uses FSI snapshot to efficiently perform the replication.
- ▶ FSI replication policies are defined on a set of directories in the source repository and target directories in the RDD.

22.4.1 Replication destination directory (RDD)

The RDD is a specially designated directory under the root of a file system at the DR site.

- ▶ The RDD is allocated and flagged upon policy creation to ensure that only one replication policy replicates to the RDD.
- ▶ After an RDD is allocated by a source policy, it cannot be used by any other policy.
- ▶ A replication policy is first created with empty RDDs.
- ▶ The source directories might not be empty when you create a policy, so part of them might not be synchronized at the destination until changes to it are applied.
- ▶ On FSI replication, a destination file or directory either exists in its synchronized state or does not exist there at all.
- ▶ Under no circumstances should a file or directory exist in the destination in an unsynchronized state. This is a crucial assumption of the design.
- ▶ Replication can be ran manually, either on a policy as a whole, or a specific directory in it.

22.4.2 ProtecTIER FSI cloning

Cloning in FSI enables you to perform DR tests at a replication destination repository while continuing replication of the original data from the source repository to the destination repository. To perform a DR test, a space-efficient clone of a read-only replication destination directory can be created on the destination repository that is read/write accessible. All existing data in the repository remains accessible during cloning. If backup data is written to the clone, the relevant modified section is written as new data on the repository.

To clone an RDD, complete the following steps:

1. From the Systems Management view, click **Replication** → **Replication Directory** → **Clone a replication directory**. The Clone a replication directory window opens.
2. In the Source directory pane, select the file system from which to clone the replication directory.
3. Click **Browse** to search for and select the directory name of the replication directory you want to clone.
4. In the Target directory pane, select the file system to which to clone the replication directory.
5. Click **Browse** to select the new directory path and enter the new directory name to which the contents of the replication directory are to be copied.
6. Click **Clone**. A message displays that it might take time for all of the directory contents to be copied.
7. Click **Yes** to confirm and begin the operation. The Cloning progress window displays the source and destination paths, the total number of files to be cloned, the start time of the operation, and the number of files that are cloned per second.
8. Click to hide the Cloning progress window.

Note: A clone cannot be deleted if it is not empty.

22.4.3 Preferred replication practices for FSI

Several preferred practices exist for network replication in combination with ProtecTIER FSI.

File system naming

Choose meaningful names when you create file systems and network shares on the ProtecTIER server. Consider using the same name for the file systems that you create and the shares that you create in these file systems. Consider some obvious scheme to be able to re-create the relationship between file shares and file systems.

Backup server separation

Each backup server output should point to a separate directory that is configured in the ProtecTIER FSI. Because network replication policies are established between the source and destination at the directory level, creating a directory per backup server ensures consistency of individual backup servers. You can use this configuration to individually configure network replication for only a subset of your backup servers.

22.5 Entering ProtecTIER DR mode

Keep the process for entering DR mode simple to avoid potential complications and human errors when you handle an emergency situation. For VTL, the ProtecTIER DR wizard in ProtecTIER Manager at the DR site facilitates the initiation of DR mode.

As shown in Figure 22-2, to enter ProtecTIER DR mode, complete the following steps:

1. From the Systems Management view, click **Replication** → **Replication Disaster Recovery** → **Enter DR Mode**.
2. Choose the repository that you want to enter DR mode and confirm the selection.ok

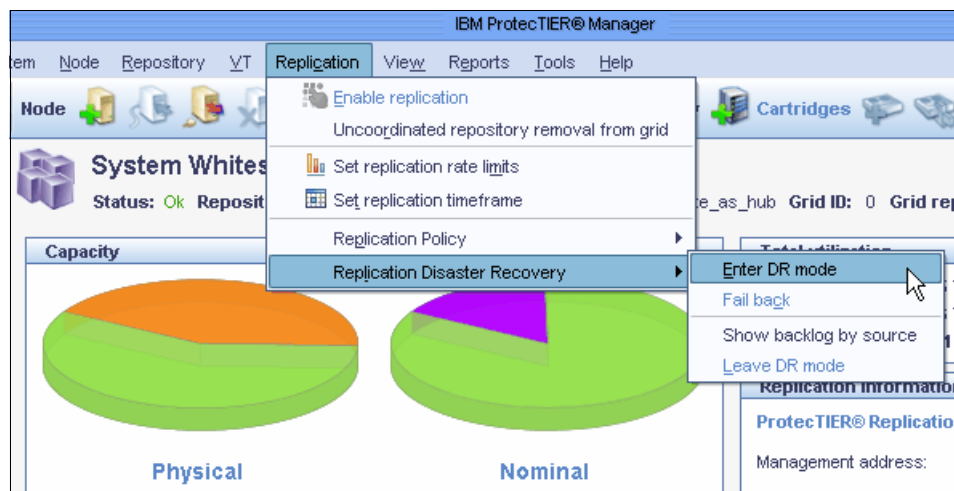


Figure 22-2 Entering ProtecTIER DR mode

An automatic procedure is run that blocks incoming replication to the DR site from the source selected.

Important: By entering and activating ProtecTIER DR mode, you have all the latest data visible at the remote site. However, the backup application must be restored manually or by using prepared semi-automatic procedures and scripts. You must be familiar with this backup application recovery process.

22.5.1 Working at the disaster recovery site

In VTL, by default, replicated cartridges at the DR (hub) repository that were created at a primary site (spoke) are in read-only mode. If it is necessary to run local backups at the DR site, then use new locally created cartridges.

22.5.2 Inventory command options for a VTL disaster recovery scenario

The inventory command options filter cartridges in a ProtecTIER repository by using various criteria. Also, these options can be used to move cartridges that match a certain criteria.

Disaster recovery with the ptcli

Use the ptcli tool command-line interface (CLI) to manage your ProtecTIER servers and replication grid. You can perform the following tasks:

- ▶ Determine the consistency point at the hub when all cartridges from a specific spoke were fully synchronized (replicated).
- ▶ Automatically move cartridges in bulk from a shelf to a library and back.
- ▶ Determine the cartridges that were created at the hub during DR mode and then move them to the shelf so they can be replicated back to the spoke.

Consider the following information when you use the ptcli commands outside of DR mode:

- ▶ The ptcli snapshot might take up to 15 minutes to be created and populated.
- ▶ The snapshot is a static one, so it reflects the status of all cartridges only at the point in time it was taken.

A DR scenario and example procedure

Assume that a disaster occurs at a spoke when replication is running and a DR condition for the specific spoke is declared at the hub. You would complete the following steps:

1. See the DR site (the hub) to determine when the last full backup occurred and recover the data from the DR to the primary site.
2. Use ptcli to produce a list of cartridges that were synchronized at the time of the disaster and a list of which cartridges were not synchronized at the time the last full backup at the spoke was completed.
3. Decide which cartridges to use at the DR site and use ptcli to move them (all or some) from the shelf to the library.
4. Save your results in a .csv file by using the **-output** command switch.

Consider the following information regarding the output file:

- ▶ The output .csv file with results can be used as an input to a CLI **move** command. The .csv file is editable, and you can remove rows (each line represents a cartridge).
- ▶ You can create bar code files and use them as an input source for a **move** command.

22.5.3 Commonly used disaster recovery queries

Note: For more information about using the ptcli tool, see the command-line interface topic in *IBM System Storage TS7600 with ProtecTIER Version 3.3*, SG24-7968.

The following list describes several commonly used DR queries that you can run:

- ▶ To create a snapshot of the current replication status, run the following command:

```
./ptcli InventoryRefresh --ip xxx.xxx.xxx.xxx -login file
```

Important: Creating a snapshot must be done before you run any other queries. Before you begin to filter the cartridges or to move cartridges by using the CLI, you must create a snapshot of the cartridges by running **InventoryRefresh**. This snapshot captures the current properties of all cartridges. Any filter or move operation is run based on the content of the snapshot. Moving a cartridge before you create a snapshot might fail if the snapshot is not up-to-date (not taken directly before the move).

- ▶ To list all in-sync cartridges, run the following command:

```
./ptcli InventoryFilter --ip xxx.xxx.xxx.xxx --querytype replica --query  
"in_sync = true" -login file -output /tmp/not_dirty_carts
```

- ▶ To list all unsynchronized cartridges marked with dirty bit, run the following command:

```
./ptcli InventoryFilter --ip xxx.xxx.xxx.xxx --querytype replica --query  
"in_sync = false" -login file -output /tmp/dirty_carts
```

- ▶ To list cartridges synchronized with the destination at a certain time range on the source, run the following command:

```
./ptcli InventoryFilter --ip xxx.xxx.xxx.xxx --querytype replica --query  
"source_time_for_last_sync_point > datetime('2009-11-13 08:00:00')" -login file
```

- ▶ To list all cartridges that are replicated to repository 18 in grid 1, run the following command:

```
./ptcli InventoryFilter --ip xxx.xxx.xxx.xxx --querytype origin --query  
"destination_repository_id = 18 and destination_grid_id = 1" -login file
```

- ▶ To list all cartridges in bar code range BR0300 - BR0330, run the following command:

```
./ptcli InventoryFilter --ip xxx.xxx.xxx.xxx --querytype all --query "bar code  
> BR0300 and bar code < BR0330" -login file -output barcodes_file
```

- ▶ To move all synchronized cartridges to the shelf, run the following command:

```
./ptcli InventoryMoveFilter --ip xxx.xxx.xxx.xxx --querytype replica --query  
"in_sync = true" -destination shelf -login file
```

22.5.4 Returning to normal operations

This section explains the meaning of failback and failback policy, and provides the steps for using the ProtecTIER Manager failback wizard.

Failback

Failback is the procedure for returning visibility to the primary site and replicating any new data that is created at the DR site to the original (or replaced) primary repository. You use failback to return to the original working mode of the primary site.

Note: The failback procedure can be initiated only while in DR mode. A mandatory step is to perform the failback process *before* you leave DR mode at the remote site.

Failback policy

The DR site uses a *one-time replication policy* during failback that places a hold on all cartridges that are replicated to the primary site. This special policy also transfers the ownership of the relevant cartridges (created at the DR site when the primary site was down) to the primary repository or to its replacement in case the original primary repository is no longer functioning.

Figure 22-3 shows the ProtecTIER Manager failback wizard.

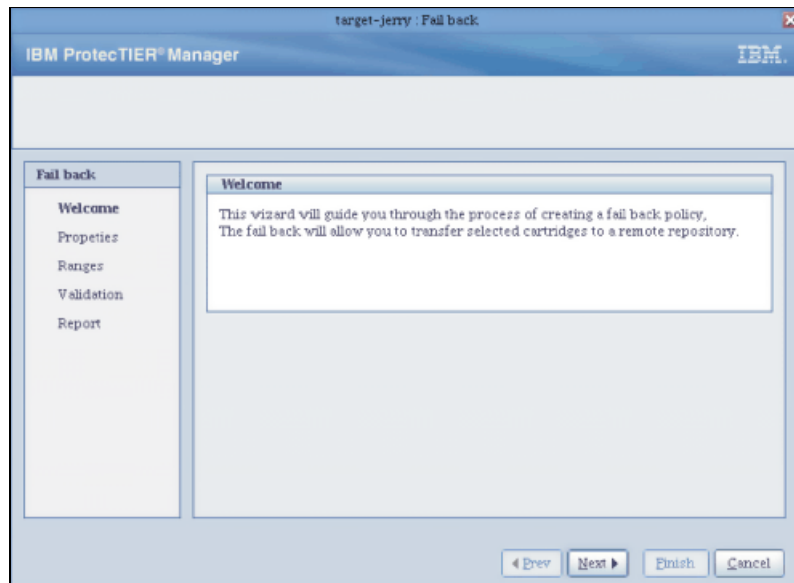


Figure 22-3 ProtecTIER Manager failback wizard

ProtecTIER Manager failback wizard

Initiate the failback process by using the ProtecTIER Manager failback wizard, as follows:

1. Cartridges must be ejected out of the library to the shelf at the DR site. These cartridges include both the original replicated cartridges and any new cartridges that are created during DR operations.
2. Define a policy with all the cartridges that must be transferred to the primary site. This policy can be run only manually. The system log ignores runtime events for this policy.
3. Approve the execution of the policy. (This approval is a manual execution of the policy in the VTL.)
4. Close the failback wizard. The system provides information about the number of pending, running, and completed cartridges. ProtecTIER Manager presents this information to the user to indicate that the failback process is complete.
5. Delete the policy after the failback task completes.

22.6 The backup application catalog

Figure 22-4 illustrates a typical backup and DR environment using the ProtecTIER product. The backup application environment is straightforward. The backup application servers are connected to storage devices (disk, real tape, or virtual tape). Every action and backup set that the backup servers process is recorded in the backup application database or catalog. The catalog is at the heart of any recovery operation. Without a valid copy of the database or catalog, restoration of data is difficult, sometimes even impossible.

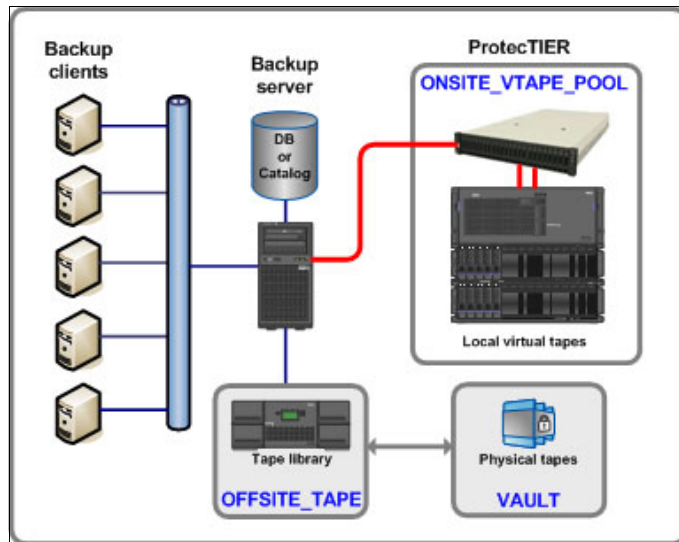


Figure 22-4 Typical backup and DR environment using the ProtecTIER product

The ProtecTIER server provides a virtual tape interface to the backup application server, which you can use for the creation of tapes, represented by ONSITE_VTAPE_POOL in Figure 22-4. The client can also maintain another tape library to create real tapes to take off-site, called OFFSITE_TAPE in Figure 22-4.

ONSITE_VTAPE_POOL is where most client recoveries and restores come from. The key advantage of this architecture is that restoration occurs much faster because the data is coming from the ProtecTIER disk-based virtual tape rather than from real tape.

Your IBM Spectrum Protect (formerly Tivoli Storage Manager) database or the backup server catalog, in general, should be backed up only after all backup sets are complete for a specific backup server. The database or catalog should be backed up in the ProtecTIER FSI-specific directory for that backup server.

Important: ProtecTIER FSI cannot be used to store the active version of the catalog that is constantly being updated. Do not place your IBM Spectrum Protect database or the backup server catalog directly in to the FSI. Only backups of supported applications are supposed to be writing to ProtecTIER FSI. For more details about IBM Spectrum Protect setup and preferred practices in an FSI environment, see 13.3, “IBM Spectrum Protect: FSI” on page 200.

22.6.1 ProtecTIER replication with IBM Spectrum Protect

The IP replication function of ProtecTIER provides a powerful tool that you can use to design a robust DR architecture. Because of data deduplication, you can now electronically place backup data into a vault and use less bandwidth at the same time. The ProtecTIER IP replication functionality can be used in an IBM Spectrum Protect environment, as shown in Figure 22-5.

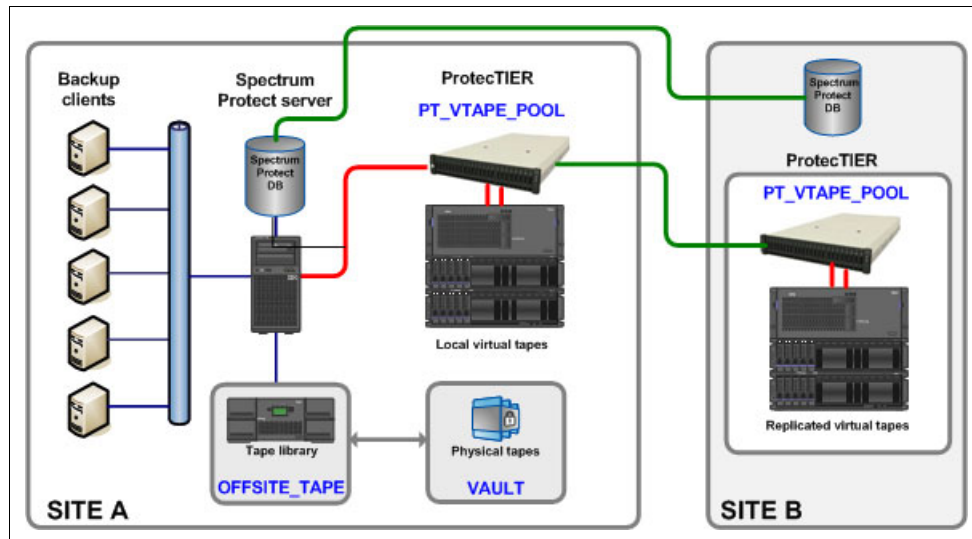


Figure 22-5 ProtecTIER replication and IBM Spectrum Protect

Figure 22-5 shows that the user chooses to replicate all virtual tapes in the PT_VTAPE_POOL to site B. The user also backs up the IBM Spectrum Protect database to virtual tapes, which are also replicated to Site B. In this way, if a disaster occurs, you can restore the IBM Spectrum Protect server in Site B, which is connected to a ProtecTIER VTL. The ProtecTIER VTL contains the IBM Spectrum Protect database on virtual tape and all of the client ACTIVE files on virtual tapes.

Note: For more information about ProtecTIER running with IBM Spectrum Protect, see Chapter 13, “IBM Spectrum Protect” on page 187 and see *Harnessing the Power of ProtecTIER and Tivoli Storage Manager*, SG24-8209.

22.6.2 Recovering the backup application catalog

Use one of the following ways to obtain a copy of the catalog at the DR site:

- ▶ From a catalog backup on a virtual cartridge that is replicated to the DR site.
- ▶ From disk-based replication, or by other means.

If the catalog is backed up to a virtual cartridge, check if the DR site on which this cartridge appears is In-Sync with the primary site. If the cartridge is not In-Sync, you need to compare the cartridge's last synchronization time with the time of the last full backup.

To recover the backup application catalog from a backup on a virtual cartridge, you must work with the replicated cartridges on the hub to get an updated copy of the catalog back to the DR site. From the Systems Management window, select the **Replica** properties view on the Cartridges tab and use the following guidelines for each cartridge before you run the procedure for recovering the catalog.

Important: The procedure for recovering the selected catalog backup depends on the backup application and should be documented in the backup application's official documentation.

If the cartridge is replicated, either a red X or a green check mark appears in the In-Sync column. If the In-Sync column has a green check mark, then nothing further needs to be verified and this cartridge is valid for recovery.

If the cartridge is not marked In-Sync, see the Last sync time column. This column displays the last time each cartridge's data was fully replicated to the DR site. The cartridge marked with the most recent last sync time date should be used to recover the backup application catalog.

The sync time is not updated only when replication for this cartridge is finished, but also *during* replication.

Assessing the cartridges' status and synchronizing with the catalog

After the DR site backup application server is recovered, the user must review the status of the replicated cartridges. This review ensures that their replication is consistent with the backup catalog or database. This section explains the process for assessing cartridges' status on the DR site and synchronizing the backup application catalog with the cartridges.

Before you run a restore for disaster recovery, you must verify that the list of associated cartridges are marked as In-Sync with the primary site. Otherwise, an earlier full backup image must be used for recovery. The easiest way to determine the time of the last full backup is if you have a specific time each day where your replication backlog is zero (that is, there is no pending data to replicate and backups are not running). If not, you can assess the cartridges by recovering the backup application catalog and scanning it to find the last full backup where its associated cartridges completed replication.

Recovering the data

After the data is recovered, scan the backup application catalog and search for the full backup image you want to recover:

- ▶ Get the start and end backup time of the full backup image.
- ▶ View the list of cartridges that are associated with this full backup.

Run the **PTCLI inventory filter** command to filter the cartridges according to the following properties:

- ▶ In-Sync
- ▶ Last update time
- ▶ Last sync time

All the cartridges that are marked as In-Sync are valid for recovery. For those cartridges not marked as In-Sync, compare the last update time, which represents the last time that the replica was updated and the last sync point destination time. If the last update time is less than or equal to the last sync point destination time, the replica cartridge has consistent point-in-time. Otherwise, the cartridge is incomplete, or in transit. If the cartridge has consistent point-in-time, ensure that this time stamp is larger than the full backup image end time. This time stamp indicates that the cartridge contains all the required data for this recovery operation. Otherwise, the user must use a previous full backup image for recovery.

You might have a case where the cartridge sync point is after the backup start time, but before the end of the backup. This situation might happen in cases where replication is working in

parallel to the backup. If the backup has many cartridges, the first cartridges might finish replicating before the backup ends and get a synchronization point earlier than the backup end time.

If the Last sync time flag on one (or more) of the cartridges indicates a time later than the backup start time, but earlier than the backup complete time, those cartridges need further inspection. Scan the backup application catalog for each of those cartridges and get the backup start time and the backup complete time.

If the Last sync time flag on all the cartridges indicates a time later than the backup complete time, your backup image was fully replicated.

Important: When you process the cartridge list to find a complete set of DR tapes, you must track the date and time discrepancies. Compare the date and time values of the source master backup server and the source ProtecTIER system. The destination environment might be in a different time zone or might be set to the incorrect date and time.

Use the *source* date and time, rather than the *destination* sync time, when you compare cartridge states to the backup catalog or database. The destination sync time should be used only to determine which cartridges are complete.

There might be a time difference between the source backup server and the source ProtecTIER server. Your administrator should be aware of the discrepancy, measure it regularly, and communicate the delta to your DR administrator or operators.

For example, if the backup server is two hours behind, a cartridge might have a sync time that precedes its backup complete time. If there is uncertainty about the time differences, compare the nominal size of the cartridge to the catalog or database value as an additional (not a substitute) layer of verification.

22.6.3 IBM Spectrum Protect reclamation and housekeeping

When you use IBM Spectrum Protect and ProtecTIER replication, configure your IBM Spectrum Protect reclamations to keep the IBM Spectrum Protect reclamation process from running during the ProtecTIER replication window. For maximum efficiency, choose dedicated windows for ProtecTIER replication and backup server housekeeping.

22.7 Single domain and multiple domains

Backup application software can be set up in many topologies. From the ProtecTIER replication standpoint, there are two general setup methods for these environments:

- ▶ Single domain
- ▶ Multiple domain

The backup application catalog or database has an entry for each cartridge that is used for backup:

- ▶ Date when the backup was performed
- ▶ List of files that are associated with the backup
- ▶ Retention period
- ▶ Other backup application-specific information

The backup application supports one catalog or database per backup server instance. In many cases, the primary and remote (secondary DR) sites have two separate backup servers, each with its own database or catalog. To efficiently read replicated cartridges at the remote site, the remote backup server needs access to the actual catalog or database of the primary backup server, or an exact copy of it.

22.7.1 Single domain environment

In a single domain environment, the same backup application catalog (or database) is shared across the separate primary and secondary sites. In these environments, the catalog is always updated in real time on the locations of the cartridges (physical and virtual). For example, this type of environment is more commonly used with Symantec NetBackup (NetBackup) and does not work with most deployments of IBM Spectrum Protect.

22.7.2 Multiple domain environment

A multiple domain approach is more widely used. The backup application does not share a catalog between the primary (local) and secondary (remote DR) sites. This scenario is the most common with IBM Spectrum Protect environments. In this type of deployment, each backup server, in both the primary and the secondary locations, has its own backup catalog.

22.8 Deploying replication with specific backup applications

In ProtecTIER, VTL mode is generally the preferred method of replication to simulate the procedure that is used with physical cartridges. Implement the time-frame mode of operation so that for every 24-hour cycle, there is a backup window, and then a replication window.

Ensure that enough bandwidth and time allotted are available so that there is no overlap and no replication backlog. Consider the following suggestions for a typical operational flow:

- ▶ Perform regular daily backups to the ProtecTIER system during the defined backup window.
- ▶ The system should be set up so that replication starts and finishes before the next backup cycle starts.
- ▶ The user should have a complete and easily recoverable set of their latest daily backup, including the backup application catalog image.
- ▶ If a disaster occurs, the user can revert to the last completed set of backups, so the recovery point objective (RPO) is in the 24-hour window, which is typical for a service-level agreement (SLA).

22.8.1 Recovery point objective (RPO)

When you design a ProtecTIER replication environment, one of the most important questions to consider is “What is the RPO?” How much lag time is acceptable for a backup that is written to virtual tape in Site A, to be replicated to Site B?

Tape-based DR

The RPO for tape-based DR is typically 24 hours. For example, consider a typical user case in which backups begin at 6 PM on Monday evening and the tape courier picks up the box of physical tapes at 10 AM Tuesday morning for transport to the vault. Therefore, on a typical

day, a 14-hour delay can occur between the time the first backup begins and when the data is safely offsite.

However, if a disaster occurs before the courier arrives, the customer recovers the applications from the Sunday replication workload, which is a day behind, providing a 24-hour RPO.

ProtecTIER replication

With ProtecTIER replication, it is possible to get the backups offsite almost immediately (if there enough bandwidth is available). Because the ProtecTIER product is always working in the backup application, the RPO typically remains 24 hours.

22.9 IBM Spectrum Protect

This section introduces various replication scenarios for when a disaster occurs at different stages of the replication process. IBM Spectrum Protect is a typical representation of a multiple domain environment, so the valid and most recent IBM Spectrum Protect server database backup must be available at the DR site.

Important: The valid and most recent IBM Spectrum Protect server database backup is a critical component of the recovery process at the disaster site.

As of Tivoli Storage Manager V6.1 ensure that you always have the file with all the records of Tivoli Storage Manager database backups (volhistory). If you do not have a copy of the volume history file, you cannot recover the Tivoli Storage Manager server database from your database backup.

22.9.1 Scenario 1: Replication complete

IBM Spectrum Protect enables the data replication to begin in a scheduled time frame after the backup cycle is completed, as shown in Figure 22-6 on page 414.

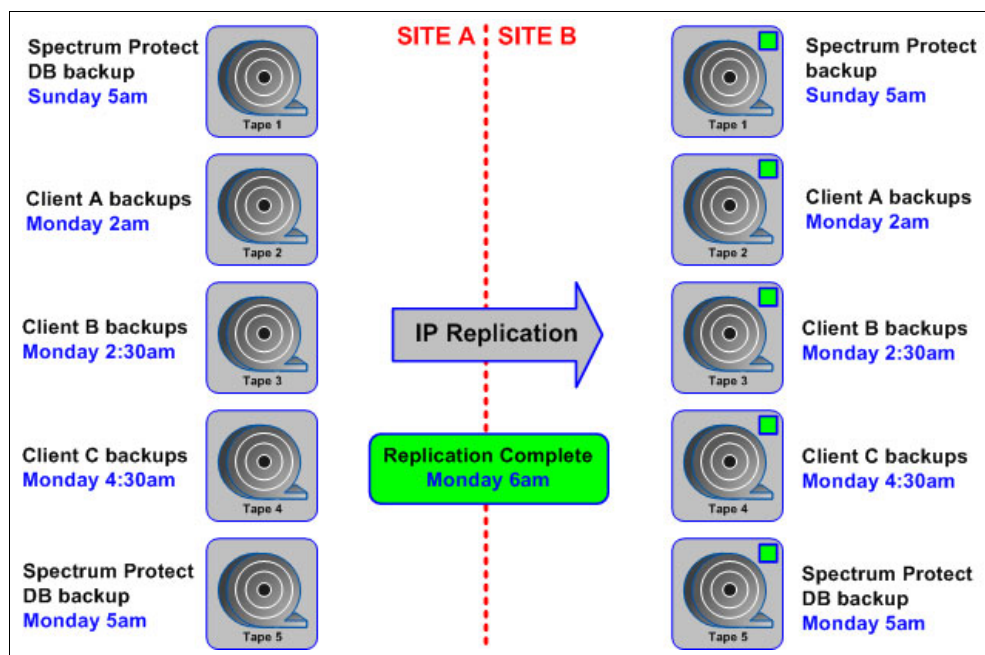


Figure 22-6 Scenario 1: replication complete

For example, assume that the replication window ended at 6 AM and the disaster occurred at 6:15 AM, 15 minutes after the replication cycle completes. (The status of all virtual cartridges is 100% replicated when the disaster event occurred.)

The user starts the IBM Spectrum Protect server by using the IBM Spectrum Protect database backup that occurred on Monday 5 AM to tape 5. The IBM Spectrum Protect database has knowledge of all the prior tapes, and restoration immediately begins from Tapes 2 - 4, which hold the most recent client backups.

22.9.2 Scenario 2: Replication incomplete

Suppose that the disaster event occurred at 5:15 AM before the nightly replication cycle completes. This scenario means that some of the backups from last night are not 100% replicated to Site B. For example, assume that the replication cycle ended at 6 AM, but the disaster event occurred at 5:15 AM Monday morning. Tapes 4 and 5 were not replicated when the link went down.

In this case, because the IBM Spectrum Protect database is not entirely replicated, the user must restore from the last available IBM Spectrum Protect database backup on Tape 1 from Sunday at 5 AM. Because Tapes 2 and 3 were also created after Sunday 5 AM, they are not in the IBM Spectrum Protect database and cannot be used. Therefore, the information that is on tapes 2 and 3 is lost.

22.9.3 Scenario 3: Auditing tapes and fixing the IBM Protect database

Another possibility is when the most recent IBM Spectrum Protect database virtual tape is replicated, but not all the associated client backup tapes completed their replication before the disaster event occurred. For example, the disaster occurred 30 minutes before the anticipated replication cycle completion. In this case, Tapes 1, 2, 3, and 5 replicated, but Tape 4, because of the size of the backup data set that is stored on it, did not finish replication before the disaster.

In this case, the IBM Spectrum Protect server is restored by using the IBM Spectrum Protect database that is stored on Tape 5. However, because Tape 4 did not complete the replication, it must be audited to fix IBM Spectrum Protect database records that are related to the missing data on tape. Use the following IBM Spectrum Protect command for auditing Tape 4:

```
audit volume Tape4 fix=yes
```

The audit and fix must be performed for every tape that was not fully replicated when the disaster occurred. If the number of tapes that need this corrective action exceed the reasonable limits to run each job manually, you can consider auditing the pool that contains the inconsistent volumes. Use the following command to have IBM Spectrum Protect server audit every tape volume that belongs to the affected pool:

```
audit stgpool PT_VTAPE_POOL fix=yes
```

Auditing hundreds of tape cartridges can take a significant amount of time to complete (healthy cartridges are also audited). You must find a compromise between the time that is processed by auditing all of the tapes at the same time, and running the audit job for each inconsistent tape separately.

22.9.4 Reclamation considerations

Another important consideration when you use ProtecTIER IP replication in an IBM Spectrum Protect environment is the efficiency that reclamation causes. Reclamation is the IBM Spectrum Protect process that frees up space on tapes, and returns empty tapes to the scratch pool. Reclamation is accomplished by deleting records in the IBM Spectrum Protect database that are related to the expired data on tapes and moving the remaining valid data to other tapes, to efficiently use tape resources.

All reclamation-related data movement is recorded in the IBM Spectrum Protect database. To restore the accurate tape environment in the event of a disaster, all the database records related to the data movement must be replicated.

If the disaster occurs when reclamation is running and the database backup that is available for restore does not reflect the status of data, an audit and fix of the volumes might be required. However, this audit process fixes only the inconsistent records in IBM Spectrum Protect database, but does not make the latest primary backup version of files available at the disaster site, unless they were replicated.

You can suppress the effect of the reclamation by using the **REUSEDELAY** parameter (valid for sequential-access storage pool, which includes virtual tapes). The **REUSEDELAY** parameter specifies the number of days that must elapse before a volume can be reused or returned to scratch status after all files are expired, deleted, or moved to another volume.

Tip: Set the **REUSEDELAY** period to the same value as the retention period of your IBM Spectrum Protect database backups. Only then, can you safely restore IBM Spectrum Protect to any available point in time.

Pending state

When you delay the reuse of such volumes and they no longer contain any files, they enter the pending state. Volumes remain in the pending state as specified by the **REUSEDELAY** parameter.

Delaying the reuse of volumes can be helpful for disaster recovery under certain conditions. When files are expired, deleted, or moved from a volume, they are not physically erased from the tape. Only the database references to these files are removed. Therefore, the files still

exist on a sequential volume if the volume is not immediately reused. This prevents a situation where the reclamation ran after the last IBM Spectrum Protect database backup to virtual tape and reclaimed tapes are replicated. If the disaster occurs at a point after the reclamation ran, you are forced to restore from the less current database backup. If so, some files might not be recoverable because the server cannot find them on replicated volumes. However, the files might exist on volumes that are in the pending state. You might be able to use the volumes in the pending state to recover data by completing the following steps:

1. Restore the database to a point in time before file expiration.
2. Use a primary or copy volume that is rewritten and contains the expired file at the time of the database backup.

If you back up your primary storage pool, set the **REUSEDELAY** parameter for the primary storage pool to 0 to efficiently reuse primary scratch volumes. For the copy storage pool, you must delay the reuse of volumes while you keep your oldest database backup.

Disabling automatic reclamation

You might need to disable automatic reclamation by changing the reclaim parameter of the sequential-access attributes by running the following command:

```
update stgpool <STG POOL NAME> reclaim=100
```

Consider disabling reclamation altogether if the user deployed the ProtecTIER product with small virtual tape cartridge sizes (for example, 50 GB or less). Statistically, there is less data to reclaim on a smaller volume than on a larger one. Everything on a smaller volume is likely to expire together, causing the tape to go immediately to the scratch pool.

The following list summarizes the information:

- ▶ Ensure that catalog and database backups are performed to virtual tape and replicated along with the daily workload each day. Perform a database backup and replicate the backup data at the end of each backup cycle.
- ▶ Create a separate tape pool for database backups.
- ▶ Consider adjusting the IBM Spectrum Protect reclamation processing to ensure that actions are in sync with the replicated database, which includes setting **REUSEDELAY** to provide a 2-day or 3-day delay in the reuse of tapes that are reclaimed.
- ▶ Consider deploying more network bandwidth to decrease the backlog. Database synchronization becomes a more serious issue if the backlog exceeds one backup cycle (typically 12 - 24 hours).

22.9.5 Determining what is available for restoration at the DR site

To determine what is available for restoration at the DR site, answer two questions:

- ▶ Which database copy at the DR site is optimal for recovery?
- ▶ What cartridges at the DR site are valid (replication completed)?

Which database copy is optimal

Before you run a restore for DR, verify that the list of associated cartridges are replicated to the DR site. Otherwise, an earlier full backup image must be used for recovery. The time of the last full backup is the specific time each day where your replication backlog is zero (there is no pending data to replicate).

If you do not know the specific time each day when your backlog is zero, assess the cartridges by recovering the backup application catalog and scanning it to find the last full

backup where its associated cartridges completed replication. There are two primary methods to obtain a copy of the catalog at the DR site:

- From a catalog backup on a virtual cartridge that is replicated to the DR site
- From disk-based replication

If the catalog is backed up to a virtual cartridge, then use the cartridge view of the library in ProtecTIER Manager to examine each of the cartridges that are used for catalog backup and find the most recent sync dates that are marked on the cartridges. If there are multiple backup copies, find the latest backup that completed replication. To recover the backup application catalog from a backup on a virtual cartridge, you must work with the replicated cartridges to get an updated copy of the catalog to the DR site.

Each cartridge contains a time stamp of a last sync time, which displays the last time that the cartridge's data was fully replicated to the DR site. The sync time is updated during the replication process, not just when the replication for this cartridge is finished. The cartridge marked with the most recent last sync time must be used to recover the backup application catalog.

Which cartridges at the DR site are valid for restore

When the IBM Spectrum Protect server at the DR site is recovered, review the status of the replicated cartridges to ensure that their replication is consistent with the IBM Spectrum Protect database. Use the available ProtecTIER Manager system Replicated Cartridges Status Report.

To eject a cartridge from a library, run the following command:

```
CHECKOUT libvolume
```

Here is an example:

```
TSM:TUCSON1>CHECKOUT LIBVOL <name of library> REMOVE=BULK FORCE=yes  
CHECKLABEL=YES VOLLIST=<volume1,volume2,...volume3>
```

To inject or import (insert) a cartridge to a library, run the following command:

```
CHECKIN libvol
```

Here is an example:

```
TSM:TUCSON1>CHECKIN libvol <name of library> search=BULK checklabel=bar code  
status=SCRATCH WAITTIME=0
```

22.10 Symantec NetBackup deployment with ProtecTIER replication

This section describes the usage of the IBM ProtecTIER IP replication system in a NetBackup environment and describes the ramifications of possible scenarios that are related to DR.

22.10.1 Scripting the inject/eject commands

Vault and inject/eject commands can be scripted to run periodically on the backup application host (see Example 22-2). These commands trigger automatic cartridge movement from the import/export slot to the library whenever the relevant cartridge is in the import/export slot. This process ensures free import/export slots.

Example 22-2 Example script of inject/eject commands

```
#!/bin/csh
while (1)
  vltinject myVault
  sleep 600
end
```

22.10.2 Scripting the inventory commands

Scripting the inventory command is *not* recommended because it scans the robot and therefore might take a long time to complete on libraries with many cartridges.

22.10.3 Setting up NetBackup for backup and restore

NetBackup deployments typically use a schema of weekly full backups and daily incremental backups. Two types of incremental backups are available:

- ▶ Cumulative: Backs up everything since the last full backup.
- ▶ Differential: Backs up everything since the last backup.

22.10.4 Setting up NetBackup for disaster recovery

When you set up NetBackup for disaster recovery, consider several key issues:

- ▶ NetBackup architecture: Does the NetBackup domain span across the primary and DR sites or are they two separate domains? This is a key step to understand and has strong implications in DR.
- ▶ Classification of clients, recovery time objective (RTO): When a company plans for DR, each server is given an RTO that is dependent on the importance of its application and the associated data to the business. Servers with short RTOs (typically less than 24 hours) generally do not use backup systems for DR. These servers typically use clustering, volume mirroring, or some form of data replication to maintain business continuity. Servers with RTO greater than 24 hours tend to use tapes for DR. Servers are then prioritized into recovery bands of 24, 36, 48, or 72 hours, depending on business requirements.

DR and production servers: Typically, only production servers are set up for DR. Test and development servers are out of scope for DR. However, the ProtecTIER product makes DR protection affordable for all applications in any environment.

- **Classification of Clients (RPO):** Running alongside RTO is an RPO. RPO is the point in time to which the server must be recovered. For most servers that use a tape DR scenario, the RPO is the point of the last complete backup before the disaster. For example, if a disaster strikes at 9:00 AM, the RPO is the previous nightly backup.

Single domain

To cater to these DR requirements, configure NetBackup with a single domain that spans both sites (NetBackup Clustered).

In this configuration, the master uses host-based replication to mirror the NetBackup databases and a clustering product to manage host failover. If there is a DR event, the NetBackup master's operations can seamlessly fail over to the DR site. As the NetBackup databases are replicated, all of the backup information is known at the DR site and therefore data restores can begin immediately.

22.10.5 Cross-site backups

The two possible options for cross-site backup scenarios are as follows:

- **Connect clients from one site through the Internet Protocol (IP) to media servers on the DR site.**

All backups are then in the DR site library and ready for restore. The primary downside is that large IP pipes are required and backups are limited to the speed of the cross-site network.

- **Stretched Tape storage area network (SAN).**

A local client backs up to a local media server, which then sends the data across the SAN to the DR site. Backups are in the DR site library and are ready for restore. The disadvantage of this option is the large SAN pipes that are required, and the back ups are limited to the speed of a cross-site SAN.

Downside of both options

Because normal backups are now in the DR library, regular restores are slower because the data must come from a DR library.

The following options and possibilities can partially eliminate this negative effect:

- **Turn multiplexing off.** To achieve the best restore performance (to meet RTOs), NetBackup must be configured without multiplexing.
- **Dedicate individual volume pools of RTO tiers or clients.** For optimum restore times (and with sufficient media in libraries), implement individual volume pools per client. In this case, there is no contention between media when you perform restores. In the physical tape environments, where the number of tape drives is limited, this configuration is often impractical.
- **Systems in current production can implement cross-site backups with client backups going to dedicated volume pools,** but this is limited to 30 clients with low RTOs. With separate volume pools, you need separate backup policies per client.

If the NetBackup configuration at the DR site is not in the same domain as the primary site, then a different strategy is required. Because the DR site has no knowledge of the backups, tapes, and so on, that are used by the primary site, you must first get a copy of the NetBackup catalog from the primary site and load it in to the NetBackup master on the DR site. This task can either be done through disk replication or tape backup.

NetBackup catalog backups: NetBackup catalog backups are different from regular backups and need special handling to restore. Not having the catalog available at the DR site means that every tape must be imported to build the catalog, which is impractical and is not considered a viable option. With the catalog in place at the DR site, the tapes can be loaded into the library, the library inventoried, and restores can commence in a short time frame.

22.10.6 ProtecTIER disaster recovery with Symantec NetBackup

Discuss the following key concepts with the NetBackup architects and senior administrators in the user organization:

- ▶ In normal operation, back up to a local VTL.
 - Backing up to a local VTL provides quick backups and quick restores.
 - Because VTL replication is at the cartridge level, and only the deduplicated data is transferred, it reduces the bandwidth that is needed, as compared to traditional cross-site replication/backups.
- ▶ Split servers for DR in to their RTO classifications.

Have servers for DR (usually production servers) split into their RTO classifications and plan for separate volume pools and backup policies. For servers with low RTO requirements, consider individual volume pools and backup policies.
- ▶ Turn off multiplexing (MPX) for all backups that require DR.

Multiplexing is accomplished at either the storage unit level or backup policy level. Disable MPX for all backups that go to the ProtecTIER VTL.
- ▶ Use large fragment sizes.

Fragment sizes are configured at the storage unit level. Large fragment sizes improve the restore performance of whole file systems.
- ▶ Disable storage checkpoints.

Storage checkpoints have an adverse effect on the deduplication ratios.
- ▶ Disable software compression.

Software compression might reduce the efficiency of the ProtecTIER deduplication and affect its factoring ratio.

22.10.7 Single domain versus two separate domains

After your architects and administrators understand the main concepts, they need to decide whether to have one domain that spans both sites or two separate domains.

Single domain approach

With a single domain approach, the same NetBackup catalog is shared across sites and is always updated with the whereabouts of all cartridges.

Replicating cartridges

ProtecTIER replicates cartridges per the policies that are set by the user. Cartridges are copied on to a virtual shelf at the DR site.

Moving cartridges

Cartridges can also be moved by using the replication policy with the visibility switch option so that they are visible to the NetBackup application at the DR site (although the actual data is available to ProtecTIER on both sites). Moving cartridges includes the following functions:

- ▶ Ejecting cartridges: Eject (export) a cartridge from a primary library.
- ▶ Injecting cartridges: Inject (import) a cartridge in to the inventory at the DR site library.

This operation can be set manually or by using the NetBackup vault. Either way, it can be automated from in the NetBackup environment.

Separate (multiple) domains approach

If a separate (multiple) domains approach is used, the items that are listed in this section apply.

Replicating cartridges

ProtecTIER replicates cartridges per the policies that are set by the user. Cartridges are copied into a virtual shelf at the DR site.

Perform a catalog backup to virtual tape at the end of the backup window. Replicate it at the end of each replication cycle to the DR site. This approach ensures that at the end of every day (assuming a 24 hour backup and replication cycle) that the DR site holds a full set of replicated cartridges with a matching NetBackup catalog.

Disaster recovery

If a disaster occurs, consider these options:

- ▶ Your first step is to get the NetBackup catalog restored on the DR site's NetBackup server by restoring the cartridges that contain the catalog.
- ▶ Your second step is to inject the cartridges from the DR shelf at the DR site ProtecTIER server into the library and perform an inventory.

After the NetBackup server is up and running with the DR repository, restores and local backup operations can resume at the DR site. After the disaster situation is cleared and the primary site is back online, the user should use the ProtecTIER failback procedure to move their main operation back to the primary site.

22.10.8 Disaster recovery scenarios

ProtecTIER replication reduces cross-site backup traffic because it replicates only deduplicated data. ProtecTIER replication also improves ease of operation (by enabling simple inject and inventory actions), and if there is a disaster or DR test, makes recovery easy to plan and implement. Deploying the ProtecTIER product in a NetBackup environment makes the business more secure and reduces the burden of NetBackup architects and administrators.

Single domain environment

In a single domain configuration, there are two possible disaster situations.

Clustered configurations with all operations complete

In a clustered configuration where all backups and all replication operations are complete and disaster occurs, no NetBackup recovery action is necessary. The NetBackup catalog at the DR site is up to date. In this case, take the following actions:

- In the ProtecTIER server, move the cartridges from the virtual shelf to the import slots.
- In NetBackup, the library inventory must be refreshed. Choose the option to import tapes.

After the inventory operation is complete, restores and local backups at the DR site can resume.

Clustered configurations with backup complete but replication incomplete

If the configuration is clustered and all backups are complete but the replication operation is incomplete, then the NetBackup catalog database at the DR site is up to date. However, because replication is not complete, roll back to the previous nightly catalog and cartridges set (RPO of one day). After the inventory operation is complete, restores and local backups at the DR site can resume.

Important: When you are working in a single domain NetBackup clustered environment *and* using the visibility switch option in the ProtecTIER server to move cartridges from the primary site directly into a DR site library, the catalog is always up to date.

Multiple domain environment

The following scenarios might occur in a multiple domain environment with NetBackup and the ProtecTIER product.

Stand-alone configurations with all operations complete

If you are working with a stand-alone configuration, and all backups and all replication operations are complete when the disaster occurs, then the catalog database at the DR site is *not* up to date. The NetBackup catalog recovery action is necessary.

Identify the latest backup catalog tape and load (import) it in to the ProtecTIER library at the DR site. After the library is inventoried, begin a standard NetBackup catalog recovery operation. After the recovery operation is complete, restores and local backups at the DR site can resume.

Stand-alone configurations with backup complete but replication incomplete

If you are working with a stand-alone configuration and all backups are complete but the replication operation is incomplete when the disaster occurs, then the catalog database at the DR site is *not* up to date. The NetBackup catalog recovery action is necessary.

Find the NetBackup backup catalog tape from the previous night and load (import) it into the ProtecTIER library at the DR site. After the library is inventoried, begin a standard NetBackup catalog recovery operation. After the recovery operation is complete, restores and local backups at the DR site can resume.

22.10.9 Determining what is available for restore at the disaster recovery site

First, determine which NetBackup database copy at the DR site is valid. Before you run a restore for disaster recovery, verify that the list of associated cartridges is replicated to the DR site. Otherwise, an earlier full backup image (usually the backup from the previous night) must be used for recovery.

Tip: The easiest way to determine the time of the last full backup is if you have a specific time each day where the replication backlog is zero (there is no pending data to replicate).

If not, then assess the cartridges by recovering the backup application catalog and scanning it to find the last full backup where its associated cartridges completed replication.

Ensuring that a copy of the catalog is available at the DR site

The preferred practice for ensuring that a copy of the catalog is available at the DR site is to use the ProtecTIER native replication function. Each day, the catalog should be backed up on a virtual cartridge after the daily backup workload completes so that the catalog is replicated to the DR site at the end of each replication cycle.

If the catalog is backed up to a virtual cartridge, use the cartridge view of the library in ProtecTIER Manager to examine each of the cartridges that are used for catalog backup to find the most recent sync dates marked on the cartridges. If there are multiple backup copies, then find the latest backup that finished replication. To recover the backup application catalog from a backup on a virtual cartridge, work with the replicated cartridges to get an updated copy of the catalog to the DR site:

- ▶ Each cartridge has a *Last Sync* time that displays the last time the cartridge's data was fully replicated to the DR site. (The sync time is updated during the replication and also when the replication for this cartridge is finished.)
- ▶ The cartridge marked with the most recent Last Sync time should be used to recover the backup application catalog.

22.10.10 Eject and inject commands from NetBackup software

Although the process can be manually scripted to enable automation, the easiest way of using the NetBackup commands for automating this process is by using the vault service in the NetBackup software.

Ejecting a cartridge from a library

Ejecting cartridges from a library can be initiated through the NetBackup GUI or by using the NetBackup **vault** option.

If you are using the **vault** command, first run your vault policy by issuing the following command:

```
/usr/opensv/netbackup/bin/vlrun<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>
```

Inserting a cartridge to a library

Inserting a cartridge can be automated through commands. Update the media manager volume by running the following command:

```
/usr/opensv/volmgr/bin/vmupdate -rt dlt -r
```

Recovering a master server from an existing database copy

You can recover the NetBackup catalog by using either the online or offline method:

- Online catalog backup

To use this method, go to the NetBackup Help menu and click **Online, hot catalog backup method**.

- Offline cold catalog backup

To use this method, go to the NetBackup Help menu and click **Offline, cold catalog backup method**.

For more information, consult the official NetBackup application documentation.

22.11 EMC NetWorker deployment with ProtecTIER replication

This section describes how to integrate the ProtecTIER product with the EMC NetWorker (NetWorker) environment for DR by using IP replication. You can find information about integrating the ProtecTIER product with NetWorker through VTL or FSI in Chapter 15, “EMC NetWorker” on page 221, and general preferred practices of ProtecTIER replication in Chapter 20, “Application considerations and data types” on page 295.

There are several benefits to deploying ProtecTIER replication in the NetWorker environment:

- It reduces the administrative effect of managing the physical transportation of tape media.
- ProtecTIER replicates only the data changes to a secondary site so minimal network bandwidth is required. This situation is especially beneficial to customers who previously refrained from implementing electronic vaulting because of the network cost.
- In addition to DR protection, ProtecTIER can be used with NetWorker to clone data to physical tape at a secondary site for third copy protection or long-term archive purposes.

Several important NetWorker terms are used in this description:

Datazone A group of computers that are administered by a NetWorker server. In a datazone, there is only one NetWorker server with several storage nodes and several NetWorker clients. A datazone can be single site or it can span across multiple sites with remote storage nodes, protecting NetWorker clients at different sites, and communicating the backup information to the NetWorker server through a LAN or WAN.

Clone pool A collection of clone volumes that are used only for cloning purposes. NetWorker cloning copies data from backup volumes in backup pools to clone volumes in clone pools.

You can use ProtecTIER virtual cartridges or Advanced File Type Device (AFTD, or FSI share) as a backup volume, and clone it to a clone volume. The clone volumes can be virtual cartridges, AFTDs, or physical cartridges.

Bootstrap A save set that is essential for NetWorker DR. A bootstrap is composed of the media database, the resource database, and NetWorker client file indexes. You should back up the bootstrap at least once a day to local storage devices on the NetWorker server.

The bootstrap backup should be included in your daily NetWorker server backup routine or whenever changes are made to the NetWorker server. Create at least two copies of bootstrap backup: keep one copy onsite for operational recovery; send another copy off-site for DR purposes.

With ProtecTIER, the bootstrap can be backed up to a virtual cartridge or an FSI share on a ProtecTIER system. It can also be replicated to a ProtecTIER system at a secondary site while using minimum TCP/IP bandwidth. This practice eases the system administration management of handling physical tapes for bootstrap backups.

22.11.1 Cloning physical tapes with ProtecTIER replication

This section describes cloning data at a secondary site. The ProtecTIER product replicates backup volumes on virtual cartridges on ProtecTIER server A to ProtecTIER server B, then NetWorker clones the replica virtual cartridges in ProtecTIER server B to physical cartridges at a secondary site. For cloning requirement and configuration details, see *EMC NetWorker Administration Guide*.¹

LTO3 media: The ProtecTIER product emulates LTO3 tape media in NetWorker, but the size of the ProtecTIER LTO3 media might vary from physical LTO3 tape media. Ensure that you have sufficient physical tape in the clone pool. NetWorker supports cloning from one tape media type to another tape media type, and from multiple tape media to single tape media.

Figure 22-7 shows a single datazone deployment where the NetWorker server is at a primary site with VTL on ProtecTIER server A attached to it. On a secondary site, the remote storage node has a VTL and a physical tape library that is attached to it. Data is backed up to virtual cartridges of ProtecTIER server A at the primary site by the NetWorker server and replicated to ProtecTIER server B at the secondary site with ProtecTIER replication. Then, the remote storage node clones the data from replica virtual cartridges in ProtecTIER server B to physical cartridges. This diagram is used for the rest of the description in this section.




Figure 22-7 How NetWorker clones data to physical tape at a secondary site

¹ You can find NetWorker documentation at the NetWorker Information Hub: <http://nsrd.info/docs.html>

Ensure that the following actions and prerequisites are completed and present before cloning:

- ▶ You must have a storage node at the secondary site that belongs to the same datazone.
- ▶ A VTL must be created on ProtecTIER server B and attached to the remote storage node.
- ▶ Define a cloning policy in NetWorker server.
- ▶ Define a replication policy at ProtecTIER server A with a target destination to a VTL of ProtecTIER server B (Figure 22-8). By default, replication takes place after a backup is started and it runs to completion. This can make it difficult to determine the replication completion time. Use a replication schedule so that the cloning process can be scheduled upon replication completion.




Figure 22-8 Set replication destination

The cloning process with the ProtecTIER product

These steps detail the cloning process with ProtecTIER in an EMC NetWorker environment:

1. The NetWorker server writes save sets to virtual cartridges in ProtecTIER server A. ProtecTIER server A then replicates the virtual cartridges to the secondary site either immediately or during the configured replication schedule.
2. Upon completion of the replication, virtual cartridges are ejected (exported) from the source VTL and moved to the virtual shelf of ProtecTIER server A. This step preserves the unique volume ID in a NetWorker datazone when the replica cartridges (with identical volume ID as source tapes) are deposited (imported) in to the target VTL at the secondary site.
3. After source virtual cartridges are exported to the virtual shelf, ProtecTIER server B moves the replica virtual cartridges to the virtual import/export slots of the target VTL. Ensure that you have sufficient virtual import/export slots in the target VTL. For more information about how to allocate virtual import/export slots, see Chapter 21, “ProtecTIER replication” on page 361.
4. Run a library inventory in the remote storage node. The replica virtual cartridges are now accessible by the remote storage node and are ready for cloning.
5. Start cloning from replica virtual cartridges to physical cartridges.
6. Upon completion of the cloning process, eject the replica virtual cartridges to the virtual shelf of ProtecTIER server B so that the source virtual cartridges can be returned to the source VTL for next backup and replication jobs.

This process can be automated with NetWorker scripts so that each step occurs in sequence.

Data deduplication: Data is rehydrated when it is cloned to physical tape. This process is transparent to the NetWorker server and storage node. Therefore NetWorker can restore the data from physical tapes without the presence of a ProtecTIER system.

22.11.2 Disaster recovery with ProtecTIER replication

Figure 22-9 shows a DR configuration where backup activity is performed at the primary site to the VTL of ProtecTIER server A. ProtecTIER server A replicates virtual cartridges to ProtecTIER server B at the secondary site. A standby server exists at the secondary site that is used for NetWorker server recovery during a disaster event.

For this description, a DR is a situation where the NetWorker server and the ProtecTIER system at the primary site are not available. During the disaster, you recover the NetWorker server on the secondary site by using the replicated cartridges of ProtecTIER server B.

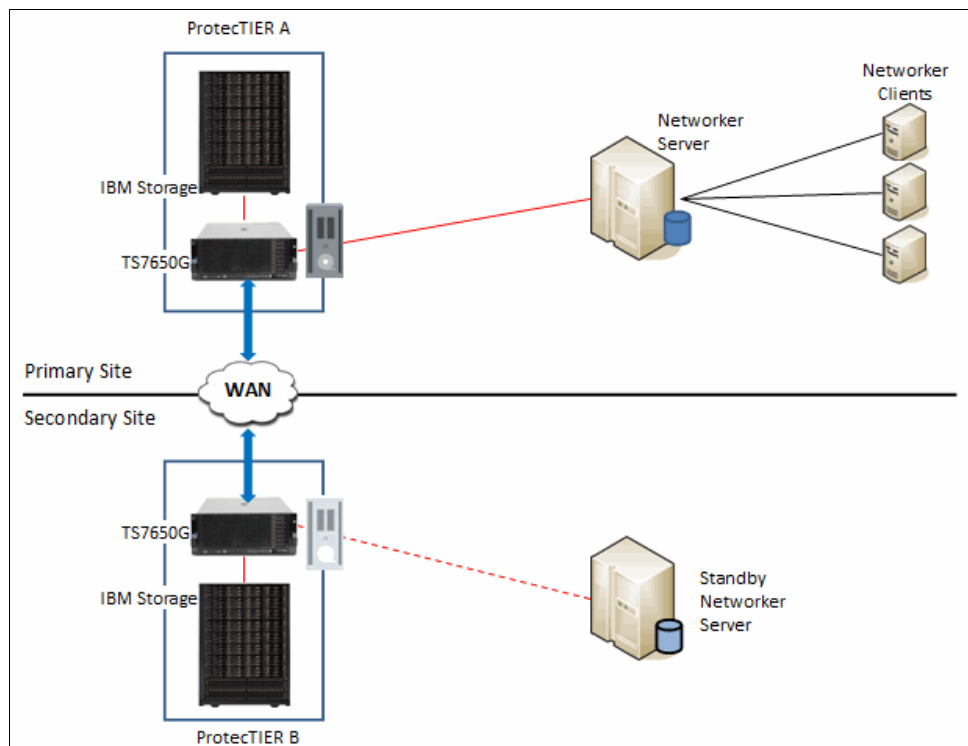


Figure 22-9 Disaster recovery with ProtecTIER replication

For more information about the NetWorker DR procedure, see *EMC NetWorker Disaster Recovery Guide*. For more information about the NetWorker CLI commands that are used in this description, see *EMC Commands Reference Guide*.

Disaster recovery preparation

The following list describes several suggestions about DR preparation:

- ▶ Configure EMC NetWorker based on the suggestions that are described in Chapter 15, “EMC NetWorker” on page 221.
- ▶ Send the bootstrap and the client file indexes to the same media pool. Create a separate media pool from a normal backup data pool for the bootstrap and the client file indexes backup.

- Write the bootstrap and client file indexes to the ProtecTIER server and replicate it over to the secondary site with ProtecTIER replication with high priority.
- During a disaster, the replica virtual cartridges on the secondary site are always in read-only mode. If write enabled cartridges are required, create virtual cartridges with a different bar code naming convention (volume ID). This action eliminates the possibility of bar code conflict (identical volume ID) during failback.

These are the NetWorker requirements for DR:

- A good backup of the bootstrap. It should include the NetWorker media database, the resource database, and client file indexes that must be available at the secondary site during DR.
- An accurate record of the NetWorker server hardware, software, network, device, and media components at the secondary site.

These are the ProtecTIER system requirements for a secondary site include:

- ProtecTIER Manager must be installed and able to access ProtecTIER server B at the secondary site.
- If the Replication Manager is not hosted at the secondary site, see *IBM System Storage TS7600 with ProtecTIER User Guide*² to recover the ProtecTIER Replication Manager. From the ProtecTIER Manager GUI, enter DR mode at ProtecTIER server B.
- If no VTL is created on ProtecTIER server B at the secondary site, then create a VTL with sufficient import/export slots.
- Ensure that the replica cartridge of the bootstrap is in a fully synchronized state.

Procedure to recover NetWorker Server

To recover an EMC NetWorker server, complete the following steps:

1. Replace the server with compatible hardware and operating system.
2. Configure the server with the identical network configuration, including the host name.
3. Install the server with the same version of NetWorker with patch levels that are equivalent to the original location.
4. If the links for the /nsr directory or any of its subdirectories, except for /nsr/res, are missing, then re-create these links.
5. Configure the NetWorker server storage device to detect the ProtecTIER VTL and tape drives by running the **jbconfig** command.
6. Make the replica cartridges or volumes so that the bootstrap is available to the NetWorker server. From ProtecTIER Manager, move the virtual cartridges from the virtual shelf to the import/export slot of the target VTL and import it into the VTL by running this command:

```
nsrib -b <volume name>
```
7. From the NetWorker server command line, inventory the VTL by running this command:

```
nsrjb -lnv -S <the slot with first volume> -f <pathname of first drive>
```
8. If you do not know the bootstrap save set ID and volume ID, see *EMC NetWorker Disaster Recovery Guide* for the steps to recover the server bootstrap.
9. Recover the bootstrap save set by running the **mmrecov -N** command, which prevents overwriting the data in the tape. NetWorker scans the volumes for the appropriate save set and recovers the NetWorker media database and resource database.

² ProtecTIER documentation is at the IBM Support Portal: <http://www.ibm.com/support/>

10. After the media database and resource database are recovered, you must stop the NetWorker services and then rename the NetWorker resource directory (/nsr/res). The directory is recovered to an alternative location, because the NetWorker service is running during the bootstrap recovery process.
11. Restart NetWorker services and start recovering all the client file indexes from the NetWorker server.
12. The NetWorker server is now ready for client and storage node recovery. Perform a test backup and recovery with the standby ProtecTIER VTL.
13. Begin normal operations at the DR site.

Recovering client data

After the NetWorker Server is recovered at the secondary site, you can restore the client data. Use the NetWorker CLI or management console to generate the list of volumes to be recovered. Based on the list of volumes, review the status of the volumes (replicated virtual cartridges) to ensure that all the data is replicated to the secondary site.

After you determine the status of the recoverable volumes, restore the save set through NetWorker.

Resuming NetWorker server operation at the primary site

If new client data backups are done at the secondary site during the disaster, the updated bootstrap and client file indexes must be recovered at the primary site before you replicate the new virtual cartridges to the primary site.

Before you leave DR mode, complete the following steps:

1. Eject all virtual cartridges from the VTL in ProtecTIER server B with NetWorker and then perform a library inventory.
2. In ProtecTIER Manager, create a failback policy to replicate all new virtual cartridges that are created at ProtecTIER server B during the disaster to a primary site.
3. Start failback replication and wait until failback replication completes.
4. At the primary site, move the virtual cartridges from the virtual shelf to a VTL on ProtecTIER server A.
5. Start the NetWorker volume deposit process and perform a library inventory.

You can now resume normal backup and restoration operation.

22.12 Commvault

This section introduces the preferred practices for managing operations by using the Commvault backup application to replicate cartridges from the primary/spoke site to a DR/hub site. Figure 22-10 shows a typical Commvault DR scenario.




Figure 22-10 Typical Commvault DR scenario

22.12.1 Prerequisites

The following list includes the prerequisites for running the Commvault backup application:

- ▶ Verify that both the primary/spoke and the DR/hub ProtectTIER servers are connected, and are identified by the Commvault backup application.
- ▶ Verify that replication activities are running and finishing successfully by checking the active jobs log. To check this log, click **PT Manager** → **Repositories View** → **Replication's Activities** and select the active jobs.
- ▶ Verify that a visibility switch is defined on all replication policies. For example, the replication destination is the DR library. To check the visibility, click **PT Manager** → **Repositories View** → **Replication's Policies** and select the policy's name.

Note: For more information about running Commvault and ProtectTIER, see Chapter 18, "Commvault" on page 257

22.12.2 Running the Commvault backup operation

To run the Commvault backup operation, complete the following steps:

1. Eject media from a library through the Commvault GUI:
 - a. From the Commvault GUI, click **Storage Resources** → **Libraries** → **Local's site library** → **Media By Groups** → **Assigned**.
 - b. Select the required cartridges and click **Export**.
 - c. Define the new outside storage location. Click **OK**.

- d. Verify the change:
 - The Commvault GUI shows that the selected media moved to the new outside storage location.
 - The ProtecTIER GUI shows that the selected cartridges moved from the local library to the shelf.
 - The ProtecTIER GUI on the DR site shows that the library cartridges are in the import/export slots.
2. Import (insert) media into a library through the Commvault GUI:
 - a. From the Commvault GUI, click **Storage Resources** → **Libraries** → **Remote's site library** → **Import Media**.
 - b. Click **Continue**.
 - c. Verify the change:
 - The Commvault GUI shows that the selected media moved to the DR site library. You can access this view by clicking **Storage Resources** → **Libraries** → **Remote's site library** → **Media by groups** → **Assigned**.
 - The ProtecTIER GUI shows that the selected cartridges moved from the local library to the shelf (subject to visibility).
3. Eject (export) and import a cartridge through Commvault CLI:
 - a. In a Windows environment, run the following command:
`<CV installation path> Simpana\Base`
 - b. Log in to the CommServe (the Commvault server) by running the following command:
`qllogin.exe`

Note: The user and password are the same as for the login to the Commvault console.
 - c. Run the following command:
`qmedia.exe export -b <barcode(s)> -el <exportlocation>`
For example:
`qmedia.exe export -b XZXZX1L3 -el Comm_Shelf`
 - d. Verify the change:
 - The Commvault GUI shows that the selected cartridge exported successfully.
 - The ProtecTIER GUI shows that the selected cartridges moved from the local library into the shelf (subject to visibility).
 - e. Run the following command:
`qlibrary.exe import -l <library>`
 - f. Verify the change:
 - The Commvault GUI shows the selected cartridge in the Assigned media window at the DR site.
 - The ProtecTIER GUI shows that the selected cartridges are located inside the DR site library.

22.12.3 Commvault resources

To obtain more technical details about Commvault, see the following resources:

- ▶ Commvault DR strategy and settings:
https://documentation.commvault.com/commvault/v10/article?p=features/disaster_recovery/c_dr_overview.htm
- ▶ Commvault recovery, restoration, and retrieval using Data Agents:
<http://www.commvault.com/solutions/by-function/data-protection-backup-and-recovery>
- ▶ Commvault CLI for configuration and command usage:
https://documentation.commvault.com/commvault/v10/article?p=products/commserve/cli_config.htm

Appendixes

This part contains the following information:

- ▶ Appendix A, “Hints and tips” on page 435
- ▶ Appendix B, “ProtectTIER compatibility” on page 457
- ▶ Appendix C, “ProtectTIER parsers” on page 461
- ▶ Appendix D, “Managing cartridge sizes with ProtectTIER” on page 473

It also contains the Glossary and Related Publications.



A

Hints and tips

This appendix provides hints, tips, and preferred practices for businesses and personnel to gain optimum results by using the ProtectTIER family of products, integrated tools, and applications.

This appendix contains the following topics:

- ▶ Remote access with IBM Endpoint Manager for Remote Control (Assist On-site)
- ▶ Operating system recovery functionality
- ▶ Online fsck command
- ▶ Dedup estimator tool
- ▶ Local directory cloning for FSI
- ▶ Commands to send files to and receive files from IBM Support
- ▶ Graphical performance tool
- ▶ ProtectTIER licensing tips
- ▶ Reconfigure an IBM Spectrum Protect instance attached to ProtectTIER after a reboot

Remote access with IBM Endpoint Manager for Remote Control (Assist On-site)

This section describes the process to enable Assist On-site (AOS) connectivity by using the IBM Endpoint Manager for Remote Control in ProtecTIER TS7620 (SM2) and ProtecTIER TS7650 GW (DD6) servers.

The functionality described here is introduced with the 3.4 version of the ProtecTIER code and will allow the ProtecTIER servers (SM2 and DD6) to open a communication channel so that IBM support can access it remotely.

Other ProtecTIER servers supported by code 3.4 such as TS7650 GW DD4 and DD5 will continue being attached to the IBM System Storage TS3000 System Console (TSSC) and the remote access to the servers will be provided by the same TSSC.

Note: Even when the ProtecTIER DD6 server is considered a Gateway, it completely eliminates the need to be attached to a TSSC for both error reporting (Call Home) and remote access (AOS) purposes. Although attaching a DD6 to a TSSC is not possible, all functionality that is provided by the TSSC is now embedded in the code of the DD6.

Assist On-site (AOS) description

AOS is the preferred remote assistance tool that support engineers can use to connect to, view, and control client systems. AOS enables IBM support representative to remotely takeover the client's endpoint quickly and resolve problems collaboratively.

In terms of security, AOS allows the client to control the session while working with the representative on the phone, and all the activities are performed with permission of the client, on the client's desktop and in full view when regular AOS connectivity is being used.

AOS is configured in *Lights Out* mode in both the TSSC and the SM2 and DD6 servers. In this mode, the system is accessible at any time for an IBM support engineer and client confirmation is not needed. The client can decide to turn off AOS and turn it on only when remote support is required.

Unattended AOS access is usually configured in remote sites where there is no human presence on a regular basis and where remote assistance might be required within a short period of time after a problem presents.

AOS terminology

The following terminology is commonly used with IBM Remote Assistance, AOS.

Console	Support engineer interface. It provides a login point, utilities toolbar, help, configuration and connection options.
Client	Customer interface. Entry point into client, which manifests itself as the ibmaos process.
Relay	Geographically based server selected by a customer when he or she initiates the connection or by support on the HTTP link form. It is also the connection point between the client and the console and comprises AOS.war file.
Controller	Entry point for the console to authenticate support credentials, and create or join a session. It consists of the Admin.war and AOS.war files. The two URLs used for the controllers are <code>us.ihost.ibm.com</code> and <code>uk.ihost.ibm.com</code> .

Note: For AOS v4.0, `uk.ihost.ibm.com` is the only valid controller.

Available session options

AOS offers ways to connect to an end point depending on the type of issue that is presented. All available options are described next; however, notice that the *only* available session option for both SM2 and DD6 is *Port forwarding*.

Chat Only	Opens a text window through which the client and the support engineer can chat. This is usually done when the appropriate type of session is still being decided.
View Only (Monitor)	Client machine is visible through the client (<code>ibmaos.exe</code>). No remote control is available in this type of session. This is the <i>Monitor</i> option, shown in Figure A-1.
Shared Keyboard and Mouse	Client machine is visible through the client (<code>ibmaos.exe</code>). Remote control is possible upon client confirmation. This is the <i>Active</i> option (arrow 1 in Figure A-1).
Collaboration	Session between two or more engineers during the client session. It requires a collaboration code generated by the primary engineer (arrow 2 in Figure A-1).
Port forwarding	This is the default type of session configured for ProtecTIER SM2 and DD6. It creates a tunnel to connect to a specific port in the server which in this case is <i>port 22</i> (ssh).



Figure A-1 AOS console session options

AOS in ProtecTIER (SM2 and DD6)

You can access various actions through the AOS service menu. You can also perform various configurations.

AOS menu

All the actions related to the AOS service that are installed in the ProtecTIER SM2 and DD6 servers are accessible from the service menu by selecting **menu** → **Problem Alerting** → **Manage Assist On Site service (...)**, which then opens the menu shown in Example A-1.

Example A-1 Accessing the menu

```
+-----+
| ProtecTIER Service Menu running on vela |
| Problem Alerting (...)                  |
| Manage Assist On Site service (...)     |
+-----+
| 1) Start the AOS (ibmtrct) service      |
| 2) Stop the AOS (ibmtrct) service       |
| 3) Configure the AOS (ibmtrct) service  |
| 4) Get status of the AOS (ibmtrct) service |
| 5) Enable the AOS (ibmtrct) service     |
| 6) Disable the AOS (ibmtrct) service    |
| 7) Test Connectivity of AOS service     |
|                                         |
| B) Back                                |
| E) Exit                                |
+-----+
Your choice?
```

AOS configuration

Most of the Assist On-site configuration is done automatically when you select the **Configure the AOS (ibmtrct) service** menu. Figure A-2 shows the operations that are performed when you select that menu option.

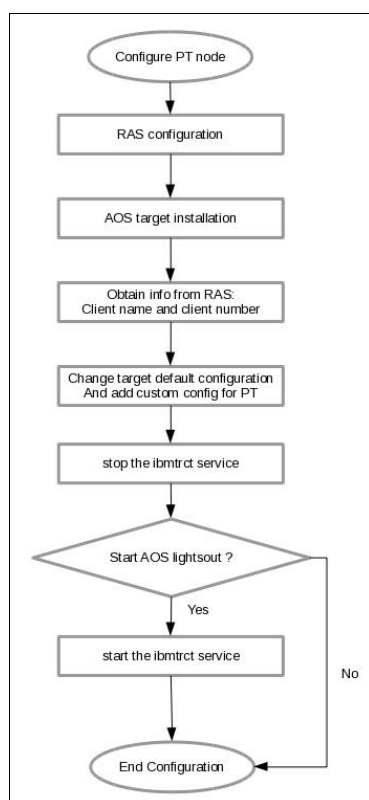


Figure A-2 AOS client configuration sequence

The AOS configuration menu option will modify the following file and will add the necessary information to configure AOS in lights-out mode:

`/var/opt/ibm/trc/target/profiles/lightsoutprofile.properties`

Part of this information includes the customer's name and customer's number, which means that to be able to successfully configure AOS, reliability, availability, and serviceability (RAS) software configuration must be completed first.

Finally, a second file, `/etc/ibmtrct.conf`, is modified to allow faster remote connections where the **PortToListen** parameter is set to zero.

Important:

- ▶ AOS is not configured automatically as a part of the ProtecTIER code configuration. It is only after selecting the **Configure the AOS (ibmtrct) service** menu option that the ProtecTIER system will be listed as a *Registered Target* and will be made available for IBM support to access.
- ▶ The following IP address and port should be open in the firewall in order to reach the AOS server:
195.171.173.165 : 443

Operating system recovery functionality

This section describes one of the new functionalities introduced in the ProtecTIER code version 3.4 and in particular for the DD6 model of the server.

The main objective of operating system (OS) recovery is to add an extra level of redundancy to the ProtecTIER server by making it able to create *images* of its current configuration (including all the operating system configuration files and ProtecTIER software configuration files) and installed programs. The images created by this functionality can eventually be used to reduce the time needed to recover a server from a catastrophic failure because both the operating system and the ProtecTIER configuration can be recovered in a single step.

OS Recovery functionality description

The OS Recovery functionality makes use of one of the design characteristics of the ProtecTIER server DD6 model, which differentiates it from all other server models that were previously released. The ProtecTIER DD6 server creates the `pt_work` partition (which is where the system logs are stored) in a disk that is not the same as the disk where the root partition resides.

Figure A-3 shows where the `pt_work` partition resides, depending on the role and physical location of the node that is configured.



Figure A-3 DD6 server front view and `pt_work` partition location based on canister that mounts it

Note: Other than slots 0 and 23, all front SAS disk slots in a DD6 server are empty.

An OS Recovery image is a binary file that contains all the files that were backed up when the image was created. In general, it maintains the exact same file system structure as the root partition at the moment of its creation.

The way the images are stored is shown in Figure A-4 on page 441. If the images are created in a two-node cluster environment, a copy of the most recent image will be sent to the peer node for redundancy purposes.




Figure A-4 OS Recovery images creation in a 2 node cluster environment

The OS Recovery menu

All operations related to OS Recovery and OS Recovery images are performed from the service menu and can be accessed by selecting **menu** → **ProtectTIER Configuration** → **OS images Management (...)**, which opens the menu shown in Example A-2.

Example A-2 Service menu

```

+-----+
| ProtectTIER Service Menu running on vela |
| ProtectTIER Configuration (...)          |
| OS images Management (...)              |
+-----+
| 1) Create an OS image                   |
| 2) List the existing OS images          |
| 3) Check the status of the OS image creation |
| 4) Check the latest OS recovery Journal  |
|                                          |
| B) Back                                |
| E) Exit                                |
+-----+
Your choice?

```

The menu has the following options:

1. Launches an OS recover image creation and is described in detail in “Launching an OS Recovery image creation” on page 442.
2. Lists the existing OS images stored in /pt_work/images directory in the local node. If this is a dual node cluster environment, then this option will also show the copies of the images residing in the peer node.

3. Checks whether an OS Recovery image is being created. An important observation is to notice that the images are always created in the background so that the regular operations of the ProtecTIER system will not be suspended or affected while this is happening. For this reason this option is the preferred way to check for the status of the images creation process.
4. Displays the most recent OS Recovery journal. The journal is a small file that keeps a record of all the sub-actions performed by the image creation process. Every row in this file indicates an action that was completed.

Note: Unless the last line in the journal file says `Finished`, a journal file with a size greater than 0 (zero) is an indication of an on-going image creation or a failed attempt. The third menu option will differentiate between various status indications automatically and will report the current one.

Launching an OS Recovery image creation

An OS Recovery image creation is automatically initiated after a code installation or after a code upgrade.

The way this works is as follows:

1. At the end of the code installation or upgrade and just before the user is asked to press Enter to reboot the node, a special line is added to the **crontab** configuration.
2. This additional line in the **crontab** configuration executes the command that launches the image creation in the background.
3. After this command runs, it will automatically remove itself from the **crontab** configuration file.

Additional to the automatic OS Recovery image creation, a user can manually initiate the creation of an image by using the **Create an OS image** option in the menu. As with an automatic image creation, the manual execution of this option sends the image creation operation to the background.

Important:

- ▶ Be sure an image is created after every major change to the node's configuration. Although an image will be created automatically after an installation or an upgrade, a good approach is to manually launch an image creation whenever a change in the configuration such as networking or other significant changes are done.
- ▶ Only one image can be created at a time. If an image is already being created when the manual execution of this option is performed, the operation will fail immediately, indicating the reason.

Recovering the operating system using an image

OS images contain a copy of the file system structure as it was when the image was created. This implies that when a system is recovered, using an image in particular, each one of the files that form the operating system will be overwritten with the versions that correspond to the image that is being used.

Figure A-5 on page 443 shows the sequence of actions needed to recover the files, and configuration stored in an OS image is described in the flowchart.

Important: The OS Recovery process should not be started unless indicated by IBM support. Doing so might lead to inconsistencies in the configuration of the server and the repository, and can potentially make it unusable and extend the service actions needed to recover it.




Figure A-5 OS Recovery sequence

The steps shown in Figure A-5 are mostly actions that are composed of automated procedures that require minimal human intervention. However there is one action that requires basic knowledge of the cluster layout on which the recovery is being performed. This action is represented by the *Respond initial questions* box in the flowchart.

The initial questions are displayed when the OS Recovery menu option is selected, as shown in Example A-3. The initial questions are used to allow the OS Recovery process to properly determine which node is being recovered which in turn is required to be able to display the appropriate OS image files in case this is a dual node cluster. For this reason, the person that is doing the recovery must be aware of the cluster configuration and also the location of the node (upper slot or lower slot) that is being recovered in order to increase the possibilities of recovering the system.

Example A-3 OS Recovery Initial questions

```
++++++ Starting OS Image Recovery ++++++
Is this a dual node cluster? [Y/N]
N <---- User response
Obtaining canister index
Canister 2 will be recovered
Is this correct? [Y/N]
Y <---- User response
```

Note: Failing to properly define the node that is being recovered can result in duplicated configurations including duplicated IP addresses and duplicated WWNs, which in turn can cause more problems.

When the OS Recovery process displays the available images, it will also include a time stamp for each of them. The time stamp can be used to determine the most suitable image to recover from (for example, the most recent image might have better chances of containing the most recent version of the installation and configuration files).

Every image that is successfully created is *marked* with a special file, which indicates that it contains all the files it was supposed to contain. In other words, it indicates that the image is complete and valid.

A validation of the selected image is performed by the OS Recovery process. If this special file is not detected, you are warned that the image might not be complete; however, you are given the option to proceed with the recovery.

Important: Do not proceed with the restoration of an incomplete image unless indicated by IBM support.

OS Recovery was designed to reduce the recovery time when a node reinstallation is needed, however the previous Node Replacement procedure (that is selected through the **Recover Configuration for a replaced server** option) is available also.

The difference between OS Recovery and Node Replacement is as follows:

- ▶ OS Recovery stores the installation and configuration files in an image file that resides in one of the front SAS disks in a DD6.
- ▶ Node Replacement rebuilds the configuration of a ProtecTIER server by reading the configuration stored in the repository of the system.

Another difference between both processes is that in order to execute Node Replacement on a node, the node must be upgraded to the same level of code as the repository which could imply to perform more than one code upgrade after the initial installation before executing the Replacement.

Online fsck command

This section describes a method to run the File System Check (**fsck**) command while backups are being sent to the ProtecTIER system.

Important: The **fsck** command must be run under the guidance of IBM support and should never be attempted without a deep analysis of the situation present in the system.

The Linux fsck tool

In Linux, everything that is managed by the operating system is represented by a file. To be able to manage such files, a file system must exist. One of the main functions of a file system is to enable read/write access to the files it contains. In addition, some file systems include security features and other more advanced characteristics that make them unique.

Most Linux file systems include a special set of tools that are used to attempt and repair inconsistencies that can potentially damage the files they contain. These tools are called *system-specific checkers*. What **fsck** does is simply call the appropriate checker according to the type of file system that is being analyzed or repaired.

Execution of fsck

Various situations might require the running the **fsck** command, and many of those situations are potential data-loss scenarios that must be treated carefully. In general, to try to repair a file system using **fsck**, the system must be unmounted, which implies that the access to the files it contains is suspended.

Although this is also true for a ProtecTIER system, it has major implications because the ProtecTIER repository is built on top of a Linux file system. Losing access to any of the components of the repository can cause adverse situations, such as from losing access to a subset of the user data, to being unable to send new backups and restoring the existing ones.

Running **fsck** consumes a large amount of RAM and, depending on the size of the file system, it can potentially allocate nearly 100% of the memory in the system. This allocation puts extra stress on the system were it is running, especially when other applications are also requesting RAM to operate. For this reason, the general suggestion when running **fsck**, traditionally, as shown in Example A-4, is to first close all other programs and then let the check run for as long as needed.

Example A-4 Traditional fsck execution

```
fsck -y -v /dev/mapper/vgfs0002-lvfs0002
```

This way of course translates into long outages in ProtecTIER systems during which no backups or restores will work.

Online fsck procedure

To reduce down time and to allow the continuation of backup and restore activities, starting with ProtecTIER code version 3.3.7, a method to run **fsck** without losing access to the data was developed.

Important: The following requirements must be met in order to run the online **fsck** command:

- ▶ ProtecTIER must be configured as a two-node cluster.
- ▶ Each ProtecTIER node must be configured with at least 64 GB of RAM.

One of the following criteria must be true in order to have access to the data:

- ▶ Medium changer is active on the node that is up during fsck execution.
- ▶ Control Path Failover (CPF) is enabled if you use IBM Spectrum Protect (formerly Tivoli Storage Manager) as backup application.

The procedure is as follows:

1. Select one of the ProtecTIER nodes.
2. At the login prompt, log in with the ID **root** and the password **admin**.
3. Shut down the **vtfd** and the **ptcluster** service in the following order:

```
#service vtfd shutdown  
#service ptcluster shutdown
```
4. Start the **cman** and **clvmd** services in the following order:

```
#service cman start  
#service clvmd start
```
5. For each of the problematic File Systems, run the **gfs_fsck -yv** command in the background and with high verbosity. Direct the output to a file as shown in Example A-5.

Example A-5 The fsck command run in the background with high verbosity

```
#fsck -yv logical_volume_name > /pt_work/logical_volume_name_fsck.out 2>&1 &
```

Note: ProtecTIER supports up to 10 parallel **fsck** instances.

6. When all the **fsck** processes are complete, collect the generated output files and make them available for IBM support to analyze. IBM support should indicate what the next steps are to regain access to the affected File Systems based on the analysis of the output files.

Dedup estimator tool

The *dedup estimator* tool can help you analyze the effect of deduplication on the data that is stored in the back-end disk arrays. The functions described in this section are an alternative to the *analyze_session* tool, which is limited to the data that currently resides in the repository. This new method can help determine possible corrective actions to improve the factoring ratio of the data, such as redirecting *dedup unfriendly* data to other backup targets or more complex actions designed with help of IBM support when possible.

Previous solution

The dedup estimator tool was introduced in the ProtecTIER v3.3.7 code package. Before that version, the ProtecTIER solution included only one tool to process the deduplication statistics to create user-readable reports. This *analyze_sessions* tool is still supported in ProtecTIER v3.4 code package. The tool can analyze the deduplication information of the data ingested by ProtecTIER and generates a report that groups this information by session, day, hour, and month. However, it does not account for data that expired or that was deleted from the repository. At a low level, this tool works by analyzing the *Compressor Logs* that are stored in the repository.

Tool functions

The dedup estimator tool runs concurrently and has no effect in the overall performance of the system when executed, even when the list of barcodes to analyze is large.

The interface of the tool is accessed through the ProtecTIER command-line interface (**ptcli**), which means that in order to call it, a valid user profile must exist. If no profile exists or if you want create a new one, complete these steps:

1. Run the following command:

```
/opt/dtc/ptcli/ptcli -p [file name with full path]
```

2. When asked, enter a valid user name and password.

Note: The user name and password must be previously defined in the ProtecTIER Manager GUI.

In general, when a valid profile already exists, the **ptcli** tool has the following syntax:

```
/opt/dtc/ptcli/ptcli [command name] --loginFile [path to the profile file]  
[command specific options]
```

Function: CalculateVtlDedupRatioByFile

This function calculates the VTL deduplication ratio for a list of cartridges defined in a plain text file.

The syntax is as follows:

```
/opt/dtc/ptcli/ptcli CalculateVtlDedupRatioByFile --loginFile /tmp/myProfile  
--file PATH
```

In this syntax, the **PATH** argument contains the full path to the location of the list of cartridges or barcodes to be analyzed. The input file must be a plain text file with one barcode per line, as in the following example:

```
A00001  
A00002  
A00003
```

The **CalculateVtlDedupRatioByFile** function has the following optional parameters:

- ▶ **[--maxErrorsPercentage NUM]**

NUM indicates the maximum errors percentage in the range of 0 - 100. The default value is taken from the **GetDedupRatioParameters** function output.

- ▶ **[--tolerancePercentage NUM]**

NUM indicates the tolerance percentage in the range of 0 - 100. The default value is taken from the **GetDedupRatioParameters** function output.

- ▶ **[--numOfSamples NUM]**

NUM indicates the number of samples used to generate the analysis. The default value is taken from the **GetDedupRatioParameters** function output.

Function: CalculateVtlDedupRatioByLibrary

This function calculates the VTL deduplication ratio by library.

The syntax is as follows:

```
/opt/dtc/ptcli/ptcli CalculateVtlDedupRatioByLibrary --loginFile /tmp/myProfile  
--libraryName LIBRARY
```

In this syntax, the LIBRARY argument contains the name of the library to be analyzed. All the cartridges residing in the library will be used to generate the final report.

The CalculateVtlDedupRatioByLibrary function has the following optional parameters:

- ▶ [--maxErrorsPercentage NUM]
NUM indicates the maximum errors percentage in the range of 0 - 100. The default value is taken from GetDedupRatioParameters function output.
- ▶ [--tolerancePercentage NUM]
NUM indicates the tolerance percentage in the range of 0 - 100. The default value is taken from GetDedupRatioParameters function output.
- ▶ [--numOfSamples NUM]
NUM indicates the number of samples used to generate the analysis. The default value is taken from GetDedupRatioParameters function output.

Function: CalculateVtlDedupRatioByRange

This function calculates the VTL deduplication ratio for a range of cartridges specified in the parameters.

The syntax is as follows:

```
/opt/dtc/ptcli/ptcli CalculateVtlDedupRatioByLibrary --loginFile /tmp/myProfile  
--fromBarcode BARCODEFROM --toBarcode BARCODETO
```

In this syntax, the BARCODEFROM argument contains the initial barcode of the range and BARCODETO contains the last barcode of the range.

This function has the following optional parameters:

- ▶ [--maxErrorsPercentage NUM]
NUM indicates the maximum errors percentage in the range of 0 - 100. The default value is taken from GetDedupRatioParameters function output.
- ▶ [--tolerancePercentage NUM]
NUM indicates the tolerance percentage in the range of 0 - 100. The default value is taken from GetDedupRatioParameters function output.
- ▶ [--numOfSamples NUM]
NUM indicates the number of samples used to generate the analysis. The default value is taken from GetDedupRatioParameters function output.

Sample output of the functions

Example A-6 shows output of using one of the dedup estimator functions.

Example A-6 Output from a dedup estimator function

```
Result:
#### WARNING!! ####
All open sessions of PT Manager must be terminated before running Dedup Estimator
utility.
To continue working with PT Manager while running the Dedup Estimator utility,
connect to PT Manager as ptooper.
Please wait while Dedup Estimator is working. This may take a few minutes to
complete.
Preparing: 100% completed
Analyzing: 100% completed
Calculating: 100% completed
Node-id = 1
Estimated compression-ratio = 1:2.09
Estimated dedup-ratio = 1:3.61
Estimated overall factoring-ratio = 1:7.54
16384 successful samples out of 16384
```

Local directory cloning for FSI

This section describes a procedure to clone a local directory to create semi-instantaneous clones of directories in ProtecTIER FSI repositories at the source site in a replication environment.

Notes:

- ▶ This procedure works equally well if the repository is not in a replication environment.
- ▶ ProtecTIER GA Version 3.4 was released with only the Virtual Tape Library (VTL) interface support. File System Interface (FSI) support was added to ProtecTIER PGA 3.4 Version. For details, see the announcement letter:

http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/6/872/ENUSAP16-0076/index.html&lang=en&request_locale=en

This feature is similar to the procedure for cloning an FSI replication destination directory described in 22.4.2, “ProtecTIER FSI cloning” on page 403. However subtle differences between both procedures exists.

Consider the following information about the procedure for cloning a local directory for FSI:

- ▶ This procedure is intended to be run at the source site.
- ▶ This procedure cannot be run from the ProtecTIER Manager GUI, only from the command line.
- ▶ The main objective of this procedure is not to test a DR scenario but to create snapshots of directories in a specific point in time

Cloning a local directory in FSI

To clone a local directory in FSI, the command shown in Example A-7 must be issued from a command line on the ProtecTIER node that is attached to the repository where the FSI directory resides.

Example A-7 Cloning a local directory using ptcli

```
/opt/dtc/ptcli/ptcli CloneDirectory --loginInline [user,password] --force  
--sourceFS [SOURCEFS] --targetFS [TARGETFS] --sourceDir [SOURCEDIR] --targetDir  
[TARGETDIR]
```

The following list describes the parameters used in Example A-7:

- ▶ **user,password**
A valid user and password with administrator privileges is required. For the PT Manager GUI, ptadmin,ptadmin are the default credentials.
- ▶ **SOURCEFS**
The name of the file system where the directory that will be cloned resides, as seen in the PT Manager GUI.
- ▶ **TARGETFS**
The name of the file system where the clone of the directory will reside. TARGETFS can be the same as SOURCEFS.
- ▶ **SOURCEDIR**
The full path to the directory that will be cloned as seen in the PT Manager GUI. If the directory is under root, then use only its name without the initial forward slash (/) character. If it is a sub-directory, continue forming the path by following the UNIX file path format.
- ▶ **TARGETDIR**
The full path to the new directory created after cloning the SOURCEDIR. If the new directory will be under root, then use only the name of the new directory without the forward slash (/) character. If it will be a sub-directory, then continue forming the full path by following the UNIX file path format.

After a local directory is cloned, the new directory will have read/write permissions and the local repository can use it as if it was any other directory residing in the local repository. The local repository can even create a CIFS share or an NFS export on top of a cloned directory without issues.

Local directory cloning benefits

Cloning the local directory for FSI provides the following benefits when applied in conjunction with a certain backup application such as IBM Spectrum Protect or Veritas NetBackup, or during a specific situation:

- ▶ A local directory can be duplicated to a new directory without any impact on the physical used space because the cloned directories have a change rate of 0% or, in other words, they are 100% deduplicated.
- ▶ Data restoration from a specific point in time can be implemented for deleted data.

- Local directory cloning can be used in test environments to restore data.
- Local directory cloning can be used together with a backup application to resolve issues related to incremental and full backups.

Note: To demonstrate this point, we use RMAN backups as an example. Normally a full backup is run followed by several incremental backups the next days. The problem is that when you apply the incremental backup to a full backup, the original full backup is lost and replaced by the merged version of the full plus the incremental backups.

To resolve this issue, create a clone of the full backup prior to an incremental update operation, which will not consume physical space in the repository. In this way, you can generate backups in fractions of time. In other words, instead of generating only one full backup and many incrementals, you can now have one full plus one incremental backup every day which will accelerate the restores in case they are needed.

Commands to send files to and receive files from IBM Support

This section includes a set of commands that can be used to upload or download files when you interact with the Enhanced Customer Data Repository (ECUREP). This repository is used by IBM support to receive files, such as ProtecTIER Service Reports, to diagnose problems in the system. It is also used to send large files to the client when needed. Such files can include ProtecTIER software packages, repaired ProtecTIER repository structures and other files that can not be attached to a regular email.

Note: If you are concerned about specific aspects of the ECUREP security standards, contact your sales representative to ask about the Blue Diamond program and how it can be applied to secure your data.

Commands to upload files from a ProtecTIER node to ECUREP

One of the most common tasks related to the support process is to upload a *service report*. This report contains all the necessary files and logs for IBM to start working on an analysis related to a support case.

Several ways are available to upload a service report to ECUREP and attach it to a problem management record (PMR). The method described here aims to reduce the total time needed to make the Service Report available in a PMR by minimizing the number of file transfers needed. Depending on the type of service report and depending on specific conditions present in the ProtecTIER node when the file is created, the final size can vary from hundreds of megabytes to several gigabytes, and this can be problematic when trying to upload it to ECUREP.

Prerequisite

The only real prerequisite is to make sure the ProtecTIER node where this procedure is being executed has access to the Internet. One way to test this access is to try to ping known IP addresses and to try to directly connect to ECUREP.

Note: The configuration of Red Hat in ProtecTIER does not include capabilities to resolve domain names. In other words, all the tests and the commands to transfer the files must be executed by using the IP addresses associated with the services that are being accessed.

Complete the following steps:

1. Make sure the system is authenticated on all the firewalls that control the access to the Internet. One way to do this is to use Telnet to connect to a known DNS as many times as needed until Telnet does not return with a prompt. Example A-8 shows how this task is performed. In this case, Telnet was able to connect, which means that a firewall is requesting authentication. The DNS used here does not respond to the Telnet protocol so one way to confirm that all the necessary firewalls were passed is by making sure Telnet *does not* connect successfully to this IP address.

Example A-8 Telnet to a known DNS

```
[root@donramon ~]# telnet 8.8.8.8
Trying 8.8.8.8...
Connected to 8.8.8.8 (8.8.8.8).
Escape character is '^]'.
```

2. Try to connect through FTP to one of the ECUREP IP addresses. The IP address used in Example A-8 might vary depending on how the fully qualified domain name is resolved.

Example A-9 shows the results of a successful attempt to ensure that the system is authenticated on all firewalls that control access to the Internet.

Example A-9 FTP connect to ECUREP IP address

```
[root@vela ~]# ftp 192.109.81.7
Connected to 192.109.81.7.
220-FTP Server (user 'vanfalen@mx1.ibm.com')
220
500-Syntax error -- unknown command
500
500-Syntax error -- unknown command
500
KERBEROS_V4 rejected as an authentication type
Name (192.109.81.7:root):
```

3. If the FTP connection is successful, then this ProtecTIER node is able to connect to ECUREP. The next step is to prepare the service report or any other file that should be attached to a PMR for the transfer. This preparation implies renaming the file to add the number of the PMR at the beginning of the file name to allow ECUREP to move it to the appropriate path associated with the PMR. Example A-10 describes how to do this task.

Example A-10 Naming convention for files sent to ECUREP

If the file is a service report, it is very possible it will be located under /pt work directory, similar to this:

```
/pt_work/ProtecTier__vela_default_Mar_6_2016_0_57_49_Report.tar
```

If the target PMR is for example 12345,000,123 then the file should be renamed as follows:

```
mv ProtecTier__vela_default_Mar_6_2016_0_57_49_Report.tar
12345,000,123_ProtecTier__vela_default_Mar_6_2016_0_57_49_Report.tar
```

4. Finally, the file can be transferred to ECUREP by issuing the command in Example A-11. Be sure to notice that the file will take some minutes to be listed in the directory in the PMR because it is automatically moved by a process running in the background in the ECUREP repository.

Example A-11 Command to initiate the transfer of a file from the PT node to ECUREP

```
curl -u anonymous:[EMAIL] -T
12345,000,123_ProtectTier__vela_default_Mar_6_2016_0_57_49_Report.tar
ftp://192.109.81.7/toibm/aix/12345,000,123_ProtectTier__vela_default_Mar_6_2016_0_57_49_R
eport.tar
```

Consider the following information about Example A-11:

- anonymous indicates that the file is being uploaded anonymously.
- [EMAIL] is a valid client email address.
- 192.109.81.7 is the IP address of ECUREP.
- /toibm/aix/ is a directory in ECUREP repository to receive files sent by customers.

Receiving files from IBM support

The ECUREP security rules have evolved over time to become more reliable. Derived from these improved security policies, the old practice that involved placing files in a specific ECUREP directory for a client to download anonymously was deprecated.

Currently, the only official procedure for IBM to send a file to a client is by making an explicit request in ECUREP to place it in a public FTP server for a limited period of time. After the request is submitted, the client will receive an email reply. A link to download the file directly will be included in that same email.

This way implies that only files associated with PMRs can be transferred from ECUREP to a client. This limits the number of users that can perform this operation. Only support engineers, technical advisors, technical advocates and people with special permissions will be able to do it.

Note: For additional information about sending or receiving files through the ECUREP service, contact IBM Support.

Graphical performance tool

Starting in ProtectTIER code 3.4, a new tool was added to the base installation of the ProtectTIER. The tool allows the creation of graphical reports based on the processing of the statistical data collected by each ProtectTIER system.

The tool generates an HTML page that can be opened in any web browser. After it is created, the report must be transferred to a work station because the ProtectTIER code package does not include programs to open such files.

This tool generates several graphs that can be used to visualize and detect problems that are related to low performance, bottle necks, high-delete backlog, high-fragmented space, and others related to replication activities.

For more details about this tool, see the topic about creating a performance analysis report in *IBM TS7650 with ProtectTIER V3.4 User's Guide for VTL Systems*, GA32-0922-10.

ProtectTIER licensing tips

This section contains tips regarding to the capacity licensing of ProtectTIER. The information described here was gathered from the most common questions that are associated with implementation of a ProtectTIER system.

For more information, contact your IBM sales representative.

Capacity representation in the ProtectTIER Manager GUI

The ProtectTIER Manager GUI displays capacity information in the following format:

- ▶ A tebibyte (TiB) is defined as a multiple of the byte data unit. More specifically, 1 TiB equals 2^{40} bytes or 1024 gibibytes (GiB), which equals 1 099 511 627 776 bytes.
- ▶ A terabyte (TB) in contrast, is defined as 10^{12} bytes, which equals 1 000 000 000 000 bytes.

Both units of measurement are usually used as synonyms; however, knowing the exact values of both scales is important. It explains why 1TB is approximately 1.1 TiB and it also explains the numbers displayed by the ProtectTIER Manager GUI.

Figure A-6 shows an example of capacity representation in the ProtectTIER Manager GUI.

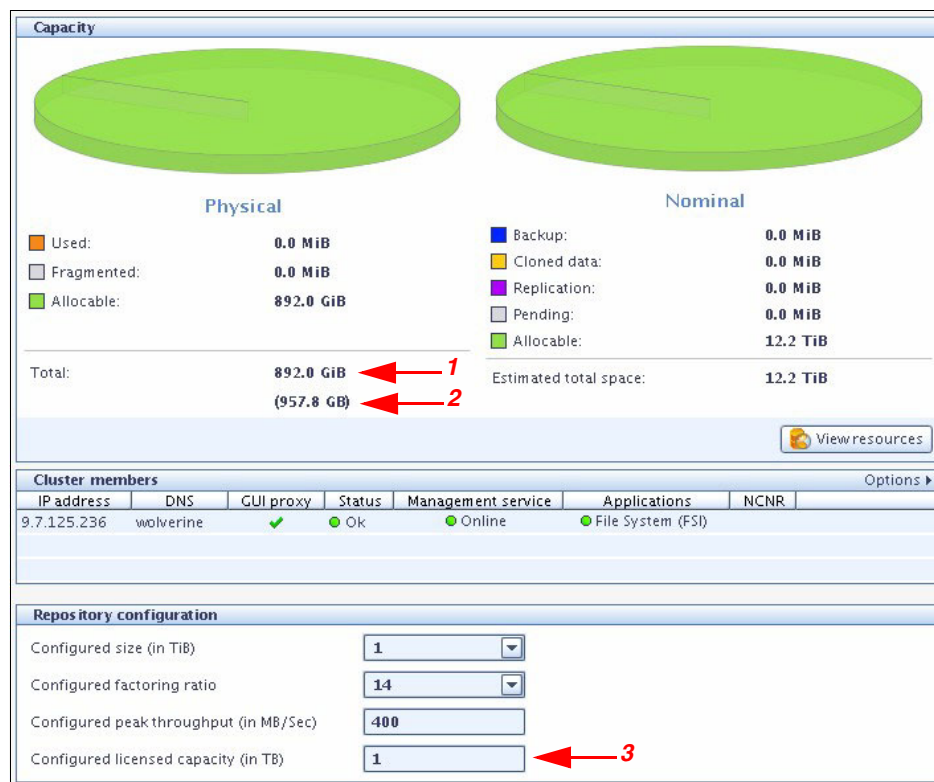


Figure A-6 Capacity view in the ProtectTIER Manager GUI

Figure A-6 on page 454 shows how the capacity is represented in the ProtecTIER Manager GUI. The highlighted numbers in the figure represent the following information:

1. Usable capacity: Licensed capacity reduced by the Linux file system overhead and other ProtecTIER internal overhead. It is displayed as a binary value (GiB or TiB).
2. Configured real licensed capacity: The licensed real capacity, which is the sum of all user data file systems allocated to the repository.
3. Configured licensed capacity: This value represents the capacity that was configured when the repository was created. This number indicates how much capacity the current metadata is able to manage. This number can be larger than the value displayed (highlighted item 2) because ProtecTIER allows preconfiguration of the metadata to leave it ready for future capacity increases.

The role of the metadata space in the licensing process

The capacity license of ProtecTIER considers only the user data portion of the repository. The user data portion is where the actual deduplicated data is stored in the storage system that provisions the LUNs that are presented to the ProtecTIER system.

The metadata space required by the ProtecTIER repository stores vital structures that are used by all the ProtecTIER internal modules to operate. Its size depends on characteristics of the repository that is being created, such as maximum throughput, expected factoring ratio, and the type and configuration (RAID configuration) of the disks that are allocated for metadata.

Reconfigure an IBM Spectrum Protect instance attached to ProtecTIER after a reboot

This section describes a procedure to recover the configuration of an IBM Spectrum Protect (formerly Tivoli Storage Manager) instance when it is attached to a ProtecTIER system. The procedure described in this section is valid for IBM Spectrum Protect instances running on Linux.

When ProtecTIER services are restarted or when the ProtecTIER server is rebooted, the serial numbers of the tape devices might differ from those that IBM Spectrum Protect is expecting to detect. To resolve this issue, the following procedure can be applied:

1. Create a file named `98-lin_tape.rules` in the `/etc/udev/rules.d/` directory.

The contents of the file are shown in Example A-12.

Example A-12 Contents of the 98-lin_tape.rules file

```
KERNEL=="IBMtape*",ATTR{serial_num}=="8515844000",SYMLINK="lin_tape/by-id/IBMtape7n"
.....
KERNEL=="IBMtape*",ATTR{serial_num}=="8515844029",SYMLINK="lin_tape/by-id/IBMtape29n"
```

2. You can query `ATTR{serial_num}` by issuing the following command:

```
udevadm info --attribute-walk --name /dev/IBMtapeX
```

3. Change the `persistent_n_device` parameter in `lin_tape` as indicated in Example A-13.

Example A-13 How to change persistent_n_device parameter

```
#lin_taped stop
#modprobe -r lin_tape
#vi /etc/modprobe.d/lin_tape.conf
options lin_tape persistent_n_device=1
```

* save and exit

```
#modprobe lin_tape
#lin_taped
#ls -l /dev/lin_tape/by_id/
```

* one or more files should be created in the `../lin_tape/by_id/` path

4. Create a library in IBM Spectrum Protect and use `IBMtape*n`.



ProtecTIER compatibility

This appendix describes the main resources that can be used to determine the compatibility of ProtecTIER with other software and hardware products, such as backup applications, operating systems, Host Bus Adapters, drivers, and so on.

This appendix describes the following topics:

- ▶ IBM System Storage Interoperation Center (SSIC)
- ▶ Independent Software Vendor Support Matrix

IBM System Storage Interoperation Center (SSIC)

The IBM System Storage Interoperation Center (SSIC) is at the following address:

<http://www.ibm.com/systems/support/storage/ssic/interoperability.wss>

To help decide what values to enter for ProtecTIER in the SSIC, the search options can be divided as follows:

- ▶ Options for the ProtecTIER:
 - Storage Family
 - Storage Model
 - Storage Version
 - Connection Protocol
- ▶ Options for the host and the back end:
 - Host Platform
 - Server Model
 - Operating System
 - Adapter (HBA, CNA, and others)
 - Switch Module
 - SAN or Networking
 - Multipathing
 - Storage Controller

Although this division is not official, values can be taken from either ProtecTIER or from the host, using the description here as guidance.

On each section, only one value can be selected at one time. After you select a value, all sections are automatically repopulated to show possible further options to continue the filtering search.

For ProtecTIER, the search can start at Storage Family, by selecting **IBM System Storage Virtualization Engine for Tape**.

Then, you see how the Storage Model options repopulate with the ProtecTIER model family: **TS7610, TS7620, TS7650, TS7650G**.

After selecting the Storage Model, refine your options for ProtecTIER by choosing values for **Storage Version** and **Connection Protocol**.

After selecting your *options for the ProtecTIER*, continue filtering by selecting *options for the host and back end*.

Figure B-1 on page 459 shows how the sections are populated after the first selection for ProtecTIER.

IBM System Storage Interoperation Center (SSIC)

• Start your search with ANY of the below selection boxes. You are NOT required to perform your query from the top down.
 • Please view the details of your selected configuration. This requires clicking the Submit button or exporting your data.

Revise Selected Criteria - click link below to change search query
 (1) [Storage Family](#)

[New Search](#)
Configuration Results= 700,442
[SSIC Education and Help](#)

Storage Family <small>Type to filter selections</small> <div style="border: 1px solid #ccc; padding: 2px;"> IBM System Storage Enterprise Tape IBM System Storage LTO Ultrium Tape IBM System Storage Virtualization Engine for Tape IBM System Storage SONAS IBM System Storage Dual Drive Enclosures </div>	Storage Model <small>Type to filter selections</small> <div style="border: 1px solid #ccc; padding: 2px;"> TS7610 TS7620 TS7650 TS7650G TS7650G Storage Controller Support </div>
Storage Version <small>Type to filter selections</small> <div style="border: 1px solid #ccc; padding: 2px;"> TS7610 (R2.1) TS7610 (R2.5) TS7610 (R3.1) TS7610 (R3.2) TS7620 (R3.2) </div>	Connection Protocol <small>Type to filter selections</small> <div style="border: 1px solid #ccc; padding: 2px;"> CIFS FCP (z Systems) Fibre Channel FICON NFS </div>

[Export Selected Storage Version \(xls\)](#)

Host Platform <small>Type to filter selections</small> <div style="border: 1px solid #ccc; padding: 2px;"> IBM z Systems IBM Power Systems (p6 and newer) IBM BladeCenter IBM Flex Systems IBM System i </div>	Server Model <small>Type to filter selections</small> <div style="border: 1px solid #ccc; padding: 2px;"> ALL x86 & x86_64 Compatible Servers HP Compatible Servers HP Servers IBM BladeCenter Servers (Intel) IBM BladeCenter Servers (POWER) </div>
Operating System <small>Type to filter selections</small> <div style="border: 1px solid #ccc; padding: 2px;"> HP HP-UX 11iv2 (11.23) HP HP-UX 11iv3 (11.31) IBM AIX 5.2 - *EOL IBM AIX 5.3 - *EOL IBM AIX 6.1 </div>	Adapter (HBA, CNA, etc) <small>Type to filter selections</small> <div style="border: 1px solid #ccc; padding: 2px;"> Brocade 415 Brocade 425 Brocade 815 Brocade 825 Emulex LP10000 </div>

Figure B-1 SSIC for ProtecTIER

Independent Software Vendor Support Matrix

The ProtecTIER Independent Software Vendor Support Matrix is in the ProtecTIER overview at the IBM website. To locate the ProtecTIER ISV Support Matrix, complete these steps:

- Go to the Tape storage web page:
<http://www.ibm.com/systems/storage/tape/index.html>
- Scroll to the Virtual tape systems section and then select any of the ProtecTIER models, as highlighted in Figure B-2.

Virtual tape systems		
Enterprise Virtual	Midrange Virtual	Entry Virtual
<ul style="list-style-type: none"> TS7650G ProtecTIER TS7700 	<ul style="list-style-type: none"> TS7620 ProtecTIER Appliance Express TS7650G ProtecTIER 	<ul style="list-style-type: none"> TS7620 ProtecTIER Appliance Express

Figure B-2 IBM ProtecTIER models

3. Under Learn more (Figure B-3), click **Independent Software Vendor (ISV) matrix**.

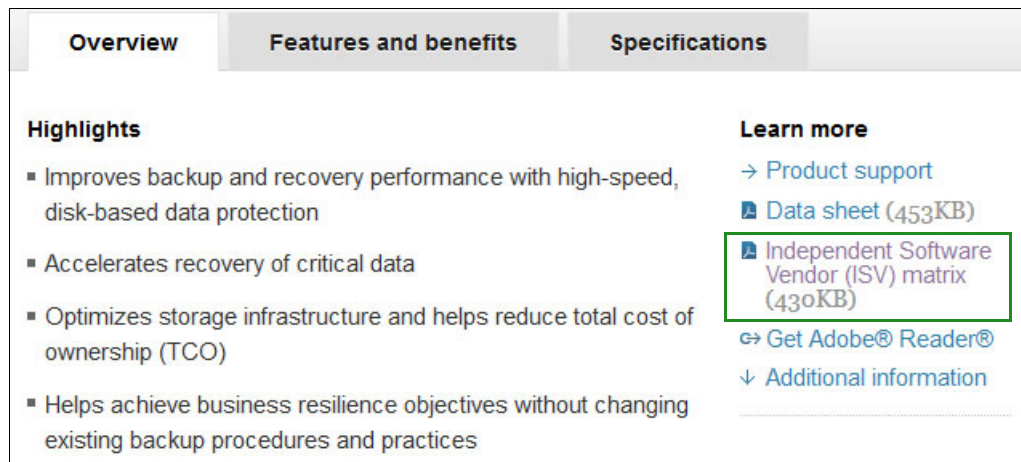


Figure B-3 ProtecTIER ISV support matrix link

The ProtecTIER ISV Support Matrix is divided into the following sections:

- ▶ TS7600 Backup Application ISV (VTL Support)
- ▶ TS7600 Backup Application ISV (FSI-NFS Support)
- ▶ TS7600 Backup Application ISV (FSI-CIFS Support)
- ▶ TS7600 Backup Application ISV (OST Support)
- ▶ TS7600 Mainframe Edition
- ▶ TS7600 NDMP Compatibility
- ▶ TS7600 OEM Cluster MGT Ethernet Switches
- ▶ TS7600 OEM Cluster MGT Fencing Switches
- ▶ TS7600 ProtecTIER Manager Platforms
- ▶ TS7600 Qualified Back-end Disk



ProtecTIER parsers

This appendix describes the ProtecTIER parsers that are used with various backup applications to improve deduplication ratios. It describes terminology, explains how metadata from backup applications hinders deduplication and what the ProtecTIER parser does to reverse this effect, and the supported parsers in ProtecTIER.

Read through this appendix to help you understand what workloads benefit from ProtecTIER parsers, and what are the causes of low deduplication ratios. It also reviews several sample environments and describes whether they benefit from parsers.

This appendix also describes the **analyze_sessions** utility to monitor a ProtecTIER parser.

The ProtecTIER parsers

The ProtecTIER product excels when it finds large *matches* that can be deduplicated. Some common backup applications add metadata, also known as *backup application headers*, to the backup stream for various purposes. This metadata interrupts the large matches, which hinders deduplication. The ProtecTIER parsers separate the metadata from the backup stream dynamically, leaving the user data to deduplicate without interruptions. When the data is restored, the ProtecTIER product adds the metadata back in to the data stream so that the backup application can use it.

The following sections describe terminology, why the ProtecTIER product needs a parser, and the causes of fragmented data, which hinders the matches and ProtecTIER performance and deduplication ratios. They also review several sample environments and describe whether they benefit from parsers.

Terminology

Consider the following terms:

- New data and old data

The users' servers send data to the ProtecTIER server, which deduplicates the data. On arrival at the ProtecTIER server, the data is first passed through the deduplication engine, which searches for the data in the repository. Some data is not found (the *new data*) and some is found (the *old data*).

- Change rate

The *change rate* is the ratio of new data to total data in the backup, that is, the percentage of data that the deduplication engine did not find.

- Old data not factored

The ProtecTIER server examines the old data and might decide that some of it cannot be deduplicated efficiently because doing so would introduce fragmentation and affect the restore performance. The percentage of data in these cases is called *old data not factored*, that is, the percent of data that could not be stored effectively.

- System change rate

In terms of deduplication, both the change rate and old data not factored represent data that was written to the ProtecTIER disk without being deduplicated. Their sum is termed the *system change rate*, that is, the percentage of data that could not be deduplicated because the data was either new to the repository or because deduplicating the data would cause unacceptable fragmentation.

Note: System change rate = change rate + old data not factored.

How metadata from the backup application hinders deduplication

Some backup applications insert metadata into the backup stream for various control purposes by prefixing each block of user data with a small header that includes items such as the sequence number, session identifier, and cartridge bar code.

Figure C-1 shows how a backup application adjusts the sequence of user data by inserting its headers (the backup application metadata) at regular intervals in the tape image.




Figure C-1 Backup applications inserting metadata at regular intervals

The headers are always at regular intervals regarding the start of the tape cartridge. Figure C-2 shows how this spacing affects the tape image if the user data changes its configuration, perhaps because of an insertion of data (for example, the A-tag in Figure C-2). The block B in the original tape image is unchanged, but in the modified tape image, this block is split between B1 and B2, separated by a header.

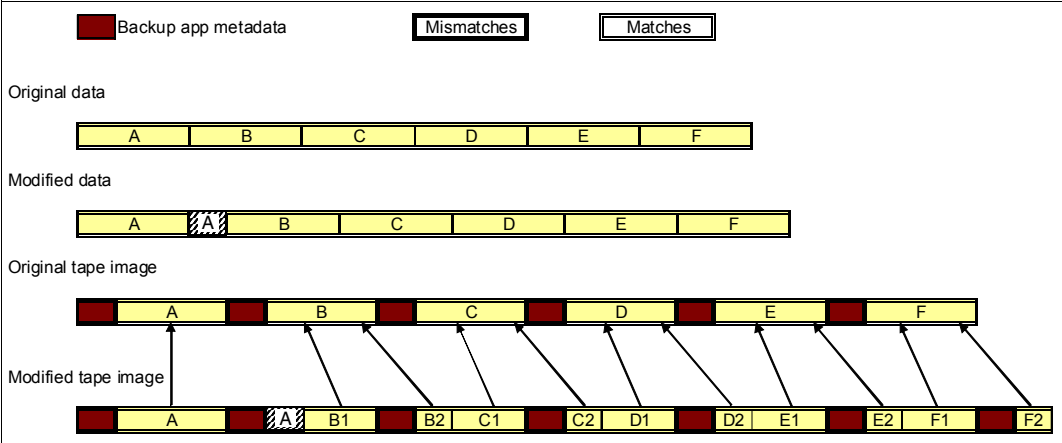


Figure C-2 Top - changes to data; bottom - fragmented data

The deduplication engine finds the B1 and B2 data blocks, but the backup application header between them interrupts their original sequence. The ProtecTIER server needs one pointer for B1 (pointing to part of the original data in B). A second pointer points to where the data of the backup application header is stored (not shown in Figure C-2). A third pointer is needed to point to the location that stores the data of segment B2.

The sequence of backup application headers introduces artificial change. Because the data in the backup image is shifted, it has a ripple effect all the way to the end of the backup or the end of the virtual cartridge. This effect multiplies the number of pointers that are needed to store the new data. Each of these pointers point to smaller data segments than the previous backup. Each generation of the backup potentially amplifies the fragmentation. The cost of adding a pointer to these fragments of old data eventually becomes so high, both in terms of extra I/Os and extra space, that the ProtecTIER server decides not to add the pointers. This old data is merged with adjacent new data and stored again as though it were new. This is how fragmentation leads to old data not factored.

ProtecTIER parser functionality

A ProtecTIER parser reverses the effect of the backup application headers by extracting them from the data stream. The backup data then deduplicates much better against previous backups and other data that is in the repository. When a user asks to restore the data, the ProtecTIER product reinserts the metadata in to the backup stream before it returns it to the backup application.

A ProtecTIER parser examines the start of each backup stream to determine whether the backup application is one that it needs to parse. If the backup application is not identified as needing a parser, the rest of the cartridge is processed normally. If the backup application needs parsing, the parser infrastructure sets up a mechanism in the data path that extracts the backup application headers from the user data. The backup application headers are compressed and stored separately. The remaining user data is now free of the backup application headers and the ripple effects that they cause, and is passed on to the deduplication engine.

The usage of a parser introduces a performance effect that is below 3%.

Deduplication ratios can increase by as much as the value of old data not factored, depending on the characteristics of the data, because the benefit is achieved by avoiding fragmentation. To estimate the expected improvement on your system, see “Estimating the benefit of a parser” on page 466.

ProtecTIER parsers support

Applications like Commvault, Legato, IBM Spectrum Protect (formerly Tivoli Storage Manager), and NetWorker add metadata to the backup stream.

Notes:

- ▶ Starting with version 7.1.3 Tivoli Storage Manager was rebranded to IBM Spectrum Protect.
- ▶ For legacy and reference purposes, we continue to refer to the parser that was built for Tivoli Storage Manager as the *Tivoli Storage Manager parser*.

For VTL models of the ProtecTIER product, the following ProtecTIER parsers are available:

- ▶ Commvault: Parser available since IBM acquired ProtecTIER.
- ▶ Legato parser
 - ProtecTIER V2.3.4 and V2.4.1 introduced a parser for Legato 7.4/7.5/7.6 (May/June 2010)
 - ProtecTIER V2.4.7 and V 2.5.5 introduced a parser for Legato 7.6.1.6 (August 2011)
- ▶ Tivoli Storage Manager): ProtecTIER V3.1.4 introduced a parser for Tivoli Storage Manager V5.5 and later (October 2011).

For CIFS models of ProtecTIER servers, the following ProtecTIER parsers are available:

- IBM Spectrum Protect (formerly Tivoli Storage Manager).

Important: When you use the Tivoli Storage Manager parser with a CIFS model of the ProtecTIER product, use the **DATAFormat = NATive** option with the disk storage pool definitions. This value is the default value and is the format that the CIFS ProtecTIER Tivoli Storage Manager parser recognizes. Do not use the **DATAFormat-NONblock** option.

- Commvault

A Legato parser was not written for CIFS because Legato settings make it unnecessary. Legato users should choose *Advanced File* rather than *File* disk storage when setting up the CIFS share on a ProtecTIER server. This setting does not require a parser.

Backup applications and parsers: Sometimes, although rarely, new versions of backup applications introduce changes to the format of headers that they insert in the tape stream. This situation might cause the ProtecTIER parser to miss headers. Although this situation does not risk backup data in any way, it can cause deduplication to drop as the old data not factored increases. ProtecTIER Quality Assurance monitors the new versions of these backup applications, so check with ProtecTIER Support before you upgrade backup applications. Support can advise whether there is a change in the efficiency of the parser at that level.

What workloads benefit from the ProtecTIER parsers

The deduplication ratio that is achieved by a backup is affected by two key factors:

- The change rate of the data (*change rate*): This rate depends on how much the data in the backup changes from day to day and some of the parameters that are set as part of the backup (for example, encryption, compression, and multiplexing).
- The amount of data that is not factored (*old data not factored*): This rate can be high in environments where changes are interspersed throughout the data, causing fragmentation, such as might happen with a backup application that inserts metadata.

The ProtecTIER parsers are designed to reduce the amount of old data not factored and therefore increase the deduplication ratio in environments where the backup application inserts metadata in the backup stream. A system with a larger percentage of data in this category (for example, 15%) benefits more from adding a ProtecTIER parser than a system with a smaller amount of data in this category (for example, 3%).

High change rate: If the system also has a high change rate, the benefit from the parser might be less noticeable.

Workloads that achieve lower deduplication ratios because of a high change rate should look at other avenues to increase the deduplication rate.

Background information: Causes of low deduplication ratios

There are two bases to good deduplication:

- ▶ The first and most important is multiple copies of the data. If multiple copies of the data do not exist, deduplication does not occur.
- ▶ The second is the similarity between (successive) copies. If the copies are similar (low-change rate), they deduplicate well.

Poor deduplication is caused by not enough copies or too much change between copies. However, sometimes successive copies are similar but the changes are small and evenly distributed. In this case, the storage system cannot effectively store just the changes, but must rewrite some of the common data as though it were new. This phenomenon is measured by old data not factored.

The objective of parsers in the ProtecTIER environment is to focus on the *old data not factored* that is caused by the backup application headers in the backup stream. By removing these headers and storing them separately, the ProtecTIER parsers remove a cause of small and evenly distributed change that interferes with the underlying user data. This action prevents some of the fragmentation in the data and the extra pointers that are needed to track the data.

Often, the number of copies of the data (*retention*) is set by company policy. The old data not factored is controlled by characteristics of the actual and imposed changes. If the deduplication ratio is still low, then reducing the change rate is the best place to concentrate your efforts. Many causes of a high change rate exist. Some common causes of high-change rate are as follows:

- ▶ Multiplexed backups: Some backup applications intermix backups from multiple sources in to the same (virtual) tape image. This situation causes the deduplication engine to try to search too many different sources for common data. A similar phenomenon occurs if the order of files in the backup does not remain stable.
- ▶ Backup of compressed or encrypted data: Sometimes unchanged compressed files deduplicate by using their previous backup as a reference, but in general, compressed or encrypted data does not deduplicate well.
- ▶ Cross-deduplication between files with similar content but a different layout: This situation can happen when different VM images are backed up. Each VM has similar files in it, but their layout in each VM's file system is different.
- ▶ Files with common data but high internal change: Some applications occasionally reorganize, defragment, or reindex their data. Backups that are made immediately after these operations are not likely to deduplicate well.

In extreme cases, the ProtecTIER product has a set of parameters that control the deduplication engine and the storage system. ProtecTIER Level 3 support can suggest parameter changes to help in cases of poor deduplication.

Estimating the benefit of a parser

The guidance in this section can help you estimate the benefit of a parser.

Consider a hypothetical system with a change rate of 20% and old data not factored of 13.3% before you implement the parser. This situation means that the system change rate is 33.3% and the user sees one-third of the data change each backup, or two-thirds of the data deduplicated. Such a system might reach a deduplication ratio of 3:1 (before compression) if sufficient backup copies are retained. If the parser reduces the old data not factored to 5%,

then the system change rate drops to 25%, and the user sees one-quarter of the data change each backup. The system might now reach a 4:1 deduplication ratio (before compression) if sufficient backup copies are retained. If compression achieves a 2:1 ratio, then this system can improve from 6:1 to 8:1 ratio.

ProtectTIER Support can review a system and provide the figures for change rate and old data not factored to help with this calculation.

Environments that benefit from parsers

The following four cases illustrate scenarios in two application environments: NetBackup and Tivoli Storage Manager. The cases that are provided demonstrate environments where a parser would be helpful in improving deduplication ratio and sometimes performance. Each case has a graph that shows the change rate and old data not factored values. The graphs were created by using data from a ProtecTIER server and a spreadsheet graphing tool.

Case 1: Low change rate and low old data not factored

This case shows a customer who is using NetBackup. The workload has a low change rate and low old data not factored values Figure C-3. This workload does not need a parser.




Figure C-3 Case 1 - low change rate and low old data not factored

Case 2: Moderate change rate and high old data not factored

Figure C-4 shows a case where a parser is effective. The backup application is Tivoli Storage Manager. The change rate is moderate, but the old data not factored is high.




Figure C-4 Case 2 - moderate change rate and high old data not factored

Case 3: High change rate, moderate to low old data not factored

In Figure C-5, the change rate is high and the old data not factored is moderate to low. A parser might help, but the benefit that it offers is marginalized by the high change rate. In this case, the best action might be to look for causes of the high change rate in the environment and try to reduce that value.




Figure C-5 Case 3 - high change rate, moderate to low old data not factored

Case 4: High change rate, low old data not factored

Taking Case 3 to an extreme, Figure C-6 shows a ProtecTIER installation that reports a high change rate and a low old data not factored. The Tivoli Storage Manager parser does not help in this situation, and the best action is to look for causes of the high change rate in the environment.

Note: Starting with version 7.1.3 Tivoli Storage Manager was rebranded to IBM Spectrum Protect. For legacy and reference purposes, we continue to refer to the parser that was built for Tivoli Storage Manager as the *Tivoli Storage Manager parser*.




Figure C-6 Case 4 - high change rate, low old data not factored

Experience from one user site

Figure C-7 tracks the change in old data not factored of a ProtecTIER installation. It shows about four weeks of history while the site was running a ProtecTIER version with no parser, and about six weeks after upgrading to a ProtecTIER version with a parser. Over a period of one week after the upgrade, the old data not factored dropped fairly rapidly from 6% to 3% (approximately).




Figure C-7 Change in old data not factored at a user site after implementing a ProtecTIER parser

Not every site upgrading to a ProtecTIER release with a parser sees such an immediate improvement in the old data not factored. The actual improvement depends on many factors outside of the ProtecTIER product.

Use the `analyze_sessions` utility to monitor the benefit of a ProtecTIER parser

You can monitor changes in the deduplication ratio at a fine-grain level, whether it is at a per cartridge, per session, daily, or even hourly level.

The `analyze_sessions` utility assists with understanding the deduplication of the workloads sent to the ProtecTIER. It examines the deduplication logs of the ProtecTIER cluster and produces a report with the deduplication split (by default) into sessions. The definition of a session is a period of system ingest activity. A session ends when there is a short break in the I/O, during which the ProtecTIER product is idle. There can be therefore one session that covers a multiple hour period, or many sessions in an hour.

The `analyze_sessions` utility is in the `/opt/dtc/app/utls` directory on a ProtecTIER server. This utility examines the deduplication logs in their default location in the repository of the ProtecTIER server and generates a report on the deduplication that is achieved by the ProtecTIER cluster. The report is in CSV format and is placed in the `/pt_work` directory. The report is best viewed by using a spreadsheet program, which can also be used to plot the deduplication ratios of the sessions and their trends.

Several parameters can be used with **analyze_sessions**:

- n <number>** Report only <number> of months back.
- s <start date>** Start reporting only from date (date format is YYYY-MM-DD-HH-MM).
- e <end date>** Report only up until the end date.
- sd <directory>** The directory that contains the deduplication logs (if it is not the default location).
- min** Report only overall statistics.
- daily** Include a daily summary.
- hourly** Include an hourly summary.
- d <minutes>** The number of minutes of idle time that defines a session boundary.
- i <file of carts>** The file that contains a list of cartridge bar codes. Only these bar codes are reported.
- c** Used with **-i**. Create a separate file for each cartridge.
- l** Used with **-i**. Add a line to the report for each cartridge.
- o <output>** Specify an output file name other than the default.

Figure C-8 shows the default output of the **analyze_sessions** utility. The important fields are the system change rate (the effect of only the deduplication, not the compression), and the compressedBytesCount (the actual amount of data that is written to the repository by this session). The deduplication of a particular session (regarding all the data already in the repository) is column C divided by column F, and includes the effects of deduplication and compression of the data.

	A	B	C	D	E	F	G	H
11	Name	Total data (TB)	Total data (GB)	System change rate	Factoring ratio	compressedBytesCount (GB)	start time	end time
12								
13	Grand totals							
14	all	15.2996	15666.8	37.27%	2.68336	5248.25	14/08/2011 10:39	04/07/2012 11:19
15								
16	By session (summary)							
17	2011-8-14 10:39:35 to 2011-8-14 11:03:10	0.0747367	76.5304	38.20%	2.61775	27.4528	14/08/2011 10:39	14/08/2011 11:03
18	2011-8-14 11:09:19 to 2011-8-14 11:12:39	0.00152501	1.56161	99.21%	1.00797	1.45244	14/08/2011 11:09	14/08/2011 11:12
19	2011-8-14 12:01:39 to 2011-8-14 12:31:29	0.0762905	78.1214	39.04%	2.56135	28.6391	14/08/2011 12:01	14/08/2011 12:31
20	2011-8-14 12:41:28 to 2011-8-14 13:05:06	0.0747367	76.5304	39.57%	2.52739	28.4324	14/08/2011 12:41	14/08/2011 13:05
21	2011-8-14 13:11:26 to 2011-8-14 13:11:26	0.000517929	0.530359	97.70%	1.02354	0.485788	14/08/2011 13:11	14/08/2011 13:11
22	2011-8-14 13:17:47 to 2011-8-14 13:17:47	0.000517929	0.530359	97.66%	1.02395	0.485602	14/08/2011 13:17	14/08/2011 13:17

Figure C-8 Output from the **analyze_sessions** command

Planning for the Tivoli Storage Manager parser

Note: Starting with version 7.1.3 Tivoli Storage Manager was rebranded to IBM Spectrum Protect. For legacy and reference purposes, we continue to refer to the parser that was built for Tivoli Storage Manager, the *Tivoli Storage Manager* parser.

When a new parser is added to a ProtecTIER system, a slight performance degradation (less than 3%) might occur.

Initially, a temporary and slight decrease in deduplication ratio might occur and an increase in the space that is used in the repository because the newly parsed user backups might not match as well with the existing unparsed backups in the repository. After a few generations, a new steady state is reached, where parsed user backups deduplicate by using other parsed user backups as their reference for deduplication.

Because there are no interfering headers, old data not factored is expected to decline, causing the deduplication ratio to reach an improved level. Consequently, used space in the repository drops. Performance might also be improved because the ProtecTIER product processes data more efficiently when the deduplication level is higher.

For environments where the backup window is tight or the repository is full, clients should work with ProtecTIER Support before they load the first ProtecTIER release that contains the applicable parser. ProtecTIER Support has tools to analyze problem reports that give an historic breakdown of actual change rate and old data not factored. These figures can help you understand both the startup effect and how effective the parser is in improving deduplication ratio over time.



D

Managing cartridge sizes with ProtecTIER

This appendix provides general information about managing cartridge sizes with ProtecTIER. With ProtecTIER, the total amount of space that is available in your repository is variable. These dynamics can make managing cartridges complex. This appendix is intended to help administrators to plan and manage their ProtecTIER cartridges and repository.

This appendix describes the following topics:

- ▶ Effects of dynamic cartridge sizes
- ▶ The mechanism behind fluctuating cartridge sizes

Effects of dynamic cartridge sizes

In a long running ProtecTIER environment, the HyperFactor ratio tends to stabilize around a certain value. If you then encounter changes in your environment, these changes have an unexpected effect on the HyperFactor ratio. When the HyperFactor ratio changes, the amount of free space that is available in your repository is recalculated. A higher HyperFactor ratio results in more free space being available. A lower HyperFactor ratio results in less free space being available.

After such events, you might find difficulty in determining how much free space is available. Also, you might see virtual cartridges that are marked as full before such an event is expected.

The mechanism behind fluctuating cartridge sizes

All virtual cartridges get an equal amount of space up to either the LIMIT size that is configured when the virtual cartridge was created (maximum cartridge size), or the calculated amount of Nominal Space/Number of Carts.

If the nominal space is not large enough to hold the total of Number of Carts times the LIMIT size of the cart, then those tapes are marked full before they reach their size (that is, there is no thin provisioning).

Knowing this behavior can help an administrator know when best to add more cartridges and whether to limit the size of the carts or leave them unlimited.

Deciding on a limit for the virtual cartridge size enables you to reserve space or divide space in libraries. With this strategy, you can prevent one of your virtual libraries from using up all the space in your repository, causing the other libraries to run out of space.

For example, if you have two libraries, one with 500 carts and the other with 1000 carts, your usage is 33% for the smaller library, and 66% for the larger library.

The decision about how many carts you have should be based on the nominal space in the repository. The number of carts should be either calculated for a wanted approximate size (unlimited) or calculated so that there is enough space for all carts to reach the limit that is set, while leaving a little room for fluctuation (the factoring ratio).

Having the wrong cartridge size has led to problems in the following cases:

- You might think that because you have low scratch tapes that adding scratch tapes at a limited size brings more space to the backup application. This situation works if there is a reserve of nominal space in the repository. However, if the factoring ratio is lower than planned, and less space results, adding cartridges results in an even smaller cartridge size.

Important: Adding more virtual cartridges does not increase the available free space in your repository. If you face an out-of-space condition, you must expand your ProtecTIER repository by adding more physical disk to the back end.

- An insufficient idle time for background jobs results in a build-up of delete and defragmentation data, which reduces allocatable space for backups. Adding cartridges in this scenario can reduce the size of the cartridges, which has a negative effect.

- ▶ An extra library with 1000 extra empty tapes results in the production library running out of room, even though the repository showed plenty of nominal space. Adding tapes in this case results in even smaller tapes as well.
- ▶ Collocation of backups with many partial tapes might make the repository appear to have much space, but the partial tapes use more space than expected because of the size and usage. In this case, limiting the tapes to a smaller size could enable the tapes to store data more efficiently.
- ▶ Finally, if the repository has more space available when you multiply the total number of tapes by the amount of space, then it is a good plan to add more limited cartridges (This situation will not happen if you do not have any limited size tapes.)

If you have a mixture of tape sizes and types, the management of these tapes becomes complicated, especially if you have a large variance in the sizes. If you want to “unlimit” all of your existing tapes, contact IBM Support and request help with running the support utility to “unlimit” the cartridge sizes.

Glossary

3958 DD1 This original server has been available since August 2008. This server is based on the IBM System x3850 M2 Type 7141. When it is used as a server in the TS7650G, its machine type and model are 3958 DD1. This machine type and model is no longer supported.

3958 DD3 This is a higher performance server, which has been available since March 2009. This server is based on the IBM System x3850 M2 Type 7233. When used as a server in the TS7650G, its machine type and model are 3958 DD3. This machine type and model is no longer supported.

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

3958 DD5 This is a higher performance server, which has been available since May 2012. This server is based on the IBM System x7145 model. When used as a server in the TS7650G, its machine type and model are 3958 DD5. Use this machine type and model for service purposes.

3958 DD6 This is the newest server available since March 2016, it was introduced with ProtecTIER Version 3.4. This server is a 2U platform based on two Intel Xeon Processor E5-2695 v2 (12 cores, 2.4 GHz), it comes with 128 GB RAM (8x 16 GB DDR3 at 1866 MHz).

asynchronously parallel system A system in which the backed up data does not return to the host (and out to file) in the same order each time.

Backup, Recovery, and Media Services for IBM i (BRMS) Helps you implement a disciplined approach to managing your backups, and provides you with an orderly way to retrieve lost or damaged data. BRMS also enables you to track all of your backup media from creation to expiration.

BRMS policies A set of defaults that is commonly used (for example, device or media class). Generally used defaults are in the BRMS system policy. Backup-related defaults are in the BRMS backup policy.

chown The **chown** command (abbreviation for **change owner**) is used on UNIX based systems to change the owner of a file. In most implementations, it can be run by only the superuser to prevent users from simply changing the ownership of files randomly.

compaction (data compaction) The reduction of the number of data elements, bandwidth, cost, and time for the generation, transmission, and storage of data without loss of information by eliminating unnecessary redundancy.

Common Internet System (CIFS) ProtecTIER emulates Windows file system behavior and presents a virtualized hierarchy of file systems, directories, and files to Windows CIFS clients. When configured for FSI-CIFS, ProtecTIER emulates a network-attached storage (NAS) backup target that can use both HyperFactor and ProtecTIER native replication bandwidth reduction techniques for storing and replicating deduplicated data.

concurrent saves and restores The ability to save or restore different objects from a single library or directory to multiple backup devices or different libraries or directories to multiple backup devices at the same time from different jobs.

control group A group of items (for example, libraries or stream files) to back up, and the attributes that are associated with how to back them up.

CSV file (.csv) Comma-separated value file, sometimes called comma-delimited. This type of file is a specially formatted plain text file that stores spreadsheet or basic database-style information in a simple format, with one record on each line, and each field in that record separated by a comma. CSV files are used by ProtecTIER Manager as a simple way to transfer a large volume of database information between programs. This type of file can be imported into most spreadsheet programs.

deduplication A data compression technique in which redundant data is eliminated. The technique improves storage usage and can also be applied to network data transfers to reduce the number of bytes that must be sent across a link.

direct attachment Refers to a digital storage system that is directly attached to a server or workstation, without a storage network in between.

dirty bit A dirty bit is a flag that indicates whether an attribute must be updated. This situation occurs when a bit in a memory cache or virtual memory is changed by a processor but is not updated in storage.

dirty bit technology The ProtecTIER system uses a “dirty-bit” feature/technology and cartridges are marked as in-sync after the data finishes replicating from the primary to the secondary site, so that at the time of synchronization, the local cartridges and their DR site replicas are identical.

disaster recovery (DR) The process of recovering production site data at a DR location. Disaster recovery is useful if a disaster occurs or a situation occurs where the production (or primary) site goes offline.

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

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

File System Interface (FSI) The File System Interface (FSI) presents ProtecTIER as a network-attached storage backup and recovery target that can use the HyperFactor algorithm and ProtecTIER native replication bandwidth reduction techniques for storing and replicating deduplicated data. The FSI configuration option enables ProtecTIER to present disk repository storage as a virtualized hierarchy of file systems.

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

failback The procedure for replicating updated cartridges, new or old, from the DR site to the original (or restored) production site to bring it up to date in case it was down, or lost and rebuilt.

hub The hub (target server) is connected, through your Ethernet replication network, to one or more spokes (source servers). The hub stores backup images of the data repositories, file system configurations, and other system settings that are on the spokes. If there is a spoke failure, the stored image can be easily retrieved from the hub and transmitted to the spoke. This action restores the spoke to its previous configuration with minimal data loss.

IBM Tivoli Assist On-site (AOS) IBM Tivoli Assist On-site (AOS) is a web-based tool that enables a remote support representative from IBM to view or control the management node desktop. For more information, see the Assist On-site web page:
<http://www.ibm.com/support/assistsite>

IP address Internet Protocol address. A numerical label that is assigned to each device that participates in a computer network.

load throttling Load throttling is a process that helps avoid dangerous overload situations. Load throttling limits the number of permitted incoming connections, enabling resources to be allocated to all processes.

logical partition (LPAR) A division of a computer's processors, memory, and storage into multiple sets of resources so that each set of resources can be operated independently with its own operating system instance and applications. The number of LPARs that can be created depends on the system's processor model and resources that are available.

logical unit number (LUN) A number that is used to identify a logical unit that is a device that is addressed by Fibre Channel. A LUN can be used with any device that supports read/write operations, such as a tape drive, but is most often used to refer to a logical disk that is created on a SAN.

LUN masking An authorization process that makes a LUN available to some hosts and unavailable to other hosts. LUN masking is used in the ProtecTIER product as a precaution against servers corrupting disks that belong to other servers. By masking (hiding) LUNs from a specific server (or servers), you effectively tell those servers that the LUN does not exist, and those servers cannot corrupt the disks in the LUN.

media A tape cartridge (volume) that holds saved data.

media class A logical grouping of media with similar physical, logical, or both of these characteristics (for example, density).

media identifier A name that is given to a physical piece of media.

Network File System (NFS) ProtecTIER emulates UNIX file system behavior and presents a virtualized hierarchy of file systems, directories, and files to UNIX based clients using the NFS protocol. When configured for FSI-NFS, ProtecTIER emulates a network-attached storage (NAS) backup target that can use both HyperFactor and ProtecTIER Native Replication bandwidth reduction techniques for storing and replicating deduplicated data.

nominal capacity of the repository The physical space and expected factoring ratio.

nominal data The original amount of backed-up data before you apply the ProtecTIER deduplication factor.

OpenStorage (OST) Enables the ProtecTIER product to be integrated with NetBackup to provide backup-to-disk without using a Virtual Tape Library (VTL) emulation. Note: The OST interface is no longer supported with version 3.4.

parallel saves and restores The ability to save or restore a single object or library or directory across multiple backup devices from the same job.

peak throughput The maximum of the ProtecTIER server capabilities.

principality The privilege to write to a cartridge (set to read/write mode). The principality of each cartridge belongs to only one repository in the grid. By default, the principality belongs to the repository where the cartridge was created.

ptcli ProtecTIER command-line interface.

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

RAID 10 Many storage controllers enable RAID levels to be nested. The elements of a RAID can be either individual drives or RAID's themselves. Therefore, a RAID 10 (or RAID 1+0) is a configuration in which multiple drives are first combined into multiple RAID arrays. Each RAID 1 array is treated as a single drive. These arrays are then combined into a single RAID 0 array.

RAID 6 The RAID 6 architecture supports block-level striping with double distributed parity. It provides fault tolerance of two drive failures; the array continues to operate with up to two failed drives. This situation makes larger RAID groups more practical, especially for high-availability systems. This situation becomes increasingly important as large-capacity drives lengthen the time that is needed to recover from the failure of a single drive. Single-parity RAID levels are vulnerable to data loss until the failed drive is replaced and its data rebuilt. Double parity gives more time to rebuild the array without the data being at risk if another drive fails before the rebuild is complete.

reclamation The Tivoli Storage Manager process that frees up space on tapes, and returns empty tapes to the scratch pool. Reclamation is accomplished by deleting expired data from tapes and moving any unexpired data to other tapes to more efficiently use tape space.

recovery point objective (RPO) How much lag time is acceptable for a backup that is written to virtual tape in Site A to be replicated to Site B.

redundant array of independent disks (RAID) A storage technology that combines multiple disk drive components into a logical unit. Data is distributed across the drives in one of several ways (RAID levels). The physical drives are said to be in a RAID array, which is accessed by the operating system as one single drive. The different schemes or architectures are named by the word RAID followed by a number (for example, RAID 0 or RAID 1).

redundant host connection The duplication of connections, or two or more paths that connect two objects in a network. The intention of redundancy is to increase the reliability of the system, or to provide a backup or failsafe if one of the connections fails.

remote cloning The process of using a secondary (DR) site to clone cartridges. ProtecTIER replication enables users to offload tape cloning to their secondary site.

replication A process that transfers logical objects, such as cartridges, from one ProtecTIER repository to another one.

replication grid A set of repositories that shares a common ID and can potentially transmit and receive logical objects through replication.

replication grid ID A number (0 - 63) that identifies a replication grid in an organization.

replication grid member A repository that is a member in a replication grid.

Replication Manager The utility in the ProtecTIER Manager GUI through which you can set replication policies, define replication time frame windows, delete replication activities, and much more.

replication pairs Two repositories in a replication grid that replicate from one to another.

replication policy A policy made up of rules that define a set of objects (for example, VTL cartridges) from a source repository to be replicated to a target repository.

replication rate control (RRC) A built-in resource-governing mechanism. RRC gives precedence to backup and restore requests and throttles down replication traffic whenever backup and restore activity increases above an idle state. When the backup and restore workload drops below that idle threshold, RRC returns to the default priority. The RRC determines the maximum replication rate for both system states, IDLE and BUSY, based on the performance limits set by the user.

replication time frame A scheduled period for replication to take place for all policies.

replication window The time frame during which replication runs.

repository A warehouse to store data for safekeeping.

repository unique ID (RID) A number that uniquely identifies the repository. The RID is created from the replication grid ID and the repository internal ID in the grid.

SAN fabric The hardware that connects workstations and servers to storage devices in a SAN is referred to as a fabric. The SAN fabric enables any server to connect to any storage device through Fibre Channel switching.

shelf A container of VTL cartridges in a ProtecTIER repository.

SLA Service level agreement.

spoke The spokes are the servers that process and store the information that is generated during daily business operations. The stored information is then replicated to a hub, according to a user-defined replication policy.

storage area network (SAN) A dedicated network that provides access to consolidated, block-level data storage. SANs make storage devices, such as disk arrays and tape libraries, accessible to servers so that the devices appear similar to locally attached devices to the operating system.

storage checkpoint Storage checkpoints are pointers added to the backup stream so that if the backup fails, a rerun of the backup starts from the last storage checkpoint, rather than the beginning of the stream.

system console The system console is an IBM TS3000 System Console (TSSC). It is the console used with a keyboard to issue commands to the server through the CLI, and monitors the functions of the server. The TS3000 combined with the keyboard is commonly referred to as the KVM (keyboard, video, monitor).

Note: The TS7650G DD6 and the TS7620 ProtecTIER Appliance Express do not support TSSC.

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

Virtual Tape Library (VTL) The ProtecTIER VTL service that emulates traditional tape libraries.

visibility switching The automated process that transfers the visibility of a VTL cartridge from its master to its replica and vice versa.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Some publications referenced in this list might be available in softcopy only:

- ▶ *Harnessing the Power of ProtecTIER and Tivoli Storage Manager*, SG24-8209
- ▶ *IBM DS8870 Architecture and Implementation (Release 7.5)*, SG24-8085
- ▶ *IBM DS8880 Architecture and Implementation (Release 8)*, SG24-8323
- ▶ *IBM Spectrum Accelerate Deployment, Usage, and Maintenance*, SG24-8267
- ▶ *IBM System Storage TS7600 with ProtecTIER Version 3.3*, SG24-7968
- ▶ *IBM System Storage Solutions Handbook*, SG24-5250
- ▶ *IBM Tape Library Guide for Open Systems*, SG24-5946
- ▶ *IBM Tivoli Storage Manager Implementation Guide*, SG24-5416
- ▶ *IBM XIV Storage System: Host Attachment and Interoperability*, SG24-7904
- ▶ *Implementing IBM FlashSystem 900*, SG24-8271
- ▶ *Implementing IBM Storage Data Deduplication Solutions*, SG24-7888
- ▶ *Implementing the IBM Storwize V7000 and IBM Spectrum Virtualize V7.6*, SG24-7938
- ▶ *Implementing the IBM Storwize V7000 Unified Disk System*, SG24-8010
- ▶ *Implementing the IBM System Storage SAN Volume Controller with IBM Spectrum Virtualize V7.6*, SG24-7933
- ▶ *Rapid Data Backup and Restore Using NFS on IBM ProtecTIER TS7620 Deduplication Appliance Express*, TIPS-0990
- ▶ *SAP Backup using Tivoli Storage Manager*, SG24-7686
- ▶ *Tivoli Storage Manager as a Data Protection Solution*, SG24-8134
- ▶ *XIV Storage System Product Guide*, REDP-5272

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Other publications

These publications are also relevant as further information sources.

Publications common to the TS7650 Appliance and TS7650G

- ▶ *IBM TS7650G ProtecTIER Deduplication Gateway Introduction and Planning Guide*, GA32-0918
- ▶ *IBM TS7650 with ProtecTIER for ProtecTIER v3.4, Software Upgrade Guide*, SC27-3643
- ▶ *IBM TS7650G ProtecTIER Deduplication Gateway Installation Roadmap Guide*, GA32-0921
- ▶ *IBM TS7650G ProtecTIER User's Guide for VTL Systems*, GA32-0922
- ▶ *IBM TS7650G ProtecTIER Deduplication Gateway Problem Determination and Service Guide*, GA32-0923
- ▶ *Best Practices Guide for IBM System Storage TS7600 ProtecTIER Deduplication Solutions*, GA32-0924
- ▶ *IBM TS7650G Deduplication Gateway for ProtecTIER V3.3.7, User's Guide for FSI Systems*, GA32-2235

TS7620 Appliance Express publications

- ▶ *IBM System Storage TS7620 ProtecTIER Deduplication Appliance Express Introduction and Planning Guide*, GA32-0913
- ▶ *IBM System Storage TS7610 and TS7620 ProtecTIER Deduplication Appliance Express ProtecTIER User's Guide for VTL Systems*, GA32-0916
- ▶ *IBM System Storage TS7610 and TS7620 - ProtecTIER Deduplication Appliance Express ProtecTIER User's Guide for OpenStorage*, GA32-2230
- ▶ *IBM System Storage TS7610 and TS7620 - ProtecTIER Deduplication Appliance Express - ProtecTIER User's Guide for FSI and NFS*, GA32-2231
- ▶ *IBM System Storage TS7610 and TS7620 with ProtecTIER Version 3.3.5, Software Upgrade Guide V3.3.5, PN 12X5236 EC M12998S, C27-3641*
- ▶ *IBM System Storage TS7620 ProtecTIER Deduplication Appliance Express Installation and Setup Guide for VTL, and OpenStorage Systems*, GA32-0914
- ▶ *IBM System Storage TS7620 ProtecTIER Deduplication Appliance Express V3.3 Increasing Capacity on the 3959 SM2 from 6 TB to 12 TB (Feature Code 9317)*, GA32-2222
- ▶ *IBM System Storage TS7620 ProtecTIER V3.3 Deduplication Appliance Express - Feature Code 9345 (3959 EXP), Field Installation of Expansion Drawer*, SC27-5413

Online resources

These pages in the IBM Knowledge Center are relevant as further sources of information:

- ▶ TS7650 V3.4, IBM Systems Storage TS7650, V3.4 documentation:
<http://pic.dhe.ibm.com/infocenter/ts7650/cust/index.jsp>
- ▶ Version 3.4 of the IBM System Storage TS7650 with ProtecTIER:
http://www.ibm.com/support/knowledgecenter/STLPPL_ent/com.ibm.storage.ts7650_3-4.service.doc/welcome/ts7650_ic_service_welcome.html
- ▶ IBM SAN Volume Controller:
<http://www.ibm.com/support/knowledgecenter/STPVGU/welcome>
- ▶ IBM Storwize V7000:
<http://www.ibm.com/support/knowledgecenter/ST3FR7/welcome>
- ▶ IBM Storwize V7000 Unified Storage:
http://www.ibm.com/support/knowledgecenter/ST5Q4U/landing/v7000_unified_welcome.htm
- ▶ IBM Tivoli Storage Manager V7.1 product suites and related products:
http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.5/tsm/r_related_products.html
- ▶ IBM XIV Storage System:
http://www.ibm.com/support/knowledgecenter/STJTAG/com.ibm.help.xivgen3.doc/xiv_kcwelcomepage.html
- ▶ IBM DS8870:
http://www.ibm.com/support/knowledgecenter/ST8NCA/product_welcome/ds8000_kcwelcome.html

These web pages are also relevant as further information sources:

- ▶ IBM Tape Device Drivers Installation and User's Guide, GC27-2130:
<http://www.ibm.com/support/docview.wss?rs=577&uid=s5sg1S7002972>
- ▶ IBM TS7620 ProtecTIER Deduplication Appliance Express overview and specifications:
<http://www.ibm.com/systems/storage/tape/ts7620/index.html>
- ▶ IBM TS7650G ProtecTIER Deduplication Gateway overview and specifications:
<http://www.ibm.com/systems/storage/tape/ts7650g/index.html>
- ▶ List of supported Fibre Channel switches:
<http://www.ibm.com/systems/support/storage/config/ssic/index.jsp>
- ▶ Red Hat Enterprise Linux:
<https://www.redhat.com/wapps/store/allProducts.html>
- ▶ IBM System Storage Interoperation Center (SSIC):
<http://www.ibm.com/systems/support/storage/config/ssic/index.jsp>
- ▶ Veritas NetBackup:
<https://www.veritas.com/product/backup-and-recovery/netbackup>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



IBM ProtecTIER Implementation and Best Practices Guide

SG24-8025-03
ISBN 0738441694



(1.0" spine)
0.875" <-> 1.498"
460 <-> 788 pages



SG24-8025-03

ISBN 0738441694

Printed in U.S.A.

Get connected

