


# Implementing IBM Storage Data Deduplication Solutions

Fitting deduplication into your  
enterprise environment

Adding deduplication to your  
storage hierarchy

Gaining the space benefits of  
IBM deduplication



Alex Osuna  
Eva Balogh  
Alexandre Ramos Galante de Carvalho  
Rucel F. Javier  
Zohar Mann

**Redbooks**





International Technical Support Organization

**Implementing IBM Storage Data Deduplication  
Solutions**

March 2011

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

**First Edition (March 2011)**

This edition applies to Version 2.5 of ProtecTIER, versions 6.1 and 6.2 of IBM Tivoli Storage Manager and Data ONTAP 7.3.4 for N series.

**© Copyright International Business Machines Corporation 2011. All rights reserved.**

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



# Contents

<b>Notices</b> .....	ix
Trademarks .....	x
<b>Preface</b> .....	xi
The team who wrote this book .....	xi
Now you can become a published author, too! .....	xiii
Comments welcome .....	xiii
Stay connected to IBM Redbooks .....	xiv
<b>Part 1. Introduction to Deduplication</b> .....	1
<b>Chapter 1. Introduction to Deduplication</b> .....	3
1.1 IBM data reduction and deduplication .....	4
1.2 Deduplication overview .....	4
1.2.1 Chunking .....	5
1.2.2 Processing .....	6
1.2.3 Consolidation .....	6
1.3 Architecture .....	6
1.3.1 Where deduplication processing occurs .....	6
1.3.2 When deduplication occurs .....	7
1.4 Benefits of data deduplication .....	8
<b>Chapter 2. Introduction to N series Deduplication</b> .....	9
2.1 How deduplication for IBM System Storage N series storage system works .....	10
2.2 Deduplicated Volumes .....	13
2.3 Deduplication metadata .....	14
2.4 Sizing for performance and space efficiency .....	15
2.4.1 Deduplication general best practices .....	15
2.5 Compressing and Deduplicating .....	17
<b>Chapter 3. Introduction to ProtecTIER deduplication</b> .....	19
3.1 Terms and definitions .....	20
3.2 Overview of HyperFactor technology and deduplication .....	22
3.3 IBM System Storage TS7600 with ProtecTIER .....	24
3.3.1 TS7610 - Entry Edition .....	24
3.3.2 TS7650 - Appliance Edition .....	24
3.3.3 TS7650G - Enterprise Edition .....	24
3.3.4 TS7680 - Gateway Edition for System z .....	24
3.4 Deduplication with HyperFactor technology .....	25
3.4.1 Impact of HyperFactor .....	25
3.4.2 ProtecTIER data ingest flow .....	26
3.4.3 ProtecTIER VTL concepts .....	28
3.4.4 ProtecTIER OST concepts .....	28
3.4.5 Steady state .....	29
<b>Chapter 4. Introduction to IBM Tivoli Storage Manager deduplication</b> .....	31
4.1 Deduplication overview .....	32
4.2 How ITSM data deduplication works .....	33
4.3 IBM Tivoli Storage Manager deduplication overview .....	34

4.3.1	ITSM server-side deduplication overview . . . . .	34
4.3.2	ITSM client-side deduplication overview . . . . .	35
<b>Part 2.</b>	<b>Planning for deduplication. . . . .</b>	<b>39</b>
<b>Chapter 5.</b>	<b>N series deduplication planning . . . . .</b>	<b>41</b>
5.1	Supported hardware models and ONTAP versions . . . . .	42
5.2	Deduplication and Data ONTAP version considerations . . . . .	42
5.3	Deduplication licensing . . . . .	43
5.4	Compatibility with native N series functions . . . . .	43
5.4.1	Deduplication and high availability pairs . . . . .	43
5.4.2	Deduplication and non-disruptive upgrades . . . . .	44
5.4.3	Deduplication and Performance Acceleration Modules . . . . .	44
5.4.4	Deduplication and Snapshot copies . . . . .	44
5.4.5	Deduplication and SnapRestore . . . . .	45
5.4.6	Deduplication and SnapMirror replication . . . . .	45
5.4.7	Deduplication and SnapVault . . . . .	49
5.4.8	Deduplication and SnapLock . . . . .	50
5.4.9	Deduplication and MultiStore (vFiler) . . . . .	51
5.4.10	Deduplication and LUNs . . . . .	51
5.4.11	Deduplication and the volume copy command . . . . .	55
5.4.12	Deduplication and FlexClone volumes . . . . .	55
5.4.13	Deduplication and read reallocation . . . . .	56
5.4.14	Deduplication with Quotas . . . . .	56
5.4.15	Deduplication with NDMP . . . . .	56
5.4.16	Deduplication with DataFabric Manager/Protection Manager . . . . .	57
5.5	Compatibility with non-native functions . . . . .	57
5.5.1	Deduplication with VMWare . . . . .	57
5.5.2	Deduplication with Tivoli Storage Manager . . . . .	58
5.5.3	Deduplication with Backup Exec . . . . .	58
5.5.4	Deduplication with Lotus Domino . . . . .	58
5.5.5	Deduplication with Microsoft Exchange . . . . .	58
5.6	Data Characteristics . . . . .	58
5.7	Deduplication and storage capacity . . . . .	60
5.7.1	Metadata . . . . .	60
5.8	Deduplication and performance . . . . .	60
5.8.1	Duration of the deduplication operation . . . . .	61
5.8.2	The I/O performance of deduplicated volumes . . . . .	61
5.9	Deduplication scheduling . . . . .	62
5.10	Aggregate and volume considerations . . . . .	62
<b>Chapter 6.</b>	<b>ProtectTIER planning . . . . .</b>	<b>65</b>
6.1	Hardware planning for the 3959-SM1, 3958-AP1, 3958-DD4, and 3958-DE2 . . . . .	66
6.1.1	Specifications of the IBM System Storage TS7600 series . . . . .	66
6.1.2	Hardware and software components of the 3959-SM1 . . . . .	68
6.1.3	Hardware and software components of the 3958-AP1 . . . . .	69
6.1.4	Hardware and software components of the 3958-DD4 . . . . .	70
6.1.5	Hardware and software components of the 3958-DE2 . . . . .	72
6.2	Planning for deduplication . . . . .	73
6.3	Planning for Open systems with VTL . . . . .	74
6.3.1	Sizing inputs . . . . .	77
6.3.2	Capacity sizing . . . . .	82
6.3.3	Performance Sizing . . . . .	92
6.4	Planning for Open systems with OST . . . . .	100

6.5 Planning for installation . . . . .	101
6.5.1 Supported backup server operating environments . . . . .	101
6.5.2 ProtecTIER manager workstation requirements . . . . .	101
<b>Chapter 7. ITSM planning . . . . .</b>	<b>103</b>
7.1 ITSM Planning overview . . . . .	104
7.2 ITSM Deduplication pre-requisites . . . . .	104
7.2.1 ITSM active log and archive log sizing . . . . .	104
7.2.2 ITSM database sizing . . . . .	105
7.3 Memory . . . . .	106
7.4 Types of ITSM deduplication . . . . .	106
7.4.1 Server-side deduplication . . . . .	106
7.4.2 Client-side deduplication . . . . .	107
7.5 ITSM deduplication considerations . . . . .	108
7.5.1 Supported versions . . . . .	109
7.5.2 Eligible storage pools . . . . .	109
7.5.3 Encrypted data . . . . .	109
7.5.4 Compressed data . . . . .	110
7.5.5 Small files . . . . .	110
7.5.6 LAN-free considerations . . . . .	110
7.5.7 Hierarchical Storage Management (HSM) . . . . .	110
7.5.8 Collocation . . . . .	110
7.5.9 Disaster recovery considerations . . . . .	111
7.6 When to use ITSM deduplication . . . . .	112
7.7 When not to use ITSM deduplication . . . . .	112
<b>Part 3. Implementing Deduplication . . . . .</b>	<b>113</b>
<b>Chapter 8. Implementing N series deduplication . . . . .</b>	<b>115</b>
8.1 Requirements Overview . . . . .	115
8.2 The IBM N series System Manager . . . . .	116
8.2.1 Overview of System Manager . . . . .	116
8.2.2 Bug Fix . . . . .	116
8.2.3 Features . . . . .	117
8.2.4 System Requirements . . . . .	117
8.2.5 Installing the IBM N series System Manager Software . . . . .	118
8.2.6 Starting IBM N series System Manager . . . . .	123
8.3 End-to-end Deduplication configuration example using command line . . . . .	152
8.4 End-to-end deduplication configuration example using IBM N series System Manager Software . . . . .	155
8.5 Configuring deduplication schedules . . . . .	164
8.6 Sizing for performance and space efficiency . . . . .	166
8.6.1 Deduplication general best practices . . . . .	166
<b>Chapter 9. Implementing ProtecTIER . . . . .</b>	<b>169</b>
9.1 Getting started . . . . .	170
9.1.1 TS7610 SMB Appliance . . . . .	170
9.1.2 TS7650 Appliance . . . . .	170
9.1.3 TS7650G Gateway . . . . .	170
9.2 Installing ProtecTIER Manager . . . . .	171
9.2.1 Prerequisites . . . . .	171
9.2.2 Installing on Windows XP . . . . .	171
9.3 Starting ProtecTIER Manager . . . . .	177
9.3.1 Adding Nodes to ProtecTIER Manager . . . . .	177

9.3.2	Logging in into ProtecTIER Manager .....	179
9.4	ProtecTIER configuration on TS7610 .....	181
9.4.1	ProtecTIER Configuration Menu .....	181
9.4.2	ProtecTIER Manager Configuration Wizard .....	185
9.5	ProtecTIER software install .....	193
9.5.1	autorun .....	193
9.5.2	ptconfig .....	195
9.5.3	fsCreate .....	198
9.6	Repository creating .....	199
9.6.1	Repository planning .....	199
9.6.2	Repository creating .....	201
9.7	OST configuration .....	206
9.7.1	The OpenStorage Operating Environment .....	206
9.7.2	Configuring Storage Server (STS) .....	207
9.7.3	Modifying STS .....	207
9.7.4	Creating Logical Storage Units (LSUs) .....	208
9.7.5	Modifying LSUs .....	212
9.7.6	Deleting LSU .....	215
9.8	Virtual library creation .....	216
9.8.1	TS7610 Library .....	225
9.9	Host implementation .....	226
9.9.1	Connecting hosts to ProtecTIER systems .....	226
9.9.2	Installing and configuring the device driver .....	230
9.10	Deduplication operation with ProtecTIER .....	234
9.10.1	Deduplication with VTL .....	235
9.10.2	Compression with VTL .....	236
9.10.3	Deduplication with OST .....	236
9.10.4	Compression with OST .....	236
9.11	Backup and restore applications .....	237
9.11.1	Suggestions for all backup servers .....	237
9.11.2	General suggestions .....	237
9.11.3	Data types .....	240
9.11.4	IBM Tivoli Storage Manager .....	242
<b>Chapter 10.</b>	<b>Implementing ITSM deduplication .....</b>	<b>259</b>
10.1	Implementing server-side deduplication .....	260
10.2	Implementing client-side deduplication .....	264
10.3	Managing data deduplication .....	265
10.3.1	Starting duplicate-identification processes automatically .....	266
10.3.2	Starting duplicate-identification processes manually .....	266
10.3.3	Enabling data deduplication .....	267
10.3.4	Disabling data deduplication .....	267
10.3.5	Disabling the copy storage pool backup requirement .....	268
10.3.6	Restoring or retrieving deduplicated files .....	269
10.4	Deduplication best practices .....	270
10.4.1	Server resources .....	270
10.4.2	Data safety .....	270
10.4.3	Software version .....	271
10.4.4	Administrative schedules .....	271
10.4.5	Restore performance .....	271
<b>Chapter 11.</b>	<b>Suggestions .....</b>	<b>273</b>
11.1	Data flow considerations .....	274

11.2 Deduplication platform features . . . . .	274
11.3 Deduplication features matrix . . . . .	276
11.4 ProtecTIER deduplication . . . . .	276
11.5 ITSM deduplication . . . . .	278
11.6 IBM N-Series deduplication. . . . .	278
<b>Part 4. Appendixes . . . . .</b>	<b>281</b>
<b>Appendix A. N series use case . . . . .</b>	<b>283</b>
Lab environment . . . . .	284
Results . . . . .	284
Summary . . . . .	284
<b>Appendix B. ProtecTIER user cases . . . . .</b>	<b>287</b>
ProtecTIER deduplication . . . . .	288
<b>Appendix C. ITSM deduplication examples . . . . .</b>	<b>291</b>
Environment . . . . .	292
Results . . . . .	292
Summary . . . . .	293
<b>Related publications . . . . .</b>	<b>295</b>
IBM Redbooks . . . . .	295
Other publications . . . . .	295
Online resources . . . . .	295
How to get Redbooks. . . . .	295
Help from IBM . . . . .	296
<b>Index . . . . .</b>	<b>297</b>



# Notices

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

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

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

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

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

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

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

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

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

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


## COPYRIGHT LICENSE:

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

# Trademarks

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

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

AIX 5L™	Informix®	System Storage®
AIX®	Lotus Notes®	System x®
DB2®	Lotus®	System z®
Diligent®	Notes®	Tivoli®
Domino®	POWER5+™	VTF®
DS4000®	ProtecTIER®	XIV®
DS8000®	Redbooks®	z/OS®
HyperFactor®	Redbooks (logo)  ®	z/VM®
IBM®	System p5®	

The following terms are trademarks of other companies:

Snapshot, WAFL, SnapVault, SnapRestore, SnapMirror, SnapLock, NearStore, MultiStore, FlexVol, FlexClone, DataFabric, Data ONTAP, and the NetApp logo are trademarks or registered trademarks of NetApp, Inc. 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.

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

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

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



# Preface

Until now, the only way to capture, store, and effectively retain constantly growing amounts of enterprise data was to add more disk space to the storage infrastructure, an approach that can quickly become cost-prohibitive as information volumes continue to grow and capital budgets for infrastructure do not.

In this IBM® Redbooks® publication, we introduce data deduplication, which has emerged as a key technology in dramatically reducing the amount of, and therefore the cost associated with storing, large amounts of data. Deduplication is the art of intelligently reducing storage needs through the elimination of redundant data so that only one instance of a data set is actually stored. Deduplication reduces data an order of magnitude better than common data compression techniques. IBM has the broadest portfolio of deduplication solutions in the industry, giving us the freedom to solve customer issues with the most effective technology. Whether it is source or target, inline or post, hardware or software, disk or tape, IBM has a solution with the technology that best solves the problem. This IBM Redbooks publication covers the current deduplication solutions that IBM has to offer:

- ▶ IBM ProtecTIER® Gateway and Appliance
- ▶ IBM Tivoli® Storage Manager
- ▶ IBM System Storage® N series Deduplication

## The team who wrote this book

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

**Alex Osuna** is a Project Leader at the International Technical Support Organization, Tucson Center. He writes extensively on IBM storage. Before joining the ITSO five years ago, Alex worked in the Tivoli Western Region supporting storage software products as a Principle Systems Engineer. Alex has over 33 years in the IT industry with the United States Air Force and assignments with IBM in Phoenix, Tucson, Gaithersburg, and San Jose.

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

**Alexandre Ramos Galante de Carvalho** is an IT Specialist for IBM Global Services in Brazil, supporting international accounts. He holds a Bachelor degree in Computer Science from PUC/SP, São Paulo, Brazil. He has 11 years of experience working with all Unix flavors, clustering solutions, IBM Tivoli Storage Manager, and IBM Tivoli Data Protection. He has nine certifications in total, including AIX®, HP-UX, and TSM.

**Rucel F. Javier** is an IT Specialist from IBM Philippines, Inc. He is a Field Technical Support Specialist for System Technology Group (STG), providing pre-sales and post-sales support. He has worked in IT Industry for 9 years and is an IBM Certified Specialist in System x®, Blade Servers/Centers, and Mid Range Systems Storage. He has worked with the implementation and support of N series, DS3000, and DS5000 storage systems, and has an extensive experience in Linux® operating systems. Rucel is also doing proof of concept (POC), product presentations, and solutions for customer requirements.

**Zohar Mann** is an IBM Storage Architect based out of San Francisco, California. He has 14 years of enterprise level IT experience with UNIX® systems administration, enterprise management systems administration, and mass Storage/SAN/NAS administration and architecture.



*Figure 1 The team: Alex, Alexandre, Eva, Zohar and Rucel*

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

Dan Edwards  
IBM Global Technology Services

Shayne Gardener  
IBM Software Group, Tivoli

Mikael Lindstrom  
GTS Services Delivery

Craig McAllister  
IBM Software Group, Tivoli

Norbert Pott  
Tivoli Storage Manager V6.1 Technical Guide authors

Dave Accordino  
IBM Systems and Technology Group, Systems Hardware Development

Denise Brown  
IBM Systems and Technology Group, Systems Hardware Development ProtecTIER Test

Joe Dain  
IBM Systems and Technology Group, Systems Hardware Development TS7600 Development

Mary Lovelace  
IBM International Technical Support Organization

Gerard Kimbuende  
IBM Systems and Technology Group, Systems Hardware Development

Carlin Smith  
IBM Advanced Technical Skills (ATS) - Americas

## Now you can become a published author, too!

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

Find out more about the residency program, browse the residency index, and apply online at:  
[ibm.com/redbooks/residencies.html](http://ibm.com/redbooks/residencies.html)

## Comments welcome

Your comments are important to us!

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

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

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

- Send your comments in an email to:

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

- Mail your comments to:

IBM Corporation, International Technical Support Organization  
Dept. HYTD Mail Station P099

2455 South Road  
Poughkeepsie, NY 12601-5400

## Stay connected to IBM Redbooks

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



# Part 1

# Introduction to Deduplication

In this part we introduce deduplication from industry and IBM points of view.





# Introduction to Deduplication

Business data growth rates will continue to increase rapidly in the coming years. Likewise, retention and retrieval requirements for new and existing data will expand, driving still more data to disk storage. As the amount of disk-based data continues to grow, there is ever-increasing focus on improving data storage efficiencies across the information infrastructure.

Data reduction is a tactic which can decrease the disk storage and network bandwidth required, lower Total Cost of Ownership (TCO) for storage infrastructures, and optimize use of existing storage assets and improve data recovery infrastructure efficiency. Deduplication and other forms of data reduction are features that can exist within multiple components of the information infrastructure. IBM offers a comprehensive set of deduplication solutions.

IBM has been the industry leader in storage efficiency solutions for decades. IBM invented Hierarchical Storage Management (HSM) and the progressive incremental backup model, greatly reducing the primary and backup storage needs of its customers. Today, IBM continues to provide its customers the most efficient data management and data protection solutions available. N-Series, ProtecTIER and IBM Tivoli Storage Manager are excellent examples of IBM's continued leadership.

## 1.1 IBM data reduction and deduplication

Data deduplication is one of the techniques for achieving data reduction. As with any technology, there are benefits and costs associated with different deduplication deployment options. IBM offers data deduplication capabilities in multiple parts of its storage hardware and software portfolio to give you more flexible deployment options:

- **As a Virtual Tape Library**

IBM ProtecTIER's deduplication technology allows excellent scalability, performance and data integrity characteristics. ProtecTIER is offered as a gateway or disk-based appliance. It is accessed today as a Virtual Tape Library (VTL). ProtecTIER offers global deduplication across a wide domain of IBM and non-IBM backup servers, applications, and disks. IBM Tivoli Storage Manager works effectively with ProtecTIER and can exploit ProtecTIER's efficient network replication capability.

- **In the data protection application**

Another option is native storage pool deduplication with IBM Tivoli Storage Manager Version 6, which offers reduction of backup and archive data. Native deduplication helps customers store more backup data on the same disk capacity, thereby enabling additional recovery points without incurring additional hardware costs. IBM Tivoli Storage Manager deduplication is especially applicable in smaller environments where scalability is not a primary requirement or where an additional deduplication appliance is not economically feasible. IBM Tivoli Storage Manager deduplication can be used in larger environments if appropriate CPU, memory, and I/O resources are available on the server.

- **As a primary or secondary storage system**

IBM N series storage systems offer native, block-level, file-level, application-level, OS-level, protocol-agnostic postprocess data deduplication. This no cost feature consolidates duplicate 4 KB blocks on the disk storage system with no additional hardware.

As a data reduction appliance, IBM Real-time Compression for Network Attached Storage (NAS) is an ideal solution for large active databases and virtual server files. Because accessing information requires decompressing only the data being used, not the entire file, active data can be optimized in real time.

## 1.2 Deduplication overview

Data deduplication is a method of consolidating redundant copies of a file or file subcomponent. Incoming or existing data are standardized into "chunks" that are then examined for redundancy. If duplicates are detected, then pointers are shifted to reference a single copy of the chunk and the extraneous duplicates are then released (see Figure 1-1 on page 5).



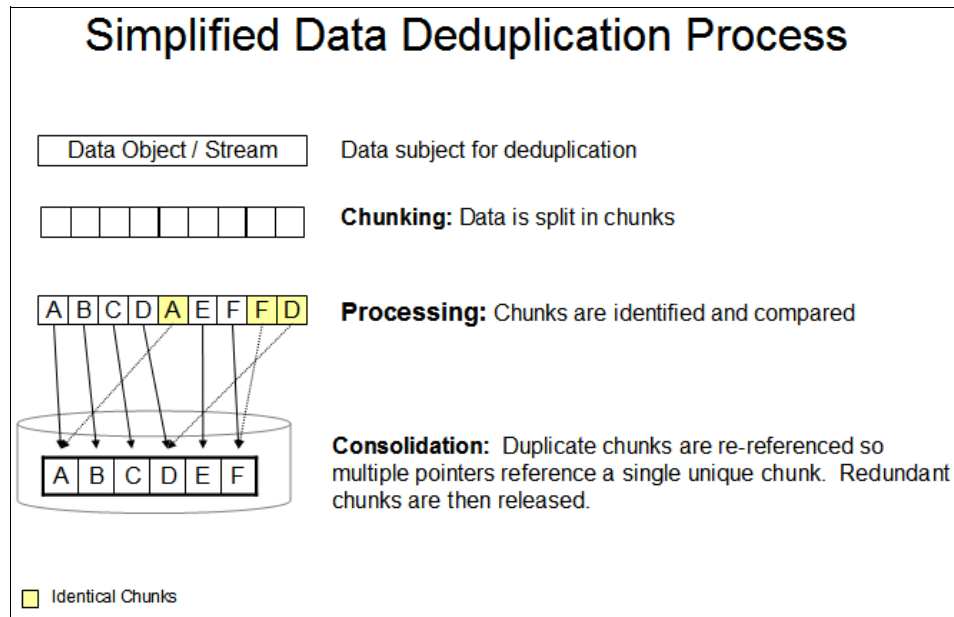


Figure 1-1 Deduplication processing

## 1.2.1 Chunking

Chunking refers to breaking data into standardized units that can be examined for duplicates. Depending on the technology and locality of the deduplication process, these units can be files or more granular components such as blocks.

File-level deduplication is generally less versatile than other means of deduplication, but if the deduplication platform is format-aware, it can potentially identify explicit components within certain files such as embedded objects.

Block-level deduplication is generally data-agnostic in the sense that it will compare blocks regardless of the file type, application, and the OS that the data originated from. Some handling of data such as compartmentalization or encryption can cause identical files to have mismatched blocks, which can reduce the efficacy of block-level deduplication. Block-based deduplication is more granular, but requires more processing power and a larger index or catalog to track the individual pieces.

There are the following methods of data chunking. Each method influences the data deduplication ratio.

<b>File based</b>	Each chunk is a single file. File-based deduplication is typically used with devices that only have file system visibility.
<b>Block based</b>	The data object is chunked into blocks of fixed or variable size. Block-based deduplication is typically used by block storage devices.
<b>Format aware</b>	This method understands explicit data formats and chunk data object according to the format. For example, format aware chunking could break a PowerPoint presentation into separate slides.
<b>Format agnostic</b>	Chunking is based on an algorithm that looks for logical breaks or similar elements within a data object.

## 1.2.2 Processing

Each chunk of data must be identified in a way that is easily comparable. Chunks are processed using a parity calculation or cryptographic hash function. This processing gives the chunks shorter identifiers known as a hash values, digital signatures, or fingerprints. These fingerprints can be stored in an index or catalog where they can be compared quickly with other fingerprints to find matching chunks.

In rare cases where hash values are identical but reference non-unique chunks, a hash collision occurs that can lead to data corruption. To avoid this scenario, a secondary comparison should be done to verify that hash-based duplicate chunks are in fact redundant before moving on to consolidation.

Certain deduplication technologies will do a byte comparison of chunks after a fingerprint match to avoid hash collisions. Processing of chunks is generally the most CPU intensive part of deduplication, and can impact I/O if done in-line.

There are three methods to differentiate a chunk. Each method influences the duplicate identification performance:

<b>Hashing</b>	This method computes a hash (MD-5, SHA-1, SHA-2) for each data chunk and compares that hash with the hashes of existing data. An identical hash means the data is most likely identical. In the case of <b>hash collisions</b> where there is an identical hash but non-identical data, data corruption is prevented through a secondary comparison using additional metadata, a second hash method, or binary comparison.
<b>Binary comparison</b>	This method compares all bytes of similar chunks.
<b>Delta differencing</b>	This method computes a <i>delta</i> between two similar chunks where one chunk is the baseline and the second chunk is the delta. Because each delta is unique, there is no possibility of collision. To reconstruct the original chunk, the deltas have to be re-applied to the baseline chunk.

## 1.2.3 Consolidation

After duplicate chunks have been compared and identified, the pointers to those chunks must be changed so they point to a single unique copy rather than multiple duplicate chunks. After the pointers are consolidated, the now extraneous data chunks can be released. In cases of inline processing, the duplicate chunks are not written to the physical disk storage.

## 1.3 Architecture

This section offers a high level description of the deduplication architecture.

### 1.3.1 Where deduplication processing occurs

Data deduplication processing can occur on the client, on an infrastructure server, or on the storage system (see Figure 1-2 on page 7). Each option has factors to consider.

### ***Client based deduplication processing***

Client based deduplication processing can reduce the amount of data being transferred over the network to the storage system, but there are often requirements for extra CPU and disk I/O processing on the client side.

Clients might be constrained to file level visibility of the data, which can reduce the granularity of the deduplicated component analysis. Also, clients might only have visibility to a limited pool of data, which can impact the duplicate frequency of the pool and reduce deduplication efficiency.

### ***Server-based deduplication processing***

Server-based deduplication processing allows you to deduplicate multiple client's data at a scheduled time, but requires extra CPU and disk I/O processing on the server. In the cases of deduplication described in this document, the server refers to dedicated ITSM or ProtecTIER infrastructure servers. Data I/O must flow through these systems for deduplication to occur. Therefore data already resident on disk storage systems will have to be routed through these systems for deduplication.

### ***Storage system deduplication processing***

With storage system based deduplication processing, the deduplication occurs at the disk storage device where the data is stored. It is generally transparent to the clients and the servers. It uses CPU and disk I/O on the storage system.

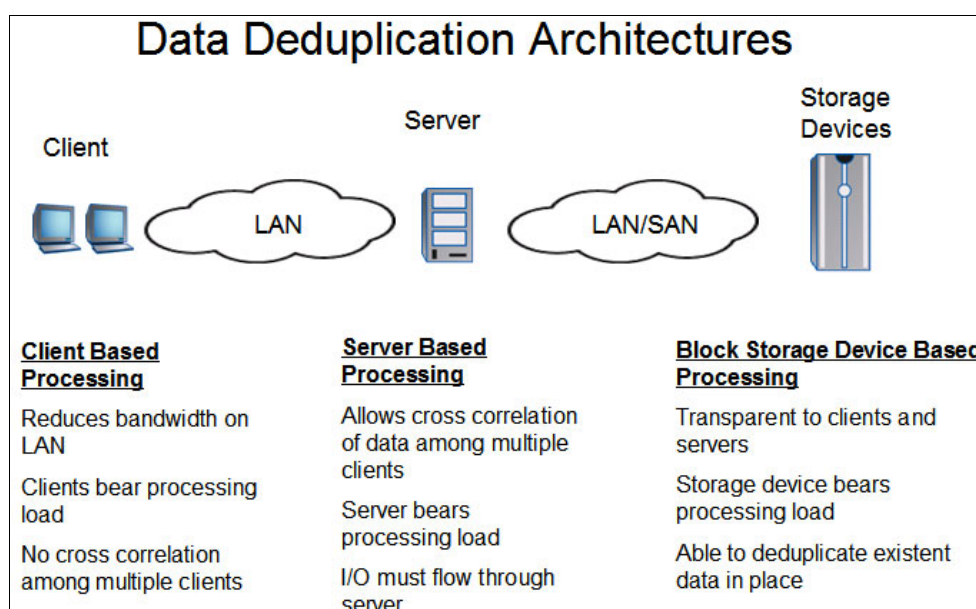


Figure 1-2 Data deduplication architecture

IBM Tivoli Storage Manager can do client based (ITSM client) and server based (ITSM server) deduplication processing. ProtecTIER and N series do storage system based deduplication processing.

## **1.3.2 When deduplication occurs**

There are two approaches for when deduplication occurs in the data flow.

### **Inline Deduplication**

Inline deduplication processes and consolidates data before it is written to disk. Hashing and the hash comparison must be done on the fly, which can add performance overhead. If a byte-for-byte comparison must be done to avoid hash collisions, even more overhead is needed. In cases of deduplicating primary data when milliseconds are significant, inline deduplication is generally not recommended. A benefit of inline deduplication is that duplicate chunks are never written to the destination disk system.

### **Postprocess Deduplication**

Postprocess deduplication refers to handling processing and consolidation after the data has been written to disk, generally with scheduled or manual runs. This allows more control over the performance impact of deduplication to avoid peaks or conflicts with other data I/O processes such as backups or DR replication. The duplicate data will consume disk space until deduplication runs. Buffer capacity must be more rigorously maintained to accommodate the more dynamic utilization. This method is the only option for processing data already in place on the disk storage systems prior to bringing deduplication technology into the environment.

## **1.4 Benefits of data deduplication**

Data deduplication can reduce storage footprint and Total Cost of Ownership (TCO). It will help you save the following:

- ▶ Storage capacity by allowing you to store more data per physical storage system
- ▶ Energy by using less data/disk, thereby requiring less energy
- ▶ The amount of data that must be sent across a network to primary storage, for backup replication, and for disaster recovery

In case you are using disks for backups, the more efficient use of disk space also allows for longer disk retention periods, which provides better recovery time objectives (RTO) for a longer time.



## Introduction to N series Deduplication

IBM N series Storage System offers deduplication as a native part of its Data ONTAP® operating system. This is a feature that can be added for no additional cost and leverages several benefits already built into the storage system's OS. Because it is post process deduplication within the storage system, it is possible to realize immediate benefits with existing storage capacity.

The IBM N series deduplication follows a risk-averse postprocess methodology and provides owners the benefits of both source and destination deduplication. For environments that rely on multi-tier storage architecture with N series, the deduplication gains will be achieved not only on primary storage, but on backup and archival disk tiers as well.

The N series system uses sub-file, fixed block granularity, which means it is generally file and OS agnostic. Data chunks being compared are at the 4K physical block level. Hashing of incoming data is done by default within the Data ONTAP OS so deduplication is able to take advantage of these existing block fingerprints by simply copying them to a catalog for later comparison. These fingerprints are initially used to detect duplicate blocks after which a byte level comparison between the blocks is done to validate duplication. This two-phase comparison eliminates risk of hash collisions and data corruption.

Because the IBM N series OS is already capable of multiple block referencing, it is simple for the deduplication process to change inode pointers to reference a single physical block rather than multiple, duplicate blocks. After the duplicate blocks are no longer referenced, they are considered to be available again and space savings have been achieved. A maximum of 255 such references can be made to a single physical block. Thus the maximum deduplication ratio that can be achieved is 255:1.

The deduplication license comes at no cost with the IBM N series Filer and Gateway systems. The Gateway systems can be added to existing non-N series disk storage systems to allow these systems to also take advantage of the benefits and features of Data ONTAP.

## 2.1 How deduplication for IBM System Storage N series storage system works

Regardless of operating system, application, or file system type, *all* data blocks are written to a storage system using a data reference pointer, without which the data could not be referenced or retrieved. In traditional (non-deduplicated) file systems, data blocks are stored without regard to any similarity with other blocks in the same file system. In Figure 2-1, five identical data blocks are stored in a file system, each with a separate data pointer. Although all five data blocks are identical, each is stored as a separate instance and each consumes physical disk space.




Figure 2-1 Non-deduplicated data

In a deduplicated N series file system, two new and important concepts are introduced:

- ▶ A catalog of all data block is maintained. This catalog contains a record of all data blocks using a "hash" or fingerprint that identifies the unique contents of each block.
- ▶ The file system is capable of allowing many data pointers to reference the same physical data block.

Cataloging data objects, comparing the objects, and redirecting reference pointers forms the basis of the deduplication algorithm. As shown in Figure 2-2, referencing several identical blocks with a single master block allows the space that is normally occupied by the duplicate blocks to be given back to the storage system.

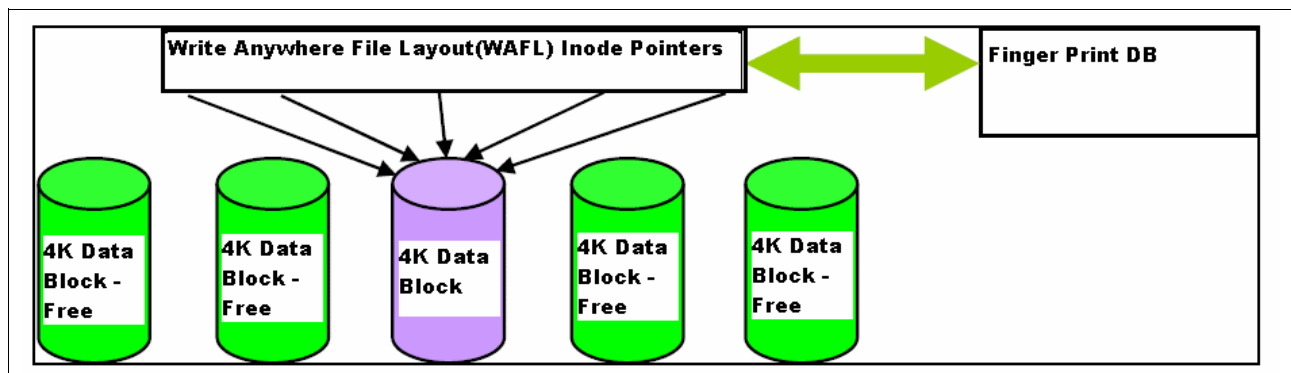



Figure 2-2 Deduplicating identical blocks creates available space

## Hashing

Data deduplication begins with a comparison of two data blocks. It would be impractical (and arduous) to scan an entire data volume for duplicate blocks each time a new data is written to that volume. For that reason, deduplication creates small hash values for each new block, and store these values in a catalog.

A hash value, also called a digital fingerprint or digital signature as shown in Figure 2-3, is a small number that is generated from a longer string of data. A hash value is substantially smaller than the data block itself, and is generated by a mathematical formula in such a way that it is unlikely (although not impossible) for two non-identical data blocks to produce the same hash value.



*Figure 2-3 A hash value is a digital fingerprint that represents a much larger object*

The deduplication process goes through the following steps:

- ▶ The fingerprint catalog is sorted and searched for identical fingerprints.
- ▶ When a fingerprint “match” is made, the associated data blocks are retrieved and scanned byte-for-byte.
- ▶ Assuming successful validation, the inode pointer metadata of the duplicate block is redirected to the original block.
- ▶ The duplicate block is marked as “Free” and returned to the system, eligible for re-use.

## Hash catalog

A catalog of hash values is used to identify candidates for deduplication. A system process identifies duplicates, and data pointers are modified accordingly. The advantage of catalog deduplication is that the catalog is used only to identify duplicate objects; it is not accessed during the actual reading or writing of the data objects. That task is still handled by the normal file system data structure as shown in Figure 2-4 on page 12.

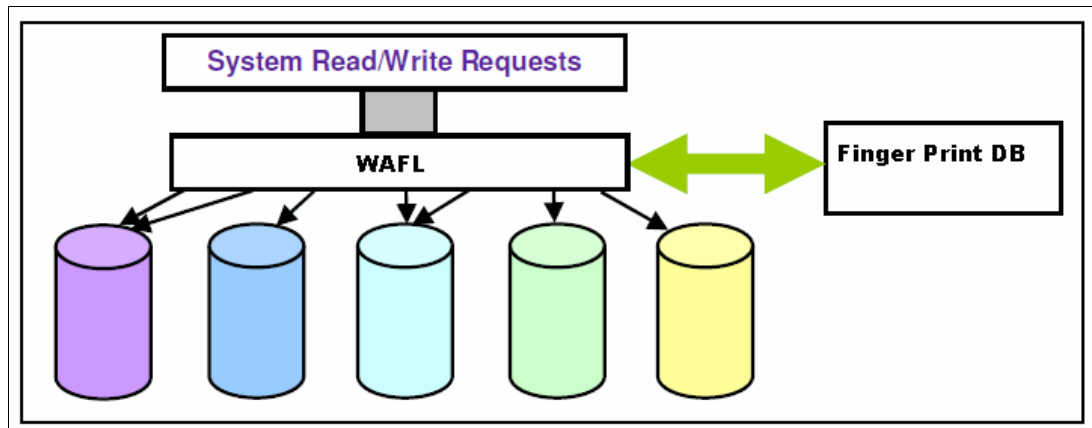


Figure 2-4 Catalog indexing: the file system controls block sharing of deduplicate blocks

Deduplication is an IBM System Storage N series storage efficiency offering that provides block-level deduplication within the entire flexible volume on IBM System Storage N series storage systems. Beginning with Data ONTAP 7.3, IBM System Storage N series gateways also support deduplication. IBM System Storage N series gateways are designed to be used as a gateway system that sits in front of third-party storage, allowing IBM System Storage N series storage efficiency and other features to be used on third-party storage. Figure 2-5 shows how IBM System Storage N series deduplication works at the highest level.

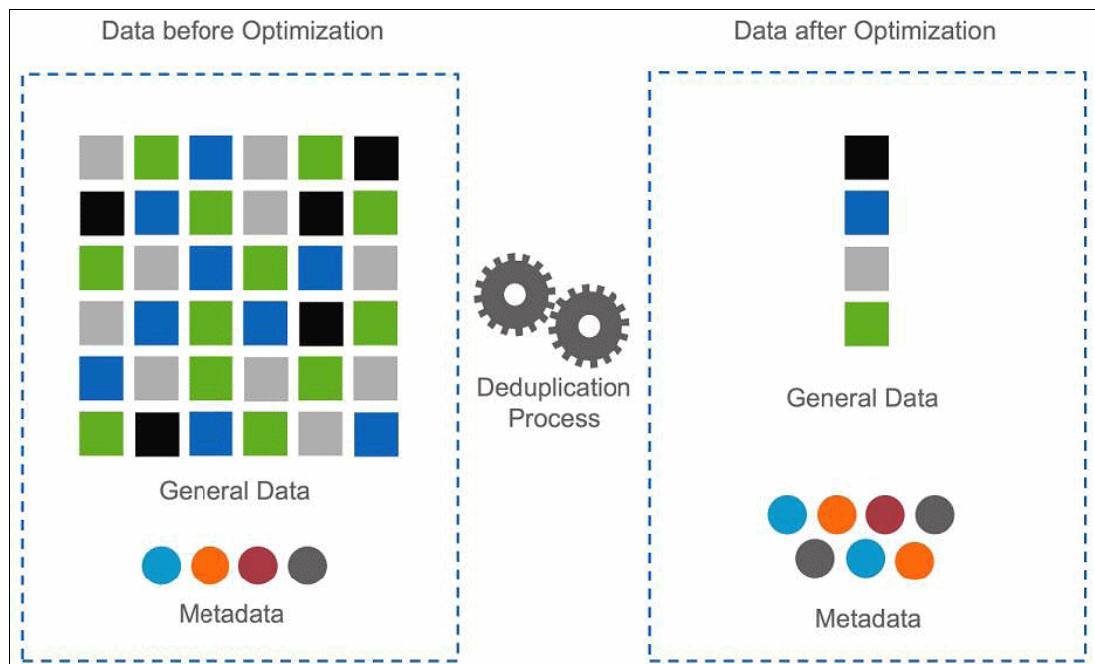


Figure 2-5 Deduplication Operations

Essentially, deduplication stores only unique blocks in the flexible volume and creates a small amount of additional metadata during the process. Deduplication has the following notable features:

- ▶ It works with a high degree of granularity (that is, at the 4 KB block level).
- ▶ It operates on the active file system of the flexible volume. Any block referenced by a Snapshot™ is not made available until the Snapshot is deleted.



- ▶ It is a background process that can be configured to run automatically, can be scheduled, or can run manually through the command-line interface (CLI).
- ▶ It is application-transparent, and therefore can be used for data originating from any application using the IBM System Storage N series storage system.
- ▶ It is enabled and managed through a simple CLI.
- ▶ It can be enabled on (and can deduplicate blocks on) flexible volumes with both new and existing data.

In summary, deduplication works as follows:

1. Newly saved data on the IBM System Storage N series storage system is stored in 4 KB blocks as usual by Data ONTAP.
2. Each block of data has a digital fingerprint that is compared to all other fingerprints in the flexible volume.
3. If two fingerprints are found to be the same, a byte-for-byte comparison is done of all bytes in the block. If an exact match is found between the new block and the existing block on the flexible volume, the duplicate block is discarded and its disk space is reclaimed.

## 2.2 Deduplicated Volumes

Despite the introduction of less expensive ATA disk drives, one of the biggest challenges for disk-based backup today continues to be storage cost. We want to reduce storage consumption (and therefore storage cost per megabyte) by eliminating duplicated data through sharing across files.

The core IBM System Storage N series technology used to accomplish this goal is the *dense volume*, a flexible volume that contains shared data blocks. The Data ONTAP file system, Write Anywhere File Layout (WAFL®) is a file system structure that supports shared blocks in order to optimize storage space consumption. Basically, within one file system tree, you can have multiple references to the same data block, as shown in Figure 2-6

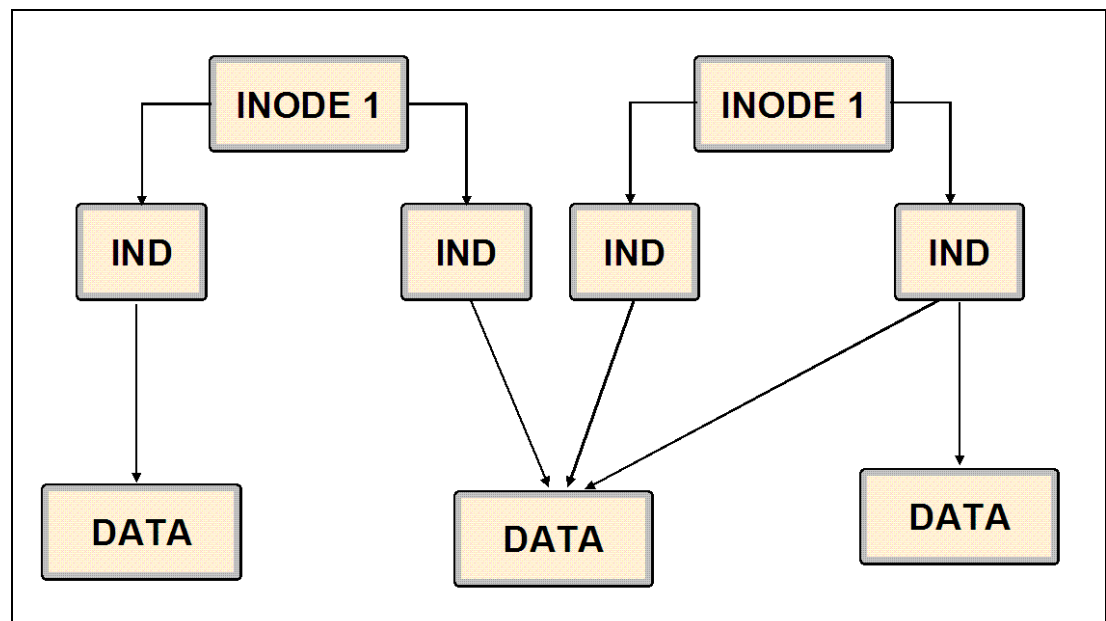


Figure 2-6 Data structure in a deduplicated volume

In Figure 2-6 on page 13, the number of physical blocks used on the disk is three (instead of five), and the number of blocks saved by deduplication is two (five minus three). In the remainder of this chapter, these blocks are referred to as *used blocks* and *saved blocks*. This is called multiple block referencing.

Each data block has a block-count reference kept in the volume metadata. As additional indirect blocks (shown as IND in Figure 2-6 on page 13) point to the data, or existing blocks stop pointing to the data, this value is incremented or decremented accordingly. When no indirect blocks point to a data block, it is released. The Data ONTAP deduplication technology therefore allows duplicate 4 KB blocks anywhere in the flexible volume to be deleted.

The maximum sharing for a block is 255. For example, that if there are 500 duplicate blocks, deduplication would reduce those blocks to only two blocks. Also note that this ability to share blocks is separate from the ability to keep 255 Snapshot copies for a volume.

## 2.3 Deduplication metadata

The core enabling technology of deduplication is *fingerprints*. These are unique digital signatures for every 4 KB data block in the flexible volume.

When deduplication runs for the first time on a flexible volume with existing data, it scans the blocks in the flexible volume and creates a fingerprint database that contains a sorted list of all fingerprints for used blocks in the flexible volume.

Although deduplication can provide substantial space savings, a percentage of storage overhead is associated with it that you should consider when sizing a FlexVol®.

After the fingerprint file is created, fingerprints are checked for duplicates and, when found, a byte-by-byte comparison of the blocks is done to make sure that the blocks are indeed identical. If they are found to be identical, the block's pointer is updated to the already existing data block, and the new (duplicate) data block is released.

Releasing a duplicate data block entails updating the indirect inode pointing to it, incrementing the block reference count for the already existing data block, and freeing the duplicate data block.

In real time, as additional data is written to the deduplicated volume, a fingerprint is created for each new block and written to a change log file. When deduplication is run subsequently, the change log is sorted and its sorted fingerprints are merged with those in the fingerprint file, and then the deduplication processing occurs.

### Notes:

- ▶ There are two change log files. As deduplication runs and merges the new blocks from one change log file into the fingerprint file, new data that is being written to the flexible volume causes fingerprints for these new blocks to be written to the second change log file. The roles of the two files are then reversed the next time that deduplication is run. If you are familiar with the Data ONTAP usage of NVRAM, this process is analogous to when it switches from one half to the other to take a consistency point.
- ▶ When deduplication is run for the first time on an empty flexible volume, it creates the fingerprint file from the change log.

### **Additional details about deduplication metadata**

A fingerprint record exists for every 4 KB data block because the fingerprints for all the data blocks in the volume are stored in the fingerprint database file.

Fingerprints are not deleted from the fingerprint file automatically when data blocks are freed; however, when a threshold of 20% new fingerprints is reached, the stale fingerprints are deleted. Deletion can also be done manually.

## **2.4 Sizing for performance and space efficiency**

This section discusses the deduplication best practices and behavior that you can expect. Information in this section comes from testing, observations, and knowledge of how deduplication functions.

### **2.4.1 Deduplication general best practices**

The following list contains deduplication best practices and lessons learned from internal tests and from deployments in the field:

- ▶ Deduplication consumes system resources and can alter the data layout on a disk. Because of the application's I/O pattern, and the effect of deduplication on the data layout, the read and write I/O performance can vary considerably. The space savings and the performance impact varies significantly depending on the application and the data contents.
- ▶ The performance impact due to deduplication should be carefully considered and measured in a test setup, and that you consider sizing before deploying deduplication in performance-sensitive solutions. The following are factors that affect the performance of deduplication.
  - Application and the type of data used
  - The data access pattern such as sequential versus random access, and the size of the pattern of the input and output
  - The amount of deduplicate data: the amount of total and average file size
  - The nature of the data layout in the volume
  - The amount of changed data between deduplication operations
  - The number of concurrent deduplication operations
  - Hardware platform including System Memory and CPU module
  - Load on the system
  - Disk types such as ATA/FC /SAS and RPM of the disk
- ▶ If the amount of new data is small, run deduplication infrequently, because there is no benefit to running it frequently in such a case, and it consumes CPU resources. How often you run it depends on the rate of change of the data in the flexible volume.

The more concurrent deduplication processes you are running, the more system resources that are consumed.

Given this information, the best option is to perform one of the following actions:

- Use the auto mode so that deduplication runs only when significant additional data has been written to each particular flexible volume (this approach tends to naturally spread out when deduplication runs).
- Stagger the deduplication schedule for the flexible volumes so that it runs on alternative days.
- Run deduplication manually.

- ▶ If Snapshot copies are required, run deduplication before creating the Snapshot to minimize the amount of data before the data gets locked in to the copies. Make sure that deduplication has completed before creating the copy. Creating a Snapshot on a flexible volume before deduplication has a chance to run and complete on that flexible volume can result in lower space savings.
- ▶ If Snapshot copies are to be used, the Snapshot reserve should be greater than zero (0). An exception to this might be in an FCP or iSCSI LUN scenario, where it is often set to zero for thin-provisioning reasons.
- ▶ For deduplication to run properly, you have to leave some free space for the deduplication metadata.

Deduplication is tightly integrated with Data ONTAP and the WAFL file structure. Because of this integration, deduplication is performed with extreme efficiency. Complex hashing algorithms and lookup tables are not required. Instead, Deduplication is able to use IBM System Storage N series storage systems with the NearStore® option internal characteristics to create and compare digital fingerprints, redirect data pointers, and free-up redundant data areas, as shown in Figure 2-7.

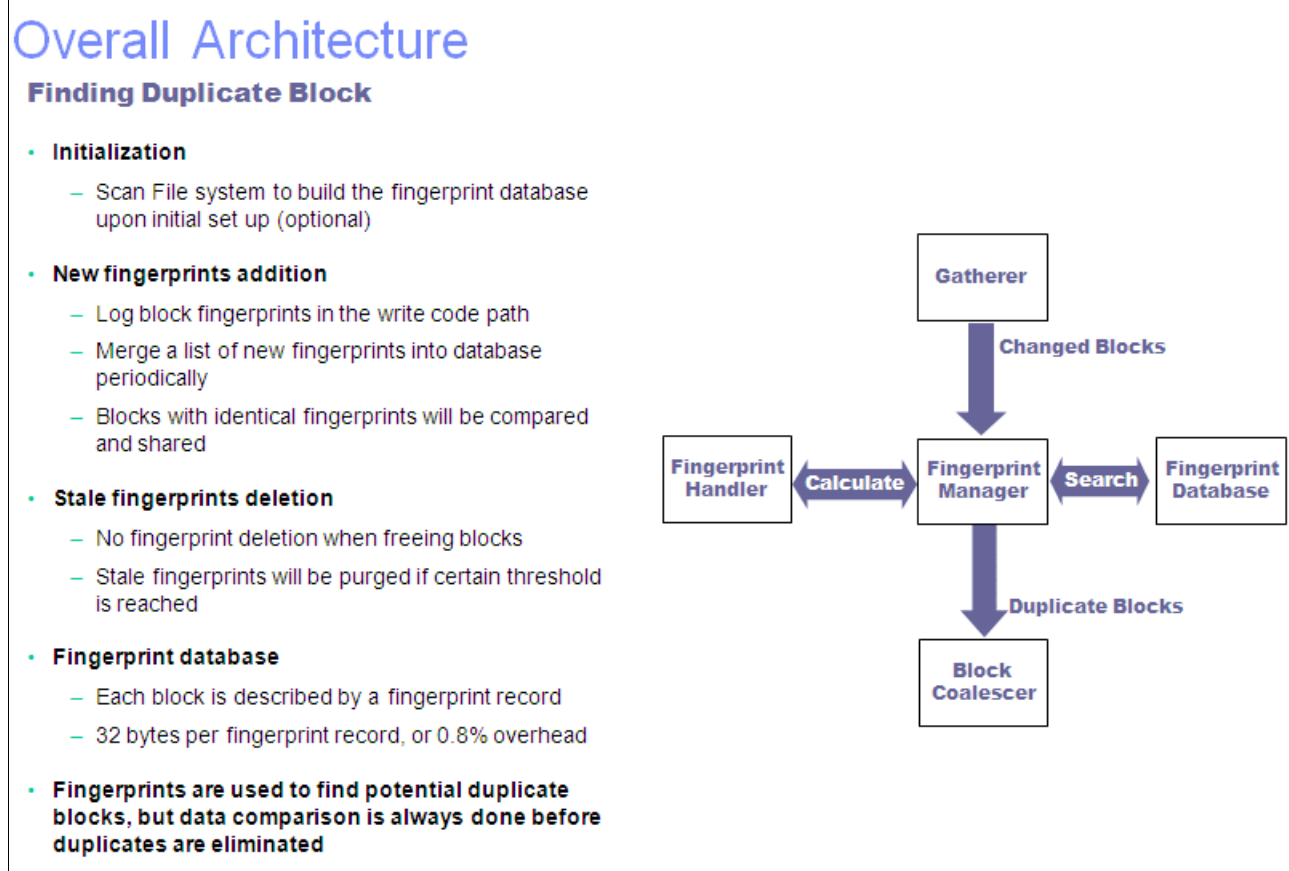


Figure 2-7 Digital fingerprinting

## 2.5 Compressing and Deduplicating

It has been shown that combined with other IBM products like ProtecTIER for compressing data, deduplicating offers the greatest space savings. IBM also has a complement to N series deduplication with IBM Real-time Compression Appliances.

Real-Time Compression is designed to sit transparently in front of your primary storage and reduce the size of every file you create up to 15x depending upon file type. Applications have random, read-write access to compressed data, and the physical capacity required to store a file, or copies and permutations of a file are significantly reduced throughout the entire life cycle including backup. Because less data is written to disk, overall network and storage performance and utilization can also be significantly enhanced.





# Introduction to ProtecTIER deduplication

In this chapter, we describe the ProtecTIER data deduplication concept called HyperFactor® data deduplication technology.

This chapter contains the following topics:

- ▶ Terms and definitions
- ▶ Overview of HyperFactor and deduplication
- ▶ IBM ProtecTIER deduplication solutions
- ▶ How HyperFactor and deduplication works

## 3.1 Terms and definitions

The following terms are used in this and later chapters:

### ***Front end***

The connection between the ProtecTIER system and the backup server is referred to as a front end connection.

### ***Back end***

The connection between the ProtecTIER system and the disk array is referred to as a back end connection.

### ***Node***

A IBM System x3850 X5 is viewed as a node from the ProtecTIER Manager software. You can have a single node (one IBM System x3850 X5) or two node clusters (two IBM System x3850 X5).

### ***Metadata***

Metadata is the data used to keep track of the data about your backup data, including where it is stored on the disk.

### ***User data***

User data is the backup files and data sets stored on the virtual tape library. It is the data that you are storing on disk.

### ***Metadata file system***

The metadata file system stores all aspects of the data that is backed up and cataloged, but not the data itself, whether it requires new disk space or not. It is critical that the performance of the metadata file system be optimal. Therefore RAID 10 groups (4+4, 6+6, or 8+8 disks) on Fibre Channel disk are optimal for the metadata file systems.

### ***User data file system***

The user data file system stores the actual data that is backed up or referenced by new generations of the data. The user data file system is stored on a RAID 5 configuration.

### ***Nominal capacity***

The amount of user data that ProtecTIER is managing.

### ***Physical capacity***

The physical disk space available within the array.

### ***Factoring ratio***

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

### ***Repository***

The repository is the physical disk that holds the ProtecTIER factored data. There are two types of file systems that make up the ProtecTIER Repository: user data and metadata.



***Disk array***

The disk array attaches to the IBM System Storage TS7600 with ProtecTIER through back end connections, and holds the repository or cache of factored backup data.

***Data retention period***

The period of time (usually measured in days) that defines how long customers will keep their disk-based backups online. This period of time typically ranges from a period of 30 to 90 days, but can be longer.

***Change rate***

The percentage of change in data from one backup set to the next. The daily change rate is the percentage of change from one day's backup cycle to the next. For example, if the daily change rate is 10%, it means that only 10% of the backed-up data changes from one day to the next.

***Disaster recovery***

Disaster recovery (DR) is the process of recovering production site data from a remote location. It includes a way to indicate to a remote repository that the production site has gone down.

***Disaster recovery test***

A simulation of the disaster recovery process.

***Failover***

The process of failover occurs when continued operations at the source location is no longer possible. A disaster is declared.

***Failback***

A process that is initiated from the remote site when the source site is now able to continue production operations and therefore back up processes. The process ensures that the paired repositories are re-synchronized using the least amount of bandwidth and maintaining the most recent copies of backups. After the failback process, the operational norms prior to the execution of a DR resume.

***ProtecTIER Replication Manager***

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

***Principality/ownership***

An attribute indicating the repository in which an individual cartridge can be updated or written on by a backup application. A cartridge at its principal repository can be read/write (R/W) or read only (R/O). At other sites it is R/O. A cartridge can have principality/ownership turned on for only one site.

***Replication***

A process that transfers logical objects like cartridges from one ProtecTIER repository to another. The replication function allows ProtecTIER deployment to be distributed across sites. Each site has a single or clustered ProtecTIER environment. The ProtecTIER server that is a

part of the replication grid has two dedicated replication ports. Eth3 and Eth4 are use for replication. Replication ports are connected to the customer's WAN and are configured on two subnets as the default.

### ***Shelf***

A container of virtual tape library (VTL) cartridges within a ProtecTIER repository. This is analogous to a shelf or a rack where physical tapes are kept when outside an automated tape library.

### ***TS7650***

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

## **3.2 Overview of HyperFactor technology and deduplication**

The cornerstone of ProtecTIER is HyperFactor, an IBM technology that deduplicates data inline as it is received from the backup application in case of Open Systems or from the z/OS® tape processing in case of Mainframe Systems. ProtecTIERs bandwidth-efficient replication, inline performance, and scalability directly stem from the technological breakthroughs inherent in HyperFactor. HyperFactor is based on a series of algorithms that identify and filter out the elements of a data stream that have previously been stored by ProtecTIER. Over time, HyperFactor can increase the usable capacity of a given amount of physical storage by up to 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. Replication is not described in this Redbooks publication.

As Figure 3-1 on page 23 shows, when new data is received by ProtecTIER deduplication technology, HyperFactor finds any similar data elements that have already been stored. This search is extremely quick using a small and efficient memory-resident index. After similar data elements are found, HyperFactor can compare the new data to the similar data to identify and store only the byte-level changes.

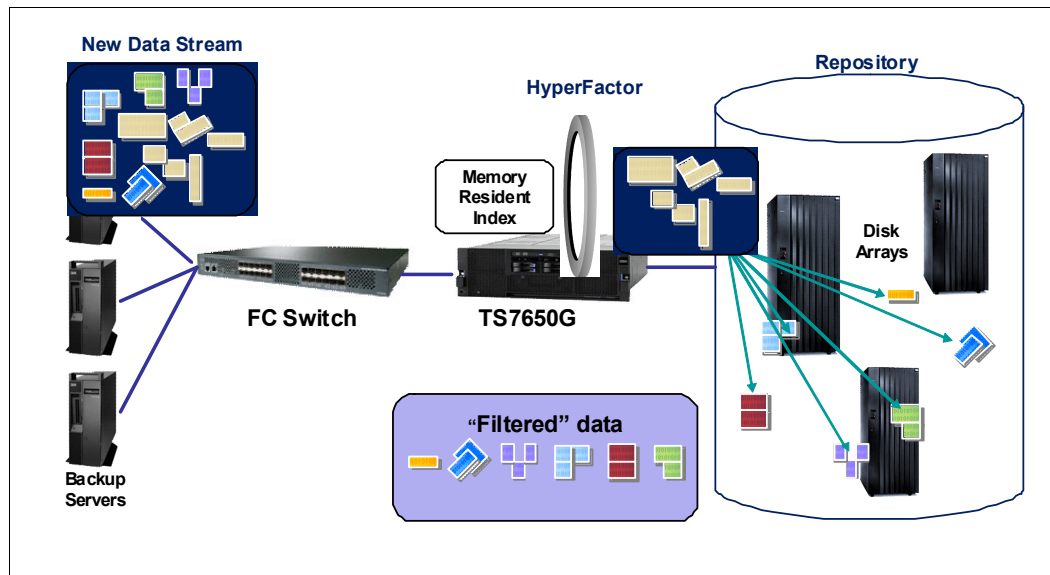


Figure 3-1 Hyperfactor in progress

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

In Figure 3-2 with HyperFactor deduplication, when new data is received, HyperFactor looks for data similarities and check those similarities in the Memory Resident Index. When similarity matches are found, the existing similar element is read from disk and a binary differential is performed on the similar elements. Unique data with corresponding pointers is stored in the repository and the Memory Resident Index is updated with the new similarities. Existing data is not stored.

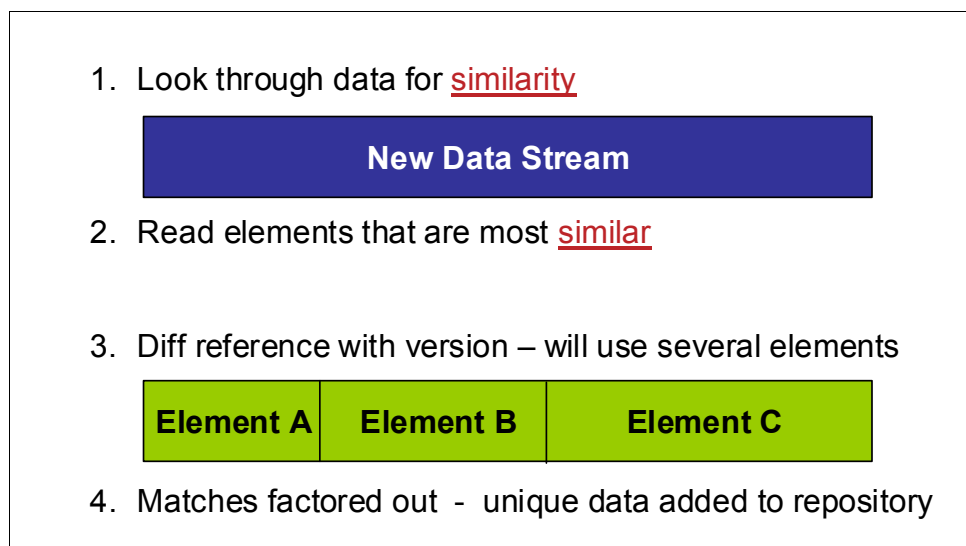


Figure 3-2 Hyperfactor data deduplication

HyperFactor data deduplication uses a 4 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 one PB of data on your disk array. For example, with a ratio of 12 to 1, you could store 12 PB on that one PB of disk array. With the Memory Resident Index, HyperFactor can identify potentially duplicate data quickly for large amounts of data and does this on data ingest, or in-line, reducing the amount of processing required for your data.

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

## **3.3 IBM System Storage TS7600 with ProtecTIER**

In this chapter we introduce the IBM System Storage TS7600 Product Family.

For open systems:

- ▶ TS7610 - Entry Edition
- ▶ TS7650 - Appliance Edition
- ▶ TS7650G - Enterprise Edition or Gateway

For mainframe systems:

- ▶ TS7680 - Gateway Edition for System z®

### **3.3.1 TS7610 - Entry Edition**

IBM System Storage TS7610 ProtecTIER Deduplication Appliance Express is the newest and smallest member of the TS7600 Product Family. For more details, go to 6.1.2, “Hardware and software components of the 3959-SM1” on page 68.

### **3.3.2 TS7650 - Appliance Edition**

IBM System Storage TS7650 ProtecTIER Deduplication Appliance is the midrange member of the TS7600 Product Family. For more details, go to 6.1.3, “Hardware and software components of the 3958-AP1” on page 69.

### **3.3.3 TS7650G - Enterprise Edition**

IBM System Storage TS7650 ProtecTIER Deduplication Gateway is the Enterprise member of the TS7600 Product Family for Open Systems. For more details, go to 6.1.4, “Hardware and software components of the 3958-DD4” on page 70.

### **3.3.4 TS7680 - Gateway Edition for System z**

IBM System Storage TS7680 ProtecTIER Deduplication Gateway for System z is another Enterprise level member of the TS7600 Product Family. For more details, go to 6.1.5, “Hardware and software components of the 3958-DE2” on page 72.

## 3.4 Deduplication with HyperFactor technology

This topic describes IBM's data factoring technology, known as HyperFactor.

ProtecTIER is the first virtual tape product to contain patented data factoring technology that IBM calls HyperFactor. This technology detects recurring data in tape applications. The common data is merged into a single instance store, saving disk space needed to store multiple instances of data without sacrificing performance or the availability of recall of that data. HyperFactor is a breakthrough on several fronts:

- ▶ It scales up to 1024 TB (1 PB).
- ▶ The algorithm used to find the common data between tape processing does not affect the tape processing performance of the virtual tape engine.
- ▶ Data integrity is not compromised, not even statistically.
- ▶ Merged data is stored in a format that preserves restore performance.

### 3.4.1 Impact of HyperFactor

HyperFactor saves space by taking advantage of the fact that only a small percentage of data actually changes from one tape processing window to the next tape processing window (see Figure 3-3).




Figure 3-3 The impact of Hyperfactor

The amount of space saved is a function of many factors, but mostly of the tape processing policies and retention periods and the variance of the data between them. Over time, the effect of HyperFactor is a system-wide factoring ratio. The factoring ratio is derived by dividing the total data before reduction by the total data after reduction.

$$\left( \frac{\text{TotalDataBeforeReduction}}{\text{TotalDataAfterReduction}} \right) = \text{HyperFactorRatio}$$

**Example:** Total Data Before Reduction: 250TB, Total Data After Reduction: 10 TB

250 TB /10 TB = 25 : 1 HyperFactor Ratio

With ProtecTIER the percentage of storage saved can be significant. For the percentage saved at various HyperFactor ratios, see Table 3-1.

*Table 3-1 HyperFactor Ratio translated to Space saved%*

HyperFactor Ratio	Space saved
2 : 1	50%
4 : 1	75%
10 : 1	90%
20 : 1	95%

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

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

For more details about the factoring ratio see “Factoring ratio” on page 20.

### 3.4.2 ProtecTIER data ingest flow

ProtecTIER performs the deduplication process on data ingest, which means data deduplication is performed inline. The data flow is shown in Figure 3-4 on page 27.

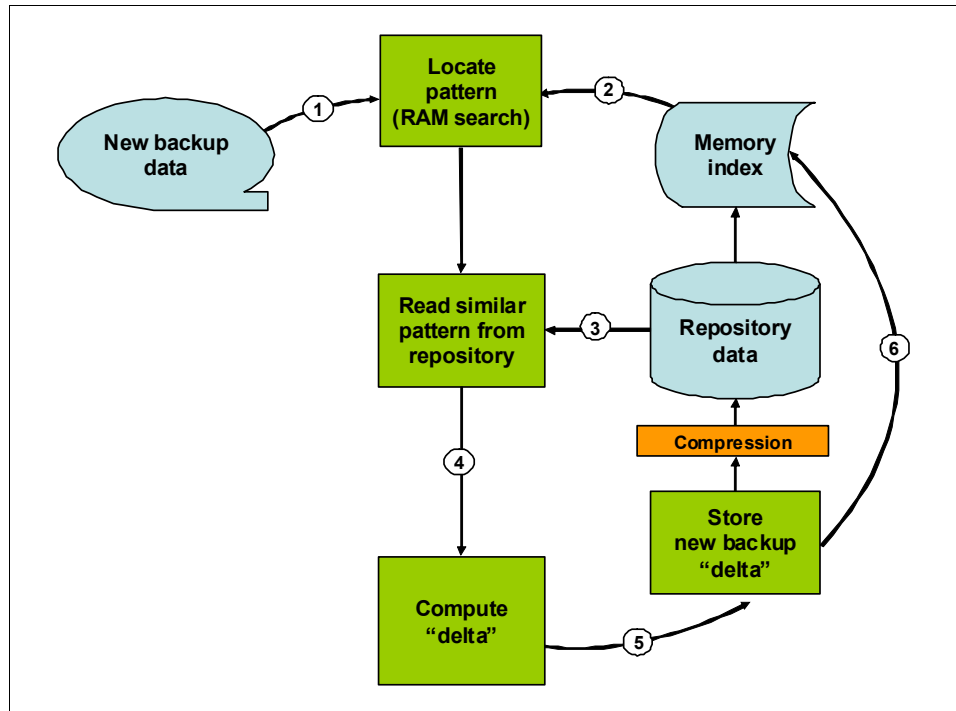


Figure 3-4 Data ingest flow of ProtecTIER

The data flow is as follows:

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

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

### 3.4.3 ProtecTIER VTL concepts

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

When duplicate data is identified by ProtecTIER, it updates a reference count in the database. ProtecTIER uses the reference count to determine when a data segment can be overwritten (deleted). As shown in Figure 3-5, sector 3 represents a segment that occurs four times within the virtual cartridges in the repository. In the lower left corner is a representation of the reference table showing that sector 3 is referenced four times.

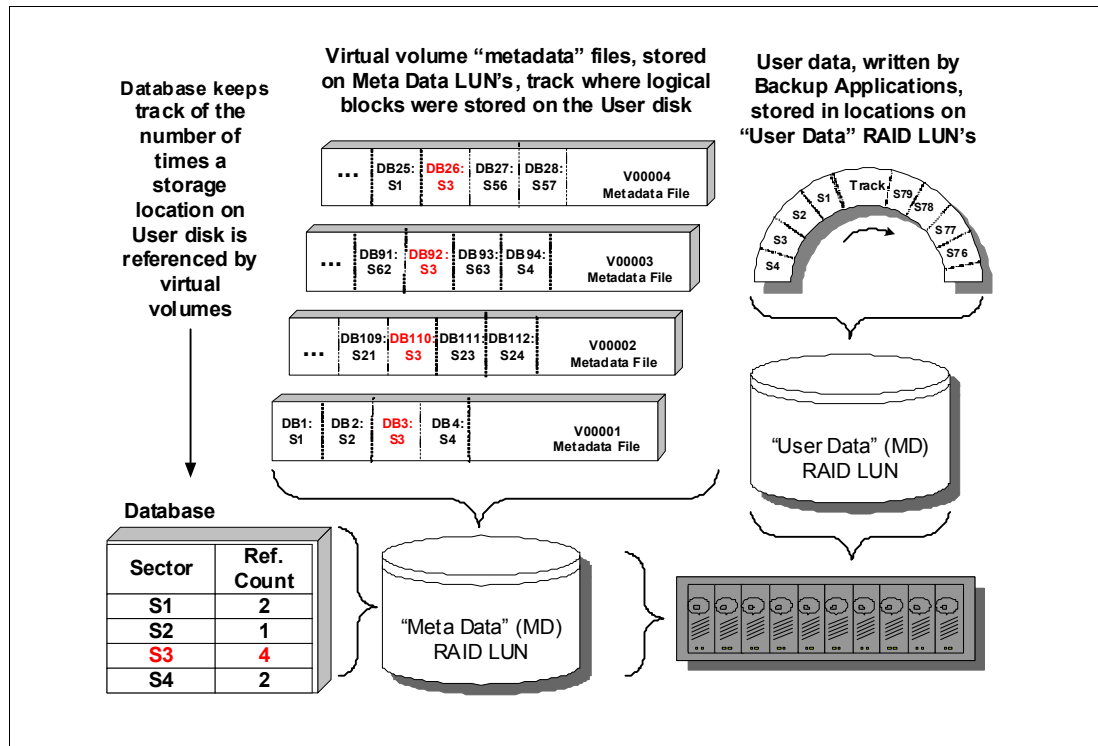


Figure 3-5 ProtecTIER virtual tape library concepts

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

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

### 3.4.4 ProtecTIER OST concepts

OpenStorage (OST) is an Application Programming Interface (API) of Semantic NetBackup.

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

- Backup-to-disk as a NetBackup client



- ▶ Single point of control of NetBackup policies, cartridge sets, and pools
- ▶ Full support through API to simplify writing and managing backup data to disk
- ▶ Connects to NBU through a plug-in on the media server
- ▶ Provides new interface for the storage device called a Logical Storage Unit (LSU)
- ▶ LSUs can be duplicated, moved, and shared by multiple NetBackup media servers

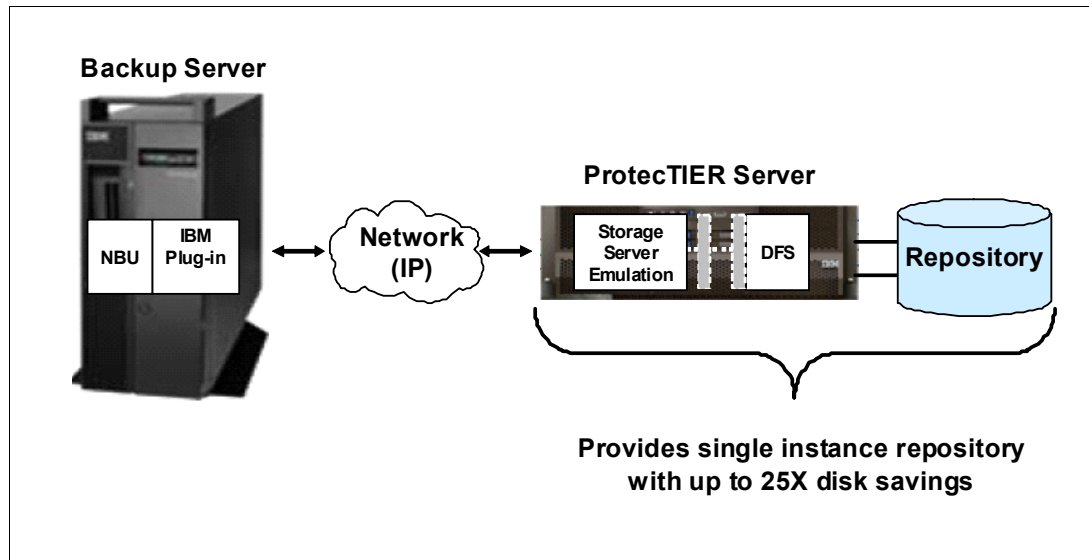


Figure 3-6 OST Overview

### 3.4.5 Steady state

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

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

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

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

The first step of that processing reads the contents of each of the old metadata files, finds every reference to the user data on disk, and decrements the reference count for each storage block identified. After all references in a metadata file are processed, the metadata

file is deleted. Each time that the reference count of a storage block goes to zero, that storage block is returned to the pool of free blocks and becomes usable as free space.

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

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

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

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

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

From a performance standpoint, when the ProtecTIER system first begins ingesting data, ProtecTIER is not matching new data to existing data and no fragmentation is occurring. This enables high performance. After steady state is achieved, performance stabilizes. Figure 3-7 shows the performance from initial implementation through steady state. The change in performance and the time to reach steady state depends on the size the repository and the operational characteristics of your data.

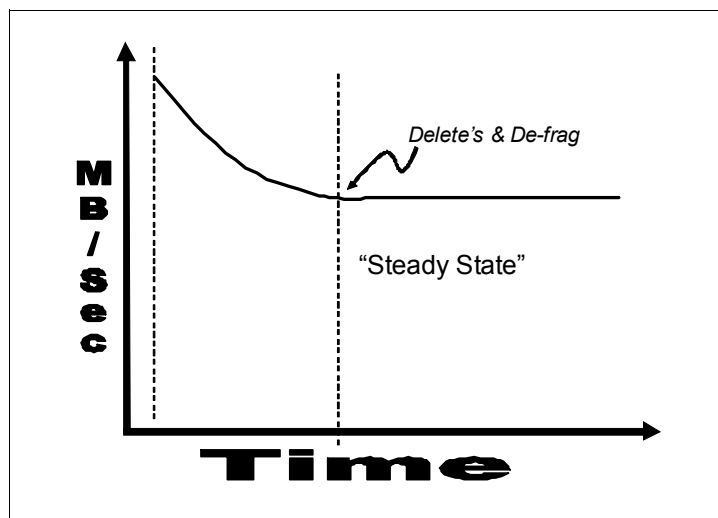


Figure 3-7 ProtecTIER performance from initial implementation through steady state



# Introduction to IBM Tivoli Storage Manager deduplication

In this chapter we describe data deduplication functionality and how deduplication works in IBM Tivoli Storage Manager (ITSM).

This chapter contains the following topics:

- ▶ Deduplication overview
- ▶ How ITSM data deduplication works
- ▶ IBM Tivoli Storage Manager deduplication overview

## 4.1 Deduplication overview

Deduplication is just one of the data reduction methods available in IBM Tivoli Storage Manager. For example, ITSM has contained a duplication avoidance strategy since its inception in 1990 - the progressive incremental backup methodology. This reduces the amount of duplicates for backup data coming into the server, although in a fairly simple fashion. It only backs up files that have changed (for example, one can simply change the modification data of a file and IBM Tivoli Storage Manager will need to back it up again). In terms of effect on stored data, this is similar to data deduplication at the file level (reducing the redundant data at source by not backing up the same file content twice).

Figure 4-1 shows all data reduction methods available on IBM Tivoli Storage Manager.

	Client compression	Incremental forever	Subfile backup	Server-side deduplication	Client-side deduplication
How data reduction is achieved	Client compresses files	Client only sends changed files	Client only sends changed subfiles	Server eliminates redundant data chunks	Client and server eliminate redundant data chunks
Conserves network bandwidth?	Yes	Yes	Yes	No	Yes
Data supported	Backup, archive, HSM, API	Backup	Backup (Windows only)	Backup, archive, HSM, API	Backup, archive, API
Scope of data reduction	Redundant data within same file on client node	Files that do not change between backups	Subfiles that do not change between backups	Redundant data from any files in storage pool	Redundant data from any files in storage pool
Avoids storing identical files renamed, copied, or relocated on client node?	No	No	No	Yes	Yes
Removes redundant data for files from different client nodes?	No	No	No	Yes	Yes

Available prior to V6
Available 6.1
Available 6.2

All of these data reduction methods conserve storage pool space

Figure 4-1 ITSM Data reduction methods

Deduplication is a technique that allows more data to be stored on disks. It works by removing duplicates in the stored version of your data. In order to do that, the deduplication system has to process the data into a slightly different form. When you need the data back, it can be reprocessed into the same form as it was originally submitted.

With deduplication, the larger the quantity being deduplicated, the more opportunity exists to find similar patterns in the data, and the better the deduplication ratio can theoretically be.

Although all workloads can benefit from deduplication, the effectiveness can vary depending on data type, data retention, change rate, and backup policies. It requires that there are similarities in the data being deduplicated: For example if a single file exists more than once in the same store, this could be reduced down to one copy plus a pointer for each deduplicated version (this is often referred to as a “Single Instance Store”). Other workloads such as uncompressible and non-repeated media (JPEGs, MPEGs, MP3, or specialist data such as geo-survey data sets) will not produce significant savings in space consumed. This is because the data is not compressible, has no repeating segments, and has no similar segments.

In many situations, deduplication works better than compression against large data sets, because even with data that is otherwise uncompressible, deduplication offers the potential to efficiently store duplicates of the same compressed file.

To sum up, deduplication typically allows for more unique data to be stored on a given amount of media, at the cost of the additional processing on the way into the media (during writes) and the way out (during reads).

Figure 4-2 illustrates a data deduplication overview.

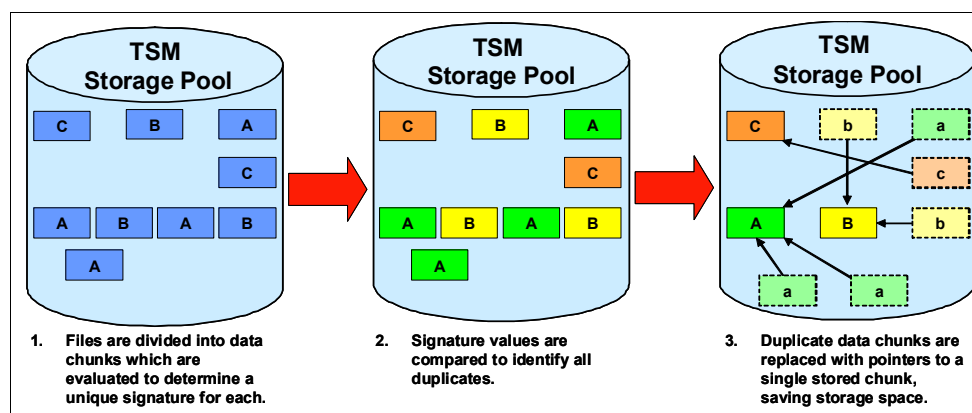


Figure 4-2 Data deduplication overview

## 4.2 How ITSM data deduplication works

There are a large number of approaches to the deduplication of data. A simple way of deduplicating would be to simply look for files of the same name and content, and remove duplicates on that level. However, this would overlook the possibility to deduplicate similar pieces of data *within* different files. For that sort of processing, we need to deal with data at a level below the files themselves, at the block or bit level. We also need a way of referring to the parts of files we consider are worth deduplicating. Some implementations use a fixed size of sub-file called a block, and others use variable sized pieces of file, which we will call *chunks* for the purposes of this book.

In order to process data quickly, many storage techniques use hash functions. A hash function is a process that reads a certain amount of input data (also referred to as a *chunk*), and returns a value that can then be used as a way to refer to that data. An example of this is demonstrated using the AIX `csum` command. We are able to return hash values for a given file with more than one algorithm. The real-world application for `csum` is to check that a file has been downloaded properly from a given Internet site, provided that the site in question has the MD5 listed with the file. Example 4-1 shows a pair of example hashes for a given file: the two hash functions are MD5 and SHA-1. In this case, MD5 produces a 128-bit digest (hash value), and SHA1 produces a 160-bit value. This is shown by the different digest lengths (notice that the result from the first command is shorter than the result from the second command).

*Example 4-1 Hash values of two commonly used functions against the same file*

```
# csum -h MD5 tivoli.tsm.devices.6.1.2.0.bff
05e43d5f73dbb5beb1bf8d370143c2a6  tivoli.tsm.devices.6.1.2.0.bff

# csum -h SHA1 tivoli.tsm.devices.6.1.2.0.bff
```

A typical method of deduplication is to logically separate the data in a store into manageable chunks, then produce a hash value for each chunk, and store those hash values in a table. When new data is taken in (ingested) into the store, the table is then compared with the hash value of each new chunk coming in, and where there's a match, only a small pointer to the first copy of the chunk is stored as opposed to the new data itself.

Typical chunk sizes could be anywhere in the range of 2 KB to 4 MB (average chunk size is 256 KB). There is a trade-off to be made with chunk size: a smaller chunk size means a larger hash table, so if we use a chunk size that is too small, the size of the table of hash pointers will be large, and could outweigh the space saved by deduplication. A larger chunk size means that to gain savings, the data must have larger sections of repeating patterns, so although the hash-pointer table will be small, the deduplication will find fewer matches.

The hashes used in deduplication are similar to those used for security products: MD5 and SHA-1 are both commonly used cryptographic hash algorithms, and both are used in deduplication products, along with other more specialist customized algorithms.

Before Tivoli Storage Manager chunks the data at a bit file object level, it calculates an MD5 of all the objects in question, which are then sliced up into chunks. Each chunk has an SHA-1 hash associated with it, which is used for the deduplication. The MD5s are there to verify that objects submitted to the deduplication system are reformed correctly, because the MD5 is recalculated and compared with the saved one to ensure that returned data is correct.

With any hash there is a possibility of a collision, which is the situation when two chunks with different data happen to have the same hash value. This possibility is extremely remote: in fact the chance of this happening is less likely than the undetected, unrecovered hardware error rate.

Other methods exist in the deduplication technology area which are not hash based, and therefore do not have any logical possibility of collisions. One such method is called hyperfactor; this is implemented in the IBM ProtecTIER storage system and explained in Chapter 3, "Introduction to ProtecTIER deduplication" on page 19.

## 4.3 IBM Tivoli Storage Manager deduplication overview

In ITSM, data deduplication is a method of eliminating redundant data in sequential-access disk primary, copy, and active-data storage pools (FILE type device class). One unique instance of the data is retained on storage media, and redundant data is replaced with a pointer to the unique data copy. With data deduplication you can save storage space and reduce the overall amount of time that is required to retrieve data by letting you store more data on disk, rather than on tape.

### 4.3.1 ITSM server-side deduplication overview

Server-side data deduplication in IBM Tivoli Storage Manager is a two-phase process. In the first phase, duplicate data is identified. During the second phase, duplicate data is removed by certain server processes, such as reclamation processing of storage-pool volumes.

You can deduplicate any type of data except encrypted data. You can deduplicate client backup and archive data, IBM Tivoli Data Protection data, and so on. You can also

deduplicate data that has already been stored. No additional backup, archive, or migration is required.

**Note:** ITSM does not deduplicate random-access disk storage pools (DISK type device class), nor tape storage pools. The primary storage pool must be a sequential-access disk storage pool (FILE type device class) that is enabled for data deduplication. Data deduplication can be enabled by the IBM Tivoli Storage Manager administrator on each storage pool individually, so it is possible to deduplicate those types of data which will benefit the most, as opposed to everything.

Figure 4-3 shows how server-side deduplication works.

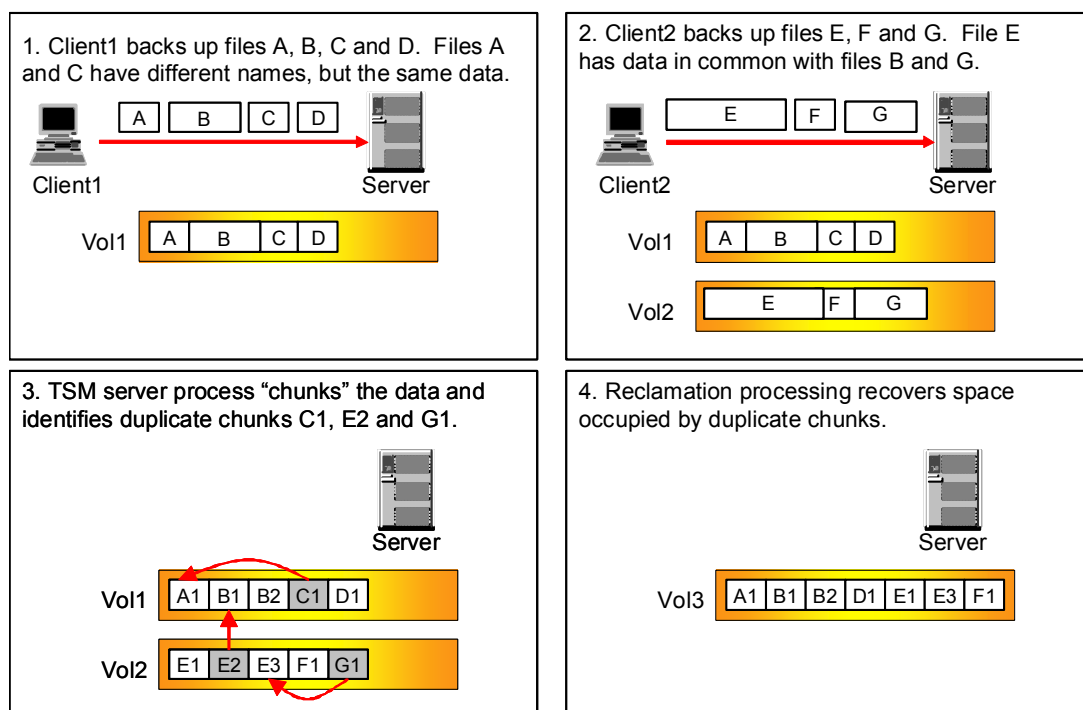


Figure 4-3 ITSM server-side data deduplication

When migrating or copying from storage pools with the FILE device class type to a tape storage pool, ITSM does not store the deduplicated version of the data on the tape device class devices. Instead, it reconstructs the data so that full copies are stored on tape. This is to aid processing for recoveries, which could otherwise become difficult to manage.

### 4.3.2 ITSM client-side deduplication overview

In client-side data deduplication, the IBM Tivoli Storage Manager backup-archive client and the server work together to identify duplicate data.

In IBM Tivoli Storage Manager V6.1, only the server could identify and remove redundant data. Starting in V6.2, you have the option of identifying and removing redundant data during backup and archive processing before data is sent to the server. This method of data deduplication is called client-side data deduplication. It is available with V6.2 backup-archive clients and the V6.2 IBM Tivoli Storage Manager application programming interface (API).

ITSM backup-archive clients earlier than V6.2 can restore deduplicated, compressed data. That means they can access existing deduplicated data and storage pools that are already set up for data deduplication.

When restoring or retrieving files, the client node queries for and displays files as normal. If a user selects a file that exists in a deduplicated storage pool, the server manages the work of reconstructing the file.

Client-side data deduplication uses the following process:

- ▶ The client creates extents. *Extents* are parts of files that are compared with other file extents to identify duplicates.
- ▶ The client and server work together to identify duplicate extents.
- ▶ The client sends non-duplicate extents to the server.

Subsequent client data-deduplication operations create new extents. Some or all of those extents might match the extents that were created in previous data-deduplication operations and sent to the server. Matching extents are not sent to the server again.

Figure 4-4 shows how client-side deduplication works.

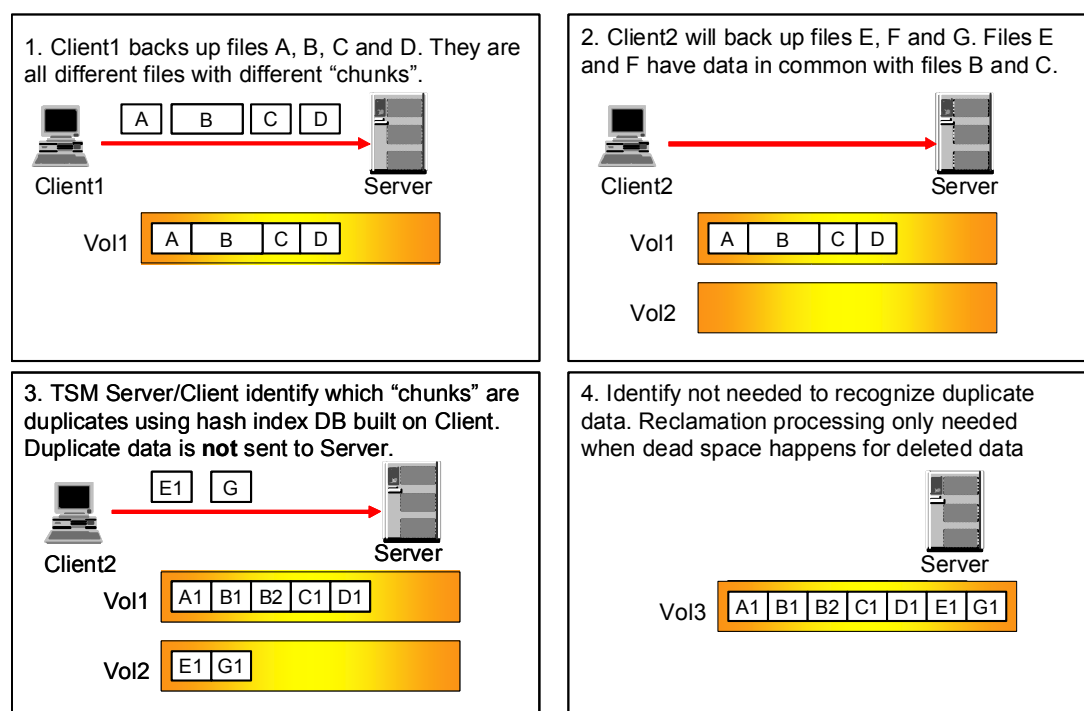


Figure 4-4 ITSM client-side data deduplication

With client-side data deduplication, you can do the following tasks:

- ▶ Exclude specific files on a client from data deduplication.
- ▶ Enable a data deduplication cache that reduces network traffic between the client and the server. The cache contains extents that were sent to the server in previous incremental backup operations. Instead of querying the server for the existence of an extent, the client queries its cache. If an inconsistency between the server and the local cache is detected, the local cache is removed and repopulated.
- ▶ Enable both client-side data deduplication and compression to reduce the amount of data that is stored by the server. Each extent is compressed before being sent to the server.



The trade-off is between storage savings and the processing power that is required to compress client data. In general, if you compress and deduplicate data on the client system, you are using approximately twice as much processing power as data deduplication alone.

You enable client-side data deduplication using a combination of settings on the client node and the ITSM server. The primary storage pool that is specified by the copy group of the management class associated with the client data must be a sequential-access disk storage pool (FILE type device class) that is enabled for data deduplication.

For more information about client-side deduplication, refer to the *Tivoli Storage Manager Backup-archive clients* documentation located at the following URL:

<http://publib.boulder.ibm.com/infocenter/tsminfo/v6r2/topic/com.ibm.itrm.doc/welcome.html>





## Part 2

# Planning for deduplication

In this part we discuss the planning activities necessary for deduplication.





## N series deduplication planning

In this chapter we discuss planning related to deployment of deduplication with N series.

The following topics are covered:

- ▶ Supported hardware models and ONTAP versions
- ▶ Deduplication and Data ONTAP version considerations
- ▶ Deduplication licensing
- ▶ Compatibility with native N series functions
- ▶ Compatibility with non-native functions
- ▶ Data Characteristics
- ▶ Deduplication and storage capacity
- ▶ Deduplication and performance
- ▶ Deduplication scheduling
- ▶ Aggregate and volume considerations

## 5.1 Supported hardware models and ONTAP versions

Table 5-1 shows the models that support deduplication and the minimum ONTAP version that must be running.

*Table 5-1 Supported models*

N series Model	Is deduplication supported?	Minimum ONTAP version
N3300	Yes	7.2.5.1
N3400	Yes	7.3.2
N3600	Yes	7.2.5.1
N3700	No	Not compatible
N5200	Yes	7.2.2
N5300	Yes	7.2.2
N5500	Yes	7.2.2
N5600	Yes	7.2.2
N6040	Yes	7.2.5.1
N6060	Yes	7.2.5.1
N6070	Yes	7.2.2
N7600	Yes	7.2.2
N7700	Yes	7.2.4
N7800	Yes	7.2.5.1
N7900	Yes	7.2.4
All Gateway models	Yes	7.3

## 5.2 Deduplication and Data ONTAP version considerations

Deduplication involves the following considerations depending on the version of Data ONTAP that you are using:

- ▶ Data ONTAP 7.2.6.1 and later: Performance of sequential read operations involving highly shared or deduplicated blocks is improved by more efficient caching.
- ▶ Data ONTAP 7.2.x: Deduplication metadata is stored in the volume being deduplicated.
- ▶ Data ONTAP 7.3 and later:
  - Part of the deduplication metadata is stored at the aggregate level rather than within the volume. The metadata location impacts capacity and replication considerations. Refer to 5.7.1, “Metadata” on page 60 for more details.
  - The configuration data for deduplication of a given volume is located in a metafile within the deduplicated volume. Previously, it was located in the registry within the root volume. This change helps volume-level operations such as Volume SnapMirror®, Snap Restore, vol copy, and vol clone keep the deduplication configuration data consistent between the source and the destination.

- ▶ Data ONTAP 7.3.1 and later:
  - Deduplication volume limits are increased over prior versions of Data ONTAP. Refer to 5.10, “Aggregate and volume considerations” on page 62, for more details.
  - Provides checkpoint restarts. If the deduplication process is interrupted, it can later restart and continue processing the same changelog from a checkpoint rather than from the beginning of the last run.
- ▶ Data ONTAP 8: No longer requires the NearStore license for Deduplication. Only the A-SIS license is needed. See 5.3, “Deduplication licensing” on page 43 for more details.

See 5.4, “Compatibility with native N series functions” on page 43 and 5.5, “Compatibility with non-native functions” on page 57 for Data ONTAP compatibility with other functions.

## 5.3 Deduplication licensing

Prior to Data ONTAP 8, two licenses must be installed for deduplication to be enabled:

- ▶ NearStore license (no longer required for deduplication if using Data ONTAP 8)
- ▶ A-SIS license

Both of the licenses are available independently at no additional cost. Contact your procurement representative for order details.

With HA-paired filers, all partnered nodes should have the licenses installed

For licensing requirements with SnapMirror source and destination filers, see 5.4.6, “Deduplication and SnapMirror replication” on page 45.

Before removing the A-SIS license, you must disable deduplication on any flexible volumes for which it is enabled. If you attempt to remove the license without first disabling deduplication, you will receive a warning message asking you to disable this feature.

If licensing is removed or expired, no additional deduplication can occur and no deduplication commands can be run. However, the deduplicated volumes remains deduplicated, the existing storage savings are kept, and all data is still accessible.

## 5.4 Compatibility with native N series functions

This section will discuss compatibility topics with deduplication

### 5.4.1 Deduplication and high availability pairs

Standard and Mirrored HA pairs support deduplication with the same hardware and ONTAP requirements as stand-alone systems. The requisite deduplication licenses must be installed on all paired nodes.

**Note:** Prior to ONTAP 8, these were referred to as Active/Active Configurations or software.

When no takeover has been done, deduplication on each node will work independently: the standard deduplication requirements and limitations are applicable to each node separately.

Upon takeover, deduplication processing stops for the unavailable partner node's volumes. Change logging continues on the active partner until the change log is full. All data continues to be accessible regardless of change log availability. Deduplication operations for volumes owned by the active partner will continue as scheduled but CPU overhead could become a bigger impact in conjunction with the increased general workload from the takeover. Because of the increased workload during takeover, disable deduplication on the active partner until after giveback.

After giveback occurs, data in the change logs is processed during the next deduplication run (scheduled or manual).

Stretch and Fabric MetroClusters support deduplication starting with Data ONTAP 7.3.1

## 5.4.2 Deduplication and non-disruptive upgrades

Non-disruptive major and minor Data ONTAP upgrades (NDU) are supported within one minute when deduplication is enabled on 100 FlexVol volumes or less, provided that no deduplication operations are active during the software upgrade.

The total number of FlexVol volumes on the storage system must not exceed the total number of FlexVol volumes supported for non disruptive upgrades on your system.

To meet the 60-second limit, users should perform the Data ONTAP upgrade during a time when deduplication operations are not scheduled to run and halt any active deduplication operations in progress.

## 5.4.3 Deduplication and Performance Acceleration Modules

The Performance Acceleration Module (PAM) card is supported with Data ONTAP 7.3 and later.

In environments where there are deduplicated blocks that are read repeatedly, the PAM card can help reduce the number of disk reads, thus improving the read performance. Performance Acceleration Modules are deduplication aware and use intelligent caching. The amount of performance improvement with the PAM card depends on the duplication rate, the access rate, the active data set size, and the data layout.

Adding a PAM card to a system does not increase the deduplication maximum volume size for that system.

**Note:** When you run a deduplication scan on a volume with the 16 GB Performance Acceleration Module installed, you obtain suboptimal space savings.

## 5.4.4 Deduplication and Snapshot copies

Deduplication processes the active file system by consolidating identical data blocks and then releasing the redundant blocks. If blocks are already referenced by snapshots, then the referenced blocks are locked, which prevents them from being released through deduplication.

There are two ways that Snapshots can decrease deduplication storage efficiency:

- Duplicate data can be locked in retained snapshots from before deduplication was enabled.



- ▶ Deduplication metadata within the volume can be locked into a snapshot when the copy is created. In Data ONTAP 7.2.x, all of the deduplication metadata resides within the volume and is susceptible to snapshot locking. Starting with Data ONTAP 7.3.0, the deduplication change logs and fingerprint database have been relocated to the aggregate level outside of the volume. The temporary metadata files that are still in the volume are deleted when the deduplication operation completes. This change enables deduplication to achieve higher space savings with Data ONTAP 7.3 if using Snapshots.

For deduplication to provide the greatest storage efficiency when used in conjunction with Snapshot copies, consider the following during planning:

- ▶ Remove old Snapshot copies maintained in deduplicated volumes to release duplicate data locked in snapshots taken prior to deduplication.
- ▶ Reduce the retention time of Snapshot copies maintained in deduplicated volumes.
- ▶ Run deduplication before creating new Snapshots and make sure deduplication has enough time to complete so the temporary metadata is cleared.
- ▶ Schedule deduplication only after significant new data has been written to the volume.
- ▶ Configure the appropriate reserve space for the Snapshot copies.
- ▶ If the snap reserve is zero, turn off the Snapshot auto-create schedule. This is the case in most LUN deployments.
- ▶ If the space used by Snapshot copies grows to more than 100%, reports obtained by running the **df -s** command will show incorrect results because some space from the active file system is being taken away by Snapshot and therefore actual savings from deduplication are not reported.

### 5.4.5 Deduplication and SnapRestore

The SnapRestore® functionality is supported with deduplication, and it works in the same way with deduplication as it does without deduplication. If you are running Data ONTAP 7.3 or later, note the following information:

- ▶ Starting with Data ONTAP 7.3, the deduplication metadata files (the fingerprint database and the change log files) do not get restored when SnapRestore is executed because they are located outside the volume in the aggregate. In this case, after the SnapRestore operation, there is not a fingerprint database file in the active file system for the data. This data, however, retains the original space savings. After SnapRestore, if deduplication is enabled on the volume, any new data written to the volume continues to be deduplicated. However, the deduplication process obtains space savings in the new data only, and does not deduplicate between the new data and the restored data.

To run deduplication for all the data in the volume (and thus obtain higher space savings), run the **sis start -s** command. This command builds the fingerprint database for all the data in the volume. Depending on the size of the logical data in the volume, this process can take a long time to complete.

- ▶ Before using the **sis start -s** command, make sure that both the volume and the aggregate containing the volume have sufficient free space to accommodate the addition of the deduplication metadata. For information about how much extra space to leave for the deduplication metadata, refer to 5.7.1, “Metadata” on page 60.

### 5.4.6 Deduplication and SnapMirror replication

Deduplication is supported with both Volume SnapMirror (VSM) and Qtree SnapMirror (QSM) but there are important differences to be aware of.

## Deduplication with Volume SnapMirror

VSM is physical replication: it functions at the block level for the entire volume. This means that deduplicated blocks at the source stay deduplicated over the network and onto destination filer. Deduplication metadata within the volume is also replicated, as are snapshots. Deduplication is managed at the source and inherited at the destination.

A flexible volume can be replicated to a secondary storage system (destination) by using Volume SnapMirror (VSM), as shown in Figure 5-1.

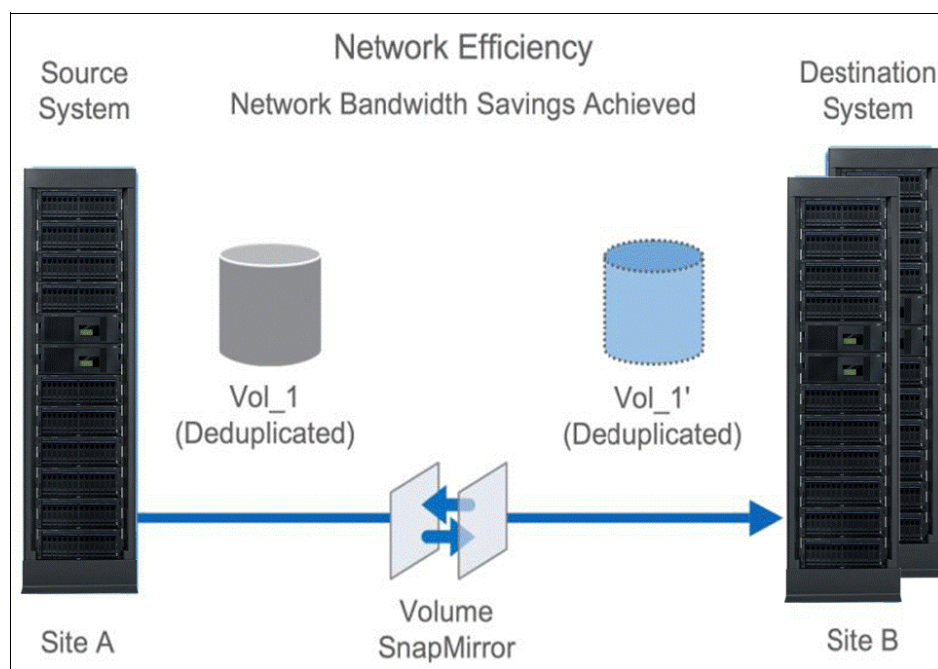


Figure 5-1 Replication to destination

Running deduplication with Volume SnapMirror has these characteristics:

- ▶ Deduplication is only enabled, run, scheduled, and managed from the primary location.
- ▶ The flexible volume at the secondary location *inherits* all of the attributes and storage savings through SnapMirror.
- ▶ Only unique blocks are transferred, so deduplication reduces network bandwidth usage as well.
- ▶ The maximum volume size limit is imposed based on the lower value of the maximum volume size limit of the source and destination volumes.

### Licensing

The A-SIS license must be installed at the primary location (source) but is not strictly required at the destination. As a best practice, the A-SIS license should also be installed on the destination filer if there is ever the intention of having it run deduplication on volumes locally (for example, if the destination filer ever becomes the primary such as in a DR scenario).

The NearStore license must be installed on both the source and destination with Data ONTAP versions prior to version 7.3.1. Starting with Data ONTAP 7.3.1, the NearStore license is no longer required on the destination system. With Data ONTAP 8, no NearStore license is required for deduplication on either the source or destination.

SnapMirror licensing requirements are unchanged from the standard SnapMirror requirements when used concurrently with deduplication.

### ***Volume SnapMirror modes***

Deduplication can be turned on for a source volume with synchronous, semi-synchronous, and asynchronous Volume SnapMirror. However, do not use deduplication on volumes with synchronous SnapMirror because, although technically it can work, there is increased potential for negative performance impact.

### ***Scheduling***

When configuring both Volume SnapMirror and Deduplication for a volume, consider the scheduling for each. The processes are not aware of each other and scheduling must be managed independently. Start VSM on the intended volume only after deduplication of that volume has completed. This approach avoids capturing undeduplicated data and deduplication metadata in the VSM Snapshot, and then replicating the extra data over the network and onto the destination filer.

Regular Snapshot and deduplication metadata storage efficiency and scheduling concerns are potentially amplified with VSM because these regular snapshots will also be replicated to the destination storage system. Refer to 5.4.4, “Deduplication and Snapshot copies” on page 44. Deduplication and VSM processes can also potentially contend for CPU cycles if running concurrently. Refer to 5.8, “Deduplication and performance” on page 60.

### ***Breaking the VSM relationship***

In case of a disaster at the primary location, you might have to break the Volume SnapMirror relationship and have the Volume SnapMirror destination start serving data. If using ONTAP 7.3 or later on the source storage system, the fingerprint database is stored at the aggregate level rather than within the volume. This reduction of data adds replication efficiency but it also means the fingerprint database will be missing from the destination volume.

If the destination data becomes primary, the existing deduplicated data retains the space savings from the deduplication operations performed earlier on the original VSM source. Also, the deduplication process continues for new data being written to the volume and creates the fingerprint database for this new data. The deduplication process obtains space savings in the new data only, and does not deduplicate between the new data and the old data. To run deduplication for all of the data in the volume (and thus obtain higher space savings), use the **sis start -s** command. This command rebuilds the fingerprint database for all of the data in the volume.

**Important:** Before running deduplication, make sure that both the volume and the aggregate containing the volume have sufficient free space to accommodate the addition of the deduplication metadata.

## **Deduplication with Qtree SnapMirror Overview**

QSM is logical replication: it functions at the file level for qtrees. This means that QSM is agnostic to some of the block and volume-specific considerations of VSM. There are no block-level deduplication savings over the network with QSM.

Neither deduplication metadata nor standard snapshots are replicated with QSM, but it does use SnapMirror-specific snapshots to assess changes related to the qtree. Any deduplicated files are considered changed files and therefore are transferred to the destination. This can result in a larger than usual transfer after first enabling deduplication. Deduplication is run and scheduled independently on the source and destination volumes containing the SnapMirrored qtrees.

A qtree can be replicated to a secondary storage system (destination) by using Qtree SnapMirror (QSM), as shown in Figure 5-2, Figure 5-3, and Figure 5-4 on page 49.

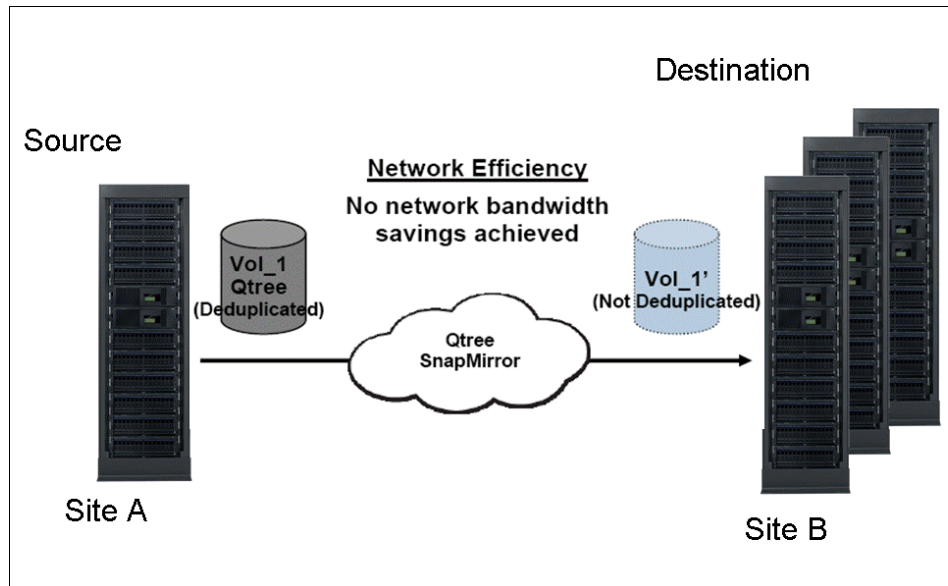


Figure 5-2 Qtree SnapMirror

Figure 5-2 shows Qtree SnapMirror replication from a deduplicated source volume to a non deduplicated destination volume.

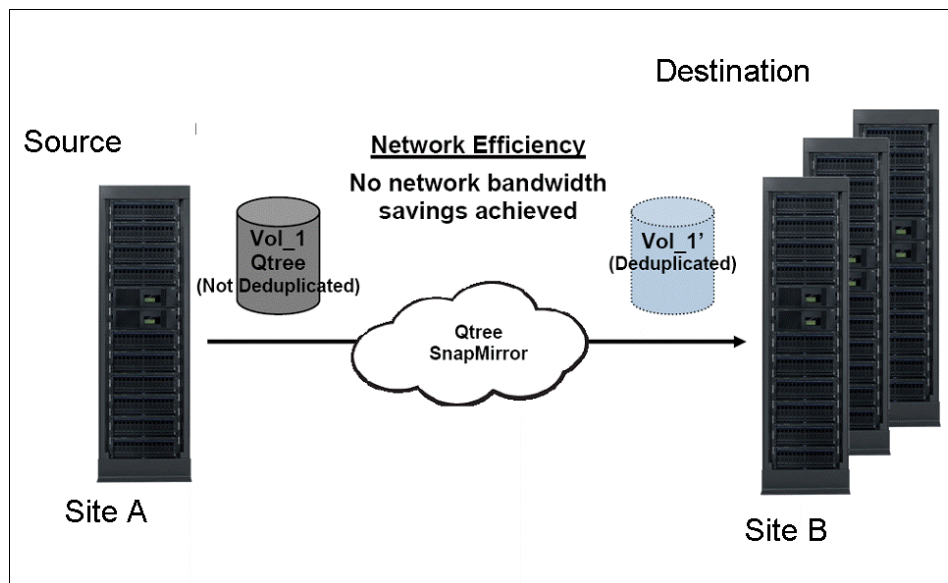


Figure 5-3 Non deduplicated to deduplicated

Figure 5-3 shows Qtree SnapMirror replication from a non deduplicated source volume to a deduplicated destination volume.

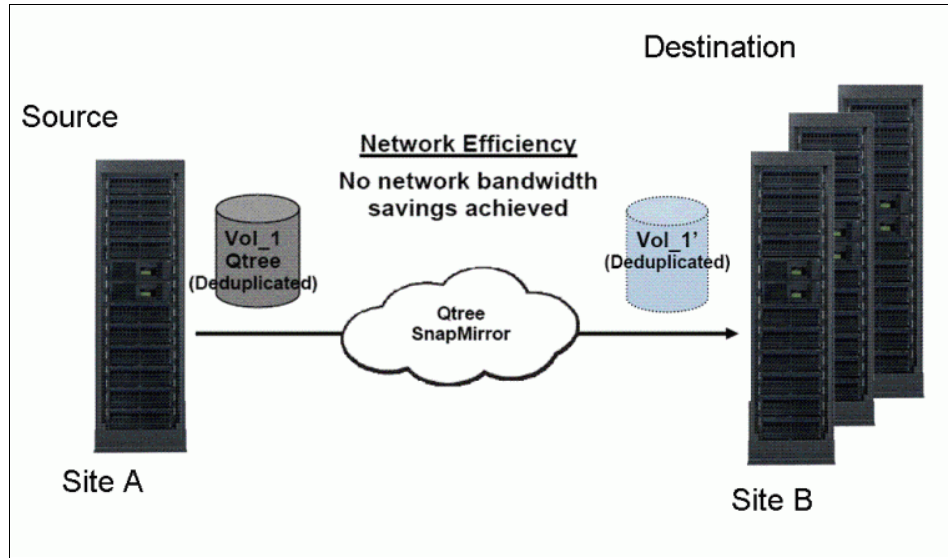


Figure 5-4 Deduplicated volume to another deduplicated volume

Figure 5-4 shows Qtree SnapMirror replication from a deduplicated source volume to a deduplicated destination volume.

### **Licensing**

The A-SIS license must be installed on any storage systems running deduplication whether source, target, or both.

The NearStore license must be installed on any storage systems running deduplication whether source, target, or both. With Data ONTAP 8, no NearStore license is required for deduplication on either the source or destination.

The SnapMirror licensing requirements are unchanged from the standard SnapMirror requirements when used concurrently with deduplication.

### **Qtree SnapMirror modes**

QSM only works with asynchronous mode.

### **Scheduling**

Deduplication savings are not transferred with Qtree SnapMirror. Deduplication needs to be configured and scheduled independently at the source and destination, if desired. It can be run at either or both independently. Also, the QSM and deduplication processes are not aware of each other and scheduling for these must be managed separately. As a best practice, start deduplication of the destination qtree only after QSM of that qtree has completed.

Deduplication and QSM processes can potentially contend for CPU cycles if running concurrently.

## **5.4.7 Deduplication and SnapVault**

Data ONTAP 7.3 or later is required to deduplicate the SnapVault® destination rather than just the source. The behavior of deduplication with SnapVault is similar to the behavior of deduplication with Qtree SnapMirror in that it is managed independently at source and destination. Distinctions about using deduplication with SnapVault include:

- ▶ The deduplication schedule on the destination volume is linked to the SnapVault schedule. Every SnapVault update (baseline or incremental) kicks off the deduplication process on the destination after the archival Snapshot is taken. The deduplication schedule on the source is still independent of SnapVault, just like qtree SnapMirror.
- ▶ The deduplication schedule on the destination cannot be configured manually, and the **sis start** command is not allowed either. However, the **sis start -s** command can be run manually on the destination.
- ▶ The archival Snapshot is replaced with a new one after deduplication has finished running on the destination. The name of this new Snapshot is the same as that of the archival copy, but the creation time of this copy is changed.
- ▶ The SnapVault update is not dependent on completion of the deduplication operation on the destination. Subsequent SnapVault incremental updates can run while the deduplication process on the destination volume from the previous backup is still in progress. In this case, the deduplication process continues to run, but the archival Snapshot does not get replaced until after deduplication has finished running.
- ▶ When using SnapVault, the maximum volume sizes for deduplication for the primary and secondary are independent of one another. Volumes on each of the systems will need to abide by their respective maximum volume size limits.
- ▶ Starting with Data ONTAP 7.3, if using SnapVault with NetBackup, block sharing is supported for partner volumes in takeover mode.
- ▶ When deduplication is run on an existing SnapVault source for the first time, all saved space is transferred to the destination system because the next SnapVault update will recognize deduplicated blocks as changed blocks. Because of this, the size of that transfer can be several times larger than the regular transfers. Running deduplication on the source system periodically will help prevent this issue for future SnapVault transfers. If possible, you should run deduplication before the SnapVault baseline transfer.
- ▶ Protection Manager 3.8 or later can be used to manage deduplication with SnapVault.
- ▶ With Data ONTAP 7.3, SnapVault integration with deduplication replaces Snapshot copies. As a result, Protection Manager has to wait for deduplication to finish before renaming Snapshot copies. During the time that Protection Manager waits, it does not allow clients to list the Snapshot copies or restore from them. This can adversely affect the recovery point objective.
- ▶ Data ONTAP 7.3 is required for use with Open Systems SnapVault (OSSV).

### 5.4.8 Deduplication and SnapLock

With Data ONTAP 7.3.1, deduplication is fully supported with SnapLock®, including both enterprise and compliance modes. The following items should be taken into consideration:

- ▶ A SnapLock volume with files committed to “Write once, Read Many” (WORM) can be deduplicated. Capacity savings will be similar to savings where the files were not committed to WORM. Both deduplication and subsequent undeduplication do not result in any changes to the SnapLock attributes or WORM behavior of the volume or the file.
- ▶ Deduplication is applied across WORM, WORM append, and non-WORM (normal) files.
- ▶ Volume restore from a Snapshot is only permitted on SnapLock enterprise volumes. When a volume restore occurs from a Snapshot with deduplicated data, the file system returns to the state at which the Snapshot was created, including the state of deduplication, and the (WORM) status of the volume and the files.
- ▶ File folding will continue to function, irrespective of the WORM and deduplication status of the files.

- ▶ Autocommit functions irrespective of the deduplication status of the files.
- ▶ When using Qtree SnapMirror, deduplication needs to be run separately on the source and destination. The WORM property is carried forward by qtree SnapMirror. Switching on WORM or deduplication on either end has no effect on the qtree SnapMirror transfers. Undoing deduplication will also have no effect when done on either the source or the destination.
- ▶ When using Volume SnapMirror, the WORM property of the files is carried forward by Volume SnapMirror. Deduplication only needs to be run on the primary. Volume SnapMirror allows the secondary to inherit the deduplication. Undoing deduplication can only be done after breaking the Volume SnapMirror relationship.
- ▶ To revert to a previous release of ONTAP on a system hosting a volume that is deduplicated and has SnapLocked data on it, the volume must first be undeduplicated.

### 5.4.9 Deduplication and MultiStore (vFiler)

Starting with Data ONTAP 7.3.0, deduplication is fully supported with MultiStore®. The deduplication commands are available only at the CLI of vFiler, but they allow any volume to be included in the command arguments regardless of which vFiler the volume is associated with.

Beginning with Data ONTAP 7.3.1, the deduplication commands are available in the CLI of each vFiler unit, allowing each vFiler unit to be configured from within itself.

### 5.4.10 Deduplication and LUNs

When using deduplication in a file-based (NFS or CIFS) environment, deduplication savings are relatively straightforward and automatic. As duplicate blocks are freed, they are marked as available and the IBM System Storage N series storage system recognizes these free blocks and makes them available to the volume.

Deduplication in a block-based (FCP or iSCSI) LUN environment is slightly more complicated because of the space guarantees and fractional reservations often used by LUNs. With space guarantees, for example, a 500 GB LUN that is created consumes exactly 500 GB of physical disk space. If the data in the LUN is reduced through deduplication, the LUN still reserves the same physical space capacity of 500 GB, and the space savings are not apparent to the user.

#### Definitions

Several definitions are used in the examples. They are defined as follows:

<b>Volume Space Guarantee</b>	This is a volume option that specifies whether the volume's space is reserved from the aggregate's free pool at the time of the volume creation. Setting this option to "none" is essentially thin provisioning of the volume.
<b>Fractional Reserve</b>	Fractional reserve is a volume option that enables you to determine how much space Data ONTAP reserves for Snapshot copy overwrites for LUNs, and for space-reserved files when all other space in the volume is used. Default is 100%. The behavior of the fractional reserve space parameter with deduplication is the same as though a snapshot has been taken in the volume and blocks are being overwritten.

<b>LUN Space Reservation</b>	This is an N series LUN option that insures 100% of the LUN's space is removed from the volume's free pool at the time the LUN is created.
<b>Volume free pool</b>	This pool refers to the free blocks in the parent volume of the LUN. These blocks can be assigned anywhere in the volume as necessary.
<b>Aggregate free pool</b>	This pool refers to the free blocks in the parent aggregate of volume containing the LUN. These blocks can be assigned anywhere in the aggregate as necessary.

Volume space guarantees, LUN space reservations, and fractional reserves can be configured so that the use of the freed blocks by the IBM System Storage N series storage system changes depending on the configuration. By varying the values, freed blocks can be returned to the LUN overwrite reserve, the volume free pool, the aggregate free pool, or a combination.

## LUN configuration examples

This section describes five common examples of LUN configurations and deduplication behavior.

Table 5-2 summarizes the five configurations (A - E).

Table 5-2 Configuration examples

Configuration	A (default)	B	C	D	E
LUN Space Reservation Value	on	on	on	off	off
Volume Fractional Reserve Value	100	1-99	0	any	any
Volume Space Guarantee	volume	volume	volume	volume	none
After deduplication and thin provisioning (if applicable), free blocks are returned to:	Fractional Overwrite Reserve	Fractional Overwrite Reserve + Volume Free Pool	Volume Free Pool	Volume Free Pool	Aggregate Free Pool

### Configuration A: The default LUN configuration

The default configuration of an IBM System Storage N series LUN is for all IBM System Storage N series LUNs is to turn off the controller Snapshot, delete all the scheduled Snapshot copies, and set the Snap Reserve to 0% as shown in Table 5-3.

Table 5-3 Configuration A settings

Setting	Default
LUN Space Reservation value = on	on
Volume Fractional Reserve Value = 100	100%
Volume Space Guarantee = volume	volume



### Description

When a LUN containing default values is deduplicated, no apparent savings are observed by the storage administrator because the LUN was, by default, *space reserved* when it was created and fractional reserve was set to 100% in the volume. Any blocks freed through deduplication are allocated to the fractional reserve area. This configuration means that overwriting to the LUN should never fail, even if it is overwritten entirely.

### Advantages and disadvantages

The advantage of this configuration is that Snapshot copies consume less space when blocks in the active file system are no longer being used. As a result, this volume can hold more Snapshot copies. The disadvantage of this configuration is that free blocks are not returned to either the free volume pool or the free aggregate pool. Moreover, there is no direct space savings in the active file system. In fact, this configuration could consume more space in the volume because of new indirect blocks being written, if no Snapshot copies exist in the volume and the Snapshot schedule is turned off.

**Note:** If Snapshot copies are turned off for the volume (or no copy exists in the volume), this is not a recommended configuration for a volume with deduplication.

### Configuration B: LUN configuration for shared volume space savings

To apply freed blocks to both the fractional overwrite reserve area and the volume free pool, use the configuration settings given in Table 5-4.

Table 5-4 Configuration B settings

Setting	Value
LUN Space Reservation value	on
Volume Fractional Reserve Value	1 - 99%
Volume Space Guarantee	volume

### Description

The only difference between this configuration and configuration A is that the amount of space reserved for overwriting is based on the fractional reserve value set for the volume. As a result, this configuration splits the free blocks between fractional overwrite reserve and volume free space. For example, if the fractional reserve value is set to 25, then 25% of the freed blocks go into fractional overwrite reserve and 75% of the freed blocks are returned to the volume free pool.

### Advantages and disadvantages

The advantage of this configuration is that the overwrite space reserve does not increase for every block being deduplicated. Freed blocks are split between volume free pool and fractional reserve. The disadvantage of this configuration is that overwrites to the LUN beyond the fractional reserve capacity can fail because freed blocks might have already been allocated. Another disadvantage of this configuration is that freed blocks stay in the parent volume and cannot be provisioned to any other volumes in the aggregate.

**Note:** If Snapshot copies are turned off for the volume (or if no Snapshot exists in the volume) and the percentage of savings is less than the fractional reserve because of deduplication, this configuration is not recommended for a volume with deduplication.

### Configuration C: LUN configuration for maximum volume space savings

To apply freed blocks to the volume free pool, use the configuration settings given in Table 5-5.

Table 5-5 Configuration C settings

Setting	Value
LUN Space Reservation value	on
Volume Fractional Reserve Value	0%
Volume Space Guarantee	volume

#### Description

In this configuration, the value of fractional reserve is set to zero. As a result, this configuration *forces* all freed blocks to the volume free pool and no blocks are set aside for fractional reserve.

#### Advantages and disadvantages

The advantage of this configuration is that all freed blocks are returned to the volume free pool. The disadvantage is that the chance of overwriting failure is higher than with configurations A and B because no freed blocks are assigned to the fractional overwrite area.

### Configuration D: LUN configuration for maximum volume space savings

To apply freed blocks to the volume free pool, use the configuration settings given in Table 5-6.

Table 5-6 Configuration D settings

Setting	Value
LUN Space Reservation value	off
Volume Fractional Reserve Value	0 - 100%
Volume Space Guarantee	volume

#### Description

The difference between this configuration and Configuration C is that the LUN is not space reserved. With LUN space guarantees off, the value for volume fractional reserve is ignored for all LUNs in this volume. From a deduplication perspective, there is no difference between this and the previous configuration, and all freed blocks go to the volume free pool.

#### Advantages and disadvantages

From a deduplication perspective, this configuration has same advantages and disadvantages as Configuration C.

### Configuration E: LUN configuration for maximum aggregate space savings

In many cases, the user might prefer to reclaim all freed blocks from the volume and return these blocks to the aggregate free pool. To accomplish this task, use the configuration settings in Table 5-7.

Table 5-7 Configuration E settings

Setting	Value
LUN Space Reservation value	off
Volume Fractional Reserve Value	0 - 100%
Volume Space Guarantee	none

### Description

This configuration *forces* the free blocks out of the volume and into the aggregate free pool, where the blocks can be reallocated for any other volumes in the aggregate.

### Advantages and disadvantages

The advantage of this configuration is that it provides the highest efficiency in aggregate space provisioning. Deduplicated blocks will be returned to the aggregate free pool. This configuration is one method of thin provisioning at the volume level. Because the volume's allocated capacity is no longer 1:1 with the physical storage in the aggregate, the volumes can easily be overallocated. Unchecked growth can result in full aggregates before the volumes fill to their allocated size. This configuration requires the storage administrator to carefully monitor and manage the free space available in the aggregates.

## 5.4.11 Deduplication and the volume copy command

When deduplicated data is copied by using the **vol copy** command, the copy of the data at the destination location inherits all the deduplication attributes and storage savings of the original data.

Starting with Data ONTAP 7.3, some of the deduplication metadata files do not get copied by the **vol copy** command because they are located outside of the volume in the aggregate. In this case, there is no fingerprint database file in the destination volume for the data. However, the data retains the space savings. The deduplication process also continues for any new data written to the destination volume, and creates the fingerprint database for the new data. The deduplication process obtains space savings in the new data only, and does not deduplicate between the new data and the old data. To run deduplication for all the data in the cloned volume (and thus obtain higher space savings), use the **sis start -s** command. Depending on the size of the logical data in the volume, this process can take a long time to complete.

## 5.4.12 Deduplication and FlexClone volumes

When a FlexClone® volume is created:

- ▶ The FlexClone volume of a deduplicated volume is also a deduplicated volume.
- ▶ The cloned volume inherits the deduplication configuration of the parent volume, such as the deduplication schedule.
- ▶ Starting with Data ONTAP 7.3, the deduplication metadata files (the fingerprint database and the change log files) do not get cloned, because they are located outside the volume in the aggregate. In this case, there is no fingerprint database file in the cloned volume for the data that came from the parent. However, the data in the cloned volume inherits the space savings of the original data. The deduplication process also continues for any new data written to the clone, and creates the fingerprint database for the new data. However, the deduplication process obtains space savings in the new data only, and does not deduplicate between the new data and the old data. To run deduplication for all the data in the cloned volume (and thus obtain higher space savings), use the **sis start -s**

command. Depending on the size of the logical data in the volume, this process can take a long time to complete.

- ▶ Beginning with Data ONTAP 7.3.1, in addition to standard FlexClone, FlexClone at the file and LUN level is available and allowed on deduplicated volumes.
- ▶ Deduplication can be used to regain capacity savings on data that was copied using FlexClone at the file or LUN level, and has been logically migrated (that is, with Qtree SnapMirror, SnapVault, NDMP dump, and so on).

### **Volume splitting**

When a cloned volume is split from the parent volume, all of the original data in the clone is undeduplicated after the volume split operation. If deduplication is running on the cloned volume, this data gets deduplicated again in subsequent deduplication operations on the volume.

## **5.4.13 Deduplication and read reallocation**

For workloads that perform a mixture of random writes, and large and multiple sequential reads, read reallocation (REALLOC) improves the file layout to achieve faster sequential read performance. When you enable read reallocation, Data ONTAP analyzes the parts of the file that are read sequentially. If the associated blocks are not already largely contiguous, Data ONTAP updates the file layout by rewriting those blocks to another location on disk where they can be contiguous. The rewrite improves the file layout, thus improving the sequential read performance the next time that section of the file is read. REALLOC might result in a higher load on the storage system. In addition, it can result in more storage use if Snapshots are used. If you want to enable read reallocation but storage space is a concern, you can enable read reallocation on FlexVols using the `space_optimized` option. The `space_optimized` option conserves space but can slow read performance through the Snapshot copies, so if fast read performance through snapshot copies is a high priority to you, do not use it.

A read reallocation scan does not rearrange deduplicated blocks. For files to benefit from read reallocation, they should be stored on volumes that are not enabled for deduplication.

## **5.4.14 Deduplication with Quotas**

Quotas are not aware of deduplicated blocks. For example, a user with a quota limit of 10 GB cannot store more than 10 GB of data in a volume even if their data could potentially deduplicate down to be less than 10 GB at the block level. If using deduplication and quotas in the same volume, storage administrators can utilize the block-level savings by over committing total quotas within the volume or by thin provisioning the volume within the aggregate.

## **5.4.15 Deduplication with NDMP**

Transferring data from a deduplicated volume using NDMP is supported, but there is no space optimization when the data is written to tape because the dump to tape is a logical file-level operation rather than a physical block-level operation. Not having deduplicated blocks on the tapes could actually be considered an advantage in some cases- if the data on the tape was deduplicated, it would then contain a proprietary format that would require Data ONTAP for restores.

## 5.4.16 Deduplication with DataFabric Manager/Protection Manager

With Data ONTAP 7.3, SnapVault integration with deduplication replaces Snapshot copies. As a result, Protection Manager has to wait for deduplication to finish before renaming Snapshot copies. During this wait time, Protection Manager does not allow clients to list the Snapshot copies or restore from them. This can adversely affect recoverability.

## 5.5 Compatibility with non-native functions

This section will cover compatibility with other product applications.

### 5.5.1 Deduplication with VMWare

VMware Infrastructure (VI) and vSphere environments deduplicate extremely well. However, while working out the Virtual Machine Disk format (VMDK) and datastore layouts, keep the following points in mind:

- ▶ Deduplication is volume-specific. Try to group similar data within the same volume to achieve optimal results.
- ▶ Operating system VMDKs deduplicate extremely well because the binary files, patches, and drivers are highly redundant between virtual machines (VMs). Maximum savings is achieved by keeping these in the same volume.
- ▶ Be alert for performance impact as virtual machine density increases. CPU resources might well run out before capacity resources become an issue.
- ▶ Application binary VMDKs deduplicate to varying degrees, depending how much block-level similarity they have. Applications from the same vendor commonly have similar libraries installed and deduplicate somewhat successfully. Applications written by different vendors often do not deduplicate at all.
- ▶ Application data sets, when deduplicated, have varying levels of space savings and performance impact based on application and intended use. Careful consideration is necessary, just as with non-virtualized environments, before deciding to keep the application data in a deduplicated volume.
- ▶ Transient and temporary data, such as VMware ESX swap files, page files, and user and system temp directories, do not deduplicate well and potentially add significant performance pressure when deduplicated. Therefore, consider keeping this data on a separate VMDK and volume that are not deduplicated.

A major factor that affects deduplication efficiency is the amount of user or application data. New installations typically deduplicate extremely well because they do not contain a significant amount of unique data.

**Important:** With VI and vSphere, the need for proper partitioning and alignment of the VMDKs is extremely important (not just for deduplication). To help prevent the negative performance impact of LUN/VMDK misalignment, refer to *IBM System Storage N series with VMware ESX Server*, SG24-7636. Also note that the applications in which the performance is heavily affected by deduplication (when these applications are run without VI) are likely to suffer the same performance impact from deduplication when they are run with VI.

A deduplication and VMware solution with file-based NFS shares is easy and straightforward. Combining deduplication and VMware with block-based LUNs requires a bit more work. For more information about this topic, refer to 5.4.10, “Deduplication and LUNs” on page 51

## 5.5.2 Deduplication with Tivoli Storage Manager

If IBM Tivoli Storage Manager (ITSM) and Deduplication will be used together, consider the following points:

- ▶ Deduplication savings with ITSM will not be optimal due to the fact that ITSM does not block align data when it writes files out to volumes. The result is that there are fewer duplicate blocks available to deduplicate.
- ▶ ITSM compresses files backed up from clients to preserve bandwidth. Compressed data does not usually yield good savings when deduplicated.
- ▶ ITSM client-based encryption results in data with no duplicates. Encrypted data does not usually yield good savings when deduplicated.
- ▶ ITSMs progressive backup methodology backs up only new or changed files, which reduces the amount of duplicate data going to the storage system.

## 5.5.3 Deduplication with Backup Exec

Deduplication savings with Backup Exec will not be optimal because Backup Exec does not block align data when it writes files out to its volumes. The result is that there are fewer duplicate blocks available to deduplicate.

## 5.5.4 Deduplication with Lotus Domino

There have been unverified reports of degradation in read performance when deduplication is used with Lotus® Domino® on primary storage. Be sure to carefully monitor performance.

## 5.5.5 Deduplication with Microsoft Exchange

In certain Exchange environments, extents are enabled to improve the performance of database validation. Enabling extents does not rearrange blocks on disk that are shared between files by deduplication on deduplicated volumes. Enabling extents does not predictably optimize sequential data block layout when used on deduplicated volumes, so there is no reason to enable extents on deduplicated volumes.

## 5.6 Data Characteristics

Deduplication savings will be based on how many identical blocks are in a single Data ONTAP volume. Because data sets are so unique, there is wide variation in the yields even for a similar data type. Even if files are the same at the application or OS level, their underlying blocks might differ.




Figure 5-5 Amount of duplicate data recovered from an actual customer's production data

Be aware of any application-level data management such as containerization or compression as this could potentially reduce the efficacy of deduplication block consolidation. Compressed, encrypted, containerized, block mis-aligned, and audio/video data will likely have low returns with deduplication. Be sure to determine if potential deduplication savings will offset deduplication capacity overhead requirements at the volume and aggregate levels.

### Encrypted Data

Encryption removes data redundancy. As a result, encrypted data usually yields extremely low deduplication savings, if any. Because encryption can potentially be run at the share level, it is possible to create a flexible volume where only part of the data on the volume is encrypted and it is still possible to deduplicate the rest of the volume effectively.

### Volatile Data

Volumes containing highly volatile data are not good candidates for deduplication because of unpredictable capacity impact. Also, deduplication might run more frequently than desired if using the **auto** schedule.

### Mission Critical Data

Volumes containing mission critical data are generally not recommended for deduplication because of possible impact to I/O performance.

**Tip:** Remember that deduplication is volume-specific. Try to group similar data within the same volume to achieve optimal results.

## 5.7 Deduplication and storage capacity

This section will cover storage capacity metrics when deduplication is involved.

### 5.7.1 Metadata

Although deduplication can provide storage savings in many environments, a small amount of storage capacity overhead is associated with it. This fact should be considered when planning for deduplication. Metadata includes the fingerprint database, the change logs, and temporary data.

- ▶ **Fingerprint database:** A fingerprint record exists for every 4 KB data block. The fingerprint records for all of the data blocks in the volume are stored in the fingerprint database file.
- ▶ **Change logs:** The size of the deduplication change log files depends on the rate of change of the data and on how frequently deduplication is run.
- ▶ **Temporary files:** When deduplication is running, it creates several temporary files. These temporary metadata files are deleted when the deduplication process finishes running.

#### **If running Data ONTAP 7.2.x:**

All of the deduplication metadata files reside within the volume. This metadata is therefore captured and locked in the snapshot copies of the volume as well. Of the volume's total data capacity (calculated pre-deduplication), 6% should be allowed for metadata overhead.

#### **If running Data ONTAP 7.3.x:**

**Volume capacity:** Up to 2% of the total data capacity (calculated pre-deduplication) within the volume can be consumed by deduplication temp files. This data will be deleted each time the deduplication of the volume is completed. If snapshots are taken before the temporary metadata is deleted, the temporary metadata will be locked into a snapshot. This same consideration applies to the snapshots used in association with most of the native data replication features of N series.

**Aggregate capacity:** Up to 4% of the volume's total data capacity (calculated pre-deduplication) can be consumed at the aggregate level by the fingerprint database and change logs. Remember that this will present for each deduplicated volume in the aggregate, so plan capacity accordingly.

**Note:** The amount of space required for deduplication metadata is dependent on the amount of data being deduplicated within the volumes rather than the size of the volumes or the aggregates.

## 5.8 Deduplication and performance

Deduplication processes can run on up to eight volumes concurrently except models N5500 and N3400 that have a maximum of five. Although deduplication runs as a low priority background processes on the system, the number of concurrent deduplication processes running and the phase of each process can affect the performance of other native functions on the system.



## 5.8.1 Duration of the deduplication operation

The performance of the deduplication operation is affected by other operations on the storage system, which in turn impacts the deduplication processing time.

On an IBM System Storage N7900 storage system with no other load on the system, deduplication performances of up to 120 MBps (running a single deduplication session) have been observed. If multiple deduplication streams are running, this total bandwidth is divided evenly into the number of streams.

To get an idea of how long the deduplication process takes to complete, in our example the deduplication process is running on a flexible volume at 25 MBps. If 1 TB of new data has been added to the volume since the last deduplication update, the deduplication operation will take about 10 to 12 hours to complete.

**Note:** There are no configurable parameters that can tune the deduplication process. The priority of this background process in Data ONTAP is fixed.

## 5.8.2 The I/O performance of deduplicated volumes

This section discusses general information about the write and read performance of deduplicated volumes.

### Write performance to a deduplicated volume

The impact of deduplication on the write performance of a system is a function of the hardware platform that is being used, and the amount of CPU load that is placed on the system.

If the load on a system is low, that is, for systems in which the CPU utilization is around 50% or lower, the performance difference is a negligible when writing data to a deduplicated volume, and there is no noticeable impact on other applications running on the system. On heavily used systems, however, where the system is nearly saturated with the amount of load on it, the impact on write performance can be expected to be around 15% for most IBM System Storage N series storage systems. The performance impact is more noticeable on higher-end systems than on lower-end systems. On the N7900 system, this performance impact can be as much as 35%. Note that these numbers are for FC drives: if ATA drives are used in a system, the performance impact would be greater.

### Read performance from a deduplicated volume

When data is read from a deduplication-enabled volume, the impact on the read performance varies depending on the difference between the deduplicated block layout compared to the original block layout. There is minimal impact on random reads.

Because deduplication alters the data layout on the disk, it can affect the performance of sequential read applications such as dump source, qtree SnapMirror, SnapVault source, SnapVault restore, and other sequential read-heavy applications. This impact is more noticeable in Data ONTAP releases earlier than Data ONTAP 7.2.6 with data sets that contain blocks with repeating patterns (such as applications that pre initialize data blocks to a value of zero). Data ONTAP 7.2.6 and later have specific optimizations, referred to as intelligent cache, that improve the performance of these workloads to be closer to the performance of non-deduplicated data sets. This is useful in many scenarios, especially in virtualized environments. In addition, the Performance Acceleration Modules are also deduplication aware and use intelligent caching.

## 5.9 Deduplication scheduling

Deduplication can be enabled on any number of flexible volumes in a storage system, but no more than eight volumes can be deduplicated concurrently (five if using an N3400 or N5500). If more than eight are scheduled to run concurrently, they will queue and round-robin. The default schedule is everyday at midnight (sun-sat@0). This generally is not the best practice. The optimal scheduling policy will depend on Filer CPU utilization peaks, volume I/O peaks, data replication schedules, and the number of volumes with deduplication turned on.

Alternately, manual execution can be used for the highest degree of control.

Use the following guidelines if using scheduling:

- ▶ Try to schedule deduplication so that it has time to complete before any replication functions run against data in the selected volume, especially those leveraging snapshots.
- ▶ If possible, avoid scheduling deduplication to run at times when system CPU is over 50% unless the deduplicated volume is not sensitive to write performance impact.
- ▶ Try to schedule less than eight to run concurrently (or less than five if using an N3400 or N5500).

### The auto schedule

Using auto scheduling is often a better way to randomize the start time and reduce schedule maintenance.

The auto schedule causes deduplication to run on that flexible volume whenever there are 20% new fingerprints in the change log. The check of the change log is done in a background process and occurs every hour. Starting with Data ONTAP 7.3.1, the 20% threshold can be adjusted.

One consideration is that reaching the set threshold that initiates deduplication will more likely happen during a time of I/O activity rather than when the volume is idle.

## 5.10 Aggregate and volume considerations

Specifically with ONTAP 7.3 and later, aggregate free space must be considered before turning on deduplication for contained volumes. A small percentage of space for deduplication metadata will be required at the aggregate level for each deduplicated volume. For more information, see 5.7.1, “Metadata” on page 60.

Deduplication is specific to blocks within a volume. It will not deduplicate across multiple volumes. Deduplication can only be used on flexible volumes, rather than traditional volumes, qtrees, or aggregates. It can be enabled on any number of flexible volumes in a storage system, but only eight volumes can be deduplicated concurrently (five if using an N3400 or N5500).

### *Maximum volume sizes for deduplication*

These limits are for the total volume size. These sizes including Snap Reserve space and are not related to the amount of data in the volume. With ONTAP versions earlier than 7.3.1, volumes that exceeded the size limit at any point cannot be deduplicated, even if the size is reduced back below the limit.

Table 5-8 Volume size limits

Data ONTAP version	N3300	N3600 N5200	N3400	N5500	N5300 N6040	N5600 N6060	N6070 N7600 N7700	N7800 N7900
ONTAP 7.3.0 and earlier	0.5 TB	1 TB	Not suppor ted	2 TB	3 TB	6 TB	10 TB	16 TB
ONTAP 7.3.1 and later	1 TB	2 TB	3 TB (must run 7.3.2 or later)	3 TB	4 TB	16 TB	16 TB	16 TB

### **Maximum volume data amounts for deduplication**

The maximum amount of data that can be deduplicated is 16 TB. Adding 16 TB to the maximum volume size gives the total amount of deduplicated data that can be in a single volume. Because the amount of stored data in a deduplicated volume exceeds these limits, no new deduplication occurs but writes to the volume are not impacted.

Table 5-9 Volume data limit table

Data ONTAP version	N3300	N3600 N5200	N3400	N5500	N5300 N6040	N5600 N6060	N6070 N7600 N7700	N7800 N7900
ONTAP 7.3.0 and earlier	16.5 TB	17 TB	Not suppor ted	18 TB	19 TB	22 TB	26 TB	32 TB
ONTAP 7.3.1 and later	17 TB	18 TB	19 TB (must run 7.3.2 or later)	19 TB	20 TB	32 TB	32 TB	32 TB





## ProtecTIER planning

In this chapter, we discuss configuration and sizing considerations, and provide detailed planning information to prepare for a smooth implementation of the TS7600 Family in your environment. You will find the following topics in this chapter:

- ▶ Hardware planning
- ▶ Planning for deduplication
- ▶ Planning for Open systems with Virtual Tape Library (VTL)
- ▶ Planning for Open systems with OpenStorage (OST)
- ▶ Planning for installation

## 6.1 Hardware planning for the 3959-SM1, 3958-AP1, 3958-DD4, and 3958-DE2

In this chapter, we describe which options and features can be configured with the following storage systems:

- ▶ IBM System Storage TS7610 ProtecTIER Deduplication Appliance Express (3959-SM1)
- ▶ IBM System Storage TS7650 ProtecTIER Deduplication Appliance (3958-AP1)
- ▶ IBM System Storage TS7650G ProtecTIER Deduplication Gateway (3958-DD4)
- ▶ IBM System Storage TS7680 ProtecTIER Deduplication Gateway for System z (3958-DE2)

This section discusses the following topics:

- ▶ Specifications
- ▶ Hardware and software components
- ▶ Configuration options

### 6.1.1 Specifications of the IBM System Storage TS7600 series

The functionality of the 3959-SM1, 3958-AP1, 3958-DD4 and the 3958-DE2 are identical, but there are differences in performance and maximum capacity.

#### **TS7610 ProtecTIER Deduplication Appliance Express (3959-SM1)**

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

- ▶ 4 TB
- ▶ 5.4 TB

For the 3959-SM1, the repository and one virtual library are already configured, as are the file systems over the disk storage. This product is customer installable. After hardware installation and software configuration, it is ready to work on. See more details in 6.1.2, “Hardware and software components of the 3959-SM1” on page 68.

#### **TS7650 ProtecTIER Deduplication Appliance (3958-AP1)**

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

The disk controller for the 3958-AP1 is an IBM System Storage DS5020 Express and the disk expansion module is an IBM System Storage DS4000® EXP810 Storage Expansion Unit. The DS5020 storage is configured for optimal performance at manufacturing. Three levels of disk capacity can be ordered:

- ▶ 7 TB
- ▶ 18 TB
- ▶ 36 TB

For the 3958-AP1 the repository is already configured, as are the file systems over the disk storage. After the 3958-AP1 is installed by a IBM SSR and the ProtecTIER Manager is

installed on a workstation, you can start to work on the 3958-AP1. After your host is attached and zoned to the 3958-AP1, you can set up your application on your host.

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

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

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

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

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

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

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

### TS7680 ProtecTIER Deduplication Gateway for System z (3958-DE3)

The IBM System Storage TS7680 ProtecTIER Deduplication Gateway for System z is a 1-rack (19 inch) high availability solution. It comes with a redundant pair of Enterprise Tape Controllers for Mainframe Host (z/OS) attachment and two clustered ProtecTIER servers running the deduplication engine. The TS7680 comes without disk storage, allowing clients to use IBM or non-IBM disk storage as back-end disk storage.

### Comparison of TS7600 series specifications

Table 6-1 shows the specifications for a single server. In case of the DE2, it shows the specifications for two ProtecTIER servers.

Table 6-1 Differences between the 3959-SM1, 3958-AP1, 3958-DD4 and 3958-DE2

Components	3959-SM1	3958-AP1	3958-DD4	3958-DE2 <sup>a</sup>
Host type	Open	Open	Open	Mainframe
Number of processors cores	4	16	32	48
Memory (GB)	24	32	64	64
Number of virtual libraries	Up to 4	Up to 12	Up to 16	1
Number of virtual tape drives	Up to 64	Up to 256	Up to 256	256
Number of virtual cartridges	Up to 8.192	Up to 128.000	Up to 500.000	1.000.000
Supported Virtual Library Emulation(s)	<ul style="list-style-type: none"> <li>▶ ATL P3000</li> <li>▶ DTC VTF® 0100</li> <li>▶ IBM TS3500</li> <li>▶ IBM V-TS3500</li> </ul>	<ul style="list-style-type: none"> <li>▶ ATL P3000</li> <li>▶ DTC VTF 0100</li> <li>▶ IBM TS3500</li> <li>▶ IBM V-TS3500</li> </ul>	<ul style="list-style-type: none"> <li>▶ ATL P3000</li> <li>▶ DTC VTF 0100</li> <li>▶ IBM TS3500</li> <li>▶ IBM V-TS3500</li> </ul>	IBM C06-PT library

Components	3959-SM1	3958-AP1	3958-DD4	3958-DE2 <sup>a</sup>
Supported Virtual Drive Emulation(s)	<ul style="list-style-type: none"> <li>▶ QUANTUM DLT7000</li> <li>▶ IBM ULT3580-TD2</li> <li>▶ IBM ULT3580-TD3</li> </ul>	<ul style="list-style-type: none"> <li>▶ QUANTUM DLT7000</li> <li>▶ IBM ULT3580-TD2</li> <li>▶ IBM ULT3580-TD3</li> </ul>	<ul style="list-style-type: none"> <li>▶ QUANTUM DLT7000</li> <li>▶ IBM ULT3580-TD2</li> <li>▶ IBM ULT3580-TD3</li> </ul>	▶ IBM 3592-J1A
Number of supported disk capacity	Up to 5.4 TB	Up to 36 TB	Up to 1 PB	Up to 1 PB
IBM Path failover technology	Yes, <sup>b</sup> with IBM TS3500 virtual tape library definition and multipath driver	Yes, <sup>b</sup> with IBM TS3500 virtual tape library definition and multipath driver	Yes, <sup>b</sup> with IBM TS3500 virtual tape library definition and multipath driver	Yes
Two-node cluster configuration	No	Yes	Yes	Yes <sup>a</sup>
IP-based replication configuration	Yes	Yes	Yes	Yes
Disaster recovery failover	Yes	Yes	Yes	Yes
Flexible disk-based storage option	No	No	Yes	Yes
Sustained in-line throughput <sup>c</sup>	Up to 80 MBps	Up to 500 MBps	Up to 500 MBps per node	Average of 500 MBps or more
Data reduction <sup>c</sup>	Up to 25:1 or more	Up to 25:1 or more	Up to 25:1 or more	Average of 25:1 or more
Preinstalled disk storage	Yes	Yes	No	No
Server(s) come(s) in a rack	No	Yes	No	Yes <sup>a</sup>

a. Only available in dual node configuration

b. Not for the Quantum P3000 virtual library

c. Depending on the back-end disk storage and workload

## 6.1.2 Hardware and software components of the 3959-SM1

The IBM System Storage TS7610 ProtecTIER Deduplication Appliance Express solution, consists of IBM System Storage ProtecTIER Enterprise Edition software and the IBM System Storage TS7610 Deduplication Appliance Express (3959-SM1) hardware, is designed to address the disk-based data protection needs of Small Medium Enterprises (SMEs).

The TS7610 ProtecTIER deduplication Appliance Express is shipped with the repository already configured, as are the file systems over the disk storage. One virtual library is created by default. This product is installable by the customer, and after initial hardware setup, the ProtecTIER software must be configured through direct console access.

The IBM System Storage ProtecTIER Entry Edition Small (or Medium) Capacity version 2.2 software and Red Hat Enterprise Linux 5.4 Advanced Platform 64-bit are installed at manufacturing. The IBM System Storage TS7610 has native Reliability, Availability, and Serviceability (RAS) software. You do not need additional hardware like TS3000 console. It is



also installed at the manufacturing, but you have to perform a final setup to provide details for email and SNMP traps. For more details, see 9.1.1, “TS7610 SMB Appliance” on page 170.

The Entry Edition Small is the 4 TB version. It can be easily upgraded to 5.4 TB, which is the Entry Edition Medium. Because the appliance contains all disk drives, no physical changes are required. You will only need to update the software. The software media is shipped based on the upgrade order, which is Feature Code #9314. The upgrade is a customer driven procedure.

Figure 6-1 shows the Appliance Express.



Figure 6-1 TS7610 Appliance Express

### 6.1.3 Hardware and software components of the 3958-AP1

The IBM System Storage TS7650 ProtecTIER Deduplication Appliance solution, comprising IBM System Storage ProtecTIER Appliance Edition (AE) V2.5 software and the IBM System Storage TS7650 ProtecTIER Deduplication Appliance (3958 Model AP1) hardware, is designed to address the disk-based data reduction needs of enterprise and distributed data centers. ProtecTIER AE V2.5 is enhanced to include IP Replication, IBM i support, a low-cost no deduplication licensing option, and additional capacity features for appliances with more than 32 TB of available storage. For more information, see 9.1.2, “TS7650 Appliance” on page 170.

The Appliance version contains the following components:

- ▶ 3958-AP1 server(s)
- ▶ Storage Disk System
- ▶ TS3000 system
- ▶ KVM kit
- ▶ Power switch (if required)
- ▶ Network Ethernet switches
- ▶ Cabling, labelling

There are three capacity levels of TS7650 Appliances:

- ▶ 7 TB
- ▶ 18 TB
- ▶ 36 TB

Figure 6-2 on page 70 shows the Appliance.



Figure 6-2 TS7650 Appliance with Dual-nodes

#### 6.1.4 Hardware and software components of the 3958-DD4

The IBM System Storage TS7650G ProtecTIER Deduplication Gateway solution, which consists of IBM System Storage ProtecTIER Enterprise Edition V2.5 software and the IBM System Storage TS7650G Deduplication Gateway (3958-DD4) hardware, is designed to address the disk-based data protection needs of enterprise data centers.

The 3958-DD4 server is available in a single node configuration and in a two-node clustered configuration. The single node configuration consists of one IBM System Storage TS7650G ProtecTIER Deduplication Gateway, which is based on an IBM System x3850 X5 server. In this book we refer to this server as IBM machine type and model 3958-DD4 for the IBM System Storage TS7650G ProtecTIER Deduplication Gateway and model 3958-AP1 for the IBM System Storage TS7650 ProtecTIER Deduplication Appliance

The two-node clustered configuration includes the following hardware components:

- ▶ Two IBM 3958-DD4 Gateway servers
- ▶ Two network Ethernet switches
- ▶ One remote network power switch

To allow installation, service, and maintenance of the 3958-DD4 a console must be provided by the customer that can be directly attached to the 3958-DD4.

A TS3000 should be included in the order with the TS7650 or TS7650G. One TS3000 console can be used to monitor and manage several ProtecTIER servers. If a TS3000 console was not included in the order, we recommend that one be ordered. Existing TS3000 models 306 or earlier are not supported ProtecTIER servers.

If you need a TS3000 System Console, it can be ordered with the 3958-DD4 ProtecTIER Deduplication Gateway. They will be shipped together. The TS3000 System Console is a one-unit (1U) System x server that allows an IBM System Service Representative (SSR) to perform maintenance and, if enabled by the customer, the TSCC can remotely monitor the installation and automatically inform IBM of any hardware errors.

**Important:** The 3958-DD4 requires the purchase of one or more frames, one or more disk arrays, and expansions to be fully functional:

- ▶ A frame is a 19-inch rack supplied by the customer and used to house the Gateway servers and the TS3000 System Console. A second frame can be used to house disk arrays and expansions.
- ▶ The disk array is the term used in this document to refer to a disk storage subsystem.
- ▶ The expansion refers to a disk expansion attached to the disk array.

The supported disk arrays have the following characteristics:

- ▶ Support for the 3958-DD4 server operating system with the correct update level.
- ▶ Dual active-active controller for compatibility with the Linux Multipath software included in the 3958-DD4 server operating system to allow path failover.
- ▶ Fibre Channel or SATA disk systems.
- ▶ Support for the Back End Fibre Channel Host Bus Adapter (HBA) brand, model, and firmware level installed on the Gateway server. The back end HBAs are used to direct or SAN attach the 3958-DD4 server to the disk array. In case of SAN attachment, the disk array must also support the fabric switches used. In certain cases such as XIV®, direct attach is not supported.
- ▶ The array should not perform its own compression by default. ProtecTIER does not require additional compression to be effective. ProtecTIER performs compression, by default, after the deduplication process.

The front view of the 3958-DD4 server is shown in Figure 6-3.



Figure 6-3 Front view of the 3958-DD4 server

The rear view is shown in Figure 6-4.




Figure 6-4 Rear view of the 3958-DD4 server

## 6.1.5 Hardware and software components of the 3958-DE2

The IBM System Storage TS7680 ProtecTIER Deduplication Gateway for System z is a parallel development to the IBM System Storage TS7650G ProtecTIER Deduplication Gateway and the IBM System Storage TS7650 ProtecTIER Deduplication Appliance. The TS7650G and TS7650 are for Open Systems attachment only but the TS7680 is for System z attachment. All of these solutions have the same ProtecTIER server hardware built in.

The IBM System Storage ProtecTIER Enterprise Edition V1.2 software and Red Hat Enterprise Linux Server Release 5.4 Advanced Platform 64-bit come preloaded on the TS7680 ProtecTIER servers.

The TS7680 comes in a two-node ProtecTIER clustered configuration. These two nodes are based on an IBM System x3850 M2 Type 7145-PBR.

To provide the System z attachment, two additional Enterprise Tape Controllers, consists of IBM System p5® 520 (model 9131-52A), also known as 3592-C06 Enterprise Tape Controller or IBM System Storage IBM TS1120 Tape Controller, each configured with two POWER5+™ 64-bit processors, are part of the TS7680. The Enterprise Controller comes with a preloaded Licensed Internal Code (LIC) to provide the internal System z to ProtecTIER attachment, the library manager, and virtualization functions.

The IBM System Storage TS7680 ProtecTIER Deduplication Gateway for System z configuration includes the following hardware components:

- ▶ Two IBM System Storage ProtecTIER Deduplication Gateway servers (lower and upper ProtecTIER servers)
- ▶ Two IBM System Storage IBM TS1120 Tape Controllers (lower and upper Enterprise Tape Controller)
- ▶ Two Ethernet network switches
- ▶ Two Ethernet network routers
- ▶ One remote network power switch

To allow installation, service, and maintenance of the TS7680, you will need a console that can be directly attached to one of the ProtecTIER servers. With an IBM TS3000 System Console you can manage up to 24 TS7680 product frames.

If you need a TS3000 System Console, it can be ordered with the TS7680 and then shipped together. The TS3000 System Console is a one-unit (1U) System x server that allows an IBM System Service Representative (SSR) to perform maintenance and, if enabled by the client, the TS3000 System Console (TSSC) can remotely monitor the installation and automatically inform IBM of any hardware errors.

See Figure 6-5 on page 73 for a logical layout of the IBM System Storage TS7680 ProtecTIER Deduplication Gateway for System z.




Figure 6-5 TS7680 logical layout

## 6.2 Planning for deduplication

HyperFactor is the core technology of ProtecTIER and consists of a series of algorithms that factor, or deduplicate, data efficiently. In each new backup, HyperFactor finds the data in common with previous backups and this common data in the new backup is effectively filtered out and pointers are used to reference existing data in the repository (see 1.2, “Deduplication overview” on page 4 for more details about the deduplication process). In this way, the entire contents of the new backup are stored and only new data is required to be stored, that is, a small fraction of the entire amount of the new backup data. See Figure 6-6 on page 74 for an outline of this process.

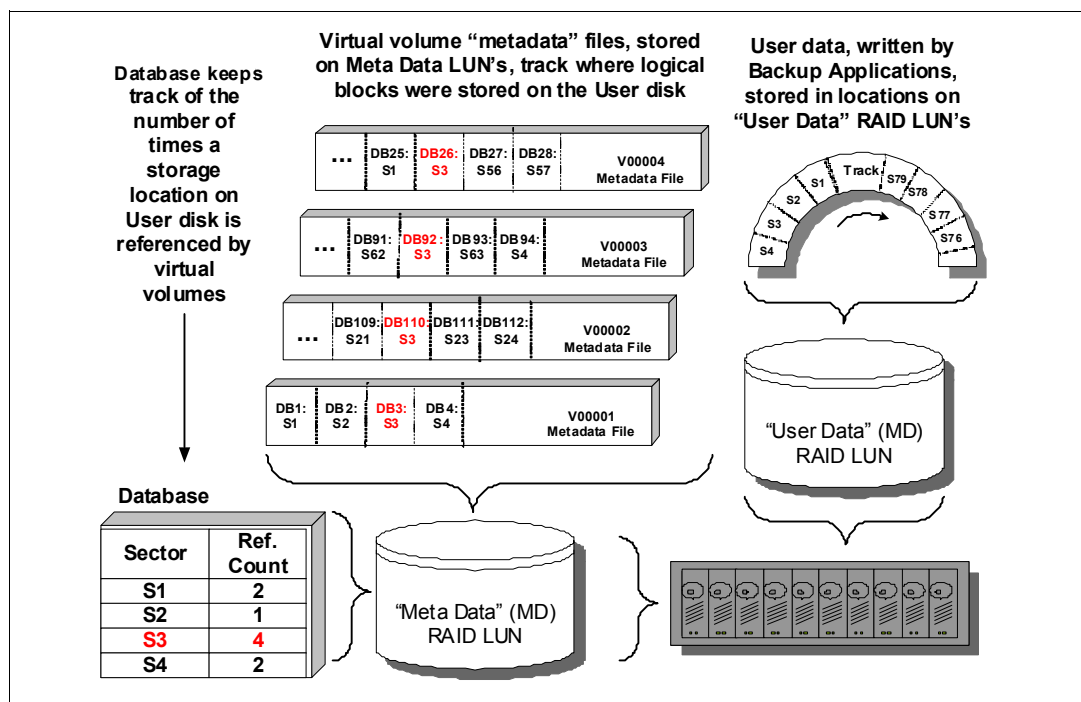


Figure 6-6 HyperFactor deduplication

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

Deduplication is most effective in environments in which there is a high degree of data redundancy. For example, a 10:1 deduplication ratio implies that the system is able to find redundancy in 90% of the data received.

Backup workloads are the best fit, as the backup process creates high levels of redundancy. Longer retention times increase the number of copies, and therefore the amount of data redundancy.

In Open Systems, 100% of ProtecTIER systems are deployed in backup environments. See details in 6.3, "Planning for Open systems with VTL" on page 74, and 6.4, "Planning for Open systems with OST" on page 100.

Mainframe workload profiles will vary significantly. Often, a repository will consist of a large variety of applications and data types, which has implications on planning and sizing the repository. See Appendix B, "ProtecTIER user cases" on page 287.

## 6.3 Planning for Open systems with VTL

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

► Sizing inputs

Section 6.3.1, “Sizing inputs” on page 77 describes a deep analysis of your requirements, workloads, backup application, and data center topology.

► Capacity sizing

In section 6.3.2, “Capacity sizing” on page 82, we provide an overview of the process that is used to estimate your physical storage requirements for your current environment and your scalability needs.

► Performance sizing

6.3.3, “Performance Sizing” on page 92 helps you to understand how many metadata and user data file systems are required, based on disk technologies and RAID configurations, to ensure proper performance in terms of backup/recovery throughput.

The methods described in this section provide the basis for assessing and understanding how the IBM System Storage TS7600 with ProtecTIER can be fully integrated in your environment. Figure 6-7 shows a flowchart of the sizing process. Although each step is described in the following sections, IBM technical personnel or Business Partner personnel will perform the sizing process before the pre-sales and pre-installation phases.

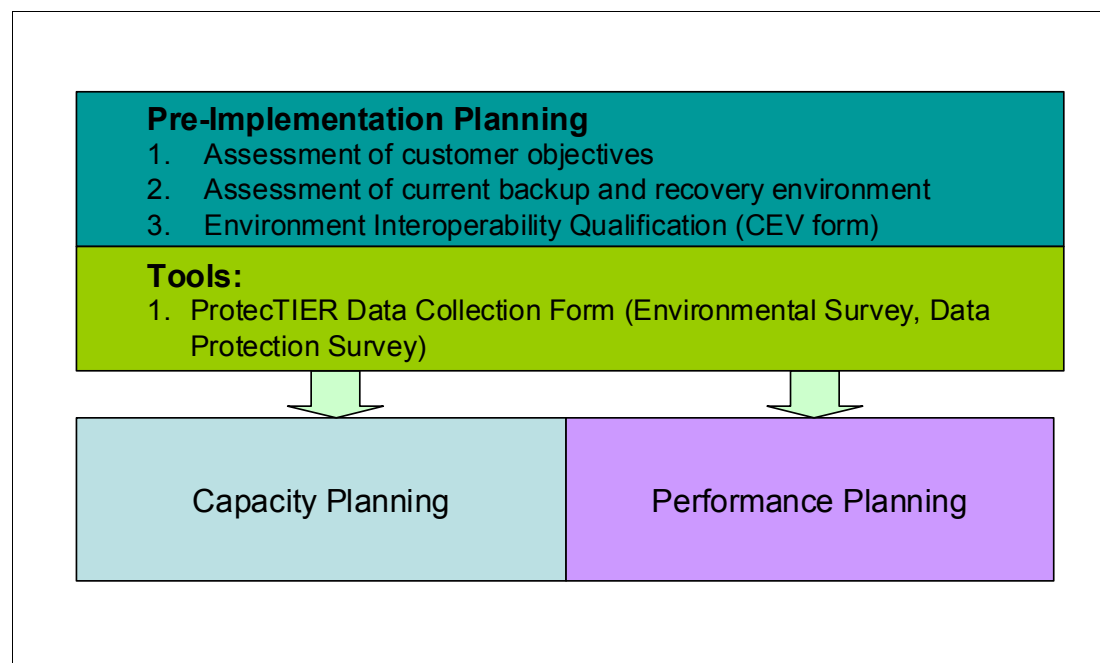


Figure 6-7 Key elements within the TS7650G implementation planning phase

ProtecTIER calculates the deduplication ratio by comparing the nominal data sent to the system to the physical capacity used to store the data. This information is displayed through the ProtecTIER Manager GUI and through other ProtecTIER utilities.

Table 6-2 on page 76 describes general Customer profiles for Open systems, giving a high level overview and suggestion for a possible ProtecTIER Product. Because Appliances have a predefined capacity and performance, you have to make sure that your needs are covered with them.



**Note:** For Appliances, upgrades can be done only within the same appliance family, either TS7610 or TS7650.

In case of TS7650G (and TS7680 for System z) you do not have predefined capacity, it has to be sized in advance. A maximum of 1 PB of physical storage can be attached to a TS7650G or TS7680 system. See more details about TS7680 in Appendix B, “ ProtecTIER user cases” on page 287.

Table 6-2 General Customer profiles for Open systems

ProtecTIER Solution	Physical Capacity	Performance	General Profile
TS7610 Small Appliance	4 TB	Up to 80 MBps	<ul style="list-style-type: none"><li>▶ 500 GB or less incremental backups per day</li><li>▶ 1-3 TBs full backups each week</li><li>▶ Experiencing solid data growth</li><li>▶ Looking to make backup and recovery improvements without making radical changes</li></ul>
TS7610 Medium Appliance	5.4 TB	Up to 80 MBps	
TS7650 Appliance	7 TB	Up to 150 MBps	<ul style="list-style-type: none"><li>▶ TB or less incremental backups per day</li><li>▶ 1-3 TB full backups each week</li><li>▶ Experiencing average data growth</li><li>▶ Need a cost effective solution</li></ul>
TS7650 Appliance	18 TB	Up to 250 MBps	<ul style="list-style-type: none"><li>▶ 3 TB or less incremental backups per day</li><li>▶ 3-6 TB full backups each week</li><li>▶ Experiencing rapid data growth</li><li>▶ Need good performance to meet backup window</li></ul>
TS7650 Appliance	36 TB	Up to 500 MBps	<ul style="list-style-type: none"><li>▶ 5 TB or less incremental backups per day</li><li>▶ 5-12 TB full backups each week</li><li>▶ Additional growth expected</li><li>▶ Higher performance is needed to meet the backup window.</li></ul>
TS7650G Gateway	Up to 1 PB	Up to 1200 MBps	<ul style="list-style-type: none"><li>▶ &gt;5 TB of backups per day</li><li>▶ Meeting backup window is an issue: need highest performance</li><li>▶ Experiencing rapid data growth: need scalability</li></ul>

Table 6-3 illustrates how physical space consumption is derived from the three general factors: backup size, retention, and data change rate. Table 6-3 is also giving a suggestion which model of appliance could be a good fit for the described cases.

Table 6-3 Suggested Appliance models for backup size, retention, and data change rate

Daily Backup (GB)	Retention (Days)	Data Change Rate	Required Physical Space (GB)	Dedup Ratio x : 1	Suggested Model
300	30	15%	1388	6.5	TS7610 4 TB
300	30	20%	1750	5.1	TS7610 4 TB
300	60	15%	2513	7.2	TS7610 4 TB
300	60	20%	3250	5.5	TS7610 4 TB



Daily Backup (GB)	Retention (Days)	Data Change Rate	Required Physical Space (GB)	Dedup Ratio x : 1	Suggested Model
300	90	15%	3638	7.4	TS7610 5.4 TB
300	90	20%	4750	5.7	TS7610 5.4 TB
500	30	15%	2313	6.5	TS7610 4 TB
500	30	20%	2917	5.1	TS7610 4 TB
500	60	15%	4188	7.2	TS7610 5.4 TB
500	60	20%	5417	5.5	TS7650 7 TB
500	90	15%	6063	7.4	TS7650 7 TB
500	90	20%	7917	5.7	TS7650 18 TB

### 6.3.1 Sizing inputs

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

#### Understanding the requirements

A key point to remember when you want to evaluate a virtual tape library (VTL) with deduplication and replication technology is to make sure that you have an understanding of your requirements and how you will be using the technology.

This analysis starts with the environment. Answer the following general questions about your requirements:

- ▶ Why do you want to introduce a VTL solution with data deduplication and replication in your environment? For example, are you going to use the solution to support disk-to-disk-to-tape (D2D2T) backups, to support disk-to-disk (D2D) with disk replication, or for archiving?
- ▶ What is the impact of a deduplication and replication solution on your current backup/recovery environment?
- ▶ Are you struggling with data proliferation?
- ▶ Are your applications good candidates for data deduplication?
- ▶ How many media servers do you need to attach and with how much storage capacity?
- ▶ What host interface technologies are you going to use (Fibre Channel, iSCSI, NAS, and so on)?
- ▶ Where are the potential bottlenecks in your current backup and recovery environment? How can a deduplication solution fix them?
- ▶ What are your current and projected performance requirements? Performance should include today's backup requirements and peak throughput needs, and should also factor in the data growth and the implications of future performance needs.
- ▶ What is your estimated annual data growth rate?
- ▶ What are your expected capacity savings with a deduplication solution?
- ▶ What are the possible changes that you plan to do in your current backup architecture? For example, you might have a current backup environment like this:
  - LAN-free backup to the physical tape library for databases
  - Disk-to-disk-to-tape (D2D2T) for file servers, web servers, mail servers, and application servers
  - A disaster recovery solution based on remote tape vaulting by a truck

You want to change your environment to this configuration:

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

## Understanding the existing environment

This section describes the information necessary to understand your current backup and recovery infrastructure. An environment questionnaire is used to collect this information and assess a pre-implementation plan. After it is completed, the questionnaire can be used to determine how ProtecTIER can fit into your environment. These questions will be asked by an IBM or Business Partner representative and will help evaluate three major areas:

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

## Environment configuration survey

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

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

## Operating systems and hardware inventory for compatibility check

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

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

## Data protection survey

Accurate capacity planning considers the behavior of each data type. A data type can be a file system, an operating system, databases, and so on. The size of one full backup is usually equal to the size of the online disk capacity. The Data Protection Survey is a worksheet that IBM technical personnel use during capacity sizing and that provides information about the number of versions, frequency, and retention for each backup. We assume that the retention of a weekly backup is associated with the retention of its incremental backup.

Important information for this survey includes:

- ▶ All the workloads that you back up in your environment.
- ▶ How much capacity is used for your current full backups to physical tape.
- ▶ How much capacity is used for the current daily backups to physical tape, including differentials, incrementals, and cumulative backups.
- ▶ The rate at which the data received from the backup application changes from backup to backup. This measurement has most relevance when like backup policies are compared. Data change rates can range from 1% to >25%, but are difficult to observe directly.
- ▶ How often the full backups are performed.
- ▶ How many cycles of full backups are kept.
- ▶ The relationship of how many daily backups are performed in between full backups, including differential, incremental, and cumulative backups.
- ▶ How many cycles of full and incremental, differential, or cumulative backups are kept.
- ▶ Whether a monthly full backup is kept for longer periods than the regular weekly full backup.

## Throughput considerations

The IBM System Storage TS7600 with ProtecTIER is a virtual tape library with enterprise scale in-band factoring, which means that all data reduction occurs in real time as the backups are running. With post process, in contrast, backups are first written to disk and then factored at a later point. The in-band factoring approach has many advantages, but also requires the appropriate level of hardware and proper configurations to achieve optimal performance. Properly configured, a single Enterprise-level IBM System Storage TS7650G with a ProtecTIER node is capable of achieving sustained throughput rates of up to 500 MBps in live production environments. Using a two-node clustered configuration, the IBM System Storage TS7650G with ProtecTIER can achieve sustained throughput rates of up to 900 MBps for backup and 1200 MBps for restore. The actual performance that any given environment achieves depends on several variables that we cover in this section.

The purpose of this section is to discuss performance considerations that can impact throughput performance, measured in megabytes per second (MBps), when testing and deploying the IBM System Storage TS7600 with ProtecTIER.

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

- ▶ SAN connectivity
- ▶ Disk array
- ▶ Data type (also called backup policy)

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

### ***SAN connectivity***

For the best SAN connectivity:

- ▶ Make sure that the fabric switches are up to the latest firmware revision of their operating system (contact manufacturer or reseller).

- ▶ IBM System Storage TS7600 with ProtecTIER front-end ports should not be in a zone with any other IBM System Storage TS7600 with ProtecTIER end ports.
- ▶ If possible, dedicated Host Bus Adapter (HBA) ports in the backup server should be zoned to a single IBM System Storage TS7600 with ProtecTIER and its front-end ports.
- ▶ Ensure that Inter Switch Links (ISL) between switches, connected to IBM System Storage TS7600 with ProtecTIER ports and backup servers or storage arrays, are not oversubscribed. ISL links for DS4000 and DS5000 are not recommended.
- ▶ Use at least 4 Gbps HBAs for the TS7650G back-end and front-end connections.
- ▶ If using SAN P2P topology to connect the TS7650G to the disk array, create dedicated zones (one zone per initiator) for ProtecTIER backend ports. For best results, each zone should have only one initiator and one target port, creating overlapping zones. Do not mix the ProtecTIER backend ports (Qlogic) with the front end ProtecTIER ports (Emulex) or any other SAN devices in the same zone.

### Disk array

A critical hardware component in a ProtecTIER implementation is the disk array that holds the ProtecTIER Repository. See Figure 6-8. The repository is the physical disk that holds the ProtecTIER HyperFactored data. There are two types of file systems that make up the ProtecTIER Repository:

- ▶ Metadata
- ▶ User data

See more details about setting up of disk array in 6.3.3, “Performance Sizing” on page 92.

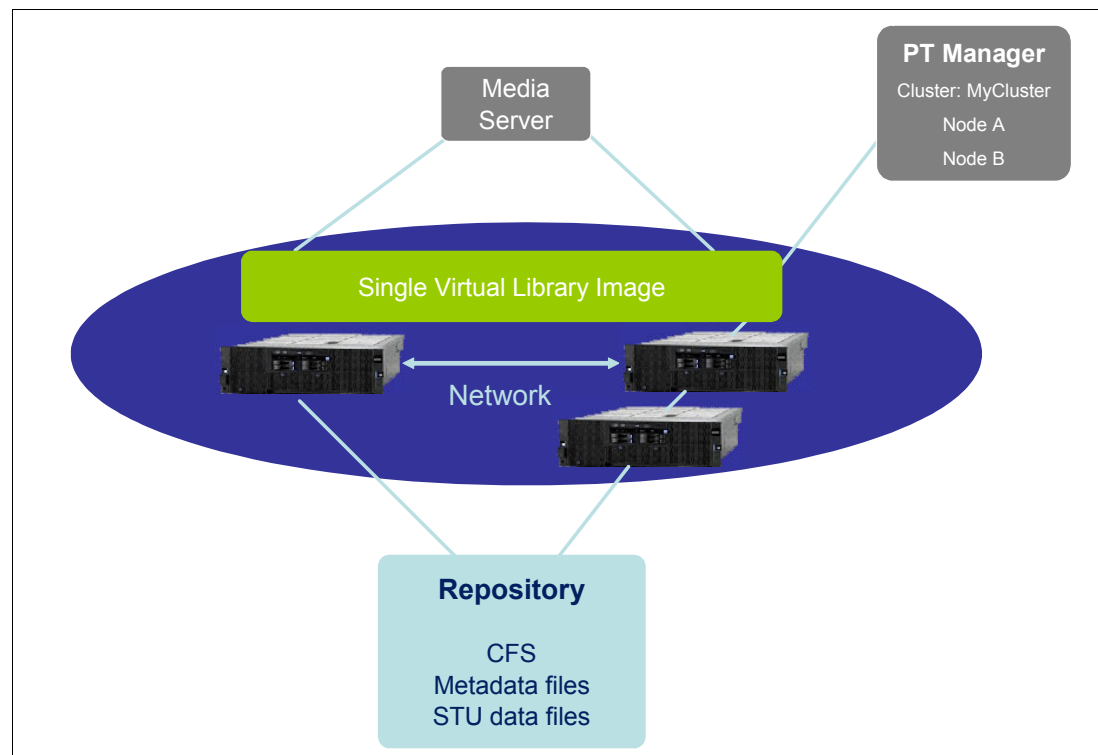


Figure 6-8 Repository of ProtecTIER: metadata and user data

Metadata file systems store all aspects of the data that is backed up and cataloged, but not the data itself, whether it requires new disk space or not. The user data file systems store the actual data that is backed up or referenced by new generations of the data. It is critical that the performance of the metadata file system be optimal. In general, we recommend RAID 10

RAID groups (4+4, 6+6, or 8+8 disks) for the metadata file systems. See details in “Metadata” on page 98.

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

Tuning the array for the unique ProtecTIER I/O pattern is critical. ProtecTIER is random read oriented. Eighty to ninety percent of I/O in a typical ProtecTIER system environment is random reads at 60KB block size. Therefore, any storage array deployed with ProtecTIER should be optimized for this I/O pattern. In all cases, the disk array manufacturer (or reseller) should be consulted to determine the best tuning and configuration parameters for the particular array being deployed. The user data LUNs should be tuned for a random-read intensive workload, and the metadata LUNs should be tuned for a random-write intensive workload. In 6.3.3, “Performance Sizing” on page 92 we describe implementation considerations related to the disk array configuration.

### Data type

The other factor that affects performance in a ProtecTIER system environment is the data that is being targeted for backup. Some data, such as databases and Lotus Notes® email, is highly compressible and also factors quite well. Other data, such as video or seismic data, cannot be compressed or factored well. Also, the various backup applications have features, such as encryption or multiplexing, that can also affect the ProtecTIER factoring ratio and performance. The type of data that ProtecTIER systems are HyperFactoring can affect both the factoring ratio and system performance. This section also includes general system testing and benchmarking considerations. See Table 6-4.

Table 6-4 Sample composition of Backup Policies, or Data types

Data Type	Example
Production Database	DB2®
Data Warehouse	Informix®
Email	Lotus Notes
File & Print Services	Windows® or Unix

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

Consider the following for the throughput:

- ▶ If multiple backup streams of the same data set are used for testing, we recommend that a single copy of the data be backed up first to populate the repository.
- ▶ If encryption features of the backup application are turned on, the factoring ratio and performance of these data sets will degrade drastically.

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

- ▶ Backup application or database backup programs should disable compression. Compression is common with SQL database backups using LiteSpeed. Data should be sent to the ProtecTIER servers uncompressed, or the factoring ratio for this data will be

low. ProtecTIER can manage multiple VTLs, each with its own configuration. For compressed data streams, create a new ProtecTIER VTL with Compression turned off. Compressing data a second time can cause data expansion, so compressed data should be segregated in ProtecTIER whenever possible.

**Note:** Compaction is a type of data compression.

- ▶ Multiplexing features of the backup application or database backup tool (RMAN or similar) should be disabled. The backup application should send only one stream to each virtual tape drive. Because ProtecTIER systems allow up to 256 virtual tape drives per node for the IBM System Storage TS7600 with ProtecTIER, the system can process many streams. For example, IBM System Storage TS7650 and TS7650G with ProtecTIER, RMAN sent nine streams (three file streams to three real tape drives each). With the IBM System Storage TS7650 and TS7650G, the parameters should be adjusted to send one file stream each to nine virtual tape drives. This does not adversely affect database backup speed, and might actually improve it.
- ▶ Small files do not factor as well as larger files. For best system performance, in test or in production, at least 24 data streams of backup data sets should be run at the same time. This takes advantage of the fact that ProtecTIER systems can process 24 storage units at the same time. The storage unit is one of the four allocation entities that ProtecTIER systems use for abstracting the physical disk to provide contiguous logical space when actual physical space is fragmented.

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

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

**Tip:** In general, the more virtual tape drives defined, the better. IBM System Storage TS7600 with ProtecTIER is optimized for a large number of virtual drives.

### 6.3.2 Capacity sizing

In this section, we document the general importance of capacity sizing, but this process requires a pre-sales engagement with IBM technical personnel. They use IBM internal tools to correctly size the IBM System Storage TS7600 with ProtecTIER through a comprehensive discussion with you about your data protection environment, requirements, and business objectives. See Figure 6-9 on page 84 for a diagram of the process. In the performance of the repository sizing, they will keep in mind the maximum throughputs of which the configurations are capable. The specifications are based on a realistic customer workload, assuming properly configured back-end disk arrays for the repository:

- ▶ VTL: Single node cluster: 900 MBps backup (up to 1200 MBps restore)
- ▶ VTL: Dual-node Cluster: 1200 MBps backup (up to 1600 MBps restore)
- ▶ DD4 OST Performance Specs: 600 MBps Single node; 1050 MBps cluster

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

- ▶ Example A (scheduled replication mode of operation)
  - There is backup activity running for 10 hours a day at a 500 MBps ingest rate.
  - The replication activity is running in a separate time slot of 12 hours a day at 500 MBps.

- The repository should support a sustained rate of 500 MBps.
- ▶ Example B (Replication activity runs in a continuous mode of operation.)
  - The backup activity is running 24x7 at an ingest rate of 400 MBps.
  - Replication runs concurrently and in parallel with the backup activity at 400 MBps.
  - The repository should support a sustained rate of 800 MBps.

Correct sizing of the ProtecTIER Repository provides the following benefits:

- ▶ Enables you to purchase the correct amount of storage
- ▶ Keeps backup operations running smoothly

Capacity sizing might rely on estimates and forecasts. Therefore, there is always a margin of error in the estimates and you should plan for additional capacity to compensate for this margin.

Table 6-5 shows the capacity sizing terminology that we use in our examples.

*Table 6-5 Capacity sizing terminology*

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

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

Figure 6-9 on page 84 shows the overview of inputs and outputs of the IBM Capacity Planner tool.

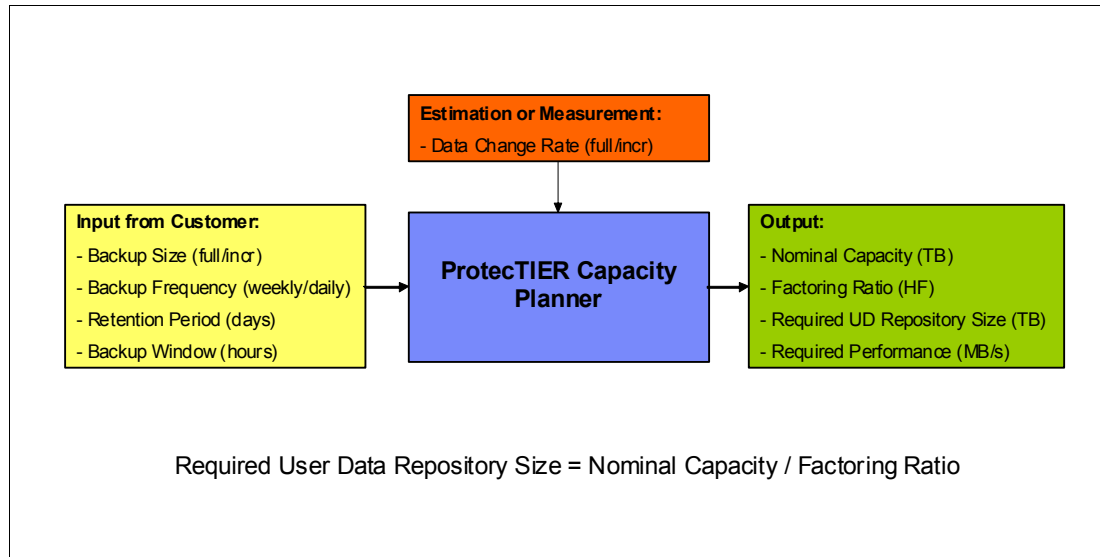


Figure 6-9 Capacity planning with IBM Capacity Planner tool

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

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

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

**Note:** Each policy can be said to have its own unique factoring ratio, and nominal and physical storage capacities.

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

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

- ▶ You add more backup policies to your environment.
- ▶ Your backup policies grow (corporate data growth).

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

**Note:** Adding 10% to the physical storage calculations is a good general rule.

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



## Factoring ratio considerations

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

- The data retention period: This is the period of time (usually measured in days) that defines how long you are going to keep your disk-based backups online. This period of time typically ranges from a period of 30 to 90 days, but can be much longer. This value is required when you compile the data protection survey, as discussed in “Data protection survey” on page 79.

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

- The data change rate: This is the rate at which the data received from the backup application changes from backup to backup. This measurement has the most relevance when like backup policies are compared. Data change rates can range from 1% to more than 25%, but are difficult to directly observe. Data change rate can be directly measured through an onsite assessment or can be estimated. Therefore, the more accurate the data change rate is, the more accurate the estimate will be about the sizing of the ProtecTIER Repository. Note that the factoring ratio is roughly the inverse of the data change rate. The data change rate is required when IBM conducts the data protection survey as part of the sizing process during an onsite, pre-sales engagement, as discussed in “Data protection survey” on page 79, and when you calculate the estimated factoring ratio, as described in “Factoring ratio considerations” on page 85.

**Note:** A small data change rate will yield a higher factoring ratio because when the level of changes are small, the same set of data will show up more often.

The following are other factors that influence the factoring ratio:

- Compression and encryption: Use of compression and encryption software prior to data being deduplicated will dramatically lessen the factoring ratio.
- Type of data: Typically data base files, operating system and application software packages, log files, email, user documents, and snapshots provide high factoring. For images, video, and seismic data the factoring ratio will be lower unless they are redundant.
- Backup application: Tivoli Storage Manager progressive incremental backups will lead to less data, thus less possibilities for deduplication. Full backups will lead to higher ratio.
- Initial factoring ratios will sometimes be lower than expected while the repository is being populated with daily/weekly backups. The expected factoring ratios will be achieved after a full 60-90 days of backups have been completed.

## Calculating deduplication factoring ratios

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

- Deduplication assumptions: Compression, full backup change rate, and incremental backup change rate. The change rate can be estimated or can be calculated, as we explain in the following method sections.
- All the data gathered in the data protection survey (see “Data protection survey” on page 79).

We describe the formula gathering the main factors to simplify and better understand the meaning of each factor. As described in “Factoring ratio considerations” on page 85, the

factoring ratio is the nominal capacity divided by the physical capacity. The formula for nominal capacity is:

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

Where:

<b>NominalCapacity</b>	This parameter represents the overall capacity stored in the repository during the retention period and is composed of all the full and incremental versions stored.
<b>FullCapacityVersions</b>	This parameter represents the overall full backup capacity (expressed in GB) stored during the retention period. In the following formula, you can see how the FullCapacityVersions depends on the FullCapacity, FullRetention, and FullFrequency parameters.
<b>FullCapacity</b>	This parameter represents the capacity (expressed in GB) stored during full backup.
<b>FullRetention</b>	This parameter represents the retention period (expressed in days) for the full backup jobs. For example, you might decide to retain your full jobs for 30 days.
<b>FullFrequency</b>	This parameter indicates how often you perform the full jobs during the retention period. For example, you might perform four versions in 30 days, that is, one full job a week, so this parameter must be set to a value of 7.

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

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

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

Where:

<b>IncrementalCapacityVersions</b>	This parameter represents the overall incremental backup capacity (expressed in GB) stored during the retention period. In the formula below, you can see how the IncrementalCapacityVersions depends on the IncrementalCapacity, IncrementalFrequency, IncrementalRetention, FullRetention, and FullFrequency parameters.
<b>IncrementalCapacity</b>	This parameter represents the capacity (expressed in GB) stored during incremental backup.
<b>IncrementalRetention</b>	This parameter represents the retention period (expressed in days) for the incremental backup jobs.
<b>IncrementalFrequency</b>	This parameter indicates how often you perform the incrementals during the retention period (this parameter must be set to the value 1 if you perform an incremental every day).
<b>FullRetention</b>	This parameter represents the retention period (expressed in days) for the full backup jobs (for example, you might decide to retain your full jobs for 30 days).

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

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

For the physical capacity, we have the following formula.

$$IncrementalCapacityVersions = IncrementalCapacity \times \left( \frac{IncrementalRetention}{IncrementalFrequency} - \frac{FullRetention}{FullFrequency} \right)$$

Where:

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

**FullPhysicalCapacity** This parameter indicates the full physical capacity (expressed in GB) effectively required in the repository. In the formula below, note that a first full backup must be entirely stored because no data is in the repository. Therefore, it is not possible to make an initial delta comparison.

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

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

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

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

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

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

$$IncrementalPhysicalCapacity = IncrementalCapacityVersions \times IncrementalChangeRate$$

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

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

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

This formula is quite complex, but it might give you an idea of the impact of the estimated data change rate on the estimated factoring ratio. Increasing the data change rate leads to a decreased factoring ratio. Also note how the compression rate is inversely proportional to the physical capacity. Another relationship involves the nominal capacity, the retention period, and the backup frequency. Increasing the retention period or decreasing the backup frequency leads to an increasing factoring ratio.

## Planning for cartridges

Prior to discussing the various elements of the TS7600 that monitor and manage capacity, we will provide a brief summary.

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

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

IBM System Storage TS7600 with ProtecTIER has internal mechanisms to manage the nominal capacity of the system and communicate with the backup administrator in the language of tape.

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

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

Over time, traditional tape libraries reach an equilibrium state. Tapes are recycled, which means that they are put back into the pool of cartridges available for new backups. In equilibrium, the number of cartridges that are returned to the available pool roughly equals the number required for the given day's payload. This makes capacity shortages easy to predict. Typically, if the number of new tapes used exceeds the number of tapes being returned to the pool, capacity shortages will happen.

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

### ***TS7600 with ProtecTIER versus traditional tape libraries***

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

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

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

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

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

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

### ***Managing capacity fluctuations***

As mentioned, cartridge capacity changes to reflect the system-wide shift in nominal capacity. In Table 6-6, the factoring ratio at equilibrium is greater than the factoring ratio that was used when the IBM System Storage TS7600 with ProtecTIER was first installed. There are 1000 tape cartridges, but because the factoring ratio has stabilized at a higher number (12:1 versus 10:1), the nominal capacity has increased from 100 TB to 120 TB. To accommodate the change, the capacity per cartridge has increased from 100 GB to 120 GB. The IBM System Storage TS7600 with ProtecTIER handles this through the learning algorithm mentioned earlier. The backup application still manages 1000 cartridges in its catalog, and because it will only change cartridges when an end of cartridge signal is sent by the IBM System Storage TS7600 with ProtecTIER, the increase in cartridge capacity is transparent to the backup application.

*Table 6-6 Effect of learning algorithm with higher than expected factoring ratio*

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

In Table 6-7, the factoring ratio at equilibrium is less than the factoring ratio that was used when the TS7650G or TS7650 was first installed. As you can see in Table 5-8, there are 1000 tape cartridges, but because the factoring ratio has stabilized to a lower value (8:1 versus 10:1), the nominal capacity has decreased from 100 TB to 80 TB. To accommodate the change, the capacity per cartridge has decreased from 100 GB to 80 GB.

*Table 6-7 Effect of learning algorithm with a lower than expected factoring ratio*

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

As cartridge size changes, the EW signal arrives sooner or later than originally. In the example shown in Table 6-6 on page 89, the EW for each cartridge occurs 20 GB later on day 30 than on day 1, allowing more data to fit on a given cartridge. In the example shown in Table 6-7, the EW for each cartridge occurs 20 GB earlier on day 30 than on day 1, allowing less data to fit on a given cartridge. As a result of the learning algorithm, more or fewer tapes will be consumed during a given day's workload.

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

### ***Capacity management implications: Initialization phase***

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

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

From an ongoing management perspective, the backup administrator must be aware that as cartridge size changes, the EW signal arrives sooner or later than originally. In the example in Table 6-6 on page 89, the EW for each cartridge occurs 20 GB later on day 30 than on day 1, allowing more data to fit on a given cartridge. As a result, more or fewer tapes will be consumed during a given day's workload. Therefore, backup administrators for IBM System Storage TS7650 and TS7650G with ProtecTIER must keep track of the number of cartridges used as a key indicator to indicate capacity fluctuations.

### ***Adding new data sets to an existing IBM System Storage TS7600 with ProtecTIER***

New data stream will have a high change rate because all of the data is new to the IBM System Storage TS7600 with ProtecTIER. This causes an increase in the system-wide change rate and a decrease in the nominal capacity because the factoring ratio is going to decrease. As the new data set runs through a full cycle, the nominal capacity might or might not return to what it was previously, depending on the data change rate of the new data set. The best way to add new data streams is to first sample the data to project the likely impact. In certain cases, this action might create a need for more physical disks not built in to the original TS7650G or TS7650 design.

### ***Space reclamation and steady state***

As mentioned previously, with real tape systems, cartridges expire and are returned to the available pool of cartridges to be used for new backups. This process is easy to manage and understand: When a cartridge is returned to the available pool, its full capacity is available for new backups. See 3.4.5, “Steady state” on page 29 for more details.

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

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

### ***Summary of TS7600 with ProtecTIER capacity management***

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

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

### ***ProtecTIER capacity management with the learning algorithm***

ProtecTIER cartridge capacity provisioned to the backup hosts is represented nominally. The total available nominal capacity within a ProtecTIER repository can be calculated as the repository's physical capacity multiplied by the projected factoring ratio. Projected factoring ratio is a weighted factoring ratio that takes into account the configured factoring ratio and actual factoring ratio relative to the physical repository consumption.

For example, a 10 TB repository with a factoring ratio of 10:1 yields 100 TB of nominal capacity.

- ▶ A 10 TB physical repository configured with factoring ratio of 10:1, and has 3 TB (30%) used at an actual ratio of 5:1 will have an 8.5:1 projected factoring ratio. We made the following calculation: (actual factoring ratio multiplied with the percentage of used space) plus (configured factoring ratio multiplied with the percentage of the unused space), which is  $([5 \times 0.3] + [10 \times (1 - 0.3)]) = 8.5$ .
- ▶ In a brand new system with 0% used space, the projected ratio will equal the configured ratio.

ProtecTIER regularly and continuously re-calculates its projected factoring ratio and adjusts the total nominal capacity value to maintain the most accurate representation of the nominal-to-physical capacity relationship.

### ***Recommendation to overcome of fluctuating cartridge size***

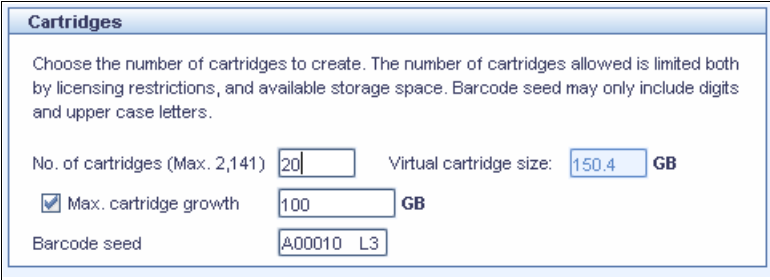
We see that one of most important part is to have the configured factoring ratio planned in advance in a manner that it will be close to the actual factoring ratio.

Because the first dataset will yield a lower factoring ratio than expected when the system is just populated with data, you will need to prepare for it.

An example:

- ▶ We have a new system with 10 TB of physical repository, with 10:1 HyperFactor ratio, leading to a 100 TB nominal capacity
- ▶ We create one library with 70 1-TB cartridges, which means 70 TB of the nominal capacity is assigned, 30 TB is reserved.
- ▶ The first full backup is run, yielding a low factoring ratio, which is normal because most of the data is new. The nominal capacity decreases to 80 TB, leaving us with 10 TB of reserve. Because we still have the reserve, the existing cartridges with 1 TB size will not change.
- ▶ After several backup generations, the factoring will increase, and our reserve will be grow, maybe back to 30 TB.
- ▶ After full backup retention/cycle is reached, we can add more cartridges to the library from the nominal reserve.

The reserve margin is established during the Cartridges window of the create-library wizard. You will see the same wizard when adding new cartridge adding to a library. This wizard gives you the option to set maximum cartridge growth as shown in Figure 6-10.



Cartridges

Choose the number of cartridges to create. The number of cartridges allowed is limited both by licensing restrictions, and available storage space. Barcode seed may only include digits and upper case letters.

No. of cartridges (Max. 2,141)  Virtual cartridge size:  GB

☒ Max. cartridge growth  GB

Barcode seed

Figure 6-10 Add cartridges wizard

The virtual cartridge size text box is dynamically calculated, depending on the number of cartridges you add, and it does not count any reserve. If you set the Max. cartridge growth to a lower value than the calculated value, you create reserve. In our example, it is 50 GB, or 50%. You can experiment with the numbers by adding more or less cartridges in the No. of cartridges field. The maximum cartridge growth means that the cartridge will report end of tape when reaching 100 GB. Otherwise your cartridge size would report end of tape around 150 GB, effectively creating a reserve for the deduplication ratio to change.

**Note:** If you create multiple libraries, you need to calculate the total nominal capacity and then compute the appropriate number of carts and max-growth ahead of time to maintain a reserve. The process only applies to a new system with a single library.

## **6.3.3 Performance Sizing**

It is important to configure a storage system in accordance with the needs of the user. An important question and primary concern for most users or storage administrators is how to configure the storage subsystem to achieve good performance. Those storage subsystems



capable of high IOPS, such as those with Fibre Channel drives, help deliver better TS7650 and TS7650G performance. The number of physical drives within a disk storage subsystem can also contribute to higher performance, as can the amount of disk storage controller memory and its fast write efficiency.

We also review other aspects of the system configuration that can help optimize the storage capacity and resilience of the system. In particular, we review and discuss the differences of SATA and FC technologies, RAID levels, array size, and array configuration for metadata and User Data.

## Selecting drives for TS7650

The speed and the type of drives used will impact the performance. Typically, the faster the drive, the higher the performance. This increase in performance comes at a cost. The faster drives typically have a higher cost than the lower performance drives.

FC drives outperform SATA drives. Table 6-8 compares the 10 K and 15 K Fibre Channel drives and SATA drives (single drive).

*Table 6-8 Comparison between Fibre Channel and SATA*

Factor	Fibre Channel	SATA	SATA difference
Spin speed	10.000 and 15.000	7.200	n/a
Command queuing	Yes 16 max	No 1 max	n/a
Single disk I/O rate (number of 512 byte IOPS) <sup>a</sup>	280 and 340	88	.31 and .25
Read bandwidth (MBps)	69 and 76	60	.96 and .78
Write bandwidth	68 and 71	30	.44

a. These IOPS and bandwidth figures are from disk manufacturer tests in ideal lab conditions. In practice, you will see lower numbers, but the ratio between SATA and FC disks still applies.

The speed of the drive is expressed in number of revolutions per minute (RPM). A 15 K drive rotates 15,000 times per minute. At higher speeds, the drives tend to be denser, as a large diameter plate driving at such speeds is likely to wobble. With the faster speeds comes the ability to have greater throughput.

Seek time is how long it takes for the drive head to move to the correct sectors on the drive to either read or write data. It is measured in thousandths of a second (milliseconds or ms). The faster the seek time, the quicker the data can be read from or written to the drive. The average seek time reduces when the speed of the drive increases. Typically, a 7.2 K drive will have an average seek time of around 9 ms, a 10 K drive will have an average seek time of around 5.5 ms, and a 15 K drive will have an average seek time of around 3.5 ms.

Command queuing allows for multiple commands to be outstanding to the disk drive at the same time. The drives have a queue where outstanding commands can be dynamically rescheduled or re-ordered, along with the necessary tracking mechanisms for outstanding and completed portions of workload. The SATA disks do not have command queuing and the Fibre Channel disks currently have a command queue depth of 16.

## RAID

We describe Redundant Array of Independent Drives (RAID), and their levels that you can build, and explain why we choose a particular level in a particular situation.

### RAID 0

RAID0 is also known as *data striping*. It is well suited for program libraries requiring rapid loading of large tables or, more generally, applications requiring fast access to read-only data or fast writing. RAID 0 is only designed to increase performance. There is no redundancy, so any disk failures require reloading from backups. Select RAID 0 for applications that would benefit from the increased performance capabilities of this RAID level. Never use this level for critical applications that require high availability.

### RAID 1

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

Here are general rules when using RAID 1:

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

**Note:** RAID 1 is actually implemented only as RAID 10 on DS4000 and DS5000 products.

### RAID 5

RAID 5 (Figure 6-11 on page 95) stripes data and parity across all drives in the array. RAID 5 offers both data protection and increased throughput. When you assign RAID 5 to an array, the capacity of the array is reduced by the capacity of one drive (for data-parity storage). RAID 5 gives you higher capacity than RAID 1, but RAID 1 offers better performance. RAID 5 is best used in environments requiring high availability and fewer writes than reads. RAID 5 is good for multi-user environments, such as database or file system storage, where typical I/O size is small and there is a high proportion of read activity. Applications with a low read percentage (write-intensive) do not perform as well on RAID 5 logical drives because of the way that a controller writes data and redundancy data to the drives in a RAID 5 array. If there is a low percentage of read activity relative to write activity, consider changing the RAID level of an array for faster performance. Use write caching on RAID 5 arrays because RAID 5 writes will not be completed until at least two reads and two writes have occurred. The response time of writes will be improved through the use of write cache, which should be backed up with a battery.

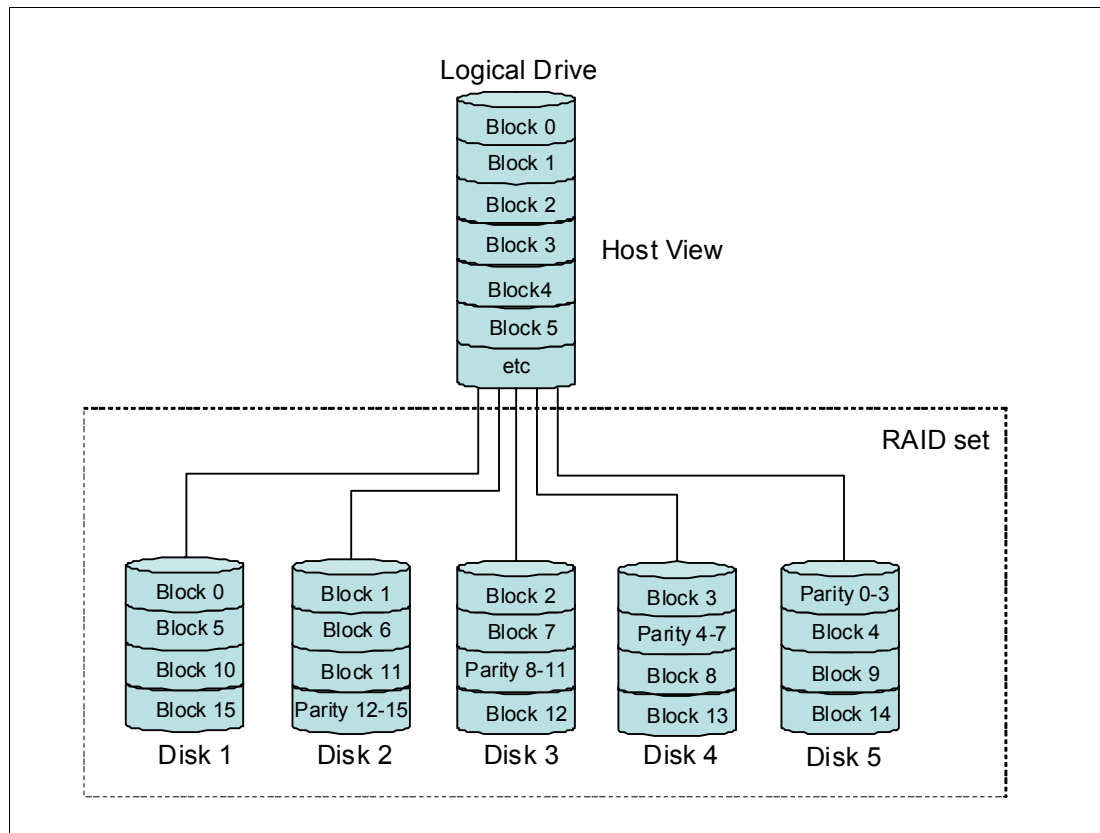


Figure 6-11 RAID 5

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

### **RAID 6**

RAID 6 (Figure 6-12 on page 96) provides a striped set with dual distributed parity and fault tolerance from two drive failures. The array continues to operate with up to two failed drives. This makes larger RAID groups more practical, especially for high-availability systems. This becomes increasingly important because large-capacity drives lengthen the time needed to recover from the failure of a single drive. Single parity RAID levels are vulnerable to data loss until the failed drive is rebuilt. The larger the drive, the longer the rebuild will take. Dual parity gives time to rebuild the array without the data being at risk if one drive, but no more, fails before the rebuild is complete. RAID 6 can be used in the same workloads in which RAID 5 excels.

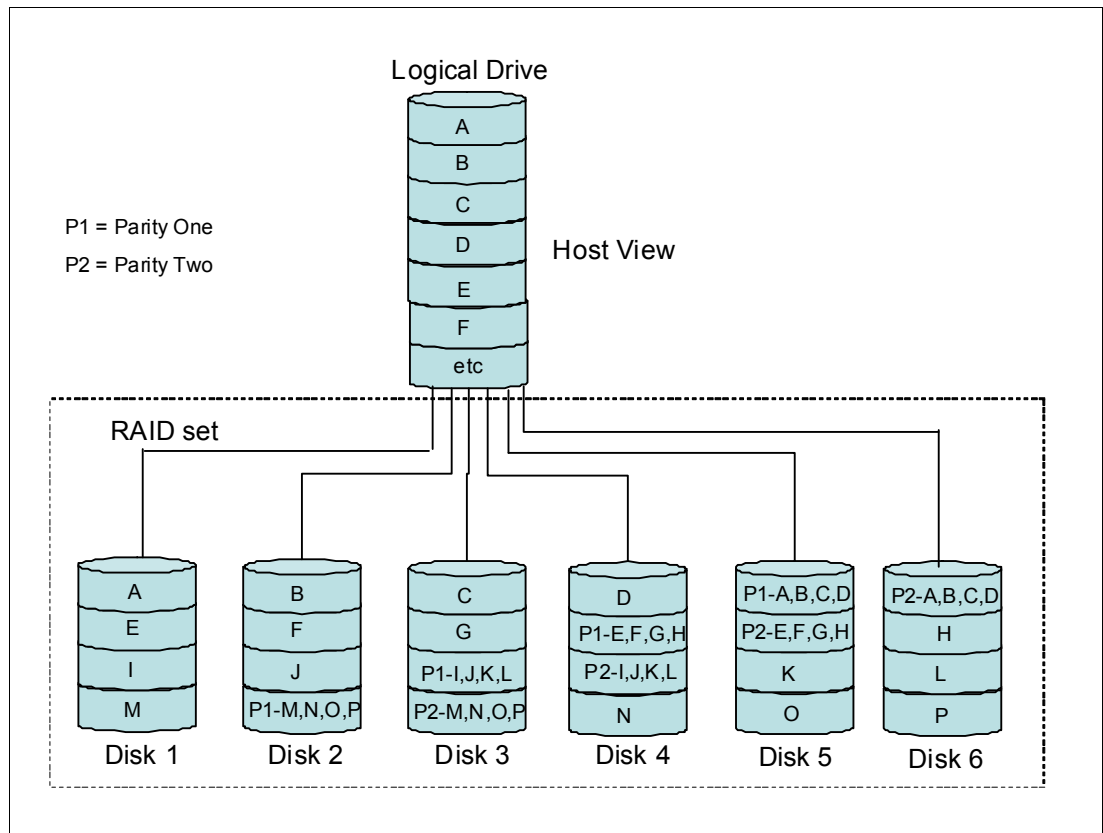


Figure 6-12 RAID 6

### RAID 10

RAID 10 (Figure 6-13 on page 97), also known as RAID 1+0, implements block interleave data striping and mirroring. In RAID 10, data is striped across multiple disk drives, and then those drives are mirrored to another set of drives. The performance of RAID 10 is approximately the same as RAID 0 for sequential I/Os. RAID 10 provides an enhanced feature for disk mirroring that stripes data and copies the data across all the drives of the array. The first stripe is the data stripe. The second stripe is the mirror (copy) of the first data stripe, but it is shifted over one drive. Because the data is mirrored, the capacity of the logical drive is 50% of the physical capacity of the hard disk drives in the array.




Figure 6-13 RAID 10

The suggestions for using RAID 10 are:

- ▶ Use RAID 10 whenever the array experiences more than 10% writes. RAID 5 does not perform as well as RAID 10 with a large number of writes.
- ▶ Use RAID 10 when performance is critical.
- ▶ Use write caching on RAID 10. Because a RAID 10 write will not be completed until both writes have been done, write performance can be improved through the use of a write cache. Be sure that the cache is backed up with a battery.

When comparing RAID 10 to RAID 5:

- ▶ RAID 10 writes a single block through two writes. RAID 5 requires two reads (read original data and parity) and two writes. Random writes are significantly faster on RAID 10.
- ▶ RAID 10 rebuilds take less time than RAID 5 rebuilds. If a real disk fails, RAID 10 rebuilds it by copying all the data on the mirrored disk to a spare. RAID 5 rebuilds a failed disk by merging the contents of the surviving disks in an array and writing the result to a spare. RAID 10 is the best fault-tolerant solution in terms of protection and performance, but it comes at a cost.

### Array size and array configuration

We advise to use only one type (FC or SATA), speed, and size of drive in any one array, and it is recommended that all the arrays in one system be identical as well.

It is critical to use RAID for data protection and performance. Use RAID5 with at least five disk members (4+1) per group should be used for Fibre-Channel User Data LUNs (4+P or 8+P for SATA disks), and RAID10 groups for metadata LUNs (with layout per planning requirements).

Even if SATA drives are used for User Data LUNs, use Fibre Channel disks for metadata LUNs.

**Note:** Do not share RAID groups or LUNs with other applications.

### ***Metadata***

We highly recommend that Fibre-channel disks should be used for metadata LUNs and RAID10 groups for metadata LUNs (with layout per planning requirements).

- ▶ **Metadata RAID Groups:** The recommended number of metadata RAID groups is determined by the performance planning tool with the help of IBM Representative. This number can range from 2 to 10 or more RAID groups (based on repository size, factoring ratio, and performance needs).

**Note:** Use only one type of RAID groups for metadata. For example: if one RAID group for metadata is 4+4, then do not create 2+2 for the second metadata RAID group.

- ▶ Only create one LUN per RAID Group: that is, one LUN that spans the entire RAID group. The only exception is for a single 1 GB metadata LUN. The 1 GB metadata LUN (Quorum or with other name Cluster filesystem) could be created on any of the metadata RAID groups.
- ▶ The size of the required metadata LUNs/file systems is a function of the nominal capacity of the repository (physical space and expected factoring ratio) and could be determined with the help of IBM Representatives using the metadata planner tool prior to the system installation.

### ***User data***

- ▶ **Fibre Channel User Data LUNs:** Use RAID5 with at least five disk members (4+1 per group).
- ▶ **SATA User Data LUNs:** Use RAID5 with at least five disk members (4+P or 8+P per group), or use RAID6.
- ▶ **User data RAID groups:** the IBM Performance Planner tool will provide the recommended amount of RAID groups for the given throughput requirements. At least 24 user data RAID groups should be created for optimal performance.
- ▶ Only create one LUN per RAID Group.
- ▶ The size of user data RAID groups/LUNs should be consistent. For example, do not mix 7+1 SATA User Data LUNs with 3+1 SATA LUNs. Smaller disk groups will hold back the performance of the larger groups and will degrade the overall system throughput.
- ▶ File system size should be maximum 8 TB. ProtecTIER will not allow to use any file system larger than 8 TB size.

**Note:** Keep in mind that hot spare disks should be separately assigned.

### ***Repository***

The repository will consists of the metadata LUNs and User Data LUNs. For an overview of the repository, see Figure 6-14 on page 99.

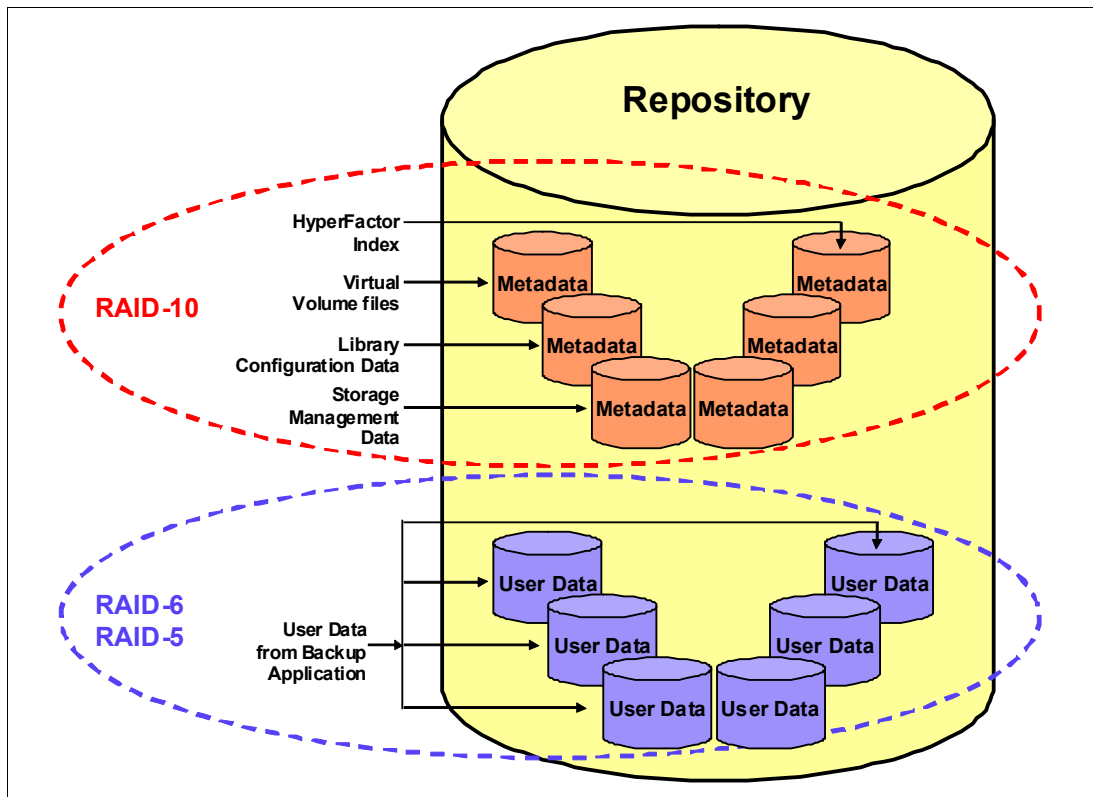


Figure 6-14 Repository

## IBM Performance Planner

Based on the data from Capacity Planning, and on the Disk technology and RAID grouping choices, the IBM Performance Planner tool provides the required amount of File Systems for metadata and user data. This planning is done by IBM Representative in the pre-sales period.

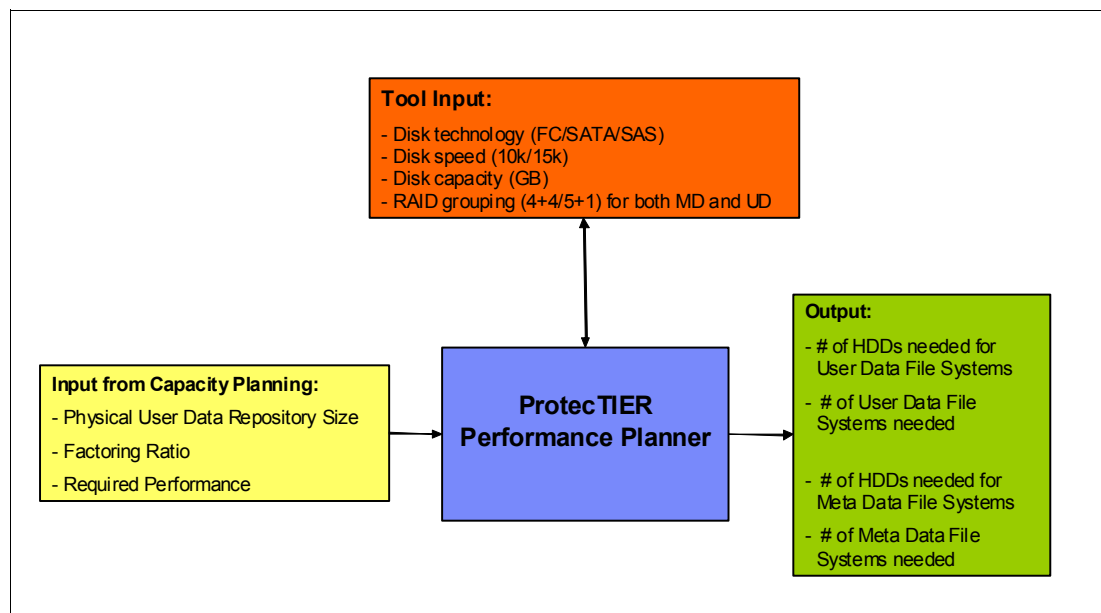


Figure 6-15 Performance planning with IBM Performance Planner tool

## 6.4 Planning for Open systems with OST

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

OST and VTL each have different topologies as follows:

- VTL: Pair (V2.3), hub and up to 12 spokes (ProtecTIER version 2.4)
- OST: Implicit hub mesh of up to 12 (ProtecTIER version 2.5).

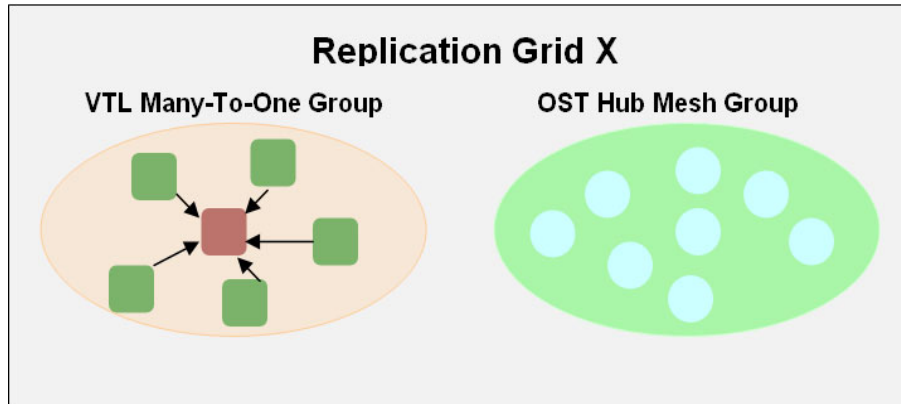


Figure 6-16 Replication topology: VTL and OST

With OpenStorage, ProtecTIER can be integrated with NetBackup to provide the means for backup-to-disk without having to emulate traditional tape libraries. Using a plug-in that is installed on an OST-enabled media server, ProtecTIER can implement a communication protocol that supports data transfer and control between the backup server and the ProtecTIER server. Therefore, to support the plug-in, ProtecTIER implements a storage server emulation as shown in Figure 6-17.

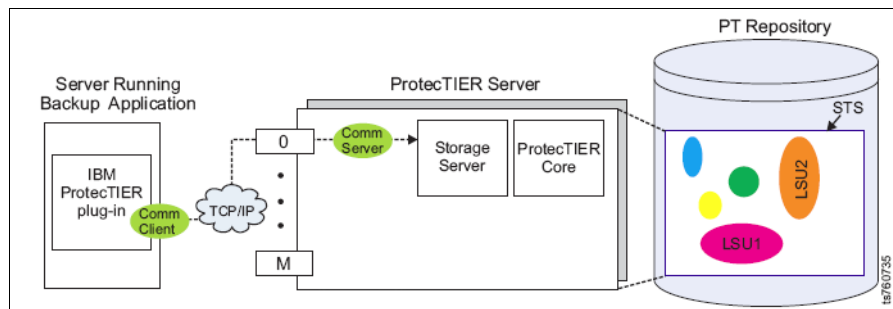


Figure 6-17 OST configuration map

OST has follow benefits:

- ▶ Treat disk as disk: This avoids the limitation of tape emulation and provides enhanced functionality only possible with disk.
- ▶ Tight integration: NetBackup is aware of all backup image copies and manage the creation, movement, and deletion of all backup images.



## 6.5 Planning for installation

In this section we present important details to help you plan for the installation and basic setup of the TS7600 Family. Planning is primarily a customer responsibility.

### 6.5.1 Supported backup server operating environments

The TS76xx can be used with a wide range of operating environments. For the most current list of supported products or for more information about support, refer to the following URL:

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

Currently, the TS76xx supports the following operating systems at the minimum levels indicated below.

For Open systems:

- ▶ AIX 5L™ V5.1, V5.2, V5.3, and V6.1 (TL4 SP2 or SP3)
- ▶ Sun Solaris 8, 9, and 10
- ▶ Microsoft® Windows 2003 and 2008 (32-bit only)
- ▶ HP-UX 11.x
- ▶ Red Hat ES3
- ▶ Red Hat ES4
- ▶ SUSE 10
- ▶ IBM i (V5R4)

And for System z:

- ▶ IBM z/OS V1R10
- ▶ IBM z/VM® 5.4

### 6.5.2 ProtecTIER manager workstation requirements

You must use the ProtecTIER Manager software to configure, manage, and monitor the operation of the TS76xx. You are responsible for obtaining the workstation where the ProtecTIER Manager software is installed. You can install the software from the supplied ProtecTIER Manager application CD.

The ProtecTIER Manager workstation must meet the following prerequisites to install and run ProtecTIER Manager effectively:

- ▶ One of the following operating systems:
  - Windows 32/64 bit (2003/XP)
  - Linux Red Hat 32/64 bit (Red Hat Enterprise 4 or 5)
- ▶ At least 1.2 GB of available disk space
- ▶ At least 256 MB of RAM
- ▶ The workstation can access the ProtecTIER service nodes' IP address (ports 3501 and 3503 are open on the firewall).

In addition, it is recommended that the monitor for ProtecTIER Manager be configured to the following settings:

- ▶ Resolution of 1024 x 768 pixels or higher (this is the minimum resolution supported, however, 1280 x 1024 is recommended).
- ▶ 24 bit color or higher

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



## ITSM planning

This chapter contains information that help you plan for a successful IBM Tivoli Storage Manager (ITSM) deduplication implementation.

In this chapter, the following topics are discussed:

- ▶ ITSM Planning overview
- ▶ ITSM Deduplication prerequisites
- ▶ Types of ITSM deduplication
- ▶ ITSM deduplication considerations
- ▶ When to use ITSM deduplication
- ▶ When not to use ITSM deduplication

## 7.1 ITSM Planning overview

Planning for data deduplication is important because there are many factors, such as data deduplication location and storage pool setup, to consider. A set of guidelines is provided to structure your planning activities.

As part of planning, you must make the following decisions:

- ▶ Determine which client nodes have data that you want to deduplicate.
- ▶ Determine whether you want to implement server-side data deduplication, client-side data deduplication, or a combination of both.
- ▶ If you choose client-side data deduplication, decide what, if any, security precautions to take.
- ▶ Decide whether you want to define a new storage pool exclusively for data deduplication or update an existing storage pool. The storage pool must be a sequential-access disk pool (FILE type device class). Data deduplication occurs at the storage-pool level, and all data within a storage pool, except encrypted data, is deduplicated.
- ▶ If you want to implement server-side data deduplication, decide how best to control the duplicate-identification processes:
  - Run duplicate-identification processes automatically all the time.
  - Start and stop duplicate-identification processes manually.
  - Start duplicate-identification processes automatically and then increase or decrease the number of processes depending on your server workload.

Whatever you decide, you can always change the settings later after the initial setup to meet the requirements of your operations. For more information, see 10.3, “Managing data deduplication” on page 265.

## 7.2 ITSM Deduplication pre-requisites

If you plan to take advantage of ITSM integrated deduplication function, you need to consider additional database and log space requirements.

Also, ensure that you have sufficient resources if you are using deduplication or expect a heavy client workload. Deduplication processing will use more CPU and disk I/O activity.

### 7.2.1 ITSM active log and archive log sizing

The effect of deduplication on the active and archive log needed for the server will vary depending upon the deduplication ratio of the data. In a scenario where there is a high percentage of data that can be deduplicated, you must have enough space on the active log and archive log to hold all of the ITSM server transactions, including the *IDENTIFY DUPLICATE* process.

Each extent identified by the *IDENTIFY* process uses approximately 1500 bytes of active log. As an example, if the number of extents that can be identified during the execution of the *IDENTIFY* process is 250,000, then the active log space needed for this operation is estimated to be:

$$(250,000 * 1500) / 1024 / 1024 = 358 \text{ MB.}$$

Now imagine we have a nightly backup load of 100,000 files per client and we have 300 clients. This represents 30,000,000 (100,000 \* 300) files for that nightly workload. If those 30,000,000 files represented 60,000,000 deduplicatable extents, the total archive log space required would be around 83.8 GB:

$$(60,000,000 * 1500) / 1024 / 1024 / 1024 = 83.8 \text{ GB.}$$

Another consideration of the active log impact from IDENTIFY processing is the size of the files. For example, if a client backs up a single object, perhaps an image backup of a file system, that is 800 GB in size. This can represent a very high number of extents because of the nature of this data; for this discussion, we assume that it is 1.2 million extents. These 1.2 million extents would represent a single transaction for an IDENTIFY process that requires an estimated 1.67 GB of active log space:

$$(1,200,000 * 1500) / 1024 / 1024 / 1024 = 1.67 \text{ GB.}$$

This 1.67 GB of active log space for that single file might be easily attainable in isolation.

If the deduplication enabled storage pool will have a mix of data (a lot of small files and a few large highly deduplicatable files), multiply the active log size estimate by two. For example, if the previous estimate recommended a 25 GB active log size, with deduplication of mixed (small and big) files the active log size becomes 50 GB and the archive log is then 150 GB (three times the size of the active log).

An easier approach to estimating the ITSM active log size is that if you deduplicate very large objects (for example, image backups), use an active log size that is 20% of the database size. If you have a 300 GB database, for example, you should use a 60 GB active log (300 \* 0.20) and a 180 GB archive log (60 \* 3).

## 7.2.2 ITSM database sizing

The size of the ITSM database depends on the number of client files to be stored and the method used by the server to manage them. Enabling data deduplication can significantly increase the size of your ITSM database.

If you can estimate the maximum number of files that might be in server storage at any time, you can estimate the database size from the following information:

- ▶ Each stored version of a file requires about 600 - 1000 bytes of database space.
- ▶ Each cached file, copy storage pool file and active-data pool file requires about an additional 100 - 200 bytes of database space.
- ▶ For deduplicated files, each extent needs 250 bytes of database space per storage pool.
- ▶ Overhead can require up to 50% in additional space.

For example, considering a backup policy that retains three versions of a file and we back up 500,000 client files for one single client node, we would have:

$$500,000 \times 3 = 1,500,000 \text{ files}$$

At 1000 bytes per file, the database space required for these files is:

$$(1,500,000 * 1000) / 1024 / 1024 / 1024 = 1.4 \text{ GB}$$

Assuming that one third of these files are deduplicated and that each deduplicated file has 10 extents:

$$((1,500,000 / 3) * 10 * 250 \text{ bytes}) / 1024 / 1024 / 1024 = 1.16 \text{ GB}$$

Therefore 2.56 GB of database space is required. Allow up to 50% additional space (or 1.28 GB) for overhead. The database should then have at least 3.84 GB for that client node.

If you cannot estimate the numbers of files that might be backed up, you can roughly estimate the database size as from 1% to 5% of the required server storage space. For example, if you need 100 GB of server storage, your database should be 1 - 5 GB.

**Note:** As shown, data deduplication can have a large impact in the ITSM database size. If you plan on using data deduplication, make sure you have enough database space left (consider additional 50% free database space for data deduplication).

## 7.3 Memory

'The minimum memory requirements to run production servers are, on 64-bit systems (which are recommended), 12 GB, or 16 GB if you use deduplication.

If you plan to run multiple instances, each instance requires the memory listed for one server. Multiply the memory for one server by the number of instances planned for the system.'

These are the minimums, For a highly used TSM 6 server with deduplication you might want to consider least 64 GB per instance.

## 7.4 Types of ITSM deduplication

Two types of data deduplication are available on IBM Tivoli Storage Manager: *client-side data deduplication* and *server-side data deduplication*.

Depending on the resources available in your environment, you can use server-side data deduplication, client-side data deduplication, or both.

### 7.4.1 Server-side deduplication

Server-side data deduplication is a data deduplication technique that is done exclusively by the server and it requires extra CPU and disk I/O resources.

Example 7-1 on page 109 shows how server-side deduplication works.

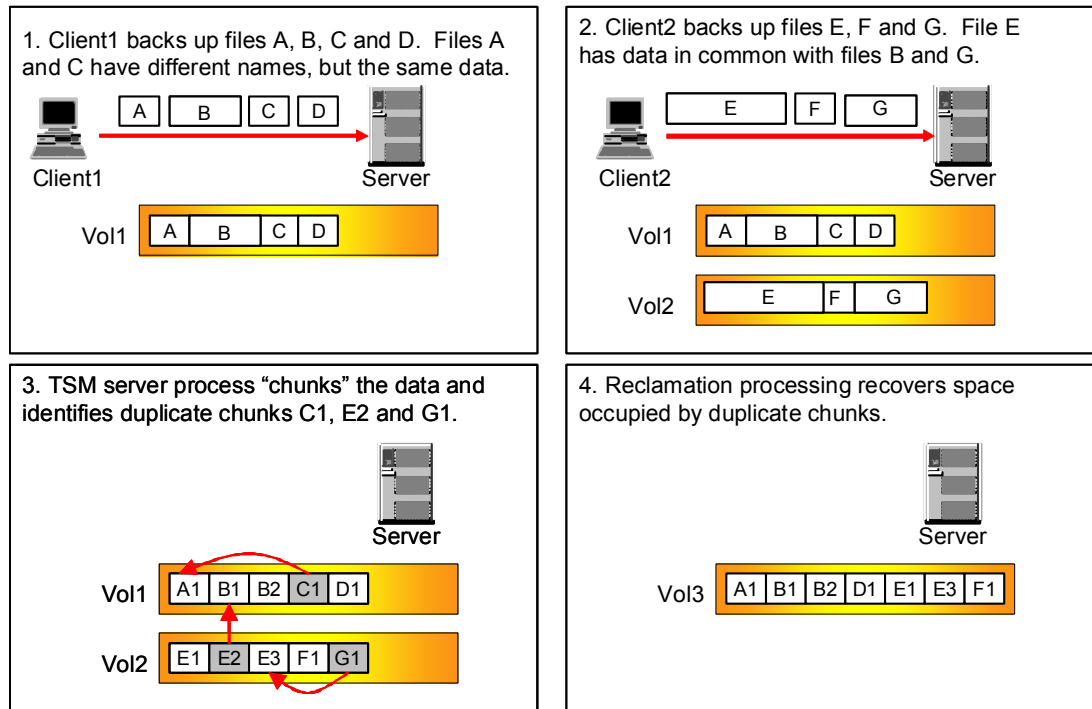


Figure 7-1 ITSM server-side data deduplication

Server-side data deduplication offers more protection against data loss. By default, primary sequential-access storage pools that are set up for data deduplication must be backed up to non-deduplicated copy storage pools before they can be reclaimed and before duplicate data can be removed. The default ensures that the server has copies of whole files at all times, in either a primary storage pool or a copy storage pool.

If your ITSM server has enough CPU and disk I/O resources available but your ITSM clients do not, you should consider using server-side deduplication.

## 7.4.2 Client-side deduplication

Client-side data deduplication is a data deduplication technique that is used on the backup-archive client to remove redundant data during backup and archive processing before the data is transferred to the IBM Tivoli Storage Manager server.

Figure 7-2 on page 108 shows how client-side deduplication works.

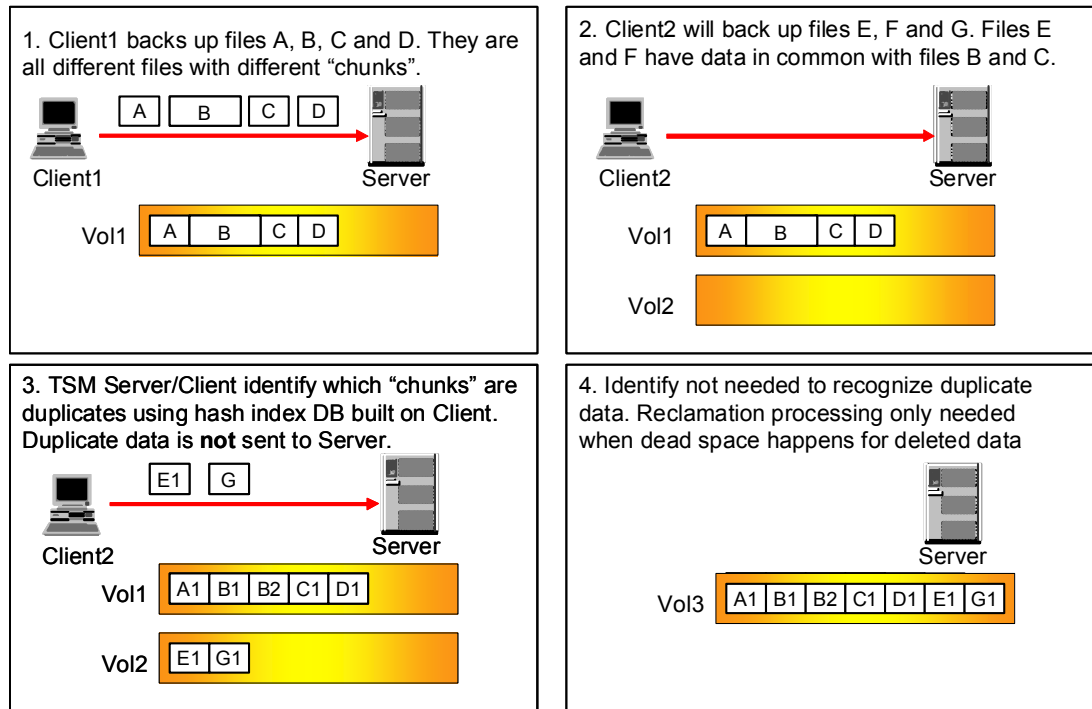


Figure 7-2 ITSM client-side data deduplication

Client-side data deduplication provides several advantages to server-side data deduplication:

- ▶ It reduces the amount of data sent over the network (LAN).
- ▶ The processing power that is required to identify duplicate data is offloaded from the server to client nodes. Server-side data deduplication is always enabled for deduplication-enabled storage pools. However, files that are in the deduplication-enabled storage pools and that were deduplicated by the client, do not require additional processing.
- ▶ The processing power that is required to remove duplicate data on the server is eliminated, allowing space savings on the server.

**Note:** For further data reduction, you can enable client-side data deduplication and compression together. Each extent is compressed before it is sent to the server. Compression saves space, but it increases the processing time on the client workstation.

For client-side data deduplication, the IBM Tivoli Storage Manager server must be Version 6.2 or higher.

If your network is highly used and you have enough CPU and I/O resources available on the ITSM client, you should consider using client-side deduplication for that client.

## 7.5 ITSM deduplication considerations

There are several considerations you need to know when planning for ITSM deduplication, especially if you are planning for client-side data deduplication.



## 7.5.1 Supported versions

Server-side data deduplication is available only with IBM Tivoli Storage Manager V6.1 or later servers. For optimal efficiency when using server-side data deduplication, upgrade to the backup-archive client V6.1 or later.

Client-side data deduplication is available only with IBM Tivoli Storage Manager V6.2 or later servers and ITSM backup-archive clients V6.2 or later.

## 7.5.2 Eligible storage pools

Data on DISK type device class or on tape cannot be deduplicated. Only data in storage pools that are associated with sequential-access disk (FILE type device class) can be deduplicated. You must enable FILE storage pools for data deduplication.

Client files must be bound to a management class that specifies a deduplication-enabled storage pool. For more information about backup policy and storage pool configuration, refer to the *Tivoli Storage Manager Administrator's Guide* documentation located at:

<http://publib.boulder.ibm.com/infocenter/tsminfo/v6r2/topic/com.ibm.itrm.ic.doc/welcome.html>

## 7.5.3 Encrypted data

The IBM Tivoli Storage Manager server and the backup-archive client cannot deduplicate encrypted files. If an encrypted file is encountered during data deduplication processing, the file is not deduplicated, and a message is logged.

**Note:** You do not have to process encrypted files separately from files that are eligible for client-side data deduplication. Both types of files can be processed in the same operation. However, they are sent to the server in separate transactions.

Data can still be encrypted on tape drives supporting those features as they currently are unaffected by deduplication.

If you need to secure your data, you can take one or more of the following steps:

- ▶ Enable storage-device encryption together with client-side data deduplication.
- ▶ Use client-side data deduplication only for node names that are secure.
- ▶ If you are uncertain about network security, enable Secure Sockets Layer (SSL).
- ▶ If you do not want certain objects (for example, image objects) to be processed by client-side data deduplication, you can exclude them on the client. If an object is excluded from client-side data deduplication and sent to a storage pool that is set up for data deduplication, the object is deduplicated on server.
- ▶ Use the **set dedupverificationlevel** command on the ITSM server to detect possible security attacks on the server during client-side data deduplication, as shown in Example 7-1. Using this command, you can specify a percentage of client extents for the server to verify (default value is 0). If the server detects a possible security attack, a message is displayed and client-side data deduplication is disabled.

*Example 7-1 Setting the dedupverificationlevel option*

---

```
tsm: TSMSRV>set dedupverificationlevel 3
ANR2668I The deduplication-verification level is set to 3.
```

## 7.5.4 Compressed data

For customers who use client-side compression (for example, due to limited network bandwidth), deduplication on the server will *not* be as effective as though the data was not compressed, but it will work with most compressed files.

Each extent is compressed on the client before being sent to the server, so if you compress and deduplicate data on the client system, you are using approximately twice as much processing power as data deduplication alone. Make sure your ITSM client has enough CPU and disk I/O resources available when using both compression and client-side deduplication.

## 7.5.5 Small files

IBM Tivoli Storage Manager does not deduplicate files smaller than 2 KB. They are backed up, but not deduplicated.

## 7.5.6 LAN-free considerations

Only V6.2 storage agents can use LAN-free data movement to access storage pools that contain data that was deduplicated by clients. V6.1 storage agents or earlier cannot use LAN-free data movement to access these storage pools.

Using a V6.1 storage agent or earlier to access a storage pool that contains client-side deduplicated data is possible, but it will cause the restore operations and retrieve operations to go over the LAN.

As part of the planning process, decide whether you want to use LAN-free data movement and whether you want to use client-side data deduplication, server-side deduplication, or both. If you decide to use LAN-free data movement and both client-side and server-side data deduplication, take one of the following steps:

- ▶ For V6.1 and earlier storage agents, store client-side deduplicated data in a *separate* storage pool. This way you can restore and retrieve deduplicated data from this storage pool over the LAN and use LAN-free data movement to restore and retrieve data from storage pools that contain data that was deduplicated *only* by the server.
- ▶ Upgrade to V6.2 storage agents. Upgrading to V6.2 storage agents provides LAN-free access to any storage pool that contains client-side deduplicated data, server-side deduplicated data, or both.

## 7.5.7 Hierarchical Storage Management (HSM)

HSM data from UNIX and Linux clients is ignored by client-side data deduplication. However, server-side deduplication of HSM data from UNIX and Linux clients is allowed.

## 7.5.8 Collocation

You can use collocation for storage pools that are set up for data deduplication. However, collocation might not have the same benefit as it does for storage pools not set up for data deduplication.

By using collocation with storage pools that are set up for data deduplication, you can control the placement of data on volumes. However, the physical location of duplicate data might be on separate volumes.

No-query-restore and other processes remain efficient in selecting volumes that contain non-deduplicated data. However, the efficiency declines when additional volumes are required to provide the duplicate data.

## 7.5.9 Disaster recovery considerations

There are a number of operational considerations when using deduplicated primary or copy storage pools. The main consideration is that deduplication is only allowed for storage pools using a device class of DEVTYPE=FILE. As such, deduplicated copy storage pools do not lend themselves to be used by DRM for the off-site protection of the IBM Tivoli Storage Manager server.

The following scenarios discuss typical and expected implementations:

- Primary pool deduplicated, single copy storage pool not deduplicated

This is probably the typical situation. In this case, the time to back up the primary storage pool to a non-deduplicated copy storage pool might take longer. The issue is that as data is copied to the copy storage pool, the deduplicated chunks representing a given file need to be read and the file “recomposed” or “unduplicated” and stored in its entirety in the copy storage pool. However, this allows for the use of the IBM Tivoli Storage Manager DRM feature and the use of off-site copy storage pool volumes to copy recovery data to a disaster recovery site.

- Primary pool deduplicated, multiple copy storage pools not deduplicated

This is also expected to be a typical situation or configuration. In this case, the time to back up the primary storage pool might also take longer because of the reads necessary to get the various chunks representing a given object being written to one of these copy storage pools. This is similar to the case above but also provides an on-site copy storage pool copy of the data in addition to the off-site copy.

- Primary pool deduplicated, single copy storage pool deduplicated

In this case, the primary goal is space savings at the primary server location. However, this does not provide for DRM exploitation and off-site rotation of copy storage pool volumes, because the deduplicated copy storage pool can only be achieved using DEVTYPE=FILE volumes and you cannot send a disk offsite. This approach has inherent risk associated with it because there is no off-site copy of the data that can be used in the event that the primary server location, hardware, or data center is damaged or destroyed.

**Important:** This approach does not provide an off-site copy of the data or the use of IBM Tivoli Storage Manager's DRM feature. If the primary product server or data center is damaged or destroyed, this might result in the loss of data or inability to recover that data.

- Primary pool deduplicated, single copy storage pool not deduplicated, single copy storage pool deduplicated

In this case, the deduplicated copy storage pool is providing space savings for the primary server location and duplication of the data that resides in the primary storage pool. Keep in mind that IBM Tivoli Storage Manager server deduplication is done at a single storage pool level so a deduplicated primary pool and a deduplicated copy storage pool will be referencing separate chunks and maintaining separate database references to track and

manage the data chunks that represent a given file. On the other hand, the non-deduplicated copy storage pool in this case, is likely real tape (something other than DEVTYPE=FILE) and is being used along with DRM for the purpose of having an off site data protection site. The off site copy storage pool volumes in conjunction with an appropriate database backup can then be used to restore the server and restore/retrieve data from the copy storage pool volumes.

In case you are using virtual volumes to store data on another ITSM server (a source ITSM server can send data to a remote ITSM server storage pool) the following scenarios apply:

- ▶ When you copy or move data to a deduplicated storage pool that uses virtual volumes, the data is deduplicated.
- ▶ When you copy or move data from a deduplicated storage pool to a non-deduplicated storage pool that uses virtual volumes, the data is reconstructed.

## 7.6 When to use ITSM deduplication

These are reasons to consider to enable IBM Tivoli Storage Manager deduplication:

- ▶ If you have a large preexisting disk storage pool and you want to deduplicate this already existing data
- ▶ If you have CPU cycles and disk I/O available. With IBM Tivoli Storage Manager's new expiration algorithm, your IBM Tivoli Storage Manager server might now have CPU/bandwidth available over IBM Tivoli Storage Manager V5.
- ▶ If you cannot consider alternatives for cost, support, power, or other reasons
- ▶ If your installation is too small to effectively use an alternative deduplication method
- ▶ If the data will reside in a disk storage pool for a long time.

## 7.7 When not to use ITSM deduplication

Here are reasons not to use ITSM deduplication:

- ▶ If your IBM Tivoli Storage Manager server is already CPU, memory, or disk IO constrained
- ▶ If you have too many IBM Tivoli Storage Manager Servers, or with large amounts of data that needs deduplication, look at platforms such as the IBM ProtecTIER product as best practice. Apart from scaling to large sizes (1 PB of non-deduplicated data), IBM ProtecTIER also opens the door to deduplication between separate IBM Tivoli Storage Manager Servers, inline deduplication, IP-based replication of stores, and a few other features.
- ▶ If you use the disk storage pool only as a temporary storage and will migrate all of the data to tape later.

For more information about IBM ProtecTIER deduplication, refer to Chapter 3, "Introduction to ProtecTIER deduplication" on page 19.



## Part 3

# Implementing Deduplication

In this part we discuss the implementation tasks required for deduplication.



# Implementing N series deduplication

This section goes through the entire typical process of creating a flexible volume and configuring, running, and monitoring deduplication on it using the command line and the IBM N series System Manager software.

## 8.1 Requirements Overview

The tables in this section identify the requirements for deduplication. Table 8-1 lists Data ONTAP versions for IBM System Storage N series models.

*Table 8-1 Supported IBM System Storage N series models*

<b>N series with NearStore option</b>	<b>Data ONTAP version</b>
N3300 A10 and N3300 A20	7.2.5.1 or later
N3400 A10 and N3600 A20	7.2.5.1 or later
N3600 A10 and N3600 A20	7.2.5.1 or later
N5000 Storage Systems	7.2.5.1 or later
N5000 Gateways	7.3
N6000 Gateways	7.3.1 or later
N6040 A10 and N6040 A20	7.3.1 or later
N6060 A20 and N6070 A20	7.3.1 or later
N7000 Gateways	7.3.1 or later
N7700 A20 and N7900 A20	7.3.1 or later

**Note:** The Gateway versions of the models require Data ONTAP 7.3.

Table 8-2 lists the software features needed.

Table 8-2 Software features needed

Requirements	Specification
Software licenses	NearStore option Deduplication(A-SIS bundled on ONTAP (see 5.3, “Deduplication licensing” on page 43)
Protocols	All
Maximum deduplication volume sizes for different Data ONTAP versions (see 5.10, “Aggregate and volume considerations” on page 62)	See <i>IBM System Storage N series Data ONTAP 7.3 Storage Management Guide</i> , GC52-1277

**Note:** A volume should never exceed the deduplicated volume size limit for the entire life of the volume. If a volume ever becomes larger than this limit, and is later shrunk to a smaller size, deduplication cannot be enabled on that volume.

## 8.2 The IBM N series System Manager

This section discusses N series System Manager including planning.

### 8.2.1 Overview of System Manager

IBM N series System Manager enables the setup and management of storage elements such as disks, volumes, and aggregates in a Microsoft Windows environment. You can use System Manager to quickly and efficiently set up IBM N series storage systems. You can use System Manager to configure protocols such as CIFS, NFS, and iSCSI; provision file sharing and applications; and monitor and manage your storage systems.

### 8.2.2 Bug Fix

System Manager 1.1R1 contains the following important bug fixes:

- ▶ During a LUN creation or the mapping change of an existing LUN, an existing initiator group (igroup) will be deleted if the same igroup name is specified in the Add Initiator Host window.
- ▶ System Manager allows lower or upper case iSCSI igroups names, which can conflict with Microsoft practices.

**Note:** System Manager 1.1R1 is an update to System Manager 1.1 to correct the issues stated above. All customers using System Manager versions 1.0, 1.0.1, or 1.1 should upgrade to System Manager 1.1R1.



## 8.2.3 Features

System Manager 1.1R1 includes the following features:

- ▶ Quota management: System Manager helps you to manage your quotas. You can create, edit, configure, and delete quotas using a wizard. You can also define quotas when you create a volume or qtree using the Quotas tab in the Volume Create and Qtree Create windows.
- ▶ System Manager provides complete support for N series storage system array LUNs. An array LUN is a group of disks or disk partitions in a span of storage space. Instead of direct-attached disks, you can view the array LUNs, create aggregates, and add array LUNs to an existing aggregate.
- ▶ Large aggregate support: When configured with Data ONTAP 8.0, System Manager supports aggregates larger than 16 TB.
- ▶ NFS protocol configuration: System Manager supports NFS protocol configuration using the System Manager GUI.
- ▶ Seamless Windows integration: System Manager integrates seamlessly into your management environment by using the Microsoft Management Console (MMC).
- ▶ Discovery and setup of storage systems: System Manager enables you to quickly discover a storage system or an active/active configuration (HA pair) on a network subnet. You can easily set up a new system and configure it for storage.
- ▶ SAN provisioning: System Manager provides a workflow for LUN provisioning, and simple aggregate and FlexVol creation.
- ▶ Network-attached storage (NAS) provisioning: System Manager provides a unified workflow for CIFS and NFS provisioning, and management of shares and exports.
- ▶ Management of storage systems: System Manager provides ongoing management of your storage system or active/active configuration (HA pair).
- ▶ Streamlined active/active configuration (HA pair) management: System Manager provides a combined setup for active/active configuration (HA pair) of IBM N series storage systems, logical grouping and management of such a configuration in the console or navigation tree, and common configuration changes for both systems in an active/active configuration (HA pair).
- ▶ Systray (Windows notification area): System Manager provides real-time monitoring and notification of key health-related events for a IBM N series storage system.
- ▶ iSCSI and FC: System Manager manages iSCSI and FC protocol services for exporting data to host systems.

## 8.2.4 System Requirements

- ▶ Windows XP Pro SP2 and later
- ▶ Windows Vista Enterprise SP1
- ▶ Windows Server 2003 SP2
- ▶ Windows Server 2008
- ▶ Microsoft Management Console (MMC) 3.0
- ▶ Microsoft Net 2.0
- ▶ IBM N series storage systems running Data ONTAP 7.2.3 or later

## 8.2.5 Installing the IBM N series System Manager Software

After you confirm that your server or environment meets the minimum requirements, download the System Manager from the IBM Storage site at the following URL:

<http://www.ibm.com/storage/support/nas>

Before you begin, close any earlier versions of system manager that is running on your system. You do not need to uninstall System Manager 1.0 or the System Manager 1.0.1 to upgrade to System Manager 1.1. Perform the following steps to install the IBM N series System Manager Software:

1. Run the executable (.exe) file (system-manager-setup-10R1-ibm.exe) and follow the instructions on the window.
2. Type or select information as required by the IBM N series System Manager setup wizard shown in Figure 8-1 and click the **Next** button.



Figure 8-1 IBM N series Software Manager setup wizard

3. Figure 8-2 on page 119 shows the license agreement of the IBM N series System Manager. Select the radio button to agree and click **Next**.

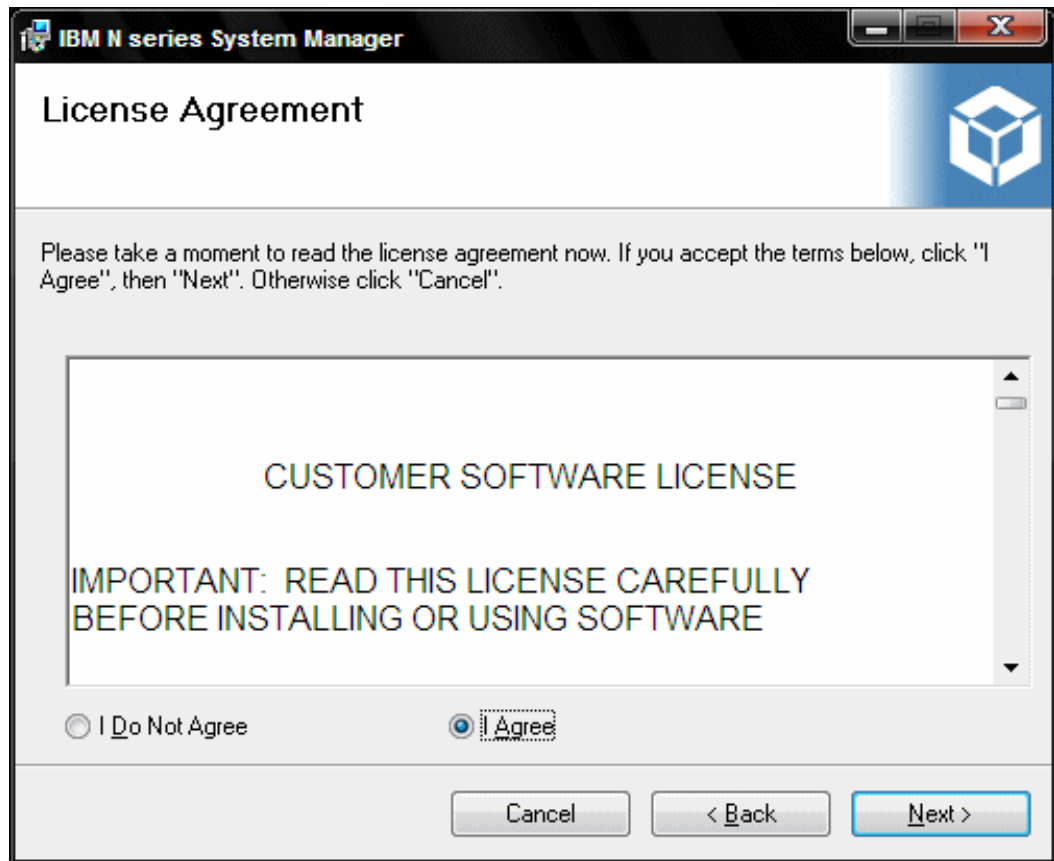


Figure 8-2 License Agreement

4. Select which folder the System Manager will be installed as shown in Figure 8-3 on page 120. Select **Everyone** if you want to allow all users to manage all IBM N series storage system. Otherwise, select **Just me**. Then click the **Next** button.

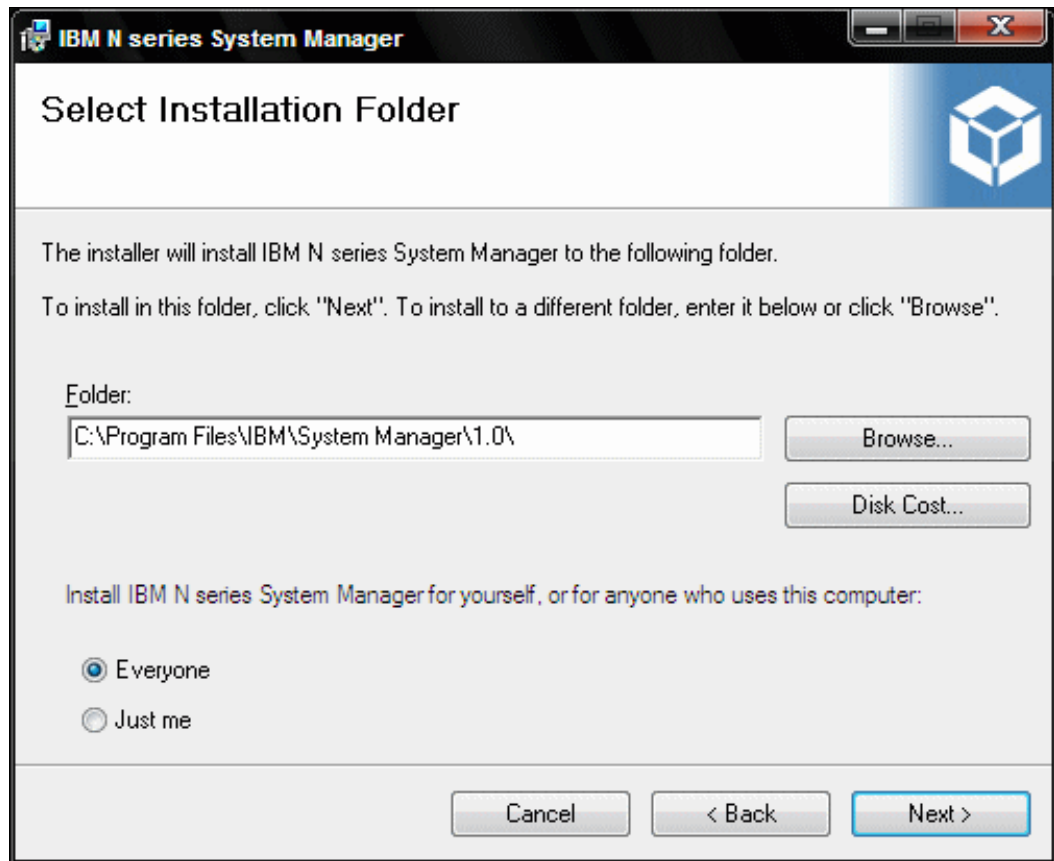


Figure 8-3 Select installation folder

5. Figure 8-4 on page 121 shows that the installation is ready to start. Press **Next** to start the installation.

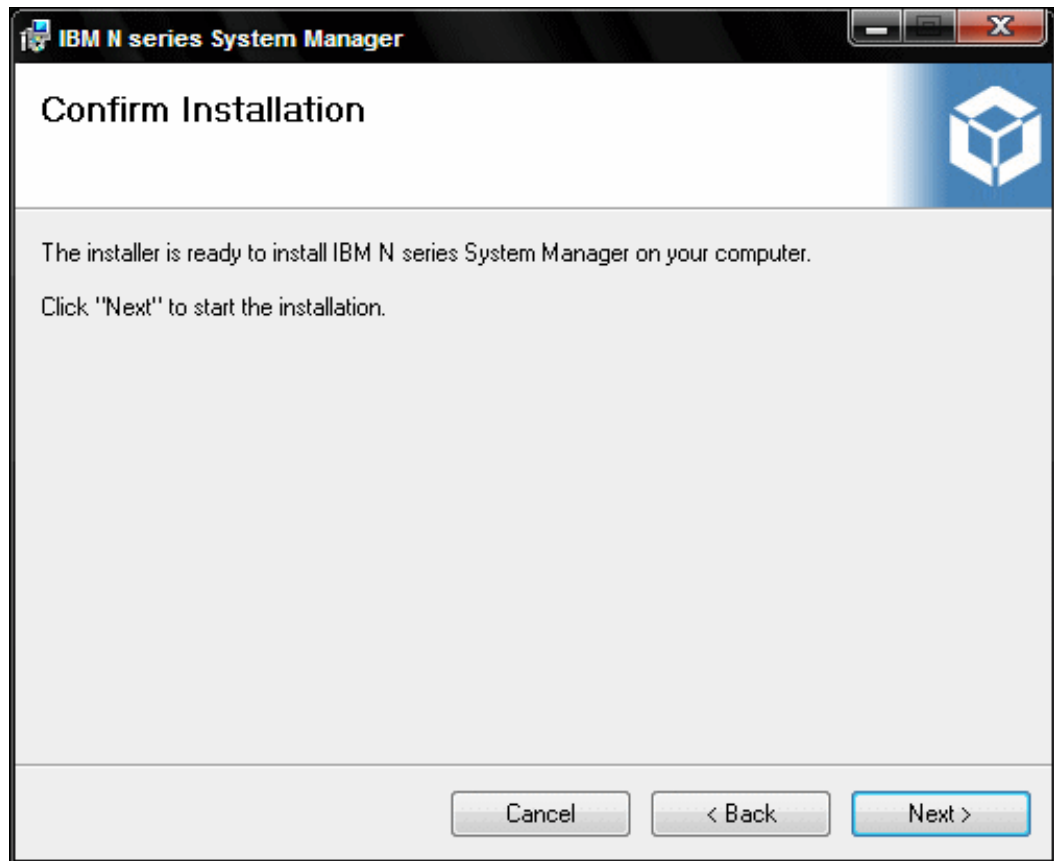



Figure 8-4 Confirm Installation

6. Figure 8-5 on page 122 shows that IBM N series System Manager Software is now being installed.



*Figure 8-5 IBM N series System Manager is now being installed*

7. Figure 8-6 on page 123 shows that IBM N series System Manager software has been installed. Click the **Close** button to finish the Setup Wizard and launch System Manager.

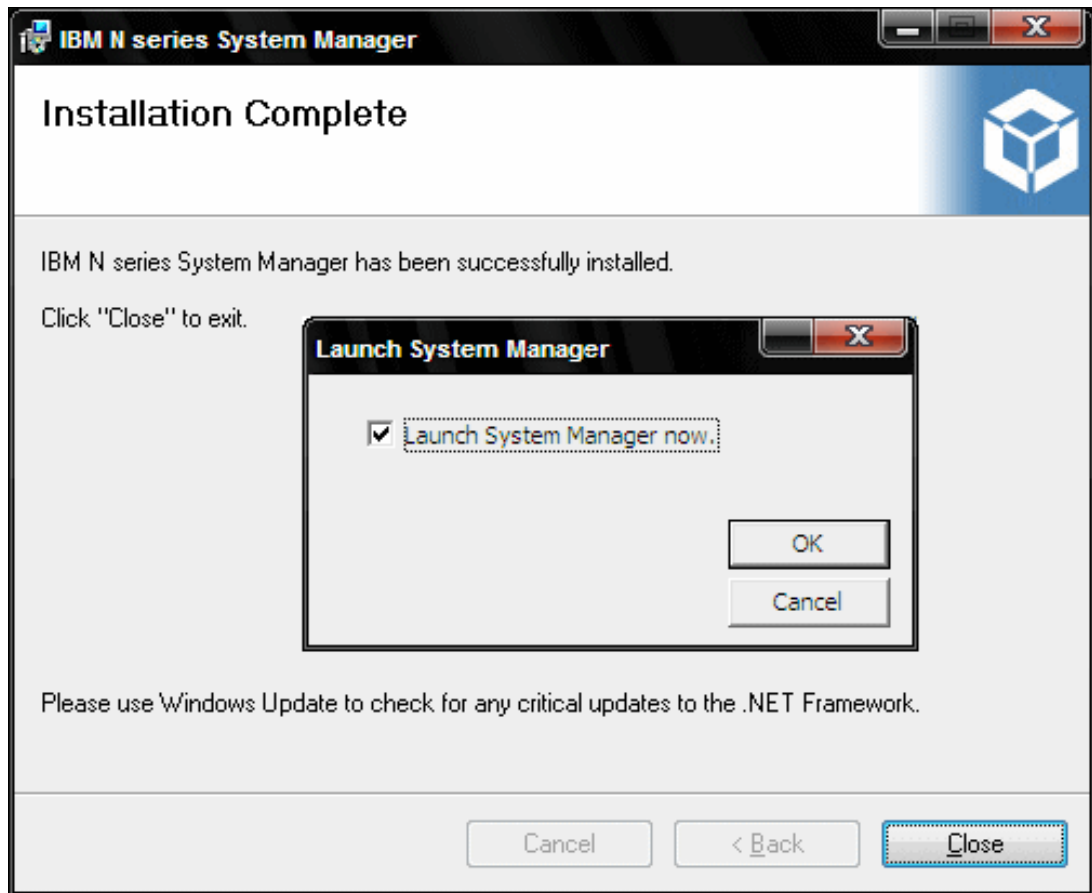


Figure 8-6 IBM N series System Manager installation complete

## 8.2.6 Starting IBM N series System Manager

Once you have installed IBM N series System Manager, start it as follows:

1. Click the **OK** button after installation, or click the IBM N series System Manager icon on your desktop. The IBM N series System Manager Software will launch. Figure 8-7 on page 124 shows the initial window of System Manager.

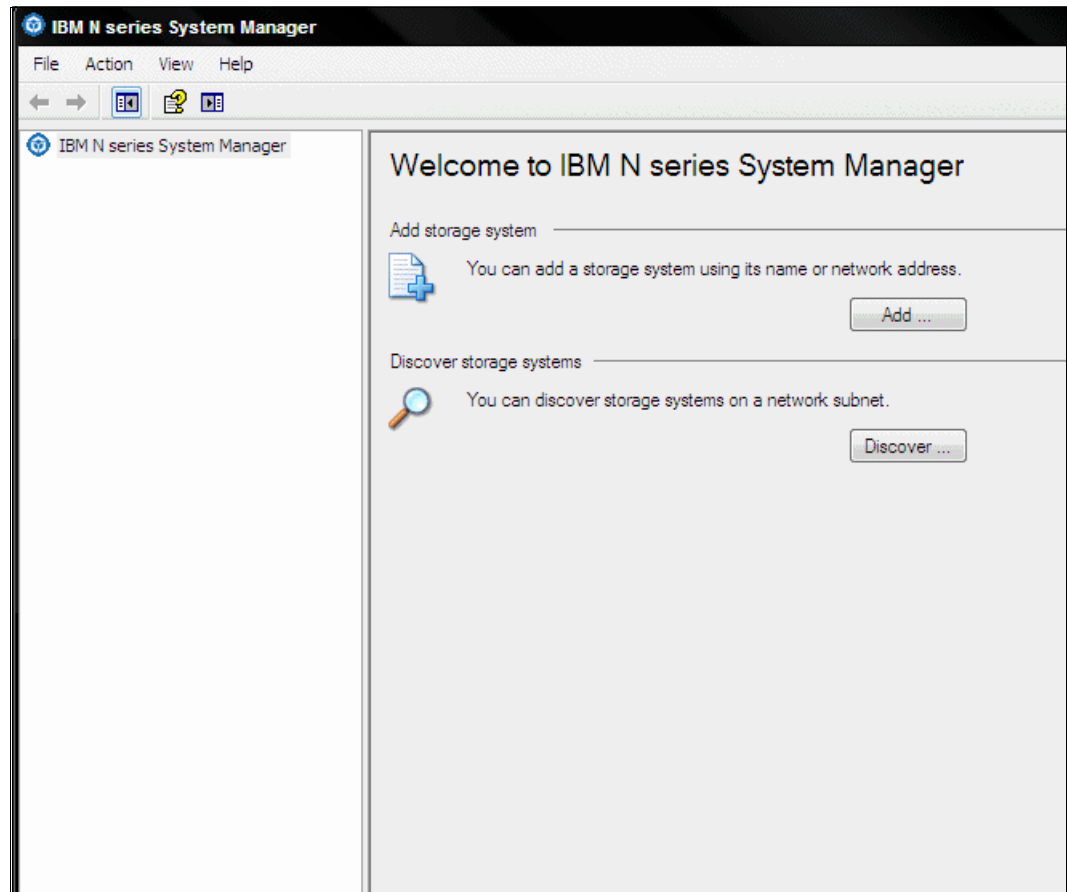


Figure 8-7 Initial window of IBM N series System Manager

2. If you know the IP address of your IBM N series Storage Systems, click the **Add** button and specify the IP address of the storage to manage as shown in Figure 8-8 on page 125.



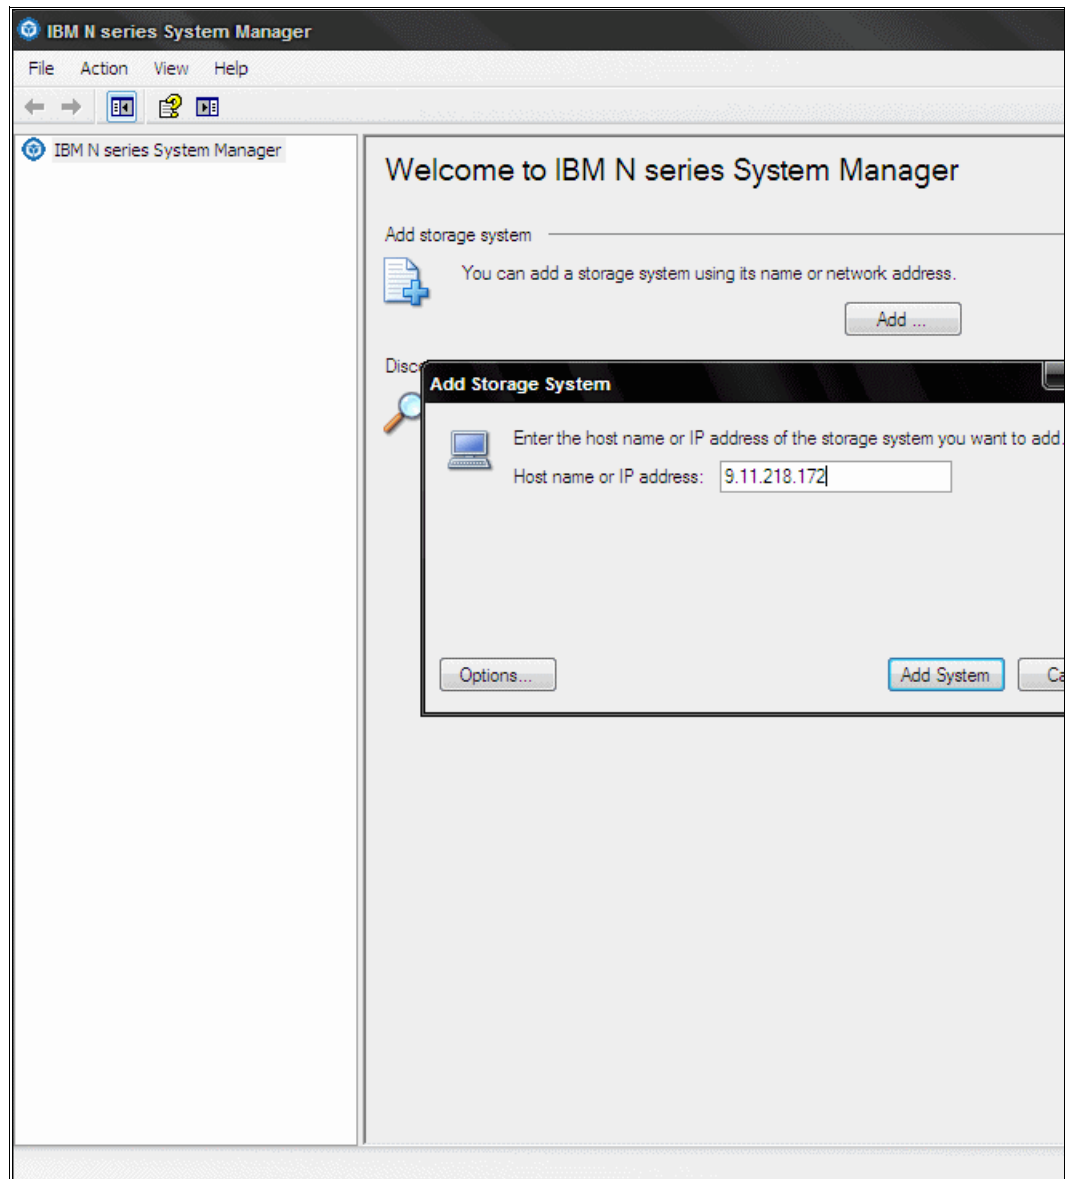


Figure 8-8 Adding IBM N series storage to System Manager

If you do not know the IP address of the IBM N series storage system, click the **Discover** button and the IBM N series System Manager will discover all IBM N series Storage System running in the network as shown in Figure 8-9 on page 126

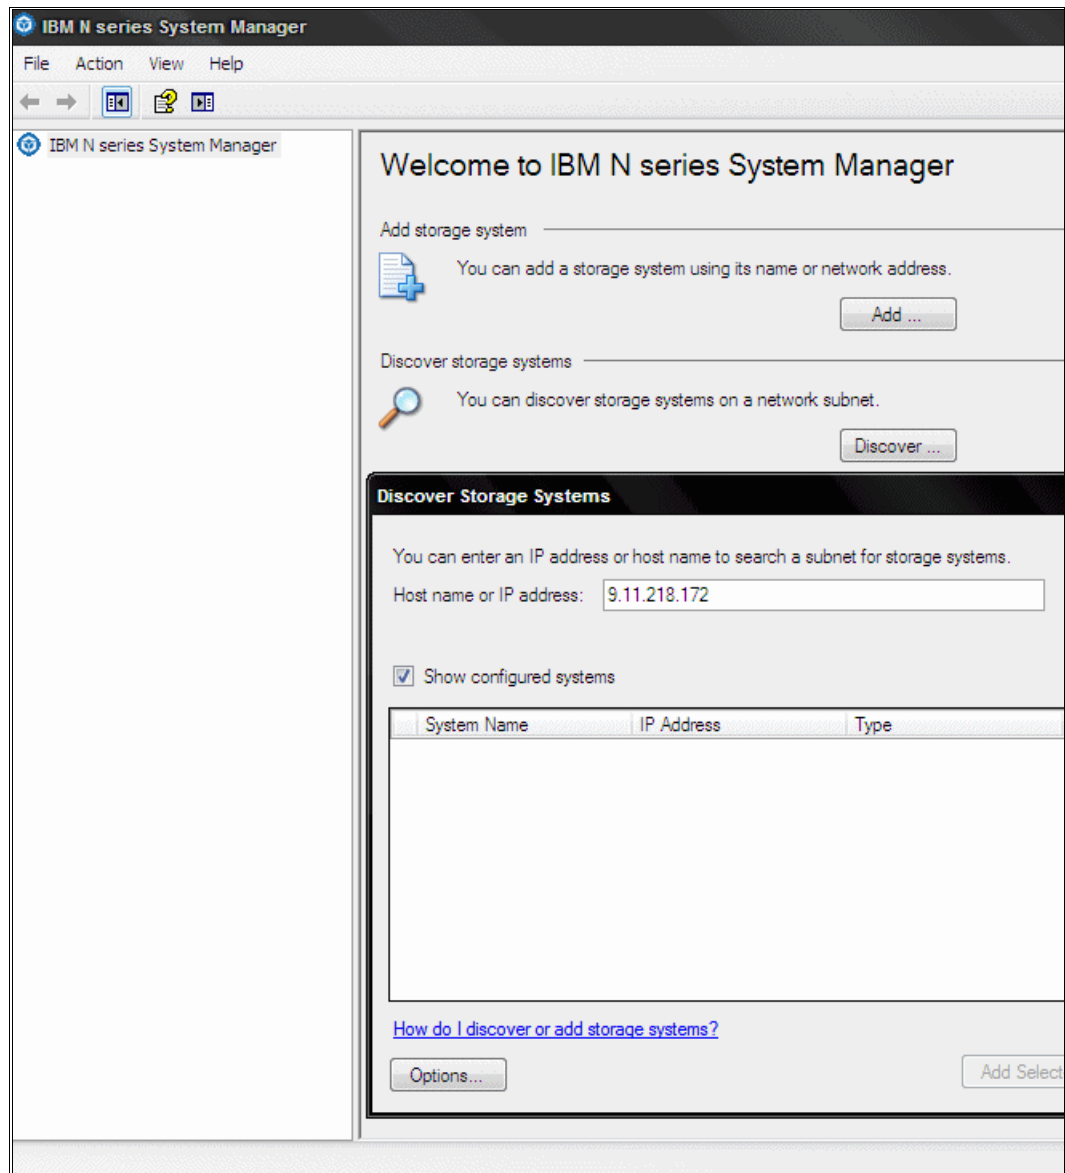


Figure 8-9 Discovering IBM N series storage system

Figure 8-10 on page 127 shows the discovered IBM N series storage system running in the network.




Figure 8-10 Discovered IBM N series storage system

3. To gain access to the storage system, a user id and password must be specified for each system as show in Figure 8-11 on page 128.

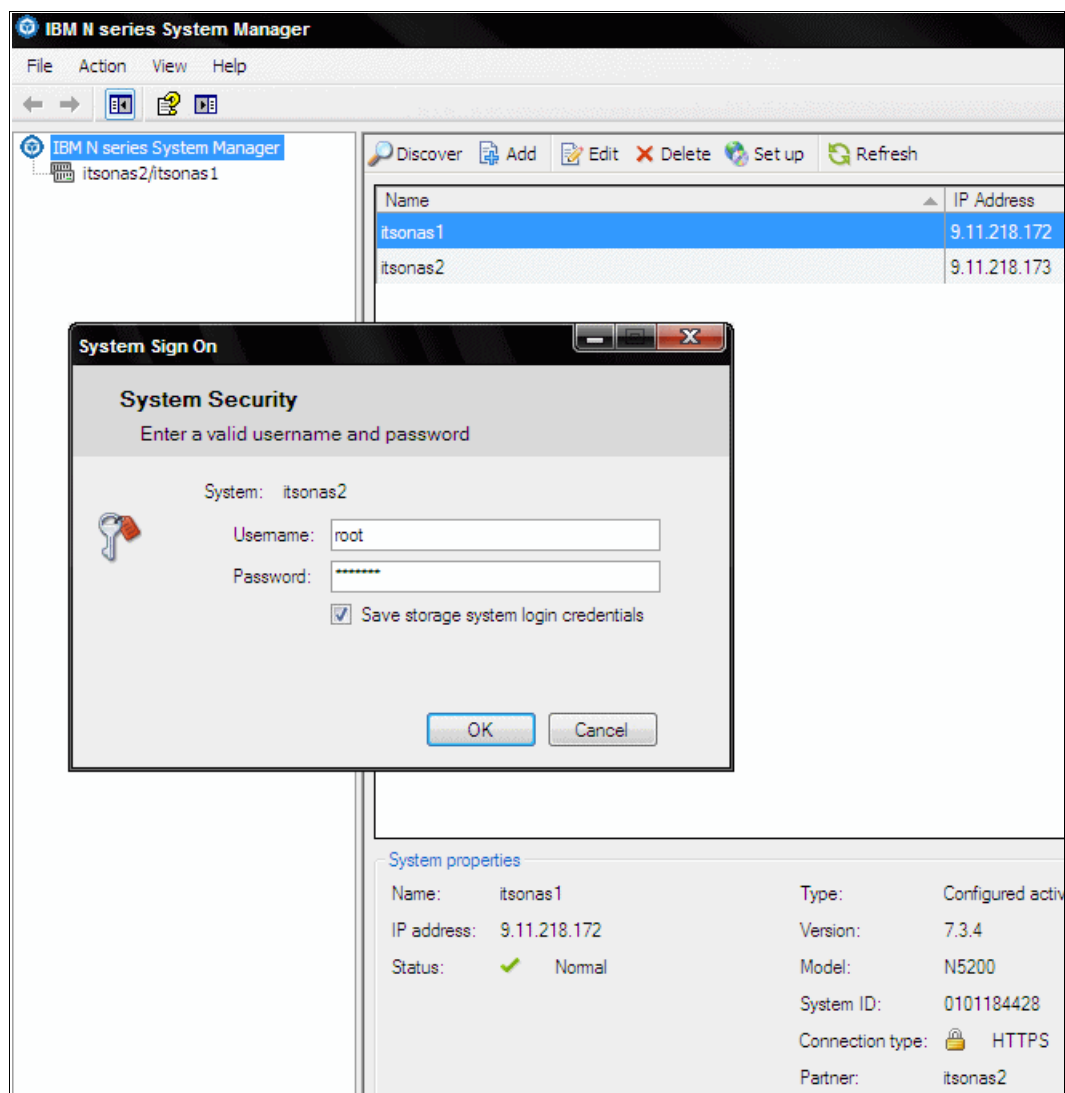


Figure 8-11 Accessing the IBM N series storage system

Figure 8-12 on page 129 shows the Managed IBM N series storage system list. Properties, Performance, Recommendations, and Reminders are displayed in the System Manager window.

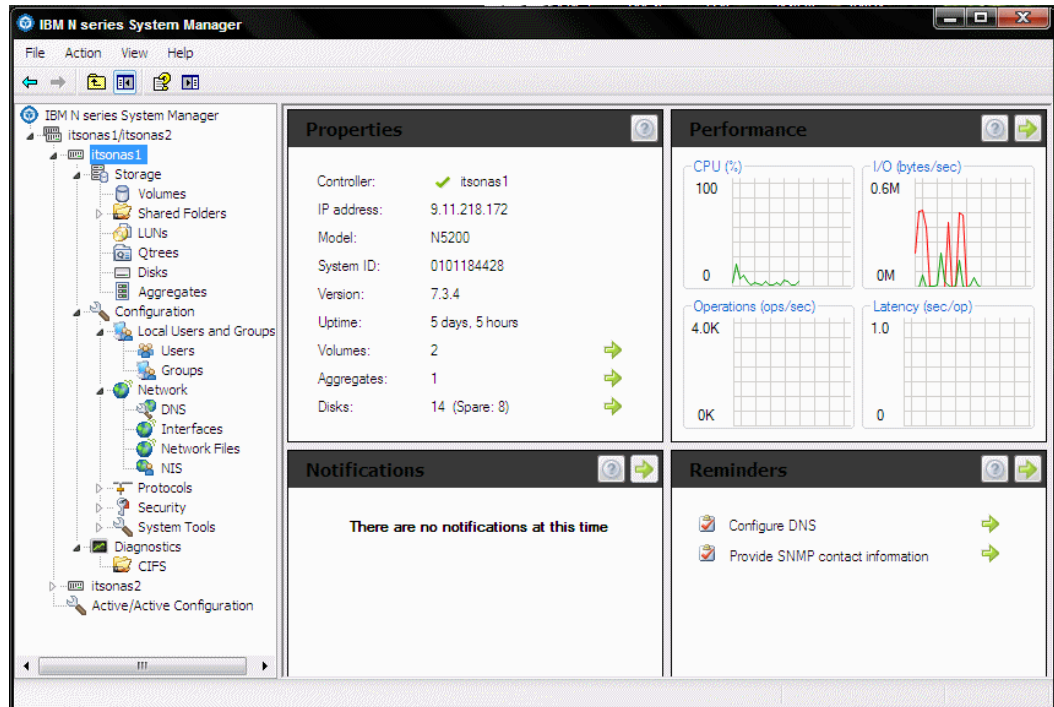


Figure 8-12 Managed IBM N series storage system

**Note:** IBM N series System Manager will be used in 8.3, “End-to-end Deduplication configuration example using command line” on page 152 in the end-to-end deduplication (Advanced Single Instance Storage) configuration example.

Frequent tasks such as Creating volume, Create LUNs, Create share/export and Create NFS datastore for Vmware can be performed using the IBM N series System Manager Software as shown in Figure 8-13 on page 130.

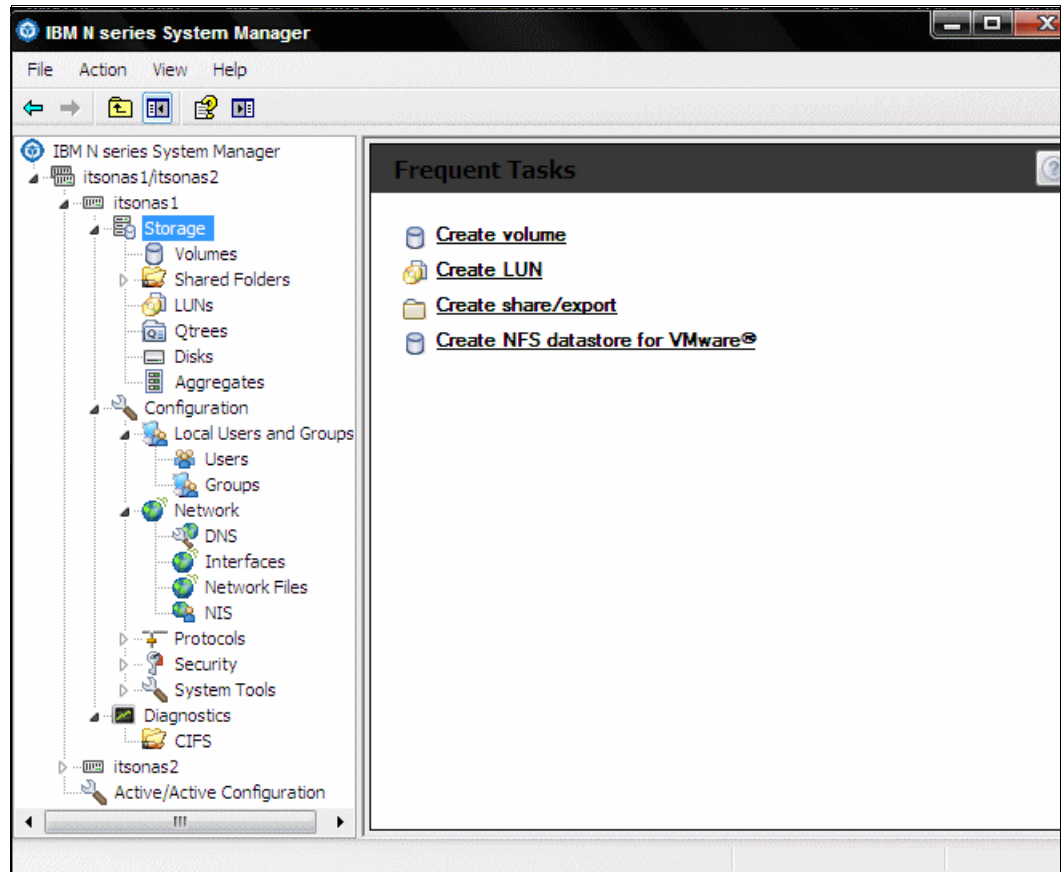


Figure 8-13 Frequent Task in IBM N series System Manager

Figure 8-14 on page 131 shows the created volume in the IBM N series storage system with details.

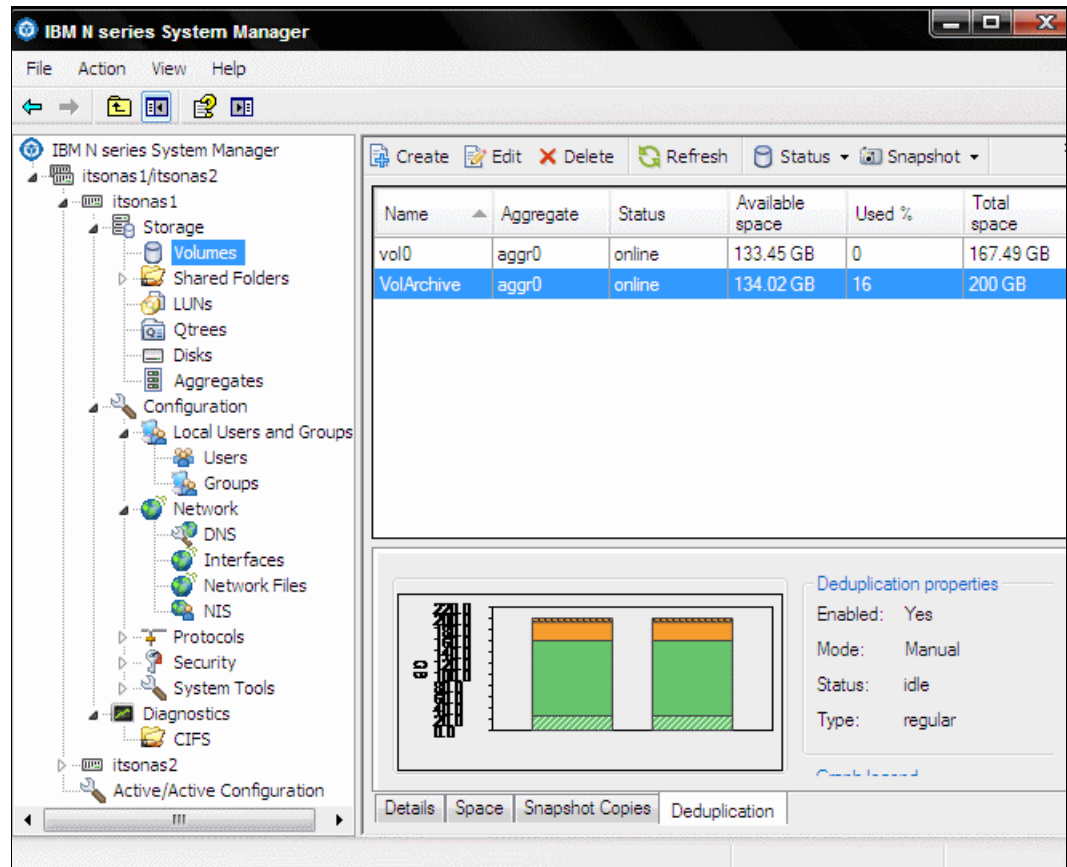


Figure 8-14 Created Volume in IBN N series storage

Figure 8-15 on page 132 shows the Shared folders being exported and details of each shared folder. If the deduplication is enabled in the volume, you will see a Deduplication tab.

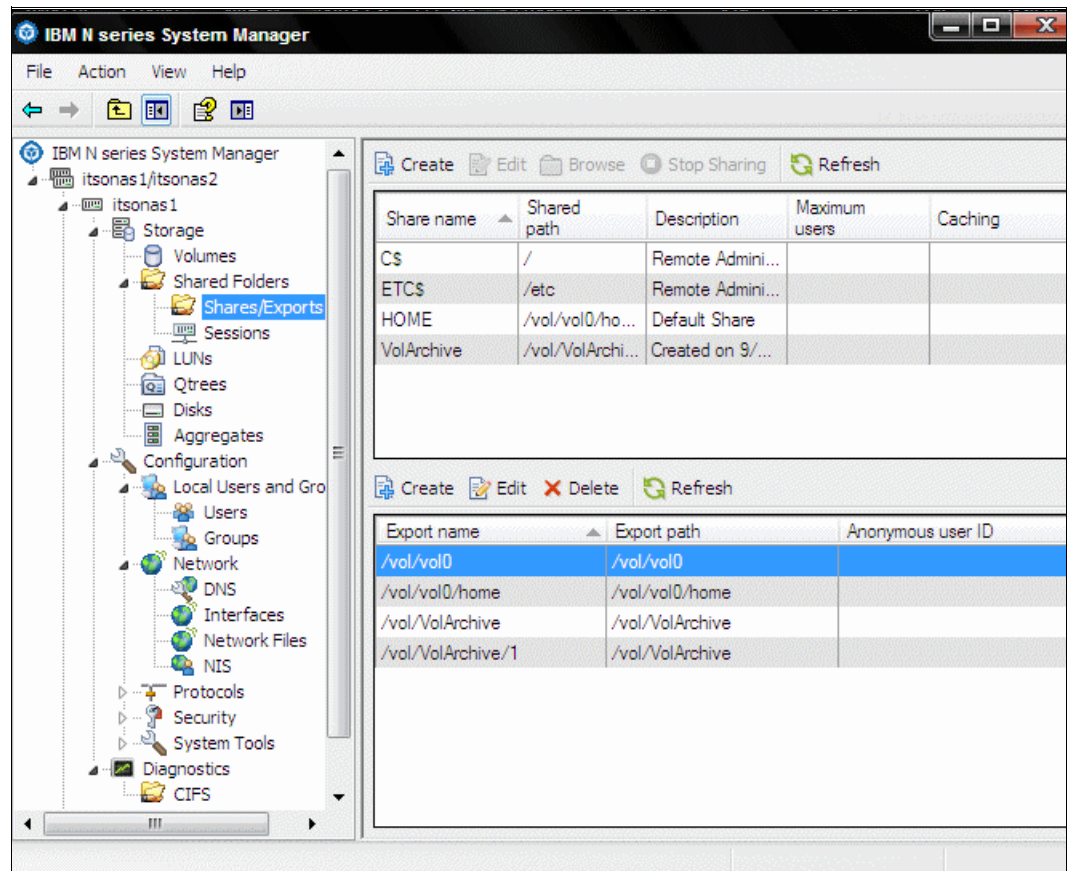


Figure 8-15 Shared Folders

Figure 8-16 on page 133 shows the session through which hosts are connected to the shared folders.






Figure 8-16 Session lists on folders

IBM N series System Manager provides views that allow you to perform tasks:

- Figure 8-17 on page 134 shows the LUNs Management view. From this window you can **Create**, **Edit**, **Delete**, show **Status**, and **Manage Snapshots**.




Figure 8-17 LUNs Management view

- Figure 8-18 on page 135 shows the Qtrees view, where you can **Create**, **Edit**, and **Delete**.

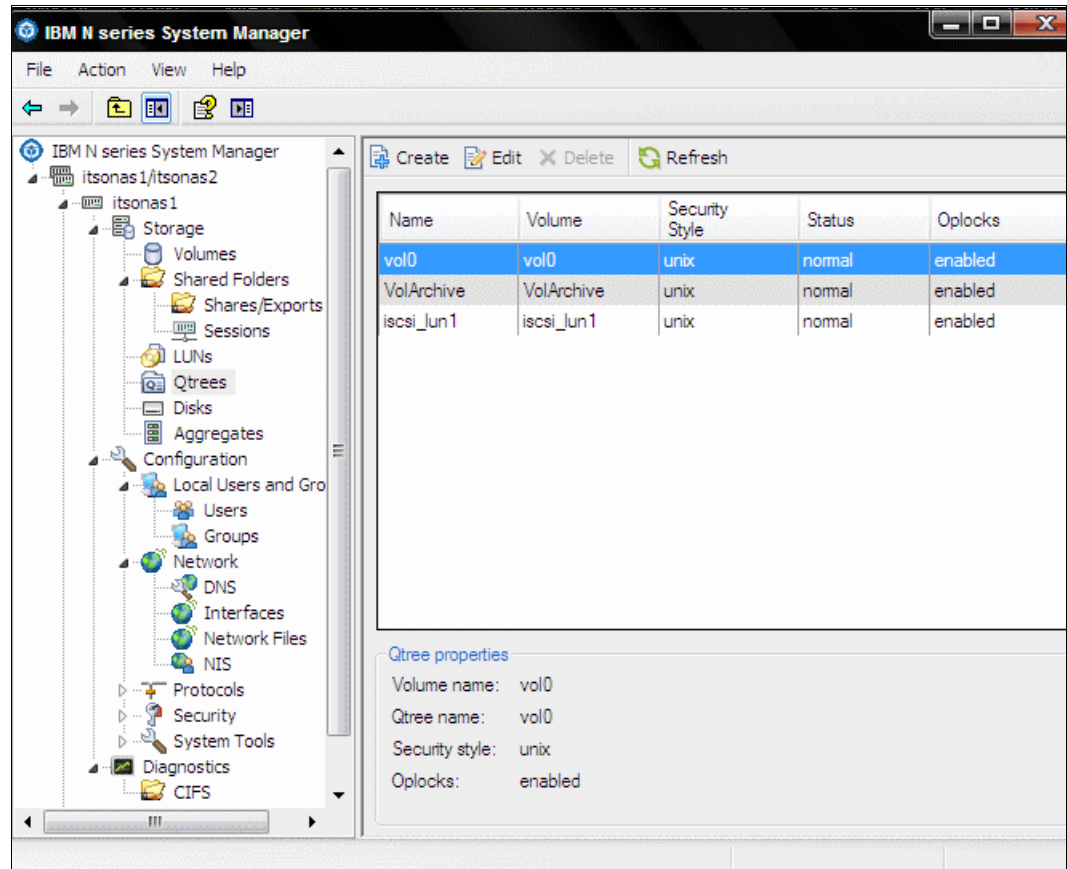


Figure 8-18 Qtrees Management view

- Figure 8-19 on page 136 shows the Disk Management view that shows the disks that are included in aggregate. You can **Create** a disk and **Add** a disk to an aggregate.

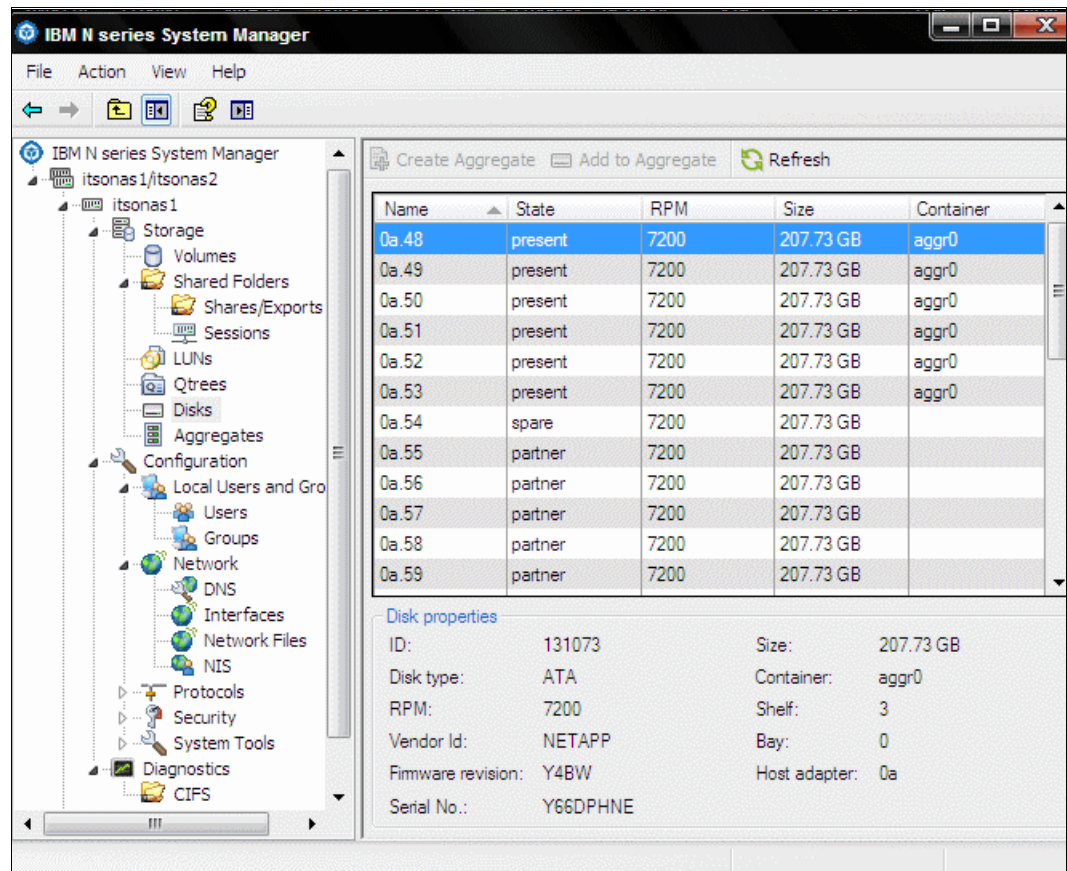


Figure 8-19 Disk Management view

- Figure 8-20 on page 137 shows the Aggregate Management view. You can **View**, **Create** and **Delete** aggregates from this view.

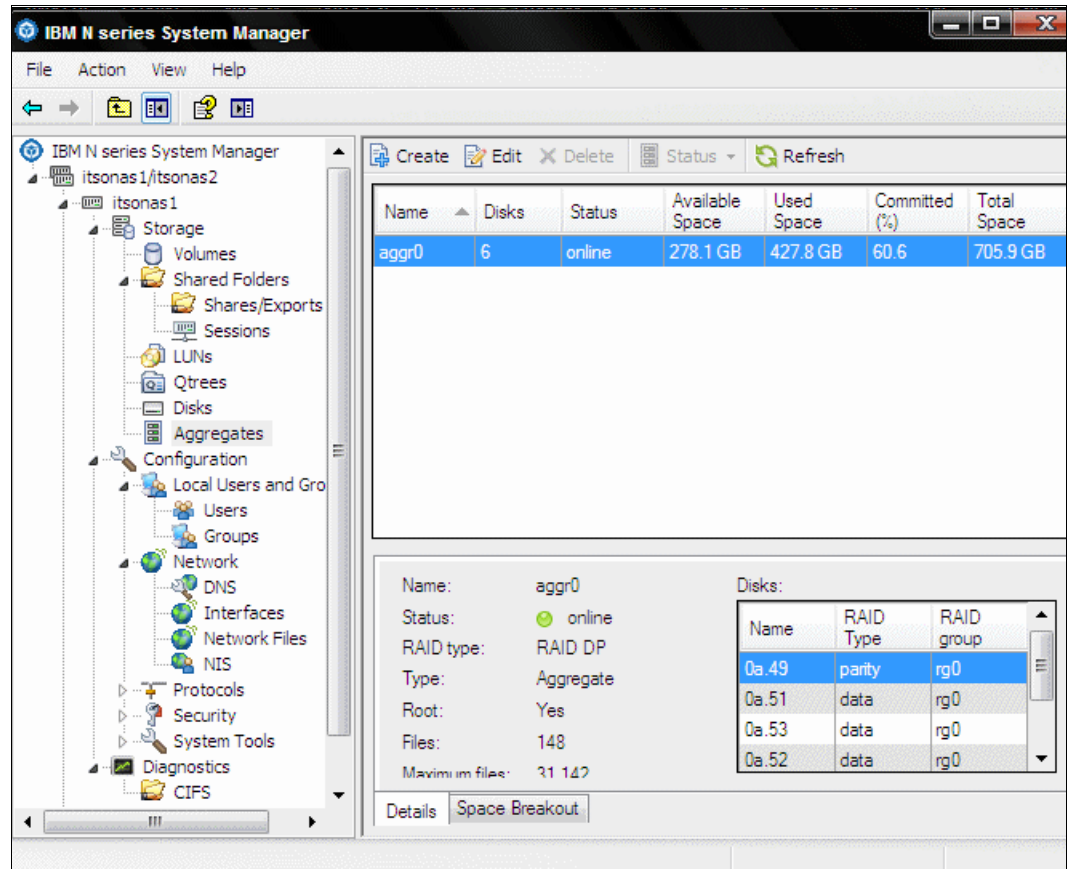


Figure 8-20 Aggregate Management view

- Figure 8-21 on page 138 shows the User Management view. You can **Create**, **Delete**, **Edit** and **Set Password** for users.




Figure 8-21 User Management

- Figure 8-22 on page 139 shows the Groups Management view. You can **Add**, **Delete**, and **View** users who belong the selected group.

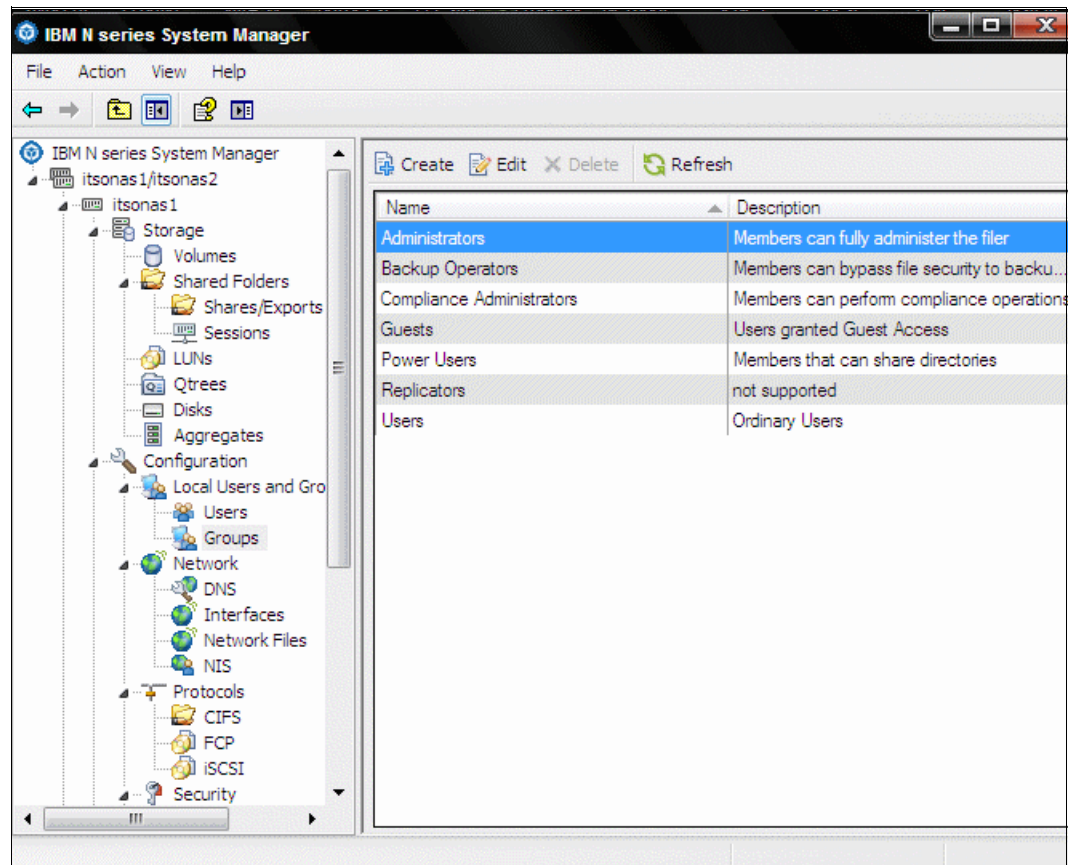


Figure 8-22 Group Management view

- Figure 8-23 on page 140 shows the DNS Setup view. In this view you can add the DNS IP address and turn on the service, the caching, and the dynamic update.

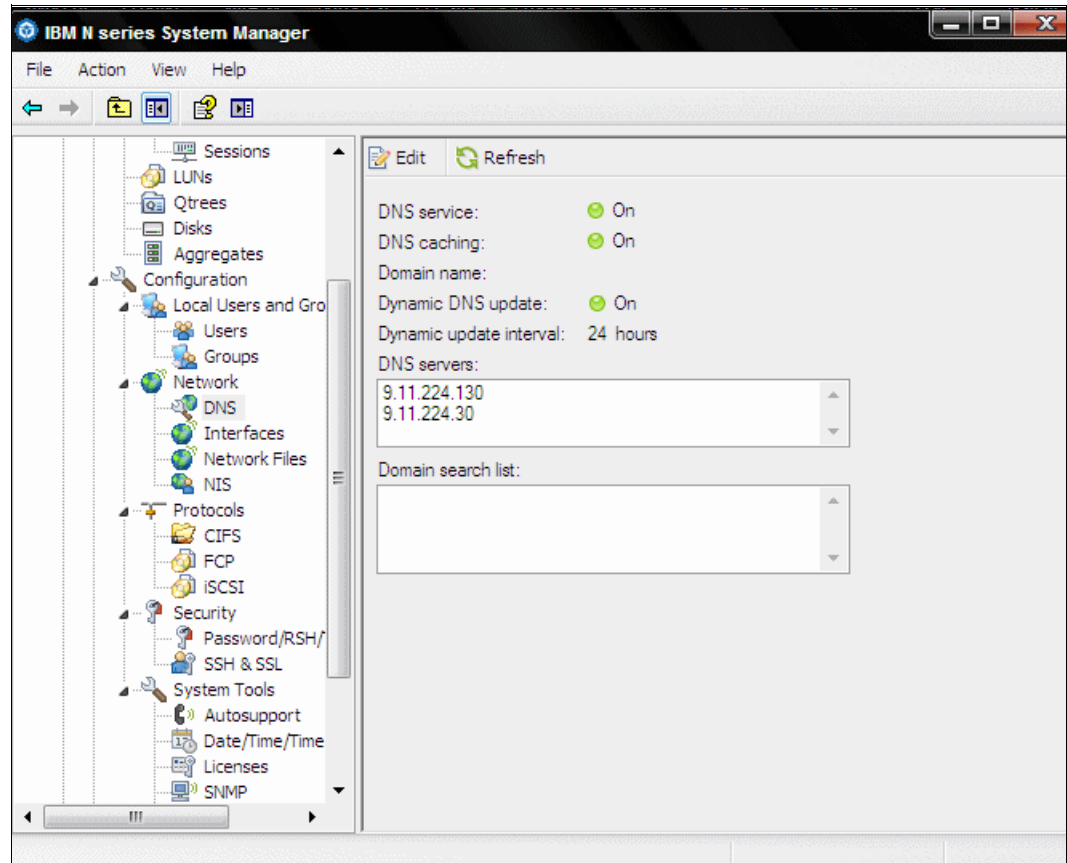


Figure 8-23 DNS Setup view

- Figure 8-24 on page 141 shows the Interfaces Management view. In this view, you can see the details of each interfaces attached to the IBM N series storage system.



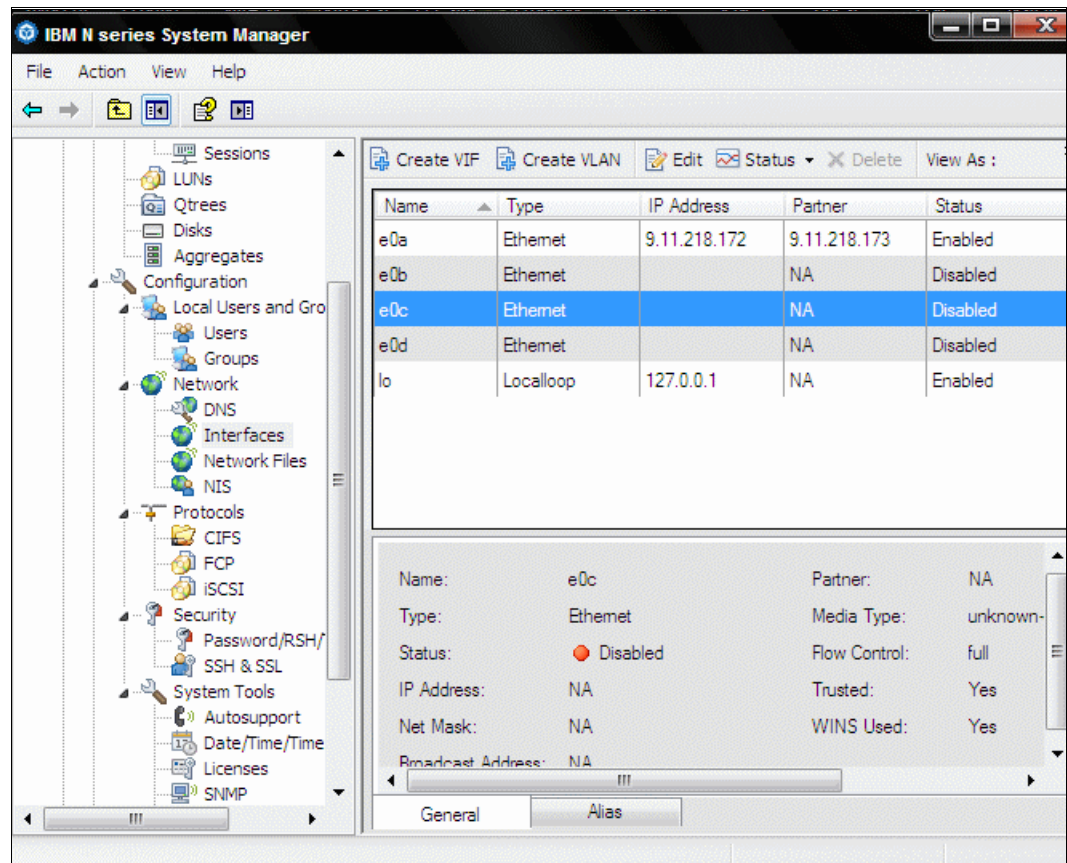


Figure 8-24 Interfaces Management

- Figure 8-25 on page 142 shows the IP address and the host name of the IBM N series system storage.




Figure 8-25 Host name of the IBM N series storage system

- Figure 8-26 on page 143 shows the Network Information Management (NIS). Can specify additional NIS servers using IP.

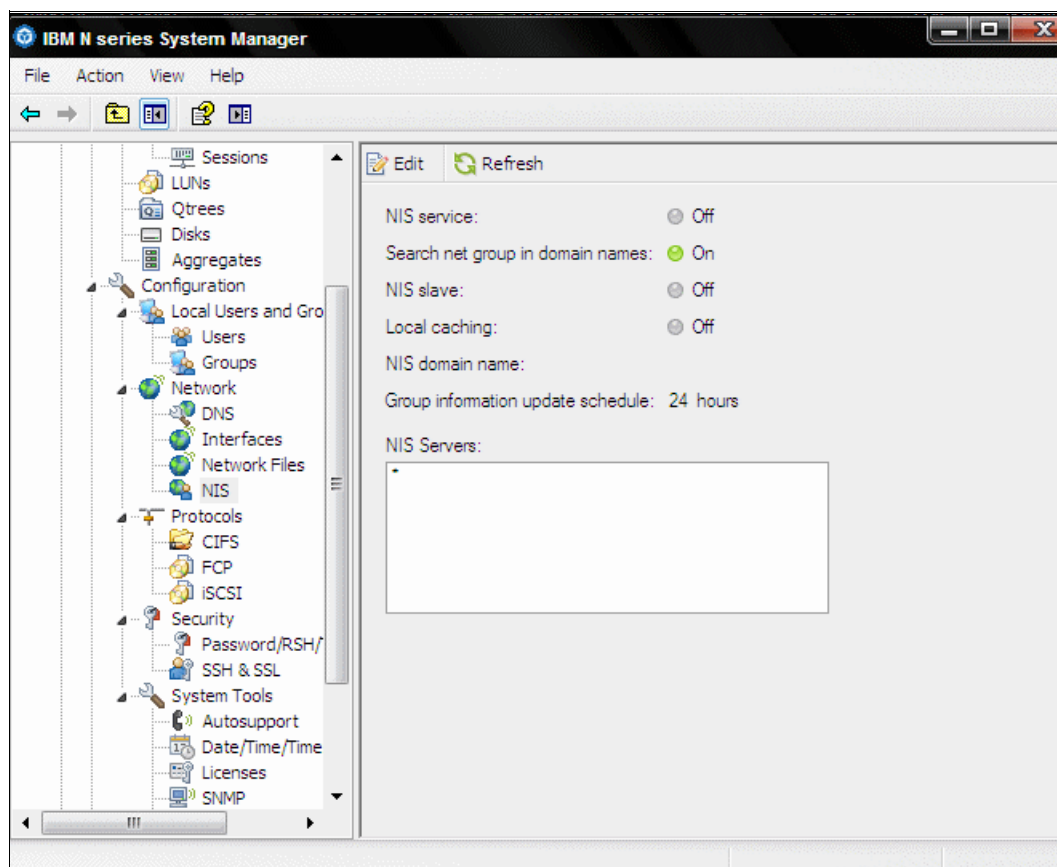


Figure 8-26 NIS Management view

- Figure 8-27 on page 144 shows the CIFS Protocol Management view. In this view you can see the details of the CIFS and the Home directories, and auditing can be enabled.

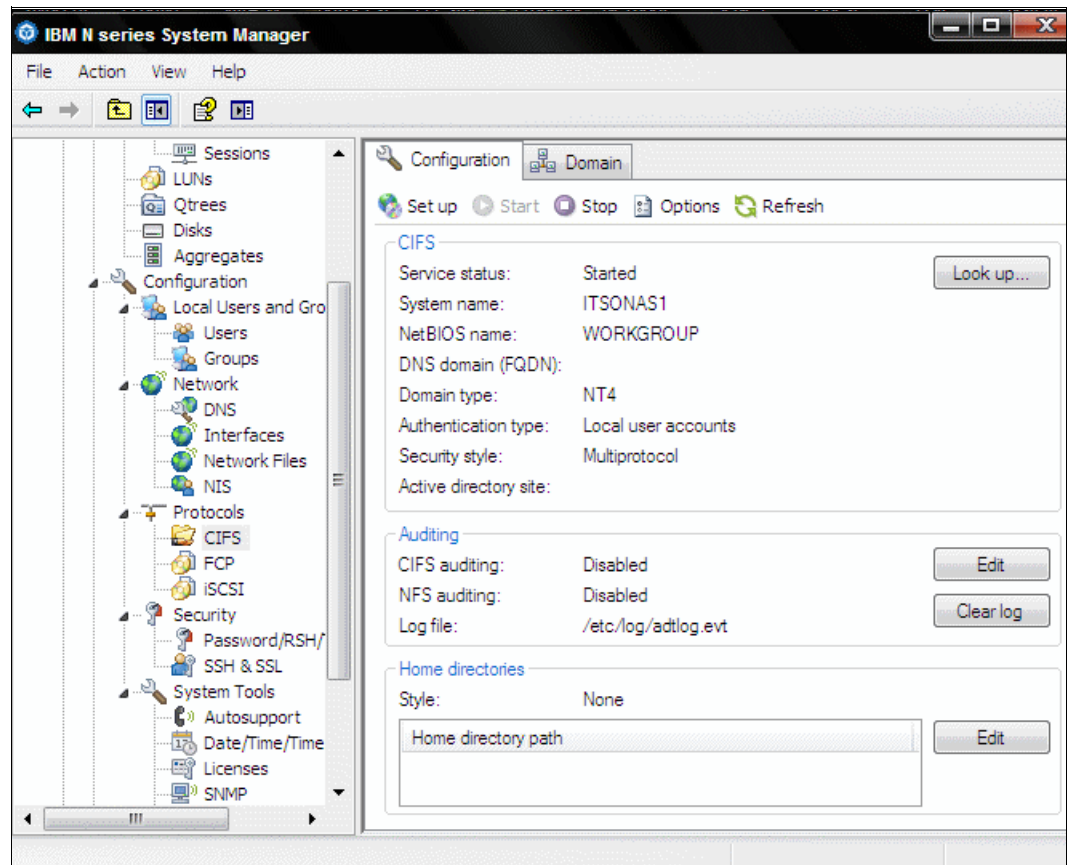


Figure 8-27 CIFS Protocol Management view

- Figure 8-28 on page 145 shows the FCP Protocol Management view, in which you can see the status of FCP Service, and **Start** and **Stop** it.




Figure 8-28 FCP Protocol Management view

- Figure 8-29 on page 146 shows the iSCSI Protocol Management view. Here you can see the details if the service is running, the target nodename, and the alias and details needed.

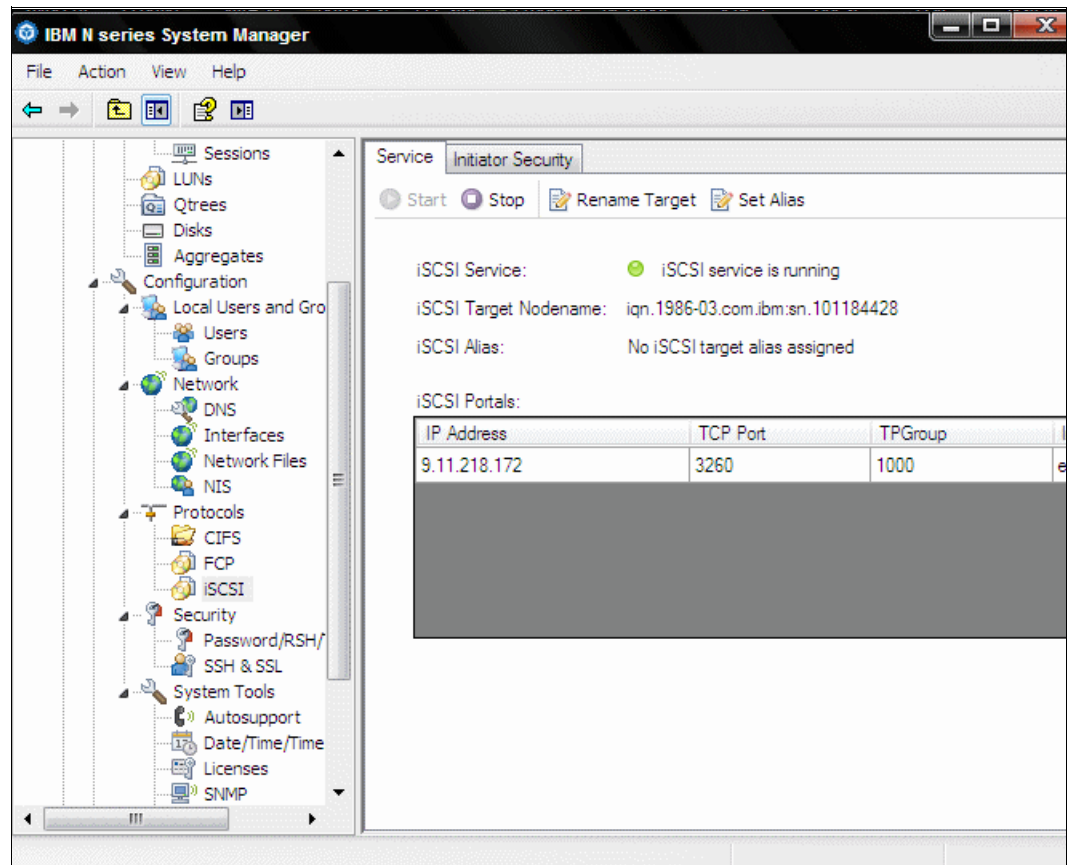


Figure 8-29 iSCSI Protocol Management view

- Figure 8-30 on page 147 shows the Security and Password Management view. It shows the trusted host and RSH settings.

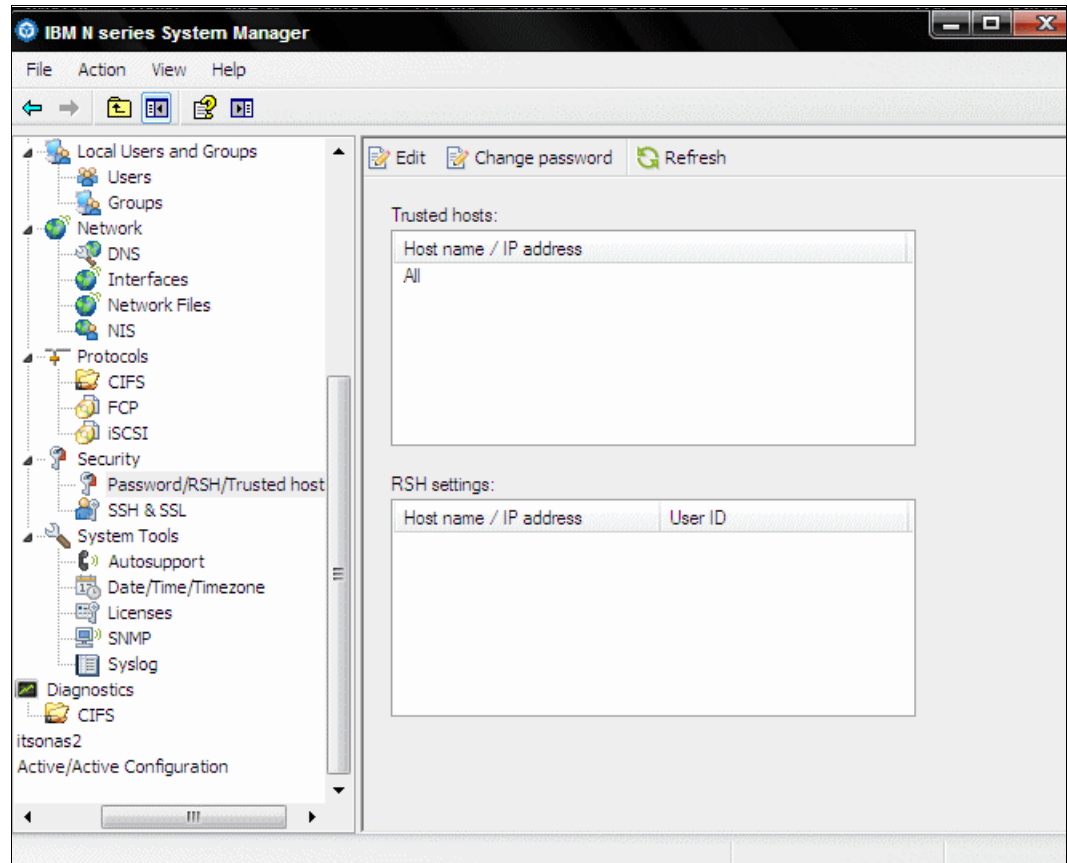


Figure 8-30 Security and Password Management view

- Figure 8-31 on page 148 shows the SSH and SSL Setup view. In this view you can generate SSH keys and certificates for secure communication.

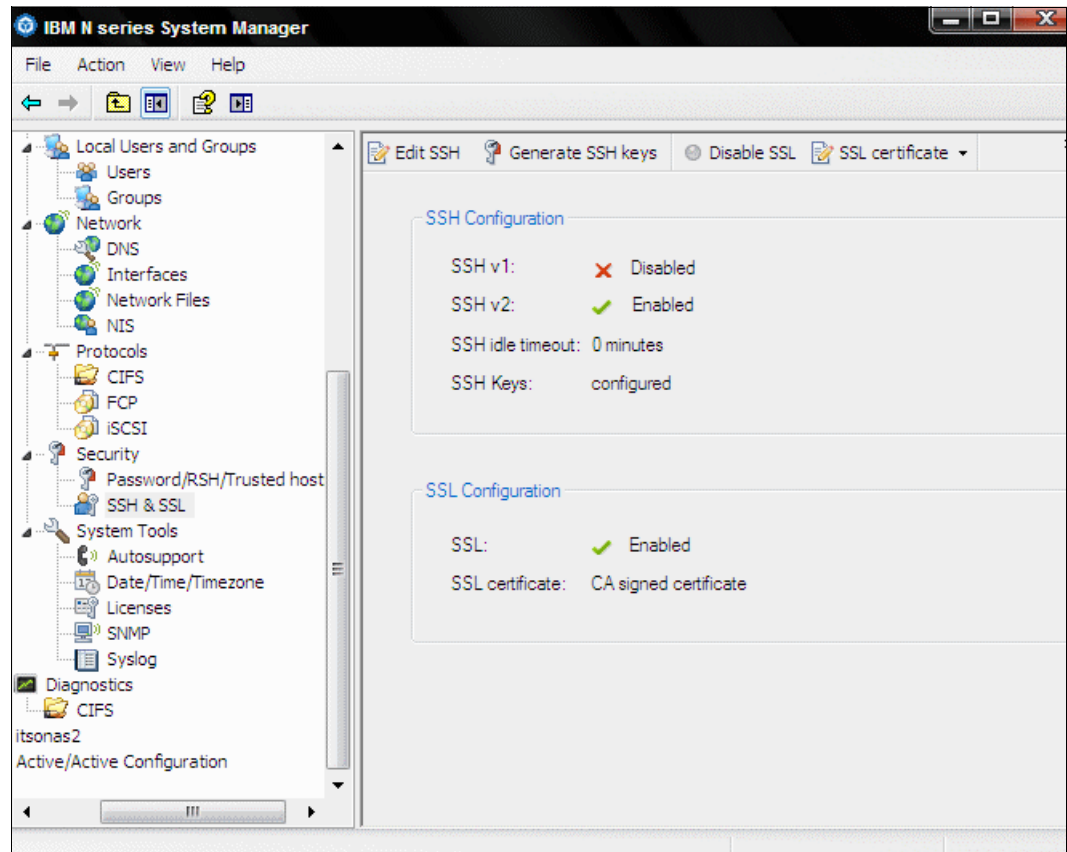


Figure 8-31 SSH and SSL Setup view

- The Auto Support Setup view is shown in Figure 8-32 on page 149. In this view you can specify the mail host name, message recipient email addresses, and the minimal message recipient email addresses.



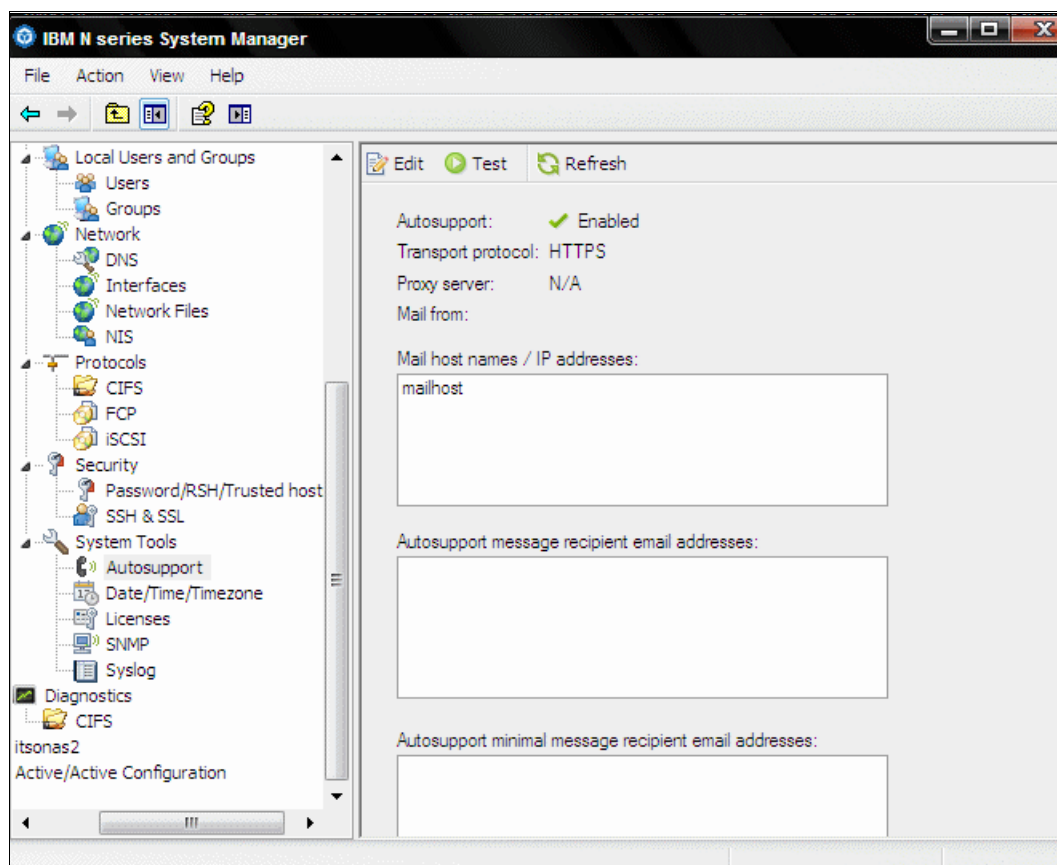


Figure 8-32 Auto Support Setup view

- The Date and Time Zone Setup view is shown in Figure 8-33 on page 150. This view shows the time servers being used by the IBM N series storage system.

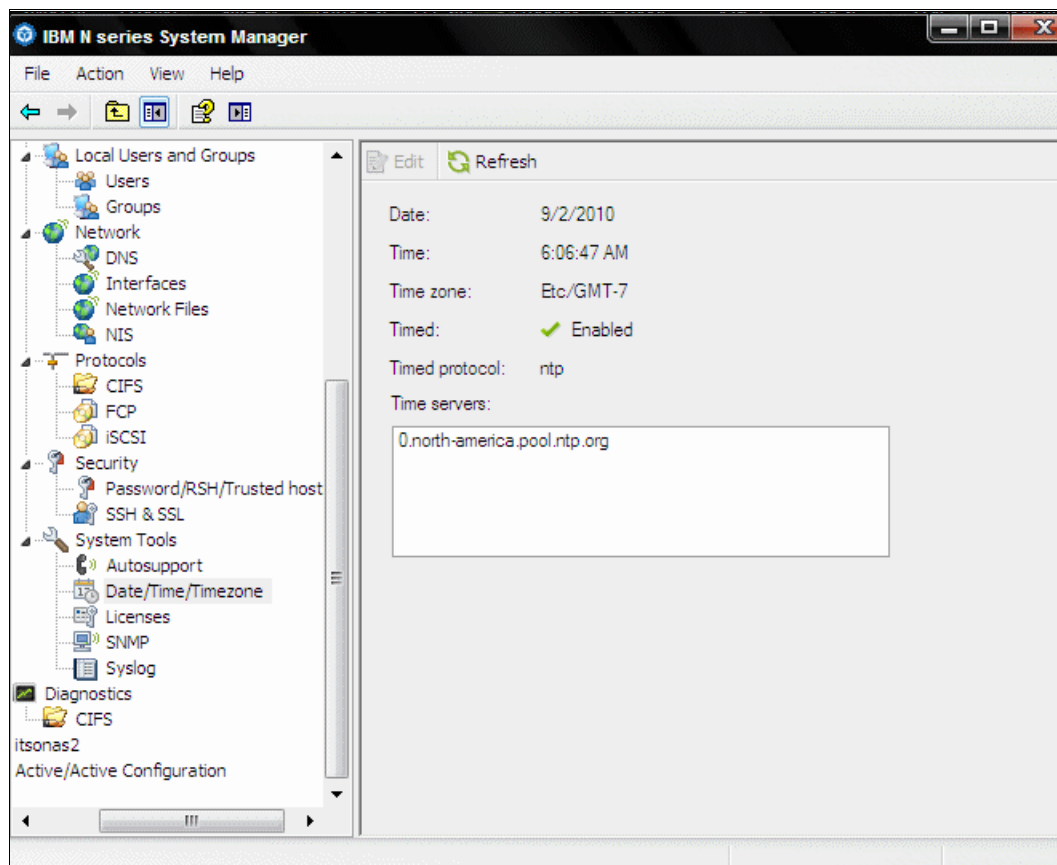


Figure 8-33 Date and Time Zone setup

- the SNMP setup can be edited in the SNMP Setup view shown in Figure 8-34 on page 151. If there are multiple SNMP IP addresses, simply add those addresses in the table shown in the figure.

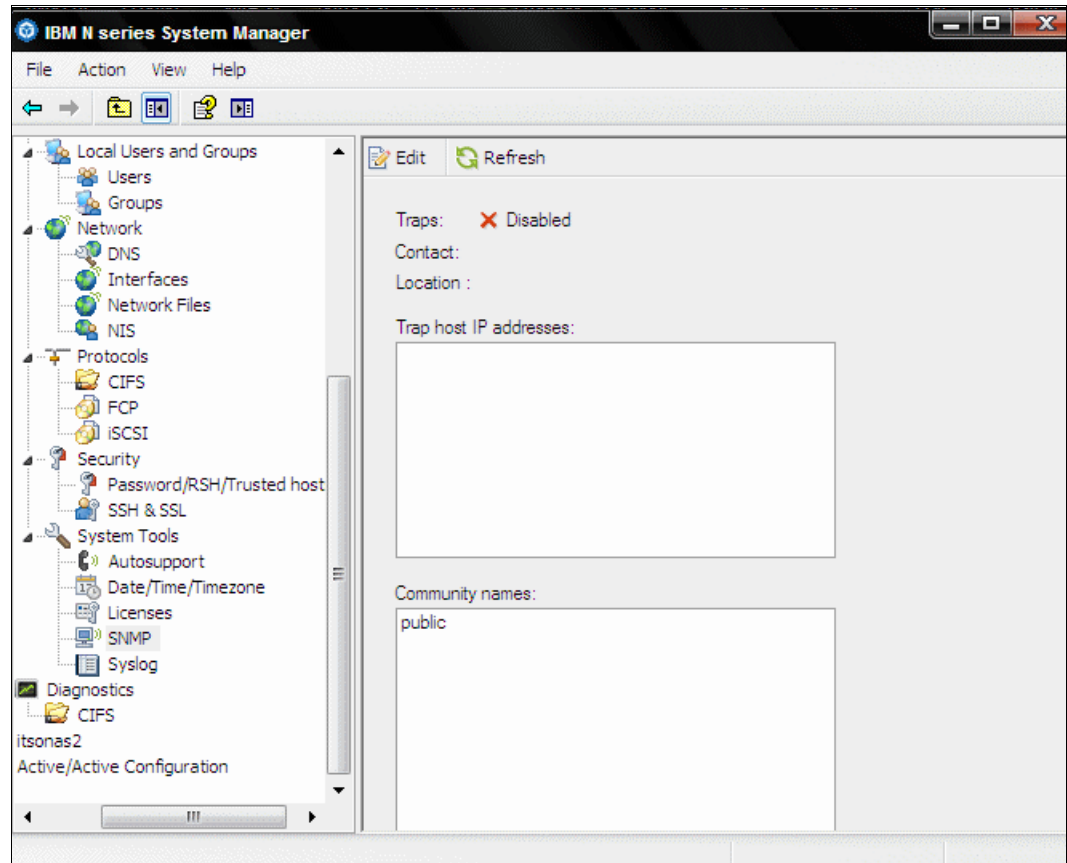


Figure 8-34 SNMP setup

- The System Log of the IBM N series storage system can be viewed as shown in Figure 8-35 on page 152. Events, severity, and messages, and when these occurred can be seen in this view.

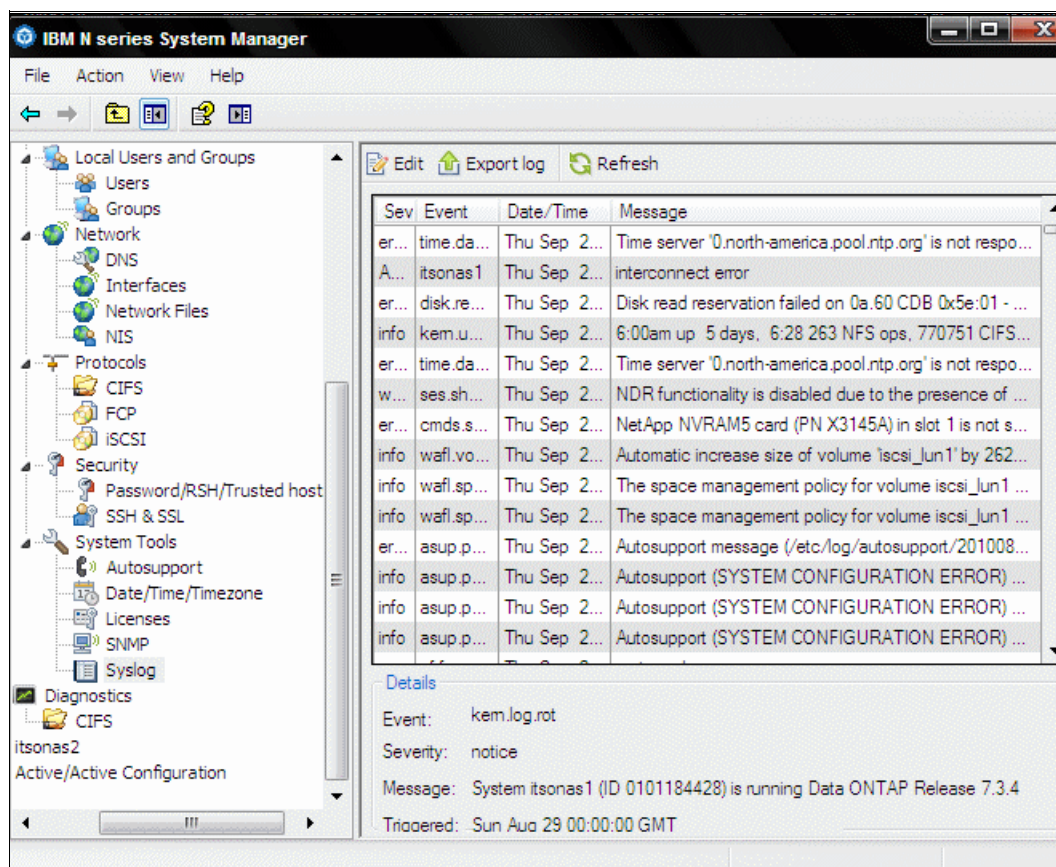


Figure 8-35 Syslog Management view

**Note:** IBM N series System Manager Software is available at no cost.

## 8.3 End-to-end Deduplication configuration example using command line

This example creates a place to archive several large data files. The destination IBM System Storage N series storage system is called itsonas1, and it is assumed that deduplication has been licensed on this machine.

Create a flexible volume with deduplication enabled by performing the following steps:

1. Create a flexible volume as shown in Example 8-1. Keep in mind the maximum allowable volume size for the platform, as specified in the Table 8-1 on page 115.

*Example 8-1 Create a flexible volume*

```
itsonas1> vol create VolArchive aggr0 200g
Creation of volume 'VolArchive' with size 200g on containing aggregate 'aggr0'
has completed
```

2. Enable deduplication on the flexible volume and verify that it is turned on. The **vol status** command shows the attributes for flexible volumes that have deduplication turned on.

After you turn on deduplication, Data ONTAP lets you know that if this were an existing flexible volume that already contained data before deduplication was enabled, you would want to run the **sis start -s** command. Example 8-2 is using a new flexible volume, so running the command is not necessary.

*Example 8-2 Using the vol status command*

---

```
itsonas1> sis on /vol/VolArchive
Deduplication for "/vol/VolArchive" is enabled.
Already existing data could be processed by running "sis start -s
/vol/VolArchive"
itsonas1> vol status VolArchive
Volume State Status Options
VolArchive online raid_dp, flex nosnap=on
                sis
Containing aggregate: 'aggr0'
```

---

3. Another way to verify that A-SIS is enabled on the flexible volume is to simply check the output of **sis status** on the flexible volume, as shown in Example 8-3.

*Example 8-3 Running the sis status command*

---

```
itsonas1> sis status /vol/VolArchive
Path           State   Status Progress
/vol/VolArchive Enabled Idle   Idle for 00:00:20
```

---

4. Turn off the default A-SIS schedule as shown in Example 8-4 by modifying the **sis config** command.

*Example 8-4 Running the sis status command*

---

```
itsonas1> sis config /vol/VolArchive
Path           Schedule
/vol/VolArchive sun-sat@0
itsonas1> sis config -s - /vol/VolArchive
itsonas1> sis config /vol/VolArchive
Path           Schedule
/vol/VolArchive -
```

---

5. Mount the flexible volume to `/testDedup` on Linux server, and copy files from the users' directories into the new archive directory flexible volume. From the host perspective, the result is shown in Example 8-5.

*Example 8-5 Result from host perspective using Linux*

---

```
[root@localhost ~]# mkdir testDedup
[root@localhost ~]# mount itsonas1:/vol/VolArchive/1 testDedup
[root@localhost ~]# df -k testDedup
Filesystem              kbytes    used    avail    capacity Mounted on
itsonas1:/vol/VolArchive/1 167772160 89353344              78418816 54%
/testDedup
```

---

6. Next, examine the flexible volume, run deduplication, and monitor the status. Use the **df -s** command to examine the storage consumed and the space saved. Note that no space savings have been achieved by simply copying data to the flexible volume, even though deduplication is turned on. What has happened is that all the blocks that have been written to this flexible volume after deduplication was turned on have had their fingerprints written to the change log file. Refer to Example 8-6 on page 154.

*Example 8-6 Examine the storage consumed and space saved*

---

```
itsonas1> df -s /vol/VolArchive
Filesystem used saved %saved
/vol/VolArchive/ 89353528 0 0%
```

---

7. Run deduplication on the flexible volume. This step causes the change log to be processed, fingerprints to be sorted and merged, and duplicate blocks to be found. Refer to Example 8-7.

*Example 8-7 Run deduplication on the flexible volume*

---

```
itsonas1> sis start /vol/VolArchive
The deduplication operation for "/vol/VolArchive" is started.
```

---

8. Monitor the progress of deduplication by using the **sis status** command, as shown in Example 8-8.

*Example 8-8 Monitor deduplication progress*

---

```
itsonas1> sis status /vol/VolArchive
Path State Status Progress
/vol/VolArchive Enabled Active 65 GB Searched

itsonas1> sis status /vol/VolArchive
Path State Status Progress
/vol/VolArchive Enabled Active 538 MB (5%) Done

itsonas1> sis status /vol/VolArchive
Path State Status Progress
/vol/VolArchive Enabled Active 1424 MB (14%) Done

itsonas1> sis status /vol/VolArchive
Path State Status Progress
/vol/VolArchive Enabled Active 8837 MB (90%) Done

itsonas1> sis status /vol/VolArchive
Path State Status Progress
/vol/VolArchive Enabled Active 9707 MB (99%) Done

itsonas1> sis status /vol/VolArchive
Path State Status Progress
/vol/VolArchive Enabled Idle for 00:00:07
```

---

9. When **sis status** indicates that the flexible volume is again in the Idle state, deduplication has completed and you can check the space savings it provided in the flexible volume. Refer to Example 8-9.

*Example 8-9 Examine the storage consumed and space savings again*

---

```
itsonas1> df -s /vol/VolArchive
Filesystem used saved %saved
/vol/VolArchive/ 79515276 9856456 11%
```

---

## 8.4 End-to-end deduplication configuration example using IBM N series System Manager Software

To create a flexible volume with deduplication, perform the following steps:

1. Create a flexible volume using the IBM N series System Manager, as shown in Figure 8-36.

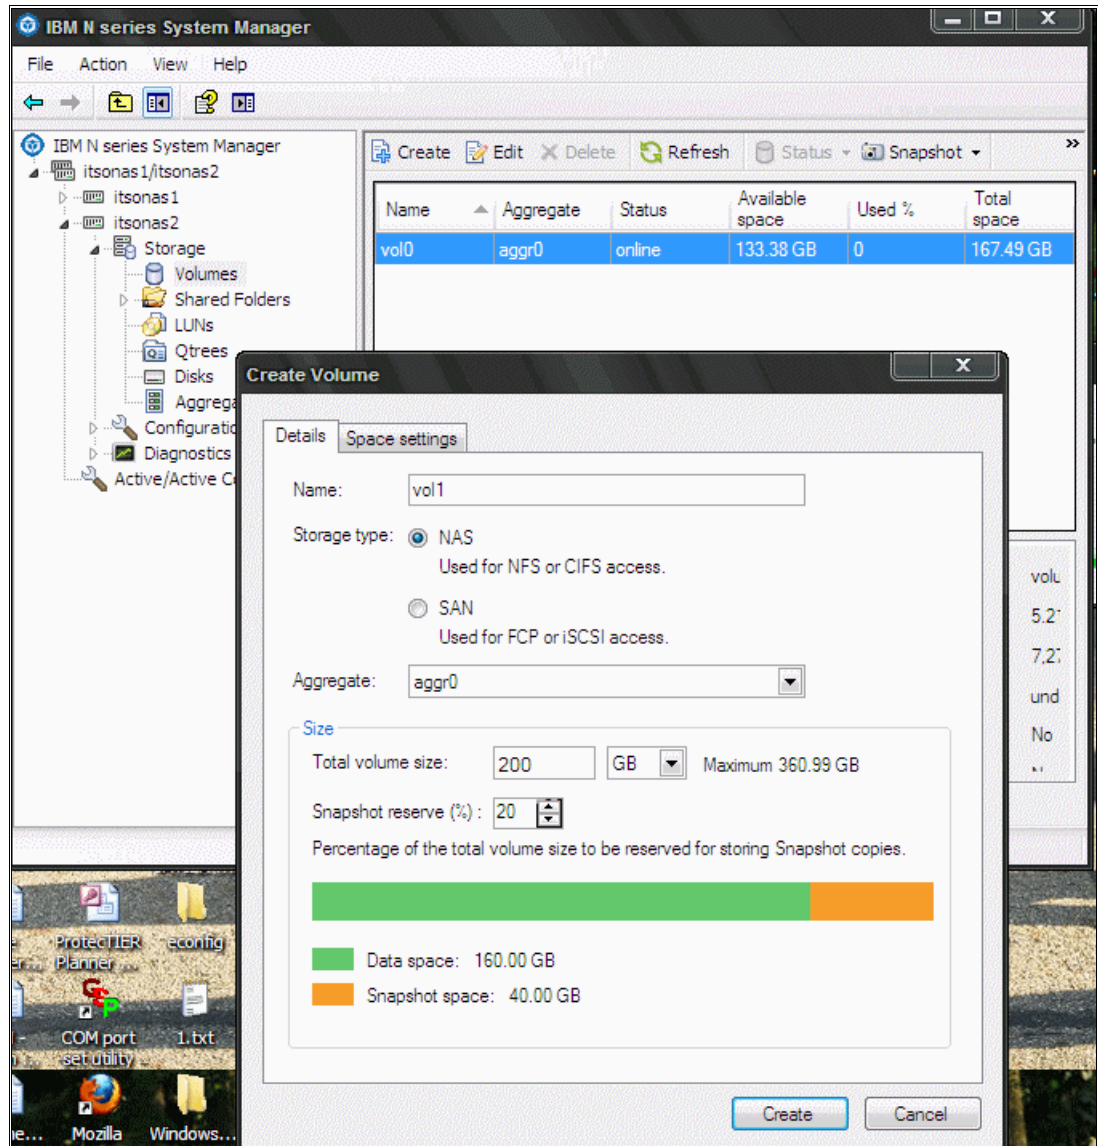


Figure 8-36 Creating a 200GB volume for deduplication

2. Enable deduplication on the flexible volume and verify that it is turned on as shown in Figure 8-37 on page 156, and click the **Next** button to create the flexible volume.



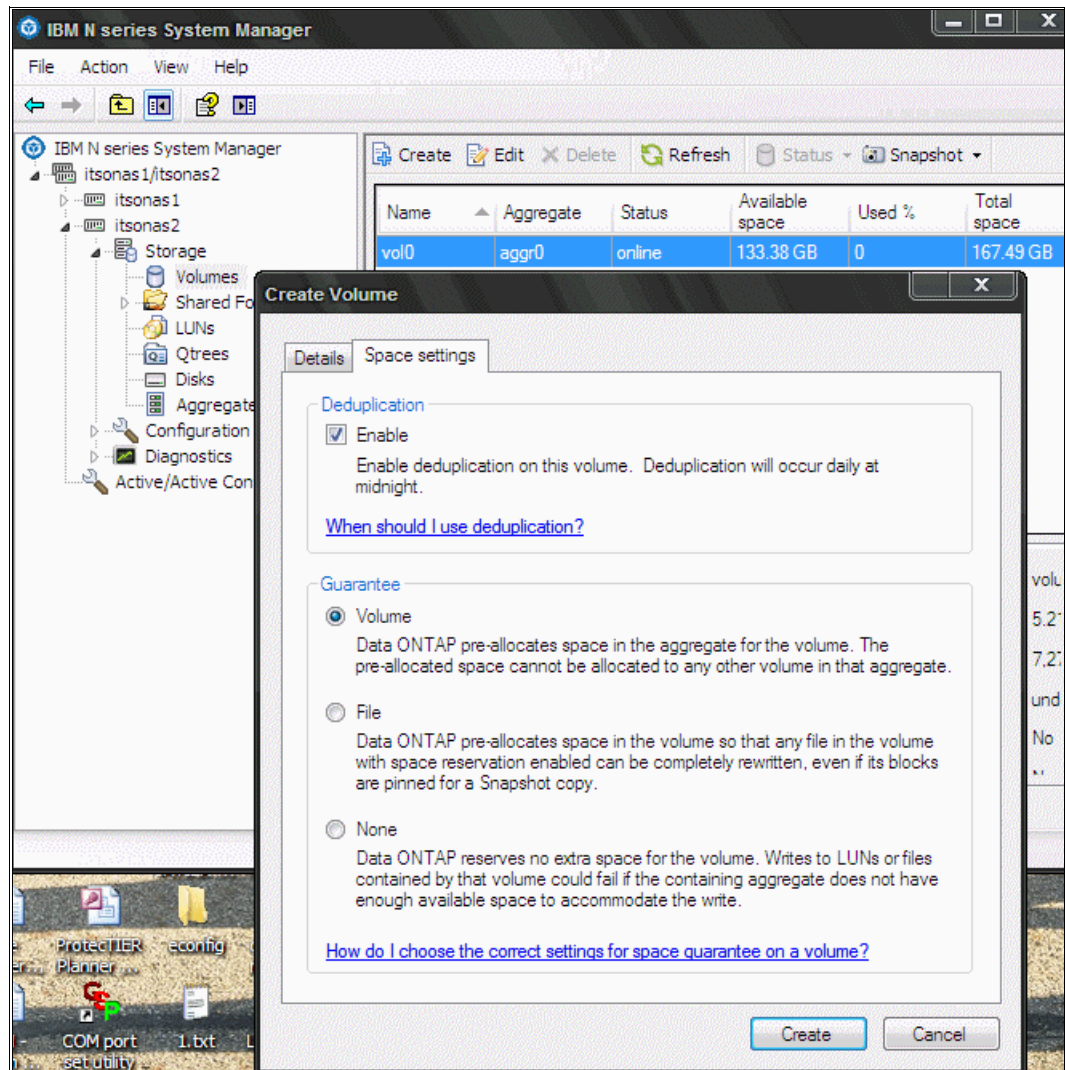


Figure 8-37 Enable Deduplication on the volume created

- After creating the flexible volume, IBM N series System Manager lets you know that it is created and deduplication is enabled, the schedule, the status, and the type of deduplication as shown in Figure 8-38 on page 157



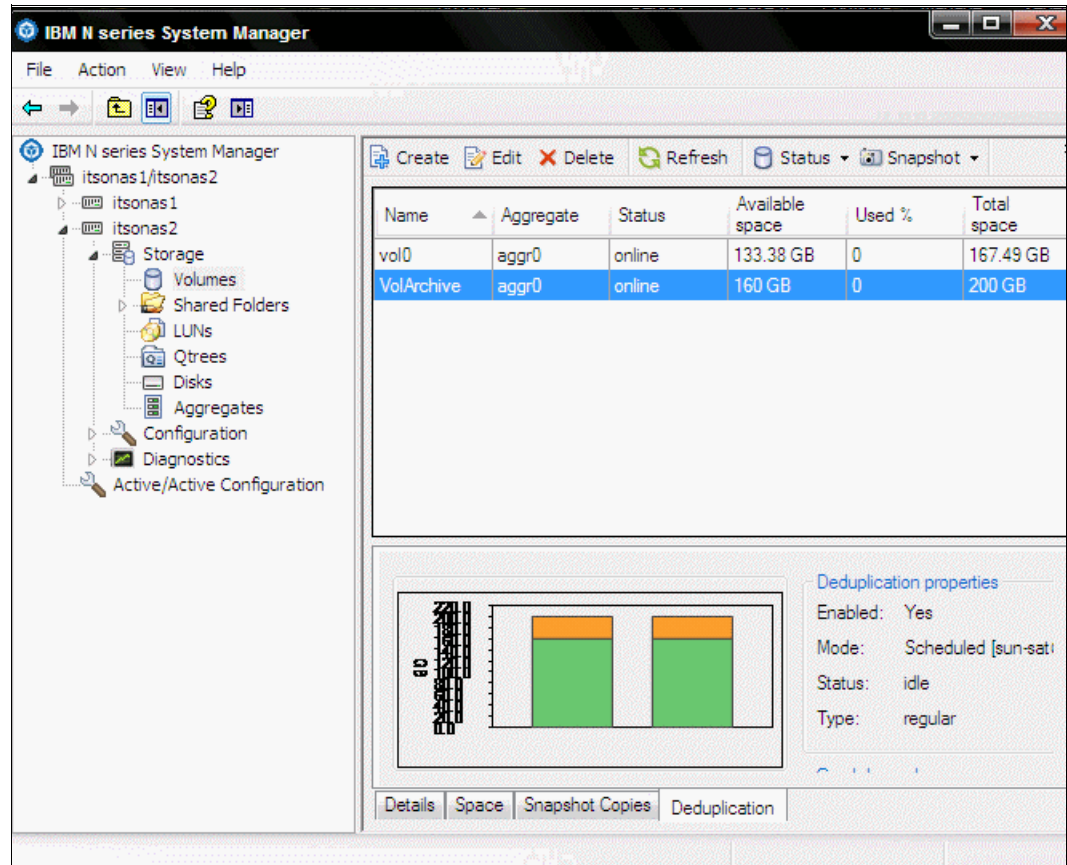


Figure 8-38 Deduplication properties on /vol/VolArchive

4. Click **Apply** and then **OK** to disable the deduplication as shown in Figure 8-39 on page 158.

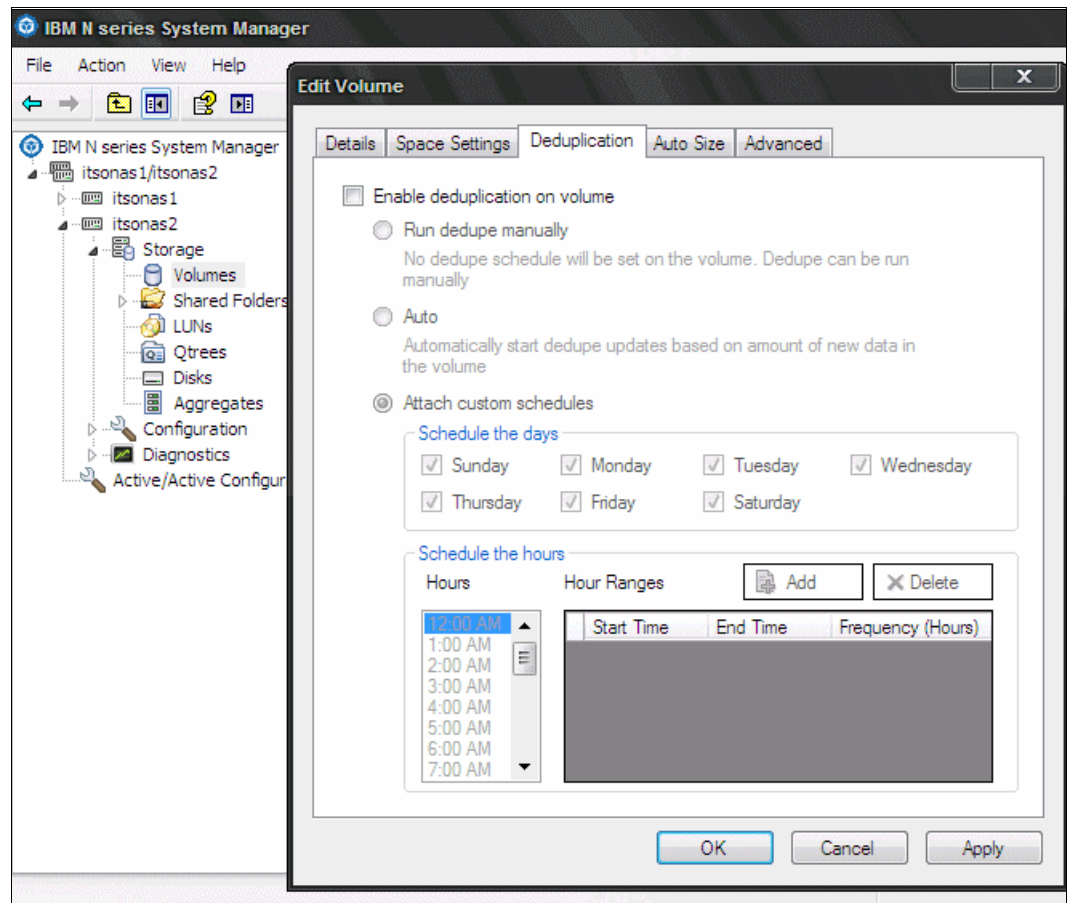


Figure 8-39 Disable the deduplication

5. Figure 8-40 on page 159 shows that the deduplication schedule is disabled.

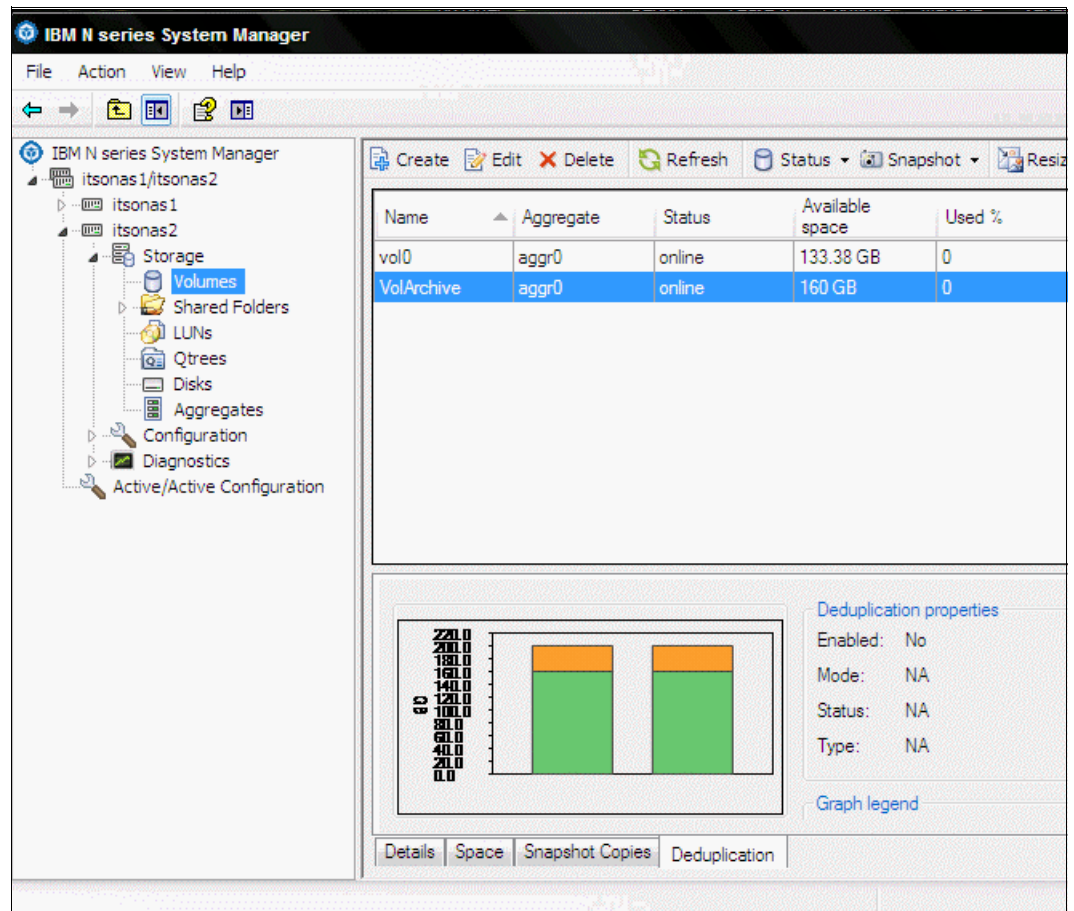


Figure 8-40 Deduplication schedule is disabled

6. Verify the export name of the flexible volume that was created as shown in Figure 8-41 on page 160




Figure 8-41 Verify the exportname of the flexible volume

7. Mount the flexible volume to /testDedup on Linux server, and copy files from the users' directories into the new archive directory flexible volume. From the host perspective, the result is shown in Example 8-10.

*Example 8-10 Result from host perspective using Linux*

```
[root@localhost ~]# mkdir testDedup
[root@localhost ~]# mount itsonas1:/vol/VolArchive/ testDedup
[root@localhost ~]# df -k testDedup
```

Filesystem	kbytes	used	avail	capacity	Mounted on
itsonas1:/vol/VolArchive/	167772160	33388384		134383776	20%
/testArchives					

8. Enable deduplication on /vol/VolArchive by ticking the button as shown in Figure 8-42 on page 161

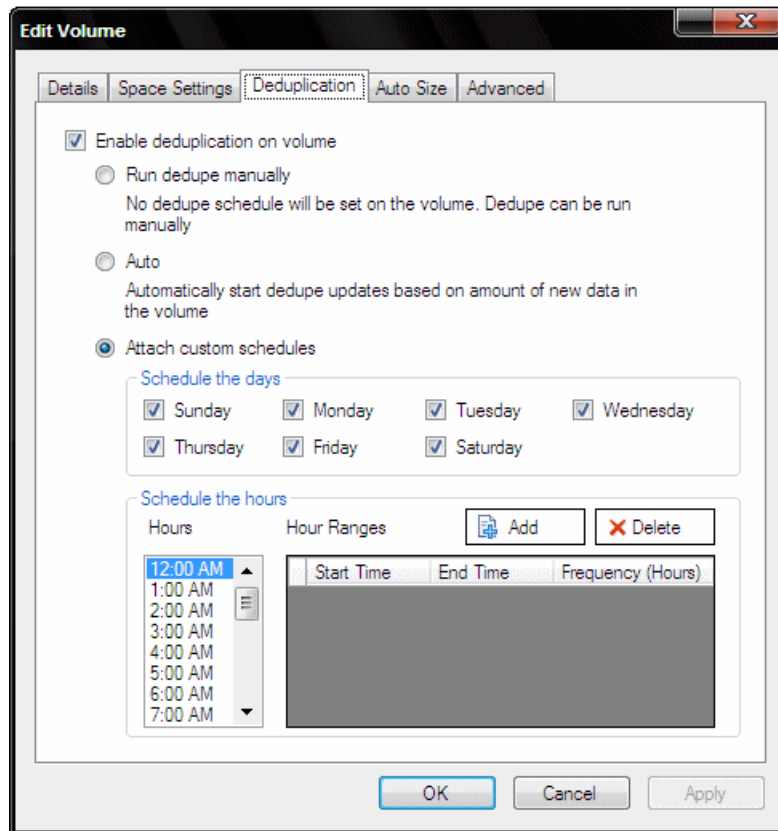


Figure 8-42 Enable deduplication

9. Take note of the percentage used of /vol/Archive as shown in Figure 8-43 on page 162 before deduplication runs.



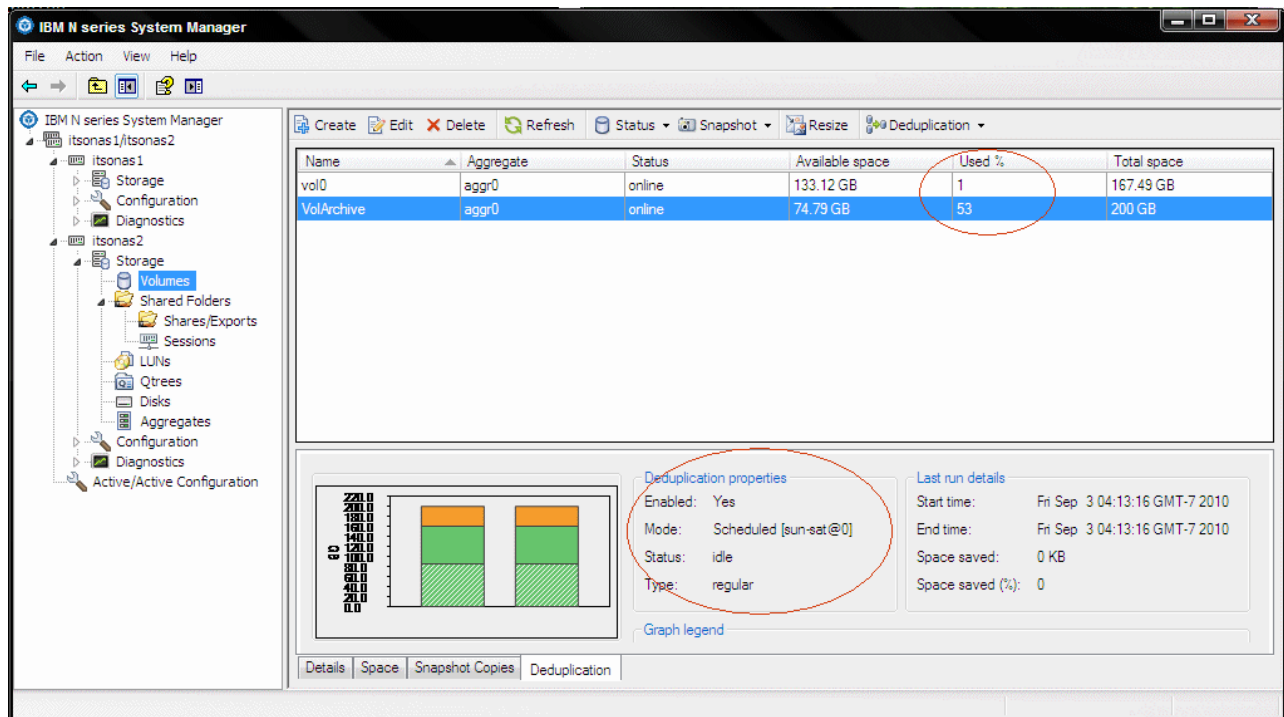


Figure 8-43 Deduplication status of /vol/VolArchive

10. Start deduplication on /vol/VolArchive as shown in Figure 8-44.

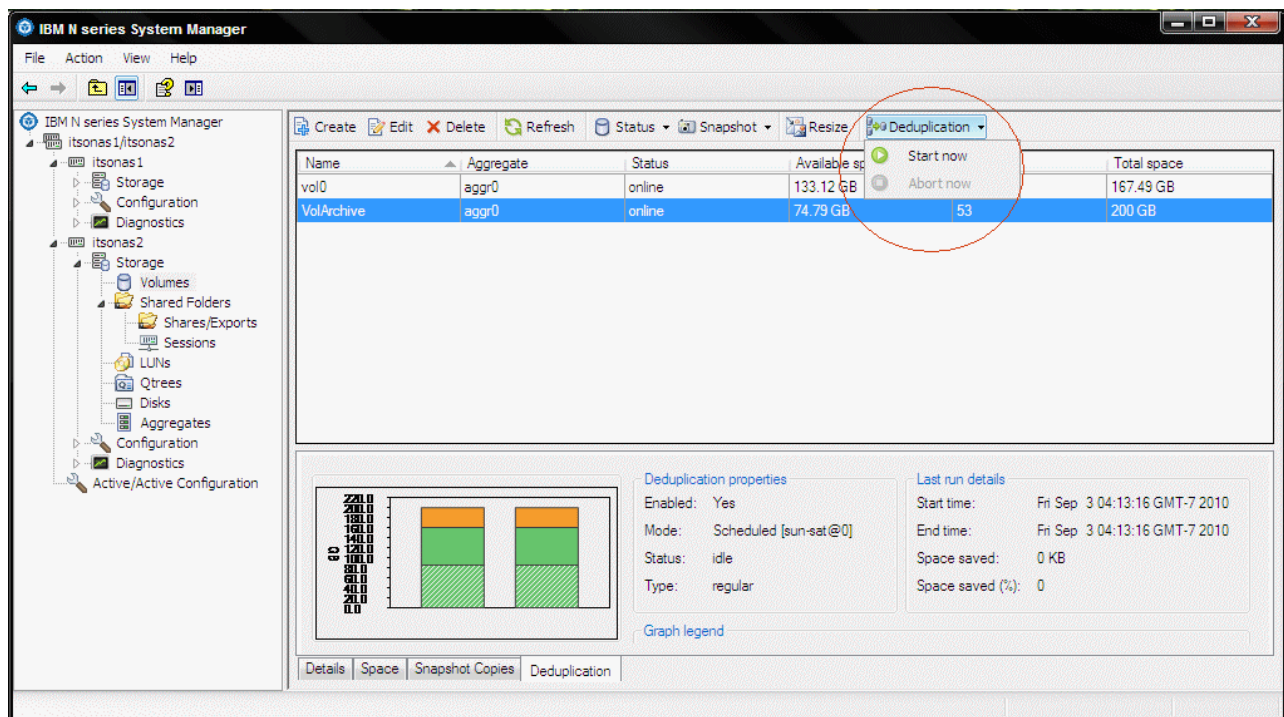


Figure 8-44 Start deduplication on /vol/VolArchive

11. After clicking the **Start** button, it will ask if you want **Partial volume deduplication** or **Full volume deduplication**. Select **Full volume deduplication** for testing purposes as shown in Figure 8-45. Click the **OK** button to start the process.




Figure 8-45 Deduplication choices

12. To verify that Deduplication is running, check the Status of Deduplication properties. It should have a value that changes after you click the **Refresh** button of the console as shown in Figure 8-46.

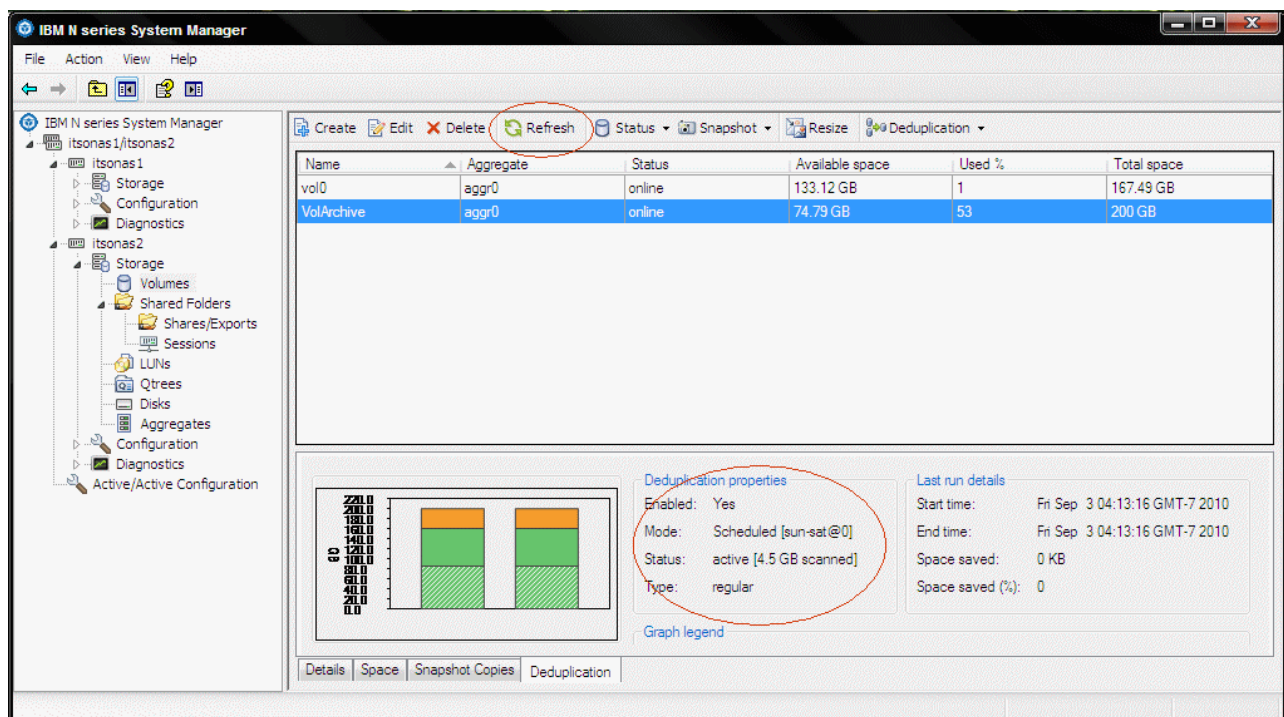


Figure 8-46 Verify Deduplication process is running

13. Figure 8-47 on page 164 shows the deduplication process is complete and shows the space saved after running the process.

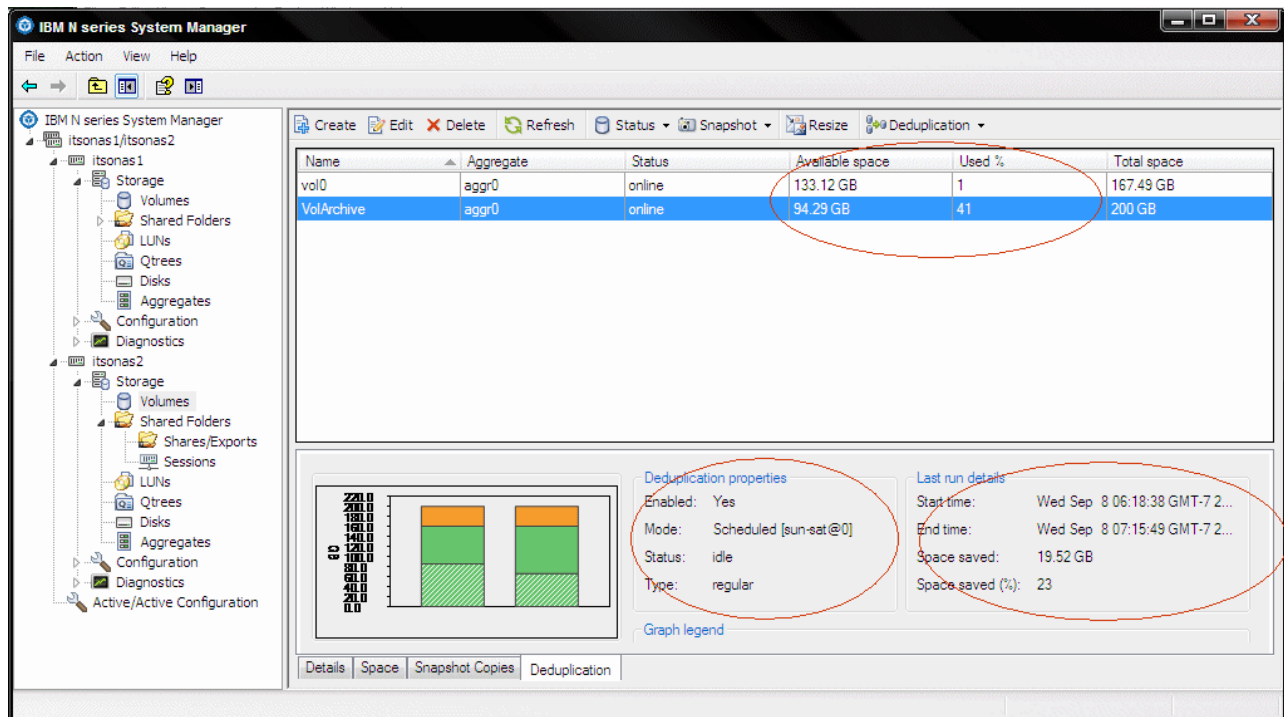


Figure 8-47 Deduplication process done

## 8.5 Configuring deduplication schedules

Generally, you will want to set up a schedule for deduplication so that you do not have to run it manually each time. For more information, see 5.9, “Deduplication scheduling” on page 62. This section provides several specifics about configuring schedules with deduplication using the command line interface and IBM N series System Manager Software.

The **sis config** command is used to configure and view deduplication schedules for flexible volumes. The usage syntax is shown in Example 8-11.

### Example 8-11 Syntax of sis config command

```
itsonas1> sis help config
sis config [ [ -s schedule ] <path> | <path> ... ]
Sets up, modifies, and retrieves the schedule of deduplication volumes.
```

If you run it without any flags, **sis config** returns the schedules for all flexible volumes that have deduplication enabled. Example 8-12 shows the four formats the reported schedules can have.

### Example 8-12 The four format types that reported schedules can have

```
itsonas1> sis config
Path                               Schedule
/vol/dvol_1                        -
/vol/dvol_2                        23@sun-fri
/vol/dvol_3                        auto
/vol/dvol_4                        sat@6
```



The meaning of each of these schedule types is as follows:

- ▶ On flexible volume dvol\_1, deduplication is not scheduled to run.
- ▶ On flexible volume dvol\_2, deduplication is scheduled to run at 11 p.m. every day from Sunday to Friday.
- ▶ On flexible volume dvol\_3, deduplication is set to auto-schedule. This means that deduplication is triggered by the amount of new data written to the flexible volume, specifically when 20% new fingerprints are in the change log.
- ▶ On flexible volume dvol\_4, deduplication is scheduled to run on Saturday at 6 a.m. When the -s option is specified, the command sets up or modifies the schedule on the specified flexible volume. The schedule parameter can be specified in one of four ways, as shown in Example 8-13.

---

*Example 8-13 The schedule parameters*

---

```
[day_list] [@hour_list]
[hour_list] [@day_list]
-
auto
```

---

The meaning of these parameters are as follows:

- ▶ `day_list` specifies which days of the week deduplication should run. It is a comma-separated list of the first three letters of the day: sun, mon, tue, wed, thu, fri, and sat. The names are not case-sensitive. Day ranges such as mon-fri can also be used. The default `day_list` is sun-sat.
- ▶ `hour_list` specifies which hours of the day deduplication should run on each scheduled day. `hour_list` is a comma-separated list of the integers from 0 to 23. Hour ranges such as 8-17 are allowed.  
  
Step values can be used in conjunction with ranges. For example, 0-23/2 means every 2 hours. The default `hour_list` is 0, which is midnight on the morning of each scheduled day.
- ▶ If a hyphen character (-) is specified, there is no scheduled deduplication operation on the flexible volume.
- ▶ The auto schedule causes deduplication to run on that flexible volume when there are 20% new fingerprints in the change log. This check is done in a background process and occurs every hour. Beginning with Data ONTAP 7.3.1, it is configurable by using the **sis config -s auto@num /vol/<vol-name> --** command, where *num* is a two digit number that specifies the percentage.

When deduplication is enabled on a flexible volume for the first time, an initial schedule is assigned to the flexible volume. This initial schedule is sun-sat@0, which means once every day at midnight.

To configure the schedules shown earlier in this section, issue the commands shown in Example 8-14.

---

*Example 8-14 Configure the schedules*

---

```
itsonas1> sis config -s - /vol/dvol_1
itsonas1> sis config -s 23@sun-fri /vol/dvol_2
itsonas1> sis config -s auto /vol/dvol_3
itsonas1> sis config -s sat@6 /vol/dvol_4
```

---

## Using the IBM N series System Manager software

To configure the schedules shown earlier in this section using System Manager, click the Deduplication tab and adjust the schedule as shown in Figure 8-48.

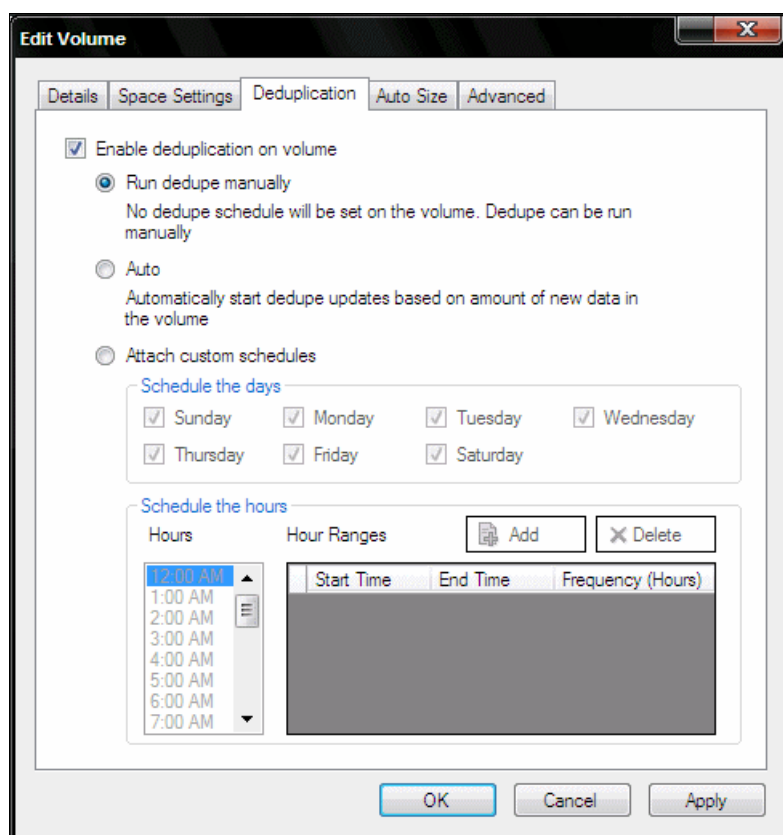


Figure 8-48 Deduplication schedules in System Manager

## 8.6 Sizing for performance and space efficiency

This section discusses the deduplication best practices and behavior that you can expect. Information in this section comes from testing, observations, and knowledge of how deduplication functions.

### 8.6.1 Deduplication general best practices

The following list contains deduplication best practices and lessons learned from internal tests and from deployments in the field:

- ▶ Deduplication consumes system resources and can alter the data layout on a disk. Because of the application's I/O pattern, and the effect of deduplication on the data layout, the read and write I/O performance can vary considerably. The space savings and the performance impact varies significantly depending on the application and the data contents.
- ▶ Performance impact due to deduplication should be carefully considered and measured in a test setup, and you should consider sizing before deploying deduplication in performance-sensitive solutions. For more information about the impact of deduplication refer to Chapter 5, "N series deduplication planning" on page 41.

- ▶ If the amount of new data is small, run deduplication infrequently because there is no benefit to running it frequently in such a case, and it consumes CPU resources. How often you run it depends on the rate of change of the data in the flexible volume.

The more concurrent deduplication processes you are running, the more system resources that are consumed.

Given this information, the best option is to perform one of the following actions:

- Use the auto mode so that deduplication runs only when significant additional data has been written to each particular flexible volume (this approach tends to naturally spread out as deduplication runs).
  - Stagger the deduplication schedule for the flexible volumes so that it runs on alternative days.
  - Run deduplication manually.
- ▶ If Snapshot copies are required, run deduplication before creating the Snapshot to minimize the amount of data before the data gets locked in to the copies. Make sure that deduplication has completed before creating the copy. Creating a Snapshot on a flexible volume before deduplication has a chance to run and complete on that flexible volume can result in lower space savings (see 5.4.4, “Deduplication and Snapshot copies” on page 44).
  - ▶ If Snapshot copies are to be used, the Snapshot reserve should be greater than zero (0). An exception to this might be in an FCP or iSCSI LUN scenario, where it is often set to zero for thin-provisioning reasons.
  - ▶ For deduplication to run properly, you must leave free space for the deduplication metadata (see 5.7.1, “Metadata” on page 60).





# Implementing ProtecTIER

In this chapter we introduce

- ▶ Getting started
- ▶ Installing ProtecTIER Manager
- ▶ Creating repositories (if applicable)
- ▶ Creating Library (if applicable)
- ▶ Adding cartridges (if applicable)
- ▶ Host implementation
- ▶ Backup applications

## 9.1 Getting started

In this section we describe implementation for each model, what activities should have happened already, and what is required for the next steps.

### 9.1.1 TS7610 SMB Appliance

The TS7610 is a customer installable appliance. The following information is required for implementation:

- ▶ IP address for the node
- ▶ System name
- ▶ Timeserver IP (if applicable)
- ▶ Identified workstation to install ProtecTIER Manager on
- ▶ Customer details, including Customer number
- ▶ SNMP Traps information
- ▶ Email alert information
- ▶ SAN Zoning of the Backup Application Host and ProtecTIER (for VTL implementation)

To complete the installation of TS7610, perform the following steps:

- ▶ Configuring ProtecTIER for TS7610, see 9.4, “ProtecTIER configuration on TS7610” on page 181.
- ▶ Installing ProtecTIER Manager, see 9.2, “Installing ProtecTIER Manager” on page 171
- ▶ Finishing configuration in PT Manager, see 9.4.2, “ProtecTIER Manager Configuration Wizard” on page 185.
- ▶ Host connection, see 9.9, “Host implementation” on page 226.

Optional steps:

- ▶ Creating Library (additional), see 9.8, “Virtual library creation” on page 216.

### 9.1.2 TS7650 Appliance

TS7650 is installed by IBM. To complete the installation of TS7650, perform the following steps:

- ▶ Configuring the ProtecTIER server, see “ptconfig for TS7650” on page 197.
- ▶ Installing ProtecTIER Manager, see 9.2, “Installing ProtecTIER Manager” on page 171
- ▶ Creating Library, see 9.8, “Virtual library creation” on page 216.
- ▶ Host connection, see 9.9, “Host implementation” on page 226.

### 9.1.3 TS7650G Gateway

For completing the installation of TS7650G, perform the following steps:

- ▶ ProtecTIER Software install, filesystem creation, see 9.5, “ProtecTIER software install” on page 193
- ▶ Installing ProtecTIER Manager, see 9.2, “Installing ProtecTIER Manager” on page 171
- ▶ Creating repository planning and repository in ProtecTIER, see 9.6, “Repository creating” on page 199.
- ▶ Creating Library, see 9.8, “Virtual library creation” on page 216.
- ▶ Host connection, see 9.9, “Host implementation” on page 226.

## 9.2 Installing ProtecTIER Manager

ProtecTIER Manager is an application that enables you to monitor the status of nodes and two-node clusters in your ProtecTIER system, along with the accompanying repositories and services. ProtecTIER Manager is used to initially configure your ProtecTIER system, and can be used to change the configuration of the system. You must install the ProtecTIER Manager application on one or more workstations.

The ability to have ProtecTIER Manager (or any additional software) installed directly on the ProtecTIER systems is not supported.

If you are installing ProtecTIER Manager on a workstation that has an older version of ProtecTIER Manager already installed, uninstall the older version first.

### 9.2.1 Prerequisites

Before you start with the installation of the ProtecTIER Manager on your workstation, make sure that the following prerequisites are met:

- ▶ One of the following operating systems:
  - Windows 32/64 bit (2003/XP/7)
  - Linux Red Hat 32/64 bit (Red Hat Enterprise 4 or 5)
- ▶ At least 1.2 GB of available disk space
- ▶ At least 256 MB of RAM
- ▶ Access to the ProtecTIER service nodes' IP address (ports 3501 and 3503 are open on the firewall)

In addition, it is recommended that the monitor for ProtecTIER Manager be configured to the following settings:

- ▶ Resolution of 1024 x 768 pixels or higher (this is the minimum resolution supported, but 1280 x 1024 is recommended)
- ▶ 24 bit color or higher

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

### 9.2.2 Installing on Windows XP

To install on Windows XP 32 bit, perform the following steps:

1. Insert the ProtecTIER Manager V2.5 CD into the CD-ROM drive of the system that will run the graphical interface.

**Note:** Always check what level of ProtecTIER Manager is required for the ProtecTIER server. This is included in the Release Notes. New updates of ProtecTIER Manager can be downloaded from IBM ftp site:

<ftp://public.dhe.ibm.com/storage/ProtecTIER/Manager/>

The Release Note is usually a txt file, uploaded with the ProtecTIER server patches:

<ftp://public.dhe.ibm.com/storage/ProtecTIER/Server/>

2. If the autorun process does not launch automatically, select **Start** → **Run**, type *D:* (where *D* is your CD-ROM drive), and press Enter. From the files listed on the CD, select the Windows version. The Introduction window is displayed (Figure 9-1).

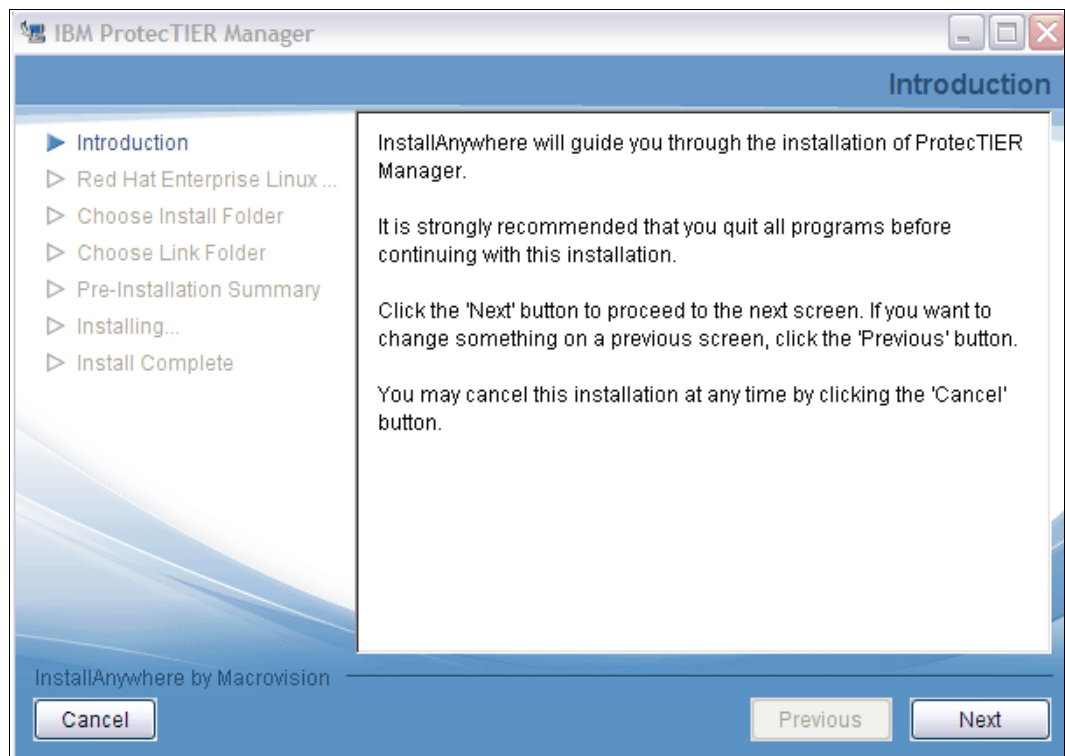


Figure 9-1 ProtecTIER Manager Install: Introduction window

Click **Next**. The Software License Agreement window is displayed (Figure 9-2 on page 173).



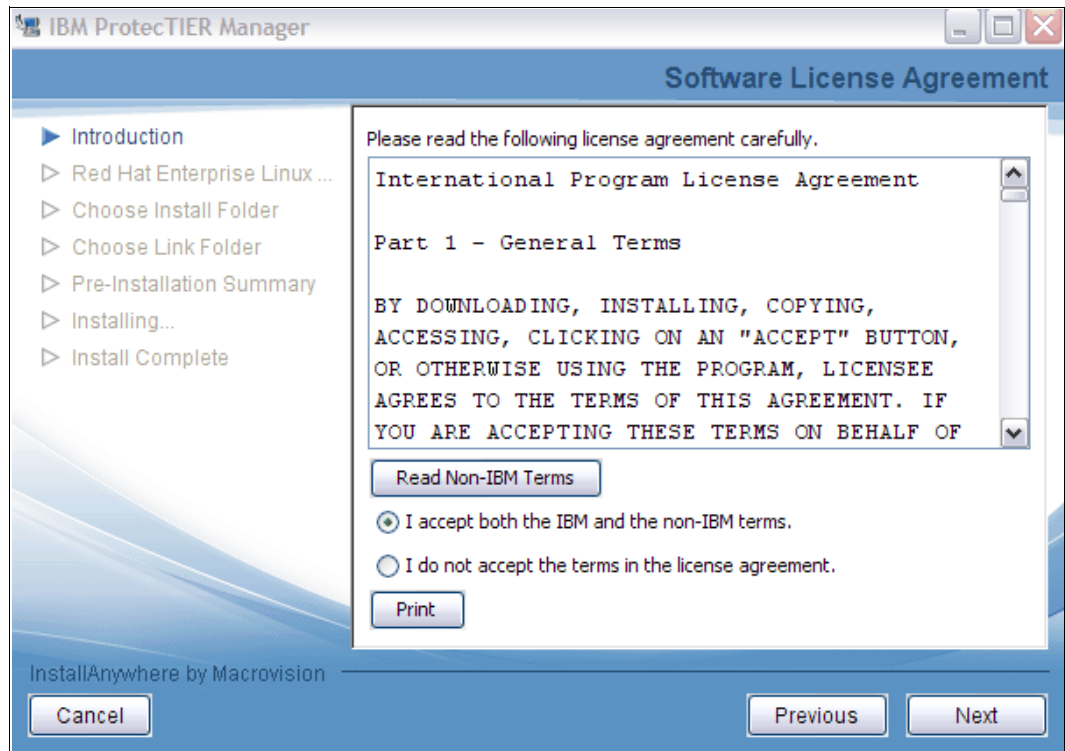


Figure 9-2 ProtecTIER Manager Install: Software License Agreement window

**Note:** You can print the License Agreement by clicking **Print**. If you want to read the non-IBM terms of the license agreement, click **Read Non-IBM Terms** and a window is displayed with the corresponding text.

3. Select **I accept both the IBM and the non IBM-terms** and click **Next**. The Red Hat Enterprise Linux License Agreement window is displayed (Figure 9-3 on page 174).

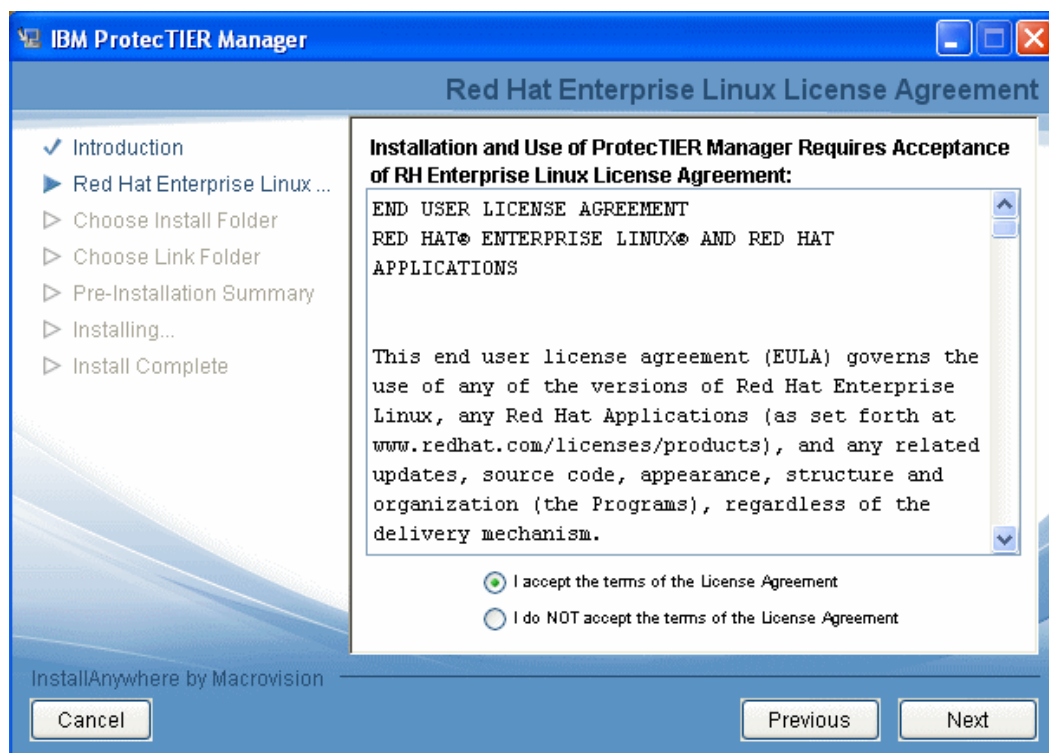


Figure 9-3 ProtecTIER Manager Install: Red Hat Enterprise Linux Licence Agreement window

4. Select **I accept the terms of the License Agreement** and click **Next**. The Choose Install Folder window is displayed (Figure 9-4).

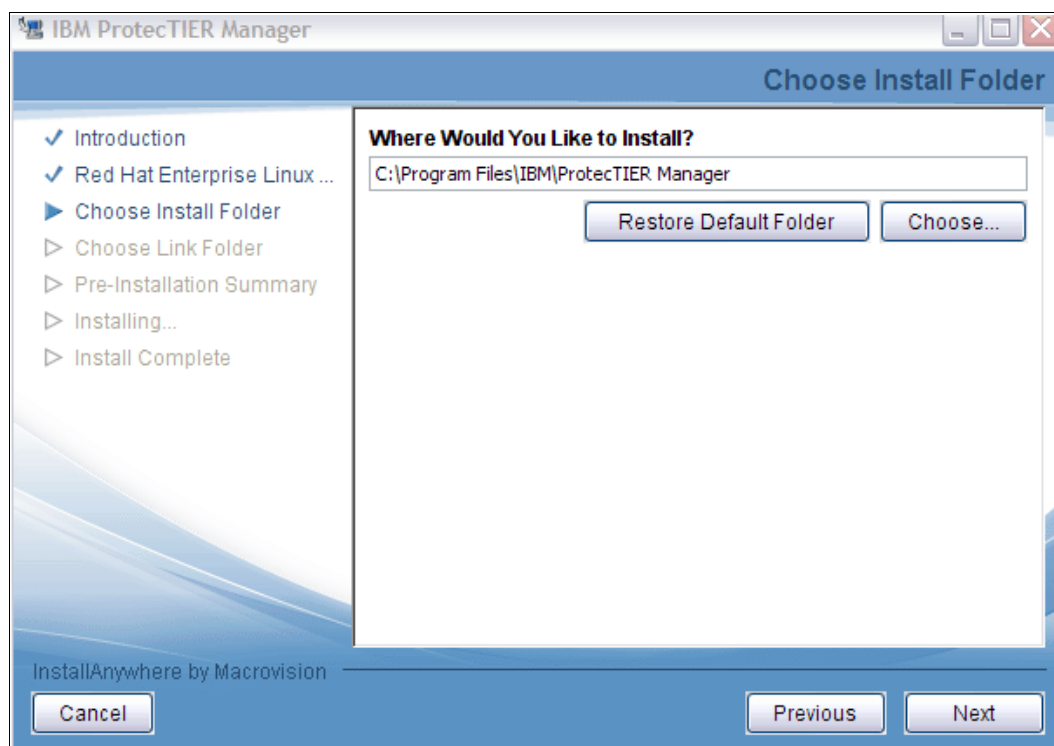


Figure 9-4 ProtecTIER Manager Install: Choose Install Folder window

5. Enter the path where you want to install ProtecTIER Manager or click **Choose** to browse for a location.

**Note:** Click **Restore Default Folder** to revert to the default path.

Click **Next**. The Choose Shortcut Folder window is displayed (Figure 9-5).

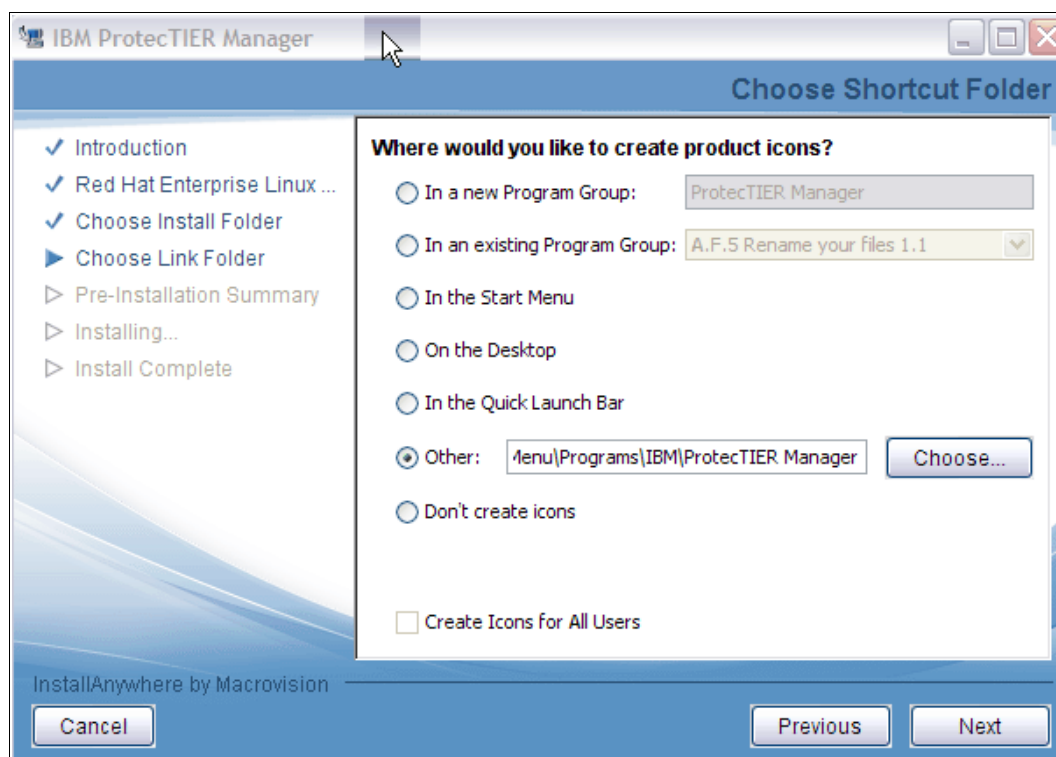


Figure 9-5 ProtecTIER Manager Install: Choose Shortcut Folder window

6. Select one of the following locations for the ProtecTIER Manager shortcut:

<b>In a new Program Group</b>	Creates a new program group in the Program list of the Start menu.
<b>In an existing Program Group</b>	Adds the shortcut to an existing program group in the Program list of the Start menu.
<b>In the Start Menu</b>	Creates shortcuts directly in the Start menu.
<b>On the Desktop</b>	Creates shortcuts on the desktop.
<b>In the Quick Launch Bar</b>	Creates shortcuts in the Quick Launch Bar.
<b>Other</b>	Enables you to enter a path location for the shortcut or to browse for a location by clicking <b>Choose</b> .
<b>Don't create icons</b>	No shortcut icons are created.

You can select **Create Icons for All Users** to create a shortcut in the defined location for all user accounts on the workstation. In our example we used the default **Other**.

Click **Next**. The Pre-Installation Summary window is displayed (Figure 9-6 on page 176).

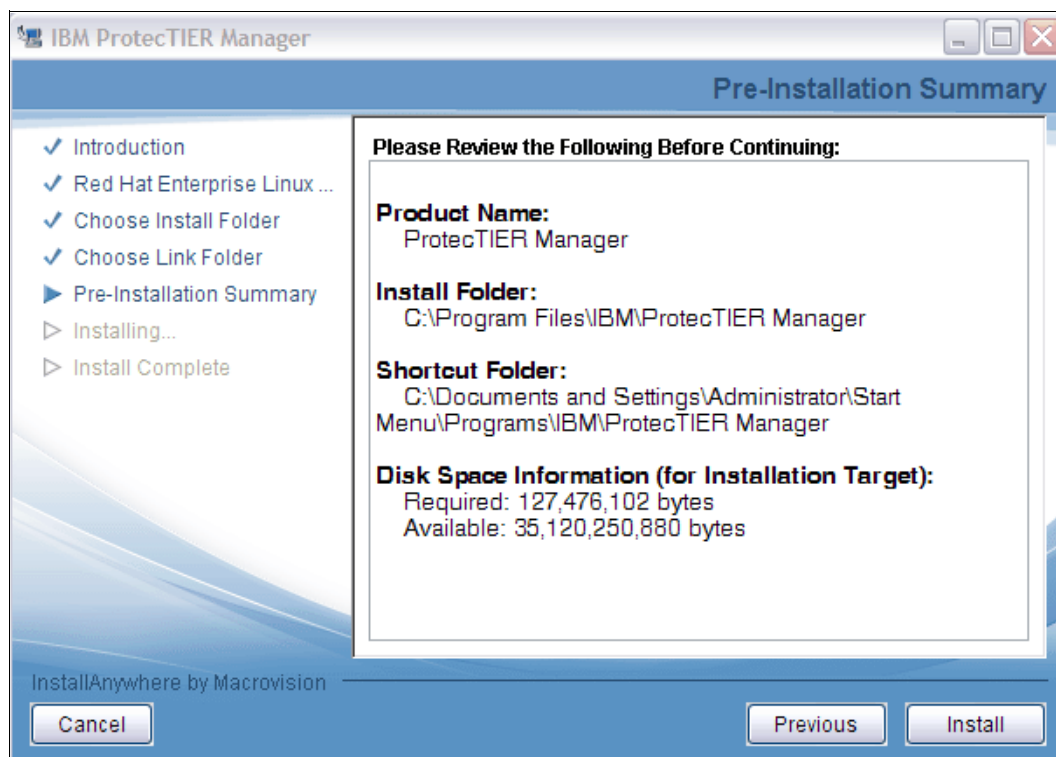


Figure 9-6 ProtecTIER Manager Install: Pre-Installation Summary window

- Click **Install**. The Installing ProtecTIER Manager window is displayed as ProtecTIER Manager is installed on your computer (Figure 9-7).

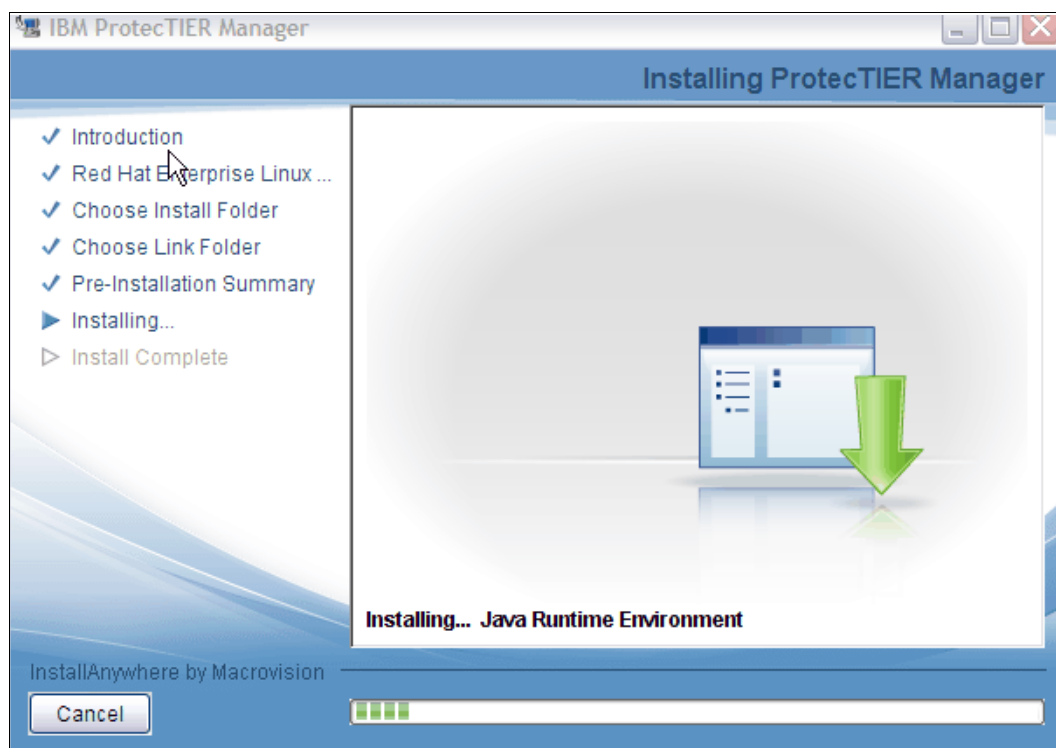


Figure 9-7 ProtecTIER Manager Install: Installing ProtecTIER Manager window

When the installation is complete and ProtecTIER Manager has been successfully installed, the Install Complete window is displayed (Figure 9-8).

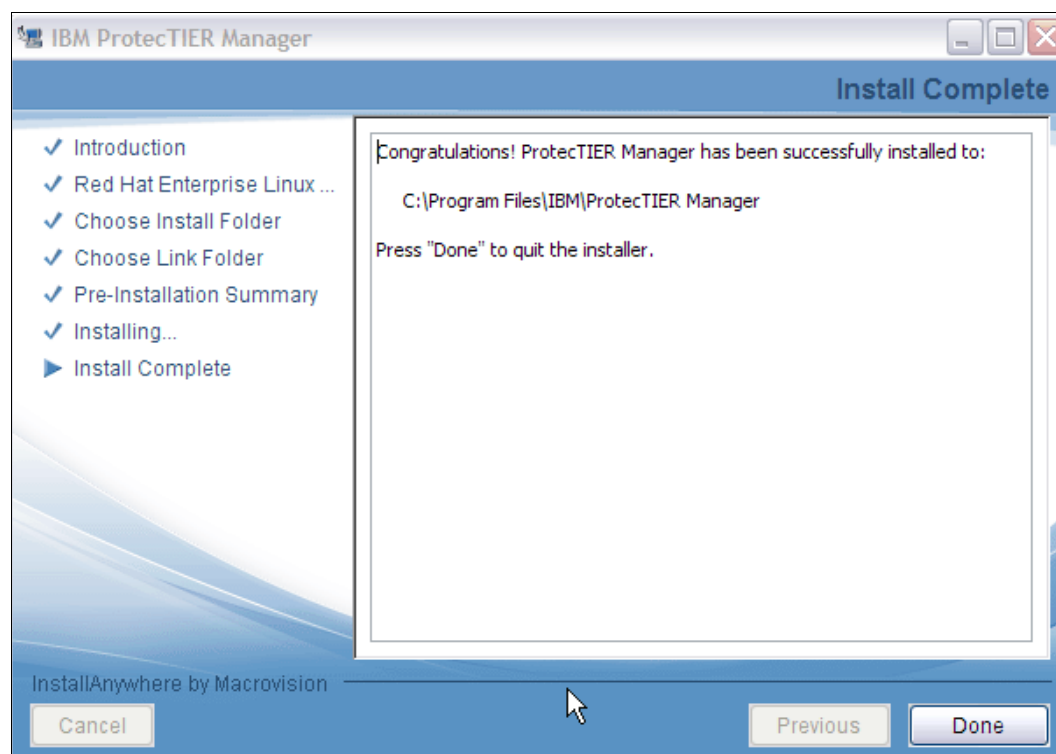


Figure 9-8 ProtecTIER Manager Install: Install Complete window

8. Click **Done**. The ProtecTIER Manager wizard closes.

## 9.3 Starting ProtecTIER Manager

Adding nodes registers the node's IP address and port number with the instance of ProtecTIER Manager at your workstation. Similarly, removing nodes removes the node's registration from ProtecTIER Manager at that workstation.

First you have to add ProtecTIER server nodes to the ProtecTIER Manager, then log into the ProtecTIER Manager. All models of the ProtecTIER family can be accessed by ProtecTIER Manager, given you are at the required level of the GUI.

**Note:** If your ProtecTIER Manager GUI is downlevel compared to the code of the ProtecTIER node, you might not be able to access the node using the GUI. You will have to uninstall and reinstall the new level of GUI. After reinstall, the previously added nodes will show up again, so you will not need to add them again.

### 9.3.1 Adding Nodes to ProtecTIER Manager

After your first install of ProtecTIER Manager, you will see the window shown in Figure 9-9 on page 178 when starting the program.

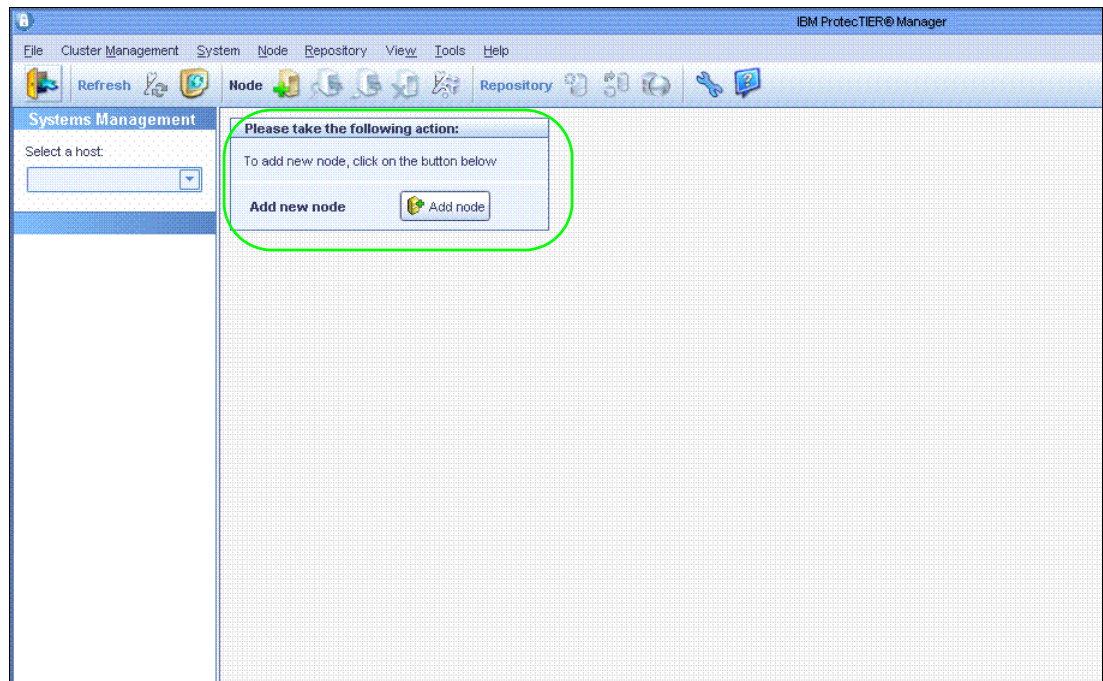


Figure 9-9 Initial window of ProtecTIER Manager

Add nodes using the following steps:

1. Click **Add Node** or click **Node** → **Add Node** if you are adding nodes. See Figure 9-10.

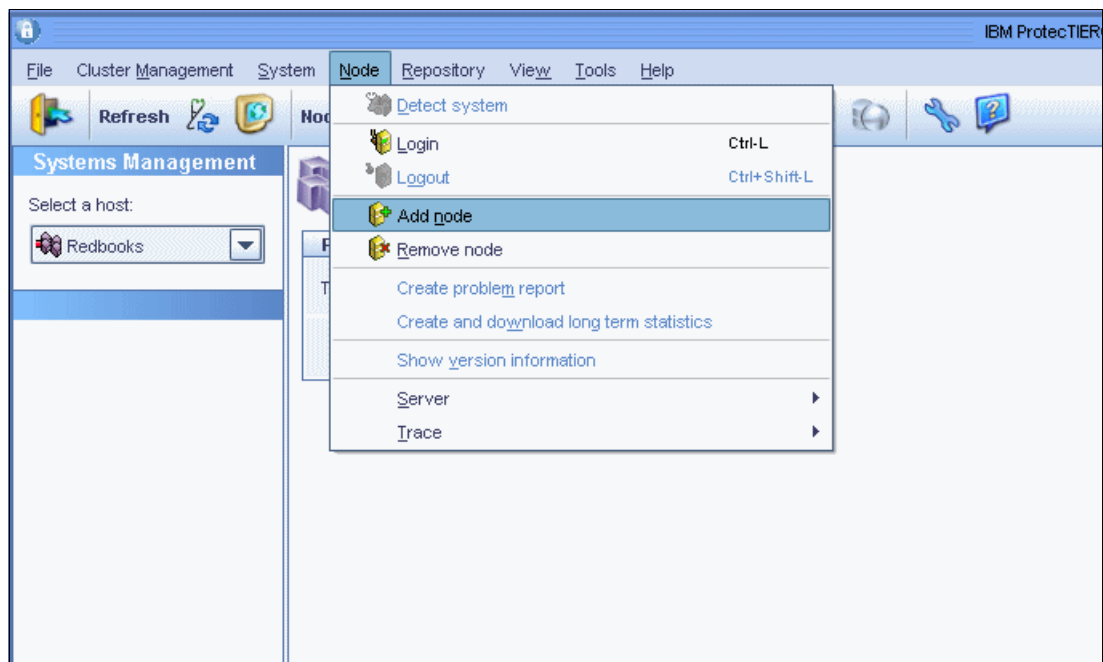


Figure 9-10 ProtecTIER Manager Node drop-down menu

It will ask to add the IP address of the node. See Figure 9-11 on page 179.

**Note:** In case of a dual-node cluster, you can add the IP of any node of the cluster. The system and both nodes will be added to the ProtecTIER Manager.

In case you just did an uninstall and upgrade of the ProtecTIER Manager, you will have all IPs of nodes still listed in ProtecTIER Manager.

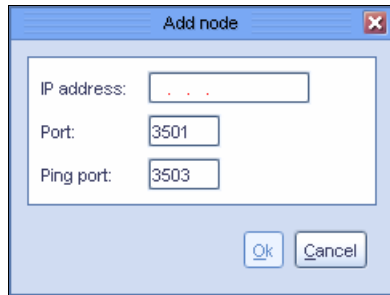


Figure 9-11 ProtecTIER Add node window

2. Enter the IP address of the node and click **OK**. The node is displayed in the Nodes section, and the **Login** button is displayed in the View section.

**Note:** Do not change the port number of the node unless directed by IBM support.

### 9.3.2 Logging in into ProtecTIER Manager

After you added nodes into the ProtecTIER Manager, you can log in.

To log in to each ProtecTIER system that you want to manage using ProtecTIER Manager, perform the following steps:

1. On the navigation pane, select a system or node that you want to log in to (Figure 9-12 on page 180).

**Note:** Even if you selected a single node to log in to, you are logged in to the system after a successful authentication check. The system can have maximum two nodes.

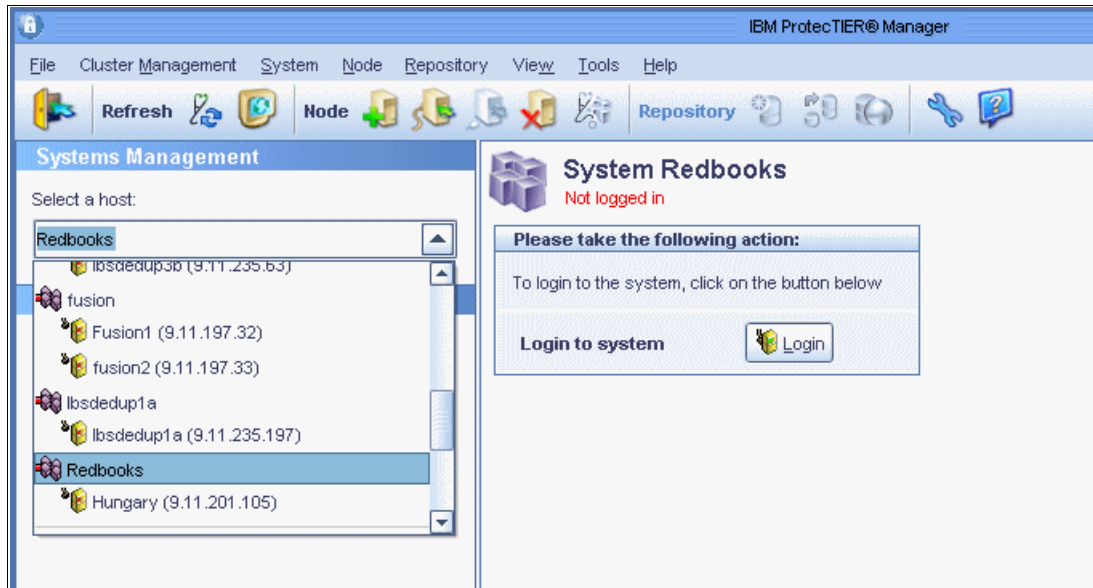


Figure 9-12 ProtecTIER Manager: Select a host view

2. Click **Login**. You are prompted for your user name and password as seen in Figure 9-13.




Figure 9-13 ProtecTIER Manager: Login window

3. Enter your user name and password.
4. Click **OK**.

ProtecTIER Manager has default user accounts corresponding to three user permission levels. See Table 9-1 for the levels, default user names, and passwords. For security, change these default values as soon as possible.

Table 9-1 Default user names and passwords

Permission level	Default user name	Default password
Administrator	ptadmin	ptadmin
Operator	ptoper	ptoper
Monitor	ptuser	ptuser

**Note:** Only one administrator can be logged into a ProtecTIER system at a time.



## 9.4 ProtecTIER configuration on TS7610

After initial HW setup, Time/Date, System Name and Customer Network must be configured through direct console access.

### 9.4.1 ProtecTIER Configuration Menu

Login to system with username ptconfig, and password ptconfig.

The menu in Example 9-1 will be displayed.

*Example 9-1 ProtecTIER configuration menu*

---

ProtecTIER Configuration Menu:

=====

1. Update Time, Date, Timezone & Timeserver(s)
2. Update System Name
3. Update Customer Network
4. Enable/Disable Call Home
5. Activate/Deactivate Replication Manager
6. Update Replication Network
7. Configure Static routes
8. Configure Application Interfaces
9. Restore Network Configuration

q. Quit

Please choose an option:

---

The first three points of the ProtecTIER Configuration Menu has to be done as a minimum setup. We describe this first three menu points in detail in the next sections.

### **Date, time, timezone and timeserver configuration**

If you select 1 in the ProtecTIER Configuration Menu, the menu shown in Example 9-2 displays.

*Example 9-2 Time selections*

---

Date, Time, Timezone & Timeserver(s) configuration

=====

1. Set date & time
2. Set Timezone
3. Set Timeserver(s)

c. Commit changes and exit

q. Exit without committing changes

Please Choose:

---

Here you can setup the date and time, the Timezone, and the Timeserver(s). We recommend this sequence: first Timezone setup, then either date & time, or Timeserver setup.

### ***Set date and time***

After selecting **1**, the Example 9-3 shows the input field for the date and time.

#### *Example 9-3 Setting date and time*

---

Please Choose:1

Please specify the date in DD/MM/YYYY format [16/09/2010]:

Please specify the time in HH:MM:SS format [21:10:42]:

---

### ***Set Timezone***

After selecting **2**, the menu shown in the Example 9-4 displays. Enter the two letter country code, and according to that the time zones will be listed. You can then select the appropriate one.

#### *Example 9-4 Country code selection*

---

Please Choose:2

Enter a 2 letter country code(or type 'm' to enter the timezone manually): us

Time zones under US:

=====

1. America/New\_York
2. America/Detroit
3. America/Kentucky/Louisville
4. America/Kentucky/Monticello
5. America/Indiana/Indianapolis
6. America/Indiana/Vincennes
7. America/Indiana/Winamac
8. America/Indiana/Marengo
9. America/Indiana/Petersburg
10. America/Indiana/Vevay
11. America/Chicago

Press <ENTER> to continue

12. America/Indiana/Tell\_City
13. America/Indiana/Knox
14. America/Menominee
15. America/North\_Dakota/Center
16. America/North\_Dakota/New\_Salem
17. America/Denver
18. America/Boise
19. America/Shiprock
20. America/Phoenix
21. America/Los\_Angeles

Press <ENTER> to continue

22. America/Anchorage
23. America/Juneau
24. America/Yakutat
25. America/Nome
26. America/Adak
27. Pacific/Honolulu

Please choose a timezone:

---

### ***Set timeserver(s)***

After selecting **3**, you can add the IP address of the timeserver(s), as seen in Example 9-5.

#### *Example 9-5 timeserver selection*

---

```
Please Choose:3
Please specify the timeserver's IP Address: 9.11.107.11
Would you like to set a secondary timeserver? (yes|no) yes
Please specify the secondary timeserver's IP Address: 9.11.107.12
```

---

### ***Commit changes and exit***

When you made the required setups, you should commit the changes before exiting as shown in Example 9-6. If you try to exit without committing these changes, you will be asked to confirm exiting without saving the changes.

#### *Example 9-6 Committing changes*

---

```
Date, Time, Timezone & Timeserver(s) configuration
=====
1. Set date & time
2. Set Timezone
3. Set Timeserver(s)

c. Commit changes and exit *
q. Exit without committing changes

Please Choose:c

Please review the following information:
=====
Date: Thu Sep 16 22:23:35 2010

Do you wish to apply those settings? (yes|no) yes
note: the cluster & VTFD services on all nodes must be stopped in order to
continue
Do you wish to continue? (yes|no) yes
Stopping RAS [ Done ]
Stopping VTFD [ Done ]
Stopping Cluster Services [ Done ]
Stopping NTPD [ Done ]
Setting Date & Time [ Done ]
Starting NTPD [ Done ]
Starting cluster [ Done ]
Cluster Started
Starting VTFD [ Done ]
Starting RAS [ Done ]
Press the ENTER key to continue...
```

---

### ***Exit without committing changes***

You can exit without committing changes if you press the q key. If there were changes, it will ask if are you sure to exit without saving. If you say no, you go back to the sub-menu. If you say yes, you will quit and go back to the main-menu. See Example 9-7 on page 184.

#### *Example 9-7 Exiting without changing*

---

Date, Time, Timezone & Timeserver(s) configuration

=====

1. Set date & time
2. Set Timezone
3. Set Timeserver(s)

- c. Commit changes and exit \*
- q. Exit without committing changes

Please Choose:q

Do you wish to exit without committing changes? (yes|no) yes

Press the ENTER key to continue...

---

### **Update System Name**

After choosing “2. Update System Name” from the Main ProtecTIER Configuration Menu, you can change the default system name, to something more meaningful for your environment. See details in Example 9-8.

**Note:** System name is not the node or host name.

#### *Example 9-8 Updating system name*

---

ProtecTIER Configuration Menu:

=====

1. Update Time, Date, Timezone & Timeserver(s)
2. Update System Name
3. Update Customer Network
4. Enable/Disable Call Home
5. Activate/Deactivate Replication Manager
6. Update Replication Network
7. Configure Static routes
8. Configure Application Interfaces
9. Restore Network Configuration

- q. Quit

Please choose an option:2

Starting Cluster, please wait

Starting cluster

[ Done ]

Cluster Started

Please enter a new system name [pt\_system]: Redbooks

Changing system name

[ Done ]

Updated system name successfully

UpdateSystemName ended successfully

Press the ENTER key to continue...

---

### **Update customer network**

The last mandatory step is updating the customer network, which can be started by selecting option 3. See Example 9-9 on page 185.

#### Example 9-9 Entering customer network

---

ProtecTIER Configuration Menu:

=====

1. Update Time, Date, Timezone & Timeserver(s)
2. Update System Name
3. Update Customer Network
4. Enable/Disable Call Home
5. Activate/Deactivate Replication Manager
6. Update Replication Network
7. Configure Static routes
8. Configure Application Interfaces
9. Restore Network Configuration

q. Quit

Please choose an option:3

Starting Cluster, please wait

Starting cluster

[ Done ]

Cluster Started

Would you like to stop the VTFD service? (yes|no) yes

Stopping RAS

[ Done ]

Stopping VTFD

[ Done ]

Please provide the following information:

-----

Customer Network, IP Address [9.11.201.105]:

Customer Network, Netmask [255.255.254.0]:

Customer Network, Default Gateway [9.11.200.1]:

Customer Network, Hostname [Hungary]:

Configuring Network

[ Done ]

Setting Hostname

[ Done ]

Saving configuration

[ Done ]

Collecting RAS Persistent configuration

[ Done ]

Updated network configuration successfully

Starting VTFD

[ Done ]

Starting RAS

[ Done ]

UpdateNetwork ended successfully

Press the ENTER key to continue...

---

**Note:** Do not specify leading zeros in any of the address numbers (e.g. 192.168.001.015).

## 9.4.2 ProtecTIER Manager Configuration Wizard

To complete the configuration after the console-based server configuration steps, connect to the appliance using ProtecTIER Manager and run the Configuration Wizard as shown in Figure 9-14 on page 186.

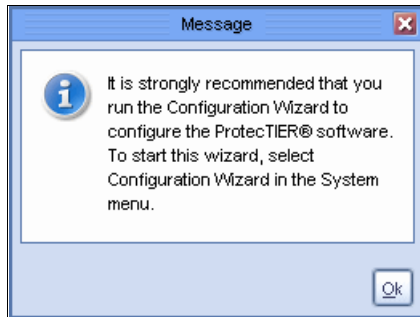


Figure 9-14 Message for running Configuration Wizard

1. To start the Configuration Wizard manually, click **Server** → **Configuration wizard**. See Figure 9-15.

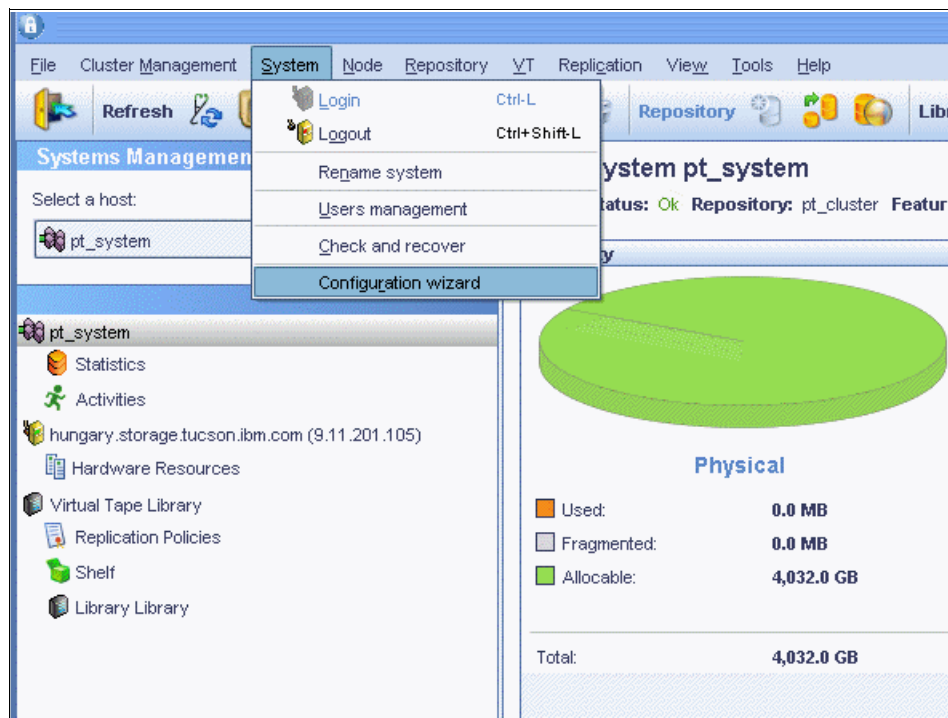


Figure 9-15 ProtecTIER Server drop-down menu

In the Configuration Wizard you will be able to set up:

- Customer information
- Enable SNMP traps
- Provide details for SNMP traps
- Enable E-mail alerts
- Provide details about E-mail alerts

The Configuration wizard starts with a Welcome window as shown in Figure 9-16 on page 187.

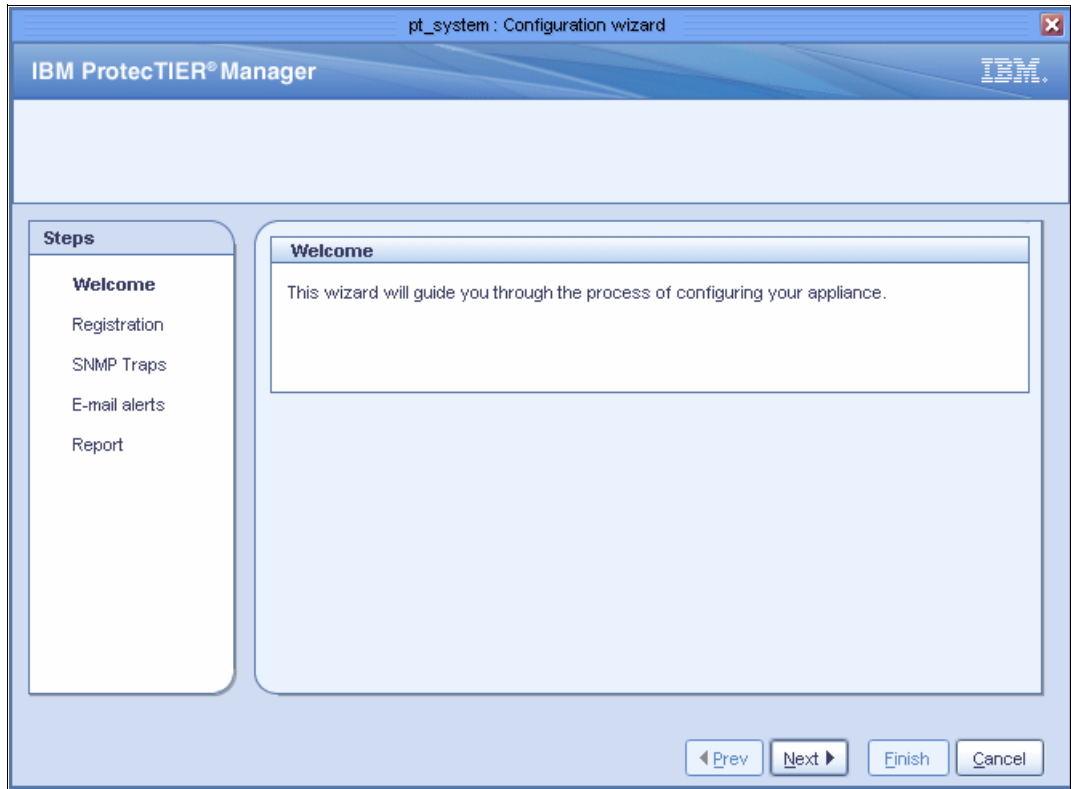


Figure 9-16 Configuration wizard: Welcome window

2. After clicking **Next**, you see the Registration window as shown in Figure 9-17.

Figure 9-17 Configuration wizard: Registration information window

3. After filling it out with your details, you click the **Next** button as seen in Figure 9-18.

The screenshot shows the 'Registration information' window of the IBM ProtecTIER Manager Configuration wizard. The window has a title bar 'Redbooks : Configuration wizard' and a header 'IBM ProtecTIER Manager' with the IBM logo. On the left, a 'Steps' panel lists 'Welcome', 'Registration' (selected), 'SNMP Traps', 'E-mail alerts', and 'Report'. The main area is divided into two sections: 'Company information' and 'System administration information'. The 'Company information' section contains fields for Name (IBM Corporation), Street (1 Street), City (City), State (State), Postal code (11111), and Customer number (1234567). The 'System administration information' section contains fields for Name (SysAdmin), Phone (555-123-456), and E-mail (sysad@email.com). At the bottom right, there are four buttons: 'Prev', 'Next', 'Finish', and 'Cancel'.

Company information	
Name:	IBM Corporation
Street:	1 Street
City:	City
State:	State
Postal code:	11111
Customer number:	1234567
Country code:	US

System administration information	
Name:	SysAdmin
Phone:	555-123-456
E-mail:	sysad@email.com

Figure 9-18 Configuration wizard: Registration information window

4. The next dialog is the SNMP trap enablement as shown in Figure 9-19 on page 189.



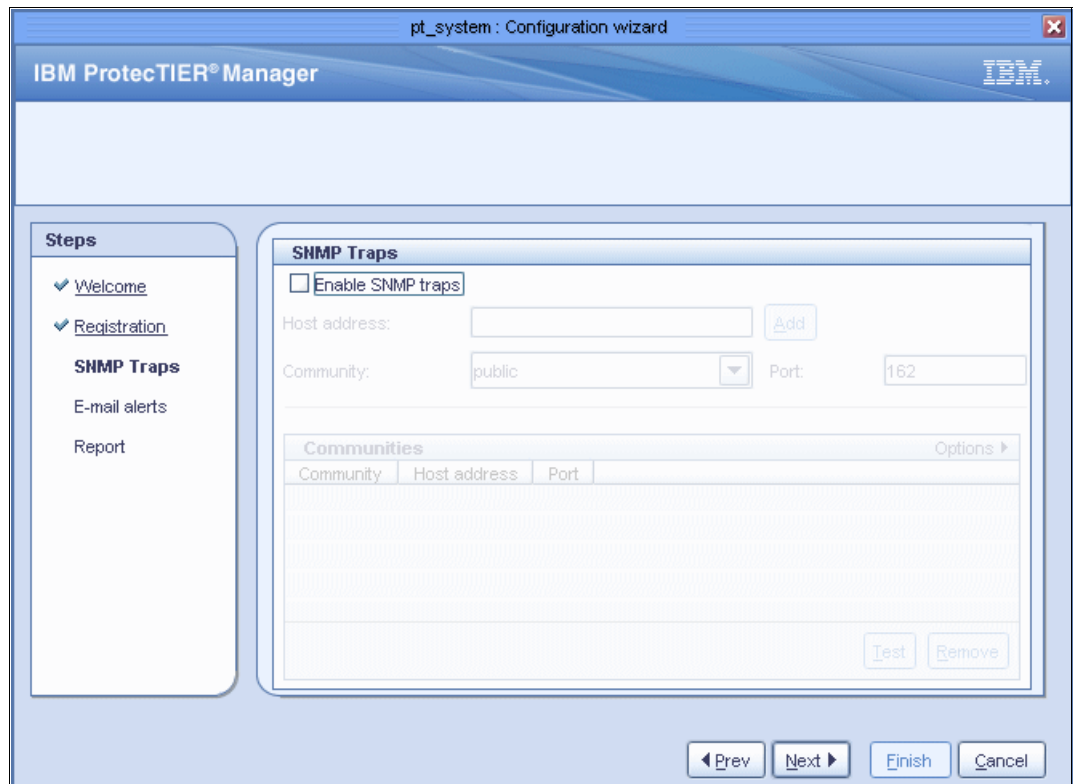


Figure 9-19 Configuration wizard: Enabling SNMP traps window

5. If you enable SNMP traps, you can fill out more details as shown in Figure 9-20 on page 190.

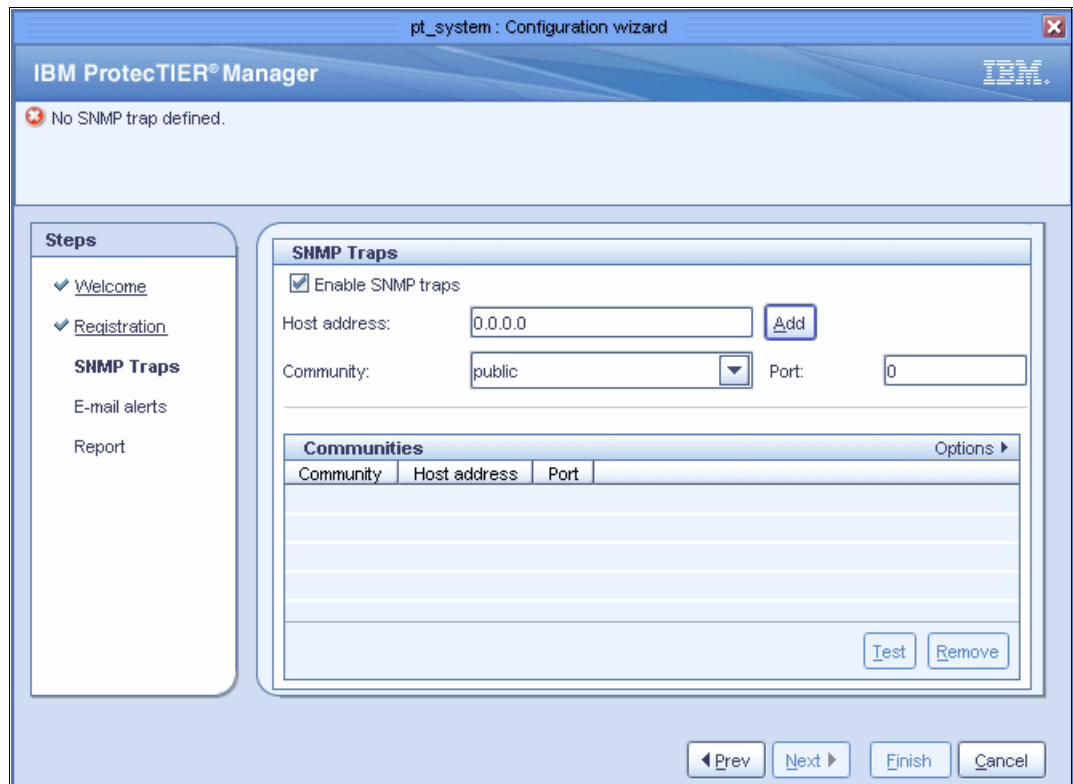


Figure 9-20 Configuration wizard: SNMP traps setup window -adding host

6. You can add multiple SNMP hosts. After clicking **Add**, the host will be listed at the Communities section. See Figure 9-21 on page 191.

pt\_system : Configuration wizard

IBM ProtectTIER® Manager

The host address can not be empty, in order to add a new SNMP trap.

**Steps**

- ✓ Welcome
- ✓ Registration
- SNMP Traps**
- E-mail alerts
- Report

**SNMP Traps**

☒ Enable SNMP traps

Host address:

Community:  Port:

Communities			Options ▶
Community	Host address	Port	
public	0.0.0.0	0	

◀ Prev Next ▶ Finish Cancel

Figure 9-21 Configuration wizard: SNMP traps setup window

7. After clicking **Next**, you can enable email notifications. See Figure 9-22.

pt\_system : Configuration wizard

IBM ProtectTIER® Manager

**Steps**

- ✓ Welcome
- ✓ Registration
- ✓ SNMP Traps
- E-mail alerts**
- Report

**Configure outgoing E-mail server**

☐ Enable E-mail alerts

Host address:  Port:

**Recipients list**

New mail recipient:

Recipients		Options ▶
E-mail address		

◀ Prev Next ▶ Finish Cancel

Figure 9-22 Configuration wizard: E-mail alert enabling window

8. If you enable email alerts, you can add email server IP and email recipients. See Figure 9-23.

The screenshot shows the 'IBM ProtecTIER® Manager' configuration wizard window. The title bar reads 'pt\_system : Configuration wizard'. The main header area displays 'IBM ProtecTIER® Manager' and the IBM logo. Below the header, a message states 'No mail recipients defined.' with a red error icon. On the left, a 'Steps' sidebar lists 'Welcome', 'Registration', 'SNMP Traps', 'E-mail alerts' (which is selected), and 'Report'. The main content area is titled 'Configure outgoing E-mail server'. It includes a checkbox for 'Enable E-mail alerts' which is checked. Below this, there are input fields for 'Host address' (containing '0.0.0.0') and 'Port' (containing '25'). A section titled 'Recipients list' contains a 'New mail recipient' field with 'recipient@email.com' and an 'Add' button. Below this is a table with the header 'Recipients' and a sub-header 'E-mail address'. The table is currently empty. To the right of the table header is an 'Options' dropdown arrow. At the bottom right of the recipients section are 'Test' and 'Remove' buttons. At the very bottom of the window are navigation buttons: 'Prev', 'Next', 'Finish', and 'Cancel'.

Figure 9-23 Configuration wizard: E-mail alert configuring window

9. When you finished, click **Next**. The Summary report window summarizing your changes displays. See Figure 9-24 on page 193.

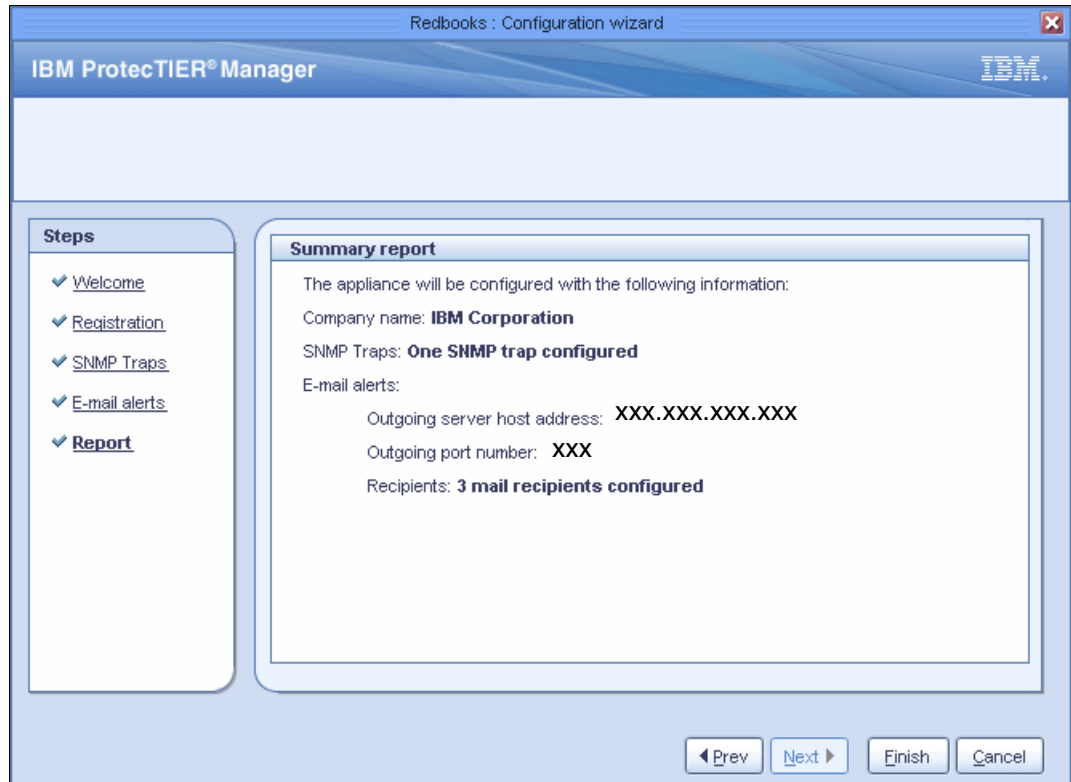


Figure 9-24 Configuration wizard: Summary report

10. If you are satisfied with the Summary, you can hit **Finish**, and your changes will be added to the system.

Your system is ready for host attachment. You can continue at 9.9.1, “Connecting hosts to ProtecTIER systems” on page 226.

## 9.5 ProtecTIER software install

New software updates are downloadable from here:

<ftp://public.dhe.ibm.com/storage/ProtecTIER/>

You have to go to the appropriate directory: Server for TS7650 and TS7650G, TS7610 for SMB appliance and TS7680 for Enterprise Gateway for System z. Look for the latest update, each of which contains a Release Note. The Release Note is usually a text file, and contains details about how to install the package for each model.

### 9.5.1 autorun

To install the software, in most cases you have to copy the tar package to the ProtecTIER server, untar it, and running **autorun**. In some cases it is **autorun -f**. Follow the details in the Release Note.

*Example 9-10 Installing software*

```
[root@tuscan y ~]# cd /Install/PT_TS7650G_V2.5.0.TST_117_3-full.x86_64
```

```
[root@tuscan y PT_TS7650G_V2.5.0.TST_117_3-full.x86_64]# ./autorun -f
```

---

To check the installed level you can do one of the following:

- Run the **get\_versions** command, as seen in Example 9-11:

*Example 9-11 get\_versions command*

---

```
[root@tuscan y ~]# /opt/dtc/app/sbin/get_versions
<?xml version="1.0" encoding="UTF-8" ?>
<version-info>
  <component name="PT" version="2.5.0.1261"/>
    <component name="ptlinux" version="7123.126-1261"/>
    <component name="dtcemulex" version="5223.004-0"/>
    <component name="ptrepmgr" version="6123.069-0"/>

    <component name="vtl" version="7.123.126"/>
    <component name="model" version="TS7650G"/>
</version-info>
```

---

- Check it from rasMenu, as seen in Example 9-12:

*Example 9-12 RAS Menu*

---

```
[root@tuscan y ~]# rasMenu
+-----+
| RAS Text Based Menu running on tuscan y |
+-----+
|  1) Check if RAS service is running    |
|  2) Run RAS environment check          |
|  3) Start RAS service                  |
|  4) Stop RAS service                   |
|  5) Get RAS Version                    |
|  6) Get PT Code Version                 |
|  7) Display Firmware Levels             |
|  8) Manage Configuration (...)          |
|  9) System Health Monitoring (...)      |
| 10) Problem management (...)            |
| 11) Call Home Commands (...)            |
| 12) Collect Logs (...)                  |
| 13) Enterprise Controller (...)         |
|  E) Exit                               |
+-----+
>>> Your choice? 6
Begin Processing Procedure
PT version      : 2.5.0.1261
Build date     : Oct_04_2010
PVT main package : ptlinux-7123.126-1261
DTC Emulex driver : dtcemulex-5223.004-0
vtl version    : 7123.126

End Processing Procedure
Press any key to continue
```

---

- Check it from the GUI. Click the Version number in the GUI, and you will see the Show version information window, as seen on Figure 9-25 on page 195.

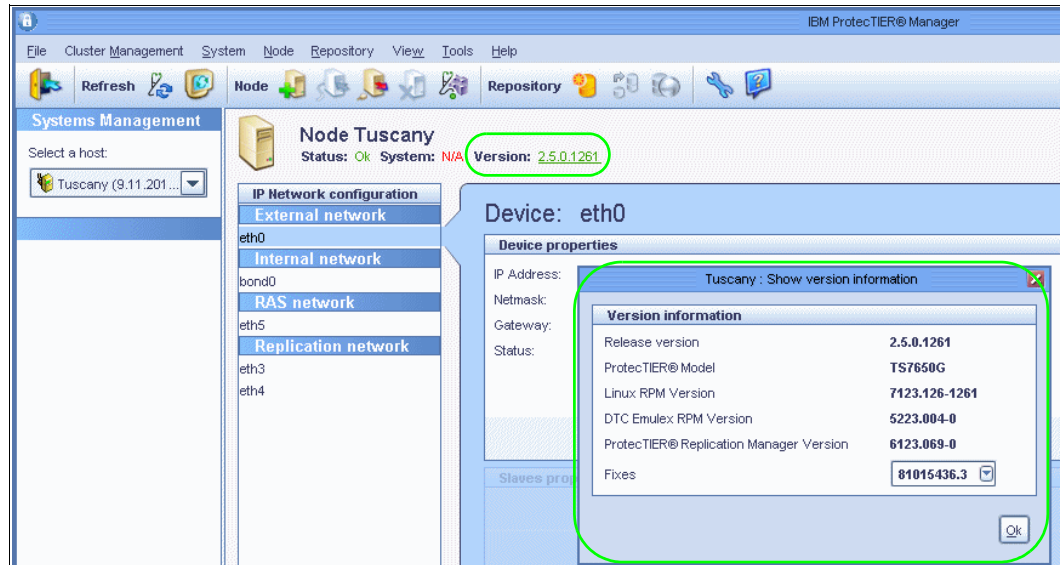


Figure 9-25 ProtecTIER software version check

## 9.5.2 ptconfig

The ptconfig script has many functions. The utility can be found in /opt/dtc/install directory as shown in Example 9-13.

### Example 9-13 ptconfig script

```
[root@tuscany ~]# cd /opt/dtc/install/
[root@tuscany install]# ./ptconfig
Main utility for installing and maintaining a ProtecTIER system
```

#### Usage:

```
./ptconfig -modelList
./ptconfig -restoreNetworkConfig
./ptconfig -updateFirmwares
./ptconfig -setClock
./ptconfig -updateReplicationIp
./ptconfig -addReplication
./ptconfig -updateSystemName
./ptconfig -activatePTRepMan
./ptconfig -upgrade -model=<model> [-usedefaults]
./ptconfig -deactivatePTRepMan
./ptconfig -configRAS
./ptconfig -validate
./ptconfig -addReplicationNic
./ptconfig -updateNetwork
./ptconfig -install -app=<app> -model=<model> [-usedefaults]
./ptconfig -appInterfaces
./ptconfig -replace -app=<app> -model=<model> [-usedefaults]
./ptconfig -appList
./ptconfig -updateDSfw
./ptconfig -staticRoutes
./ptconfig -restoreSerialNumber
./ptconfig -modifyNic
```

-modelList	list the available models (*)
-restoreNetworkConfig	Restores network configuration (*)
-updateFirmwares	Update xSeries firmwares (*)
-setClock	Manage the time, date, timeserver(s) and timezone settings (*)
-model	specify model type
-updateReplicationIp	update ip addresses used by replication (*)
-addReplication	add IP configuration for replication (*)
-updateSystemName	replace the system name (cluster name) (*)
-usedefaults	don't ask the user for configuration attributes
-activatePTRepMan	Install the ProtecTIER Replication Manager (*)
-upgrade	upgrade an existing model to a different model (*)
-deactivatePTRepMan	Uninstall the ProtecTIER Replication Manager (*)
-configRAS	Configure the RAS service (*)
-validate	run validation test on an installed cluster (*)
-app	specify application to install (use the appList option to see all available applications)
-addReplicationNic	replace replication NIC location after upgrading DD1 to DD3 (*)
-updateNetwork hostname (*)	replace the external (management) ip address and/or hostname (*)
-install (*)	install and configure a PT system with a specific model (*)
-appInterfaces	Configure Application Interfaces (*)
-replace repository (*)	install and configure a PT system with an existing repository (*)
-appList	list the available applications (*)
-updateDSfw	Update the DS disk array firmware (*)
-staticRoutes	edit static routes configuration (*)
-restoreSerialNumber	Restore the serial number for the server (*)
-modifyNic	change Network interface card configuration (*)

Options marked with a '\*' are mutually exclusive

---

In this example, we use the **ptconfig -install -model=TS7650G -app=VTL** command to install a TS7650G Gateway model with VTL as shown in Example 9-14.

*Example 9-14 installation example*

---

```
[root@tuscanyc install]# ./ptconfig -modelList
Available models:
=====
                TS7650G                TS7650G Gateway

[root@tuscanyc install]# ./ptconfig -appList
Available applications:
=====
                OST_10G                OST using 10G NICs
                VTL_OLD                VTL Application on M2
                OST_1G                OST using 1G NICs
                VTL                  VTL Application

[root@tuscanyc install]# ./ptconfig -install -model=TS7650G -app=VTL
```

---



For each software update you can list the available models.

**Note:** Do not specify leading zeros in any of the address numbers (e.g. 192.168.001.015).

**Note:** If your system is OST, you have to run the **ptconfig -install** differently, for example:

```
[root@tuscanys install]# ./ptconfig -install -model=TS7650G -app=OST_1G
```

## ptconfig for TS7650

If you have a TS7650 appliance, you have to run different option under **ptconfig**. Because TS7650 is an appliance with storage, everything is preconfigured in the manufacturing. You will have to run the commands shown in Example 9-15 to configure it for your environment.

### Example 9-15 Appliance install

```
[root@italy ~]# cd /opt/dtc/install/
[root@italy install]# ./ptconfig
Main utility for installing and maintaining a ProtecTIER system
```

#### Usage:

```
./ptconfig -modelList
./ptconfig -restoreNetworkConfig [-usedefaults]
./ptconfig -updateFirmwares
./ptconfig -setClock
./ptconfig -updateReplicationIp
./ptconfig -addReplication
./ptconfig -updateSystemName
./ptconfig -activatePTRepMan
./ptconfig -upgrade -model=<model> [-usedefaults]
./ptconfig -deactivatePTRepMan
./ptconfig -configRAS
./ptconfig -validate
./ptconfig -addReplicationNic
./ptconfig -updateNetwork
./ptconfig -install -app=<app> -model=<model> [-usedefaults]
./ptconfig -appInterfaces
./ptconfig -replace -app=<app> -model=<model> [-usedefaults]
./ptconfig -appList
./ptconfig -updateDSfw
./ptconfig -staticRoutes
./ptconfig -restoreSerialNumber
./ptconfig -modifyNic
```

-modelList	list the available models (*)
-restoreNetworkConfig	Restores network configuration (*)
-updateFirmwares	Update xSeries firmwares (*)
-setClock	Manage the time, date, timeserver(s) and timezone settings (*)
-model	specify model type
-updateReplicationIp	update ip addresses used by replication (*)
-addReplication	add IP configuration for replication (*)
-updateSystemName	replace the system name (cluster name) (*)
-usedefaults	don't ask the user for configuration attributes

-activatePTRepMan	Install the ProtecTIER Replication Manager (*)
-upgrade	upgrade an existing model to a different model (*)
-deactivatePTRepMan	Uninstall the ProtecTIER Replication Manager (*)
-configRAS	Configure the RAS service (*)
-validate	run validation test on an installed cluster (*)
-app	specify application to install (use the appList option to see all available applications)
-addReplicationNic	replace replication NIC location after upgrading DD1 to DD3 (*)
-updateNetwork hostname (*)	replace the external (management) ip address and/or hostname (*)
-install (*)	install and configure a PT system with a specific model (*)
-appInterfaces	Configure Application Interfaces (*)
-replace repository (*)	install and configure a PT system with an existing repository (*)
-appList	list the available applications (*)
-updateDSfw	Update the DS disk array firmware (*)
-staticRoutes	edit static routes configuration (*)
-restoreSerialNumber	Restore the serial number for the server (*)
-modifyNic	change Network interface card configuration (*)

Options marked with a '\*' are mutually exclusive

```
[root@italy install]# ./ptconfig -updateNetwork
```

```
[root@italy install]# ./ptconfig -updateSystemName
```

---

**Note:** Do not specify leading zeros in any of the address numbers (e.g. 192.168.001.015).

### 9.5.3 fsCreate

This step is required only when you have a separately attached storage, for example in case of TS7650G or TS7680. The script is not required for the TS7610 or TS7650 appliance models.

The fsCreate utility is creating file systems on the attached storage devices. You need to run fsCreate, which is in the /opt/dtc/app/sbin directory. See Example 9-16.

*Example 9-16 fsCreate utility*

---

```
[root@tuscanly install]# cd ../app/sbin/
[root@tuscanly sbin]# pwd
/opt/dtc/app/sbin
```

```
[root@tuscanly sbin]# ./fsCreate
```

Syntax:

fsCreate

Options (mutually exclusive):

-n	#create GFS file systems for all mpath devices during first time installation
----	---

```

-e          #create GFS file systems for new mpath devices
            during capacity upgrade
-t          #create mount points and register GFS file
            systems to /etc/fstab
-r          #display all repository GFS file systems
-u          #display unused devices
-g          #display all non-repository GFS file systems

```

Optional parameters:

```

-s          #script mode, removes header from output and
            disables user prompts

```

```
[root@tuscanysbin]# ./fsCreate -n
```

---

In our example we use `-n` because this is the first time installation. The `fsCreate` script removes any existing data on the disk array as a result of creating the file systems. You will have to type `data loss` to ensure that all data will be deleted.

You will need to use this script again in case of a capacity increase, but with the `-e` parameter.

**Note:** `fsCreate` is used only for TS7650G and TS7680 models, as the storage is attached separately.

## 9.6 Repository creating

This step is needed only for TS7650G and TS7680 because the storage subsystem is not included with the product. TS7610 and TS7650 appliances have the storage subsystem configured and a repository is created on them.

After the filesystems are ready, which is described in Example 9-16 on page 198, you can make planning for the repository.

### 9.6.1 Repository planning

You should click **Repository** → **Create Repository Planning**. You will get the wizard shown in Figure 9-26 on page 200.

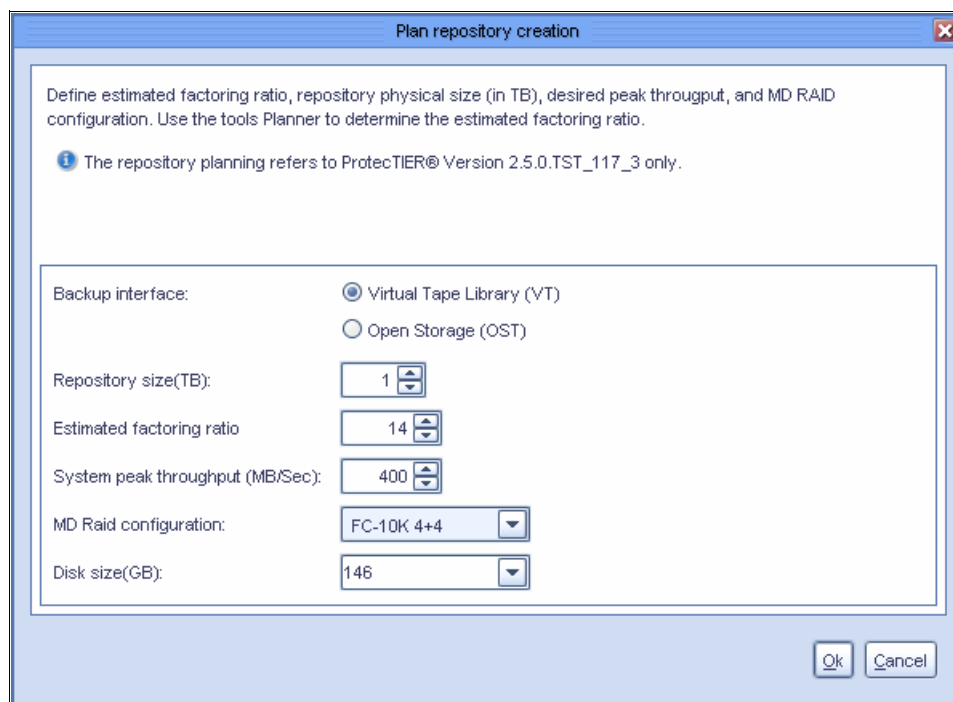


Figure 9-26 ProtecTIER Plan repository creation dialog window

Here you can select if you are planning for Virtual Tape Library (VTL) or Open Storage (OST), and give the inputs you have from the Planning, in 6.3, “Planning for Open systems with VTL” on page 74 or the planning for OST. These numbers are critical for the Repository because after you create the Repository, you will be able to change only the Repository size.

1. In the Repository size field select the size value of the repository that you want to create.
2. In the Estimated factoring ratio field, enter the value estimated for your environment based on your data change rate, backup policies, and retention period.
3. In the System peak throughput field, specify the rate of system peak throughput that your metadata file systems can support.
4. In the MD RAID configuration field, select the RAID configuration of the logical volumes on which the repository metadata file systems are to be created. For example, select FC-15K 4+4 for a configuration of RAID 10 4+4 with Fibre Channel 15K RPM disks.
5. In the Disk size field enter the size of the disks that you use in your storage array.
6. Click **OK**. The Repository metadata storage requirements dialog is displayed with a list of file system arrangement options that are suitable for your needs.

After you are giving the inputs, you will get a window describing the metadata requirements, as seen on Figure 9-27 on page 201.

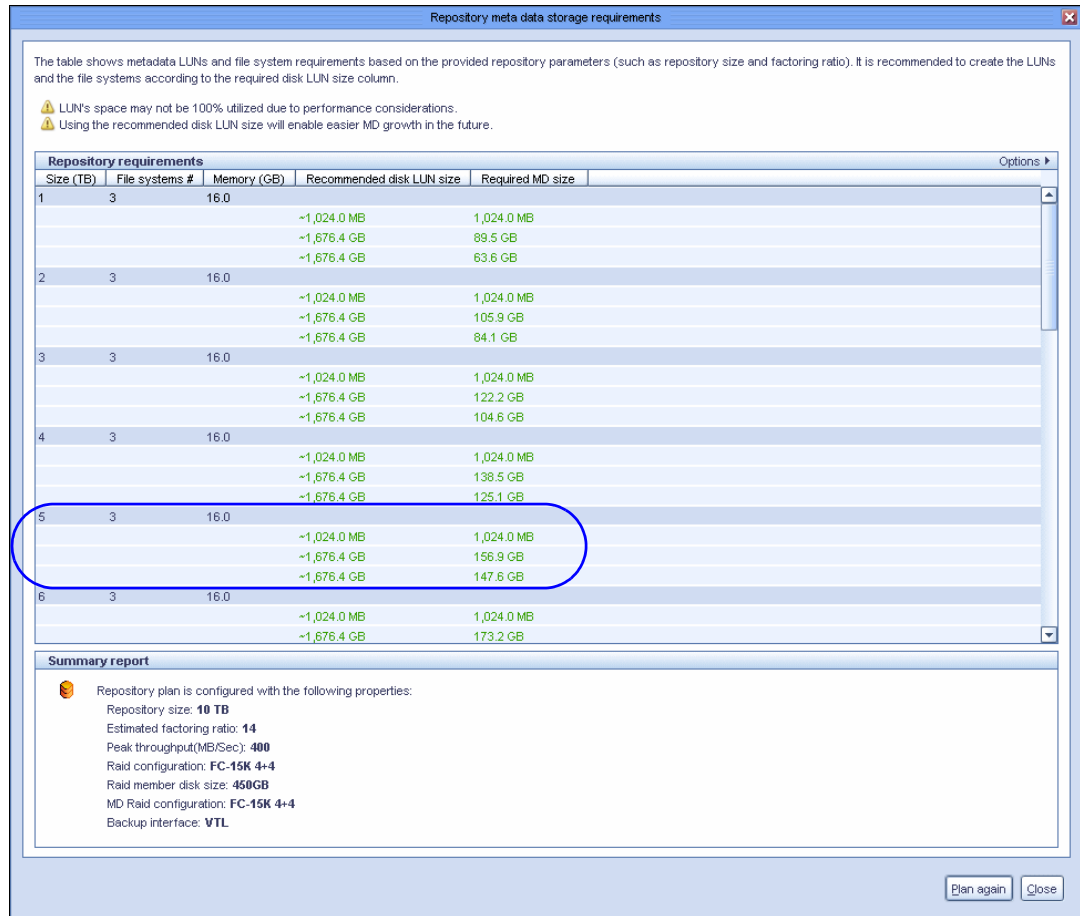


Figure 9-27 ProtecTIER metadata storage requirements

On Figure 9-27 you can see what are the requirements from a metadata point of view. As an example: if you have a 5 TB of Repository, estimated factoring ratio is 14, peak throughput is 400 MBps, use FC 15K 450GB drives in RAID 10 4+4 for metadata, and you will need three file systems and 16 GB of Memory in the ProtecTIER server. The required metadata (MD) size is 1024 MB, 156.9 GB, and 147.6 GB. The 1024 MB requirement is always present, whereas the others are depending on factors like HyperFactor ratio. When you allocate the LUN, use all of the space in the RAID group. In our case we use RAID 10 4+4 for the MD and 450 GB drives, therefore creating approximately  $4 \times 450 \text{ GB} = 1.6 \text{ TB}$  LUNs. For more information about metadata, see “Metadata” on page 98. Also, create a slightly different size of LUNs for metadata and user data, so you are able to distinguish them from each other.

## 9.6.2 Repository creating

A repository can only be created on a one-node cluster.

**Note:** Creating a repository is a prerequisite for adding a second node to a cluster. The repository creation must be done from node A.

1. In the Nodes window, select the node on which to create the repository.
2. Click **Create new repository**. The Create repository wizard Welcome window is displayed as shown in Figure 9-28 on page 202.

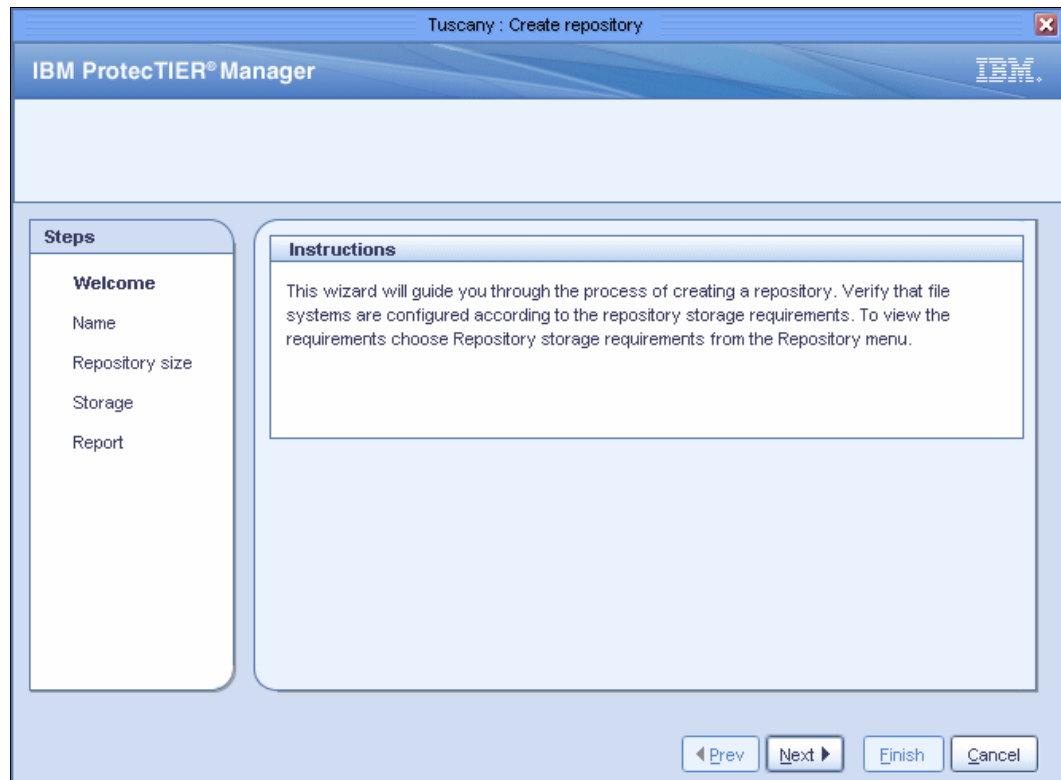


Figure 9-28 Create repository wizard: Welcome instructions

3. Click **Next**. The Repository Name window is displayed as shown in Figure 9-29.

The screenshot shows the 'Tuscany : Create repository' window of the IBM ProtecTIER Manager, now at the 'Name' step. The title bar includes the IBM logo. The main area is divided into two sections: 'Steps' on the left and 'Name' on the right. The 'Steps' section lists 'Welcome', 'Name', 'Repository size', 'Storage', and 'Report', with 'Welcome' marked with a checkmark. The 'Name' section contains two text input fields: 'System name:' with the value 'Redbooks' and 'Repository name:' with the value 'Redbooks\_repository'. Above these fields, there are two informational messages: 'System name field limited to 16 characters.' and 'Repository name field limited to 32 characters.' At the bottom right, there are four buttons: 'Prev', 'Next', 'Finish', and 'Cancel'.

Figure 9-29 Create repository wizard: Name window

4. In the System name field, type the name of the system on which the repository will be created. The Repository name field will be populated automatically. If you would like to have a different name for the repository, you can change that after you entered the System name.
5. Click **Next** and the Repository size window is displayed as shown in Figure 9-30.

The screenshot shows the 'Create repository' wizard in IBM ProtecTIER Manager. The 'Steps' pane on the left indicates the current step is 'Repository size'. The 'Properties' section contains the following fields:

- Repository size (TB):** 5
- Estimated factoring ratio:** 14
- System peak throughput (MB/Sec):** 320
- MD Raid configuration:** FC-15K 4+4
- Disk size(GB):** 300

Navigation buttons at the bottom include 'Prev', 'Next', 'Finish', and 'Cancel'.

Figure 9-30 Create repository wizard: Repository size window

6. In the Repository size field, type the repository size in terabytes that you determined using the Create repository planning wizard.
7. In the Estimated factoring ratio field, enter the estimated factoring ratio value that was determined with the assistance of your IBM System Services Representative (SSR), Lab-based Services (LBS), and Field Technical Service and Support (FTSS) personnel.
8. In the System peak throughput field, specify the rate of system peak throughput that your metadata file systems can support.
9. In the metadata RAID configuration field, select the RAID configuration of the logical volumes on which the repository metadata file systems are to be created. For example, select FC-15K 4+4 for a configuration of RAID 10 4+4 with Fibre Channel 15K RPM disks. This depends on the way that your storage arrays are configured. Other choices are Serial Attached SCSI (SAS) storage arrays and Serial Advanced Technology Attachment (SATA) storage arrays.
10. In the Disk size field enter the size of the disks that are in your storage array.
11. Click **Next**. You will go into the Resources menu, which shows each filesystem that was created before.
12. Verify that the correct file systems are selected for metadata and user data, based on the metadata file system sizes indicated by the repository planning process. The system automatically distributes the filesystems between metadata file systems, available file systems and user data file systems. If the file systems selected by ProtecTIER for metadata and user data do not match the file systems that you created for those purposes, change the assignment by selecting file systems from the available file systems





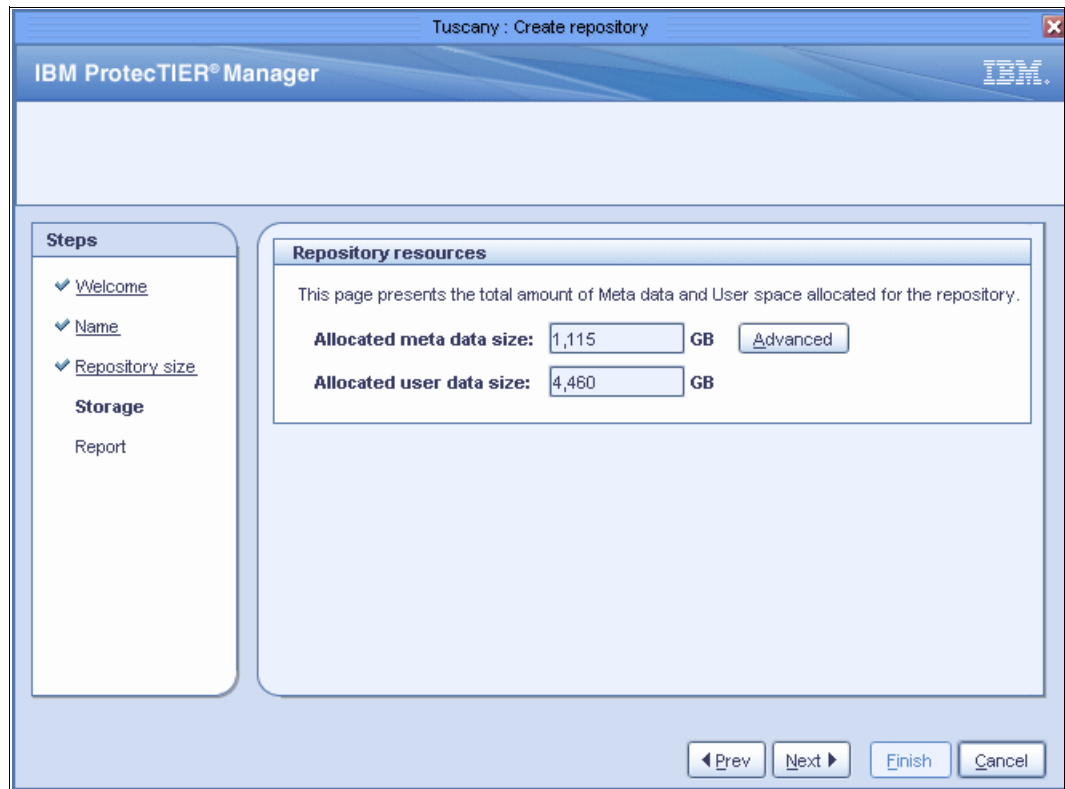


Figure 9-32 Create Repository wizard: Repository resources summary

15. Click **Next**. The Create repository Report window is displayed as shown in Figure 9-33.

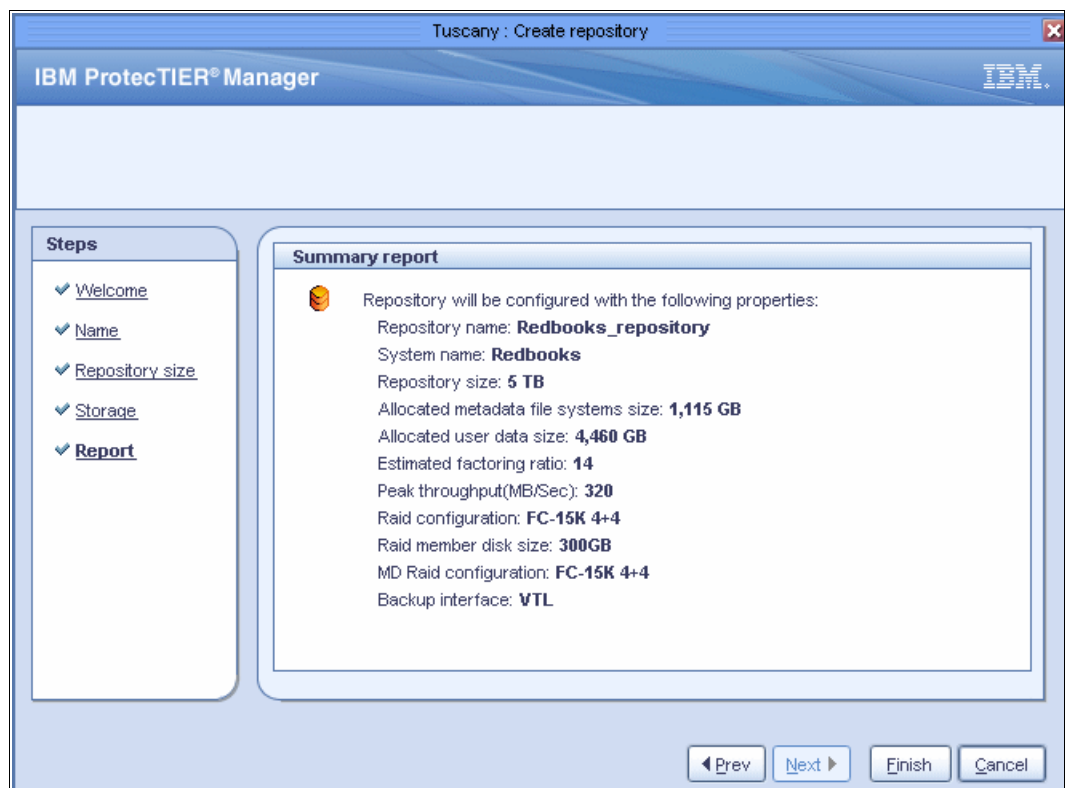


Figure 9-33 Create Repository wizard: Summary report

16. Click **Finish**. The Create repository wizard closes and a confirmation window is displayed (Figure 9-34).

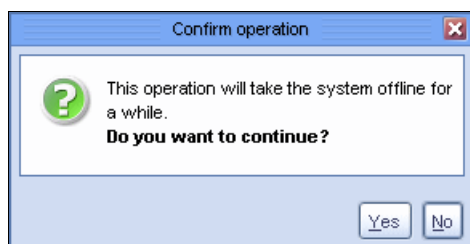


Figure 9-34 Create Repository wizard: Confirm operation

17. Click **Yes**. The ProtecTIER system temporarily goes offline to create the repository. This operation might take a while. The Create repository window is displayed until the repository is created as shown in Figure 9-35.

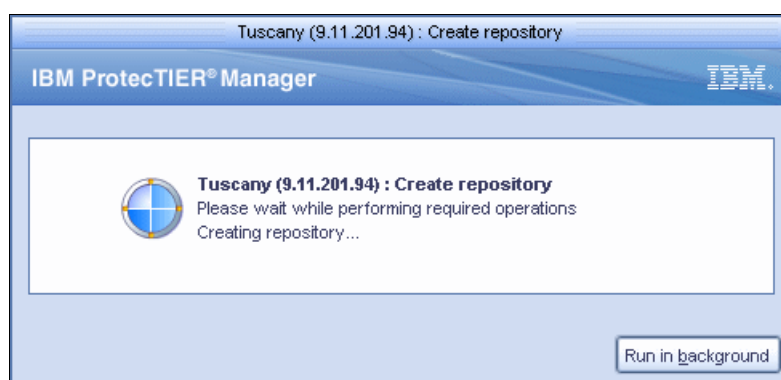


Figure 9-35 Create repository window

## 9.7 OST configuration

If you are using your system for OST, you will need to run **ptconfig -install** with **-app=OST**, and you need to create the Repository with the OST option. Here we describe the next steps to set up Open Storage for NetBackup.

### 9.7.1 The OpenStorage Operating Environment

There are two major components that comprise the OpenStorage operating environment and communicate through a TCP IP network:

- ▶ The storage server
- ▶ The plug-in

The storage server is an entity which runs on the ProtecTIER servers and uses the major internal functionality of the ProtecTIER platform (such as DHF, clustering, and replication).

The plug-in is a shared library (i.e. a stateless software component) that resides on the NetBackup machine and is dynamically linked to the NetBackup application for data transfer to the ProtecTIER storage server emulation.

The storage server (STS) is a high-level component defined by the OST API. In simple terms, it is a “container” of logical storage units (LSUs) and provides the means to access the logical

storage units and their contents. Currently, only one STS can be defined for each ProtecTIER OST storage appliance.

## 9.7.2 Configuring Storage Server (STS)

The first step is to configure Storage Server to OST. Log in to the ProtecTIER Manager and perform the following steps:

1. Click OST → **Storage Server** → **Add Storage server** as shown in Figure 9-36.




Figure 9-36 Add Storage server for OST

2. The STS definition window will show up, where you will need to fill out the STS name and credentials (user name and password) as shown in Figure 9-37. NetBackup uses these credentials so that the media server can log in to the storage server for storage access.

A screenshot of a dialog box titled 'Define a new storage server name and credentials.' The dialog box contains four error messages, each preceded by a red star icon: 'STS Name field is empty.', 'User name field is empty.', 'New password field is empty.', and 'Verify password field is empty.' Below these messages are four text input fields: 'STS Name:', 'User name:', 'New password:', and 'Verify password:'. The 'STS Name' field is currently empty. Below the input fields is a section header 'Credentials' followed by a horizontal line. At the bottom right of the dialog box are 'Ok' and 'Cancel' buttons.

Figure 9-37 Storage Server definition

**Note:** The STS unique name cannot be modified after the STS has been created. Only the credentials can be changed later.

3. You can create only one STS. You can delete the STS, but if you have LSUs existing already you will get a warning.

## 9.7.3 Modifying STS

You can modify the STS credentials, as seen in the Figure 9-38 on page 208, by going to OST → **Storage Server** → **Modify Storage server credentials**.

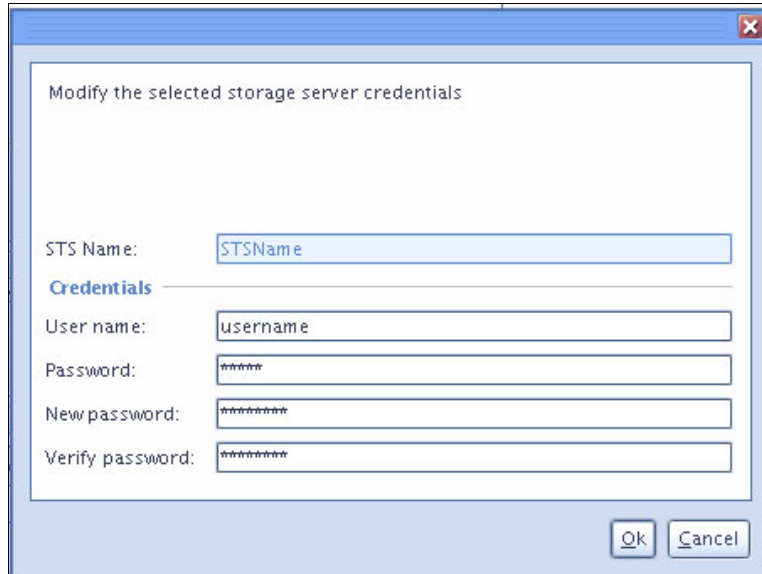


Figure 9-38 Modify Storage server (STS) credentials

**Note:** You cannot modify the STS name, only the credentials.

## 9.7.4 Creating Logical Storage Units (LSUs)

Configuring logical storage units, or LSUs, on a storage server divides the appliance into one or more logical units of space. An LSU, like an STS, is also defined by the OST API and is a “container” of storage and images (which consume storage). Up to 256 LSUs can be defined per STS and are identified by a name string that is unique within the storage server.

In ProtecTIER for OST, an LSU's storage properties are defined in nominal terms. This means that the LSU's storage capacity is defined as a nominal percentage of the repository's overall nominal capacity, taking into consideration the configured percentages of the other LSUs that all share the repository's physical storage.

**Note:** The overall capacity of all LSU's together can be less than or equal to the repository's nominal capacity.

When you are defining LSU, a unique name should be used, which cannot be modified later. The other parameters can be modified later.

Do these steps:

1. Click **OST** → **Logical Storage Unit management** as shown in Figure 9-39.




Figure 9-39 OST -Logical Storage unit management

2. You will see the management window, as on Figure 9-40 on page 209.



**Note:** After an LSU has been created and saved, the LSU name can no longer be modified.

5. In the Nominal percentage field, configure the LSU to the percentage of the total repository. A maximum of 100% can be allocated, and you can define it up to two precision digits in the %, like 10.75%. If you try to overallocate, a warning will display as shown in Figure 9-42.

[illegible]

Figure 9-42 Adding LSU - trying to overcommit

6. Set the HyperFactor mode for each LSU as shown in Figure 9-43 on page 211.

The screenshot shows a dialog box titled "Adding an LSU". Inside, there's a section "Define the LSU properties" with two messages: "LSU Name field is empty." (with a red error icon) and "Note that LSU name can not be modified after creation" (with a blue info icon). Below these are five input fields: "LSU Name:" (empty text box), "Description:" (empty text box), "Nominal percentage:" (text box with "0.0" and a spinner), "HyperFactor:" (dropdown menu showing "Enabled"), and "Compression:" (dropdown menu showing "Enabled", "Disabled", and "Baseline"). At the bottom right are "Ok" and "Cancel" buttons.

Figure 9-43 LSU HyperFactor modes

7. You can set up Compression mode for each LSU as shown in Figure 9-44.

This screenshot is identical to the previous one, showing the "Adding an LSU" dialog box. The "HyperFactor:" dropdown is still set to "Enabled". The "Compression:" dropdown menu is now open, showing the options "Enabled", "Disabled", and "Baseline". The "Enabled" option is highlighted in blue.

Figure 9-44 LSU Compression modes

8. After you filled out all fields and selected the appropriate modes, click **OK**. The LSU displays in the management window as New as shown in Figure 9-45 on page 212.

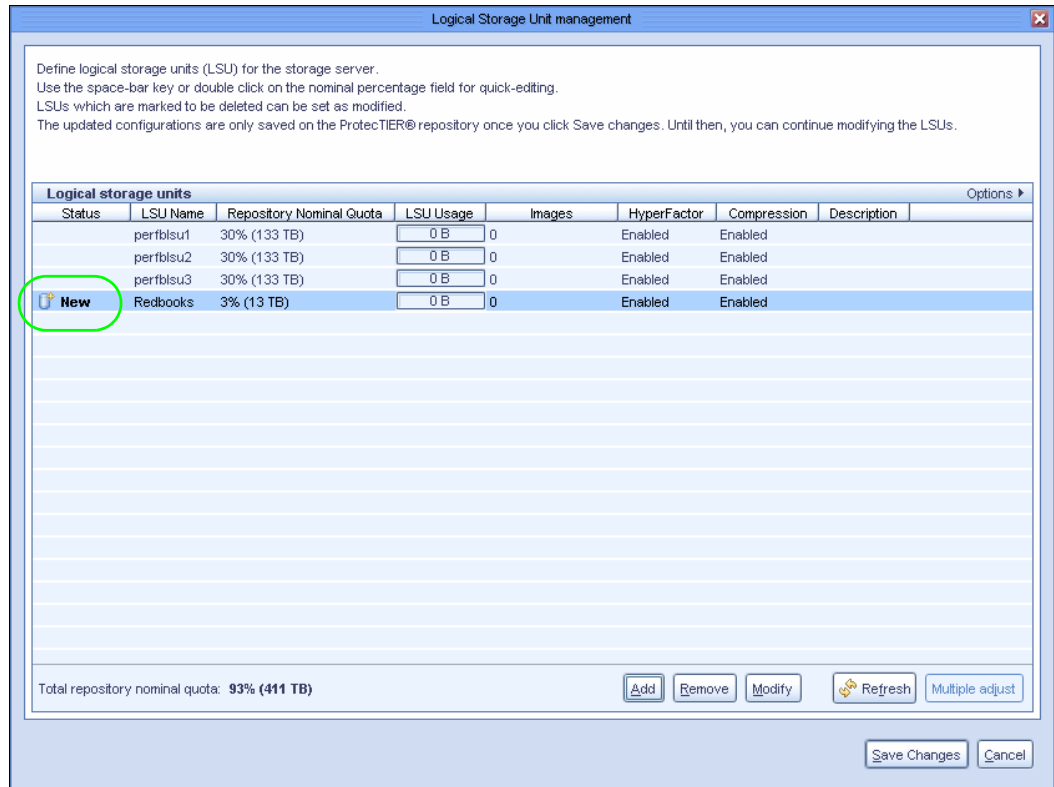


Figure 9-45 LSU Management: New LSU

9. To create this LSU, click **Save changes**.

### 9.7.5 Modifying LSUs

You can modify LSUs later in the LSU Management window by performing the following steps:

1. Select the LSUs you want to change and click **Modify** as shown in Figure 9-46 on page 213.



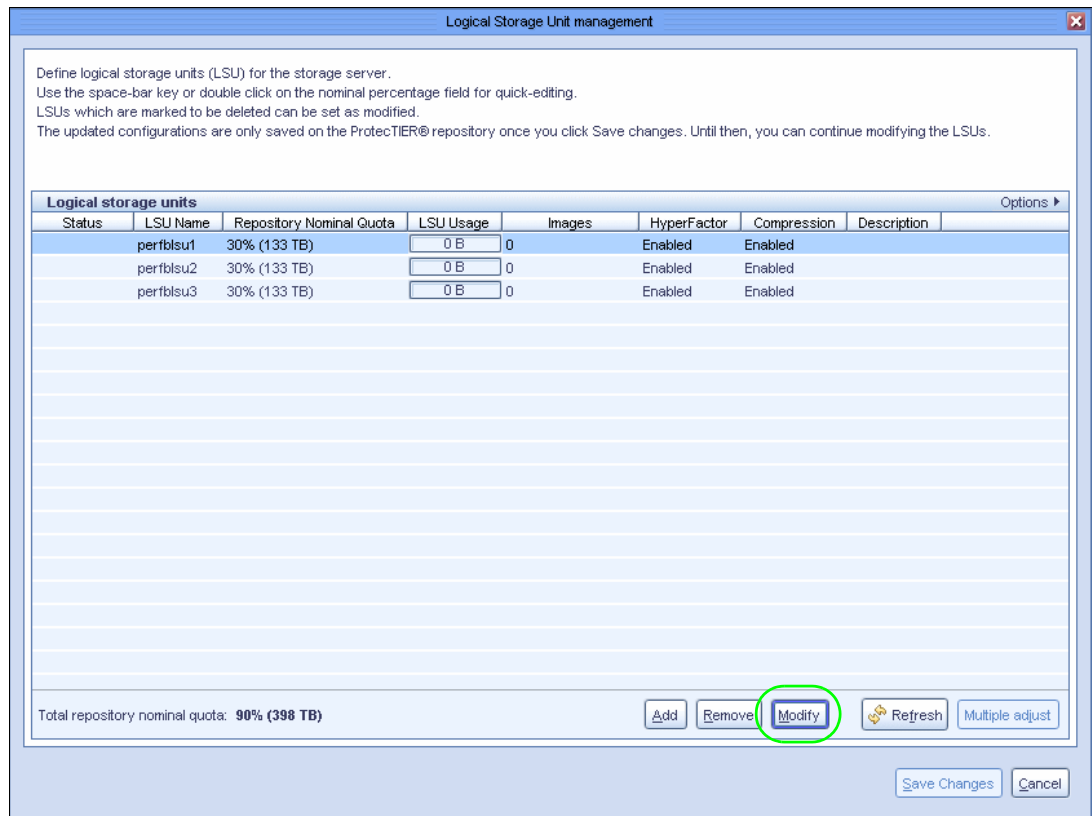


Figure 9-46 LSU management window: Modify

After selecting the LSU and clicking **Modify**, the window shown in Figure 9-47 displays.

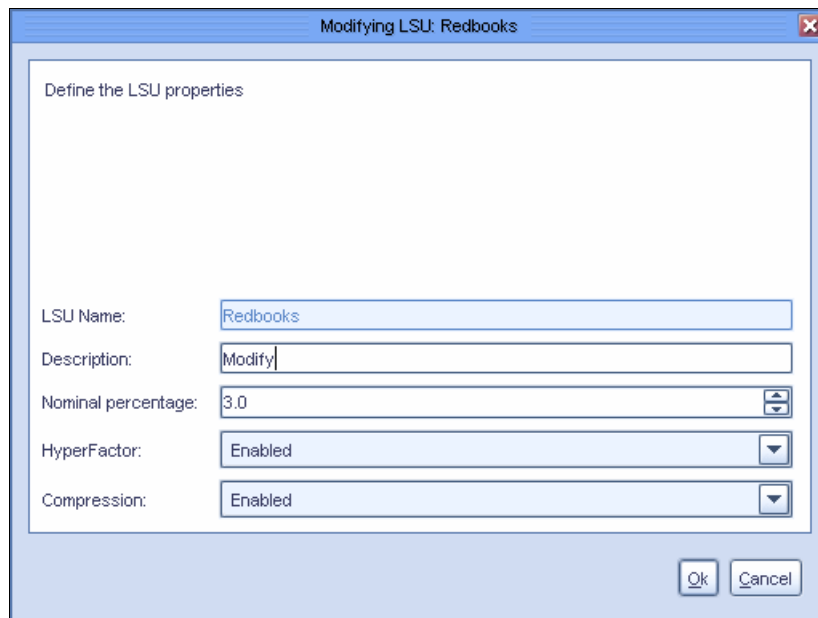
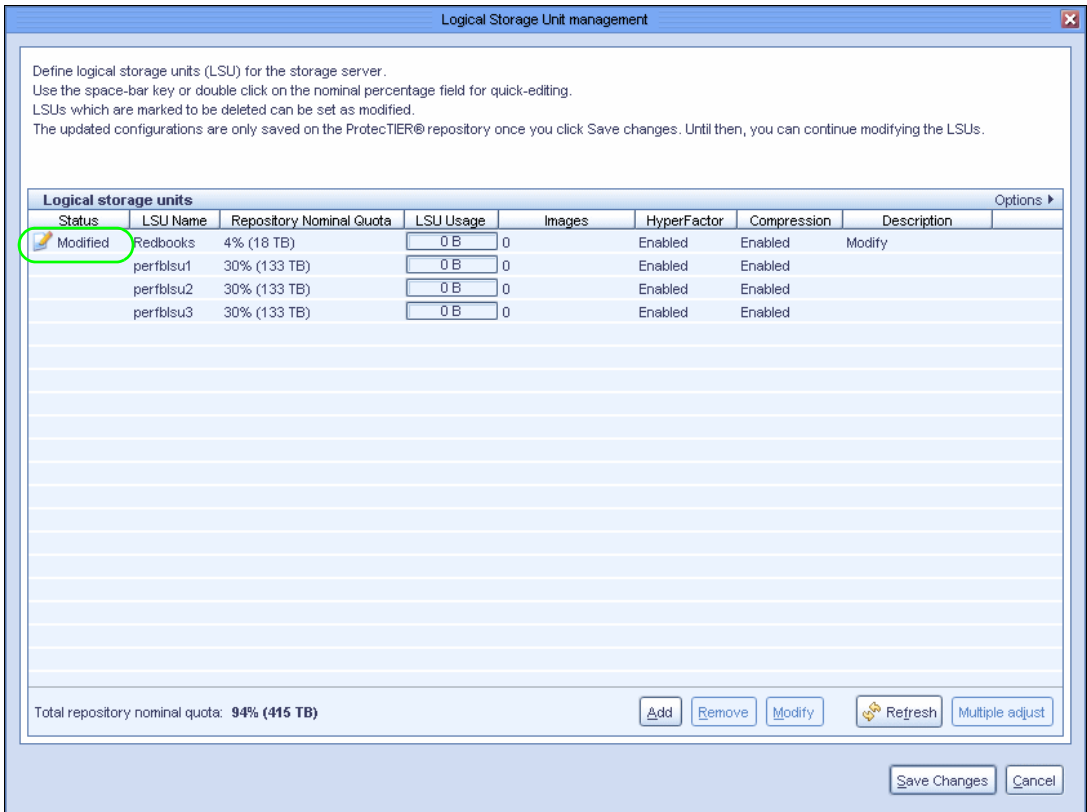


Figure 9-47 LSU Modify

2. Change the description, the % size of LSU, mode of HyperFactor, and mode of Compression as desired.

**Note:** You cannot change the LSU name.

3. When you complete the modifications, click **OK** to go back to the LSU management window. The modified LSU displays as shown in Figure 9-48.



**Figure 9-48** *LSU Management: Modified LSU*

4. To save your modifications, click **Save changes**.

**Note:** If you change the HyperFactor or Compression mode to enabled or disabled, the changes will take effect only on the next image.

You can select more than one LSU by pressing and holding the CTRL key and clicking each as shown in Figure 9-49 on page 215.

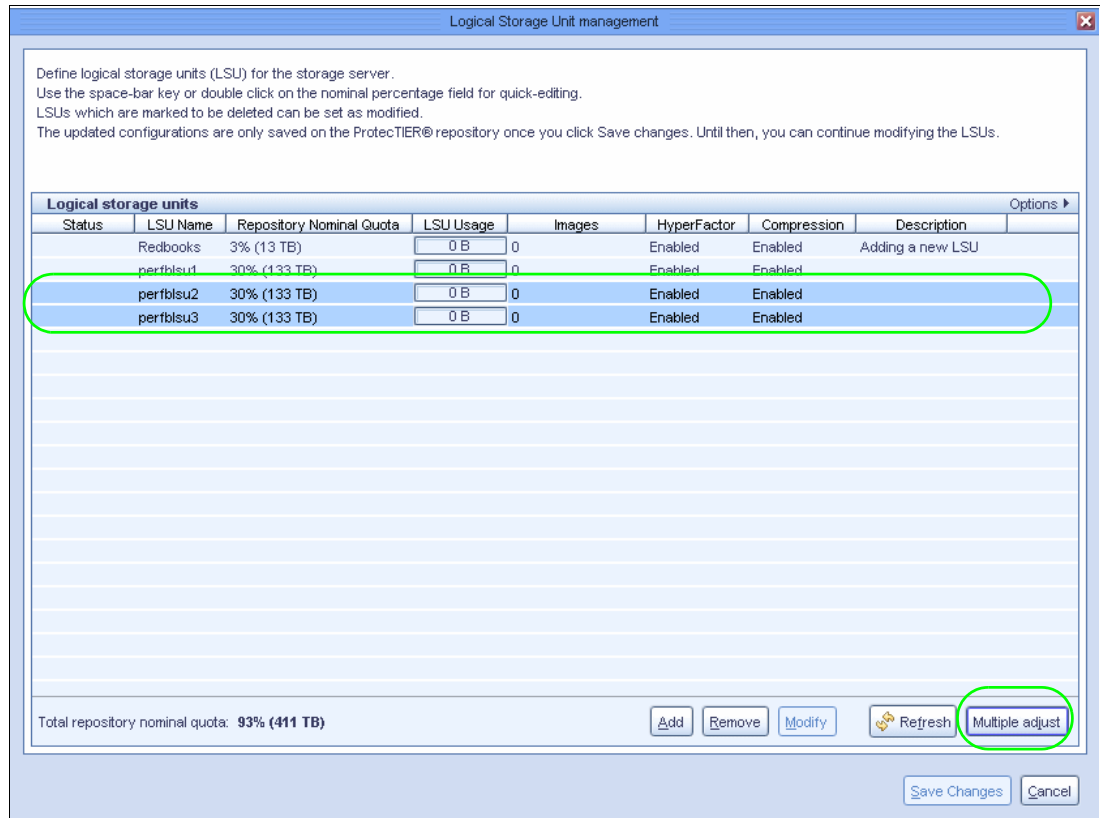


Figure 9-49 LSU Management: Multiple selection

In this case the Multiple adjust function becomes available. If you click it, you can distribute the remaining quota of the repository equally between the selected LSUs, or you set all of them to a specified percentage of the quota as shown in Figure 9-50.

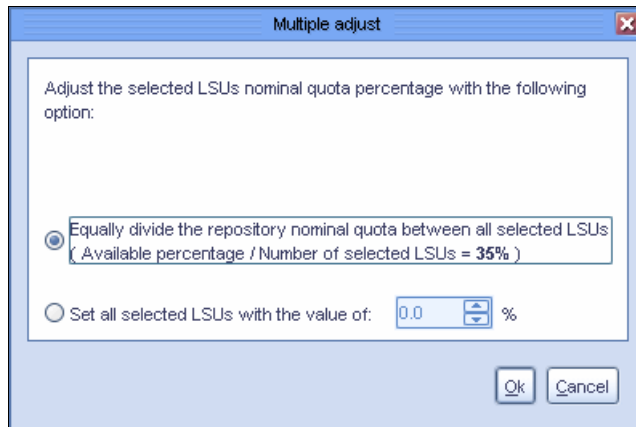


Figure 9-50 LSU Management: Multiple adjust

## 9.7.6 Deleting LSU

You can delete an LSU by selecting it in the LSU management window, and clicking **Delete**. The LSU will be marked as Delete as shown in Figure 9-51 on page 216.

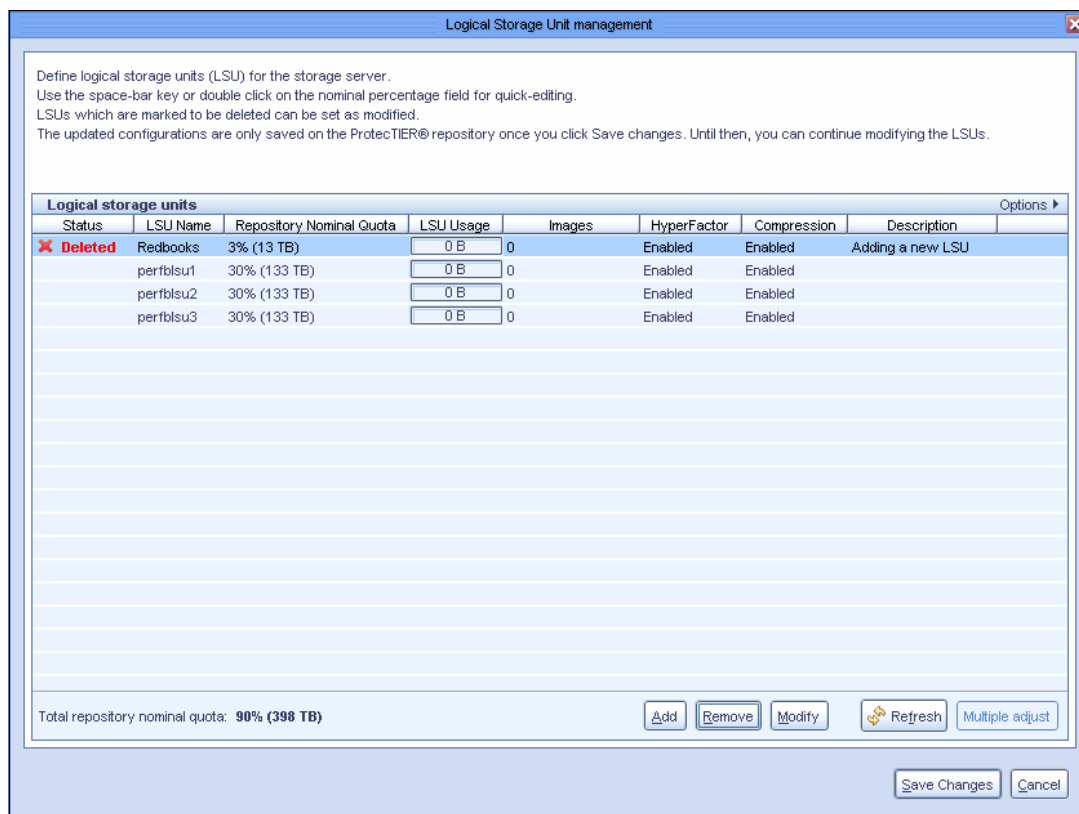


Figure 9-51 LSU Management: Delete

To make the changes, click **Save Changes**. You will be asked to confirm the delete by typing data loss as shown in Figure 9-52.

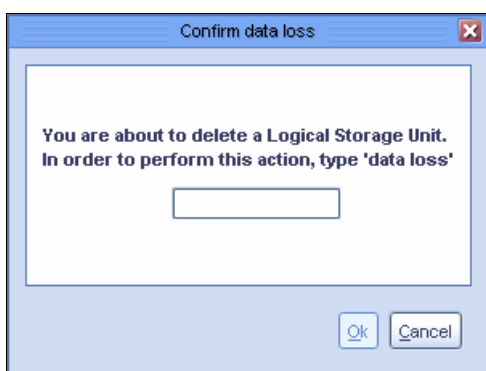


Figure 9-52 LSU Management: Delete confirmation

After clicking **OK**, the LSU will be deleted.

## 9.8 Virtual library creation

The ProtectTIER system enables you to create virtual tape libraries on which the backup application stores your data. These libraries and their components are part of the virtual tape service.

This step is not necessary for TS7610 because it has preconfigured library, but if you want to create more libraries, perform the following steps.

A library can be created on a ProtecTIER system of either a one-node cluster or a two-node cluster.

**Note:** Use the **Scan** button of the Port attributes pane to verify that the ports of the ProtecTIER system to which the virtual devices of the library are to be assigned are connected to the correct host.

If LUN masking is enabled, make sure to create LUN masking groups for the new library too.

1. Log in to the system you want to create the new library on.
2. Click **VT** → **VT Library** → **Create new library**, and the Create library wizard will start as shown in Figure 9-53.

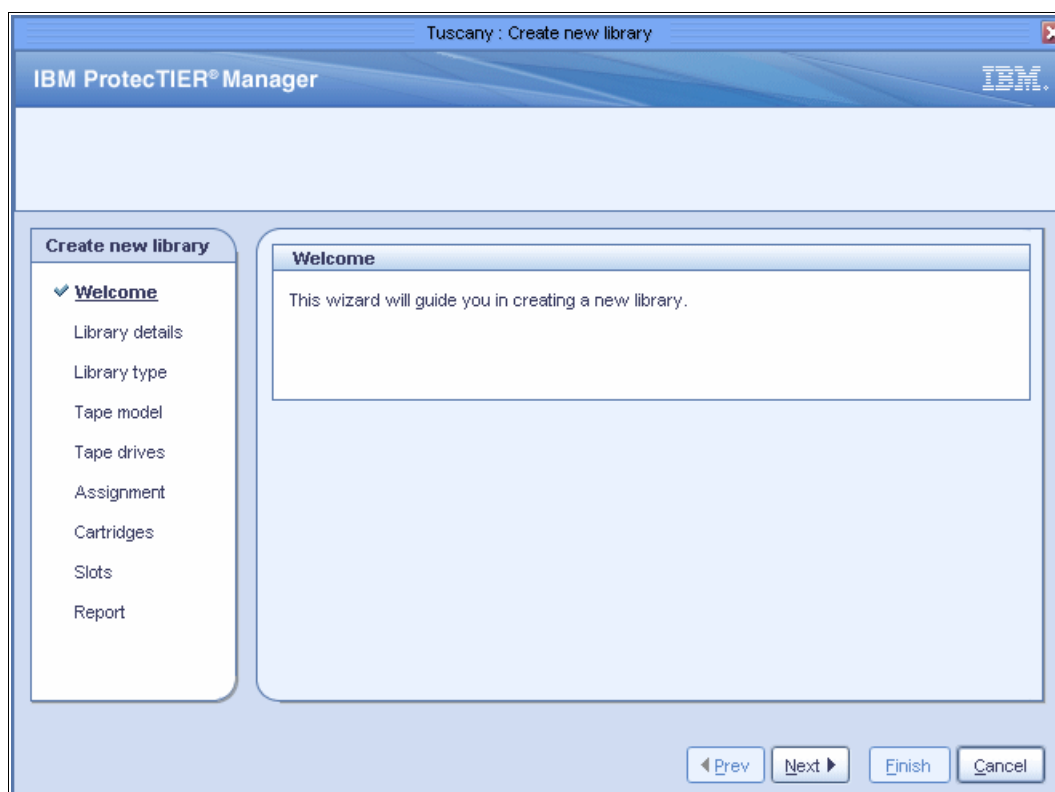


Figure 9-53 Create new library Welcome window

3. Click **Next**. The Library details window is displayed as shown in Figure 9-54 on page 218.
4. In the ProtecTIER VT name field, type a name for the library.

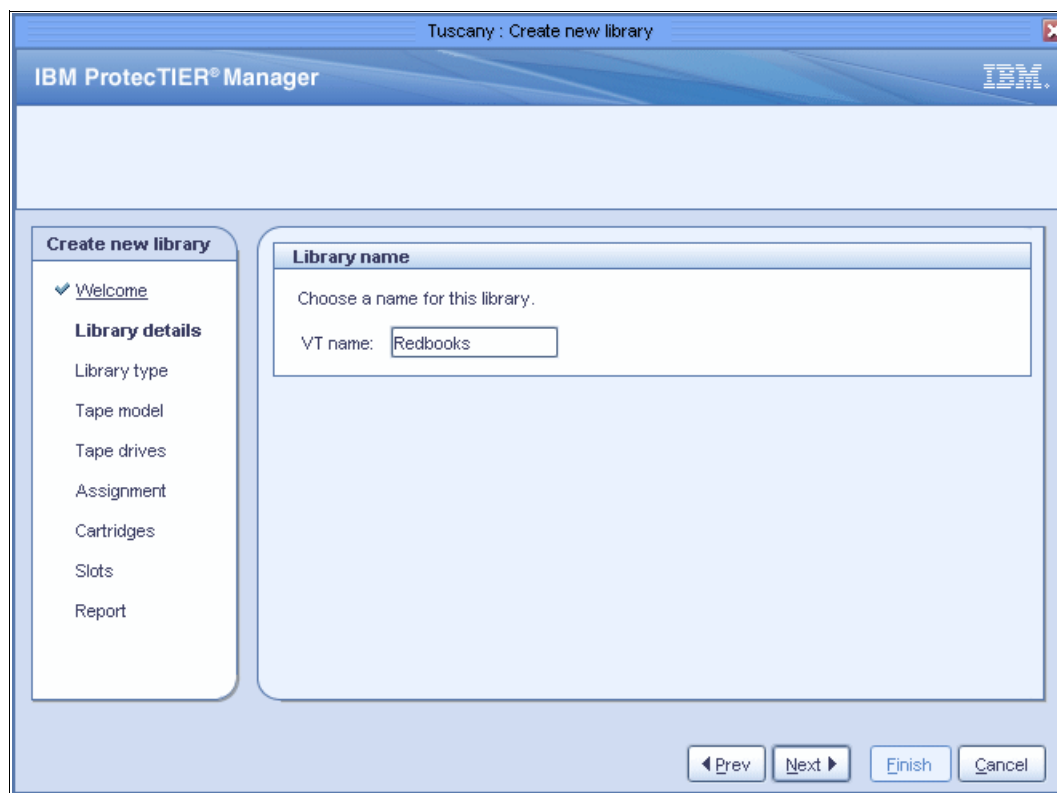


Figure 9-54 Create new library: defining library name

5. Click **Next**. In the library details specify the type of library that you want to use for your application as shown in Figure 9-55 on page 219.

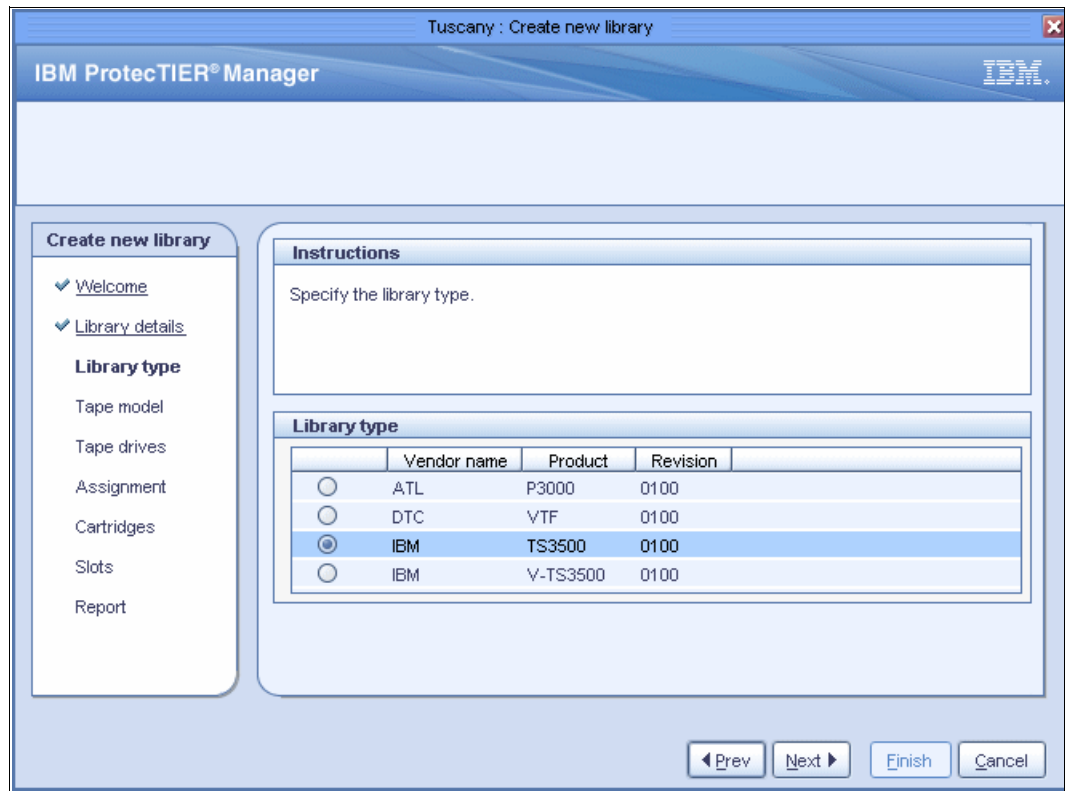


Figure 9-55 Create new library: selecting Library type

By default the IBM TS3500 is selected.

**Note:** V-TS3500 is Symantec Veritas NetBackup (NBU) requirement. The functionality of the IBM TS3500 and the IBM V-TS3500 are the same.

**Note:** Verify that the backup application that you are using supports the type of library model that you select.

- Click **Next**. The Tape Model window is displayed. Select the tape drive model that you to use for your virtual library as shown in Figure 9-56 on page 220.

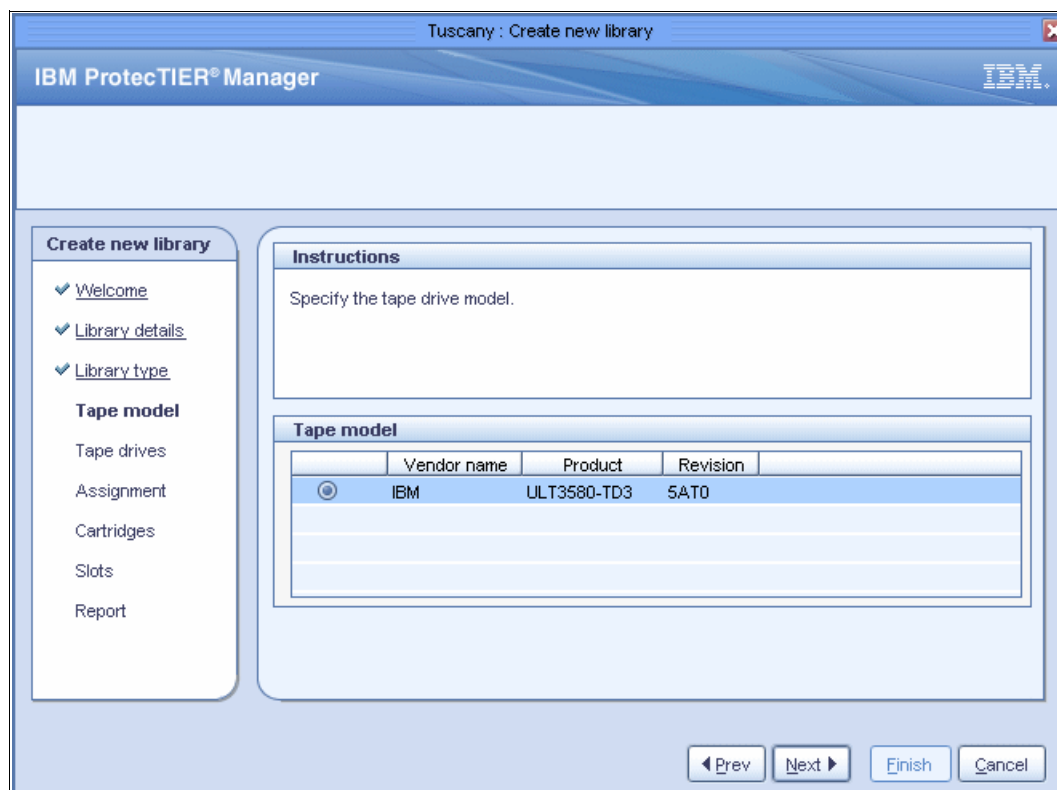


Figure 9-56 Create new library: selecting tape drive model

- Click **Next**. The Tape drives window opens. In the Number of tape drives field for each node, type the number of tape drives to assign to the node. To maximize load balancing, distribute tape drives across the nodes in a two-node cluster based on the relative power of the nodes. See the window in Figure 9-57 on page 221.



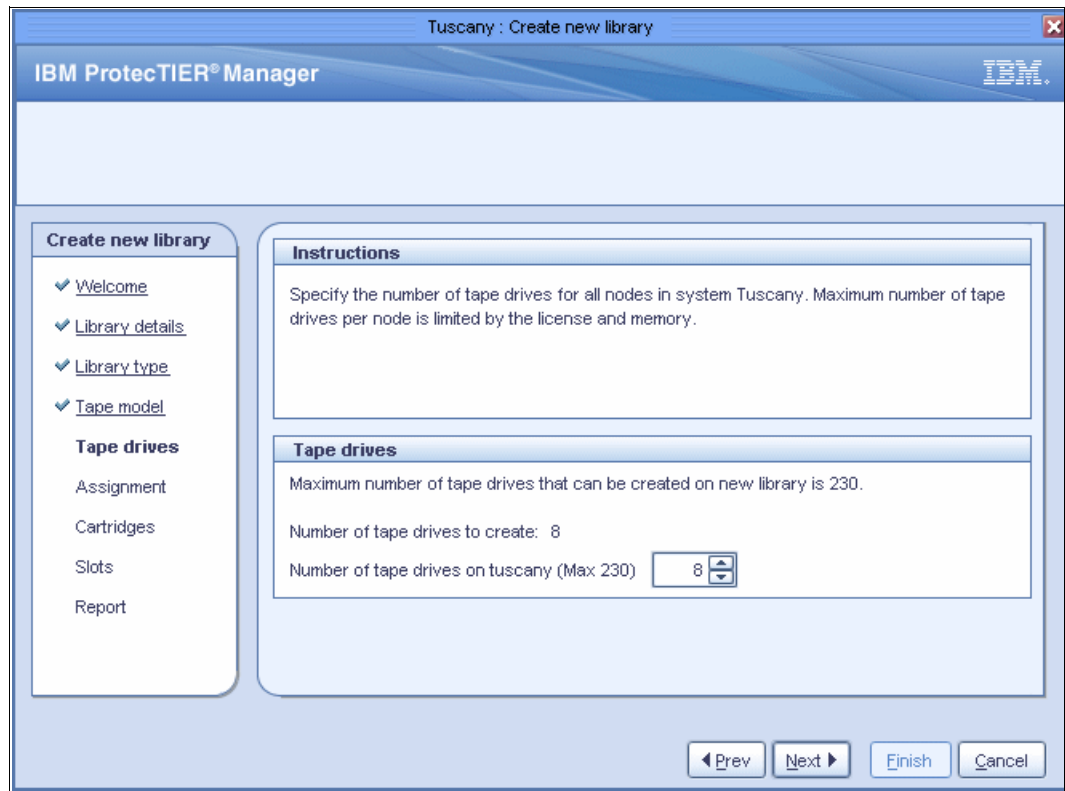


Figure 9-57 Creating new library: defining number of drives

In this example we create eight drives on node Tuscany. We can create maximum 230 drives because there are other libraries already using tape drives.

**Note:** The maximum number of drives is 256 per node. These drives are divided between the virtual libraries as you are configuring them.

8. Click **Next**. The Assignment window opens

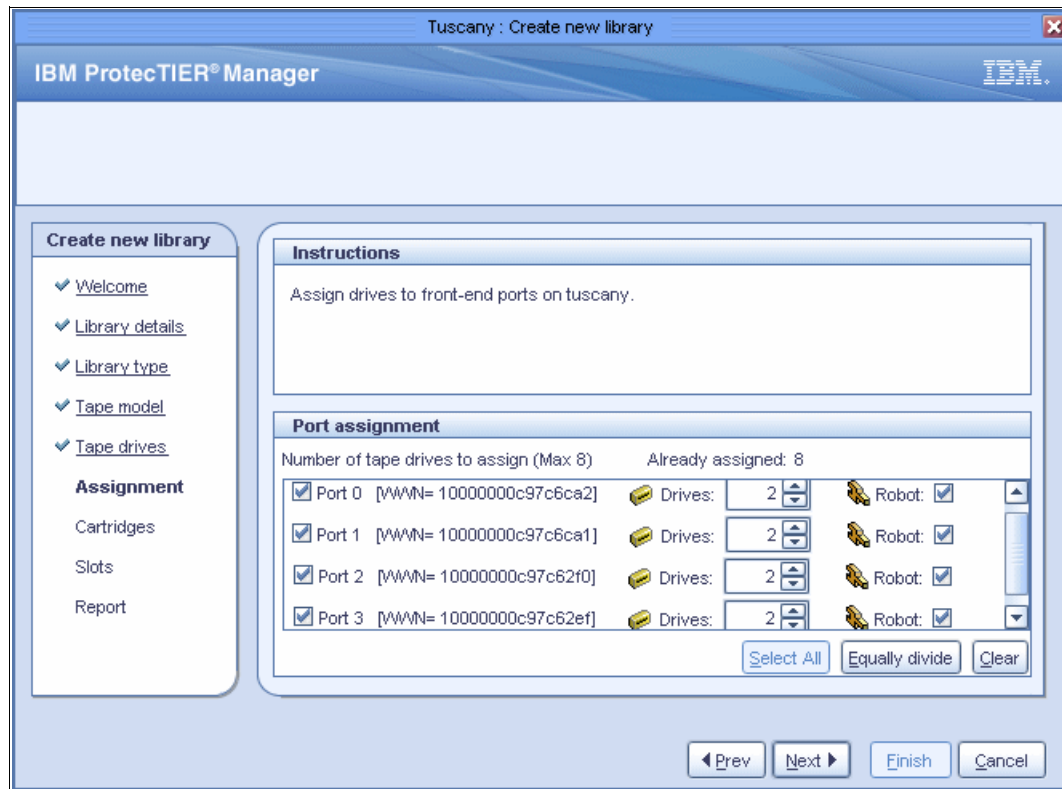


Figure 9-58 Create new library: assigning drives and robots to ports

9. Select or deselect the check boxes next to each port to define which of the node's ports are assigned virtual devices. In our example we set up a IBM TS3500 and by default all the robots are selected and enabled. If you have chosen a library model other than IBM, the robots are not checked and only one must be chosen.
10. In the Drives fields corresponding to each selected port, select the number of virtual tape drives that are assigned to each port.  
 Optionally, click **Select All** to automatically select both ports. Click **Equally divide** to evenly divide the number of drives between the ports.
11. Check the Robot check box if you want the library virtual robot to be accessible through this port.

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

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

If you are using LUN masking, make sure that all required devices are listed in the LUN masking group.

12. Click **Next**. The cartridge windows is displayed as shown in Figure 9-59 on page 223.

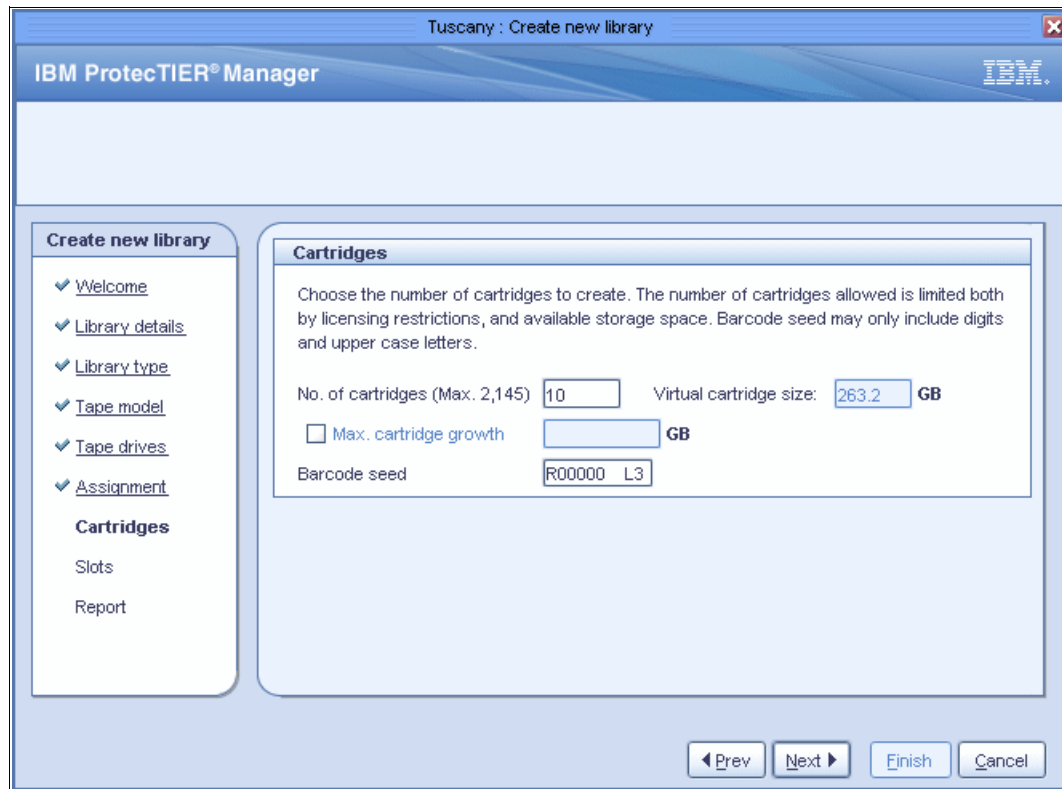


Figure 9-59 Create new library Cartridges window

13. In the No. of cartridges field, enter the number of cartridges that you want to have in the library. The Virtual cartridge size field automatically displays the maximum possible size for virtual cartridges for your system, based on the number of cartridges entered, the total amount of available storage space in your repository, and the current HyperFactor ratio.

Optionally, select the Max. cartridge growth check box. When selected, you can limit the maximum amount of nominal data that a cartridge can contain. The value of the maximum number of cartridges possible on a system depends on the amount of storage space available on your system. For more details, see “Planning for cartridges” on page 88.

In the Barcode seed field, enter a value for the barcode seed. The barcode seed is the barcode that is assigned to the first cartridge created. Every cartridge added after the first cartridge is assigned a barcode following the initial barcode seed.

**Note:** The barcode seed must contain only numbers and capital letters and be only six characters in length (for example, R00000).

The maximum quantity of cartridges depends on the amount of storage space available on your system.

14. Click **Next**. The Slots window opens as shown in Figure 9-60 on page 224.

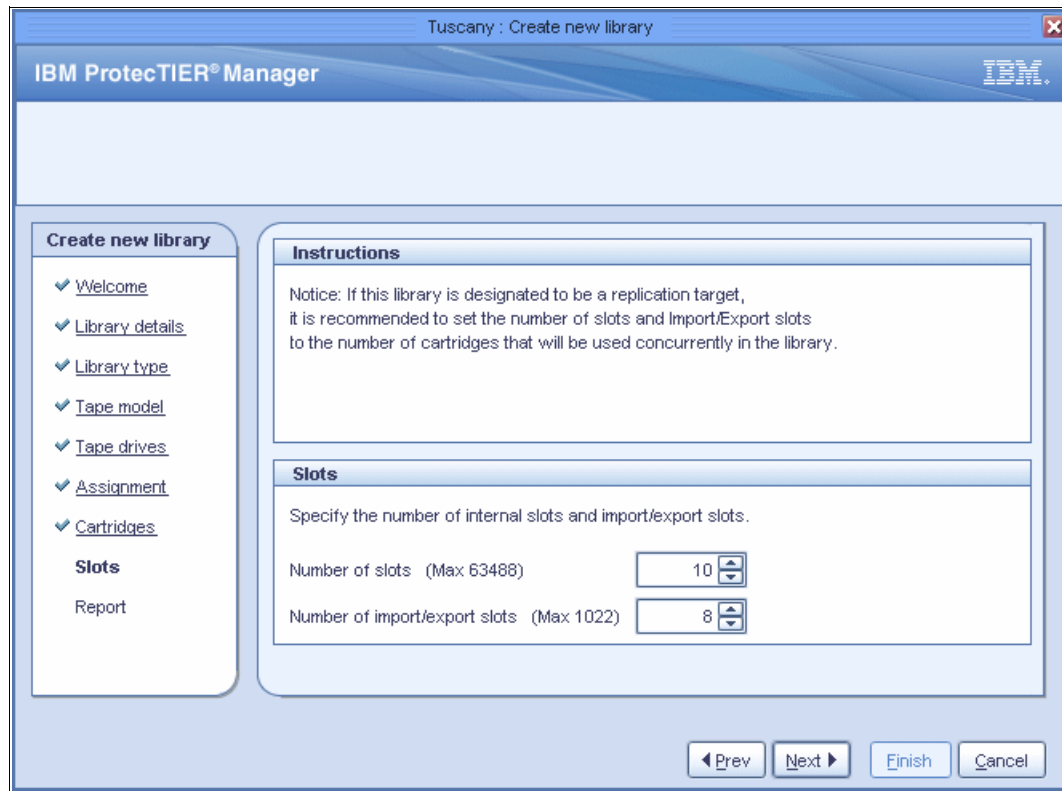


Figure 9-60 Create new library: Defining slots

In the Number of slots field, select the number of cartridge slots that you want to have in the library. By default the number of slots will be equal to the number of cartridges added in previous step.

**Note:** Create more slots than cartridges created in the previous step so that you have empty slots. If you have empty slots, you are able to add cartridges later easily, without any disruption.

If you do not have empty slots when adding new cartridges, the action will be disruptive. Because the dimensions of the library have to be changed in this case, the system will go offline creating it.

In the Number of import/export slots field, enter the number of import/export slots that you want to have in the library. The maximum number of import/export slots that can be defined in the entire system is 1022. The I/O slots will be used to move cartridges between the libraries and shelf.

15. Click **Next** → **Finish**. The Create new library wizard closes and a summary report is displayed as shown in Figure 9-61 on page 225.

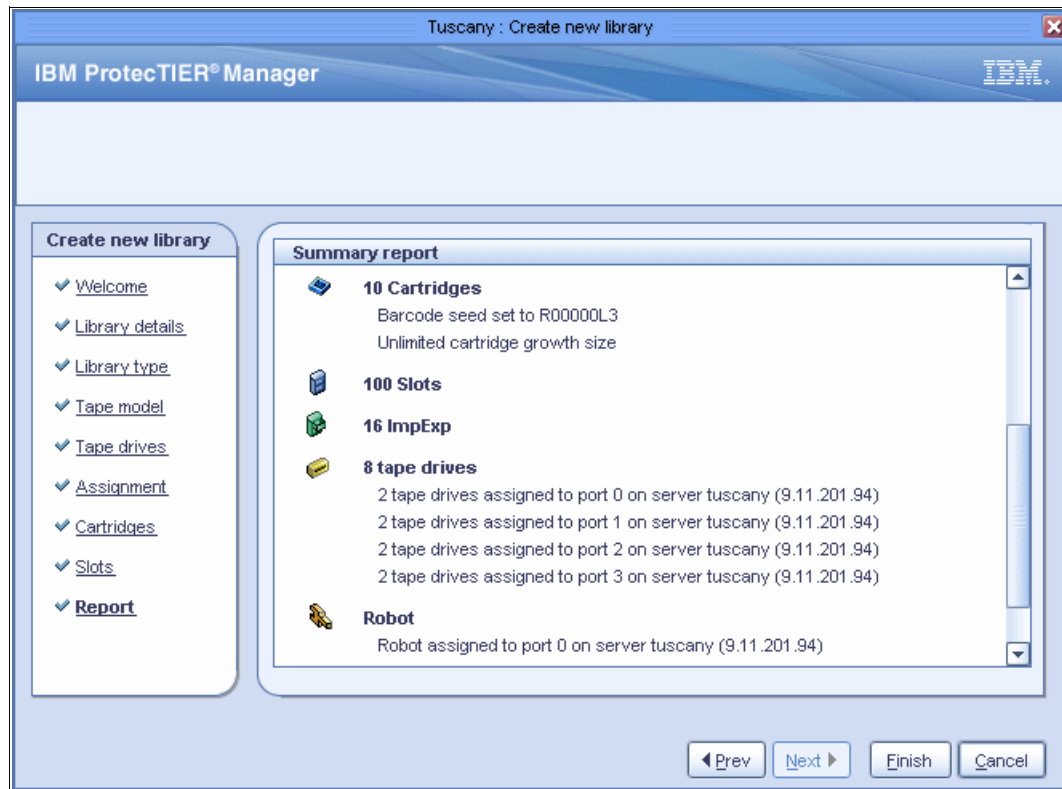


Figure 9-61 Create new library: summary report

16. The ProtecTIER system temporarily goes offline (Figure 9-62) to create the library. The library is displayed in the Services window and the VT monitoring window opens.




Figure 9-62 Create new library: system offline confirmation

Your virtual library becomes available for your host.

### 9.8.1 TS7610 Library

IBM TS7610 ProtecTIER SMB appliance has a preconfigured library, depending on the model you ordered, with the following specifications:

- ▶ TS3500 Virtual Library
- ▶ 16 LTO3 virtual drives, balanced evenly across both FC ports
- ▶ 200 GB cartridge size
- ▶ 16 virtual import / export slots
- ▶ In case of Small model:
  - 400 virtual slots
  - 400 cartridges
- ▶ In case of Medium model

- 540 virtual slots
- 540 cartridges

## 9.9 Host implementation

In this section, we discuss configuring host systems for attachment to ProtecTIER systems and what is required to achieve optimum performance from ProtecTIER systems.

Prerequisite is to have a working Virtual Tape Library or OST.

### 9.9.1 Connecting hosts to ProtecTIER systems

You can connect your backup server to the SAN and then create a zone to include both the host HBA WWN and the WWN of ProtecTIER Front End ports. But if you intend to take advantage of path redundancy or you have a two-node clustered system, then the connect should be well designed, not only for redundancy purposes but also for performance reasons.

In this section we discuss connecting and assigning for optimum performance.

#### What bandwidth you need for ProtecTIER

Many SANs are designed for disk I/O, which means that they are optimized for high input/output per second (IOPS) and are not necessarily for the high bandwidth that is needed for tape and backup I/O. During backup sessions, high volumes of data are generally transferred from disk to tape or virtual tape system. High-speed transfer is also a necessity to minimize backup time. Current tape or virtual tape drives, like the IBM ProtecTIER with up to 500 MBps per node and up to 1000 MBps of two-node clustered system, can easily saturate a single FC link if several virtual drives operate concurrently. Using a SAN fabric configuration, you can potentially attach and access many tape drives through one or more HBAs. But how will this affect performance if several tape virtual drives run concurrently? The theoretical maximum data transfer rate for one FC connection in a 2 Gbit SAN is 200 MBps. In reality, we typically see an effective data transfer rate of about 160 MBps. In a 4 Gbit SAN we see an effective data transfer rate of just around 300 MBps. So if you have only one 4 Gbit SAN HBA in your backup media server, you cannot maximize the power of your ProtecTIER system.

As well as the number of HBAs, the overall SAN design must be capable of supporting the volume of data transferred during the backup process. The SAN might be able to easily sustain the overall load and normal data traffic, but still have an insufficient bandwidth when backups take place. This is especially true in complex fabrics with multiple switches where the inter-switch links (ISLs) can become saturated. You can increase your bandwidth if you install additional HBAs and increase the number of ISLs.

Figure 9-63 on page 227 shows a typical SAN layout with ISL between the servers and virtual tapes.




Figure 9-63 SAN bandwidth

### Multiple paths to tape drives

For HA or performance reasons, you might use more than one HBA on the server to access virtual tape libraries and drives. However, if you implement this configuration, you might see duplicate virtual tape drive definitions or robotics definitions on your server. To solve this problem, use one of the following solutions:

- ▶ Enable *LUN masking* in ProtecTIER, which allows you to exactly define for one HBA to see specific devices
- ▶ Enable zoning, which allows one HBA to see only certain tape drives
- ▶ Use the alternate path function within the IBM Tape Device Driver
- ▶ Use persistent binding for the HBAs to see only certain WWNs
- ▶ Ignore the extra copy of each device

### Relationship of virtual tape resources and device name in OS

Regardless of the solution that you choose, be aware of the relationship between the virtual tape resource and the device name in the OS. Unlike a physical tape library, you cannot distinguish the tape devices by WWN because we can get more than one device from one FC port on the virtual tape library because they might have the same WWN. In order to identify them, you must use the serial number. You can get the serial number by using the **itdt** command in Linux, as shown in Example 9-17. The **itdt** command is installed with the **lin\_tape** driver, which is described in 9.9.2, “Installing and configuring the device driver” on page 230.

#### Example 9-17 itdt command

```
[root@frankfurt tmp]# ./itdt -f /dev/IBMtape0 inquiry 80
Issuing inquiry for page 0x80...
```

Inquiry Page 0x80, Length 14

```
      0 1 2 3 4 5 6 7 8 9 A B C D E F 0123456789ABCDEF
0000 - 0180 000A 3134 3937 3531 3530 3030      [...1497515000 ]
```

Exit with code: 0

```
[root@frankfurt tmp]# ./itdt -f /dev/IBMtape1 inquiry 80
```

Issuing inquiry for page 0x80...

Inquiry Page 0x80, Length 14

```
      0 1 2 3 4 5 6 7 8 9 A B C D E F 0123456789ABCDEF
0000 - 0180 000A 3134 3937 3531 3530 3032      [...1497515002 ]
```

Exit with code: 0

```
[root@frankfurt tmp]# ./itdt -f /dev/IBMtape2 inquiry 80
```

Issuing inquiry for page 0x80...

Inquiry Page 0x80, Length 14

```
      0 1 2 3 4 5 6 7 8 9 A B C D E F 0123456789ABCDEF
0000 - 0180 000A 3134 3937 3531 3530 3034      [...1497515004 ]
```

Exit with code: 0

We recommend to identify the virtual tape devices and create a table with the data, as shown in Table 9-2.

Table 9-2 Virtual Tape Library worksheet

Device in OS	Type	VTL system	VTL node	VTL port	VTL port WWN	Serial number	Element number

## SAN zoning

The ProtecTIER Gateway and appliance have specific suggestions about how SAN zones are created:

- Use zones based on World Wide Port Name (WWPN).
- Use two-member zones, that is, one initiator port and one target port per zone.
- For each backup server, create a separate zone for each HBA that will access ProtecTIER virtual resources.

Before creating WWN zones at a SAN switch, you must get the WWPN of each port for both your ProtecTIER and your host computer.



For ProtecTIER, you need the WWPN of each front end port. To do this, you can use the ProtecTIER Manager GUI as shown in Figure 9-64. In ProtecTIER the Front End ports are Emulex HBAs, the Back End ports are Qlogic HBAs.

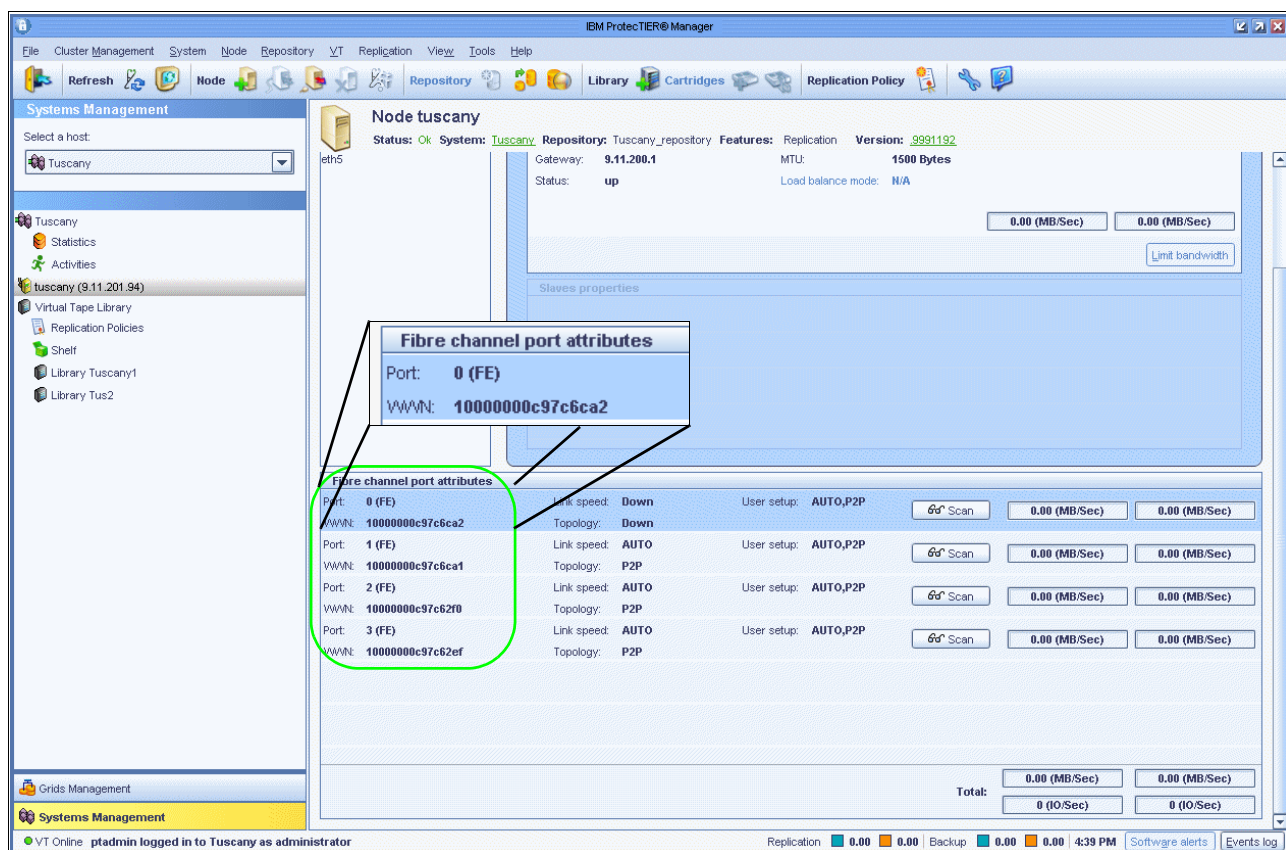


Figure 9-64 ProtecTIER Front-end port WWPN

For your host computer, you can issue a command to get WWPN of each HBA. In our case, we worked with Linux as shown in Example 9-18.

#### Example 9-18

```
[root@frankfurt ~]# cat /sys/class/fc_host/host6/port_name
0x10000000c97c70cf
```

After successful zoning, if you click the **Scan**, you can check what WWPN's are seen by which Port as shown in Figure 9-65. In this example the setup is correct, and this front end port of ProtecTIER sees the WWPN of our Frankfurt node.




Figure 9-65 Scan Port

## Share FE port consideration

In certain product environments, a ProtecTIER system might be shared by a lot of backup servers or LAN-free clients. You can connect or zone each host computer HBA to different front end ports of a ProtecTIER system. But when the host computer has more than the ports' ProtecTIER system, the front end ports must be shared.

In this case LUN-masking should be enabled for ProtecTIER. With LUN-masking you are able to define exactly that what devices should be available for a host.

## 9.9.2 Installing and configuring the device driver

In this section we describe about device driver installation in Linux. For other OS types the information can be found in the following manuals:

- ▶ IBM Tape Device Drivers - Installation and User's Guide, GC27-2130
- ▶ Implementing IBM Tape in UNIX Systems, SG24-6502

### Getting device drivers

The IBM device driver must be installed on all operating systems that will use ProtecTIER virtual drives and libraries. The latest device driver code can be downloaded from the IBM website at this URL:

<http://www-01.ibm.com/support/docview.wss?rs=577&uid=ssg1S4000784>

Or you can go to <http://www.ibm.com> and select **Support & Downloads** → **Download** → **Fixes, updates & drivers**, choose the product, click **Go** on the Support for TS7650G or TS7650 with ProtecTIER page, then select **Device drivers**.

There are additional pages for tape device driver downloads:

- ▶ AIX Page
- ▶ HP-UX Page
- ▶ Linux Page
- ▶ Solaris Page
- ▶ Microsoft Windows Page

Go to the page for the OS that you are using. From here you can download the following documentation:

- ▶ Fixlist
- ▶ README
- ▶ IBM Tape Device Drivers Installation and User's Guide
- ▶ IBM Tape Device Drivers Programming Reference

Then you can go to the device driver ftp site by clicking **IBM Tape Device Drivers FTP site** from the list of directories displayed and selecting your specific device driver.

### Installing IBM lin\_tape driver for Red Hat Linux

In this section we describe how to install the IBM Linux Tape and Medium Changer Device Driver (lin\_tape) on a RedHat Linux Platform and connect it to the ProtecTIER. For tape diagnostic and utility functions, we provide a Tape Utility Program (IBMtapeutil). For more details refer to the following guides:

- ▶ IBM Tape Device Drivers: Installation and User's Guide, GC27-2130
- ▶ Implementing IBM Tape in Linux and Windows, SG24-6268

## Interoperability

Check the IBM Interoperability Matrix to ensure that the version of backup server and the operating system that you are running on are supported for ProtecTIER. Also check the HBA and FW level from the host platform to ensure that your end-to-end environment is supported.

**Note:** Not all backup applications require the IBM Tape Device Driver installation. For certain vendors, the SCSI pass-through or native OS driver is used. Check vendor requirements and ISV.

IBM maintains the latest levels of System Storage tape drive and library device drivers and documentation on the Internet. Obtain them by accessing the following URL:

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

Refer to the SSIC and ISV websites for release information:

- ▶ System Storage Interoperation Center (SSIC) website:  
[http://www-03.ibm.com/systems/support/storage/config/ssic/displayesssearchwithoutjs.wss?start\\_over=yes](http://www-03.ibm.com/systems/support/storage/config/ssic/displayesssearchwithoutjs.wss?start_over=yes)
- ▶ Independent Software Vendors (ISV) website:  
<http://www-03.ibm.com/systems/storage/tape/library.html#interoperabilityv>

## The Linux Device Driver (lin\_tape)

The lin\_tape and medium changer device driver is designed specifically to take advantage of the features provided by the IBM tape drives and medium changer devices. The goal is to give applications access to the functions required for basic tape operations (such as backup and restore) and medium changer operations (such as mount and demount the cartridges), and to the advanced functions needed by full tape management systems. Whenever possible, the driver is designed to take advantage of the device features transparent to the application.

The software described in this section covers the Linux Device Driver (lin\_tape device driver) and the interface between the application and the tape device.

Figure 9-66 illustrates a typical data flow process.




Figure 9-66 Data flow for Linux Device Driver (lin\_tape)

The lin\_tape device driver for Linux is provided in a source rpm package. The utility tools for lin\_tape are supplied in binary rpm packages. They will be downloaded with the driver.

**Note:** Latest packages can be accessed here, at Fix Central:

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

To install your device driver Linux lin\_tape, see details in the IBM device driver users guide.

## Scan all devices

Go to the directory `lin_tape` was installed (in our case its `/tmp`), and run `itdt` as shown in Example 9-19. When you run it for the first time, you will need to read the User Licence Agreement and agree to it.

### Example 9-19 Scanning all devices

```
[root@frankfurt ~]# cd /tmp
[root@frankfurt tmp]# ./itdt
```

You will go to the Entry Menu as shown in Example 9-20.

### Example 9-20 Entry Menu

IBM Tape Diagnostic Tool Standard Edition - V4.1.0 Build 026

Entry Menu

[S] Scan for tape drives (Diagnostic/Maintenance Mode)  
[U] Tapeutil (Expert Mode)

[H] Help  
[Q] Quit program

Notes:

- During a test, user data on the cartridge will be erased!
- Make sure no other program is accessing the devices used by ITDT!
- A device scan may take several minutes in some cases!
- Q + Enter will always close this program.
- H + Enter will display a Help page.

<[H] Help | [Q] Quit | Command > s

Enter `S` to scan for tape drives. During the scan, you will see the window shown in Example 9-21.

### Example 9-21 Scan execution

IBM Tape Diagnostic Tool Standard Edition - Device List

	Host	Bus	ID	LUN	Model	Serial	Ucode	Changer	
0									
1									
2									
3									
4									
5									
6									
7									
8									

9												
10												
11												
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+												

Scan running...

+-----+

+-----+

<[Q] Quit | Command >

After the scan is finished you will see the populated window shown in Example 9-22. In our case the Library has two robots and ten drives.

Example 9-22 post results

IBM Tape Diagnostic Tool Standard Edition - Device List

	Host	Bus	ID	LUN	Model	Serial	Ucode	Changer	[#]
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									
0	6	0	0	0	03584L32	75159990402	0100		
1	6	0	0	1	ULT3580-TD3	1497515000	5AT0	75159990402	
2	6	0	0	2	ULT3580-TD3	1497515002	5AT0	75159990402	
3	6	0	0	3	ULT3580-TD3	1497515004	5AT0	75159990402	
4	6	0	0	4	ULT3580-TD3	1497515006	5AT0	75159990402	
5	6	0	0	5	ULT3580-TD3	1497515008	5AT0	75159990402	
6	8	0	0	0	03584L32	75159990402	0100		
7	8	0	0	1	ULT3580-TD3	1497515001	5AT0	75159990402	
8	8	0	0	2	ULT3580-TD3	1497515003	5AT0	75159990402	
9	8	0	0	3	ULT3580-TD3	1497515005	5AT0	75159990402	
10	8	0	0	4	ULT3580-TD3	1497515007	5AT0	75159990402	
11	8	0	0	5	ULT3580-TD3	1497515009	5AT0	75159990402	
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									

[S] Scan

[T] Test

[D] Dump

[F] Firmware Update

[E] Encryption

[W] Full Write

[U] Tape Usage

[O] Other...

<[H] Help | [Q] Quit | + | - | Line # | Command >

After the scan, you can view the same output by checking /proc/scsi/IBMtape and /proc/scsi/IBMchanger, or /proc/scsi/IBM\* as shown in Example 9-23.

Example 9-23 device display

[root@frankfurt /]# cat /proc/scsi/IBMtape

lin\_tape version: 1.41.1

lin\_tape major number: 251

Attached Tape Devices:

Number	model	SN	HBA	FO Path
0	ULT3580-TD3	1497515000	lpfc	NA
1	ULT3580-TD3	1497515002	lpfc	NA
2	ULT3580-TD3	1497515004	lpfc	NA
3	ULT3580-TD3	1497515006	lpfc	NA
4	ULT3580-TD3	1497515008	lpfc	NA
5	ULT3580-TD3	1497515001	lpfc	NA

```

6      ULT3580-TD3 1497515003      lpfc      NA
7      ULT3580-TD3 1497515005      lpfc      NA
8      ULT3580-TD3 1497515007      lpfc      NA
9      ULT3580-TD3 1497515009      lpfc      NA
[root@frankfurt /]# cat /proc/scsi/IBMchanger
lin_tape version: 1.41.1
lin_tape major number: 251
Attached Changer Devices:
Number  model      SN              HBA              FO Path
0       03584L32    0014975159990402 lpfc              NA
1       03584L32    0014975159990402 lpfc              NA

```

---

You will find the same devices and serials in the ProtecTIER GUI Library General view shown in Figure 9-67.

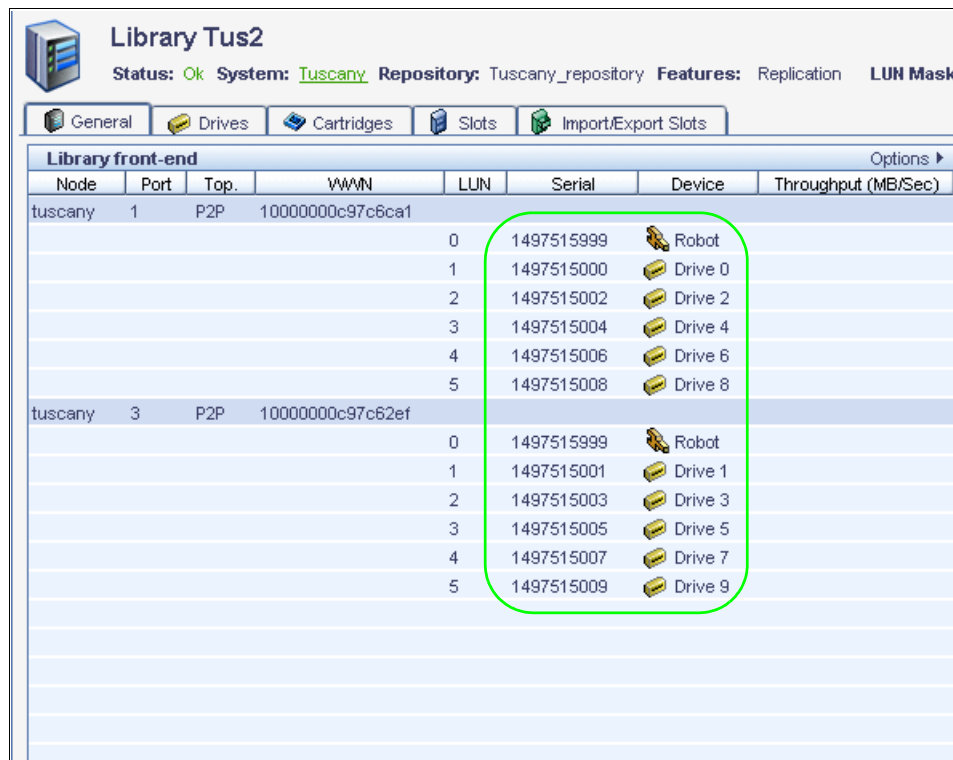


Figure 9-67 Library General view

**Note:** The robots has the same serial number. It is important to use the appropriate driver in case Control Path Failover (CPF) is enabled, as ISV matrix is describing it.

## 9.10 Deduplication operation with ProtecTIER

In this section we describe deduplication setup for VTL and OST, how to tune deduplication, and where you can expect better deduplication.

## 9.10.1 Deduplication with VTL

By default, ProtecTIER factors all new incoming data, detecting recurring data and storing only the data segments that have not previously been written to the repository. By default deduplication is enabled when each library is created as shown by **Hyperfactor mode** being **Enabled** as seen in Figure 9-68.

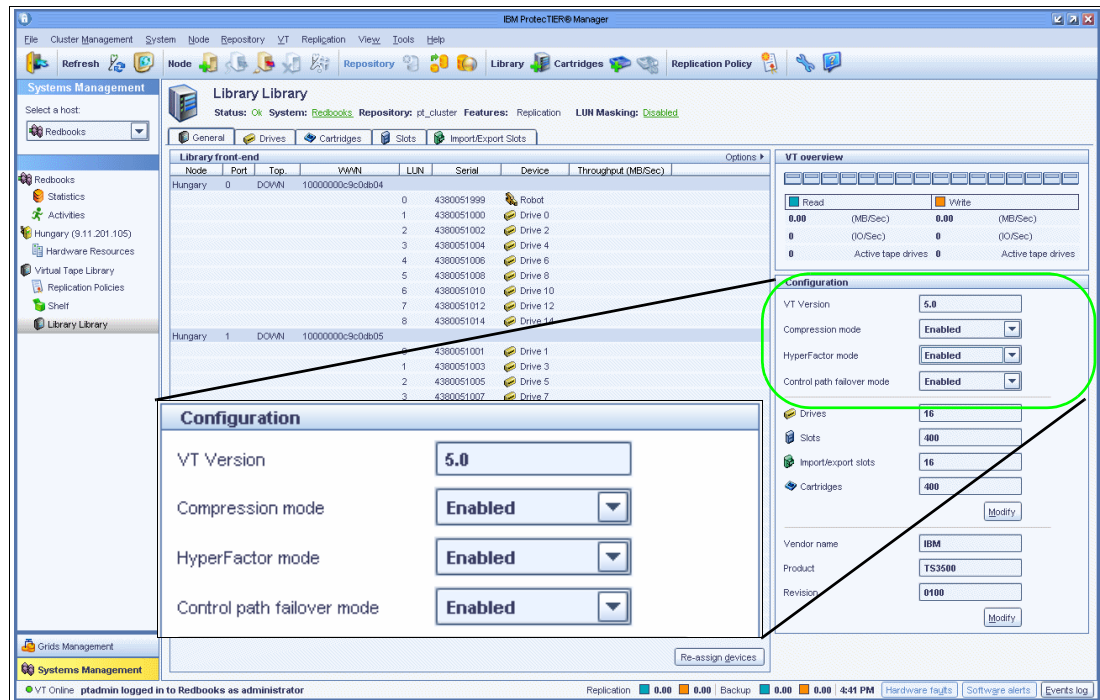


Figure 9-68 ProtecTIER Manager: Library configuration view, top part

1. Change this by clicking the drop down button next to HyperFactor mode in the Library view or select the appropriate library in the navigation pane, then in the top menu click **VT** → **VT Library** → **Set HyperFactor mode**.

You will see the dialog shown on Figure 9-69.




Figure 9-69 HyperFactor mode dialog window

2. Select one of the following options, as directed by IBM Support:
  - HyperFactor enabled: HyperFactor operates as normal.
  - HyperFactor disabled: HyperFactor stops. When you restart HyperFactor, the HyperFactor process proceeds as normal based on the data stored from before HyperFactor stopped.

- Baseline: HyperFactor stops factoring incoming data and uses the newly stored non-factored data as the reference for factoring new data after HyperFactor is resumed.
3. Click **OK**. The ProtecTIER VT Hyperfactor mode window closes.

### 9.10.2 Compression with VTL

Under normal circumstances, the ProtecTIER system compresses data. You can stop compression on a specific virtual library to free the resources usually demanded by the compression process by performing these steps:

1. Select **VT** → **VT Library** → **Set compression mode**. The ProtecTIER compression mode window is displayed (Figure 9-70).

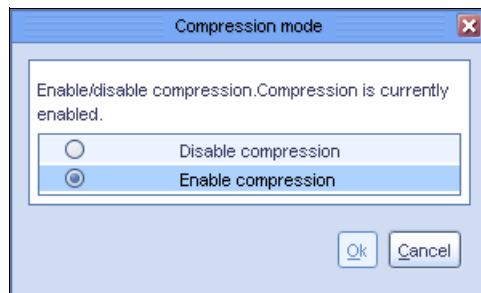


Figure 9-70 Compression mode window

2. Select **Disable compression** and click **OK**. The ProtecTIER compression mode dialog closes and compression is stopped.
3. Selecting **Enable compression** on the ProtecTIER compression mode dialog resumes data compression.

**Note:** The compression done by ProtecTIER is called Diligent® Compression.

For compressed data streams, create a new ProtecTIER VTL with compression turned off. Compressing data a second time can cause data expansion, so compressed data should be segregated in ProtecTIER whenever possible.

### 9.10.3 Deduplication with OST

With OST, you will be able to control the HyperFactor in a similar way, like for VTL.

For Open Storage (OST), in case of each LSU, you can select the HyperFactor mode:

- ▶ enable
- ▶ disable
- ▶ baseline

See more details in 9.7.4, “Creating Logical Storage Units (LSUs)” on page 208.

### 9.10.4 Compression with OST

For Open Storage (OST), in case of each LSU, you can select the Compression mode:

- ▶ enable



- disable

For more details, see 9.7.4, “Creating Logical Storage Units (LSUs)” on page 208.

## 9.11 Backup and restore applications

We will describe general setups for backup servers, and will describe one example. In our case it will be IBM Tivoli Storage Manager.

### 9.11.1 Suggestions for all backup servers

Many backup servers have features and settings that can be used to optimize performance when writing data to physical cartridges. Because ProtecTIER provides a virtual tape library with virtual drives and cartridges to the backup server, some of the settings that are optimized for real tape are no longer required, and might even have a detrimental effect on the ProtecTIER factoring ratio and performance.

The following suggestions are fairly generic and are common to all backup servers. Check the current settings of your backup server and apply the settings that can be implemented.

The backup server specific sections of this chapter have more detailed suggestions than these general topics. When this is the case, the specific recommendation should take precedence and be followed closely.

### 9.11.2 General suggestions

As a general rule, the preferred method of operation is to imitate the procedure used with physical cartridges.

Implement the time frame mode of operation so that for every 24-hour cycle there is a backup window and then a replication window. Make sure that there is enough bandwidth and time allotted so that there will be no overlap and no replication backlog.

Here is a typical operational flow:

1. Perform regular daily backups to the ProtecTIER system during the defined backup window.
2. After the daily backups are complete, perform a full catalog/DB backup to cartridge into ProtecTIER.
3. The system should be set up so that replication will start and be finished before the next backup cycle starts.
4. This process creates a complete and easily recoverable set of your latest daily backup, including the backup application catalog image.
5. In case of a disaster, you can revert back to your last completed set of backups, so the recovery point objective (RPO) is within the 24-hour window that is typical for the service level agreement (SLA).

These are general suggestions that are not specific to any backup server. More specific suggestions are made in the following sections.

## Interoperability

Check the IBM Interoperability Matrix to ensure that the version of the backup server and the operating system that you are running on are supported for ProtecTIER. You can view the matrix at the following URL:

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

## Software compatibility

Make sure that your backup server version, platform, and operating system version are on the supported hardware and software list for ProtecTIER. You can view the list at the following URL:

<http://www-03.ibm.com/systems/storage/tape/library.html#compatibility>

## Software currency

Ensure that the backup server has the latest patches or maintenance level to improve the overall factoring performance.

Ensure that the operating system of the platform that the backup server is running on has the latest patches or maintenance level to improve the overall HyperFactor performance.

## Tape library zoning

The backup server should have dedicated HBA ports for the ProtecTIER virtual tape library. These ports can be shared with a physical tape library. However, the physical tape library must not be in the same SAN zone as the virtual tape library.

## Compression

Standard compression will effectively scramble the data sent to ProtecTIER, making pattern matching difficult; real-time compression in databases and IBM Real-Time Compression does not scramble this data. As can be expected, this situation will have an effect on data-matching rates, even if the same data is sent each time. ProtecTIER will compress the data that it sends to the back-end physical disk after it has been received by the virtual tape drives and deduplicated.

If a data stream is not already compressed, our experience suggests that it is most efficient to let ProtecTIER compress data after deduplication. Compressed data can have fewer duplicate blocks than uncompressed data, so the effectiveness of deduplication can be diminished. Workloads and results vary, so we encourage experimentation.

Compressed data can reach backup systems and ProtecTIER in a number of ways. Consider the following use cases where the data stream going to ProtecTIER remains compressed:

- ▶ File backups may contain files from remote systems that enable compression in their backup client software to conserve network bandwidth. File system backups may also contain compressed file formats, such as GIF images, which remain compressed until opened by a GIF reader.
- ▶ Block-level database backups may contain compressed database objects. Database compression features deliver about 5:1 data reduction for Oracle, DB2, and Informix, and a bit less for databases with free compression.
- ▶ NDMP image backups may contain data compressed by IBM Real-time compression or another process. NDMP uses a back channel to create backups, so IBM Real-time compression or other rehydration processes are bypassed.

Actions you can take regarding data being compressed before reaching ProtecTIER include:

- ▶ ProtecTIER does not require additional compression to be effective. ProtecTIER performs compression, by default, after the deduplication process. Do not allocate disk for ProtecTIER VTLs that compresses data by default.
- ▶ ProtecTIER can manage multiple VTLs, each with its own configuration. For compressed data streams, create a new ProtecTIER VTL with compression turned off. Compressing data a second time can cause data expansion, so compressed data should be segregated in ProtecTIER whenever possible.
- ▶ File systems with small files (under 32 KB), whether or not they are compressed, should not be sent directly to ProtecTIER. The following options should be considered to prevent ProtecTIER from bypassing small files:
  - Large NAS systems with small files should use NDMP for image backups and then send those files to ProtecTIER.
  - File level backups should first back up to backup application Disk Storage Pools, and those Disk Storage Pools can be copied to ProtecTIER.
- ▶ If a data stream is not already compressed, it is usually most efficient to let ProtecTIER compress data after deduplication. Compressed data can have fewer duplicate blocks than uncompressed data, so the effectiveness of deduplication can be diminished. Workloads and results vary, so we encourage experimentation.
- ▶ Always encrypt last. Deduplicating encrypted data is ineffective. Compressing encrypted data can decrease security. Drive-level encryption has no performance impact, and it ensures that encryption occurs last.

## Encryption

Encryption makes each piece of data sent to ProtecTIER unique, including duplicate data. As can be expected, this has an effect on data matching rates and the factoring performance because even if the same data is sent each time, it will appear different to the deduplication engine. Therefore, disable any encryption features for the ProtecTIER storage pool in the backup server.

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

## Multiplexing

Do not use the multiplexing feature of any backup application with the ProtecTIER storage pool. Although ProtecTIER will work with these features, the benefits (disk savings) of the HyperFactor algorithm and compression will be greatly reduced.

Therefore, disable any multiplexing features in the backup server for the ProtecTIER storage pool.

## Tape block sizes

In order to optimize the backup server, set the block size for data sent to the (virtual) tape drives to be at least 256 KB.

## Operating system clock time

If possible, ensure that the system time on the backup servers and ProtecTIER systems are synchronized by the same source, such as an NTP server or other means. This makes any problem diagnosis activities easier to conduct, should they be required.

## Ports

If possible, provision the HBA ports by connecting to the IBM System Storage TS7650 or TS7650G storage system with no other devices (disk or real tape) connected or zoned to those ports.

Ensure that the HBAs in the backup server are spread across all the PCI buses.

## Additional factors

Another factor that affects performance in a ProtecTIER environment is the type of data being targeted for backup. Some data is well suited to data deduplication and other data is not. For example, small files (less than 32 KB in size) commonly found in operating systems do not factor well, although the built-in compression might reduce their stored size. You might want to consider some of the following options:

- ▶ Larger NAS systems should use NDMP for image backups and then be sent to ProtecTIER.
- ▶ File level backups should first back up to backup application Disk Storage Pools, and then the Disk Storage Pool can be copied to ProtecTIER.

For known configuration changes suited to specific data types, see 9.11.3, “Data types” on page 240.

### 9.11.3 Data types

One factor that affects performance in a ProtecTIER environment is the data that is being targeted for backup.

Some data, such as databases and email, are highly compressible and also factor quite well. Other data, such as video or seismic data, cannot be compressed or factored well.

The following suggestions are for the data type or application that sends backup data to your backup server.

#### Oracle database data

When using database backup tools such as Oracle Recovery Manager (RMAN), make sure to check the multiplexing-like options. RMAN by default has settings to enable multiplexing, MAXOPENFILES and FILESPERSET controls the level of multiplexing.

Multiplexing means that several files are sent into one stream, weaving together, creating one backup set that will usually have a high change rate. For example, prior to ProtecTIER being installed, RMAN sends a total of nine files requiring backup to three channels. This equates to three file streams, each containing three files ( $3 \times 3 = 9$ ) to three physical tape drives.

For performance reasons encryption and compression should be turned off.

#### ***Most effective performance and optimal factoring***

Use the following setups to enable the maximum performance and optimal factoring:

- ▶ MAXOPENFILES =1  
Controls the number of data files to be read from the DB concurrently through a single backup, therefore it needs to be set to 1 preventing multiplexing.
- ▶ FILESPERSET=4  
Defines the number of data files to be included in a backup set.

- ▶ **PARALLELISM=32** (up to 64)

Controls the number of tape drives that will be used for backup. Use more parallel backup streams to achieve higher performance. Make sure to have the amount of virtual tape drives available within ProtecTIER library to match the number of parallel streams configured within RMAN.

### ***Most effective factoring***

Use the following setup for the best factoring because this will disable multiplexing completely: for each backup command, set the value to **FILESERSET=1**, **MAXOPENFILES=1**, and **PARALLELISM=32** (up to 64).

### **SQL Litespeed data**

SQL Litespeed backups might not factor well. This is because Litespeed compresses and encrypts the SQL database before backing it up. This is for your information only. There are no actions that can be taken to improve this process.

### **Domino**

Compaction is a function of Domino to reduce the capacity of primary space the NSF files take. This is a best practice used in almost all Domino environments. Compaction shuffles the data in the NSF file so it takes less space. The data will be considered new because it will have almost 100% of change rate. There could be two solutions to overcome this:

- ▶ **Disable compaction**

This solves the problem, but is contrary to Domino's best practice.

- ▶ **Upgrade Domino to 8.5 and enable DAOS (Domino Attachment and Object Service)**

DAOS allows to store only one instance of an attachment as opposed to having multiple copies stored and written to disk. Of course, if the spreadsheet or attachment is modified, DAOS will store the newly modified copy as a separate version. Moreover, attachment consolidation is not limited to mail—it occurs as soon as an attachment is saved in any document of any database on the server on which the feature is enabled.

In a standard Notes database (NSF), the attachments are stored inside of the NSF file itself, and the database is self-contained. In order to back up a standard Notes database, only the NSF file itself needs to be backed up. After you introduce DAOS, the NSFs that participate in DAOS contain only references to the NLO files where the attachment content is stored. As a result, backing up the NSF alone is no longer enough. The NLO data needs to be backed up as well.

These NLO files could be a good candidate for deduplication.

### **SQL**

Revise the index defragmentation timing for SQL, and create a monitor procedure for it. Do *not* do index defragmentation daily, but change it to weekly, monthly or defragmentation based on threshold setup, where you define exactly at what level or percentage it should run. This would lead to a better factoring ratio because the data change rate will not be high.

For blocksize, use larger block sizes above 256KB. Because in SQL native backup you might not have the option to change it higher than 64KB, you will have to use the backup application options for that. For example if you use ITSM rather than SQL native backup, you can set up a larger block size in ITSM.

You should check settings for **BUFFERCOUNT**, **MAXTRANSFERSIZE**, and **STRIPES**. By tuning these you can eliminate some of the data shuffling associated with SQL.

## DB2

The current DB2 version involves built-in multiplexing which significantly reduces a deduplication solution's factoring ratio (and performance). Improvements to optimize factoring can be managed with larger buffer sizes. Backup performance, especially for large database files, can be achieved by increasing the parallelism of the backup streams. These split files can then be backed up with multiple ITSM backup sessions. We recommend to use maximal buffers, minimum parallelism, and minimum number of sessions as seen in Example 9-24.

### *Example 9-24*

---

```
db2 backup db abc use tsm open 8 sessions with 18 buffers buffer 16384 parallelism 8
```

---

### **Buffer size**

The size of the buffer used for a backup stream. A value of 16384 will provide best factoring, but may require more memory than is available. Adjust if necessary, for example the session in Example 9-24 requires 1.2G of memory. If that is too much, use buffer 4097 instead.

### **Parallelism**

Determines the number of DB2 threads used to handle the backup process. Set this value to the minimum number that still allows backups to complete in time. Start with 24 and adjust as necessary.

### **Sessions**

Determines the number of ProtecTIER tape drives used to write the backup. Set this value to the minimum number that still allows backups to complete in time. Start with 8 and adjust as necessary.

### **With [number] buffers**

The number of buffers to be used. DB2 will automatically choose an optimal value for this parameter unless you explicitly enter a value. Change value to be equal with (parallelism + sessions + 2).

## 9.11.4 IBM Tivoli Storage Manager

IBM Tivoli Storage Manager (ITSM) is a storage software suite that addresses the challenges of complex storage management in distributed heterogeneous environments. It protects and manages a broad range of data, from workstations to the corporate server environment.

When performing standard backups, IBM Tivoli Storage Manager uses the incremental backup method. Progressive backup methodology (often referred to as incremental backup) saves time and storage space by backing up only new and modified files. The progressive backup feature uses the IBM Tivoli Storage Manager DB2 database to track data wherever it is stored, delivering direct one-step file restore. Progressive backup eliminates the need for traditional full-plus-incremental or full-plus-differential backup and restore procedures commonly used by other storage management products.

There are some additional IBM Tivoli Storage Manager products designed specifically to interact with third-party software products, such as email (Microsoft Exchange) and databases (Oracle or MS-SQL). These are known as Tivoli Data Protection (TDP) modules.

Other suggestions are:

- ▶ For improved performance in terms of factoring, use IBM Tivoli Storage Manager V5.4 or later.
- ▶ Client compression should be disabled.

- ▶ When using Windows-based IBM Tivoli Storage Manager servers, the IBM Tivoli Storage Manager driver for tape and libraries for Windows must be used. Native Windows drivers for the emulated P3000 libraries and DLT7000 drives are not supported.

Figure 9-71 illustrates a typical ITSM environment using ProtecTIER. The ITSM environment is straightforward. The ITSM servers are connected to storage devices (disk, real tape, or virtual tape) that are used to store data backed up from the clients it is serving. Every action and backup set that ITSM processes is recorded in the ITSM database. Without a copy of the ITSM database, a ITSM server cannot restore any of the data that is contained on the storage devices.

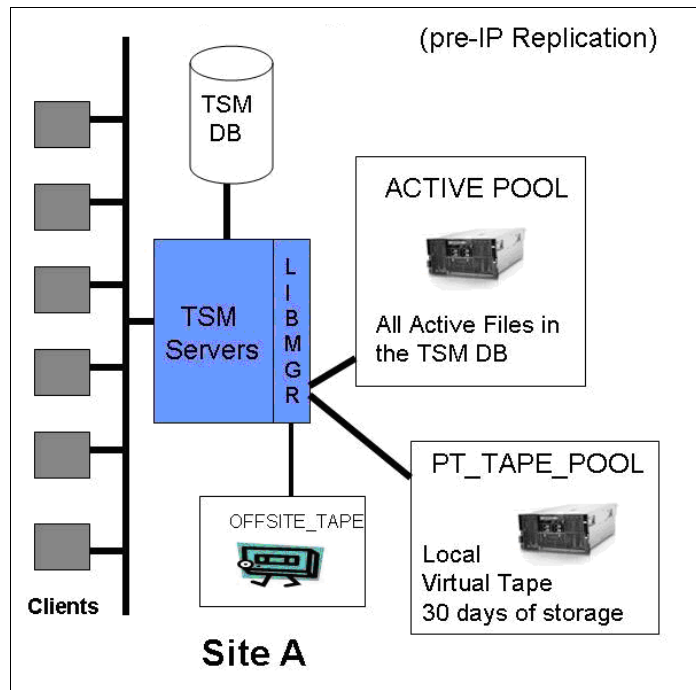


Figure 9-71 Typical ITSM environment with ProtecTIER

ProtecTIER provides a virtual tape interface to the ITSM servers and allows the creation of two storage pools:

- ▶ The ACTIVE TSM pool
- ▶ The ONSITE TAPE pool (called PT\_TAPE\_POOL in Figure 9-71)

The user can also maintain another storage pool to create real physical tapes to take offsite (called OFFSITE\_TAPE in our example). The user has sized the ProtecTIER system to store all active and about 30 days of inactive client files on virtual tape. The user also created an ACTIVE TSM pool, which is also hosted on the ProtecTIER system, which contains the most recent (active) file backed up from all client servers. The ACTIVE pool is where client restores will come from. The advantage of this architecture is that it has eliminated the use of physical tape in the data center and allows restores to occur much faster, as they are coming from the ProtecTIER disk-based virtual tape versus real tape.

### Implementing a virtual tape library

In the following sections, we walk through the steps required to define and connect a virtual tape library (VTL) in ProtecTIER to a IBM Tivoli Storage Manager server and make it usable.

For the purposes of this example, the ITSM server is named server1 and is running 6.1 as shown in Example 9-25. The host server is called frankfurt.storage.tucson.ibm.com and is running Red Hat Enterprise Linux 5.4.

*Example 9-25*

---

```
tsm: SERVER1>q stat
Storage Management Server for Linux/x86_64 - Version 6, Release 1, Level 2.0

Server Name: SERVER1
Server host name or IP address:
Server TCP/IP port number: 1500
Crossdefine: Off
Server Password Set: No
Server Installation Date/Time: 08/30/2010 16:56:11
Server Restart Date/Time: 09/24/2010 16:22:03
Authentication: On
Password Expiration Period: 90 Day(s)
Invalid Sign-on Attempt Limit: 0
Minimum Password Length: 0
Registration: Closed
Subfile Backup: No
Availability: Enabled
Accounting: Off
Activity Log Retention: 30 Day(s)
Activity Log Number of Records: 6134
Activity Log Size: <1 M
Activity Summary Retention Period: 30 Day(s)
more... (<ENTER> to continue, 'C' to cancel)
```

---

### **The ProtecTIER virtual tape library definition**

Using ProtecTIER Manager, a virtual tape library named Tus2 has already been created for use by IBM Tivoli Storage Manager (Figure 9-72 on page 245). This library name is an internal name used only within ProtecTIER and is not visible anywhere outside of it. For details about how to create a virtual tape library using ProtecTIER Manager, refer to 9.8, “Virtual library creation” on page 216.



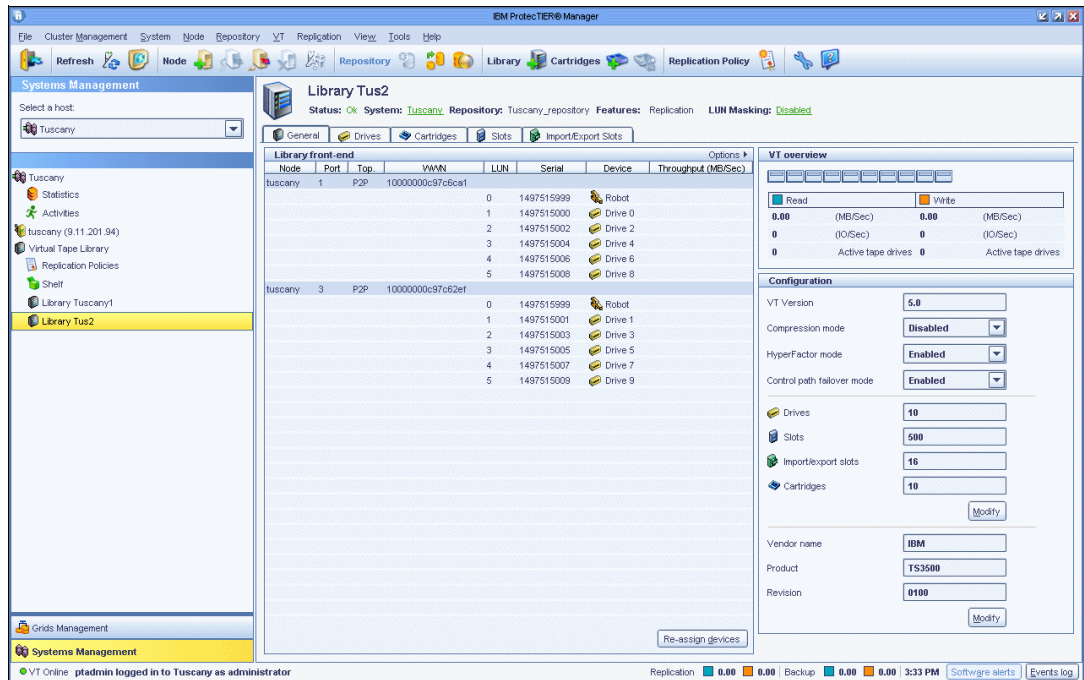


Figure 9-72 ProtecTIER Library view

The library was given ten virtual LTO3 cartridge drives. The tape drives are divided between two ports. The virtual robot is defined on the ProtecTIER node called tuscany.

Five hundred library slots were created along with sixteen import/export (I/E) slots for completeness of the library emulation. Ten virtual cartridges were created also. See more details of the virtual LTO3 drive definitions in Figure 9-73.

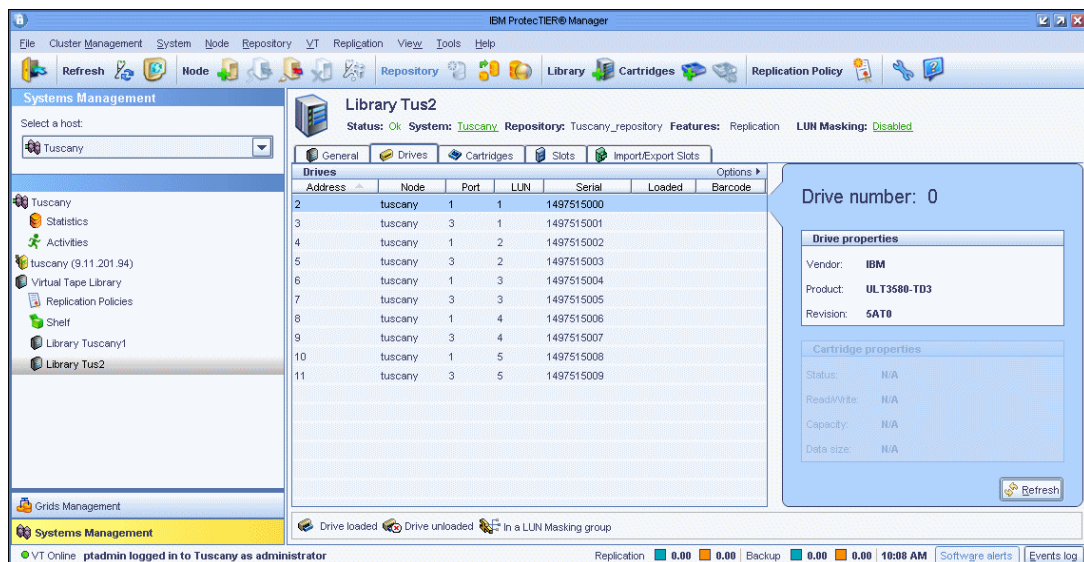


Figure 9-73 ProtecTIER Drive view in Tus2 Library

Note the logical unit number (LUN) assigned to each virtual tape drive. It is important to know that this is the ProtecTIER LUN number only and the host operating system will almost certainly assign each drive a different LUN than that which appears here, as the host will have more logical units or resources than just these tape drives.

During library definition, ProtecTIER has assigned serial numbers to each drive, seeded from a random number. These can be important later when defining IBM Tivoli Storage Manager paths to link the host devices and the IBM Tivoli Storage Manager tape drive definitions. It can be helpful to know the drive serial number if you should experience any problems when defining the paths. You can use the serial number as a common reference point when matching a drive with its host device file name.

Element number is the other address of each object in the tape library. Everything that can be located (such as robot, drive, slot, and io slot) has its unique element number in the library. ITSM (and any other tape library related software) uses this number as the address in the SCSI command to drive the robot to work. The element number can also be important later when defining the IBM Tivoli Storage Manager tape drives.

Ten virtual cartridges are added into the library. They are set up with ‘Max. cartridge growth’ of 100 GB. One of the cartridge has data in this view (A00001L3). All of them are R/W: read and write capable as shown in Figure 9-74. The capacity, Data size and Max size all have the nominal values.

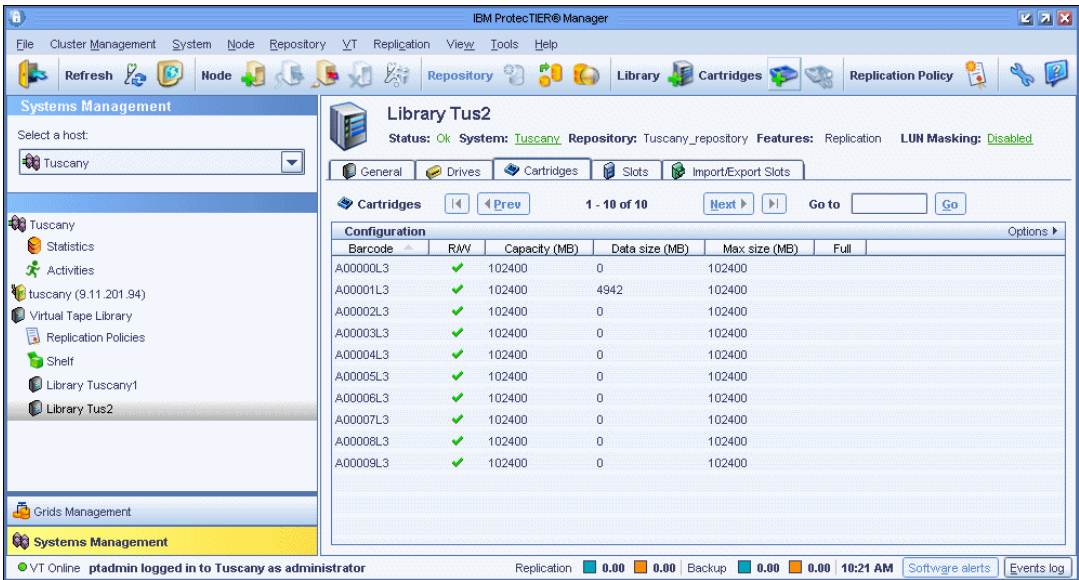


Figure 9-74 virtual cartridges

## Defining the virtual tape library to Linux with IBM Tivoli Storage Manager

The steps needed to define a ProtecTIER virtual tape library and drives to IBM Tivoli Storage Manager are identical to those required for the corresponding physical tape library and drives.

To define a physical or virtual library in ITSM, perform the following steps:

1. Install Device Drivers for tape, as demonstrated in 9.9.2, “Installing and configuring the device driver” on page 230.
2. Run the scan from itdt, as shown in “Scan all devices” on page 232. You will get the output shown in Example 9-26.

Example 9-26 Running scan from itds

IBM Tape Diagnostic Tool Standard Edition - Device List									
Host	Bus	ID	LUN	Model	Serial	Ucode	Changer	[#]	

0	6	0	0	0	03584L32	75159990402	0100		
1	6	0	0	1	ULT3580-TD3	1497515000	5AT0	75159990402	
2	6	0	0	2	ULT3580-TD3	1497515002	5AT0	75159990402	
3	6	0	0	3	ULT3580-TD3	1497515004	5AT0	75159990402	
4	6	0	0	4	ULT3580-TD3	1497515006	5AT0	75159990402	
5	6	0	0	5	ULT3580-TD3	1497515008	5AT0	75159990402	
6	8	0	0	0	03584L32	75159990402	0100		
7	8	0	0	1	ULT3580-TD3	1497515001	5AT0	75159990402	
8	8	0	0	2	ULT3580-TD3	1497515003	5AT0	75159990402	
9	8	0	0	3	ULT3580-TD3	1497515005	5AT0	75159990402	
10	8	0	0	4	ULT3580-TD3	1497515007	5AT0	75159990402	
11	8	0	0	5	ULT3580-TD3	1497515009	5AT0	75159990402	

[S] Scan            [T] Test            [D] Dump            [F] Firmware Update  
[E] Encryption    [W] Full Write    [U] Tape Usage    [O] Other...

<[H] Help | [Q] Quit | + | - | Line # | Command >

3. Obtain information about the tape devices as seen in Example 9-27.

*Example 9-27 Device information*

```
[root@frankfurt /]# cat /proc/scsi/IBMtape
lin_tape version: 1.41.1
lin_tape major number: 251
Attached Tape Devices:
Number  model      SN              HBA              FO Path
0       ULT3580-TD3 1497515000      lpfc              NA
1       ULT3580-TD3 1497515002      lpfc              NA
2       ULT3580-TD3 1497515004      lpfc              NA
3       ULT3580-TD3 1497515006      lpfc              NA
4       ULT3580-TD3 1497515008      lpfc              NA
5       ULT3580-TD3 1497515001      lpfc              NA
6       ULT3580-TD3 1497515003      lpfc              NA
7       ULT3580-TD3 1497515005      lpfc              NA
8       ULT3580-TD3 1497515007      lpfc              NA
9       ULT3580-TD3 1497515009      lpfc              NA
[root@frankfurt /]# cat /proc/scsi/IBMchanger
lin_tape version: 1.41.1
lin_tape major number: 251
Attached Changer Devices:
Number  model      SN              HBA              FO Path
0       03584L32   0014975159990402 lpfc              NA
1       03584L32   0014975159990402 lpfc              NA
```

4. You can read the Device IDs or run an Element inventory after opening the device in itdd.
  - a. Start itdd, and choose tapeutil (U) as shown in Example 9-28.

*Example 9-28 Element inventory*

```
[root@frankfurt tmp]# ./itdd
Please wait for startup completion.... (Q to quit)
```

# IBM Tape Diagnostic Tool Standard Edition - V4.1.0 Build 026

## Entry Menu

[S] Scan for tape drives (Diagnostic/Maintenance Mode)  
 [U] Tapeutil (Expert Mode)

[H] Help  
 [Q] Quit program

## Notes:

- During a test, user data on the cartridge will be erased!
- Make sure no other program is accessing the devices used by ITDT!
- A device scan may take several minutes in some cases!
- Q + Enter will always close this program.
- H + Enter will display a Help page.

<[H] Help | [Q] Quit | Command > u

b. After that Open a device (1) as shown in Example 9-29.

## Example 9-29 Command menu

----- General Commands: -----		
[1] Open a Device	[5] Reserve Device	[9] Mode Sense
[2] Close a Device	[6] Release Device	[10] Query Driver
Ver.		
[3] Inquiry	[7] Request Sense	[11] Display All
Paths		
[4] Test Unit Ready	[8] Log Sense	
----- Tape Drive Commands: -----		
[20] Rewind	[28] Erase	[36] Display Message
[21] Forward Space Filemarks	[29] Load Tape	[37] Report Density
Supp		
[22] Backward Space Filemarks	[30] Unload Tape	[38] Test Encryp.
Path		
[23] Forward Space Records	[31] Write Filemarks	[39] Config. TCP/IP
Port		
[24] Backward Space Records	[32] Synchronize Buffers	
[25] Space to End of Data	[33] Query/Set Parameter	
[26] Read and Write Tests	[34] Query/Set Tape Position	
[27] Read or Write Files	[35] Query Encryption Status	
----- Tape Library Commands: -----		
[50] Element Information	[55] Initialize Element Status	
[51] Position to Element	[56] Prevent/Allow Medium Removal	
[52] Element Inventory	[57] Initialize Element Status Range	
[53] Exchange Medium	[58] Read Device IDs	
[54] Move Medium	[59] Read Cartridge Location	

```
----- Service Aid Commands: -----
[70] Dump/Force Dump/Dump      [71] Firmware Update
```

```
<[H] Help | [Q] Quit | Command > 1
```

- c. Type in the correct device, in our case **/dev/IBMchanger0**. See Example 9-30.

*Example 9-30 Selecting device*

```
ITDT- Open a Device
```

```

Device Name      +-----+
                  | /dev/IBMtape0 |
                  +---+---+
Mode              | 1 | 1=Read/Write, 2=Read Only, 3=Write Only,
4=Append)         +---+
                  +---+
```

```
<Specify Device Name | [enter] for /dev/IBMtape0> /dev/IBMchanger0
```

- d. Select **1**, for the mode Read/Write as shown in Example 9-31.

*Example 9-31 Mode selection*

```
ITDT- Open a Device
```

```

Device Name      +-----+
                  | /dev/IBMchanger0 |
                  +---+---+
Mode              | 1 | 1=Read/Write, 2=Read Only, 3=Write Only,
4=Append)         +---+
                  +---+
```

```
<Specify Mode | [enter] for 1> 1
```

- e. Press **Enter** to execute, and you get the command result as shown in Example 9-32.

*Example 9-32 Execution*

```
Command Result
```

```

+-----+
--+
| Opening Device...
| Open Device PASSED
|
| Device:/dev/IBMchanger0 opened
|
+-----+
--+
```

< [Q] Quit | [N] Next | [P] Previous | + | - | [Enter] Return >

- f. Press **Enter** to return to main menu, and select **50** Element information as shown in Example 9-33. Here you will see the dimensions of the library and the starting element numbers.

*Example 9-33 Element information*

```

Command Result
+-----+
| Getting element information...
| Read Element Information PASSED
|
|
| Number of Robots ..... 1
| First Robot Address ..... 0
| Number of Slots ..... 500
| First Slot Address ..... 1026
| Number of Import/Exports ..... 16
| First Import/Export Address .... 64514
| Number of Drives ..... 10
| First Drive Address ..... 2
|
+-----+
< [Q] Quit | [N] Next | [P] Previous | + | - | [Enter] Return >

```

- g. Press **Enter** to go back to the menu. Now you can run **58** Read Device IDs. You will get the results, shown in Example 9-34. To go to next page within tapeutil, you have to press **n** (next). In our example we show it only for two drives.

*Example 9-34 Read device ids*

```

Command Result
+-----+
| Reading element device ids...
| Read Device Ids PASSED
|
|
| Drive Address 2
| Drive State ..... Normal
| ASC/ASCQ ..... 0000
| Media Present ..... No
| Robot Access Allowed ..... Yes
| Source Element Address Valid .. No
| Media Inverted ..... No
| Same Bus as Medium Changer .... Yes
| SCSI Bus Address Valid ..... No
| Logical Unit Number Valid ..... No
| Volume Tag .....
| Device ID, Length 34
|
|           0 1 2 3 4 5 6 7 8 9 A B C D E F 0123456789ABCDEF
| 0000 - 4942 4D20 2020 2020 554C 5433 3538 302D [IBM    ULT3580-]
| 0010 - 5444 3320 2020 2020 3134 3937 3531 3530 [TD3    14975150]
| 0020 - 3030                                     [00          ]
|
| Drive Address 3
| Drive State ..... Normal
| ASC/ASCQ ..... 0000

```

```

Media Present ..... No
Robot Access Allowed ..... Yes
Source Element Address Valid .. No
Media Inverted ..... No
Same Bus as Medium Changer .... Yes
SCSI Bus Address Valid ..... No
Logical Unit Number Valid ..... No
Volume Tag .....
Device ID, Length 34

      0 1 2 3 4 5 6 7 8 9 A B C D E F 0123456789ABCDEF
0000 - 4942 4D20 2020 2020 554C 5433 3538 302D [IBM      ULT3580-]
0010 - 5444 3320 2020 2020 3134 3937 3531 3530 [TD3      14975150]
0020 - 3031                                     [01          ]

```

**Note:** You can get the serial number as well for each device with the following command:

```

[root@frankfurt tmp]# ./itdt -f /dev/IBMtape0 inquiry 80
Issuing inquiry for page 0x80...
Inquiry Page 0x80, Length 14

0 1 2 3 4 5 6 7 8 9 A B C D E F 0123456789ABCDEF
0000 - 0180 000A 3134 3937 3531 3530 3030 [....1497515000 ]

```

- h. You can choose **52** Element Inventory to see the element numbers. You will get the following results as shown in Example 9-35.

*Example 9-35 Viewing element numbers*

```

Command Result
+-----+
| Reading element status... |
|                             |
| Element Inventory PASSED  |
|                             |
| Robot Address 0           |
| Robot State ..... Normal |
| ASC/ASCQ ..... 0000      |
| Media Present ..... No   |
| Source Element Address Valid .. No |
| Media Inverted ..... No   |
| Volume Tag .....         |
|                             |
| Drive Address 2           |
| Drive State ..... Normal |
| ASC/ASCQ ..... 0000      |
| Media Present ..... No   |
| Robot Access Allowed ..... Yes |
| Source Element Address Valid .. No |
| Media Inverted ..... No   |
| Same Bus as Medium Changer .... Yes |
| SCSI Bus Address Valid ..... No |
| Logical Unit Number Valid ..... No |
| Volume Tag .....         |
|                             |
| Drive Address 3           |

```

```

Drive State ..... Normal
ASC/ASCQ ..... 0000
Media Present ..... No
Robot Access Allowed ..... Yes
Source Element Address Valid .. No
Media Inverted ..... No
Same Bus as Medium Changer .... Yes
SCSI Bus Address Valid ..... No
Logical Unit Number Valid ..... No
Volume Tag .....

Slot Address ..... 1026
Slot State ..... Normal
ASC/ASCQ ..... 0000
Media Present ..... Yes
Robot Access Allowed ..... Yes
Source Element Address ..... 1038
Media Inverted ..... No
Volume Tag ..... A00000L3

Slot Address ..... 1027
Slot State ..... Normal
ASC/ASCQ ..... 0000
Media Present ..... Yes
Robot Access Allowed ..... Yes
Source Element Address ..... 8
Media Inverted ..... No
Volume Tag ..... A00001L3

Slot Address ..... 1028
Slot State ..... Normal
ASC/ASCQ ..... 0000
Media Present ..... Yes
Robot Access Allowed ..... Yes
Source Element Address ..... 3
Media Inverted ..... No
Volume Tag ..... A00002L3

```

5. Using the above information we filled out the tape library worksheet shown in Table 9-3.

Table 9-3 Tape library worksheet

Device in OS	Type	VTL system	VTL node	VTL port WWN	Serial number	Element number
IBMchanger0	3584	Tuscany	Tuscany	10000000c97c6ca1	0014975159990402	0
IBMtape0	LT3	Tuscany	Tuscany	10000000c97c6ca1	1497515000	2
IBMtape1	LT3	Tuscany	Tuscany	10000000c97c62ef	1497515001	3

6. Now you can define the library in IBM Tivoli Storage Manager, using the commands **DEFINE LIBRARY** and **QUERY LIBRARY**. See Example 9-36.

*Example 9-36 Defining library*

```

tsm: SERVER1>def library vtl libtype=scsi relabelscratch=yes
ANR8400I Library VTL defined.
tsm: SERVER1>q libr f=d

```



```

Library Name: VTL
Library Type: SCSI
ACS Id:
Private Category:
Scratch Category:
WORM Scratch Category:
External Manager:
Shared: No
LanFree:
ObeyMountRetention:
Primary Library Manager:
WWN:
Serial Number: 0014975159990402
AutoLabel: No
Reset Drives: No
Relabel Scratch: Yes
Last Update by (administrator): ALEXC
Last Update Date/Time: 09/24/2010 16:20:51

```

---

**Note:** The **RELABELSCRATCH** parameter allows you to automatically relabel volumes when they are returned to scratch. This will ensure that after a cartridge is expired, and returns to scratch, ITSM will relabel it. In this case the space will be freed up in ProtecTIER as well.

7. Define path to the IBM Tivoli Storage Manager library using the **DEFINE PATH** and **QUERY PATH** commands. See Example 9-37.

*Example 9-37 Define path*

```

tsm: SERVER1>def path server1 vtl srct=server destt=library
device=/dev/IBMchanger0
ANR1720I A path from SERVER1 to VTL has been defined.

```

---

8. Define the drives to IBM Tivoli Storage Manager using the **DEFINE DRIVE** and **QUERY DRIVE** commands. See Example 9-38.

*Example 9-38 Define drive*

```

tsm: SERVER1>def drive vtl drive00
ANR8404I Drive DRIVE00 defined in library VTL.

```

```

tsm: SERVER1>def drive vtl drive01
ANR8404I Drive DRIVE01 defined in library VTL.

```

```

tsm: SERVER1>q dr

```

Library Name	Drive Name	Device Type	On-Line
-----	-----	-----	-----
VTL	DRIVE00	LTO	Yes
VTL	DRIVE01	LTO	Yes

```

tsm: SERVER1>q dr vtl DRIVE00 f=d

```

```

Library Name: VTL
Drive Name: DRIVE00

```

```

Device Type: LTO
On-Line: Yes
Read Formats:
ULTRIUM3C,ULTRIUM3,ULTRIUM2C,ULTRIUM2,ULTRIUMC,ULTRIUM
Write Formats:
ULTRIUM3C,ULTRIUM3,ULTRIUM2C,ULTRIUM2
Element: 2
Drive State: EMPTY
Volume Name:
Allocated to:
WWN: 20010000C97C6CA1
Serial Number: 1497515000
Last Update by (administrator): ALEXC
Last Update Date/Time: 08/31/2010 13:42:09
Cleaning Frequency (Gigabytes/ASNEEDED/NONE): NONE

```

```
tsm: SERVER1>q dr vt1 DRIVE01 f=d
```

```

Library Name: VTL
Drive Name: DRIVE01
Device Type: LTO
On-Line: Yes
Read Formats:
ULTRIUM3C,ULTRIUM3,ULTRIUM2C,ULTRIUM2,ULTRIUMC,ULTRIUM
Write Formats:
ULTRIUM3C,ULTRIUM3,ULTRIUM2C,ULTRIUM2
Element: 3
Drive State: UNKNOWN
Volume Name:
Allocated to:
WWN: 20010000C97C62EF
Serial Number: 1497515001
Last Update by (administrator): ALEXC
Last Update Date/Time: 09/24/2010 16:46:19
Cleaning Frequency (Gigabytes/ASNEEDED/NONE): NONE

```

---

9. Define the paths to the IBM Tivoli Storage Manager drives. See Example 9-39.

*Example 9-39 Defining paths*

---

```
tsm: SERVER1>def path server1 drive00 srct=server destt=drive library=vt1
device=/dev/IBMtape0
ANR1720I A path from SERVER1 to VTL DRIVE00 has been defined.
```

```
tsm: SERVER1>def path server1 drive01 srct=server destt=drive library=vt1
device=/dev/IBMtape1
ANR1720I A path from SERVER1 to VTL DRIVE01 has been defined.
```

```
tsm: SERVER1>q path
```

Source Name	Source Type	Destination Name	Destination Type	On-Line
SERVER1	SERVER	VTL	LIBRARY	Yes
SERVER1	SERVER	DRIVE00	DRIVE	Yes
SERVER1	SERVER	DRIVE01	DRIVE	Yes

---

10. Define an IBM Tivoli Storage Manager device class using the commands **DEFINE DEVCLASS** and **QUERY DEVCLASS** as shown in Example 9-40.

*Example 9-40 Defining device class*

```
tsm: SERVER1>def devc ltoclass libr=vtl devtype=lto format=drive mountret=5
ANR2203I Device class LTOCLASS defined.
```

```
tsm: SERVER1>q devclass
```

Device Class Name	Device Access Strategy	Storage Pool Count	Device Type	Format	Est/Max Capacity (MB)	Mount Limit
DISK	Random	1				
LTOCLASS	Sequential	1	LTO	DRIVE		DRIVES

11. Define an IBM Tivoli Storage Manager storage pool that uses the device class defined in previous step using the commands **DEFINE STGPPOOL** and **QUERY STGPPOOL** as shown in Example 9-41.

*Example 9-41 Defining storage pool*

```
tsm: SERVER1>def stg ptpool ltoclass maxscr=100
ANR2200I Storage pool PTPOOL defined (device class LTOCLASS).
```

```
tsm: SERVER1>q stgpool
```

Storage Pool Name	Device Class Name	Estimated Capacity	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
PTPOOL	LTOCLASS	0.0 M	0.0	0.0	90	70	
SPACEMGPOOL	DISK	0.0 M	0.0	0.0	90	70	

12. Label the cartridges in the library using these commands: **LABEL LIBVOLUME** to label the cartridges, and **QUERY LIBVOL** to list the volumes we labeled as shown in Example 9-42.

*Example 9-42 Querying the volumes*

```
tsm: SERVER1>label libv vtl search=yes labels=barcode checkin=scratch
ANS8003I Process number 2 started.
```

```
tsm: SERVER1>q libv
```

Library Name	Volume Name	Status	Owner	Last Use	Home Element	Device Type
VTL	A00000L3	Scratch			1,026	LTO
VTL	A00001L3	Scratch			1,027	LTO
VTL	A00002L3	Scratch			1,028	LTO
VTL	A00003L3	Scratch			1,029	LTO
VTL	A00004L3	Scratch			1,030	LTO
VTL	A00005L3	Scratch			1,031	LTO
VTL	A00006L3	Scratch			1,032	LTO
VTL	A00007L3	Scratch			1,033	LTO
VTL	A00008L3	Scratch			1,034	LTO

The virtual library is now defined to your IBM Tivoli Storage Manager server, and the virtual drives and cartridges are ready for use.

You now must use standard methods to alter your management class copy groups to change the destination value to point to the storage pool created for the virtual library.

**Note:** If you do not label the virtual cartridges before use, when ITSM attempts to write data to them, the process will fail, and ITSM will issue an error message saying that it could not read the internal label of the cartridge. If this error occurs, issue **CHECKOUT LIBVOLUME REMOVE=NO** option to check out and reset the library status of all the cartridges (include the **REMOVE=NO** option so that they do not leave the virtual library) and label them again with the **LABEL LIBVOLUME** command.

If you forget to include the **REMOVE=NO** option in your **CHECKOUT LIBVOLUME** command, the library will place the virtual cartridges in the virtual import/export slots. You can view cartridges stored in these slots on the Import/Export tab of the Library window. Using the menu options accessed by right-clicking the desired cartridge, you can relocate the cartridges back to standard virtual slots. After they are relocated to standard slot locations, use the **LABEL LIBVOLUME** command to label and check them in again. Alternatively, you can label them directly from the Import/Export slots by using the **SEARCH=BULK** option on the **LABEL LIBVOLUME** command.

**Note:** If IBM Tivoli Storage Manager has **SANDISCOVERY ON** when you are using the ProtecTIER Virtual Tape Library, it can cause problems with the tape drive path descriptions of a node if the node goes offline or the path to the port is broken. With this option on in this scenario, you must reconfigure all the tape paths or determine which device belongs to which path by serial number (which will take much longer). With a lot of virtual drives, this could be time consuming.

## Recommended configuration changes

The following IBM Tivoli Storage Manager server and client options should be checked and, if necessary, changed to enable the optimum performance of ProtecTIER:

- ▶ Any IBM Tivoli Storage Manager client compression that will be stored in a ProtecTIER storage pool should have that compression disabled.
- ▶ Ensure that server option **MOVEBATCHSIZE** is set at 1000 (the default value).
- ▶ Ensure that server option **MOVESIZETHRESHOLD** is set at 2048 (the default value).
- ▶ When running IBM Tivoli Storage Manager on a Windows platform, the ITSM driver for tape drives and libraries must be used. The native Windows drivers for the emulated P3000 and DLT7000 drives will not function.
- ▶ Ensure that during defining the libraries **RELABELSCRATCH** is enabled.

Given that ProtecTIER acts as a virtual tape library and a data deduplication device, the advantages associated with disk backup over tape backup apply here too. The following points should also be considered when using ProtecTIER with IBM Tivoli Storage Manager:

- ▶ ITSM disk pools: For some large environments with several IBM Tivoli Storage Manager servers in place, you do not need to assign dedicated ITSM disk storage pools to each server. With ProtecTIER, you can either share a virtual library or you can create virtual libraries for every server.
- ▶ LAN-free backups are easier: Because ProtecTIER is a virtual tape library, it has the advantage of presenting greatly increased tape resources to the backup server. This

capability allows you to perform LAN-free backups to ProtecTIER without much regard for the limitations normally applied to these backups, such as tape drive availability. If you have many LAN-free clients already, then it is possible that your LAN-free backup windows were dictated not entirely by business needs but also by hardware availability. With ProtecTIER and its maximum of 256 virtual tape drives per ProtecTIER node, you can almost completely eliminate any hardware restrictions that you might have faced previously, and schedule your backups when they are required by your business needs.

- **Data streams:** You might be able to reduce your current backup window by taking full advantage of ProtecTIER's throughput performance capabilities. If tape drive availability has been a limiting factor on concurrent backup operations on your ITSM server, you can define a greater number of virtual drives and reschedule backups to run at the same time to maximize the number of parallel tape operations possible on ProtecTIER systems.


**Note:** If you choose to implement this strategy, you might need to increase the value of the **MAXSESSIONS** option on your ITSM server.

- **Reclamation:** You should continue to reclaim virtual storage pools that are resident on ProtecTIER. The thresholds for reclamation might need some adjustment for a period until the system reaches steady state (refer to 3.4.5, "Steady state" on page 29 for an explanation of this term). When this point is reached, the fluctuating size of the virtual cartridges should stabilize and you can make a decision about what the fixed reclamation limit ought to be.
- **Number of cartridges:** This is a decision with several factors to be considered:
  - In ProtecTIER, the capacity of your repository is spread across all your defined virtual cartridges. If you define only a small number of virtual cartridges in ProtecTIER Manager, you might end up with cartridges that hold a large amount of nominal data each. Although this might reduce complexity, it could also affect restore operations in that a cartridge required for a restore might be in use by a backup or housekeeping task. Preemption can resolve this issue, but it might instead be better to define extra cartridges so that your data is spread over more cartridges and drives to make the best use of your virtual tape environment.
  - **Reuse delay period for storage pool cartridges:** When deciding how many virtual cartridges to define, remember to consider using the current storage pool **REUSEDDELAY** value. This is usually equal to the number of days that your ITSM database backups are retained before they expire. The same delay period should apply to your storage pools that store data on ProtecTIER virtual cartridges and you might need to increase the number defined to ensure that you always have scratch cartridges available for backup.
  - **Collocation:** When using a virtual library, you should consider implementing collocation for your primary storage pools. If you begin a restore while another task (for example, a backup or cartridge reclamation) is using the virtual cartridge, you might not be able to access the data on it immediately. Using collocation means that all your data is contained on the same set of virtual cartridges. Because you do not have any of the restrictions of physical cartridges normally associated with this feature (such as media and slot consumption), you can enable the option quite safely.

Consider these points when determining how many virtual cartridges are to be created. Remember that you can always create additional virtual cartridges at any time.

- **Physical tape:** Depending on your data protection requirements, it might still be necessary to copy the deduplicated data to physical tape. This can be achieved by using standard ITSM copy storage pools that have device classes directing data to physical libraries and drives.





## Implementing ITSM deduplication

This chapter provides step by step instructions on how to implement data deduplication on IBM Tivoli Storage Manager (ITSM).

This chapter provides, describes, discusses, or contains the following:

- ▶ Implementing server-side deduplication
- ▶ Implementing client-side deduplication
- ▶ Managing data deduplication
- ▶ Deduplication best practices

## 10.1 Implementing server-side deduplication

An IBM Tivoli Storage Manager V6.x server can deduplicate data from any currently supported IBM Tivoli Storage Manager V5.x client, and data from the IBM Tivoli Storage Manager V6.x client. If possible, you should use the IBM Tivoli Storage Manager V6.2 or later client because it contains changes specifically designed to make the deduplication process more efficient.

Because the IBM Tivoli Storage Manager server performs the deduplication at a storage pool level, it is also possible to deduplicate data already backed up (potentially under older versions of IBM Tivoli Storage Manager). For more information about deduplicating previous stored data refer to 10.3, “Managing data deduplication” on page 265.

1. To implement data deduplication on ITSM, create a FILE device class and a storage pool with the **deduplicate=yes** parameter.

In this example, we also create a non-deduplicated copy storage pool that will be a copy pool for the duplicated data on the primary pool (see Example 10-1).

---

### *Example 10-1 Setting up device classes and storage pools*

---

```
tsm: TSMSRVR>define devc dedup devtype=file mountl=20 maxcap=1G dir=/tsmpool1
ANR2203I Device class DEDUP defined.
```

```
tsm: TSMSRVR>define devc nondedup devtype=file mountl=20 maxcap=1G dir=/tsmpool2
ANR2203I Device class NONDEDUP defined.
```

```
tsm: TSMSRVR>define stgpool deduppool dedup maxscr=200 deduplicate=yes
identifyprocess=0
ANR2200I Storage pool DEDUPPOOL defined (device class DEDUP).
```

```
tsm: TSMSRVR>define stgpool nondeduppool nondedup maxscr=200 deduplicate=no
pooltype=copy
ANR2200I Storage pool NONDEDUPPOOL defined (device class NONDEDUP).
```

---

**Note:** Make sure that the filesystem your device class configuration points to is large enough to store all of your backups (in our example, /tsmpool1 and /tsmpool2).

2. Create a simple backup policy structure that points to the deduplication-enabled storage pool created in Example 10-1 (in our scenario, DEDUPPOOL). See Example 10-2.

---

### *Example 10-2 Creating the backup policy structure*

---

```
tsm: TSMSRVR>def domain dedup descript="Deduplicated Data Domain"
ANR1500I Policy domain DEDUP defined.
```

```
tsm: TSMSRVR>def policyset DEDUP PS_DEDUP
ANR1510I Policy set PS_DEDUP defined in policy domain DEDUP.
```

```
tsm: TSMSRVR>def mgmt DEDUP PS_DEDUP MC_DEDUP
ANR1520I Management class MC_DEDUP defined in policy domain DEDUP, set
PS_DEDUP.
```

```
tsm: TSMSRVR>def copyg DEDUP PS_DEDUP MC_DEDUP dest=deduppool1
ANR1530I Backup copy group STANDARD defined in policy domain DEDUP, set
PS_DEDUP, management class MC_DEDUP.
```



```
tsm: TSM:SRVR>assign defmgmt DEDUP PS_DEDUP MC_DEDUP
ANR1538I Default management class set to MC_DEDUP for policy domain DEDUP, set
PS_DEDUP.
```

```
tsm: TSM:SRVR>validate policyset DEDUP PS_DEDUP
ANR1554W DEFAULT Management class MC_DEDUP in policy set DEDUP PS_DEDUP does
not have an ARCHIVE copygroup: files will not be archived by default if this
set is activated.
ANR1515I Policy set PS_DEDUP validated in domain DEDUP (ready for activation).
```

```
tsm: TSM:SRVR>activate policyset DEDUP PS_DEDUP
ANR1554W DEFAULT Management class MC_DEDUP in policy set DEDUP PS_DEDUP does
not have an ARCHIVE copygroup: files will not be archived by default if this
set is activated.
```

```
Do you wish to proceed? (Yes (Y)/No (N)) y
ANR1554W DEFAULT Management class MC_DEDUP in policy set DEDUP PS_DEDUP does
not have an ARCHIVE copygroup: files will not be archived by default if this
set is activated.
ANR1514I Policy set PS_DEDUP activated in policy domain DEDUP.
```

---

For additional information about policy domain configuration and backup or archive copy group options, refer to the *Tivoli Storage Manager Administrator's Guide* located at:

<http://publib.boulder.ibm.com/infocenter/tsminfo/v6r2/topic/com.ibm.itsm.ic.doc/welcome.html>

3. Create a new ITSM client nodename using the DEDUP domain as shown in Example 10-3.

*Example 10-3 Creating an ITSM nodename*

---

```
tsm: TSM:SRVR>reg node TSM_CLIENT <password> domain=dedup
ANR2060I Node TSM_CLIENT registered in policy domain DEDUP.
ANR2099I Administrative userid TSMNODE defined for OWNER access to node
TSMNODE.
```

---

4. Run an incremental backup of the TSM\_CLIENT machine, copying 101 GB to the DEDUPPOOL storage pool as shown in Example 10-4.

*Example 10-4 Displaying storage pool usage*

---

```
tsm: TSM:SRVR>q stg
```

Storage Pool Name	Device Class Name	Estimated Capacity	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
-----	-----	-----	-----	-----	-----	-----	-----
DEDUPPOOL	DEDUP	133 G	78.0	78.0	90	70	
NONDEDUPPOOL	NONDEDUP	0.0 M	0.0				

---

5. With data in the storage pool, run the **identify duplicates** process to identify duplicated data on the DEDUPPOOL storage pool (see Example 10-5).

*Example 10-5 Running the identify duplicates process*

---

```
tsm: TSM:SRVR>identify duplic deduppool numproc=3
ANR1913I IDENTIFY DUPLICATES successfully started 3 IDENTIFY processes.
```

```
tsm: TSMSRVR>q proc
```

Process Number	Process Description	Status
299	Identify Duplicates	Storage pool: DEDUPPOOL. Volume: /tsmpool/00000318.BFS. State: active. State Date/Time: 08/31/10 15:59:23. Current Physical File(bytes): 986,619,600. Total Files Processed: 134. Total Duplicate Extents Found: 9. Total Duplicate Bytes Found: 1,763,771.
300	Identify Duplicates	Storage pool: DEDUPPOOL. Volume: /tsmpool/00000319.BFS. State: active. State Date/Time: 08/31/10 15:59:23. Current Physical File(bytes): 252,246,757. Total Files Processed: 28. Total Duplicate Extents Found: 4. Total Duplicate Bytes Found: 1,170,719.
301	Identify Duplicates	Storage pool: DEDUPPOOL. Volume: /tsmpool/0000031A.BFS. State: active. State Date/Time: 08/31/10 15:59:23. Current Physical File(bytes): 93,628,221. Total Files Processed: 241. Total Duplicate Extents Found: 125. Total Duplicate Bytes Found: 14,430,666.

After all data in the storage pool is identified, the **identify duplicates** processes will be in an idle state and will remain like that until you cancel the processes or there is new data to be deduplicated.

Even if we run the reclamation process now on the DEDUPPOOL storage pool, the data will not be deleted because it was not copied yet to a copy pool (see Example 10-6). By default, ITSM will only delete deduplicated data during reclamation *after* they have been copied to a copy storage pool. If you want to disable this feature, see 10.3.5, “Disabling the copy storage pool backup requirement” on page 268.

#### Example 10-6 Showing storage pool usage

```
tsm: TSMSRVR>q stg
```

Storage Pool Name	Device Class Name	Estimated Capacity	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
DEDUPPOOL	DEDUP	133 G	78.0	78.0	90	70	
NONDEDUPPOOL	NONDEDUP	0.0 M	0.0				

- Back up the storage pool to a non-deduplicated copy storage pool as shown in Example 10-7.

#### Example 10-7 Backing up the primary storage pool

```
tsm: TSMSRVR>backup stg DEDUPPOOL NONDEDUPPOOL
ANS8003I Process number 302 started.
```

After the **backup stg** process finishes, we can see that all data was copied as shown in Example 10-8 on page 263.

#### Example 10-8 Displaying storage pools usage after backup stg

```
tsm: TSMSRVR>q stg
```

Storage Pool Name	Device Class Name	Estimated Capacity	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Stora- ge Pool
-----	-----	-----	-----	-----	-----	-----	-----
DEDUPPOOL	DEDUP	133 G	78.0	78.0	90	70	
NONDEDUPPOOL	NONDEDUP	133 G	78.0				

As you can see, even after backing up the deduplicated storage pool to a copy storage pool, the deduplicated data is not deleted. This will only happen after you run reclamation on the deduplicated storage pool (see Example 10-9).

#### Example 10-9 Running reclamation

```
tsm: TSMSRVR>reclaim stg DEDUPPOOL thre=1
ANR2110I RECLAIM STGP00L started as process 44.
ANR4930I Reclamation process 44 started for primary storage pool DEDUPPOOL
manually, threshold=1, duration=300.
ANS8003I Process number 44 started.
```

After the **reclaim stg process** finishes, the deduplicated data is deleted and you can see the amount of free space as shown in Example 10-10. Note the **Duplicate Data Not Stored** parameter on the storage pool listing.

#### Example 10-10 Showing storage pool usage after deduplication

```
tsm: TSMSRVR>q stg
```

Storage Pool Name	Device Class Name	Estimated Capacity	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Stora- ge Pool
-----	-----	-----	-----	-----	-----	-----	-----
DEDUPPOOL	DEDUP	133 G	58.5	58.5	90	70	
NONDEDUPPOOL	NONDEDUP	133 G	78.0				

```
tsm: TSMSRVR>q stg deduppool f=d
```

```
Storage Pool Name: DEDUPPOOL
Storage Pool Type: Primary
Device Class Name: DEDUP
Estimated Capacity: 133 G
Space Trigger Util: 98.6
Pct Util: 58.5
Pct Migr: 58.5
Pct Logical: 99.9
High Mig Pct: 90
Low Mig Pct: 70
Migration Delay: 0
Migration Continue: Yes
Migration Processes: 1
Reclamation Processes: 1
Next Storage Pool:
```

```

Reclaim Storage Pool:
Maximum Size Threshold: No Limit
    Access: Read/Write
    Description:
    Overflow Location:
    Cache Migrated Files?:
    Collocate?: Group
    Reclamation Threshold: 100
    Offsite Reclamation Limit:
Maximum Scratch Volumes Allowed: 200
Number of Scratch Volumes Used: 77
    Delay Period for Volume Reuse: 0 Day(s)
    Migration in Progress?: No
    Amount Migrated (MB): 0.00
Elapsed Migration Time (seconds): 0
    Reclamation in Progress?: No
    Last Update by (administrator): ALEXC
    Last Update Date/Time: 08/31/10 17:48:58
    Storage Pool Data Format: Native
    Copy Storage Pool(s):
    Active Data Pool(s):
    Continue Copy on Error?: Yes
    CRC Data: No
    Reclamation Type: Threshold
    Overwrite Data when Deleted:
    Deduplicate Data?: Yes
Processes For Identifying Duplicates: 0
    Duplicate Data Not Stored: 26,152 M (25%)
    Auto-copy Mode: Client
Contains Data Deduplicated by Client?: No

```

---

**Note:** Not all of the duplicated data might expire at once. Several reclamation processing might be needed to expire all of the duplicated extends.

7. Use the **show deduppending <stgpool\_name>** command on the ITSM server to show the amount of data that has been identified as duplicate data, but has not yet been deleted (see Example 10-11).

*Example 10-11 Running show deduppending command*

```

tsm: TMSRVR>show deduppending DEDUPPOOL
ANR1015I Storage pool DEDUPPOOL has 775,340 duplicate bytes pending removal.

tsm: TMSRVR>

```

---

## 10.2 Implementing client-side deduplication

Starting with IBM Tivoli Storage Manager 6.2, client-side deduplication is also available. To implement client-side deduplication, follow the same steps as described in 10.1, “Implementing server-side deduplication” on page 260. The only difference is that when creating the ITSM nodename, specify the option **dedup=clientorserver** to allow the nodename to be client-side deduplicated (see Example 10-12 on page 265).

*Example 10-12 Creating an ITSM nodename for client-side deduplication*

```
tsm: TSMSEVR>reg node TSM_CLIENT <password> domain=dedup dedup=clientorserver
ANR2060I Node TSM_CLIENT registered in policy domain DEDUP.
ANR2099I Administrative userid TSMNODE defined for OWNER access to node
TSMNODE.
```

The next step is to edit the dsm.opt (or dsm.sys if unix) file on the TSM\_CLIENT machine, and include the line

```
deduplication yes
```

You can also enable client-side deduplication from the ITSM Backup-Archive Client GUI, using the following steps:

- ▶ Click **Edit** → **Client Preferences**.
- ▶ Click the **Deduplication** tab.
- ▶ Select the **Enable Deduplication** check box.
- ▶ Click **OK** to save your selections and close the Preferences Editor.

After saving the dsm.opt (or dsm.sys for unix/ aix) file, the machine is ready for client-side deduplication.

Table 10-1 summarizes when client-side-deduplication is enabled, because the setting on the ITSM server ultimately determines whether client-side data deduplication is enabled or not.

*Table 10-1 Data deduplication settings*

Value of the “deduplication” option in the ITSM client configuration file	Value of the nodename’s “deduplication” parameter on the ITSM server	Data deduplication method
Yes	ClientOrServer	client-side
Yes	ServerOnly	server-side (Yes value ignored)
No	ClientOrServer	server-side
No	ServerOnly	server-side

You can also combine both client-side and server-side data deduplication in the same production environment. For example, you can specify certain nodenames for client-side data deduplication and certain nodenames for server-side data deduplication. You can store the data for both sets of client nodes in the same deduplicated storage pool.

## 10.3 Managing data deduplication

Because duplication identification requires extra disk I/O and processor resources, IBM Tivoli Storage Manager lets you control when identification begins, and the number and duration of processes. For more information about data deduplication, refer to the *Tivoli Storage Manager Administrator’s Guide* located at the ITSM 6.2 Information Center:

<http://publib.boulder.ibm.com/infocenter/tsminfo/v6r2/topic/com.ibm.itsm.ic.doc/welcome.html>

### 10.3.1 Starting duplicate-identification processes automatically

If you create a new storage pool for data deduplication, you can specify from 1 to 20 duplicate-identification processes to start automatically. If you do not specify a value, the server starts one process automatically. Example 10-13 shows how you can create a storage pool with three duplicate-identification processes to start automatically.

*Example 10-13 Storage pool creation with identification processes=3*

---

```
tsm: TSMSRVR>define stgpool deduppool dedup maxscr=200 deduplicate=yes
identifyprocess=3
ANR2200I Storage pool DEDUPPOOL defined (device class DEDUP).
```

---

**Note:** The identify duplicates processes will always be active and will always appear in the **q proc** command output. If there's no data to be deduplicated they will continue to run, but will show as in the idle state (see Example 10-16 on page 266).

You can change the number of the duplicate-identification processes at any time, using the **update stgpool** command as shown in Example 10-14.

*Example 10-14 Updating identification processes in a storage pool*

---

```
tsm: TSMSRVR>upd stg DEDUPPOOL identifyprocess=2
ANR2202I Storage pool DEDUPPOOL updated.
```

---

### 10.3.2 Starting duplicate-identification processes manually

The IBM Tivoli Storage Manager server does not start any duplicate-identification process if you specify zero in the **identifyprocess** option when creating a storage pool. Example 10-15 shows how you can create a storage pool when you do not want duplicate-identification processes started automatically.

*Example 10-15 Storage pool creation with identification processes=0*

---

```
tsm: TSMSRVR>define stgpool deduppool dedup maxscr=200 deduplicate=yes
identifyprocess=0
ANR2200I Storage pool DEDUPPOOL defined (device class DEDUP).
```

---

If you are creating a copy storage pool or an active-data pool and you do not specify a value, the server does not start any processes automatically.

After the storage pool has been created, you can increase or decrease the number of duplicate-identification processes as shown in Example 10-14. Remember that any value above zero will cause the duplicate-identification processes to start automatically.

Because the identify duplicates processes will not be running (we specified **identifyprocess=0**), you need to start the processes either manually or using an ITSM administrative schedule, specifying the number of processes and the duration as shown in Example 10-16.

*Example 10-16 Running identify duplicates process*

---

```
tsm: TSMSRVR>identify dup DEDUPPOOL numproc=3 duration=120
ANR1913I IDENTIFY DUPLICATES successfully started 3 IDENTIFY processes.
```

---

```
tsm: TSMSRVR>q proc
```

Process Number	Process Description	Status
58	Identify Duplicates	Storage pool: DEDUPPOOL. Volume: NONE. State: idle. State Date/Time: 09/07/10 16:23:21. Current Physical File(bytes): 0. Total Files Processed: 0. Total Duplicate Extents Found: 0. Total Duplicate Bytes Found: 0.
59	Identify Duplicates	Storage pool: DEDUPPOOL. Volume: NONE. State: idle. State Date/Time: 09/07/10 16:23:21. Current Physical File(bytes): 0. Total Files Processed: 0. Total Duplicate Extents Found: 0. Total Duplicate Bytes Found: 0.
60	Identify Duplicates	Storage pool: DEDUPPOOL. Volume: NONE. State: idle. State Date/Time: 09/07/10 16:23:21. Current Physical File(bytes): 0. Total Files Processed: 0. Total Duplicate Extents Found: 0. Total Duplicate Bytes Found: 0.

When the amount of time that you specify as a duration expires, the number of duplicate-identification processes always reverts to the number of processes specified in the storage pool definition. If that number is zero, all duplicate-identification processes are killed.

### 10.3.3 Enabling data deduplication

Before enabling IBM Tivoli Storage Manager deduplication on a previously existing storage pool, we must understand the consequences. During initial deduplication processing, the IBM Tivoli Storage Manager log file and database will grow. The log grows because the database will be processing lots of transactions, and the database grows to accommodate the table for hashes of objects and hashes for the chunks that enable deduplication to function.

The size of this table is proportional to the number of chunks and objects that IBM Tivoli Storage Manager is processing, so it is important to have enough available database and log space when turning on deduplication in a live system. If you have too many files in previous storages pools and you want them to be deduplicated, you might end up with a ITSM database with twice the size it was without deduplicated files.

To turn on data deduplication for a storage pool, use the **update stgpool** command and specify **deduplicate=yes** as shown in Example 10-17.

*Example 10-17 Turning on data deduplication*

```
tsm: TSMSRVR>upd stg DEDUPPOOL deduplicate=yes identifyproc=3
ANR2202I Storage pool DEDUPPOOL updated.
```

### 10.3.4 Disabling data deduplication

If you turn data deduplication off for a storage pool by updating the storage pool definition, new data that enters the storage pool is not deduplicated.

Deduplicated data that was in the storage pool before you turned off data deduplication is not reassembled. All information about data deduplication for the storage pool is retained,

meaning that your ITSM database will remain with the same size after you disable data deduplication for a storage pool. Deduplicated data continues to be removed due to normal reclamation and deletion.

To turn off data deduplication for a storage pool, use the **update stgpool** command and specify **deduplicate=no** as shown in Example 10-18.

Example 10-18 Turning off data deduplication

```
tsm: TSMSRVR>upd stg DEDUPPOOL deduplicate=no
ANR2202I Storage pool DEDUPPOOL updated.
```

### 10.3.5 Disabling the copy storage pool backup requirement

There is a server option available on ITSM to ensure that only data that is already backed up into a copy storage pool can be deduplicated. This option is **deduprequiresbackup**. The default value is **yes**, which means that data will only be processed for deduplication if it has been safeguarded previously by a storage pool backup. You can change this value by using the **setopt** command as shown in Example 10-19.

Example 10-19 Disabling the deduprequiresbackup option

```
tsm: TSMSRVR>setopt deduprequiresbackup no

Do you wish to proceed? (Yes (Y)/No (N)) y
ANR2119I The DEDUPREQUIRESBACKUP option has been changed in the options file.

tsm: TSMSRVR>
```

**Note:** Setting this option to **no** means that no storage pool backup is required for deduplicated data to be deleted. This is not the default value and is not recommended.

To display all ITSM server configuration options, you can use the **q opt** command on the administrative command line prompt as shown in Example 10-20.

Example 10-20 ITSM server options

```
tsm: TSMSRVR>q opt
```

Server Option	Option Setting	Server Option	Option Setting
-----	-----	-----	-----
CommTimeOut	60	IdleTimeOut	15
MessageFormat	1	Language	AMENG
Alias Halt	HALT	MaxSessions	25
ExpInterval	24	RunStats Interval	10
ActiveLogSize	16,384	DatabaseMemPerce-	AUTO
		nt	
ActiveLogDir	/tsmlog	MirrorLogDir	
ArchFailOverLogD-		ArchiveLogDir	/tsmarchlog
ir			
DbDiagLogSize	1024	CheckTapePos	Yes
ExpQuiet	No	EventServer	Yes
ReportRetrieve	No	DISPLAYLFINFO	No
VolumeHistory	/home/tsm/tsm/volhi-	Devconfig	/home/tsm/tsm/devco-
	st.out		nfig.out



TxnGroupMax	4096	MoveBatchSize	1000
MoveSizeThresh	4096	RestoreInterval	1,440
DisableScheds	No	AuditStorage	Yes
REQSYSauthoutfile	Yes	MsgStackTrace	On
QueryAuth	None	ThroughPutDataThreshhold	0
ThroughPutTimeThreshhold	0	NOPREEMPT	(No)
Resource Timeout	60	TEC UTF8 Events	No
AdminOnClientPort	Yes	NORETRIEVEDATE	No
IMPORTMERGEUsed	Yes	DNSLOOKUP	Yes
NDMPControlPort	10,000	NDMPPortRange	0,0
SearchMPQueue	No	SHREDding	Automatic
SanRefreshTime	0	RetentionExtension	365
DedupRequiresBackup	No	VARYONSERIALLY	No
NumOpenVolsAllowed	20	ClientDedupTxnLimit	300
ServerDedupTxnLimit	300	TcpAdminport	1500
TCPPort	1500	TCPBufsize	32768
TCPWindowSize	64512	CommMethod	TCPIP
TCPNoDelay	Yes	MsgInterval	1
CommMethod	ShMem	FileExit	
ShmPort	1510	FileTextExit	
UserExit		AcsAccessId	
AssistVCCRRecovery	Yes	AcsLockDrive	No
AcsTimeoutX	1	SNMPSubagentPort	1521
AcsQuickInit	Yes	SNMPHeartBeatInt	5
SNMPSubagentHost	127.0.0.1	TECPort	0
TECHost		UNIQUETDPTCEvents	No
UNIQUETECEvents	No	SHAREDLIBIDLE	No
Async I/O	No	CheckTrailerOnFree	On
3494Shared	No	SSLTCPPort	
SANdiscovery	Off	SANDISCOVERYTIMEOUT	15
SSLTCPADMINPort			

### 10.3.6 Restoring or retrieving deduplicated files

When restoring or retrieving files, the client node queries for and displays files as it typically does. If a user selects a file that exists in a deduplicated storage pool, the server reconstructs the file for the restore or retrieve. The process is transparent to the user.

However, restore or retrieve operations from a sequential-access disk storage pool (FILE type device class) that is set up for data deduplication have different performance characteristics than a FILE storage pool that is not set up for data deduplication.

In a FILE storage pool that is not set up for data deduplication, files on a volume that are being restored or retrieved are read sequentially from the volume before the next volume is

mounted. This process ensures optimal I/O performance and eliminates the need to mount a volume multiple times.

In a FILE storage pool that is set up for data deduplication, however, extents that comprise a single file can be distributed across multiple volumes. To restore or retrieve the file, each volume containing a file extent must be mounted. As a result, the I/O is more random, which can lead to slower restore or retrieve times. These results occur more often with small files that are less than 100 KB. In addition, more processor resources are consumed when restoring or retrieving from a deduplicated storage pool. The additional consumption occurs because the data is checked to ensure that it has been reassembled properly.

Although a restore or retrieve operation of deduplicated small files might be relatively slow, these operations are still typically faster than a restore or retrieve operation from tape because of the added tape mount time. If you have data where the restore or retrieve time is critical, you can use a sequential-access disk storage pool that is not set up for data deduplication.

**Note:** To reduce the mounting and unmounting of FILE storage pool volumes, the server allows for multiple volumes to remain mounted until they are no longer needed. The number of volumes that can be mounted at a time is controlled by the **NUMOPENVOLSALLOWED** option on the ITSM server (see Example 10-21 on page 271).

## 10.4 Deduplication best practices

Here are a few best practices you should be aware of when managing a ITSM server with deduplication-enabled storage pools.

### 10.4.1 Server resources

ITSM deduplication processing requires a lot of CPU and disk activity. Use a storage media with fast access characteristics for deduplication, like SSD or SAS/FC, if deduplication performance is an issue. Random read I/O is most of the activity seen during deduplication processing, so allocate your fastest disks for the deduplication-enabled storage pools.

Each identify process can use one entire CPU. You should run no more than “*N*” identify processes for an *N*-Way CPU. For example, if you have an ITSM server with four CPUs, you should run no more than four identify processes. If you need CPU for other processes, run less identify processes.

### 10.4.2 Data safety

There is a server option available to ensure that only data which is already backed up into an IBM Tivoli Storage Manager copy storage pool can be deduplicated. If **deduprequiresbackup** is set to **yes** (which is the default), then data will only be processed for deduplication if it has been safeguarded previously by a storage pool backup. Do not change this default value. For more information, check 10.3.5, “Disabling the copy storage pool backup requirement” on page 268

### 10.4.3 Software version

To get better benefits from deduplication, always install ITSM Client V6.2.1.0 or later on the servers you plan to deduplicate. ITSM Client V6.2 includes client-side deduplication and more efficient deduplication processing.

### 10.4.4 Administrative schedules

When IBM Tivoli Storage Manager deduplication runs, it de-references objects no longer required by IBM Tivoli Storage Manager in a similar way as expiration does, and in common with expired objects on sequential media, we must run reclamation to recover the space on those volumes. As part of an administrative schedule or maintenance plan, run the deduplication processes *after* backup storage pool and *before* reclamation.

Also, note that the following operations take precedence over client-side data deduplication:

- ▶ LAN-free data movement
- ▶ Simultaneous-write operations
- ▶ Data encryption

Therefore, do *not* schedule or enable any of those operations during client-side data deduplication. If any of those operations occur during client-side data deduplication, client-side data deduplication will be turned off, and a message is written to the ITSM client error log.

### 10.4.5 Restore performance

To obtain the different extents that make up a file from a deduplicated storage pool, client restore operations and certain server processes might require opening and closing FILE volumes multiple times. The frequency with which FILE volumes are opened and closed during a session can severely affect performance.

To reduce the number of times a volume is opened and closed, IBM Tivoli Storage Manager allows multiple input FILE volumes in a deduplicated storage pool to remain open at the same time during a session. To specify the number of FILE volumes in deduplicated storage pools that can remain open, use the **NUMOPENVOLSAALLOWED** server option. You can set this option in the server options file or by using the **setopt** command as shown in Example 10-21.

*Example 10-21 Changing the numopenvolsallowed server option*

---

```
tsm: TMSRVR>setopt NumOpenVolsAllowed 20
```

```
Do you wish to proceed? (Yes (Y)/No (N)) y
```

```
ANR2119I The NUMOPENVOLSAALLOWED option has been changed in the options file.
```

---

Each session within a client operation or server process can have as many open FILE volumes as specified by this option. A session is initiated by a client operation or by a server process. Multiple sessions can be started within each.

During a client-restore operation, volumes can remain open for the duration of a client-restore operation and as long as a client session is active. During a no-query restore operation, the volumes remain open until the no-query restore completes. At that time, all volumes are closed and released. However, for a standard restore operation started in interactive mode,

the volumes might remain open at the end of the restore operation. The volumes are closed and released when the next classic restore operation is requested.

For any node backing up or archiving data into a deduplicated storage pool, set the value of the **MAXNUMMP** parameter in the client-node definition to a value at least as high as the **NUMOPENVOLSALLOWED** option. Increase this value if you notice that the node is failing client operations because the **MAXNUMMP** value is being exceeded.



## Suggestions

In this chapter we discuss which deduplication option to choose between N series, IBM Tivoli Storage Manager (ITSM), and ProtecTIER.

The following topics are covered:

- ▶ Data flow considerations
- ▶ Deduplication platform features
- ▶ Deduplication features matrix
- ▶ ProtecTIER deduplication
- ▶ ITSM deduplication
- ▶ IBM N series deduplication

## 11.1 Data flow considerations

It is important to understand the movement of data within the existing infrastructure as this will help guide decisions about where and when to deduplicate. Primary disk storage systems will have higher I/O requirements and fewer duplicates when compared to secondary disk storage systems.

- ▶ What is the source and destination of the data?
- ▶ What infrastructure systems will the data I/O pass through?
- ▶ What replication tools are in place?
- ▶ Are deduplication-capable systems already part of the infrastructure?
- ▶ Is a disk-based backup system already in place?
- ▶ Is the intended duplication processing system capable of providing the overhead required for deduplication?
- ▶ What will be the impact of deduplication to the current I/O?
- ▶ Is the intention to deduplicate data already archived on disk? How will this data be presented to the deduplication processing system?

## 11.2 Deduplication platform features

This section covers deduplication differences for each product (see Table 11-1).

Table 11-1 Deduplication capabilities for IBM products

Deduplication processing	ITSM	ProtecTIER	IBM N-Series
Where?	Server based or client based processing	Storage system based processing	Storage system based processing
When?	Inline (client-side) or postprocess (server-side)	Inline	Postprocess
Which data?	All data (new and existing)	New data	All active data on the storage system
How?	File-level deduplication (incremental backup) and block-level deduplication (hash based using SHA-1)	HyperFactor	Block-level deduplication (hash + byte comparison)
Impact	CPU and disk I/O intensive and requires additional ITSM database space	CPU ( ProtecTIER server) and disk I/O intensive (attached storage)	Up to 6% capacity overhead and CPU impact on storage system

Which solution is best for a given environment is commonly a trade-off between performance and cost. Firms should consider the economic and operational benefits of deduplication in a multi-tier primary storage system versus only using deduplication for backup/archive data resident in secondary storage. However, there is no one silver bullet or one solution that fits. IT managers will have to choose solutions that address their biggest and their largest number of pain points with a clear understanding of the trade-offs that they face with the choice they make. Considerations include:

- ▶ Performance versus cost is not completely mutually exclusive: Solutions available in the market today represent a continuum that strikes a balance between performance and cost (cost of storage and network bandwidth). Decision makers will have to make purchase

decisions after careful analysis of their I/O requirements, data growth rates, network/processing capacity, and composition of data types across storage tiers. Certain suppliers offer tools that allow customers to access the performance impact of deduplication specific to their own environment before deploying in their production environment.

The balance between performance and cost savings is created primarily by using different permutations of performing deduplication such as:

- ▶ Primary storage deduplication
  - In-line deduplication
  - Post process deduplication
- ▶ Secondary storage deduplication
  - Inline deduplication
  - Post process deduplication

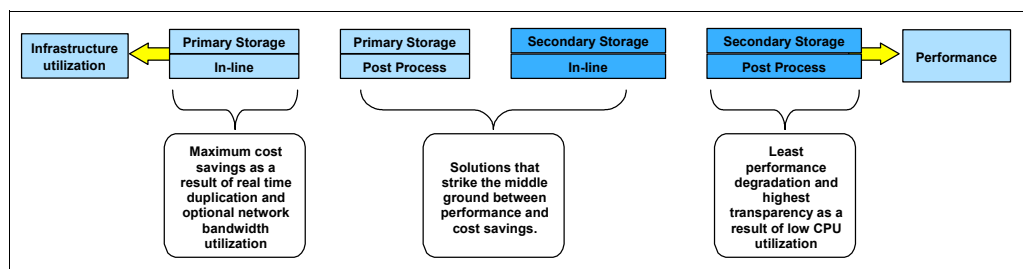


Figure 11-1 Deduplication Approaches Presenting a Continuum of Trade-offs

Other variables that play a role in balancing these trade-offs at different levels include:

- ▶ Optimization technique: Although some products use either deduplication or compression, others use both to provide higher level of optimization. The choice of optimization technique and level of optimization by itself is a balancing act between performance and cost. Solutions that perform intense optimization using both compression and deduplication on primary data tend to place heavy overhead on the computational resources and thus might not be suitable for primary storage IT environments.
- ▶ Speed versus selective approach: In the primary storage deduplication and compression space, some products attempt to address performance issues by simply being fast whereas others meet performance objectives by being prudent in the choice of data they optimize. The choice between the two approaches and the technology used to pursue either of the paths will have an effect on the performance versus cost equation.
- ▶ Features to enable, disable, revert and schedule deduplication: You should have the option to turn off deduplication or revert their data to its original state if the consequences of deduplication turn out to be undesirable. Other configurable items should include the ability to enable deduplication on specific applications or workloads while disabling it for others. Lastly, post-process deduplication should be able to be scheduled during less processing intensive periods in the day.

From the discussion and Figure 11-1, it is evident that solutions that perform in-line deduplication at the source rank high from an efficiency perspective. These solutions are well suited to specific use cases within IT environments. Applications that require random read and write I/O induce substantial latencies in environments that employ inline deduplication. Such latencies cannot be tolerated in majority of primary storage environments.

## 11.3 Deduplication features matrix

ProtecTIER, N-Series, and IBM Tivoli Storage Manager deduplication provide three options for deduplication of data. Which one is best for your environment?

There is no single answer to this question because it depends on the customer's environment and objectives. However, there are a number of decision points and best practices that provide guidance towards the best solution for your circumstances. Here is a summary decision table followed by more detailed discussions of the decision points.

*Table 11-2 Deduplication features*

Deduplication features	ProtecTIER	N-Series	ITSM
For medium to large enterprise environments requiring highest inline deduplication performance (over 1 GBps) and scaling (up to 1 PB storage representing 20+PB of data prior to deduplication)	X		
For global deduplication across multiple backup servers (IBM Tivoli Storage Manager and others)	X		
When a VTL appliance model is desired	X		
For large or small environments requiring deduplication to be completely incorporated within IBM Tivoli Storage Manager without separate hardware or software. Sufficient server resources must be available for required deduplication processing			X
For environments where deduplication across a single IBM Tivoli Storage Manager server is sufficient (for example, small or single server environments)	X		X
Deduplication with no additional licensing costs		X	X
For environments that already have N-Series as part of their infrastructure and want to enable deduplication at no extra cost		X	
Allows deduplication of data on primary storage systems (production servers)		X	
Allows deduplication of backup data on disk storage systems	X	X	X
Allows deduplication of existing data		X	X

## 11.4 ProtecTIER deduplication

Most large enterprise environments demanding high deduplication performance and scaling should choose ProtecTIER's industry leading capabilities. In addition to its unmatched speed and scalability, ProtecTIER's unique technology is extremely efficient in its utilization of memory and I/O, allowing it to sustain performance as the data store grows. This is a major area of concern with most other deduplication technologies.

So, what is a large environment? This also is subjective, but here are guidelines to consider. If you have 10 TB of changed data to backup per day, then you would need to deduplicate at 121 MBps over the full 24 hours to deduplicate all that data.



$$121 * 60 \text{ seconds} * 60 \text{ minutes} * 24 \text{ hours} = \sim 10 \text{ TB}$$

The rate of 121 MBps is moving towards the practical upper limits of throughput of most deduplication technologies in the industry with the exception of ProtecTIER. ProtecTIER, however, can deduplicate data at 900 MBps (1.2 GBps with a 2 node cluster), and can sustain those types of rates over time.

Another consideration is that it is unlikely that any environment will have a full 24 hour window every day to perform deduplication (either inline during backup data ingest, or post processed after data has been stored). Typical deduplication windows will be more like 8 hours or less per day (for example, during, or immediately after, daily backup processing). For a specific environment, you can calculate the required deduplication performance rate by dividing the total average daily amount of changed data to be backed up and deduplicated, by the number of seconds in your available daily deduplication window:

$$\text{Dedup\_Rate} = \text{Amount\_of\_MBs\_Daily\_Backup\_Data} / \text{Number\_Seconds\_in\_Dedup\_Window}$$

For example, we saw 10 TB of data in a full 24 hour window requires a deduplication rate of 121 MBps. Another example is 5 TB of daily data in a 8 hour deduplication window would require a 182 MBps deduplication rate.

$$5,242,880 \text{ MB} / 28,800 \text{ sec} = 182 \text{ MBps.}$$

Although deduplication performance rates can vary widely based on configuration and data, 100 to 200 MBps seems to be about the maximum for most deduplication solutions except ProtecTIER. Any scenario requiring greater than that performance would require ProtecTIER.

ProtecTIER can deduplicate 10 TB of data in 2.8 hours. Under ideal conditions, it would take the fastest of other deduplication solutions 13.9 hours to deduplicate 10 TB of data. Rather than utilize ProtecTIER, you could deploy multiple, distinct deduplication engines to handle a larger daily load like this, but that would restrict the domain across which you deduplicate and minimize your deduplication ratios.

The equation above can assist in determining if you need to go with ProtecTIER for highest performance. The discussion only considers deduplication processing. Note that you also need to consider that IBM Tivoli Storage Manager native deduplication will introduce additional impact to reclamation processing. Also, remember to plan for growth in your average daily backup amount.

Here are considerations for choosing ProtecTIER:

- ▶ ProtecTIER supports data stores up to 1 PB, representing potentially up to 25 PB of primary data. ProtecTIER deduplication ratios for IBM Tivoli Storage Manager data is lower due to IBM Tivoli Storage Manager data reduction efficiencies, but some ProtecTIER customers have seen up to 10:1 or 12:1 deduplication on IBM Tivoli Storage Manager data. Environments needing to deduplicate PBs of represented data should likely choose ProtecTIER.
- ▶ Environments that require global deduplication across the widest domain of data possible should also use ProtecTIER. ProtecTIER deduplicates data across many IBM Tivoli Storage Manager (or other) backup servers and any other tape applications. IBM Tivoli Storage Manager's native deduplication operates only over a single server storage pool. Therefore, if you desire deduplication across a domain of multiple IBM Tivoli Storage Manager (or other backup product) servers, you should employ ProtecTIER.
- ▶ ProtecTIER is the right choice also if a Virtual Tape Library (VTL) appliance model is desired.

Different deduplication technologies use different approaches to guarantee data integrity. However, all technologies have matured their availability characteristics to the point that availability is no longer a salient decision criteria for choosing deduplication solutions.

## 11.5 ITSM deduplication

Another way of deciding between ProtecTIER and IBM Tivoli Storage Manager deduplication is to evaluate IBM Tivoli Storage Manager server system resources (CPU, memory, I/O bandwidth, database backup size, and potential available window for deduplication). If sufficient server resources can be made available for daily deduplication and additional reclamation processing, IBM Tivoli Storage Manager deduplication is a great option.

ITSM deduplication is ideal for smaller environments and for customers who do not want to invest in a separate deduplication appliance. IBM Tivoli Storage Manager can also be used in larger environments if appropriate CPU, memory, and I/O resources are available on the server.

Like most deduplication technologies, ITSM deduplication performance rates vary greatly based on data, system resources applied, and other factors. In our labs, we have measured server deduplication rates of 300 to 400 MBps with large files (greater than 1 MB) on 8 processor AIX systems with 32 GB of memory. On 4 processor systems we've seen rates around 150 to 200 MBps on large files. Rates with mostly small files or running on single processor systems were much lower. If your environment has mostly large files, our benchmarking would suggest that 100 to 200 MBps deduplication rates are possible with IBM Tivoli Storage Manager using 4 or 8 processors.

**Note:** IBM Tivoli Storage Manager native deduplication increases the size of the ITSM database, limiting the number of objects that can be stored in the server.

IBM Tivoli Storage Manager deduplication is also the right choice if you desire the benefits of deduplication completely integrated within IBM Tivoli Storage Manager, without separate hardware or software dependencies or licenses. Native deduplication provides minimized data storage completely incorporated into IBM Tivoli Storage Manager's end to end data life cycle management.

Another important reason for choosing native deduplication is that it comes as part of IBM Tivoli Storage Manager at no additional cost.

## 11.6 IBM N-Series deduplication

IBM N series Storage System offers deduplication as a native part of its Data ONTAP operating system. This is a feature that can be added for no additional cost and leverages several benefits already built into the storage system's licensed internal code. Because it is postprocess deduplication within the storage system, it is possible to realize immediate benefits with existing storage capacity.

The IBM N series deduplication follows a risk-averse postprocess methodology and provides owners the benefits of both source and destination deduplication. For environments that rely on multi-tier storage architecture with N series, the deduplication gains can be achieved not only on primary storage but on backup and archival disk tiers as well.

The deduplication license comes at no cost with the IBM N series Filer and Gateway systems. The N series Gateway systems can be added to existing non-N series disk storage systems to allow these systems to also take advantage of the benefits and features of Data ONTAP, including deduplication.





# Part 4

## Appendixes

These appendixes will provide examples of test environments before and after deduplication of data





# A

## N series use case

This appendix provides statistics about IBM N series Storage System data deduplication ratios. Testing of various data sets has been performed to determine typical space savings in different environments. These results were obtained in two ways:

- ▶ Running deduplication on various production data sets in the lab
- ▶ Using IBM System Storage N series storage systems deployed in the real world running deduplication.

We tested different types of data and obtained different data deduplication ratios depending on the type of data such as regular user files, compressed files, mp3 files, picture files, and so on.

## Lab environment

The IBM N series Storage System used on this user case consisted of:

- ▶ Data ONTAP version 7.3.4
- ▶ 1 Aggregate
- ▶ 7 Disks and 1 Spare
- ▶ Licenses enabled:
  - A-SIS(deduplication)
  - CIFS
  - FCP
  - Flex Clone
  - iSCSI
  - Near Store Option
  - NFS
  - Snap Restore
- ▶ IBM N series System Manager

## Results

The best result achieved in our lab was a 48% space saved using VMware host images.

The best result achieved in our lab was a 92% data deduplication ratio, backing up several versions of the same presentation file type, or backing up several presentation files with minor differences between them.

**Note:** On a customer's primary storage environment, we currently see VMware data deduplication savings at **92%** and mixed application data deduplication savings at **64%**.

## Summary

The data deduplication ratio varies depending on the type of data you are backing up. The more similar blocks a volume contains, the better your deduplication ratio will be. In our tests when we backed up just one version of each file (an initial full backup) we achieved an overall deduplication ratio around 30%. When backing up similar files several times, we achieved a 92% deduplication ratio.

You might achieve even better results in a real production environment because you might have more similar files and more versions of the same file stored in your system storage.


More space savings can occur if you combine with N series. The IBM Real-time Compression Appliances increase the capacity of existing storage infrastructure helping enterprises meet the demands of rapid data growth and also enhance storage performance and utilization. All IBM Real-time Compression Appliances apply IBM patented real-time data compression techniques to primary and existing storage, delivering optimization and savings throughout



the entire storage life-cycle. The result is exceptional cost savings, ROI, and operational and environmental efficiencies.

IBM Real-time Compression Appliance Technology is based on proven Lempel-Ziv (LZ) data compression algorithms. Real-time compression enables the IBM Real-time Compression Appliances to deliver real-time, random access and lossless data compression, maintaining reliable and consistent performance and data integrity. The IBM Real-time Compression Appliance provides real-time data compression of up to 80% in N series environments, which increases storage efficiency with no performance degradation.

This product (see Figure A-1) is transparent and easy to deploy and manage. There is no change to performance and requires no change to applications, networks, or storage. It also preserves high availability.



*Figure A-1 Real Time Compression*





## **ProtecTIER user cases**

This section also contains ProtecTIER deduplication ratios for sample data.

## ProtecTIER deduplication

In our lab environment we used IBM System Storage TS7650G with 1 TB of Repository. We created one TS3500 Virtual Tape Library with 10 LTO3 drives. We used V2.5 of ProtecTIER software. With our tests we confirmed that compressed data is not a good candidate for deduplication. Because ProtecTIER does compression, we recommend you send uncompressed data to ProtecTIER.

Here are examples in backup environments (Figure B-1).




Figure B-1 HyperFactor ratios in backup environments

If you would like to see the same data in terms of space saved (Figure B-2).




Figure B-2 Space saved in backup environments

Any performance data contained in this document was determined in a controlled environment. Therefore, the results obtained in other operating environments might vary significantly. These measurements quoted in this document were made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Users of this document should verify the applicable data for their specific environment.





# ITSM deduplication examples

This appendix provides statistics about IBM Tivoli Storage Manager (ITSM) data deduplication ratios based on sample data.

We performed several backups of different types of data and we obtained different data deduplication ratios. Depending on the type of data that you will back up (regular user file, compressed file, music file, image file, etc.) and the options used (compression enabled or disabled on the ITSM client) you will achieve different data deduplication ratios.

This appendix contains the following topics:

- ▶ Environment
- ▶ Results
- ▶ Summary

## Environment

In our environment we used the following components:

- ▶ IBM Tivoli Storage Manager server and client version 6.2.1.0 running on a AIX 6.1 server with maintenance level 05
- ▶ IBM Tivoli Storage Manager client 6.2.1.0 running on a Windows 2008 server
- ▶ Client-side data deduplication enabled on the ITSM client
- ▶ Back ups running over the LAN (TCP/IP) to a deduplication-enabled storage pool

On the ITSM server we configured a FILE type device class and an ITSM storage pool with deduplication enabled. A standard backup policy retention was created, pointing to the storage pool with deduplication enabled. Finally, we created one ITSM nodename to back up our Windows 2008 server data and configured the nodename for client-side data deduplication. Incremental backups were performed.

## Results

In our tests ITSM client-side data deduplication and server-side data deduplication provided the same results regarding deduplication ratios. The only difference is where the deduplication processing will occur. Because client-side data deduplication deduplicates the files *before* sending them to the ITSM server, the overall backup duration will be longer than using server-side data deduplication. On the other hand, with client-side data deduplication you do not need to run the identify duplicates process on the ITSM server after the backup finishes.

From our tests, we obtained better results using *no* compression on the ITSM client. With compression enabled, we still had data deduplicated, but at a lower rate than without using compression. For example, we achieved a 26% deduplication ratio backing up common user files (including a lot of .mp3 files) on a windows file server with ITSM client compression disabled and 23% deduplication ratio backing up the same data using ITSM client compression.

Unless you have duplicate .mp3 or .zip files, which is rare, you will find that these files do not offer any deduplication savings. You can get better deduplication ratios with flat files rather than with compressed files (.mp3, .zip, and so on).

We achieved a 34% data deduplication ratio when backing up common user files on a windows file server (presentations, documents, executables, pictures, and so on) without ITSM client compression and a few compressed files (.mp3).

The best result achieved in our lab was a 72% data deduplication ratio, backing up all the presentation files on the windows file server (see Figure C-1 on page 293) without ITSM client compression.






Figure C-1 Deduplication results

**Note:** These results are compared to the incremental backup data stored on the ITSM storage pool. It is not compared to the full amount of data stored on the client. Incremental backups already reduce the amount of data stored on ITSM because it copies only files that have changed since the last backup.

## Summary

As you can see, the data deduplication ratio varies depending on the type of data you are backing up. These examples indicate the possible storage savings and differences in deduplication percentages based on file type. The more similar files you have, the better your deduplication ratio will be.

Any performance data contained in this document was determined in a controlled environment. Therefore, the results obtained in other operating environments might vary significantly. These measurements quoted in this document were made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Users of this document should verify the applicable data for their specific environment.



# 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

For information about ordering these publications, see “How to get Redbooks” on page 295. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *TS7680 Deduplication ProtecTIER Gateway for System z*, SG24-7796-00
- ▶ *IBM System Storage N series Software Guide*, SG24-7129-04
- ▶ *Tivoli Storage Manager V6.1 Technical Guide*, SG24-7718-00

## Other publications

These publications are also relevant as further information sources:

- ▶ *IBM System Storage TS7610 ProtecTIER Deduplication Appliance Express - ProtecTIER User's and Maintenance Guide*, GA32-0779-01
- ▶ *IBM System Storage TS7610 ProtecTIER Deduplication Appliance Express Introduction and Planning Guide*, GA32-0776-02
- ▶ *IBM System Storage TS7610 ProtecTIER Deduplication Appliance Express Installation and Setup Guide for VTL Systems*, GA32-0777-01

## Online resources

These Web sites are also relevant as further information sources:

- ▶ IBM ProtecTIER Deduplication Solutions  
<http://www-03.ibm.com/systems/storage/tape/protectier/>
- ▶ IBM System Storage TS7650G ProtecTIER Deduplication Gateway  
<http://www-03.ibm.com/systems/storage/tape/ts7650g/>
- ▶ TS7610 ProtecTIER Deduplication Appliance Express  
<http://www-03.ibm.com/systems/storage/tape/ts7610/>

## How to get Redbooks

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks publications, at this Web site:

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

## Help from IBM

IBM Support and downloads

[ibm.com/support](https://ibm.com/support)

IBM Global Services

[ibm.com/services](https://ibm.com/services)

# Index

## Numerics

3958-AP1 66–67, 69–70  
3958-DD4 66–67, 70–71  
3958-DE2 66–67, 72  
3959-SM1 66–68  
7145-PBR 72

## A

active log 104  
active/active 117  
administrative schedule 266, 271  
Aggregate 52, 60, 136, 284  
AIX xi, 230  
    csum command 33  
algorithm 5, 10, 25, 27, 33, 89–91, 112, 239  
alternate path function 227  
API 28–29, 35, 206, 208  
application data sets 57  
application-transparent 13  
archival disk tiers 9, 278  
archive 35–37, 104–105, 107, 109, 152–153, 160, 261  
archive data 4, 34, 274  
archive log 104  
A-SIS 43, 46, 49, 115–116, 153, 284  
A-SIS license 43, 46  
Assignment window 221  
Asynchronous mode 49  
Asynchronous Volume SnapMirror 47  
ATA 13  
auto mode 15, 167  
available file systems list 203

## B

Backup Exec 58  
backup policies 200  
backup set 240, 243  
backup stream 242  
bandwidth 3, 21–22, 46, 58, 61, 78, 93, 110, 112, 226–227, 237–238, 274, 278  
Barcode seed field 223  
baseline 6, 50, 236  
baseline chunk 6  
best practice 15, 166, 259  
binary comparison 6  
binary differential 23, 27  
block  
    pointer 14  
    reference 14  
block referencing 9, 14  
block size 239, 241  
block-level 4–5, 12, 47, 56–57, 274  
block-level database backups 238  
block-level deduplication 5

buffer capacity 8  
buffer size 242  
BUFFERCOUNT 241  
byte comparison 6, 8, 13–14, 274  
byte level comparison 9  
byte-for-byte comparison 14

## C

cartridge windows 222  
catalog deduplication 11  
changelog 43  
checkpoint 43  
chunk 5–6, 33–34  
chunk data object 5  
chunking 5  
chunks 6, 33, 267  
CIFS 51, 116–117, 143, 284  
CLI 13, 51  
client compression 242, 256  
client node queries 269  
client-based deduplication processing 7  
client-restore operation 271  
client-side data deduplication 35–36, 106, 108–109, 259, 292  
collision 6, 34  
collocation 110, 257  
command queuing 93  
command-line interface (CLI) 13  
compaction 241  
compressed data 110  
compression 33, 110  
CompressionRate 87  
Configuration Wizard 185–186  
copy pool 260, 262  
country code 182  
CPU 4, 6–7, 15, 44, 47, 49, 57, 61–62, 104, 106–108, 110, 112, 167, 270, 274, 278  
Create repository Report window 205  
Create repository window 206  
Create repository wizard 201, 206  
csum 33  
csum command 33  
customer network 184

## D

D2D2T 77  
DAOS 241  
data xi, 3–11, 13–17, 19–37, 42–45, 47, 50–51, 55–63, 69–70, 73–85, 87–97, 99, 104–112, 116–117, 152–153, 165–167, 259–272, 274–278, 281, 283–284, 287–289, 291–293, 306  
data access pattern 15  
data change rates 26, 79, 83, 85, 91  
data chunking 5

- data deduplication processing 6
- data integrity 285
- data object 5
- Data ONTAP 9, 12–14, 16, 42–46, 50–51, 55–58, 60–63, 115–117, 152–153, 165, 278–279, 284
- data pointers 10–11, 16
- data protection 3–4, 22, 68, 70, 82–83, 85, 94, 97, 112, 257
- data reduction 3, 32, 68
- data retention period 21, 26
- data segment 28
- data-matching 238
- DB2 81, 238, 242
- dedup=clientorserver 264–265
- deduplicate data 15, 35, 37, 110, 260, 277
- deduplicated blocks 42, 44, 46, 50, 56
- deduplicated chunks 111
- deduplicated component analysis 7
- deduplicated small files 270
- deduplicated volume 14
- deduplicating primary data 8
- deduplication xi, 1, 3–15, 17, 19, 22–24, 26–28, 31–37, 41–47, 49–63, 65, 67–69, 71, 73–75, 77–78, 85, 92, 103–112, 115–116, 131, 152–157, 163–167, 259–260, 263–271, 273–278, 281, 283–284, 287–288, 291–293, 306
  - hash functions 33
  - HyperFactor 34
  - Single Instance Store 32
  - workloads 32
- deduplication appliance 4, 278
- deduplication architecture 6–7
- deduplication efficiency 7, 57
- deduplication license 9, 279
- deduplication processing 270
- deduplication schedule 15, 50, 55, 164, 167
- deduplication solution 77
- deduplication technology 4, 8, 14, 22, 24, 34
- deduprequiresbackup 268
- dedupverificationlevel 109
- DEFINE LIBRARY 252
- DEFINE STGPPOOL 255
- delta comparison 87
- Delta Compression 27, 87
- delta difference 27
- delta differencing 6
- device driver ftp site 230
- device driver installation 230
- Device IDs 247–248, 250
- DHF 206
- disable compaction 241
- disable compression 236
- disable multiplexing 241
- Disaster Recovery 111
- disk array 20–21, 24, 71, 80–81, 196, 198–199
- disk I/O 7, 93, 104, 106–107, 110, 112, 226, 265, 274
- Disk Storage Pools 239–240
- disk-based backups 13, 21, 26, 83, 85
- disk-to-disk-to-tape 77
- distributed heterogeneous environments 242

- domain 260–261, 265
- Domino 58, 241
- Domino to 8.5 241
- DR replication 8
- drive serial number 246
- Drives fields 222
- DRM 111
- DS4000 66, 80, 94
- DS5020 66
- dsm.opt 265
- duplicate blocks 9–11, 14, 51, 58, 154, 238–239
- duplicate chunks 6, 8
- duplicate data 13
- duplicate extents 36
- duplicate-identification processes 266–267
- duplicates 4–5, 11, 14, 32–33, 36, 58, 261–262, 266, 274, 292
- dynamic utilization 8

## E

- element information 250
- Element Inventory 248, 251
- element number 228, 246, 252
- email alerts 186
- email notifications 191
- email recipients 192
- email server IP 192
- Enable compression 236
- encrypted data 34, 104, 109, 239
- encryption xi, 5, 58–59, 81, 85, 109, 233, 239–240, 247–248, 271
- Enterprise Linux License Agreement window 173
- EW 88–90
- examples
  - configuration 115
- extraneous data chunks 6

## F

- factoring ratio 20, 25–26, 29, 74, 81, 83–92, 98, 200–201, 203, 237, 241–242
- failback process 21
- failover 21, 68, 71
- FC-15K 200, 203
- FCP 16, 51, 144, 167, 284
- Fibre Channel 20, 71, 77, 93, 98, 200, 203
- Fibre Channel disks 93
- file backups 238
- FILE device class 35, 260
- file extent 270
- file level visibility 7
- file system 12
- file system arrangement options 200
- FILE volumes 271
- File-level deduplication 5
- FILESERSET 240–241
- fingerprint 6, 10–11, 13–15, 45, 47, 55, 60
- fingerprint catalog 11
- fingerprint database 14, 60
- fingerprint database file 15

- fingerprint record 15
- fingerprints 6, 9, 11, 13–16, 62, 153–154, 165
- Fixlist 230
- FlexClone 55
- flexible volume 12–16, 46, 59, 61–62, 152–156, 159–160, 165, 167
  - examples 115
- FlexVol 14, 44, 117
- format agnostic 5
- fractional overwrite reserve 53
- Fractional Reserve 51–55
- free volume pool 53
- front end ports 230
- fsCreate 198–199
- full backup 284
- FullCapacity 86–87
- FullCapacityVersions 86–87
- FullFrequency 86–87
- FullPhysicalCapacity 87
- full-plus-incremental 242
- FullRetention 86

## G

- Gateway versions 116
- get\_versions command 194
- granularity 12
- Grid Manager 21

## H

- HA 43, 117, 227
- hash 6, 8–11, 23, 33–34, 274
- hash collisions 6, 8
- hash functions 33
- hash method 6
- hash values 6, 11, 33–34
- hash-based 6
- Hashing 6, 8–9, 11
- HBA ports 238, 240
- HBAs 71, 80, 226–227, 229, 240
- Hierarchical Storage Management 110
- HP-UX xi, 101, 230
- HSM 110
- HW setup 181
- HyperFactor 19, 22–27, 34, 73, 92, 201, 210–211, 213–214, 223, 235–236, 238–239, 274
- HyperFactoring 81

## I

- I/O performance 15, 59, 61, 166, 270
- IBM device driver 230–231
- IBM Linux Tape 230
- IBM Performance Planner tool 98–99
- IBM ProtecTIER 226
- IBM Real-Time Compression Appliance Technology 285
- IBM Real-Time Compression Appliances 284–285
- IBM Storage System N series 283–284
- IBM System Services Representative 203
- IBM System Storage N series 13

- storage systems 283
- technology 13
- IBM System Storage ProtecTIER Appliance Edition 69
- IBM System Storage ProtecTIER Enterprise Edition 68, 70, 72
- IBM System Storage ProtecTIER Entry Edition 68
- IBM System Storage TS7610 ProtecTIER Deduplication Appliance Express 24, 66, 68
- IBM System Storage TS7650 240
- IBM System Storage TS7650 ProtecTIER Deduplication Appliance 66, 69–70, 72
- IBM System Storage TS7680 ProtecTIER Deduplication Gateway 66–67, 72
- IBM Tivoli Storage Manager xi, 237, 242–244, 246, 252–256, 259–260, 264–265, 267, 270–271
- IBM Tivoli Storage Manager DB2 database 242
- IBM Tivoli Storage Manager V5.4 242
- IBM TS3500 219, 222
- identical hash 6
- identification performance 6
- IDENTIFY DUPLICATE 104
- IDENTIFY process 104
- identify processes 270
- identifyprocess option 266
- idtd 247
- import/export slots 224, 256
- inactive client files 243
- IncrementalCapacity 86
- IncrementalCapacityVersions 86–87
- IncrementalChangeRate 87
- IncrementalFrequency 86
- IncrementalPhysicalCapacity 87
- IncrementalRetention 86
- index defragmentation timing 241
- indirect blocks 14
- information infrastructure 3
- initial barcode seed 223
- initiator port 228
- in-line 6, 24, 68, 275
- inline deduplication 8, 112, 275–276
- inline processing 6
- inode pointers 9, 11
- interactive mode 271
- IOPS 30, 93, 226
- IP address 170–171, 177–179, 183, 244
- iSCSI 16, 51, 77, 116–117, 145, 167, 284
- ISL 226
- ISV 78, 231, 234
- itdt 227–228, 232, 246–247, 251
- ITSM 241–244, 246, 253, 256–257, 259–262, 264–266, 268, 270–271
- ITSM client 261, 265, 271
- ITSM database 105–106, 243, 257, 267–268, 278
- ITSM database sizing 105
- ITSM disk pools 256
- ITSM disk storage pools 256
- ITSM Planning overview 104
- ITSM server 243, 257, 264–265, 268, 270

## L

- Lab-based Services 203
- LABEL LIBVOLUME 255–256
- LAN-free 78, 110, 230, 256, 271
- LAN-free backup 77, 257
- LAN-free clients 257
- library emulation 245
- lin\_tape 230–234, 247
- lin\_tape device driver 231
- lin\_tape driver 227, 230
- Linux xii, 171, 173–174, 229–231, 244, 246
- Linux Device Driver 231
- Linux lin\_tape 231
- Linux page 230
- load balancing 220
- lossless data compression 285
- Lotus Domino 58
- LSUs 29, 206–216, 236
- LTO3 225, 245, 288
- LUN masking 217, 222
- LUN Space Reservation 52, 54–55
- LUNs 16, 29, 45, 51–54, 56–58, 81, 97–98, 116–117, 129, 133, 167, 201, 217, 222, 227, 230, 232–233, 245–246
  - configuration examples summary 52
- LZH compression 27

## M

- maintenance level 238
- management class 37, 109, 256, 260–261
- manual deduplication 15, 167
- matching chunks 6
- max. cartridge growth check box 223
- maximum sharing 14
- MAXNUMMP 272
- MAXOPENFILES 240–241
- MD5 33–34
- media server 207, 226
- medium changer 231
- Memory Resident Index 23–24, 27
- metadata 6, 11–12, 14–16, 20, 24, 27–30, 42, 45–47, 55, 60, 62, 74–75, 80–81, 83, 93, 97–99, 167, 200–201, 203–204
- metadata LUNs 98
- metadata planner tool 98
- metadata RAID 203
- Metadata resources advanced dialog 204
- metafile 42
- MetroClusters 44
- Microsoft Exchange 58, 242
- Microsoft Windows Page 230
- MOVEBATCHSIZE 256
- MOVESIZETHRESHOLD 256
- Multiple adjust function 215
- multiple duplicate blocks 9
- multiple references 13
- multiplexing feature 239
- multiplexing features 239
- MultiStore 51

- multi-tier storage architecture 9, 278

## N

- N5200 284
- NAS systems 239–240
- Native Windows drivers 243
- navigation pane 179, 235
- NBU 29, 219
- NDMP 238–240
- NearStore 16, 43, 46, 49, 115–116
- NearStore license 43, 46, 49
- NetBackup 28–29, 50, 100, 206–207, 219
- NetBackup policies 29
- network replication 4
- NFS 51, 58, 116–117, 129, 284
- NLO 241
- Nodes window 179, 201
- NominalCapacity 86
- non-identical data 6
- non-unique chunks 6
- no-query restore 271
- NSF 241
- NTP server 239
- Number of slots field 224
- NUMOPENVOLSALLOWED 270–272

## O

- one-node cluster 201, 217
- OpenStorage 28, 65, 100, 206
- Oracle 238, 240, 242
- OST 28–29, 65, 100, 196–197, 200, 206–208, 226, 234, 236
- OST storage appliance 207
- OST-enabled media server 28, 100

## P

- P3000 67–68, 243, 256
- PAM 44
- parallel backup 241
- PARALLELISM 241
- parallelism 242
- parity calculation 6
- performance 42, 44, 47, 56–62
  - degradation 285
  - I/O 61
  - write 61
- persistent binding 227
- PhysicalCapacity 87
- plug-in 206
- policy domain 260–261
- postprocess deduplication 8
- Pre-Installation Summary window 175–176
- primary storage pool 35, 37, 107, 111, 262–263
- principality 21
- problem diagnosis activities 239
- program group 175
- progressive backup 242
- progressive incremental 3, 32



- ProtecTIER xi, 112, 169–171, 175–177, 179–181, 184, 193, 195–198, 201, 203–204, 206, 208, 216–217, 222, 225–231, 234–246, 253, 256–257, 306
- ProtecTIER compression mode 236
- ProtecTIER Configuration Menu 181, 184–185
- ProtecTIER Deduplication Gateway 24, 66–67, 70, 72
- ProtecTIER factored data 20
- ProtecTIER family 177
- ProtecTIER infrastructure servers 7
- ProtecTIER Manager 20–21, 66, 75, 101–102, 169–171, 175–180, 185, 207, 229, 244, 257
- ProtecTIER node 177, 257
- ProtecTIER Replication Manager 21, 196, 198
- ProtecTIER Repository 20–22, 74, 80, 83–85, 91
- ProtecTIER server configuration 170
- ProtecTIER server nodes 177
- ProtecTIER service nodes 171
- ProtecTIER storage server emulation 206
- ProtecTIER VT name field 217
- ProtecTIER VTL 28, 82, 236, 239
- Protection Manager 50, 57
- ptconfig 181, 195–198, 206

## Q

- q opt command 268
- Qlogic 80, 229
- QSM 45, 47–49
- Qtree SnapMirror 45, 47–49, 51, 56
- Qtree SnapMirror replication 48
- qtrees 47, 62, 134
- QUERY DEVCLASS 255
- QUERY LIBVOL 255
- QUERY PATH 253
- Quick Launch Bar 175
- Quorum 98
- quota management 117

## R

- RAID 20, 29, 75, 80–81, 93–99, 200–201, 203
- RAID group 201
- RAM 171
- random access 285
- RAS 68, 183, 185, 194, 196, 198
- rasMenu 194
- REALLOC 56
- Real-Time Compression 17, 238, 285
- Real-Time Compression Appliances 17, 284–285
- real-time data compression 284–285
- reclaim stg process 263
- reclamation 257, 262–264, 268, 271
- recovery points 4
- Red Hat 68, 72, 101, 171, 173–174, 230, 244
- Redbooks Web site 295
  - Contact us xiii
- redundant copies 4
- reference count 14, 28–30
- reference pointers 10
- Registration window 187
- rehydration processes 238

- RELABELSCRATCH 253, 256
- Release Note 171, 193
- repository 170, 199–206
- repository planning 199
- repository planning wizard 203
- Repository size field 200, 203
- Repository size window 203
- requirements 115
- Resources menu 203
- restore operations 257, 271
- retrieve operation 270
- risk-averse postprocess methodology 9, 278
- RMAN 82, 240–241
- Robot check box 222
- ROI 285
- root volume 42
- round-robin 62
- RPM 15, 93, 200, 203
- RSH 146

## S

- SAN xii, 71, 79–80, 117, 170, 226, 228, 238
- SAN HBA 226
- SAS 203, 270
- SATA 71, 93, 97–98, 203
- schedules
  - deduplication 164
  - parameters 165
  - viewing 164
- secondary comparison 6
- seismic data 240
- sequential-access disk 34
- server option 256, 268, 270–271
- server-based deduplication processing 7
- server-side data deduplication 34, 106
- server-side deduplication 259
- Services window 225
- setopt command 268, 271
- SHA-1 6, 33–34, 274
- show deduppending command 264
- Show version information 194
- simultaneous-write operations 271
- Single Instance Store 32
- sis config 153, 164–165
- sis start -s 45, 47, 50, 55, 153
- sizing and space 15, 166
- Slots window 223
- small files 110
- SnapLock 50
- SnapMirror 42–43, 45–51, 56, 61
- SnapMirror licensing 49
- SnapRestore 45
- Snapshot reserve 16, 167
- Snapshots 12, 14, 16, 44–45, 47, 50–53, 56–57, 133, 167
- SnapVault 49–50, 56–57, 61
- SNMP 69, 150, 170, 186, 188–190
- SNMP traps 186, 189, 191
- Software License Agreement 172–173
- Solaris Page 230

- space and sizing 15, 166
- space reserved 53–54
- space savings 283–284
- space\_optimized option 56
- split volume 56
- SQL 81, 241–242
- SSH keys 147
- stale fingerprints 15
- standard compression 238
- steady state 29–30, 91, 257
- storage
  - consumption 13
  - storage efficiency 285
  - storage infrastructure xi, 284, 306
  - storage life-cycle 285
  - storage performance 284
  - storage pool 35–37, 78, 104–105, 107, 109–112, 239, 243, 255–257, 260–263, 265–272, 277, 292–293
  - storage pool backup 268, 270
  - storage server emulation 206
  - storage system based deduplication processing 7
  - STS 206–208
  - STS credentials 207
  - sub-file 9, 33
  - supported
    - models 115
  - Sync SnapMirror 47
  - System Manager 115–119, 121–125, 128–130, 152, 155–156, 164, 166, 284
  - System name field 203
  - System peak throughput field 200, 203
  - System Storage tape drive 231

## T

- Tape drives window 220
- tape management systems 231
- Tape Model window 219
- tar package 193
- TCP/IP 206
- thin-provisioning 16, 167
- timeserver 181, 183, 196–197
- timeserver IP 170
- timezone 181–182, 196–197
- timezone setup 181
- traditional volumes 62
- transient and temporary data 57
- TS1120 Tape Controllers 72
- TS3000 68–72
- TS3000 System Console 70, 72
- TS3500 Virtual Library 225
- TS7610 24, 66, 68, 76, 170, 181, 193, 198–199, 217, 225
- TS7650 170, 193, 197–199, 230, 240
- TS7650 Appliance 197
- TS7650G 24, 66–67, 70, 72, 75–76, 78–80, 82, 88–91, 93, 170, 193–194, 196–199, 230, 240, 288
- TS7680 24, 66–67, 72, 76, 193, 198–199
- TSCC 70
- two node clusters 20
- two phase comparison 9
- two-node cluster 217, 220

- two-node clusters 171

## U

- Ultrium Linear Tape Open 78
- untar 193
- update stgpool command 266–268
- user accounts 175, 180
- User Data LUNs 98
- user data LUNs 98
- user data segment 28

## V

- vFiler 51
- VI 57
- view deduplication schedules 164
- Virtual cartridge size field 223
- virtual devices 217, 222
- virtual drives 225–226, 230, 237, 256–257
- virtual tape drive definitions 227
- virtual tape environment 257
- virtual tape interface 243
- Virtual Tape Library 4, 65, 200, 226, 256, 277
- virtual tape resource 227
- VMDK 57
- VMDKs 57
- VMware 57, 284
- volume
  - copy 42, 45, 51, 53, 55
  - metadata 14
  - splitting 56
- Volume SnapMirror 45–47, 51
- Volume Space Guarantee 51–55
- VSM 45–47
- VTL 170, 196, 200, 228, 234–236, 239, 243, 252–255

## W

- WAFL 16
- WAN 22
- Windows 32 171
- WORM 50–51, 253
- Write Anywhere File Layout (WAFL) 13
- write performance 61
- WWN 226–228, 252–254
- WWPN 228–229

## X

- x3850 20, 70, 72
- XIV 71

## Z

- zoning 227–229, 238



## Implementing IBM Storage Data Deduplication Solutions

(0.5" spine)  
0.475" <-> 0.873"  
250 <-> 459 pages







# Implementing IBM Storage Data Deduplication Solutions

**Fitting deduplication  
into your enterprise  
environment**

**Adding deduplication  
to your storage  
hierarchy**

**Gaining the space  
benefits of IBM  
deduplication**

Until now, the only way to capture, store, and effectively retain constantly growing amounts of enterprise data was to add more disk space to the storage infrastructure, an approach that can quickly become cost-prohibitive as information volumes continue to grow and capital budgets for infrastructure do not.

In this IBM Redbooks publication, we introduce data deduplication, which has emerged as a key technology in dramatically reducing the amount of, and therefore the cost associated with storing, large amounts of data. Deduplication is the art of intelligently reducing storage needs through the elimination of redundant data so that only one instance of a data set is actually stored. Deduplication reduces data an order of magnitude better than common data compression techniques. IBM has the broadest portfolio of deduplication solutions in the industry, giving us the freedom to solve customer issues with the most effective technology. Whether it is source or target, inline or post, hardware or software, disk or tape, IBM has a solution with the technology that best solves the problem. This IBM Redbooks publication covers the current deduplication solutions that IBM has to offer:

- ▶ IBM ProtecTIER Gateway and Appliance
- ▶ IBM Tivoli Storage Manager
- ▶ IBM System Storage N series Deduplication

## **INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION**

### **BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE**

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:**  
**[ibm.com/redbooks](http://ibm.com/redbooks)**