


TS7680 Deduplication ProtectTIER Gateway for System z

Understand the concepts of data deduplication

Realize bandwidth savings with the TS7680

Reduce your storage hardware requirements



Alex Osuna
Luciano Cecchetti
Erik Franz
Roberto Mencarelli

Redbooks



International Technical Support Organization

TS7680 Deduplication ProtecTIER Gateway for System z

August 2010

Archived

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

Archived

First Edition (August 2010)

This edition applies to Version 2.4 of ProtecTIER.

© Copyright International Business Machines Corporation 2010. 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

Notices	ix
Trademarks	x
Preface	xi
The team who wrote this book	xii
Now you can become a published author, too!	xiii
Comments welcome	xiii
Stay connected to IBM Redbooks	xiv
Part 1. Introduction to the TS7680.	1
Chapter 1. Concepts of data deduplication	3
1.1 What data deduplication is	4
1.2 Type of data deduplication	5
1.2.1 Hash-based	5
1.2.2 Content aware	6
1.2.3 HyperFactor and deduplication	7
1.3 Data deduplication processing	8
1.3.1 In-line method	8
1.3.2 Post-processing method	9
1.4 Components of a data deduplication system	9
1.4.1 Server	10
1.4.2 Data deduplication software	10
1.4.3 Disk array	10
1.5 Benefits of data deduplication	11
1.5.1 Reducing storage requirements	11
1.5.2 Reducing of environmental costs	11
Chapter 2. The TS7680 ProtecTIER Deduplication Gateway architecture	13
2.1 TS7650/TS7680 Differences	14
2.1.1 Shelf / visibility	14
2.1.2 Replication policies by VOLSER (barcode)	14
2.2 IBM System Storage TS7680 ProtecTIER Deduplication Gateway for System z technology overview and concepts	14
2.2.1 Enterprise controllers	17
2.2.2 ProtecTIER Virtual Tape (VTL) service	17
2.2.3 ProtecTIER servers	17
2.2.4 TS3000 System Console	18
2.3 ProtecTIER virtual tape (VTL) service and concepts	18
2.3.1 Virtual tape device	18
2.3.2 Virtual tape volume	18
2.3.3 Cache management	19
2.3.4 HyperFactor	19
2.3.5 Cartridge eject	20
2.3.6 Scratch volume processing	20
2.3.7 Grace period for scratched volumes	21
2.4 ProtecTIER virtualization and deduplication	21
2.4.1 HyperFactor	23
2.4.2 ProtecTIER Virtual Tape Library (VTL) concepts	26

2.5 Library Manager Integration	28
2.6 High Availability	29
Chapter 3. ProtecTIER native replication overview and operation.	35
3.1 Definition of terms	37
3.2 How replication works	37
3.2.1 Extended Volume Attributes	38
3.2.2 Replication features	39
3.2.3 Typical deployment	41
3.2.4 ProtecTIER's native replication Management Interface	42
3.3 Normal operation concepts	43
3.3.1 Replication	43
3.3.2 Replication data transfer	43
3.3.3 Visibility switch control	44
Chapter 4. Hardware planning for the TS7680	47
4.1 General overview of the TS7680	48
4.2 Hardware and software components for the TS7680	49
4.2.1 TS7680 ProtecTIER server characteristics	51
4.2.2 TS7680 Enterprise Tape Controller characteristics	52
4.2.3 TS7680 feature codes	53
4.2.4 TS7680 software	54
4.2.5 TS7680 rack setup	56
4.3 Installation planning	59
4.3.1 Installation worksheets	59
4.3.2 Prerequisites	60
4.3.3 Installation tasks	60
Part 2. Planning and setup.	63
Chapter 5. Planning for deduplication	65
5.1 Capacity planning for the TS7680 - key concepts	65
5.2 HyperFactor overview	66
5.2.1 Sizing inputs	68
5.2.2 Local repository sizing	75
5.2.3 Factoring ratio considerations	77
5.2.4 Storage sizing	80
Chapter 6. TS7680 setup	91
6.1 PowerUP	92
6.1.1 TS3000 System Console, keyboard, video, mouse; and keyboard, video, and mouse switch	93
6.1.2 ProtecTIER servers	93
6.1.3 Enterprise controllers	94
6.2 Setting up the TSSC for use with the 3958-DE2	95
6.2.1 Setting Console Date, Time and Network Time Protocol (NTP) server	98
6.3 Configuring the RAS package	99
6.3.1 ProtecTIER RAS package configuration	100
6.3.2 Verifying the cluster's Ethernet connections	101
6.4 Verify or create a complex and testing Call Home	103
6.4.1 Testing Call Home	105
6.5 Installing the ProtecTIER Manager GUI on external TSSC	107
6.6 Configuring ProtecTIER using ptconfig	109
6.6.1 Logging in to the server	110

6.6.2	Configuring Lower ProtecTIER Server	110
6.7	Adding a node and creating a repository	111
6.7.1	Planning the repository	114
6.7.2	Creating the repository	117
6.8	Configuring Upper PT Server, creating a library and port configuration	123
6.8.1	Add Upper Server to the cluster (using the ProtecTIER Manager)	124
6.8.2	Create a library using ProtecTIER Manager	126
6.8.3	Port configuration	126
6.9	Testing a clustered system configuration	130
6.10	ProtecTIER Library attachment to the Enterprise controllers	130
Chapter 7.	Host attachment	137
7.1	TS7680 is fully integrated with z/OS software support	138
7.2	Planning for software implementation	139
7.2.1	Software requirements	139
7.2.2	z/OS software environments	139
7.2.3	Host configuration definition	140
7.2.4	Sharing a TS7680 Virtualization Engine	140
7.3	Host implementation considerations	141
7.3.1	Channels, adapters, and protocols	142
7.3.2	System z tape controllers	142
7.4	Remote installations and switch support	142
7.4.1	Factors that affect performance at a distance	142
7.4.2	FICON Director support	143
7.4.3	FICON channel extenders	144
7.4.4	Supported distances	144
7.4.5	Implementing cascaded switches	145
7.5	Hardware I/O configuration definition	146
7.5.1	Defining a TS7680	147
7.5.2	Control unit definition	148
7.5.3	Device definition	152
7.5.4	HCD support for LIBRARY-IDs and LIBPORT-IDs	155
7.6	Cache management	156
7.7	Cartridge entry and eject	157
7.8	Tape initialization	157
7.8.1	DFSMSrmm for initialization	158
7.8.2	Scratch pooling	159
7.8.3	Defining the library through ISMF	159
7.8.4	Defining SMS constructs through ISMF	160
7.8.5	Defining Data Classes	161
Part 3.	Managing the TS7680	167
Chapter 8.	Managing and administering TS7680	169
8.1	General management	169
8.1.1	About the TS3000 System Console and its related components	169
8.1.2	User login on IBM TS3000 System Console	170
8.1.3	Logging into and out of the ProtecTIER Manager	172
8.1.4	Changing user account passwords in the ProtecTIER Manager system	173
8.1.5	Adding a user to the ProtecTIER Manager system	174
8.1.6	Logging in to the ProtecTIER server	174
8.1.7	Host reporting	176
8.2	Power management	177
8.2.1	Shutting down the 3958-DE2 (TS7680)	177

8.2.2	Shutting down the Enterprise controller	179
8.2.3	Shutting down the ProtecTIER server	181
8.2.4	Powering on the 3958-DE2 (TS7680)	182
8.2.5	Powering on the ProtecTIER servers	182
8.2.6	Powering on the Enterprise controllers	183
8.3	Managing the Enterprise controller	184
8.3.1	Varying devices offline	184
8.3.2	Notifying the Enterprise controller of change in ProtecTIER server status	185
8.3.3	Performing a standalone mount	186
8.3.4	Performing a standalone demount	187
8.4	Managing the ProtecTIER server	187
8.4.1	Adding and removing nodes from ProtecTIER Manager	187
8.4.2	Monitoring nodes using the ProtecTIER Manager	189
8.4.3	Upgrading the ProtecTIER software	191
8.4.4	Setting up the ProtecTIER system	195
8.5	Managing the ProtecTIER Manager	199
8.5.1	Installing ProtecTIER Manager	199
8.5.2	Uninstalling the ProtecTIER Manager	202
8.5.3	Managing users of the ProtecTIER Manager system	202
8.6	Managing the ProtecTIER Virtual Tape system	205
8.6.1	Managing cartridges	205
8.6.2	Monitoring system functions using ProtecTIER	209
8.6.3	Setting Control Path Failover	209
8.7	Managing virtual libraries	210
8.7.1	Creating libraries	211
8.7.2	Renaming libraries	211
8.7.3	Deleting libraries	212
8.8	ProtecTIER system	212
8.8.1	Viewing the ProtecTIER system logs	212
8.9	Wizard error messages	214
8.9.1	Generating a ProtecTIER system status report	214
8.9.2	Creating a long-term statistics report	216
8.9.3	Turning on the WTI power switch outlets	216
8.9.4	Removing a cluster member	216
8.9.5	Adding a cluster member	217
8.9.6	Changing World Wide Node Name	218
8.9.7	Disabling defragmentation	219
8.9.8	Disabling data compression	220
8.9.9	Enabling data compression	221
8.9.10	Changing the HyperFactor mode	221
8.9.11	Dumping the trace buffer contents for a node	222
8.9.12	Changing the trace levels	223
8.9.13	Resetting the trace buffer	224
8.9.14	Resetting virtual robots	225
8.9.15	Resetting the virtual tape drives	225
8.9.16	Moving cartridges	226
8.9.17	Checking and repairing errors	227
Chapter 9	Monitoring the system	231
9.1	Monitoring ProtecTIER	232
9.1.1	The Status line	233
9.1.2	The Navigation pane	234
9.1.3	The Systems window	235

9.1.4 The Nodes window	240
9.1.5 The Repository window	244
9.2 Monitoring the ProtecTIER VTL service	248
9.2.1 The Library window	248
9.3 Reporting on ProtecTIER activity	255
9.3.1 The analyze_sessions utility	255
9.4 Monitoring ProtecTIER through ptmon	258
9.5 Monitoring the virtual library through the z/OS host	263
9.5.1 Display SMS Library	263
9.5.2 Display SMS Volume	263
9.5.3 Library Display Drive	264
9.5.4 z/OS host messages	264
9.6 Other user notifications	266
9.7 Enterprise controller tracing, logs and tools	266
Chapter 10. Disaster recovery and failover scenarios	267
10.1 Disaster failover to remote site	268
10.1.1 DR Failover procedure	268
10.1.2 Disaster Recovery failback	268
10.1.3 Disaster Recovery takeover	269
10.2 Enterprise controller (3592-C06) contribution to Disaster Recovery	270
10.2.1 ProtecTIER Server (3958-DD3) contribution to disaster recovery	271
10.2.2 Failover scenarios	271
Part 4. Appendixes	289
Appendix A. Reliability, availability, and serviceability update	291
A.1 Vary virtual devices offline	292
A.2 Notify the Enterprise controller that the ProtecTIER server will be offline	292
10.2.3 RAS licensed internal code update option	293
10.2.4 Notify the Enterprise controller that ProtecTIER server is back online	295
10.2.5 Vary virtual devices online	296
Appendix B. Checklists	297
Client installation responsibilities	298
Client prerequisites for preinstallation tasks	298
Your responsibilities for the 3958-DE2 (TS7680) installation	298
3958-DE2 (TS7680) physical specifications	299
3958-DE2 (TS7680) electrical power ratings	300
Planning worksheets	300
IP address worksheet	300
Customer information worksheet	302
Appendix C. WTI Network Power Switch	307
Related publications	311
IBM Redbooks	311
Other publications	311
Online resources	311
How to get Redbooks	312
Help from IBM	312
Index	313

Archived

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®	Parallel Sysplex®	System Storage®
CICS®	POWER5+™	System x®
DB2®	POWER®	System z®
DS4000®	ProtecTIER®	Tivoli®
DS8000®	RACF®	XIV®
FICON®	Redbooks®	z/OS®
HyperFactor®	Redbooks (logo)  ®	z/VM®
IBM®	RETAIN®	zSeries®
IMS™	S/390®	
MVS™	System p5®	

The following terms are trademarks of other companies:

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

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

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

This IBM® Redbooks® publication introduces the IBM System Storage® TS7680 ProtecTIER® Deduplication Gateway for System z® (3958-DE2) hardware and the IBM System Storage ProtecTIER Deduplication Gateway for System z V1.1 software. These are designed to help address the tape processing needs of System z data centers by improving the data protection infrastructure and more cost effectively managing and protecting critical client data.

Managing this growth has become the primary source of pain for storage professionals, who are grappling with the following challenges:

- ▶ Growing storage acquisition and management costs
- ▶ Data processing administration
- ▶ Shrinking batch windows
- ▶ Demanding service levels

The TS7680 helps alleviate these challenges with the following capabilities.

Virtualization

The TS7680 includes the following virtualization capabilities:

- ▶ Two enterprise control units and two deduplication engines are combined to support a high-availability enterprise deduplication solution with drive and library virtualization.
- ▶ Up to one million virtual tape volumes.
- ▶ FICON® attach up to 8 ports with 256 logical paths/port to 1 PB of raw storage scalability per Gateway.
- ▶ High performance in-line data deduplication.
- ▶ Support for up to 256 virtual devices per Gateway.
- ▶ 128 drive images per node (control unit).
- ▶ O/S support: z/OS® V1R9 and above with PTFs, and z/VM® 5.3 and above including DFSMS/VM FL221 with PTFs.
- ▶ Library Manager functionality is integrated for the virtualization configuration.

Virtual tape volumes

Like the TS7650, the TS7680 also uses virtual tape volumes and offers the following capabilities:

- ▶ Can contain up to one million virtual tape volumes.
- ▶ Each virtual tape cartridge can hold up to 100 GB of capacity (host capacity).
- ▶ Places an enterprise control unit in front of each ProtecTIER node to provide a FICON tape device interface.
- ▶ Uses the ProtecTIER Gateway cluster for fast, efficient deduplicated back-end storage.

Key functions

The key functions of the TS7680 are as follows:

- ▶ Appears as automated tape library with the 3592-J1A (3590 emulation) device.
- ▶ Retains DFSMS (SMSTAPE) functionality.

- Features integrated library management functionality.
- Includes the scratch processing delete function.

The team who wrote this book

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

Alex Osuna is a Project Leader at the International Technical Support Organization, Tucson Center. He writes extensively about all areas of storage. Before joining the ITSO five years ago, Alex worked in the Tivoli® Western Region as a Principal SE on Tivoli storage. He has 32 years of experience in the IT industry, focused mainly on storage. Alex holds certifications from IBM, Microsoft®, Red Hat, and the Open Group.

Luciano Cecchetti is an IBM Senior Accredited Product Service Representative, working since 1985 in technical support roles in Rome, Italy. In 1988, he provided second-level support for ES/9000 during an international assignment at the IBM European Product Support Group in Montpellier, France. As a specialist member of the RMSS SWAT Team, he supported and coordinated the first 3494 VTS/Library installations in Portugal and Greece. Since 2002, he has held a Specialist (CS) degree in High End Tape Solutions from the IBM Professional Certification Program. Since 2007, he has held a nomination for Senior membership in the World-Wide Product Services Profession. In 2010 he joined the Mainz VET/PFE (Virtual EMEA Team), working as Back Office Specialist for EMEA Tape Support. His areas of expertise include CPUs, disks, and tapes in complex mainframe environments.

Erik Franz is currently working in the IBM Tape Support Front Office in Mainz supporting IBM System Storage Tape products in the Maintenance and Technical Support (MTS) organization. Previously he worked in the IBM Systems Lab in Mainz, which is part of the European Storage Competence Center (ESCC). Erik was responsible for performing customer workshops and proof of concepts for IBM disk storage, SAN, and all IBM server products. Erik received his diploma from the University of Applied Sciences in Worms and wrote his thesis on “GPFS - Configuration, Implementation, Service Setup and Design” for the STG Lab Services Europe team.

Roberto Mencarelli is an IBM Senior Accredited Product Service Representative working since September 2006 in technical support roles in Rome, Italy. Roberto joined IBM in 1981 as HW Customer Engineer mainly dedicated to management of big clients in the banking sector. In 2005 he became the coordinator for the Mediosystem Strategic Outsourcing center of IBM. Next year he will be the technical contact person responsible for the Monte dei Paschi di Siena one of the largest banks in Italy.



Figure 0-1 Erik, Luciano, Roberto, Alex

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

Thomas Grave
IBM Tape and ProtecTIER Product Management Tucson

Avner Kedmi
Sr. Product Manager, IBM ProtecTIER Deduplication Solutions

Jennifer Mason
Tucson, Arizona

James R. Whelan
Tape Subsystem SLT Test

Abraham Faibish
Technical Product Manager - ProtecTIER

Gerard Kimbunde
IBM Systems & Technology Group, Systems Hardware Development

Justin Hildebrandt
Regional Sales Executive, IBM ProtecTIER Solutions IBM Asia Pacific and Japan

Carl Bauske
Americas Advanced Technical Sales Specialist

Ericka M. Dawson
z/OS Tape Storage Software Architecture

Ralph Beeston
IBM Storage Systems Division

Now you can become a published author, too!

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

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

Comments welcome

Your comments are important to us!

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

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

ibm.com/redbooks

- ▶ Send your comments in an e-mail to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

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

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:

<http://www.facebook.com/IBMRedbooks>

- ▶ Follow us on Twitter:

<http://twitter.com/ibmredbooks>

- ▶ Look for us on LinkedIn:

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

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

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

- ▶ Stay current on recent Redbooks publications with RSS Feeds:

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



Part 1

Introduction to the TS7680

In this part we introduce architecture, planning and give a basic introductions to the TS7680. .

Archived



Concepts of data deduplication

This chapter describes general data deduplication concepts and type, data deduplication methods, system components, and benefits of data deduplication.

1.1 What data deduplication is

Data deduplication essentially refers to the elimination of redundant data. In the deduplication process, duplicate data is deleted, leaving only one copy of the data to be stored. However, indexing of all data is still retained should that data ever be required. Deduplication is able to reduce the required storage capacity since only unique data is stored.

End-to-end data deduplication for System z is illustrated in Figure 1-1.

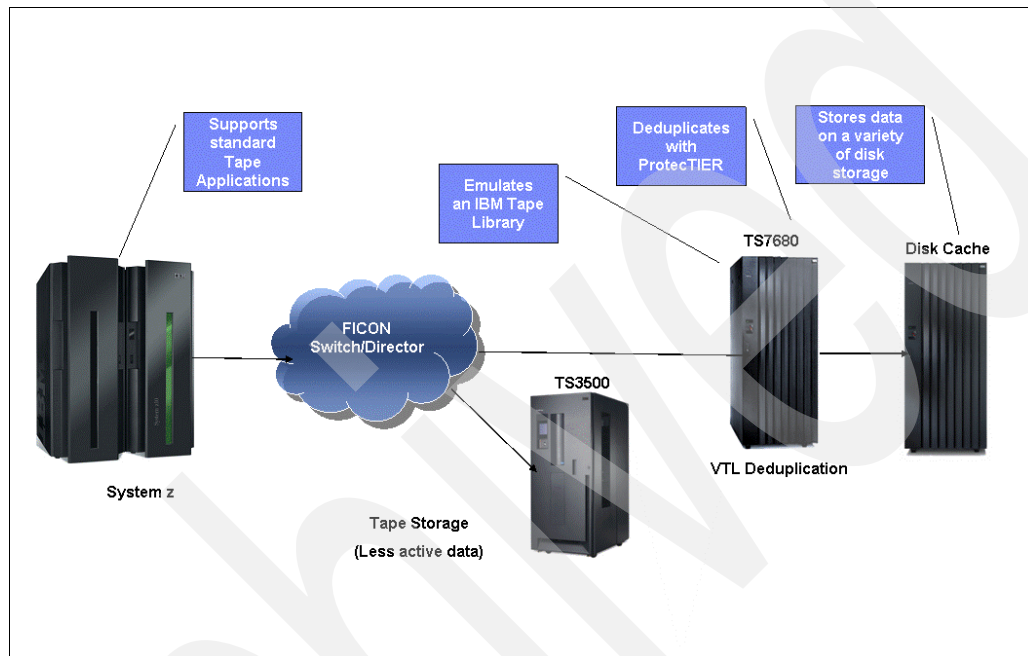


Figure 1-1 End-to-end data deduplication for System z

With data deduplication, data is read by the data deduplication system while it looks for duplicate data. Different data deduplication products use different methods of breaking up the data into elements, but each product uses some technique (see 1.2, “Type of data deduplication” on page 5) to create a signature or identifier for each data element. Whether using in-line or post processing data deduplication (see 1.3, “Data deduplication processing” on page 8), data element signature values are compared to identify duplicate data. After the duplicate data is identified, one copy of each element is retained, pointers are created for the duplicate items, and the duplicate items are not stored.

The effectiveness of data deduplication is dependent upon many variables including the data change rate, the number of tape processing tasks, and the data retention period. For example, if you back up the exact same uncompressible data once a week for six months, you save the first copy and do not save the next 24, which would provide a 25 to 1 data deduplication ratio. If you store an uncompressible file on week one, storing the exact same file again on week two and never back it up again, you have a 2 to 1 deduplication ratio.

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

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

Data deduplication can reduce your storage requirements, but the benefit you derive is determined by your data and your tape processing policies. Workloads with a high database content generally have the highest deduplication ratios. Compressed, encrypted, or otherwise scrambled workloads typically do not benefit from deduplication.

1.2 Type of data deduplication

Many vendors offer products that perform deduplication. Various methods are used for deduplicating data. Three methods frequently used are hash-based, content aware, and HyperFactor®.

1.2.1 Hash-based

Hash-based data deduplication methods use a hashing algorithm to identify “chunks” of data.

Commonly used algorithms are Secure Hash Algorithm 1 (SHA-1) and Message-Digest Algorithm 5 (MD5). When data is processed by a hashing algorithm, a hash is created that represents the data. A hash is a bit string that represents the data processed. If you processed the same data through the hashing algorithm multiple times, the same hash is created each time, adding to the overhead. Here are some examples of hash codes:

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

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

In Figure 1-2 on page 6, data chunks A, B, C, D, and E are processed by the hash algorithm, which creates hashes Ah, Bh, Ch, Dh, and Eh; for purposes of this example, we assume this is all new data. Later, chunks A, B, C, D, and F are processed. F generates a new hash, Fh. Since A, B, C, and D generated the same hash, the data is presumed to be the same data, so it is not stored again. Since F generates a new hash, the new hash and new data are stored.

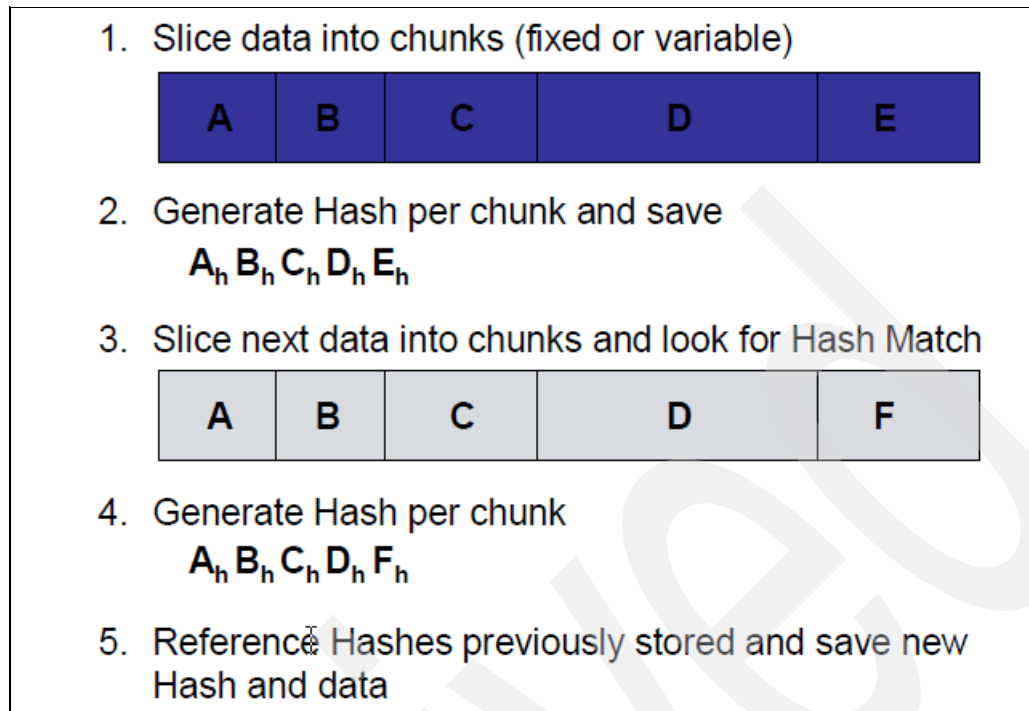


Figure 1-2 Hash-based deduplication

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

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

1.2.2 Content aware

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

In Figure 1-3 on page 7, with content-aware deduplication, the system looks through the data to find the fully qualified names and then looks for previous versions of the same file. If a match is found, the system performs a bit-by-bit comparison and updates the reference points as needed.

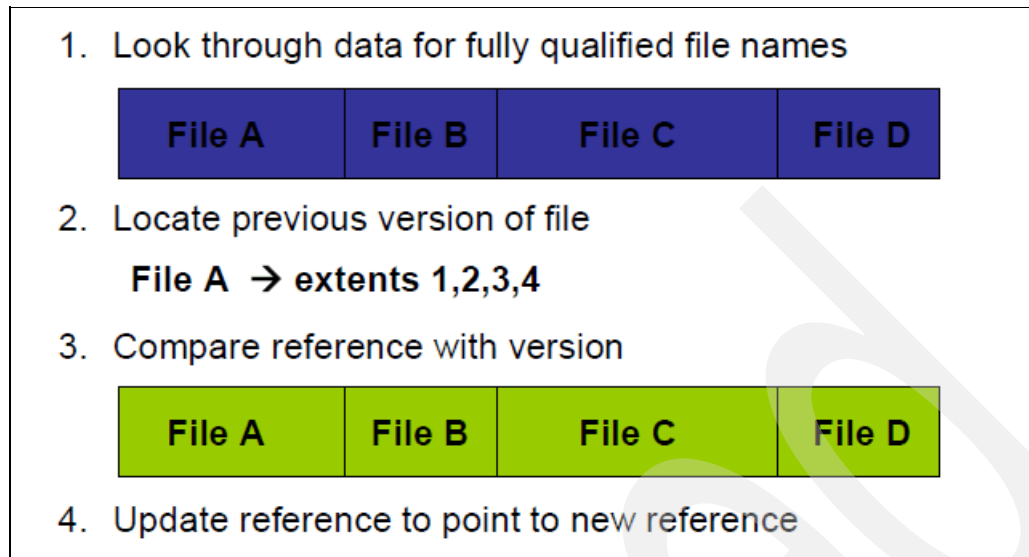


Figure 1-3 Content-aware data deduplication

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

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

1.2.3 HyperFactor and deduplication

The cornerstone of ProtecTIER is HyperFactor, an IBM technology that deduplicates data inline as it is received from the System z tape processing. ProtecTIER's 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.

When new data is received by ProtecTIER deduplication technology, HyperFactor finds any similar data elements that have already been stored. This search is very quick, using a small and efficient memory-resident index. Once the similar data elements are found, HyperFactor can then compare the new data to the similar data to identify and store only the byte-level changes.

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 very well.

In Figure 1-4, 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.

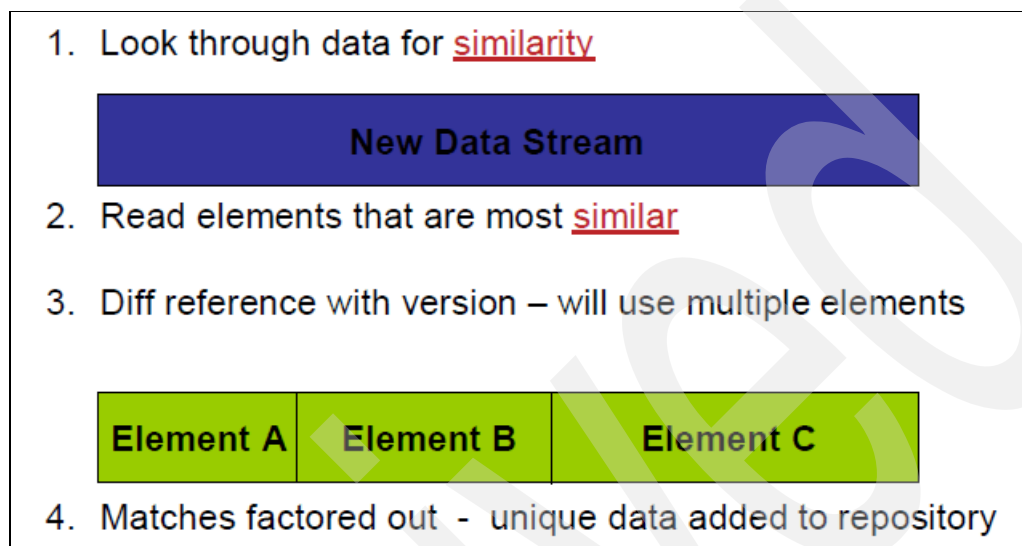


Figure 1-4 HyperFactor data deduplication

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, since 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 meta data files and fetch the data according to the pointers they contain.

1.3 Data deduplication processing

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

1.3.1 In-line method

In-line data deduplication is not dependent on the type of data deduplication used. An advantage of in-line data deduplication is that the data is only processed once and there is no additional processing after the tape processing window. In-line data deduplication requires less disk storage since the native data is not stored prior to data deduplication. Depending on the implementation, a disadvantage of in-line data deduplication is that the data deduplication processing could slow down the tape processing data stream. Algorithms used for data

deduplication can be processor intensive and data deduplication may require additional read or write access if the index is disk-based.

1.3.2 Post-processing method

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

There are disadvantages to the post-processing method:

- ▶ Increased input/output (I/O) to the storage device. Since the data is written during the tape processing, reads to identify duplicate data and the pointer(s) must be updated if there is duplicate data. Overall, the data deduplication cycle will likely be longer than if performed in-line.
- ▶ More disk is required than with an in-line method, because all the data must be stored prior to deduplication.

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

1.4 Components of a data deduplication system

Data deduplication systems implementations vary. Figure 1-5 on page 10 gives an example of common components. You can have a hash-based data deduplication system that uses the post-processing method and another hash-based data deduplication system that uses the in-line method. Some systems may integrate all components required for data deduplications and others may need to be integrated by you.

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

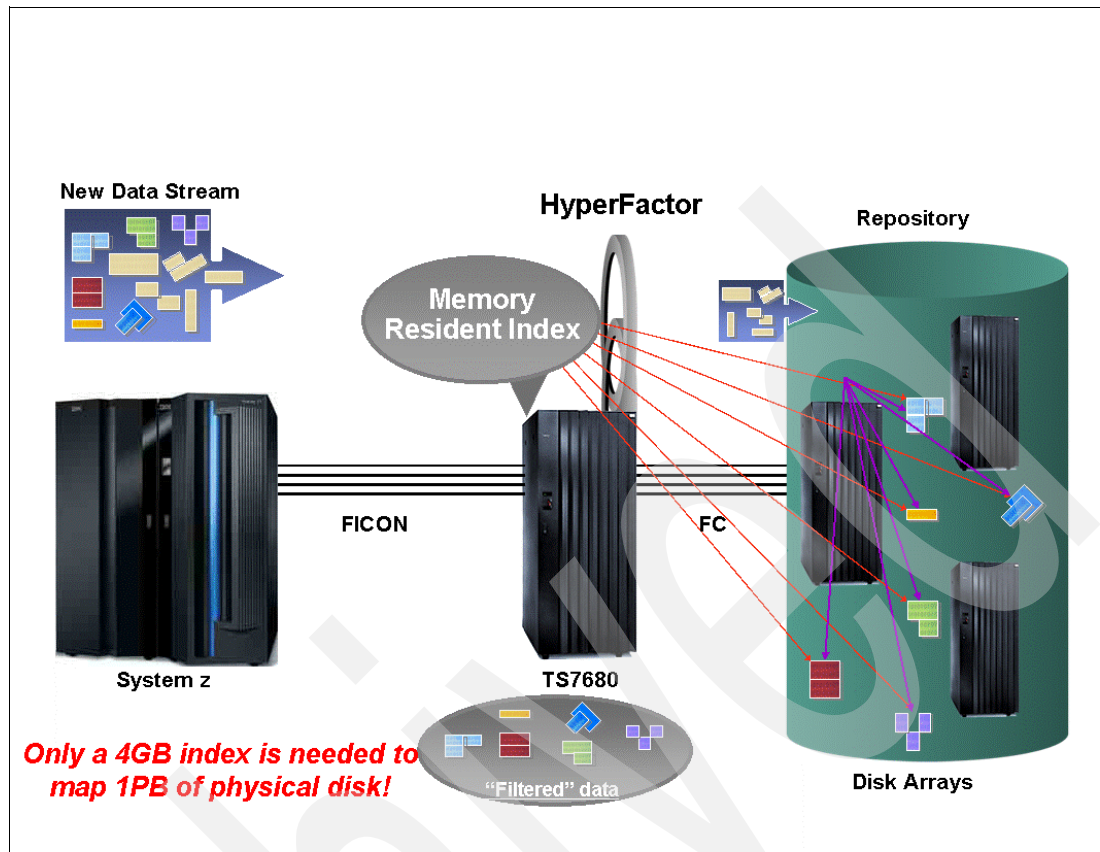


Figure 1-5 Components of a data deduplication system

1.4.1 Server

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

1.4.2 Data deduplication software

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

1.4.3 Disk array

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

1.5 Benefits of data deduplication

When appropriately used, data deduplication can provide benefits over traditional tape processing needs to virtual tape libraries. Data deduplication enables remote vaulting of tape processing data using less bandwidth, because only changed data is shipped to the remote site. Figure 1-6 gives some examples of bandwidth savings. Long term data retention for local or offside storage may still be achieved most economically with physical tape.

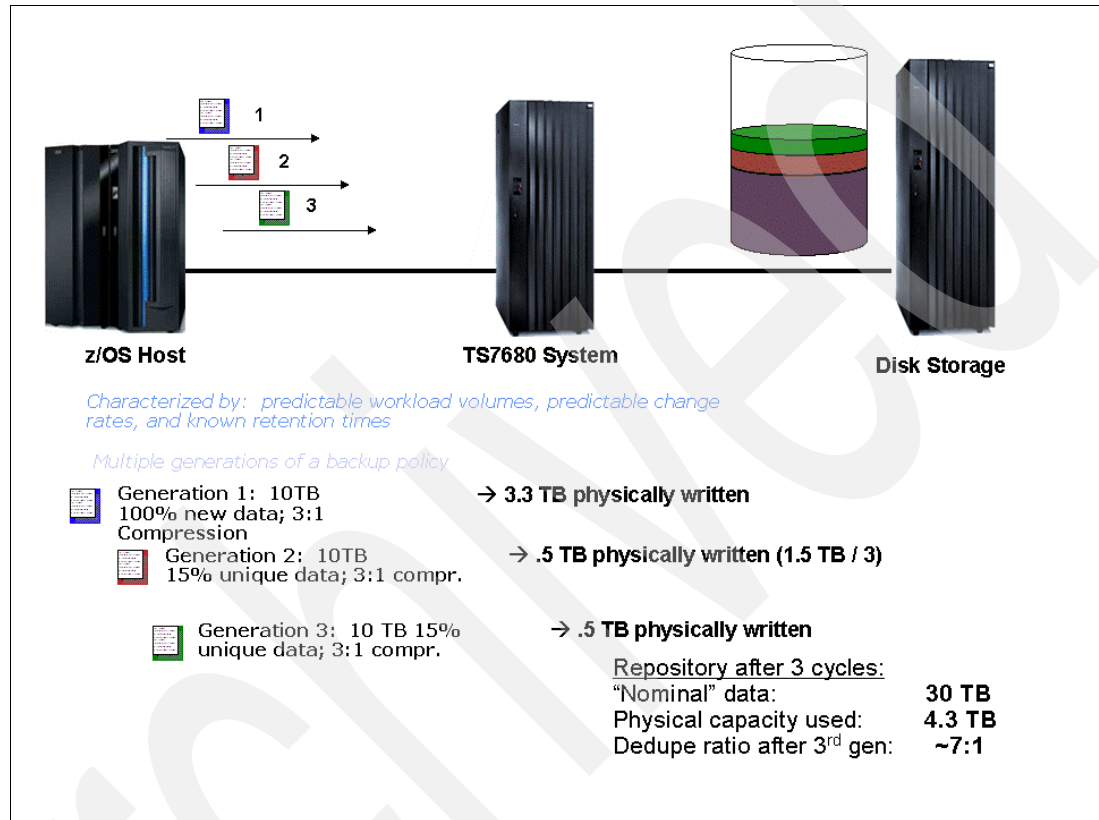


Figure 1-6 Deduplication example in a "normal" tape processing environment

1.5.1 Reducing storage requirements

With the amount of data corporations are required to maintain for compliance with government regulations and for normal business operations, data deduplication can reduce the amount of disk storage required to store data and keep it "online". Performing restores from disk can be faster than restoring from tape and having the data online for longer periods reduces the possibility that the data required may have been shipped offsite.

1.5.2 Reducing of environmental costs

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

Archived

The TS7680 ProtecTIER Deduplication Gateway architecture

In this chapter we describe the TS7680 ProtecTIER Deduplication Gateway architecture (see Figure 2-1) and discuss the following topics:

- ▶ TS7680 ProtecTIER Deduplication Gateway technology overview and concepts
- ▶ Virtual Tape Library with HyperFactor
- ▶ Library Manager Integration
 - Current enterprise control units require a Library Manager to process library commands from the host.
 - TS7680 has integrated this function into the control unit.
- ▶ High availability



Figure 2-1

2.1 TS7650/TS7680 Differences

Some of the differences in architecture between the TS7650 and TS7680 are as follows:

2.1.1 Shelf / visibility

- ▶ TS7650 – Volumes visible at one site at a time only. Think of them as being moved by truck.
- ▶ TS7680 – Volumes replicated to DR site are visible and R/O.

2.1.2 Replication policies by VOLSER (barcode)

- ▶ TS7650 – Applications use specific barcodes and replication policies can be created for specific uses.
- ▶ TS7680 – z/OS applications use one common scratch pool for all applications and use of the VOLSER for creating replication policies is not practical.

2.2 IBM System Storage TS7680 ProtecTIER Deduplication Gateway for System z technology overview and concepts

IBM System Storage TS7680 ProtecTIER Deduplication Gateway for System z is a virtual tape library (VTL) that appears to applications as one automated tape library. It is designed to help address back-up application needs of System z data centers. The TS7680 ProtecTIER Deduplication Gateway for System z V1.1 software is typically able to support implementations with minimal change to existing data center policies, practices, or procedures. The host application can access virtual tape drives and cartridges just as it would in a physical tape library environment. Typical System z VTL solutions provide about 2.5:1 compression of the data. The IBM System Storage TS7680 ProtecTIER Deduplication Gateway for System z, using HyperFactor data deduplication technology, can help significantly (by up to 10 times) reduce the amount of storage capacity required.

Data protection pain points in the data center: Enterprise data centers require a high-performance, high-capacity, reliable, and scalable disk-based data protection solution. Delivering high-level data protection is increasingly difficult for many reasons:

- ▶ Information volume is growing almost exponentially.
- ▶ Shrinking tape's data processing windows require more powerful and efficient high-performance data storage solutions.
- ▶ Pressure to lower capital expenditures for data infrastructure is driving the need for a simpler way to deploy and manage data protection solutions.
- ▶ Increased need to more efficiently manage disk space of existing infrastructures by improving utilization of current hardware.
- ▶ Increasing need to support more efficient, easier to manage, disaster recovery strategies.

The TS7680 combines a virtual tape library solution, inline data deduplication powered by IBM's unique and patented HyperFactor technology, and disk-based storage options to provide users an optimal disk-based target for System z applications that traditionally use

tape. Designed to simplify tape processing operations and improve tape application performance while reducing infrastructure costs, the TS7680 offers high-performance inline data deduplication, highly available two node clustering and up to 1 petabyte (PB) of physical storage capacity per system.

The IBM System Storage ProtecTIER for System z Gateway (TS7680) provides a connection to System z hosts to transfer your data to and from the TS7680 through FICON Channel connections. The TS7680 combines two Enterprise controllers and two ProtecTIER servers to provide a high-availability enterprise deduplication solution with drive and library virtualization. While the Enterprise controllers provide the connectivity to your System z hosts and control the operation of the system, the ProtecTIER servers provide connectivity to the disk repository and control the identification and tagging of duplicate data.

The IBM System Storage ProtecTIER for System z Gateway solution, comprised of IBM System Storage TS7680 ProtecTIER Deduplication Gateway for System z (3958-DE2) hardware and the IBM System Storage ProtecTIER Deduplication Gateway for System z V1 software, is designed to address back-up applications and data protection needs of enterprise data centers, providing a high degree of scalability and flexibility and best of breed inline deduplication for enterprise environments. Some of the applications that could benefit from the TS7680 would be IBM DFSMSdss, IBM DB2® or FDR, and DFSMShsm.

The TS7680 ProtecTIER Deduplication Gateway for System z is designed to offer:

Virtualization

- ▶ Two enterprise control units and two deduplication engines are combined to provide a high-availability enterprise deduplication solution with drive and library emulation.
- ▶ Up to one million virtual tape volumes.
- ▶ FICON attach up to 8 ports with 256 logical paths per port.
- ▶ Up to 1 PB of raw storage scalability per Gateway system.
- ▶ In-line data deduplication.
- ▶ Support for up to 256 virtual devices per two-node cluster.
- ▶ Can contain up to one million virtual tape volumes.
- ▶ Each virtual tape cartridge can hold up to 100 GB capacity (host capacity).
- ▶ Places an enterprise control unit in front of ProtecTIER to provide a FICON tape device interface.
- ▶ Uses the ProtecTIER Gateway cluster for fast, efficient deduplicated backend storage.
- ▶ Uses the availability of the ProtecTIER Gateway to support highly available access to the virtual volumes on disk.

Key functions

- ▶ Appears as automated tape library with 3592-J1A (3590 emulation) device.
- ▶ Retains DFSMS (SMSTAPE) functionality.
- ▶ Features integrated library management functionality.
- ▶ Includes scratch processing delete function.

Library function high availability

The library function is highly available and is fully functional even with the failure of a server.

- ▶ Monitor process running in each Enterprise Controller monitors the health of its own system and communicates status to the other Enterprise Controller.

Monitors:

- DB2 High Availability Disaster Recovery (HADR) status
- TCP/IP between C06's
- Library function status
- TS7680 ProtecTIER Deduplication Gateway status

Other features

- ▶ Clustering for higher performance and availability.
- ▶ FICON attach for host and fibre channel attached disk repository.
- ▶ Flexible storage choices and options.
- ▶ IBM System Storage ProtecTIER for System z V1 version is only available with a single virtual tape library.
- ▶ Up to 25x factoring (deduplication ratios will be data-dependent).
- ▶ HyperFactor Deduplication technology for multiplying effective disk capacity.
- ▶ Supports both z/OS and z/VM.
- ▶ Elimination of tape mounting time. This allows your applications that utilize tape to complete much faster. It can also reduce processor wait time, freeing up processing power for other applications.
- ▶ Virtual tape also speeds up restores.

The IBM System Storage TS7680 ProtecTIER Deduplication Gateway for System z solution is an enterprise-class data protection platform designed to protect business information while reducing the amount of physical disk storage space required and improving the reliability and quality of tape processing. See Figure 2-2 for a logical overview.




Figure 2-2 TS7680 logical layout

The TS7680 ProtecTIER Deduplication Gateway is the first IBM product to offer data deduplication technology in a z/OS and z/VM environment.

2.2.1 Enterprise controllers

Two Enterprise controllers are configured within the 3958-DE2 (TS7680). They are equipped with two fibre channel host bus adapters (HBA), four FICON card slots, two Serial ports, and two on-board Ethernet ports. The Enterprise controllers provide System z host attachment to the virtual tape library (VTL).

As part of the manufacturing process, the Licensed internal code is preinstalled on the Enterprise controllers.

The Enterprise controllers are designed to attach to FICON channels on System z servers or through a FICON/FC switch with appropriate levels of system software.

The Enterprise controller provides an interface between Host FICON and Fibre Channel (FC) communication to the ProtecTIER Server.

Note: Use machine type 3958 and model DE2 for service purposes. You can find serial number on sticker located in the rear of the frame.

2.2.2 ProtecTIER Virtual Tape (VTL) service

The ProtecTIER Virtual Tape (VTL) service emulates a TS3500.

By emulating tape libraries, ProtecTIER VTL enables you to transition to disk data storage without having to replace your entire tape processing environment. Your existing tape processing application can access virtual robots to move virtual cartridges while ProtecTIER actually stores data on a virtual volume on a deduplicated disk repository on the storage fabric.

2.2.3 ProtecTIER servers

The ProtecTIER servers are equipped with two Emulex fibre channel host bus adapters, two Qlogic host bus adapters, two gigabit Ethernet adapters, and two on-board Ethernet ports.

As part of the manufacturing process, Red Hat Enterprise Linux®, ProtecTIER software, and the Reliability, Availability, and Serviceability (RAS) package are preinstalled on the ProtecTIER servers.

Only one type of server is used in the TS7680 ProtecTIER Deduplication Gateway:

3958 DD3; see Figure 2-3 as front view reference.




Figure 2-3 3958 DD3 front view

This is a newer, higher performance server available since March 2009. It is based on the IBM System x3850 M2 Type 7233. It is used as a server in the TS7680; its machine type and model are 3958 DD3.

Note: Use machine type 3958 and model DE2 for service purposes. You can find serial number on sticker located in the rear of the frame.

2.2.4 TS3000 System Console

When a TS3000 System Console (TSSC) is purchased with the 3958-DE2 (TS7680), the TSSC and its related components (keyboard, video, mouse [KVM], KVM switch, and 16-port network switch) are factory-installed in the frame. As a result, the KVM switch provides input and video connections for each ProtecTIER server in the frame. A keyboard or graphics-capable monitor is needed for the 3958-DE2 (TS7680) configuration and no USB. It is an option to purchase the TSSC with the TS7680 or attach to another TSSC already purchased, but one or the other is required.

2.3 ProtecTIER virtual tape (VTL) service and concepts

This section gives more details on the operations of the TS7680.

2.3.1 Virtual tape device

The TS7680 emulates a 3592-J1A tape drive. The 3592-J1A drive does not support encryption nor will the TS7680.

The TS7680 will both deduplicate and compress data. However, due to the variability of the deduplication, the actual DD/compression ratio will not be reported for each individual drive. Each drive will report a ratio of 1:1.

2.3.2 Virtual tape volume

The virtual tape cartridges emulated by the TS7680 will be JA media with a capacity set at 100 GB. The 3592-J1A tape drive only supports reading and writing in one format. Therefore, multi-format support is not required for the TS7680.

The Enterprise Controller generates pseudo labels for each volume inserted or scratch in the TS7680. The volume label follows the software standard for virtual tape devices. There are two 80-byte records, the first containing VOLxxxxxx, the remainder padded with zeros. The second 80-byte record contains the value HDR1 at the beginning of the record,.

The virtual volumes each allow up to 100 GB of data storage. Early warning for End Of Medium is reported 10 MB prior to the full capacity of the volume.

Note: The TS7680 does not support back-end tape directly.

2.3.3 Cache management

When the TS7680 detects that available cache space has fallen below preset thresholds, the TS7680 sends CBR3792E (limited cache free space warning -state reached) and CBR3794A (out of cache resource critical state reached) -attention messages to the attached hosts. The CBR3792E message can be used to trigger return to scratch processing or the copying of data to another library. If the amount of available cache subsequently reaches the critical state, all fast ready (scratch) mounts are failed and any specific mount operations are allowed; however, any attempt to write to the volume will be failed. Mount operations that have been accepted before entering this state complete and volumes currently mounted can continue to perform host I/O operations. As appropriate, the VARY I SMS,STORGRP operator command can also be used to steer scratch allocations to another library that is eligible for the scratch request. The DISPLAY I SMS,LIBRARY command with DETAIL, can also be used periodically to display the CACHE PERCENTAGE USED.

When the host returns a volume to scratch, the TS7680 applies (by default) a nine day grace period to the volume. The default grace period can be overridden at install time and set to a value from 0-9 days (0 indicating no grace period). After the grace period elapses, the data associated with the scratch volume is deleted from the library to free up back-end disk space, rendering the contents of the tape volume unusable. When the TS7680 reaches the critical Out of Cache Resources state (CBR3794A), it can also be configured (at install time) to automatically delete data associated with scratch volumes that are in the “grace period.” Volumes with the shortest time remaining in the grace period are deleted first. The default behavior is to honor the grace period.

2.3.4 HyperFactor

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 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.

HyperFactor saves space by taking advantage of the fact that only a very small percentage of data actually changes from one tape processing window to the next tape processing window

(see Figure 2-4 on page 20). 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.

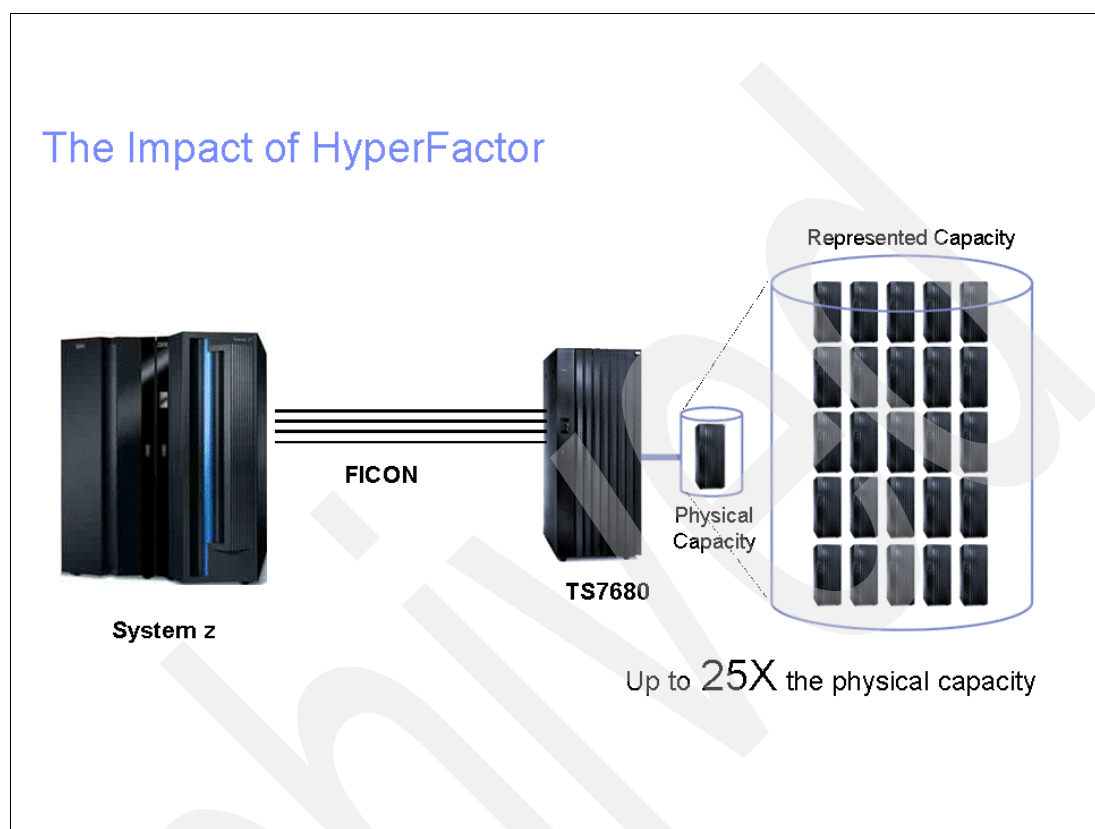


Figure 2-4 The impact of HyperFactor

2.3.5 Cartridge eject

In a virtual tape library, ejecting a cartridge results in the total loss of the data that is stored on that volume. As a safety practice, the Enterprise Controller requires a volume to be empty before allowing the z/OS host to eject a cartridge. In order to eject a volume, the volume must either be empty (never been written to before) or it needs to be known as a scratch volume in the TS7680.

The recommended approach is to eject volumes from the host. This also deletes the volumes going outboard. It is not recommended to initiate the eject through PT Manager. Doing so will create an out-of-sync condition with the host. Also, ejecting from the host requires fewer steps.

2.3.6 Scratch volume processing

The 3590 New Tape Architecture specification does not contain all of the advanced library functions that are supported in the TS7700 family of virtual tape libraries. One of the functions not available to the TS7680 is the knowledge of which categories are defined by the host as scratch. Therefore, the TS7680 does not have the knowledge necessary to know when the host is moving a volume to the scratch category.

With APAR OA27786, the z/OS host will notify the control unit that the volume being returned to category X'xxxxx' is being moved to scratch. z/VM also provides the same capability at 5.3 and above, including DFSMS/VM FL221 with PTFs.

2.3.7 Grace period for scratched volumes

Unlike scratching a physical volume which does not destroy data on the volume, scratching a virtual volume has the effect of erasing the data in the disk cache repository. As a safety feature against the risk of accidentally scratching volumes and losing data, the Enterprise controller provides a grace period between the time the host indicates its intention to scratch a volume and when the data is actually destroyed and space freed in the disk cache repository. The grace period, also referred to as the scratch delay option, is set by the SSR at the time of install.

The grace period may have the adverse effect of keeping the cache in a full usage condition when space could be freed for volumes that have been scratched. There are two methods for working around this issue.

The host may issue the **eject** command on volumes that have been scratched but are still in the grace period. The data for these volumes will then be immediately destroyed and the space made available.

The Enterprise Controller may reduce or waive the grace period on scratched volumes in the critical cache full condition (see 2.3.3, "Cache management" on page 19). This is only true if it sets the scratch delay override option.

2.4 ProtecTIER virtualization and deduplication

A virtual tape library is a software application that emulates physical tape while storing your data on a disk subsystem. Some vendors only sell the software and you have to assemble your own system by procuring the supported components (server, disk subsystem, and switches), and then integrating them (a la carte). Other vendors may provide an appliance in which all components are included and integrated. Still other vendors may provide a solution somewhere between the "a la carte" and the appliance solution.

Regardless of the procurement and integration method, all VTLs provide the basic function of emulating physical tape. To perform this function, the VTL software must accept the normal tape commands issued by a tape processing application and provide the appropriate responses back to that application. The VTL software must also store your data on the random access disk subsystem that allows it to retrieve the data and present it to the tape processing application as a sequential stream. The basic concept of VTL data storage is shown in Figure 2-5 on page 22.




Figure 2-5 Virtual tape library concepts

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

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

A virtual tape library functions in the same way. When the tape processing application expires a virtual tape cartridge, the virtual tape cartridge stays in the library and the data contained on the virtual tape cartridge remains until the tape processing application performs a scratch write (write from BOT). When viewing the space utilization for the virtual library, the expired virtual tape cartridges continue to use allocated space until the scratch write is performed or the scratch delay time elapses, and then the data is automatically deleted. If the utilized space on your virtual tape library is high but you have sufficient virtual tape cartridges in expired or scratch status in the tape processing application, you may still be able to work within your installed capacity. If the utilized space on your virtual tape library is high and you

do not have sufficient virtual tape cartridges in expired status in your tape processing application, you may need to consider increasing your capacity.

2.4.1 HyperFactor

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

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

HyperFactor takes a different approach and therefore reduces the phenomenon of missed factoring opportunities. Rather than depending on exact chunk matches, HyperFactor finds matches between “similar data.” When new data is received by ProtecTIER, HyperFactor finds any similar data elements that have already been stored (see Figure 2-7 on page 25). This search is extremely quick, using a small and efficient memory-resident index. Once the similar data elements are found, HyperFactor can then compare the new data to the similar data to identify and store only the byte-level changes.

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 very well.

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

The core technology in ProtecTIER-HyperFactor is a series of algorithms that factor, or deduplicate, data efficiently. In each new data store operation, HyperFactor finds the data in common with data store operations. This common data in the new data store operation is effectively “filtered out” and pointers are used to reference existing data in the repository (see Figure 2-6 on page 24). The net effect is that the entire content of the new data is stored in the space required to only store the small fraction of it that is truly new data.

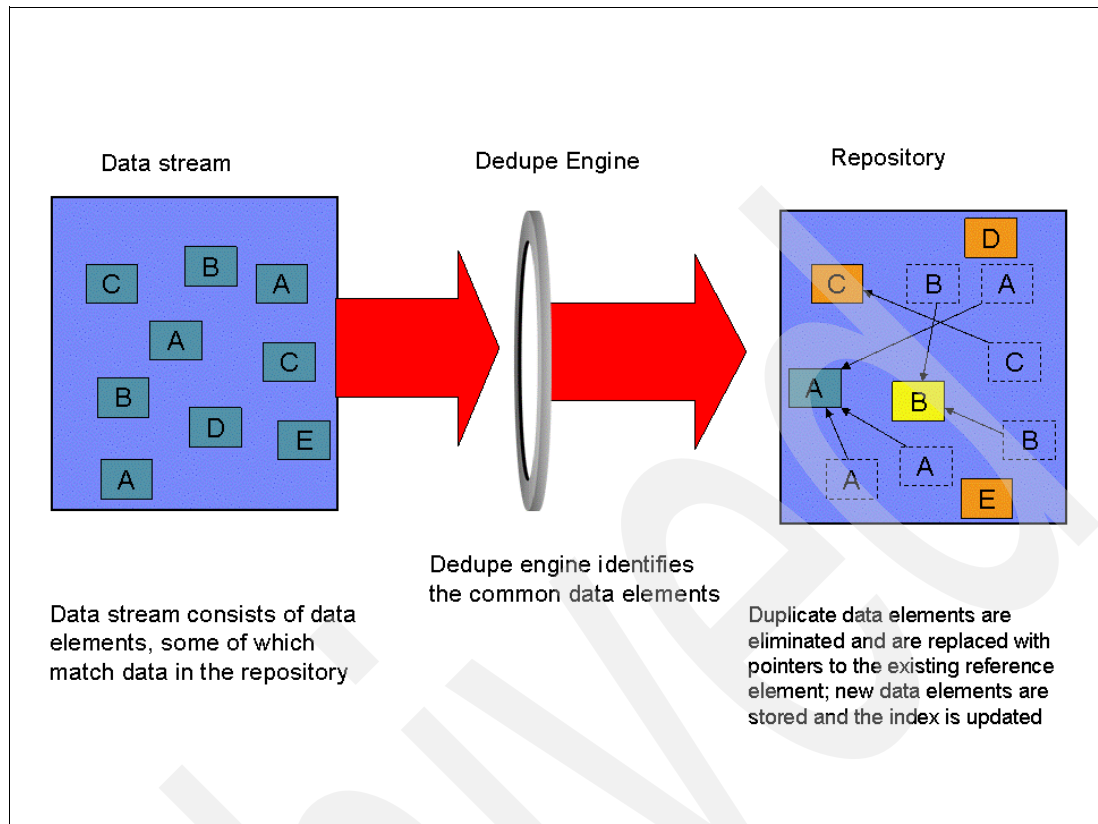


Figure 2-6 Dedupe engine

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

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

- **Data Retention Period**

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

- **Data Change Rate**

The rate at which the data received from the tape processing application changes from a data store to the next data store. This measurement has most relevance when “like” tape processing policies are compared. (Data change rates may range from 1% to >25%, but are difficult to observe directly.)

ProtecTIER data ingest flow

ProtecTIER performs the deduplication process on data ingest, which means data deduplication is performed in-line. The data flow is shown in Figure 2-7 on page 25.

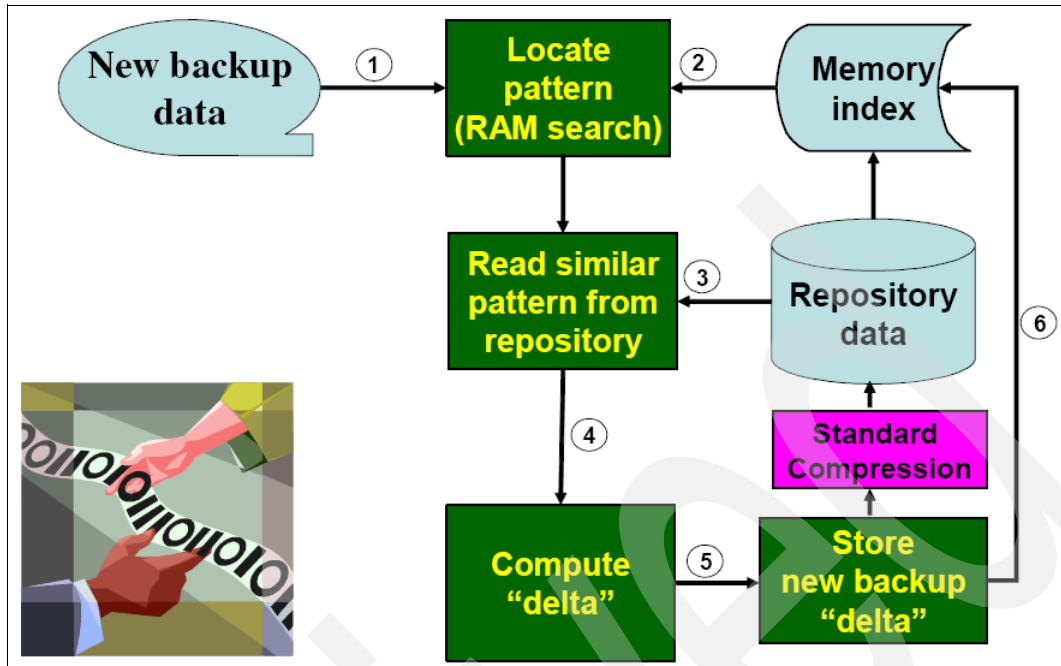


Figure 2-7 ProtecTIER software data ingest flow

The data flow is as follows:

1. A new data stream is sent to the ProtecTIER Server, where it is first received and analyzed by HyperFactor.
2. For each data element in the new data stream, HyperFactor searches the Memory Resident Index in ProtecTIER to locate the data in the repository that is most similar to the data element.
3. The similar data from the repository is read.
4. A binary differential between the new data element and the data from the repository is performed, resulting in the delta difference.
5. The delta from Step 4 is now written to the disk repository after being processed with the Lempel-Ziv-Haruyasu (LZH) compression algorithm. With LZH compression, additional size reduction may be achieved for delta data. Some size reduction may be accomplished for new data (such as the initial tape processing 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 Meta Data file system frequently.

Once 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. Since the Memory Resident Index is only used for data deduplication on data ingest, data accessibility remains if the index is lost. The Memory Resident Index is restored from the Meta Data 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 very short window. The only impact from this would be a slight, probably unmeasurable, reduction in the overall deduplication ratio.

2.4.2 ProtecTIER Virtual Tape Library (VTL) concepts

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

When duplicate data is identified by ProtecTIER, it updates a reference count in the database. ProtecTIER uses the reference count to determine when a data segment can be overwritten (deleted). As shown in Figure 2-8, Sector 3 represents a segment that occurs four times within the virtual cartridges in the repository. In the lower left corner is a representation of the reference table showing that Sector 3 is referenced four times.

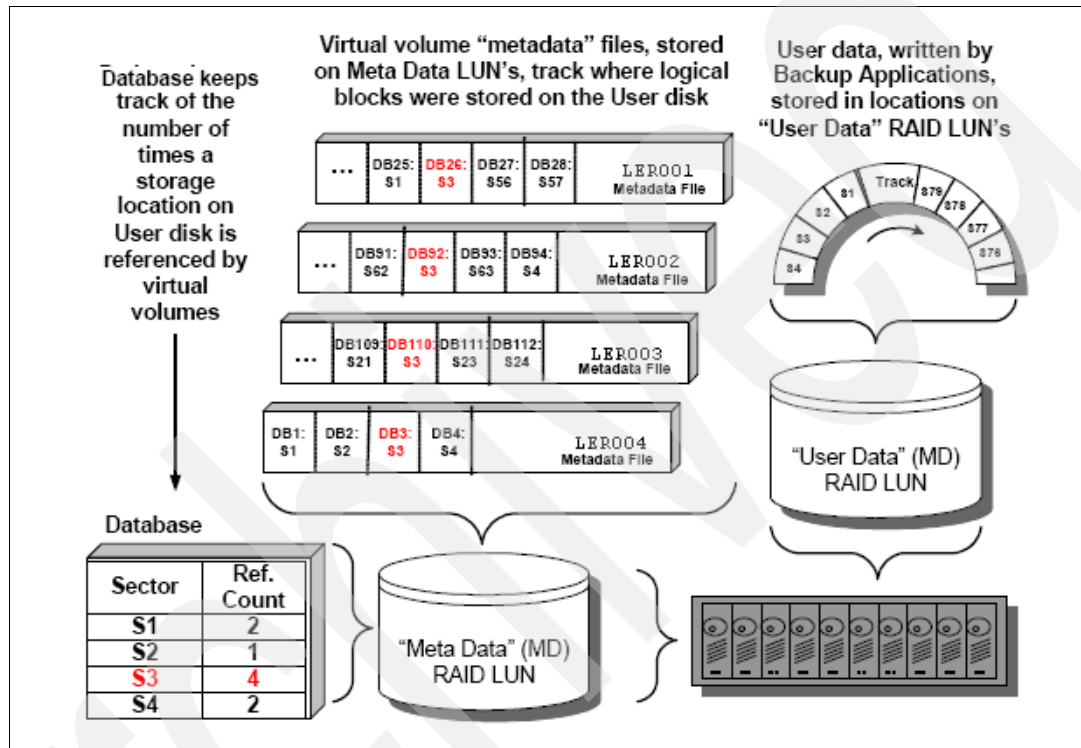


Figure 2-8 Virtual Tape Library concepts_2

ProtecTIER uses the Meta Data files to read back the virtual cartridges. When a virtual cartridge is overwritten or deleted, the reference count for each segment of User Data on a virtual cartridge is decremented. Once 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.

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 tape processing operations. To understand the concept, it may 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 scratch every night to provide the media required for the next day's tape processing window.

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

Until all available scratch cartridges are used, ProtecTIER is only performing two types of Input/Output (I/O) to the RAID. While doing tape processing, it is performing random reads to the User Data disk to “prove” the duplicate data (90%) and it is performing writes of “new” user data (10%). As that tape processing data is being written to the User Data disk, ProtecTIER is also doing roughly 90% random write I/O to the Meta Data LUNs, as it updates the virtual cartridge Meta Data files, to record where the user data is stored for each virtual cartridge.

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

The first step of that processing reads the contents of each of the old Meta Data 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 Meta Data file are processed, the Meta Data file is deleted. Each time the reference count of a storage block goes to zero, that storage block gets returned to the pool of “free blocks” and becomes usable as free space.

Not all units of free space are usable in ProtecTIER. The smallest unit of usable, or “allocatable space, in the repository is 1 MB. A storage block is 16 K. As storage blocks are freed as a result of an overwrite of a virtual cartridge, some of the space freed will be in amounts that are less than 1 MB of contiguous blocks, and the system needs to 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 Meta Data files associated with the virtual cartridge being overwritten.

Once 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 tape processing activity, as described previously), plus the new additional work of:

1. Reading and deleting old Meta Data files
2. 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, it is not matching new data to existing data and no fragmentation is occurring this enables high performance. Once steady state is achieved, performance stabilizes. Figure 2-9 on page 28 is a conceptualization of the performance from initial implementation through steady state. The change in performance and the time to reach steady state depends on the size of the repository and the operational characteristics of your data.




Figure 2-9 ProtecTIER performance from initial implementation through steady state

2.5 Library Manager Integration

The current Enterprise controller requires an outboard library manager to support the Automated Tape Library function. This function will be integrated into the Enterprise controller control unit licensed internal code. The outboard library manager function may include an option for dual redundant library managers. The TS7680 will always be configured with redundant library manager function, one cluster acting as the primary and the other in standby mode.

See Figure 2-10 as reference of Enterprise controller schematic blocks and how they are interconnected for Library Manager Integration functions.

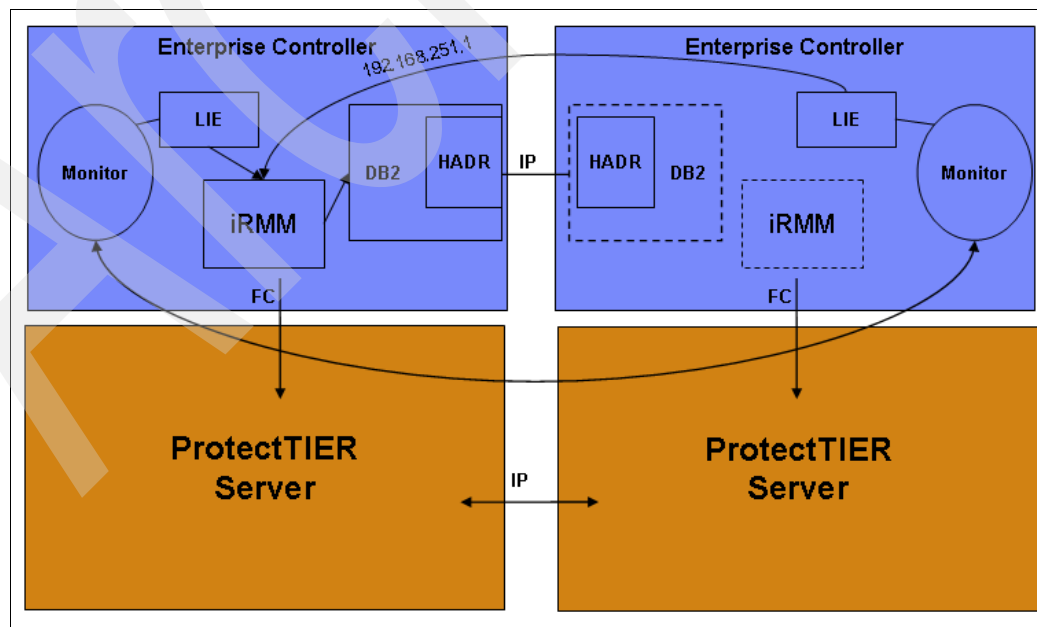


Figure 2-10 Library Manager Integration

Following is a short description of Library Manager Integration components.

Library Integration and Emulation (LIE)

- ▶ Receives Library commands from the host and converts them to IRMM (Integrated Removable Media Manager) commands.
- ▶ Monitors the health of the library and reports any problems.
- ▶ Cleans up unused scratch storage.
- ▶ Ensures that all volumes are available to the host.

Integrated Removable Media Manager (IRMM)

The IBM Integrated Removable Media Manager for the Enterprise (IRMM) package will be used as the media manager, which replaces the function previously maintained in the outboard library manager.

- ▶ Communicates to the library devices on the Virtual Tape Library (VTL)
- ▶ Stores category information in DB2.
- ▶ Calculates statistics.

DB2 High Availability Disaster Recovery (HADR)

The database that was maintained by the outboard library manager will now be contained in a DB2 database on the Enterprise controllers. One controller will act as the primary database server and the other will be the standby. In the event of a failure of the primary server, the standby will take over and become the server.

2.6 High Availability

In the 3958-DE2 (TS7680), two ProtecTIER servers, or nodes, are arranged in a cluster.

A two-node system uses two servers in an active-active cluster and enables a sophisticated system with the High Availability benefit.

High Availability (HA), a requirement in a System z environment, has also been implemented.

Two Enterprise controllers in front of the clustered ProtecTIER Servers are part of the HA architecture.

- ▶ The Library function runs on either controller and switches over as required.
- ▶ All devices on both control units have access to all volumes in the library.
- ▶ Concurrent maintenance of individual components within the C06.
- ▶ Vary off a portion of the host devices to perform concurrent maintenance of some components.
- ▶ Concurrent licensed internal code load by updating one controller at a time.

Clustered ProtecTIER Servers provides high availability access to the virtual volumes on disk and hardware redundancy in the event of a node failure.

- ▶ Both ProtecTIER Servers have access to all volumes in the library.
- ▶ Vary off a portion of the host devices to perform concurrent maintenance.
- ▶ Concurrent licensed internal code load by updating one server at a time.

Refer to Figure 2-11 for a quick view of High Availability components.

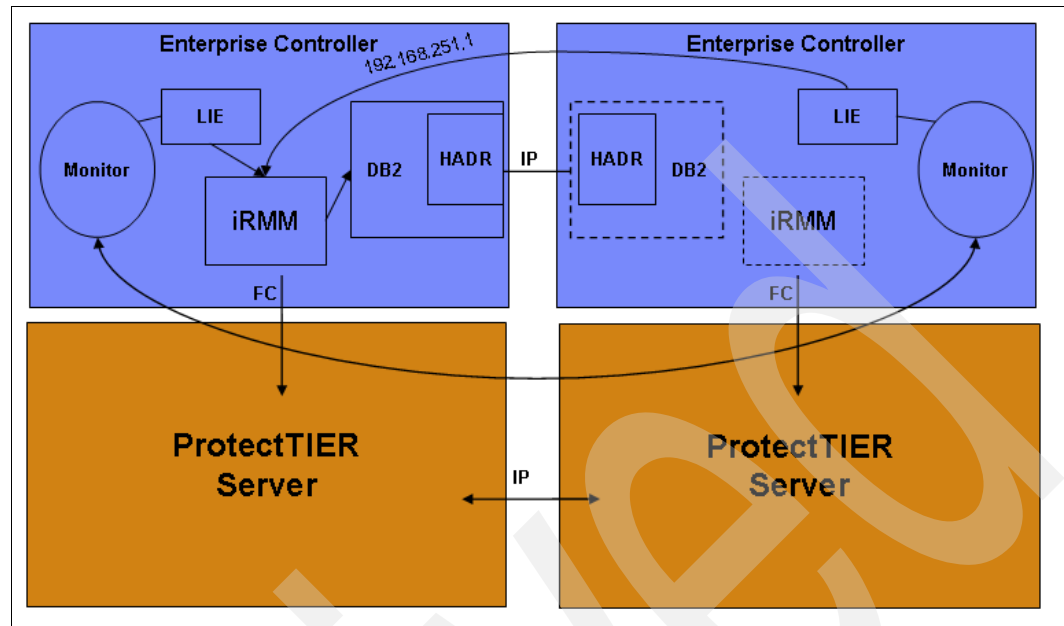


Figure 2-11 TS7680 High Availability

The control units use Integrated Removable Media Manager (iRMM) to carry out library activity. Since DB2 only allows one active database in a high availability disaster recovery (HADR) configuration, iRMM only runs on one control unit at a time. Both control units use the 192.168.251.1 address to communicate to iRMM. If the primary control unit fails, the 192.168.251.1 address is configured on the standby control unit when it initiates a takeover.

A process is running on both control units that monitors the health of each system and detects when the primary control unit is no longer capable of running iRMM. These processes use buffers in the SCSI Enclosure Services (SES) device provided by the ProtectTIER to communicate health status with one another.

The following conditions cause the monitor process on the standby control unit to initiate a takeover:

- ▶ **Loss of fibre channel communication to the ProtectTIER server**
This occurs if all fibre channel links are lost, or if the ProtectTIER server goes down. In this case, the standby detects that the primary control unit can no longer run iRMM and takeover is initiated. Also, a Service Information Message (SIM) is raised by the ermmonMonitor process, and the error log entry indicates FIBER is no longer valid.
- ▶ **A timeout occurred in communicating to iRMM**
This occurs if iRMM has not responded to commands for too long (this is currently 7 minutes). In this case, the standby detects that iRMM is no longer responding and initiates a takeover.

The following conditions do not cause a takeover, but they do cause the system to run in a degraded state that is not highly available:

- ▶ **Loss of Ethernet communication between the control units**
This occurs if both Ethernet paths between the control units are lost. In this case, the primary control unit continues to run iRMM. However, the standby control unit can no longer communicate. All library activity fails on the drives connected to the standby control

unit. Also, a SIM is raised by the ermmMonitor process, and the error log entry indicates IP is no longer functional.

- ▶ A failure in the HADR component of DB2

This should only happen if the Ethernet connection has been lost, but it is possible that it could be caused by an internal DB2 error. When this occurs, the primary continues to run iRMM, and both control units continue to communicate to iRMM using the 192.168.251.1 address if they can. Also, a SIM is raised by the ermmMonitor process, and the error log entry indicates HADR is no longer running.

Network communication between Enterprise controllers

Each control unit in a 3958-DE2 frame has two Ethernet adapters. Each of these adapters is attached to a router that provides connectivity to the TSSC and to the other control unit. The two adapters are aggregated into one logical interface using AIX®'s EtherChannel support. On each control unit, the ent0 and ent1 Ethernet adapters are combined into a single EtherChannel adapter, ent2. If one Ethernet port fails, AIX seamlessly switches to the other port with no interruption to Ethernet communication.

The EtherChannel link aggregation provides a single network, 192.168.251.XX, for communication between the control units. On this network, the lower control unit is configured as 192.168.251.10, and the upper control unit is configured as 192.168.251.20. There is also an alias IP address for the control units to communicate with iRMM on this network. This alias IP address is 192.168.251.1 and is dynamically configured on whichever control unit is currently running as the primary.

Integrated Removable Media Manager (iRMM) HA aspects

The following are the key features implemented on iRMM:

- ▶ Stores library information in a DB2 database.
 - This database is kept in sync on both CUs with the HADR function of DB2.
 - Only one active DB, so the standby will issue a takeover if there is a problem.
- ▶ It is only running on one control unit.
 - When a takeover is necessary, the standby must start iRMM to become the primary.
- ▶ Uses the 192.168.251.1 virtual IP address for communication.
 - Since this is a virtual address, it can be moved to the standby if a takeover occurs.
 - EtherChannel link aggregation is used for high availability.

Library Integration and Emulation (LIE) HA aspects

Library Manager Integration with High Availability is done as follows:

- ▶ LIE is running on both Enterprise controllers.
 - Both CUs have an active LIE component working on library commands.
 - Both LIE components share iRMM and its database.
- ▶ Heartbeat communication between the Enterprise controllers.
 - Done via the existing communication path used by the monitor process.
 - If a failure is detected, the surviving LIE process will demount the cartridges from the drives that have lost communication.

TS7680 High Availability (Monitoring)

The monitor process running in each Enterprise controller monitors the health of its own system and communicates the status to the other Enterprise controller monitor process through the ProtecTIER Server.

Following are the key features introduced by monitor processes:

- ▶ Monitors:
 - DB2 / HADR status
 - TCP/IP between C06s
 - Library function status
 - ProtecTIER Server status
- ▶ When failure of the primary node occurs, the monitor process on standby will take over the primary role.
- ▶ When service is initiated on the primary node, the monitor process on the primary will request standby to take over.
- ▶ Monitoring process mechanism
 - The monitor process on each CU puts information into an assigned buffer on the PT server.
 - This is both a heartbeat message, and a detailed report of the health of each CU as seen by the monitor thread.
 - The monitor process also sends information on behalf of LIE. This information is then read by the monitor process on the other CU and passed to the LIE component.
- ▶ Components checked by the monitor process
 - iRMM
 - If iRMM communication has been lost on both CUs, a takeover is initiated.
 - If iRMM communication is only lost on standby, no takeover will occur.
 - Fibre channel
 - Results in a SIM.
 - If fibre channel communication is lost on the primary, a takeover is initiated.
 - If fibre channel communication is lost on the standby, all takeovers are disabled.
 - TCP/IP communication
 - Results in a SIM.
 - If only TCP/IP communication is lost, all takeovers are disabled.
 - HADR synchronization
 - If the primary and standby databases lose sync, all takeovers are disabled.
 - Failure of an entire CU
 - Results in a SIM.
 - This is detected when we see a double failure. In other words, if iRMM or TCP/IP communication is lost on the standby *and* the primary CU is no longer responding with a heartbeat message, a takeover occurs.

TS7680 High Availability (scenarios)

Following are some scenarios showing how TS7680 reacts in case one component fails or must be taken out of service:

- ▶ Primary CU loses power.
 - The standby will see a TCP/IP failure, and the primary will no longer send the heartbeat message.
 - A SIM will be generated by the standby CU, and the corresponding error log entry reports “IP is no longer functional.”
 - The standby CU initiates a takeover of the database, configures the 192.168.251.1 address, and starts iRMM.
 - All drives on the down CU are demounted.
 - All takeovers are disabled until the down CU is back.
- ▶ TCP/IP failure
 - A SIM is generated by the standby CU, and the corresponding error log entry reports “IP is no longer functional.”
 - The standby LIE loses communication to iRMM and notifies the primary.
 - All drives on the standby CU are demounted.
 - Primary CU continues to function as normal.
 - All takeovers are disabled since the databases will get out of sync.
- ▶ Standby CU loses its PT server.
 - The standby CU loses communication to all fibre channel devices.
 - A SIM is generated by the standby CU, and the corresponding error log entry reports “FIBER is no longer valid.”
 - The monitor process on the standby is no longer able to send status.
 - All drives on the standby CU are demounted.
 - All takeovers are disabled since the standby cannot communicate to the VTL.
- ▶ Maintenance done on the primary CU.
 - When the primary CU is taken offline, it requests the standby CU to start a takeover.
 - If the system is in a good state, the takeover is allowed and should complete successfully.
 - Once the CU is offline, maintenance can be done without interrupting the host on the other CU.

Archived

ProtectTIER native replication overview and operation

ProtectTIER with replication enables virtual tape cartridges to be replicated from the primary site to a secondary location for enhanced disaster recovery (DR) and business continuity (BC) capabilities. By eliminating the need to transport physical tape cartridges, data can be recovered faster and more reliably, enabling users to get back online more rapidly in the event of a disaster or major system outage. The dramatic reduction in the required network bandwidth between the primary and secondary sites enabled by ProtectTIER's deduplication technology radically reduces the costs associated with electronically transmitting data to a remote location for disaster recovery purposes.

By dramatically reducing costs, ProtectTIER with replication enables IT organizations to easily expand the coverage of replication to all of the applications in their environment, as opposed to deploying replication for only a select few applications and tolerating significant risks of data loss and slow recovery for most other applications. Figure 3-1 on page 36 graphically demonstrates the prohibitive costs of traditional replication, and contrasts this to the increased level of DR protection enabled by ProtectTIER. Figure 3-1 on page 36 represents a generic IT environment in which the cost to protect 30% of their data with traditional replication is equal to the cost of covering 100% of the data with ProtectTIER replication.




Figure 3-1 Relationship between cost and percent of environment protected

3.1 Definition of terms

These are some of the more common terms with replication for the TS7680:

- ▶ Spoke - Production site
- ▶ Hub - Disaster recovery site
- ▶ Unidirectional replication - Replication from the spoke to the hub only (“Preferred” mode in TS7700)
- ▶ Native replication - The ProtecTIER replicates deduped data at a volume level
- ▶ EXVA - Extended Volume Attributes, category and scratch information (future includes use for outboard policies)

3.2 How replication works

Replication is an additional feature built into the VTL so that you can pick and choose some or all of your cartridges to be replicated to the DR site. Since ProtecTIER deduplicates data before storing it, only the changes, or unique *elements* of data, are transferred to the DR site over the replication link. This translates into substantial savings in the bandwidth needed for the replication link. Data transfer is started based on several trigger points such as policy-based transfer windows, or movement of the virtual tape to a VTL *export slot* (the VTL emulates import and export slots, as well as opening or closing of the *library door* to *insert* or *eject* a cartridge). Data verification and validation is done at the DR site to ensure integrity of the transferred data prior to making the virtual cartridge or tape available. Figure 3-2 is an example of a deployment for two systems using replication.

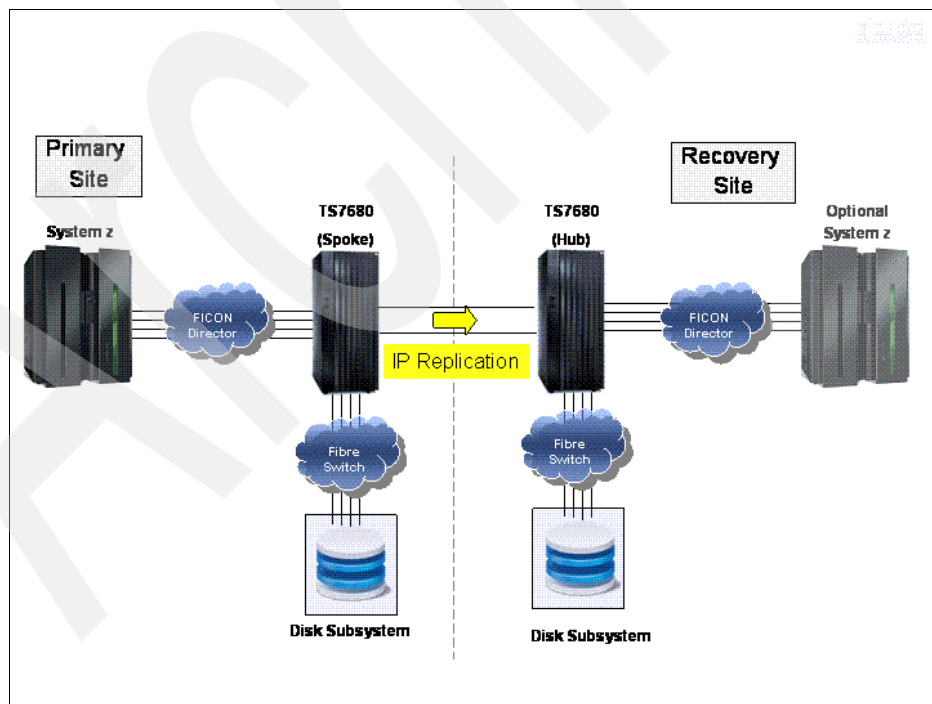


Figure 3-2 Typical deployment for two systems using replication

3.2.1 Extended Volume Attributes

The Enterprise Tape Controllers maintain EXVA (category and scratch information) in a local database for high performance processing. EXVA migration (writing the EXVA information from the ETC database to the repository) occurs when V1R2 code is installed. EXVAs are maintained in the repository for retention and replication. System z modifies EXVA with the PLF LSVC command. EXVAs are replicated from the spoke to the hub. The hub ETC database is updated with EXVA from the repository (when EXVAs are replicated).

Note: The ETCs are stateless; they do not know whether they are a spoke or hub.

EXVA data flow

Category information is set on a volume through PLF LSVC. EXVA data is stored in the local database of Enterprise Tape Controller; see Figure 3-3. The EXVA data is written to the repository through the ProtecTIER server.

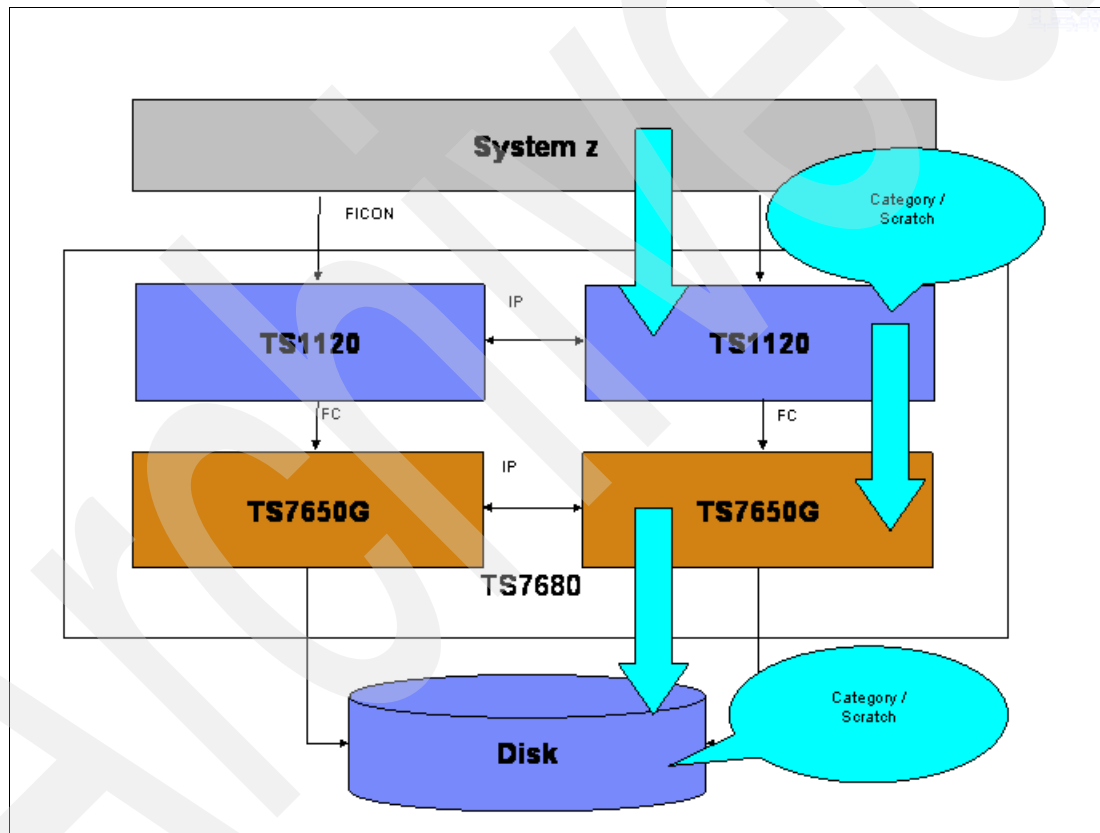


Figure 3-3 EXVA data flow

Figure 3-4 on page 39 shows EXVA category information replicated. EXVA data is written to the repository through the ProtecTIER server.

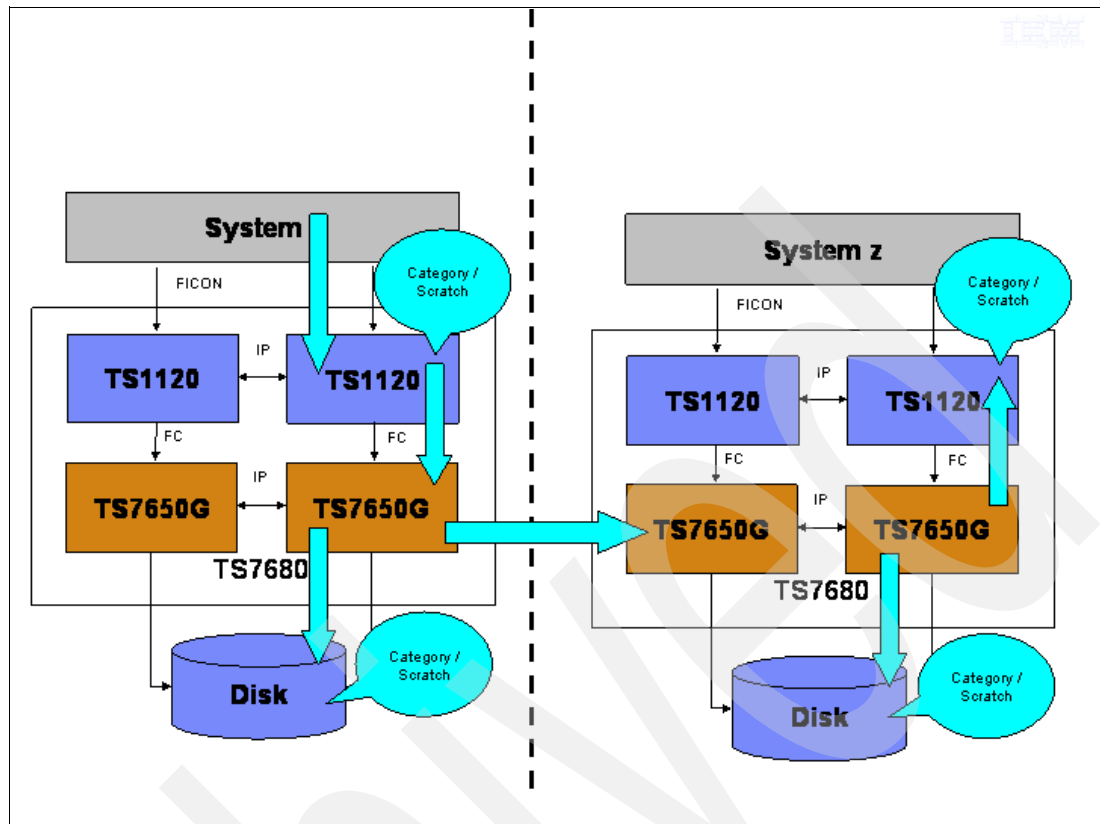


Figure 3-4 EXVA data flow continued

3.2.2 Replication features

ProtectTIER's replication is designed to provide you with great flexibility, to work seamlessly with any of the leading backup applications, and to fit easily within the overall tape operations paradigm. Replication is policy-based, allowing you to define several different policies for replicating cartridges from one system to another. The granularity of the replication policies enables you to set policies for an individual cartridge, a pool of cartridges, or an entire virtual tape library.

Replication is performed asynchronously at the logical cartridge level, and replication progress can be tracked and monitored at the cartridge level through the ProtectTIER management GUI. Full data validation is performed at the secondary site to ensure enterprise-class integrity of the transferred data prior to making the virtual cartridge available. Once replicated to the DR site, you may choose to clone these cartridges to real physical tape utilizing their backup application. In the event of a disaster, the DR site's TS7680 can become the production site until the main site comes back online. At that point the user may replicate or move the newly created tapes back to the main production site.

In the following sections we discuss a few of the key features and design points that demonstrate the flexibility and synergy with the tape operations paradigm.

Virtual tape management framework

As a target of the backup application, a ProtectTIER system presents itself as a tape library (or many libraries) to the network. The tape processing application manages the cartridges within a ProtectTIER system as though they were real cartridges, including read, write, import/export, tracking media with barcodes, and many other operations. Because replication

at the ProtecTIER level is transparent to the backup application, ProtecTIER's replication function is designed to allow synchronization with the backup application by way of normal tape management methodologies.

Virtual shelf

The ProtecTIER replication feature introduces the concept of a virtual shelf. As a general rule, replication always occurs from the source shelf to a destination shelf. As with physical tape shelves, there is a limit to the number of tapes that can be put on a virtual shelf. The limit is the result of subtracting the total number of tapes in all libraries from the total number of tapes supported in a single repository.

Visibility switch control

Visibility switch control is the means by which ProtecTIER can determine *where* cartridges actually exist, since from a backup application standpoint, any specific cartridge or barcode can *exist* in only one location at a given time. ProtecTIER native replication provides a virtual *shelf* (Figure 3-5) that is visible internally only at the ProtecTIER level (once exported by the backup application) and provides more flexibility in managing tapes and cartridges and where they are kept—similar to keeping physical tapes on an actual shelf outside of the tape library.

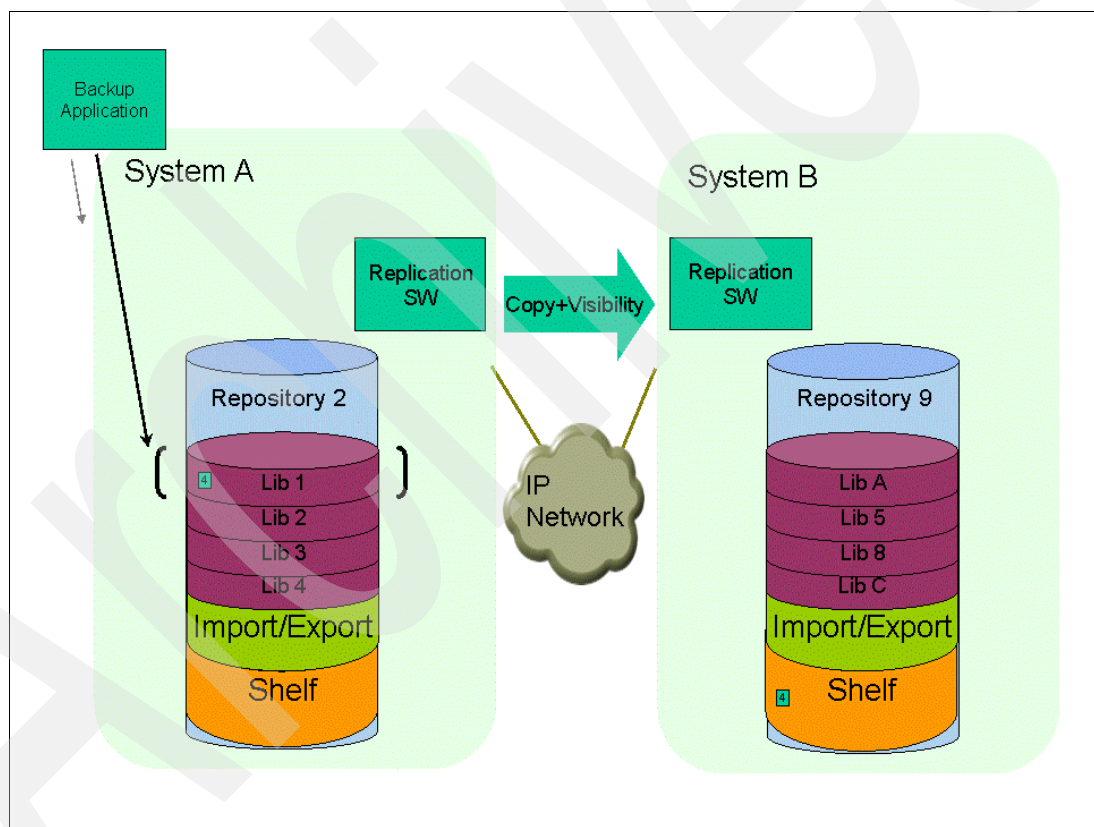


Figure 3-5 Virtual shelf

To ensure that a given cartridge is only visible to the backup application in one location despite the fact that a replica exists, ProtecTIER offers the *visibility control* function that allows you to determine in which location the cartridge should be accessible to the backup application. This is achieved by utilizing the import/export slots of the virtual libraries and exactly mimics the operation of physical tape management.

Cartridge cloning and visibility control

A common use case for ProtecTIER users is to first replicate data from the primary site to the secondary site, and then to move the data from the disk-based repository onto physical tape cartridges for long-term retention. With the visibility control function, ProtecTIER makes this operation simple for you. Once the cartridges complete replication to the secondary (DR) site, you may clone these cartridges to physical tape using their backup application tape copy function. This allows the backup application to remain in control of the end-to-end process and maintain its catalog of all cartridges, the associated data, and the location.

Policy management

Replication policies allow for flexibility in controlling replication and enable a high degree of automation. Policies consist of rules that dictate the transfer of data from one repository to another, based on events such as writing new data to an existing cartridge that belongs to a policy. By setting policies for ranges of barcodes, you may implement differing degrees of protection for different applications. You can also assign various priority levels to replication policies, which determine the order in which the data is transferred across the network.

Replication network management

ProtecTIER repositories belong to a replication grid, which is a framework for managing and monitoring the replication activities between ProtecTIER systems. Through ProtecTIER Manager, you can monitor the status of the overall replication network, the relationship between grid members, and the data throughput rate of replication. Further statistics about the cartridges involved in replication policies, as well as the statistics of each repository, are available in a single display from ProtecTIER Manager to enable ease of management and use.

Recovery management

When the need to fail over to the second site arises, whether due to a full disaster or a lower-level disruption, the ProtecTIER system enables rapid recovery of the data and restoration of the production applications such that business operations can continue with minimal downtime.

ProtecTIER is designed to enable rapid recovery of data from cartridges using the media server at the remote site. Once data has been restored, the ProtecTIER system can be brought online as the production system and be used for backup and recovery until the primary site has been restored. During the time period that the disaster recovery site acts as the primary production site, its data can be protected by replicating it to another secondary site. ProtecTIER Manager should be used to *pair* the DR site (now acting as the primary) to the other secondary site. (See more details in Chapter 10, “Disaster recovery and failover scenarios” on page 267.) When ready, you perform a failback operation to move production activity back to the primary site. At that time you may replicate any new data at the secondary site back to the primary and return the primary to its original status. All of these steps are enabled through the user interface.

3.2.3 Typical deployment

ProtecTIER native replication enables data replication capability across repositories, between ProtecTIER nodes connected to the same WAN. This WAN capability allows you to replicate data at any distance to remote cities, states, territories, and so on, for the purposes of disaster recovery. The deduplication software combined with replication reduces the amount of bandwidth required.

ProtectTIER native replication provides high availability in accessing backed-up data for backup and restore operations across sites. By having a current copy of critical data in a remote site along with a hot configuration of a TS7680 and your tape processing application ready to restore on a remote machine, data restoration can commence in a matter of minutes.

Configuring Replication

From a summary point of view to enable replication you will have to perform the following steps:

1. Capacity planning at DR site for local processing and DR test mode
2. Migration from prior code levels including Extended Volume Attribute (EXVA) migration process
3. Create a replication grid
4. Establish replication policies

3.2.4 ProtecTIER's native replication Management Interface

ProtecTIER native replication Management Interface provides an option-rich platform and is easy to use. It performs the following tasks:

- ▶ Defines policies and assigns priorities (Figure 3-6).
- ▶ Chooses which cartridges will be replicated.
- ▶ Schedules the replication window—the time frame in which replication takes place for all policies.
- ▶ Defines cartridge *visibility* features:
 - Determines *where* virtual cartridges *exist* from ProtecTIER or a backup application standpoint.
 - ProtecTIER native replication emulates moving tapes in and out of VTL's import/export slots.
 - Allows you to use the VTL export/import slots via the backup application to change the visibility of the cartridges from one library to another.




Figure 3-6 Replication policy

3.3 Normal operation concepts

This section describes a normal ProtecTIER operation sequence when replication is invoked.

3.3.1 Replication

When the tape processing application is writing to a cartridge that is part of a replication policy, ProtecTIER checks when it needs to be replicated and what its priority is so that the cartridge can be put in the correct order in the replication queue. Cartridges are always being replicated from the local library at the primary site to the virtual shelf of the repository at the secondary (DR) site. By definition, cartridges that are created at the primary (local) site repository are set to be read-write enabled, so that the backup application at the primary site has full control over them and their content. By the same token cartridges that were replicated to the secondary (remote) site are set to be in a read-only mode. By default, only one cartridge instance is located in a library. The replica is located on the virtual shelf.

Note: At any time, through ProtecTIER Manager, the user can override the default *location* of any given cartridge and manually move the replica from the virtual shelf to a library in the secondary site repository.

Cartridges are marked as sync'ed once the data finishes replicating from the primary to the secondary site, so that at the time of sync, the local cartridges and their remote replicas are exactly identical. Up to 128 cartridges can be replicated simultaneously. Before replication starts running, the system ensures that only unique, new data will be transferred over the wire. In order to achieve that, each side holds sync data per each of its cartridges. This sync data is used by the destination (secondary, remote site) to figure out which data (if any) should be replicated, as only new and unique *elements* of data are sent over the wire to the remote site. The replication mechanism has two types of data to transfer:

- ▶ Meta data, which is the data that describes the actual data and carries all the information about it
- ▶ User data, which is the actual backed-up data

Network failures, if and when they occur while the replication operation is being carried out, lead to retries of up to seven consecutive days to replicate any specific cartridge that did not finish its replication due to the line failure.

3.3.2 Replication data transfer

When the replication action is started either manually or based on a policy, the source (primary) ProtecTIER system carries out the following procedures:

- ▶ Initiates the sync-cartridge function between its own (source, primary site) repository and the destination (DR site) repository.
- ▶ Reads the unique replication data units upon requests from the remote ProtecTIER system based on what it is missing.
- ▶ Sends the unique data elements, using TCP protocol, over the WAN to the remote (DR) site.

At the same time the destination ProtecTIER system performs the following handshake actions in this order:

1. Calculates the relevant cartridges' sync point from which the replication should start.

2. Receives many data units concurrently as part of the replication action.
3. Verifies CRC for all replicated data before it becomes available as part of the cartridge.
4. Once the CRC check proves successful, the system moves each of the verified data elements into the cartridge scope and makes it available for the user.

Once all verified pieces are inserted into the cartridge, it becomes available as a complete cartridge. However, as a replica in the destination (DR) repository it is set as read only and cannot be used for tape processing purposes. This is an important factor in failover situations when the DR system may temporarily become the production site and can accept local tape processing requests. At that time you should create new tapes to accept this local backed-up data.

These new local DR site tapes can be replicated to the primary site once it becomes available and ready for production again. During the failback process, which is when you move operations back to the primary site, the newly created cartridges from the DR site can be replicated to the primary site. Under these circumstances the system grants read and write permissions to these replicated cartridges at the primary site, which becomes the *owner* of these tapes from that point on, just as though they were created there.

Note: The replication data transfer process requires that the replicated cartridge reside at the remote (DR) site ProtecTIER VTL shelf. As long as the data transfer is being performed, these cartridges cannot be moved from the shelf.

3.3.3 Visibility switch control

This is an automated process of the ProtecTIER replication feature that transfers the visibility of a cartridge from its master to its replica and vice versa. Just like a *real* physical tape, a virtual cartridge can only reside in or be visible to the tape processing application in one location at a time. This is carried out by the ProtecTIER replication by utilizing the VTL Import/Export slots. The system uses the export slots and eject operation as triggers to begin processing a tape move. The tape processing application can eject the cartridge into one of the export slots. As soon as that cartridge replication action is completed, the cartridge appears at one of the import slots of the secondary (remote, DR) site and can be imported into a library.

Replication visibility is an attribute defined in a replication policy and can be set by the user. A specifically defined library in the remote repository must be selected for the visibility transfer to be carried out. Once the cartridge is ejected at the primary (local) site, it moves to the local virtual shelf. Then it is replicated as part of a policy to the remote site virtual shelf. Once the replication is completed and verified, that replica moves to the respective library's import slot and can be imported into that library.

At this point that cartridge is only visible to the tape processing application at the remote site. A copy of the cartridge stays at the primary ProtecTIER system for fast recovery. However, it is hidden from the tape processing application. The way to move the cartridge back and make it visible to the tape processing application at the primary site is to eject it from the remote library and import it back into the local one (the same library it came from). Since the system keeps a hidden copy at the primary site, this move back is instantaneous.

Remote cloning

Remote cloning is the process of using a secondary site to clone cartridges. ProtecTIER replication enables you to offload tape cloning to your secondary site. Many users replicate

their data from the primary site to the secondary (DR) site, and then move it from the disk-based repository onto physical tape cartridges for long-term retention.

One of the advantages of performing this practice at the secondary site is to take the burden of cloning to physical tape from the production environment to the remote location, which is where tapes will most likely be kept once cloned. The remote cloning operation uses the cartridge replicas residing at the destination's ProtecTIER VTL shelf for cloning. The process imitates the commonly used physical process involving the transportation of physical cartridges from the primary site to a remote site either after being cloned or in order to clone them there for long-term archival purposes.

Hardware planning for the TS7680

In this chapter, we describe which options and features can be configured with the IBM System Storage TS7680 ProtecTIER Deduplication Gateway for System z.

We provide information about which configuration options are available and useful. We also discuss the installation tasks required and the responsibilities assigned to IBM System Service Representative (SSR), Field Technical Sales and Support (FTSS), Lab Based Services (LBS), and the client.

The following topics are covered:

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

4.1 General overview of the TS7680

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.

See Figure 4-1 for a logical layout of the IBM System Storage TS7680 ProtecTIER Deduplication Gateway for System z.




Figure 4-1 TS7680 logical layout

The TS7680 comes without disk storage, allowing clients to use IBM or non-IBM disk storage as back-end disk storage. See “Disk configurations for TS7680” on page 80 for the supported list of IBM and non-IBM disk storage. One of the advantages of a configuration with IBM disk storage is that the installed Reliability, Availability, and Serviceability (RAS) package is able to call home with information of a failing disk component.

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 above mentioned solutions have the same ProtecTIER server hardware built in.

Table 4-1 on page 48 shows the specifications of a TS7680 installation with consolidated values for the two ProtecTIER servers.

Table 4-1 TS7680 specifications

Component	TS7680
Number of processors	48

Component	TS7680
Number of memory	64 GB
Number of virtual libraries	1
Number of virtual drives	256
Number of virtual cartridges	1.000.000
Number of supported disk capacity	Up to 1 PB
IBM's path failover technology	Yes
Two-node cluster configuration	Yes
IP-based replication configuration	Yes
Flexible disk-based storage options	Yes
Sustained in-line throughput	Up to 500 MB/s and more depending on the back-end disk storage and workload
Data reduction	Up to 25:1 and more depending on the workload
Preinstalled disk storage	No
Servers come in a rack	Yes

4.2 Hardware and software components for the TS7680

The IBM System Storage ProtecTIER Enterprise Edition V2.4 software, and the Red Hat Enterprise Linux Server Release 5.2 Advanced Platform 64-bit, comes 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 server (model 7233).

To provide the System z attachment, two additional Enterprise Tape Controllers, comprised 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, a console must be provided by the client that can be directly attached to one of the ProtecTIER servers. We recommend ordering an IBM TS3000 System Console. With this 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 they are shipped together. The TS3000 System Console is a one-unit (1U) System x® server that allows an IBM System Service Representative (SSR) to perform maintenance and, if enabled by the client, the also called TS3000 System Console (TSCC) can remotely monitor the installation and automatically call home with any hardware errors.

A wide variety of disk-based storage can be attached to the TS7680. Check the System Storage Interoperation Center (see Notes box below) for specific environmental specifications and supported back-end disk configurations. Currently, the TS7680 can be connected to the following back-end disk arrays:

- ▶ IBM DS3400
- ▶ IBM DS4200
- ▶ IBM DS4300
- ▶ IBM DS4700
- ▶ IBM DS4800
- ▶ IBM DS5020
- ▶ IBM DS5100
- ▶ IBM DS5300
- ▶ IBM DS8100
- ▶ IBM DS8300
- ▶ IBM XIV®, GEN1, NEXTRA
- ▶ IBM XIV, GEN2
- ▶ IBM SVC
- ▶ EMC CX300
- ▶ EMC CX600
- ▶ EMC CX700
- ▶ HDS AMS1000
- ▶ HDS AMS2300
- ▶ HDS USP-V
- ▶ Hewlett Packard EVA8100
- ▶ Hewlett Packard EVA8400
- ▶ IBM N6040
- ▶ IBM N6060
- ▶ IBM N6070
- ▶ IBM N7600
- ▶ IBM N7700
- ▶ IBM N7800
- ▶ IBM N7900

Note: For a list of disk subsystems that are supported by the TS7680, refer to the *interoperability matrix*:

ftp://service.boulder.ibm.com/storage/tape/TS7600_iop_matrix.pdf

Referring to the TS7680G in the System Storage Interoperation Center (SSIC) provides a list of supported environments:

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

The supported disk arrays have the following characteristics:

- ▶ Support for the ProtecTIER server operating system with the correct update level.

- ▶ Dual active-active controllers for compatibility with the Linux Multipath software included in the ProtecTIER server operating system to allow path failover.
- ▶ Fibre Channel or SATA disk systems.
- ▶ Support for the back-end Fibre Channel (FC) Host Bus Adapter (HBA) brand, model, and firmware level installed on the Gateway servers. The back-end HBAs are used to direct or Storage Area Network (SAN) attach the ProtecTIER servers to the disk array; in case of SAN attachment, the disk array must also support the SAN fabric switches used.

We describe in detail the hardware components and the corresponding feature codes in the following section.

4.2.1 TS7680 ProtecTIER server characteristics

The built-in ProtecTIER servers of the TS7680 are based on a 4U rack-drawer IBM System x3850 M2; its IBM machine type is 7233. It supports the necessary performance required for enterprise-level data protection and hardware redundancy.

The TS7680 ProtecTIER servers are each shipped with the following characteristics:

- ▶ Four Intel® Six Core Xeon Processors x7350 (2.66 GHz) (24 cores total)
- ▶ 32 GB memory PC2-5300 CL5 ECC DDR2 SDRAM RDIMM (3644)
- ▶ Two 146 GB 2.5-inch hot swap Serial Attached SCSI (SAS) Hard Disk Drive (HDD) running with 10,000 revolutions per minute (rpm). The disks are protected with RAID 1 (mirrored disks) by an LSI SAS controller. This RAID array is used to store the operating system and the ProtecTIER deduplication software.
- ▶ Two dual-port QLogic FC 4 Gb HBAs (4x FC back-end ports connected to the storage subsystem)
- ▶ Two dual-port Emulex FC 4 Gb HBAs (4x FC front-end ports connected to the lower or upper Enterprise Tape Controller)
- ▶ Two dual-port Ethernet cards
- ▶ One Remote Supervisor Adapter II (RSA)
- ▶ Two 1440 W standard hot swap redundant power supplies, hot swap redundant fans, hot swap HDDs, and hot swap memory modules
- ▶ UltraSlim Enhanced Multi-Burner

The front view of one of the TS7680 ProtecTIER Deduplication Gateway servers is shown in Figure 4-2 on page 51.




Figure 4-2 TS7680 ProtecTIER Deduplication Gateway server front view

The rear view of the TS7680 ProtecTIER Deduplication Gateway servers is shown in Figure 4-3.




Figure 4-3 TS7680 ProtecTIER Deduplication Gateway server rear view

4.2.2 TS7680 Enterprise Tape Controller characteristics

The built-in Enterprise Tape Controllers, known as 3592-C06 Enterprise Tape Controller or IBM System Storage IBM TS1120 Tape Controller, are based on a 4U rack-drawer IBM System p5 520 (model 9131-52A).

The Enterprise Tape Controllers are each shipped with the following characteristics:

- ▶ Two POWER5+ 64-bit capable processors
- ▶ 2 GB memory
- ▶ Four mirrored internal disks running AIX 5.3
- ▶ Two dual port Fibre Channel HBAs connected to the ProtecTIER servers
- ▶ Four one-port FICON HBAs connected to the System z host
- ▶ Two onboard Ethernet ports

The front view of one of the Enterprise Controllers is shown in Figure 4-4 on page 52.




Figure 4-4 TS7680 Enterprise Tape Controller front view

The rear view of the Enterprise Controller is shown in Figure 4-5.




Figure 4-5 TS7680 Enterprise Tape Controller rear view

4.2.3 TS7680 feature codes

This section lists the required and optional feature codes for the TS7680.

TS3000 System Console feature codes

An IBM TS3000 System Console is required for installation, service, and maintenance of the TS7680. You may use an existing TS3000 System Console if it is within line of sight of the TS7680. If you need a TS3000, it can be ordered with the TS7680 ProtecTIER Deduplication Gateway for System z and they are shipped together.

- ▶ FC2714 - Console Expansion
Provides an Ethernet hub and cable for attaching to an existing TS3000.
- ▶ FC2722 - TS3000 System Console
Provides a rack mountable TSSC (1U server, keyboard, display, and mouse).
- ▶ FC2733 - Internal modem
Provides an internal modem installed in the TS3000 System Console.

Note: Previously ordered models of TS3000 System Consoles have feature codes of either FC2720, FC2721, or FC2730, while the current model of TS3000 System Console has feature code FC2722.

The TS3000 System Console centralizes maintenance and service terminals, and enhances remote support capability. It is the only method for electronic call home, hardware problem determination, and repair of the TS7680.

In Figure 4-6, we show the rack-optimized TS3000 System Console server.




Figure 4-6 TS3000 System Console server

Disk tracking feature codes

The following features designate whether the gateway is attached to a new IBM disk subsystem, to an existing IBM disk subsystem, or to an OEM disk subsystem. The type of disk subsystem is indicated for each feature code.

- ▶ FC9030 - Attached to DS8000®
- ▶ FC9031 - Attached to DS5000
- ▶ FC9032 - Attached to DS4000®
- ▶ FC9033 - Attached to XIV
- ▶ FC9038 - Attached to other IBM disk

If you do not intend to order a new IBM disk array, or if you are using new or existing non-IBM disk systems, you should order the following feature code:

- ▶ FC9039 - No new IBM disk

FICON attachment

All internal cables such as Ethernet and Fibre Channel are installed by manufacturing. The external connection between the TS7680 and the System z host can be single-mode long wave laser or multimode short wave laser. The following feature codes are available.

Lower Enterprise Controller

- ▶ FC3441 - FICON Shortwave
- ▶ FC3442 - FICON 4 km Longwave
- ▶ FC3443 - FICON 10 km Longwave

Upper Enterprise Controller

- ▶ FC3444 - FICON Shortwave
- ▶ FC3445 - FICON 4 km Longwave
- ▶ FC3446 - FICON 10 km Longwave

Note: For the most recent list of supported FICON directors, refer to the IBM Techdocs Library web page at:

<http://www-03.ibm.com/support/techdocs/atsmastr.nsf/webindex/FQ115356>

4.2.4 TS7680 software

This section gives an overview of the IBM System Storage TS7680 ProtecTIER Deduplication Gateway for System z Version 1 Release 1 software.

ProtecTIER Enterprise Edition V2.4 Base Software

The ProtecTIER Enterprise Edition V2.4 Base Software is the software that provides the base functionality for ProtecTIER Enterprise Edition. The software supports any capacity of FC or SATA disk drives in its disk storage pool. ProtecTIER licensing is calculated on the used capacity of the attached disk storage. The capacity ordering is by cumulative tier (from tier 1 to tier 7) and is expressed in terabytes. Starting with tier 1 (1-12 TB), you have to order the quantity from each successive tier required to reach the total TB managed by the server, filling each tier in succession. Contact your IBM Sales person for calculating the licensing for your installed disk storage.

Note: ProtecTIER licensed internal code can be client-installed.

ProtecTIER Manager V2.4 console software

The TS7680 must have a console workstation for installation and maintenance. This workstation runs the ProtecTIER (PT) Manager application. PT Manager is a Graphical User Interface (GUI) that is used to install, configure, and monitor ProtecTIER. You install the ProtecTIER Manager on a workstation running either Windows® or Linux. The ProtecTIER Manager workstation must be connected to the TS7680 ProtecTIER servers through your network.

The console must be capable of operating one of the following operating systems:

- ▶ Windows 2003
- ▶ Windows XP
- ▶ Red Hat Enterprise 4 or higher

The console must also have:

- ▶ At least 100 MB of available disk space
- ▶ At least 256 MB of RAM
- ▶ Access to ProtecTIER service node's IP address
- ▶ Keyboard, mouse, and CD-ROM drive

In addition, we recommend that the monitor for ProtecTIER Manager software be configured to the following settings:

- ▶ Resolution of 1024 x 768 pixels or higher
- ▶ 24-bit color or higher

Note: The console must have access to the ProtecTIER nodes through the IP addresses (port 3501 must be open in the firewall).

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

Enterprise Tape Controller Licensed Internal Code

The Enterprise Tape Controller Licensed Internal Code (LIC) is a software bundle based on the operating system AIX 5.3.

This software bundle is responsible for translating any Perform Library Command (PLF) functions arriving via FICON protocol from the System z host to Small Computer System Interface (SCSI) commands sent out via Fibre Channel connections to the ProtecTIER servers.

In addition, this software maintains an inventory of the ProtecTIER-provided virtual tape library, drives and cartridges information in a redundant manner and provides this data to the System z host.

The Enterprise Tape Controllers can be seen as an interface between the System z host and the TS7680 Deduplication ProtecTIER servers, and are completely transparent to the System z host.

4.2.5 TS7680 rack setup

The IBM System Storage TS7680 ProtecTIER Deduplication Gateway for System z is shipped from manufacturing with all internal cables connected. You have to provide electrical connections (preferably on different circuits) and Ethernet connections to the ProtecTIER servers and the TS3000 System Console. This section describes the internal cable layout of power, Ethernet, Fibre Channel, FICON, and KVM video connections. In Figure 4-7 on page 56 you can see that you have to add 14 U empty space between the upper ProtecTIER server and the Ethernet Switch 2 to get the physical view of the rack.

Power connections layout

Figure 4-7 shows the two external power lines for the internal Power Control Assemblies (PCA). In addition, the power layout of all components as well as the power connections from the ProtecTIER servers to the WTI Network Power Switch are shown.

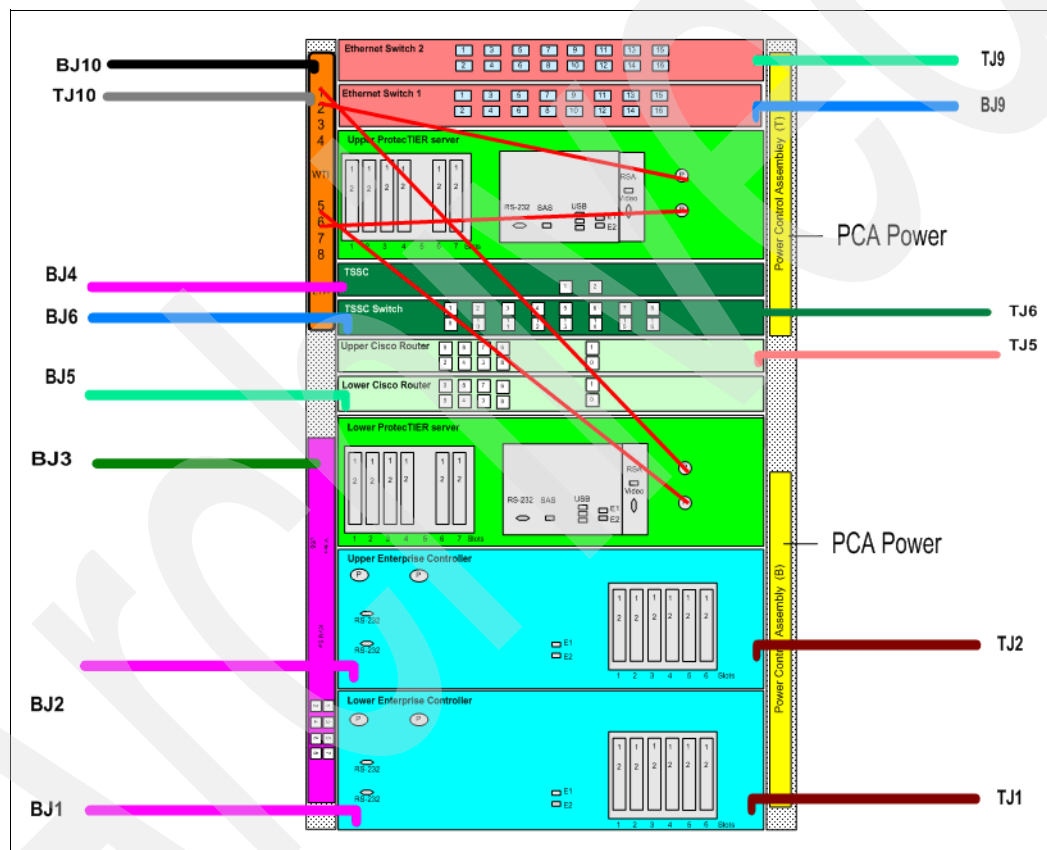


Figure 4-7 TS7680 frame power layout

Ethernet connections layout

All necessary Ethernet connections can be seen in Figure 4-8 on page 57.

The blue lines show the Ethernet connections for the lower and upper ProtecTIER server. The Ethernet ports on card 3 / port 2 and card 4 / port 1 on each ProtecTIER server are bonded for high availability and load balancing reasons. The WTI Network Power Switch is connected to both ProtecTIER servers through port 3 of the Ethernet Switch 2.

The black lines show how the two Enterprise Tape Controllers are attached to the internal network via the two Cisco Routers. Each Enterprise Tape Controller has two Ethernet connections bonded for high availability and load balancing.

The pink lines show the external interface to the TS7680. You must provide, therefore, three IP addresses for the TS7680 installation (two for the ProtecTIER servers and one for the TS3000 System Console).

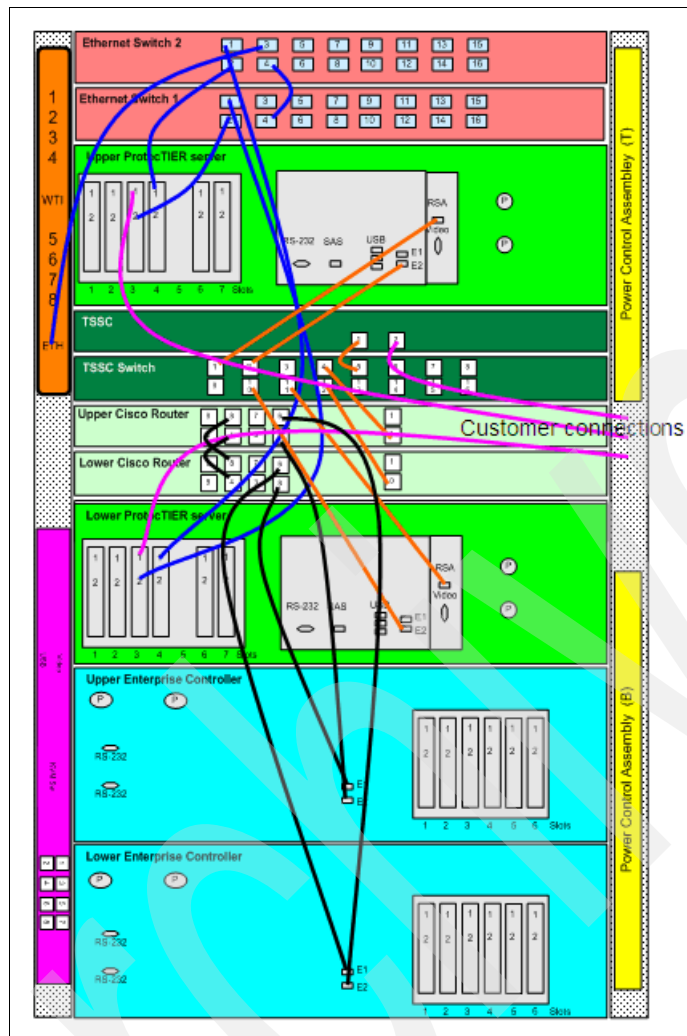


Figure 4-8 TS7680 Ethernet layout

Fibre Channel and FICON connections layout

Figure 4-9 on page 58 shows all Fibre Channel and FICON connections of the TS7680.

The orange lines are the Fibre Connections from lower and upper ProtecTIER servers to lower and upper Enterprise Controllers. The cards in slots 6 and 7 of the ProtecTIER servers are used to connect to the back-end disk storage.

The cards in slots 2, 3, 5, and 6 of the Enterprise Controllers are the FICON connections to the host (yellow lines).

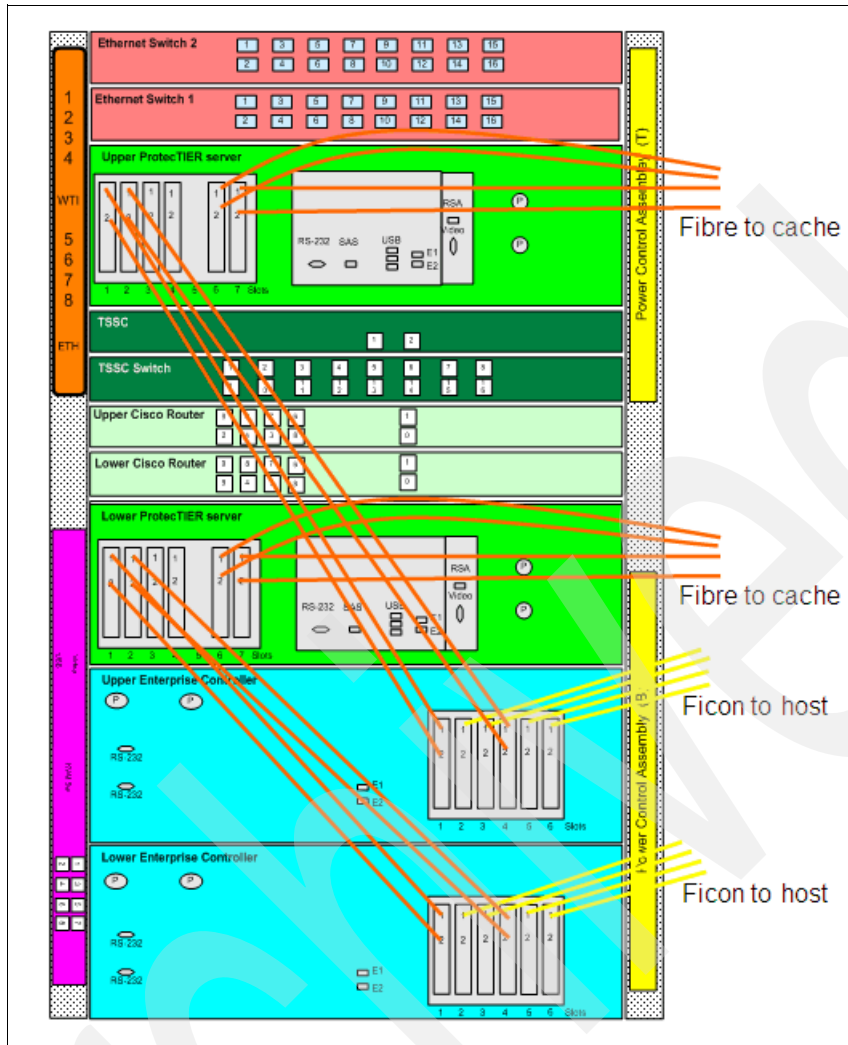


Figure 4-9 TS7680 Fibre Channel and FICON layout

KVM video cables layout

Figure 4-10 on page 59 shows the KVM video cable connections that are installed in manufacturing to provide a user interface through the TS3000 System Console to the client. The internal KVM switch connects the TS3000 System Console and the two ProtectTIER servers via VGA and USB lines.




Figure 4-10 TS7680 KVM video and USB layout

TS3000 System Console

The recommended TS3000 System Console is a one-unit high IBM System x server with a one-unit high keyboard, video, and mouse (KVM). This TS3000 System Console is attached to a network switch (see TSSC Switch in Figure 4-8 on page 57) providing the network connections to the lower and upper ProtecTIER servers. In addition the Remote Supervisor Adapter (RSA) II port for system management of the ProtecTIER servers is connected to the TSSC Switch for example to send SNMP alerts to the TS3000 System Console. Another port of the TS3000 System Console provides connection for the client via an external IP address.

4.3 Installation planning

In this section, we present details to help you plan for the installation and basic setup of the TS7680. Planning is primarily your responsibility.

4.3.1 Installation worksheets

We provide implementation worksheets (see Appendix B, “Checklists” on page 297) for you to

obtain the required information to implement a TS7680. Most of the worksheets need to be completed prior to the installation of TS7680, because the values inserted are needed to complete the installation.

4.3.2 Prerequisites

In this section, we describe the prerequisites necessary to install ProtecTIER, focusing on:

- ▶ Hardware prerequisites
- ▶ Software prerequisites

You must use the ProtecTIER Manager software to configure, manage, and monitor the operation of the TS7680. 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 hardware and operating environment requirements for the ProtecTIER Manager workstation are:

- ▶ x86 (Pentium® or higher) microprocessor
- ▶ 256 MB memory
- ▶ 100 MB of disk space
- ▶ Access to the ProtecTIER service node's IP address (port 3501 is open on the firewall)
- ▶ Keyboard, mouse, and CD-ROM drive
- ▶ Screen resolution of 1024 x 768 or higher
- ▶ 24-bit color or higher
- ▶ Operating environments supported:
 - Windows 2000
 - Windows Server 2003
 - Windows Server 2008
 - Windows XP
 - Red Hat Enterprise 3 or higher

Before installing the ProtecTIER servers, verify that the following prerequisites are met:

- ▶ All cabling and physical connections between hardware components must be done prior to installation. This includes the local LAN connections with an assigned IP address.
- ▶ A workstation connected to the client network.
- ▶ RAID groups previously created on the disk subsystem.
- ▶ The network power switch is configured.

4.3.3 Installation tasks

Installation involves the IBM System Services Representative (SSR), Field Technical Sales and Support (FTSS), Lab Based Services (LBS), and client personnel.

Client installation responsibilities

You are responsible for preparing the installation site prior to the installation of the TS7680. All physical planning for the TS7680 is a client responsibility. In summary, the client is responsible for:

- ▶ Completing the planning and preparation tasks described in the TS7680 Customer Information Center.
- ▶ Meeting the preinstallation requirements outlined in the TS7680 Customer Information Center.
- ▶ Hosting a meeting between the FTSS install team and the back-end disk storage install team to discuss the ProtecTIER configuration requirements.
- ▶ Completing the company information and IP address worksheets provided in the TS7680 Customer Information Center.
- ▶ If feature code 2722 (TS3000 System Console) is not installed in the TS7680 rack, confirming that the existing TS3000 System Console being used with the TS7680 is one of the approved models and has feature code 2719 applied.
- ▶ It is recommended that the client has one or more PCs designated to run the ProtecTIER Manager software.
- ▶ Providing Ethernet (Cat5e or higher), Fibre Channel and FICON cables for the external cable connections.
- ▶ Providing cooling, telephone service, safety, and security.
- ▶ Client-supplied frames, disk arrays, and expansions must be fully installed and operational before the installation of the TS7680 can begin.

IBM Service installation responsibilities

The IBM System Services Representative (SSR), Field Technical Sales and Support (FTSS), or IBM Lab Based Services (LBS) can install and configure the TS7680 and complete the following tasks.

SSR tasks

- ▶ Responsible for the physical installation of the system, including external cable connections.
- ▶ Install and secure the TS7680 rack in the designated location at the client site.
- ▶ Label and connect external power, Ethernet, and Fibre Channel cables, as necessary.
 - The Fibre Channel cables come from the ProtecTIER servers and go to the client-supplied disk system.
 - Label and connect the Ethernet cable to the ProtecTIER servers and to the external TS3000 System Console.
 - Label and connect the FICON cables from the Enterprise Controllers to the client's host.
- ▶ Power up the system.
- ▶ Verify the accuracy of hardware installation and cabling.
- ▶ Perform a visual check of fault indicator LEDs.
- ▶ Configure the TS3000 System Console for use with the TS7680 and as an NTP server.
- ▶ Verify that each subsystem is in time sync with the TS3000 System Console.
- ▶ Configure the Reliability and Serviceability package on each ProtecTIER server.
- ▶ Perform RAS verification tasks.

- ▶ Configure the TS7680 frame cluster definition on the TS3000 System Console.
- ▶ Perform a Call Home test on each ProtecTIER server.
- ▶ Configure Enterprise Controllers for use with the TS7680.
- ▶ Verify that all System z connections are up.
- ▶ Run a Call Home test on each Enterprise Controller.
- ▶ Verify that each Call Home sent reaches RETAIN®.

FTSS and LBS tasks

- ▶ Oversee project management for the installation and the integration of the engagement.
- ▶ Oversee change management and process control for the installation.
 - Tasks of the client
 - FTSS and/or LBS representative
 - System Service Representatives (SSRs)
- ▶ Coordinate and schedule IBM resources for client installations, and act as a focal point of contact for coordination of installation services.
- ▶ Schedule and facilitate planning and solution assurance conference calls.
- ▶ Create and document the installation service process.
- ▶ Preinstall check for Technical Delivery Assessment (TDA)
 - An important TDA check was added that requires a preinstallation conference call between the FTSS team and the disk systems installation team for discussing and finalizing of the disk configuration and results of the sizing tool to come to an agreement on the layout.
 - For every installation the disk installer needs to be contacted and involved.
 - The FTSS needs to identify the disk installer right away. Get contact information.
 - Validate that the disk was actually set up the way it was discussed.
- ▶ Configure the ProtecTIER software on the ProtecTIER servers.
- ▶ Install ProtecTIER Manager on the ProtecTIER Manager workstation and register each ProtecTIER server as a new node.
- ▶ Validate the cluster configuration.
- ▶ Log on to each Enterprise Controller and perform the Initial 3958 Configuration menu on the Enterprise Controllers.
- ▶ Initiate System Checkout.
- ▶ Verify that System Checkout runs with no errors.
- ▶ Verify that the host system is attached.
- ▶ Release the system to the client.
- ▶ Document and report installation results.



Part 2

Planning and setup

In this part we discuss planning and setup topics.

Archived

Planning for deduplication

This chapter gives an overview of planning, sizing, and configuring the IBM TS7680 with consideration of your requirements in terms of architecture, scalability, and performance.

5.1 Capacity planning for the TS7680 - key concepts

The physical capacity required is primarily a function of:

- ▶ The size and frequency of the client workloads
- ▶ Retention period
- ▶ Data change rate*

The deduplication ratio is a function of both the data change rate and retention period. Longer retention periods yield a higher factoring ratio.

Workload analysis is the key:

- ▶ Mainframe tape workloads typically include both tape backup and recovery and production operations (such as logging and journaling from DB2 or IMS™).
- ▶ Backup operations inherently include a lot of data redundancy and therefore produce predictable deduplication rates.

Production workloads include far less data redundancy, and therefore do not benefit as much from deduplication. Figure 5-1 on page 66 gives an example of deduplication scenarios.

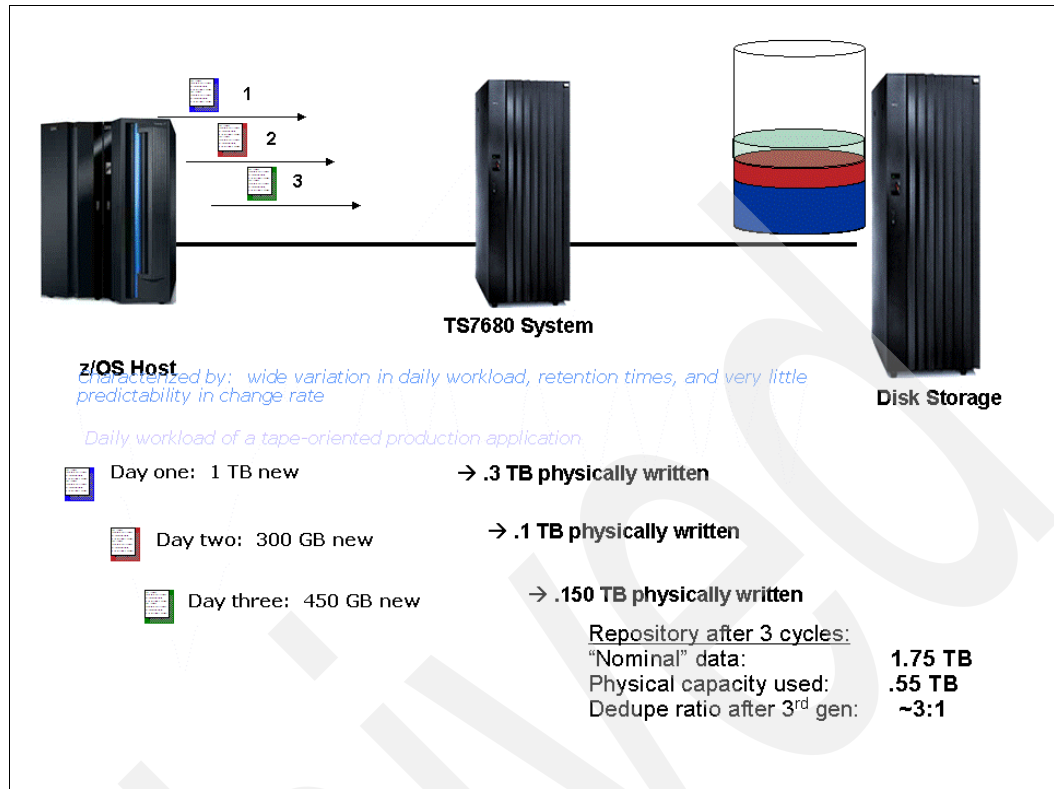


Figure 5-1 Production workloads example

- When production and backup data are mixed in the 7680 repository, the analysis should assess the deduplication expectations separately.

5.2 HyperFactor overview

HyperFactor is the core technology of ProtecTIER and consists of a series of algorithms that factor, or deduplicate, data efficiently. In each new data store operation, HyperFactor finds the data in common with previous data store operations and this common data in the new operation is effectively "filtered out." Pointers are used to reference existing data in the repository (see Chapter 1, "Concepts of data deduplication" on page 3 for more details about the deduplication process). In this way, the only new data is required to be stored, is a small fraction of the entire amount of new data. Figure 5-2 on page 67 shows the basic data handling of ProtecTIER.

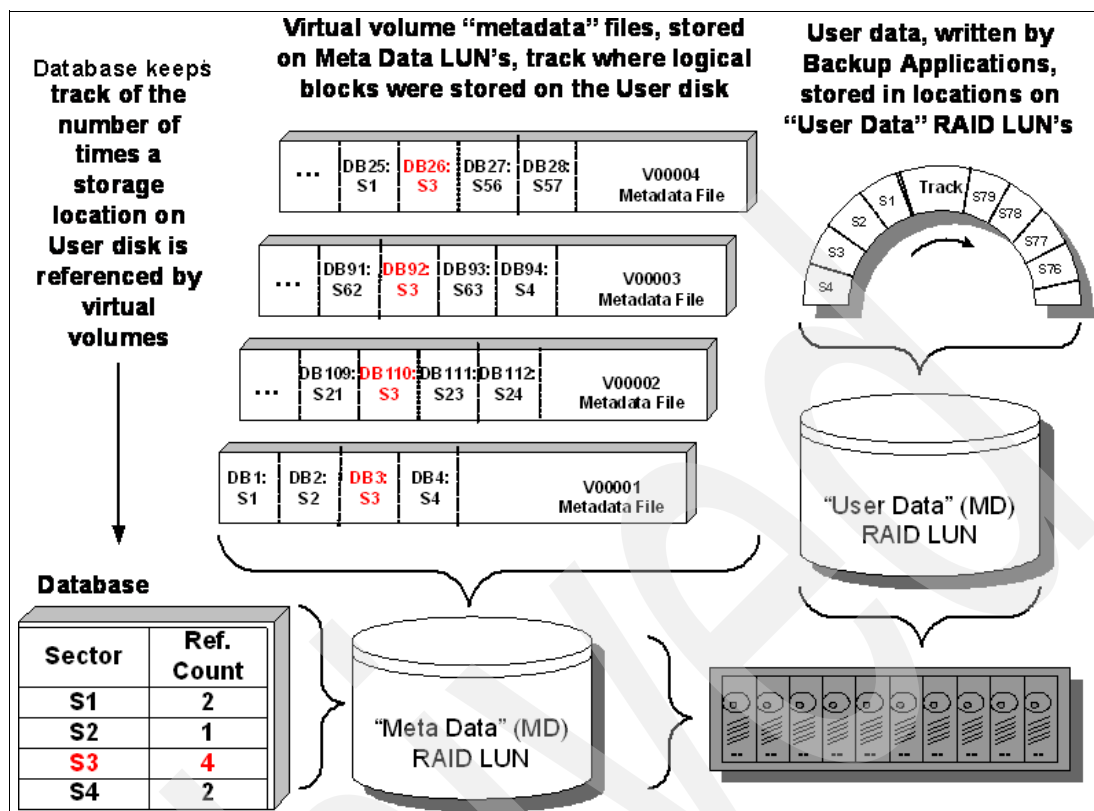


Figure 5-2 Data handling of ProtecTIER software

The capacity of the ProtecTIER repository consists of the factored tape processing streams as well as the Meta Data that describes the factored tape processing 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 tape processing streams) to the physical storage used (including all User Data, Meta Data, and spare capacity, that is, the total amount of disk storage that the user pays for).

In order to figure out the factoring ratio, we recommend that you have a detailed picture of your tape processing environment. In this chapter, we help you identify your goals and objectives, consider all the variables, and suggest some useful best practices. This chapter has four sections:

- *Sizing inputs* describes a deep analysis of your requirements, workloads, tape processing application, and data center topology.
- *Capacity sizing* gives you an overview of the process that is used to estimate your physical storage requirements for your current environment and your scalability needs.
- *Performance sizing* helps you understand how many Meta Data and User Data file systems are required, based on disk technologies and RAID configurations, to ensure proper performance in terms of tape processing throughput.
- *Storage sizing* describes in detail the performance and capacity tuning of your cache controller to match your actual amount of data and your future requirements. This tuning is closely related to performance sizing.

The methods described in this chapter give you the basis for assessing and understanding how the TS7680 can be fully integrated in your environment. In Figure 5-3, you can see the flowchart of the sizing process. Although each step is described in the following sections, IBM technical personnel and or Business Partner personnel will perform the sizing process before the presales and preinstallation phases.

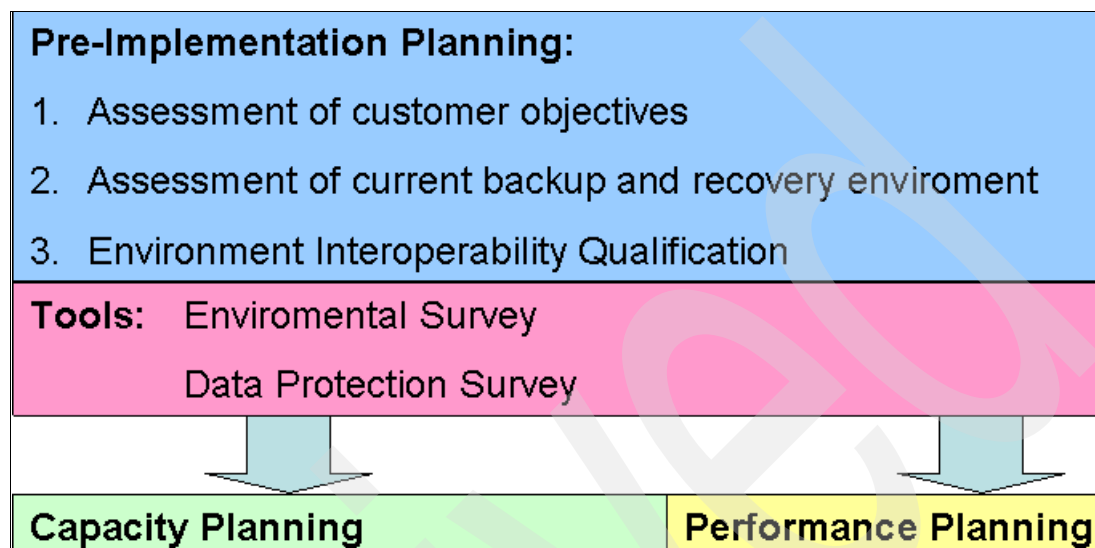


Figure 5-3 Flowchart for sizing a ProtectTIER solution

You should understand how the product calculates the deduplication ratio and presents and displays the data. ProtectTIER 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 ProtectTIER Manager GUI and through other ProtectTIER utilities.

5.2.1 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 ProtectTIER system sizing depends on many factors directly related to tape processing policies.

Understanding the requirements

A key point to remember when you want to evaluate a virtual tape library (VTL) with deduplication technology is making sure you have an understanding of your requirements and how you will be using the technology. This analysis starts with the environment and must answer the following general questions about your requirements:

- ▶ Why do you want to introduce a VTL solution with data deduplication in your environment? For example, are you going to use the solution to support disk-to-disk-to-tape (D2D2T) backups, disk-to-disk (D2D), or for archiving?
- ▶ What is the impact of a deduplication solution on your current tape processing environment?
- ▶ Are you struggling with data proliferation?
- ▶ Are your applications good candidates for data deduplication?
- ▶ How many mainframe hosts do you need to attach and with how much storage capacity?
- ▶ Where are the potential bottlenecks in your current tape processing environment? How can a deduplication solution fix them?

- ▶ What are your current and projected performance requirements? Performance *should* include today's tape processing requirements and peak throughput needs, and should also factor in the data growth and the implications of future performance needs.
- ▶ What is your estimated annual data growth rate?
- ▶ What are your expected capacity savings with a deduplication solution?
- ▶ What are the possible changes that you plan to make in your current tape processing architecture?

Once the requirement needs 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, and actual testing of the system in a live environment.

Understanding the existing environment

This section describes the information necessary to understand your current tape processing infrastructure. An environment questionnaire is used to collect this information and assess a preimplementation plan. Once 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:

- ▶ General questions:
 - How often is defrag being done to the disks?
 - Does anything happen to the data between ingestion and backup cycles that changes its nature on disk (reorgs, defrags, and so on)?
 - Does anything imbed or change any type of Meta Data between ingestion and backup cycles?
- ▶ Customer environment

Used to gather the characteristics of the tape processing application, the tape processing tiers, and the tape processing architecture. Since the TS7680 has an ISV compatibility matrix that specifies the supported tape application software and version, it is fundamental to verify that your current tape application software and version matches the ProtecTIER system requirements. If not, an infrastructure and economy impact analysis is required.
- ▶ Existing tape processing infrastructure

Used to examine how you back up your data, the current amount of data, the estimated annual data growth rate, and the type of technology used.

For general tape processing:

 - What data does this workload back up? Does it include VSAM files, DB2 files, IMS files, or CICS® files?
 - Does the general tape processing application use its own software compression?
- ▶ Disaster recovery and high availability

Used to determine whether a ProtecTIER solution can be integrated into your current disaster recovery environment.
- ▶ VSAM files
 - Which applications are used to back up VSAM files (IDCAMS, CA-Faver, and so on)?
 - Does this application use any software compression?
 - What is the frequency of the re-org being done to the VSAM file?

- ▶ DB2 files
 - Which applications are used to back up DB2 files (image copies, BMC Copyplus, and so on)?
 - Does this application use any software compression?
 - What is the policy for backing up the DB2? Is it weekly full and daily incremental, or daily full?
 - Are the DB2 files compressed on the disk?
 - What is the frequency of the re-org being done to the DB2 file?

Environment configuration survey

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

- ▶ Total amount of tape processing jobs per night and at peak (whether you do more tape processing during the weekends)
- ▶ Your tape processing window
- ▶ The length of time that you would like to retain data
- ▶ The profile of applications that are being backed up
- ▶ Other unique elements of requirements

Software and hardware inventory for compatibility check

Verify that all elements of your operating environment that will interact with TS7680 are qualified by IBM or the Business Partner to work effectively with TS7680. You can check the following for a complete list of supported systems:

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

System z host qualification

For each System z host or logical partition (LPAR) that will connect to the TS7680, provide the data requested in Table 5-1.

Table 5-1 Characteristics of the deployed System z hosts

Item	LPAR1	LPAR2	LPAR3
System z host operating system version			
Backup software and version			
FICON HBA models			
FICON HBA firmware version			
FICON HBA driver version			

Front-end fabric connectivity

For each FICON switch or director that will be connected to the front-end ports (System z host facing) of TS7680, provide the information requested in Table 5-2.

Table 5-2 Characteristics of the FICON attachment

characteristic	Director 1	Director 2
Director model		
Director release		

Storage

For each storage array connected to the TS7680, provide the information requested in Table 5-3 after an accurate assessment of the disk sizing.

Table 5-3 Characteristics of the disk arrays

Item	Disk Array 1	Disk Array 2
Disk array type and model		
Disk capacity on implementation		
Number of hard disk drives (HDDs)		
Size of HDDs		
HDDs revolutions per minute (RPM)		
Number of controllers		
Controller cache size		
Connection between TS7680 and disk array is loop or switch topology		

Data protection survey

Accurate capacity planning considers the behavior of each data type. A data type can be a file system, an operating system, databases, and so on. The size of one full tape processing operation 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 that provides information about the number of versions, frequency, and retention for each tape processing operation. It is assumed that the retention of a weekly tape processing operation is associated with the retention of its incremental tape processing operations.

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 tape processing to physical tape.
- ▶ How much capacity is used for the current “daily” tape processing to physical tape including differentials, incrementals, and cumulative tape processing.
- ▶ The rate at which the data received from the tape processing application changes from tape processing to tape processing operation. This measurement has most relevance when “like” tape processing policies are compared. (Data change rates may range from 1% to >25%, but are difficult to observe directly.)
- ▶ How often full tape processing is performed.
- ▶ How many cycles of full tape processing operations are kept.

- ▶ The relationship of how many daily tape processing tasks are performed in between full tape processing tasks including differential, incremental, and cumulative tape processing operations.
- ▶ How many cycles of full and incremental, differential, or cumulative tape processing tasks are kept.
- ▶ Whether a monthly full tape processing operation is kept for longer periods than the regular weekly full tape processing operation.

Key inputs for performance planning

In order to plan for a successful TS7680 implementation, the following need to be considered:

- ▶ Desired throughput performance level (MB/s) to achieve tape processing objectives
- ▶ Deduplication ratio (estimated)
- ▶ Repository size and estimated growth (driven by retention periods)
- ▶ Disk Technology (FC/SATA/SAS)
 - Capacity
 - RAID grouping for both Meta Data and User Data areas of the repository

Throughput considerations

The TS7680 is a virtual tape library with enterprise scale *in-band factoring*, which means all data reduction occurs in real time, as the tape processing tasks are running (this is in contrast to a “post process,” in which data is first written to disk and then factored at a later point. The in-line factoring approach has many advantages, but also requires the appropriate level of hardware and proper configurations to achieve optimal performance. Properly configured, a TS7680 is capable of achieving sustained throughput rates of up to 500 MB/s, or more, in live production environments. 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 (MB/s), when testing and deploying the TS7680.

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

- ▶ SAN/FICON connectivity
- ▶ Disk array
- ▶ Data type (also called tape processing policy)

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

SAN/FICON connectivity

For the best SAN Back End disk connectivity and FICON Front End host connectivity, do the following:

- ▶ Make sure that used fabric switches and directors for SAN and FICON are up to the latest supported firmware revision of their operating system.
- ▶ Use 4 Gbps HBAs for the TS7680G Back End connections.

Disk array

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

HyperFactored data. Two types of file systems make up the ProtecTIER repository (see Figure 5-4):

- ▶ Meta Data
- ▶ User Data




Figure 5-4 The ProtecTIER repository

Meta Data 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 Meta Data file system be optimal. In general, we recommend RAID 10 RAID groups (4+4, 6+6, or 8+8 disks) for the Meta Data file systems.

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. 80-90% of I/O in a typical ProtecTIER system environment is random reads; therefore, any storage array deployed with ProtecTIER should be optimized for this I/O pattern. In all cases, the disk array manufacturer (or reseller) should be consulted to determine the best tuning and configuration parameters for the particular array being deployed. The User Data LUNs should be tuned for a random-read intensive workload, while the Meta Data LUNs should be tuned for a random-write intensive workload. In 5.2.4, “Storage sizing” on page 80, 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 tape processing. Other data, like video or seismic data, cannot be compressed or factored very well. Also, the various tape processing 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 as well as system performance. This section also includes some general system testing and benchmarking considerations.

Consider the following:

- ▶ If multiple tape processing streams of the same data set are used for testing, we recommend that a single copy of the data be first to populate the repository.
- ▶ If encryption features of the tape processing application are turned on, the factoring ratio and performance of these data sets will degrade drastically.
- ▶ Backup application or database tape processing programs should disable compression. Data should be sent to ProtecTIER uncompressed, or the factoring ratio for this data will be very low.
- ▶ Multiplexing features of the tape processing application or database tape processing tool should be disabled. The tape processing application should send only one stream to each virtual tape drive. Since these ProtecTIER systems allow up to 256 virtual tape drives for the TS7680, the system can process many streams. This will not affect database tape processing speed (it may improve).
- ▶ In evaluation or test scenarios, do not create big data sets that repeat the same data. In addition, try to avoid copying the same data set over and over to create multistream tests (for example, using the same 100 MB to create a 40 GB data set and copy this data set several times to allow multistream tests). Although this will increase the factoring ratio, it will reduce performance, because all data is referenced to the same blocks.
- ▶ You can check the storage array performance capability from one of the ProtecTIER nodes by using a **dd** command to read from disk to /dev/null on UNIX®, for example:

```
dd if=/vol1/test.txt of=/dev/null
```

- ▶ **vmstat** and **iostat** are two system commands that can really help when checking the performance of the ProtecTIER systems.
 - If **vmstat** shows that the processor is waiting, this means the disk is not fast enough to service the I/O queues.
 - If **vmstat** shows that the processor system and user are reaching 90-95% capacity, then we might be on the edge of the processor's capability.
 - In optimal conditions, **vmstat** should show the following processor utilization parameters: processor idle 0, processor system and user are high, but not exceeding 85%, and processor wait is 0.
 - **vmstat** can show you if you are in the READ or WRITE state; if mostly READ, that means you are already in the factoring process, and if mostly WRITE, then you are writing your new data or baseline.
 - **iostat** shows you the activity on each LUN (for example, /dev/sdX). If while doing the tests not all of your /dev/sd devices are active in READ, then probably not all of the RAID and disk spindles are participating; see Example 5-1. You have optimized your system if, during tape processing, all of your RAID groups are active.

Example 5-1 iostat

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

sdh	0.00	0.00	0.00	1373	0
sdi	0.00	0.00	0.00	1373	0
sdj	0.36	6.20	3.80	4293030	2630829

- ▶ The first tape processing baseline is very important. If, during the tape processing, your entire RAID is not participating in the I/O write, it means that the sample data sets you are using might not be adequate and you are already in READ.
- ▶ Small data sets do not factor as well as larger data sets. For best system performance, in test or production, *at least* 24 data streams of tape processing data sets should be run at the same time. This takes advantage of the fact that the ProtecTIER system 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.

Tip: In general, the more virtual tape drives defined, the better. ProtecTIER and the TS7680 performance is optimized for a large number of virtual drives.

Studying the environment

Batch Magic is an application tool for tape and batch window tuning and capacity planning available to IBM'ers and IBM Business Partners. It enables the user to understand all aspects of the z/OS tape workloads as a key input to making wise hardware investments. It uses SMF and TMC data from the z/OS operating system as inputs to generate reports on the tape usage by workload in the environment.

A Batch Magic study assesses the Tape Management catalog to determine the total capacity of the live volumes required for all the relevant workloads. The primary output of this phase of the 7680 sizing process is a report that shows the breakdown of the workloads and associated capacities

The tool assesses the Tape Management catalog to assess the current capacity associated with the relevant programs, and assesses the compression rates from the SMF data to determine how much capacity (TB) is needed to store all the live data within the system, as well as the required throughput of the system.

5.2.2 Local repository sizing

In this section, we document the general importance of capacity sizing, but this process requires a presales engagement with IBM technical personnel. They use IBM internal tools to correctly size the IBM System Storage TS7680 with ProtecTIER through a comprehensive discussion with you about your data protection environment, requirements, and business objectives. In the performance of the repository sizing they will keep in mind the maximum throughputs the configurations are capable of.

The correct sizing and management of the ProtecTIER repository is important, because it:

- ▶ Enables you to purchase the correct amount of disk.
- ▶ Keeps tape processing operations running smoothly.
- ▶ Enables full utilization of repository.
- ▶ Enables host to monitor and predict repository usage.
- ▶ Enables host to react to warning conditions and shift work load prior to critical resource limits.

- Enables the user to react when critical resource limits are reached by cancelling lower priority jobs to allow higher priority jobs to complete successfully.
- Prevents write job failures due to a full repository.

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

Table 5-4 gives an explanation of the often used terms such as Data Change rates, the Factoring Ratio, and so on.

Table 5-4 ProtecTIER terms definitions

Nominal Capacity	That amount of user data that the ProtecTIER system is protecting
Physical Capacity	The physical capacity used in the array
Factoring ratio	The ratio of nominal to physical capacity
Data Change rate	The rate at which data received from the tape processing application changes from task to task. This measurement is more relevant when “like policies” are compared. Data change rates can be from 1% to 25%.
Data Retention period	The period in time (usually in days) which defines how long customers will keep their disk-based tape processing online. Retention periods on average are 30-90 days but can be longer depending on business and government regulations.
Repository	The storage pool within the deduplication storage system
Index	The knowledge map that allows for the comparison of a new incoming data stream to existing data in the repository.

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 provide an estimate or measurement for the *factoring ratio*, which is defined as the ratio of nominal capacity (the sum of all User Data tape processing streams) to the physical capacity used (including all User Data, Meta Data, and spare capacity, that is, the total amount of disk storage that the user pays for).

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. The TS7680 repository will usually be a blend of different data types from different applications such as DSS, FDR, DB2, and so on. A dedupe ratio, change rate range and storage requirement will have to be estimated for each data type. After calculating these requirements for each data type you will sum them for total storage capacity for the repository.

5.2.3 Factoring ratio considerations

In this section, we discuss the factoring ratio in more detail, focusing on the parameters it depends on. 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 tape processing data 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 71.

Note: A longer retention period will yield a higher factoring ratio because the data change rate decreases and therefore less “new data” comes into the repository, reducing the physical capacity required.

- The Data Change rate

This is the rate at which the data received from the tape processing application changes from task to task. This measurement has the most relevance when “like” tape processing policies are compared (data change rates may range from 1% to > 25%, but are difficult to directly observe). The data change rate can be directly measured through an onsite assessment or it can be estimated, so the more accurate the data change rate is, the more accurate the estimate will be about the sizing of the ProtecTIER Repository (see “Calculating deduplication factoring ratios” on page 78). Note that the factoring ratio is roughly the inverse of the data change rate. The data change rate is required when IBM conducts the Data Protection Survey as part of the sizing process during an onsite presales engagement, as discussed in “Data protection survey” on page 71, and when you calculate the estimated factoring ratio, as described in “Calculating deduplication factoring ratios” on page 78.

In production environments, the ProtecTIER Repository is a blend of many tape processing policies (data types), which protect many different application and data environments. Each tape processing policy has two variables that primarily influence the realized factoring ratio (and subsequent physical storage requirements for the ProtecTIER Repository): The Data Change rate, and the Data Retention period. The values of these variables will differ across the various tape processing 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 for all data types. Once 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 tape processing policy.

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

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

The size of this “padding” will vary from situation to situation.

Note: Adding 10% to the physical storage calculations is a good rule of thumb.

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

Calculating deduplication factoring ratios

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

- ▶ Deduplication assumptions: Compression¹, tape processing change rate, and incremental tape processing 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 71 for more details).

We describe the formula gathering the main factors to simplify and better understand the meaning of each factor. As described in 5.2.3, “Factoring ratio considerations” on page 77, the factoring ratio is the nominal capacity divided by the physical capacity, and so we are going to explicitly discuss these two factors:

$$NominalCapacity = FullCapacityVersions + IncrementalCapacityVersions$$

where:

NominalCapacity	Represents the overall capacity stored in the repository during the retention period and is composed of all the full and incremental versions stored.
FullCapacityVersions	Represents the overall full tape processing capacity (expressed in GB) stored during the retention period. In the following formula, you can see how FullCapacityVersions depends on three parameters: FullCapacity, FullRetention, and FullFrequency.
FullCapacity	Represents the capacity (expressed in GB) stored during full data storage.
FullRetention	Represents the retention period (expressed in days) for the full data storage jobs (for example, you may decide to retain your full jobs for 30 days).
FullFrequency	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 has to be set to a value of 7).

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

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

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

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

IncrementalCapacityVersions

Represents the overall incremental data storage capacity (expressed in GB) stored during the retention period. In the formula below, you can see how the incrementalCapacityVersions depends on five parameters: Incremental Capacity, IncrementalFrequency, IncrementalRetention, FullRetention, and FullFrequency.

IncrementalCapacity Represents the capacity (expressed in GB) stored during incremental data storage.

IncrementalRetention Represents the retention period (expressed in days) for the incremental tape processing jobs.

IncrementalFrequency

Indicates how often you perform the incrementals during the retention period (this parameter has to be set to the value 1 if you perform an incremental every day).

FullRetention Represents the retention period (expressed in days) for the full tape processing jobs (for example, you may decide to retain your full tape processing jobs for 30 days).

FullFrequency 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 has to be set to a value of 7).

Note: In the next formula, you can see that you have to remove the number of full versions because during full tape processings, incremental tape processing is not performed.

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

For the Physical Capacity, we have the following formula:

$$PhysicalCapacity = \frac{(FullPhysicalCapacity + IncrementalPhysicalCapacity)}{CompressionRate}$$

where:

PhysicalCapacity Represents the physical capacity (expressed in GB) effectively required in the repository to satisfy the nominal capacity of your environment.

FullPhysicalCapacity Indicates the full physical capacity (expressed in GB) effectively required in the repository. In the next formula, note that a first full data store has to be entirely stored since no data is in the repository, so it is not possible to make an initial “delta” comparison.

IncrementalPhysicalCapacity

Indicates the incremental physical capacity (expressed in GB) effectively required in the repository.

CompressionRate

Describes the compression rate obtainable in ProtecTIER through its Delta Compression. Note that it is possible to reduce the initial tape processing of unique new data as well.

In the next formula, you can calculate the FullPhysicalCapacity parameter, where:

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

FullChangeRate

Indicates the estimated change rate between data store in your current environment. Again, note that a first full data store has to 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, where:

IncrementalChangeRate

Indicates the estimated change rate between incremental tape processing in your current environment. Note that a first full tape processing operation has to 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 may 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 decreasing 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 tape processing frequency: increasing the retention period or decreasing the tape processing frequency leads to an increasing factoring ratio.

5.2.4 Storage sizing

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

Disk configurations for TS7680

We discuss the possible storage array configurations and performance considerations for the ProtecTIER system.

Selecting drives for TS7680

The speed and the type of drives used 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 5-5 compares the 10K and 15K Fibre Channel drives and SATA drives (single drive).

Table 5-5 Comparison between Fibre Channel and SATA

Factor	Fibre Channel	SATA	SATA difference
Spin speed	10K and 15K	7.2K	
Command queuing	Yes 16 Max	No 1 Max	
Single disk I/O rate (number of 512 byte IOPS) ^a	280 and 340	88	.31 and .25
Read bandwidth (MB/s)	69 and 76	60	.96 and .78
Write bandwidth (MB/s)	68 and 71	30	.44

a. Note that 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 the number of revolutions per minute (RPM). A 15K 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 data can be read from or written to the drive. The average seek time gets slower when the speed of the drive increases. Typically, a 7.2K drive will have an average seek time of around 9 ms, a 10K drive will have an average seek time of around 5.5 ms, and a 15K 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 reordered, 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.

Planning the storage structure

It is important to configure a storage system in accordance with the needs of the user. An important question and primary concern for most users or storage administrators is how to configure the storage subsystem to achieve good performance. Storage subsystems capable of high IOPS like those with Fibre Channel drives will help deliver better TS7680 performance. The number of physical drives within a disk storage subsystem can also contribute to higher performance. The amount of disk storage controller memory and its Fast Write efficiency are also factors in overall performance. There is no simple answer, no best guideline for storage performance optimization that is valid in every environment and for every particular situation. You can find some preliminary (and less detailed) performance discussions in this section.

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

RAID levels

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

RAID 0

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 for 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 level 1 for applications with a high percentage of read operations and where cost is not a major concern. Because the data is mirrored, the capacity of the logical drive when assigned RAID 1 is 50% of the array capacity. Here are some recommendations when using RAID 1:

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

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

RAID 5

RAID 5 (refer to Figure 5-5 on page 83) 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 multiuser 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 a controller writes data and redundancy data to the drives in a RAID 5 array. If there is a low percentage of read activity relative to write activity, consider changing the RAID level of an array for faster performance. Use write caching on RAID 5 arrays, because RAID 5 writes will not be completed until at least two reads and two writes have occurred. The response time of writes will be improved through the use of write cache (be sure it is battery-backed up).

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

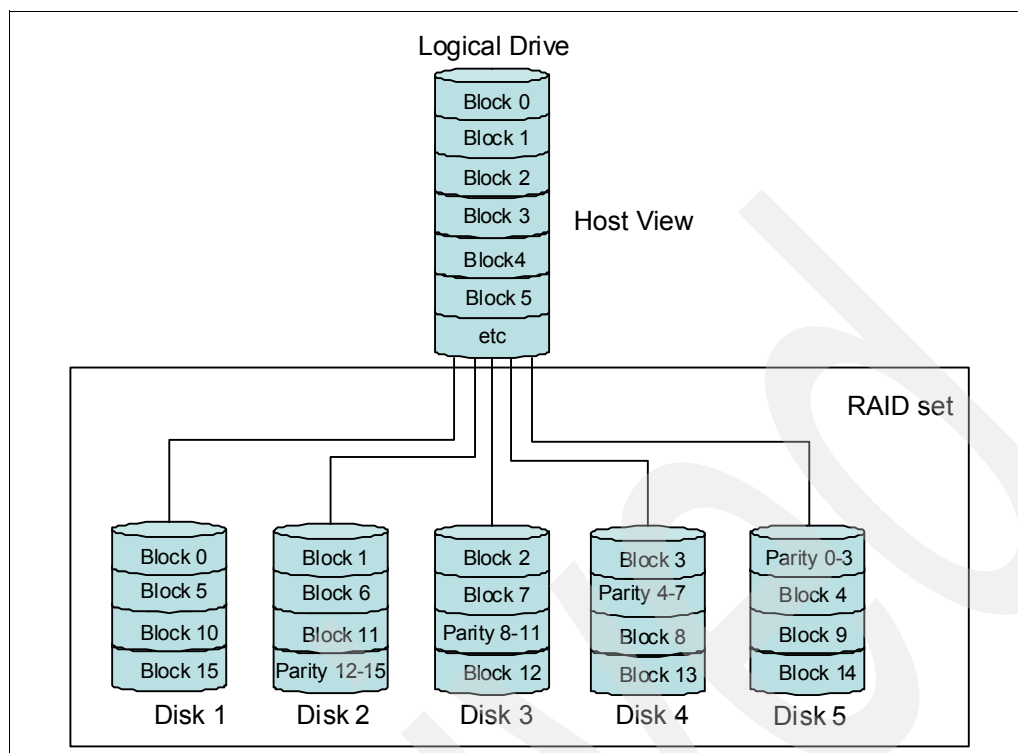


Figure 5-5 RAID 5

RAID 6

RAID 6 (refer to Figure 5-6 on page 84) 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.

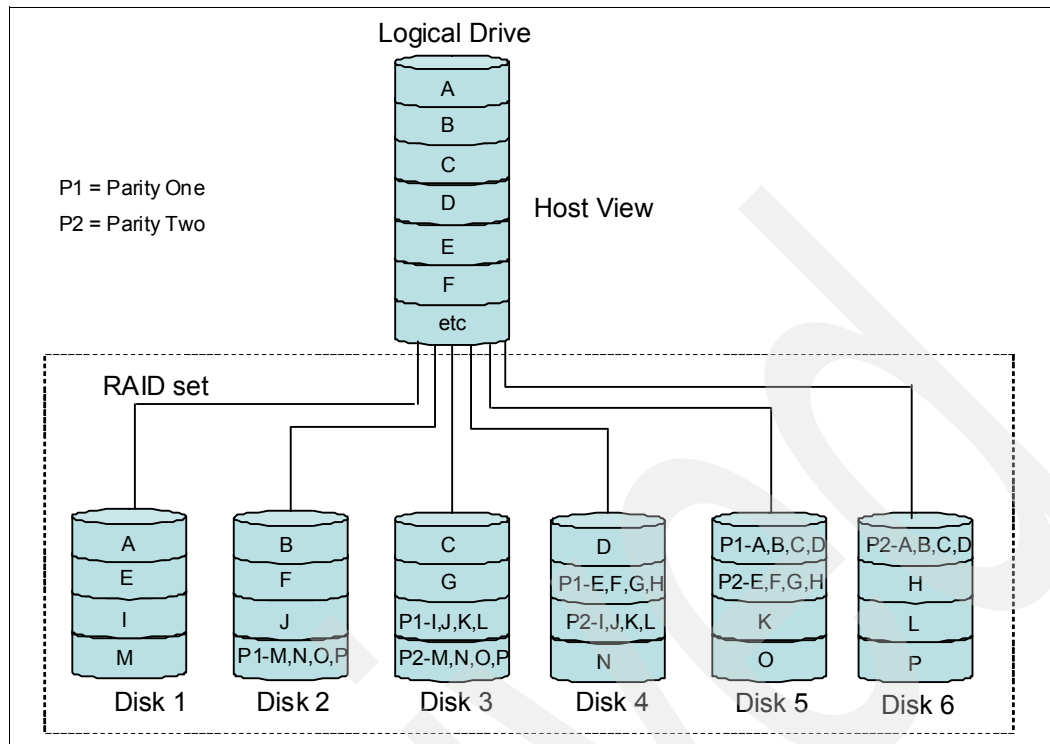


Figure 5-6 RAID 6

RAID 10

RAID 10 (refer to Figure 5-7 on page 85), 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.

The recommendations for using RAID 10 are as follows:

- ▶ 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 RAID10 when performance is critical. Use write caching on RAID10. Because a RAID10 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 it is battery-backed up).
- ▶ 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.

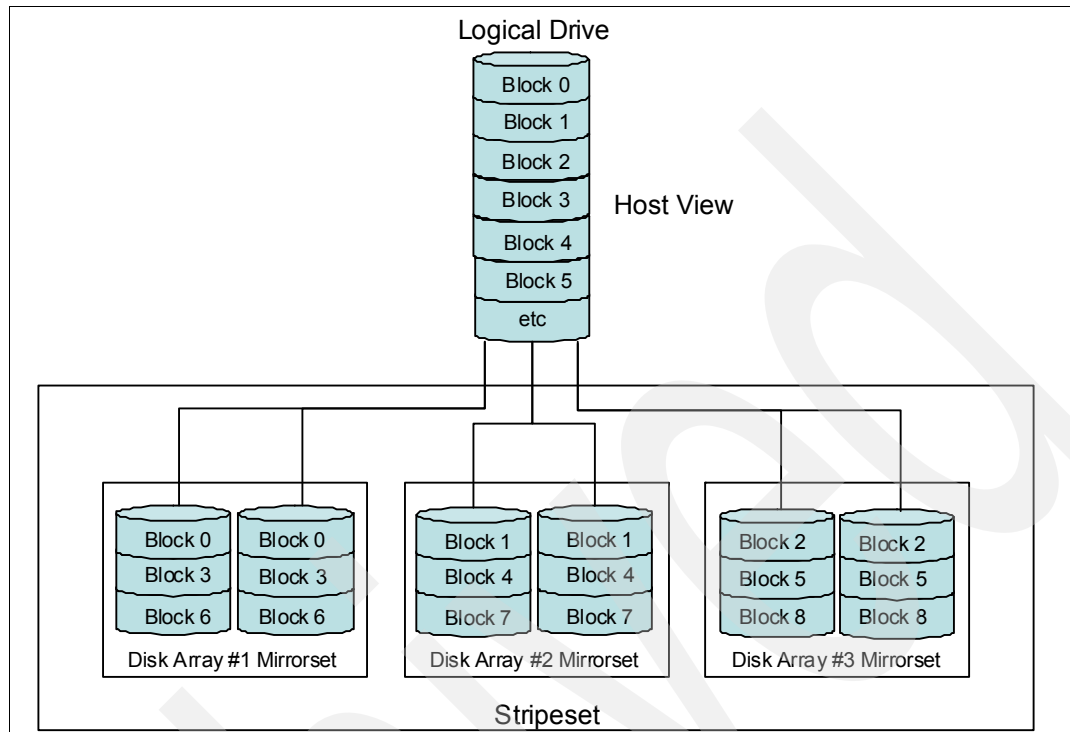


Figure 5-7 RAID 10

Planning for cartridges

Prior to discussing the various elements of the TS7680 that monitor and manage capacity, some background discussion is warranted to help provide a foundation for understanding.

The TS7680 is designed to mimic the behavior of (and management of) a traditional tape library as closely as possible. It is designed to be used intuitively by the storage administrator, who is typically trained and experienced in managing tape capacity. The TS7680 introduces some challenges not associated with traditional tape libraries:

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

The TS7680 has internal mechanisms to manage the nominal capacity of the system and communicate with the storage administrator “in the language of tape.”

Capacity management in the traditional tape library paradigm

To explain how the TS7680 adheres to tape management principles, we review some key points of managing tape capacity.

The objective of the storage administrator is to make sure that there is enough capacity for the foreseeable future. In the tape processing world, storage administrators pay close attention to the number of tapes available for the day’s operations. Through the tape processing application console, the number of available cartridges is visible. The tape processing application cannot directly see the capacity of a cartridge, but the administrator knows the capacity based on cartridge type and compression ratios. The administrator must do the math each day to answer the following question: How many tapes are available, and

what is the capacity of each? By calculating the total available capacity, the administrator knows whether the daily payload will fit on available tapes.

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

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

One other key point to note: In the physical tape world, there are "Early Warning" (EW) signals provided to the tape processing application by the tape drive when a tape cartridge is nearing its end. This signal allows the tape processing application to change to a fresh cartridge efficiently. This EW signal is relevant to understanding the TS7680 capacity management.

The TS7680 versus traditional tape libraries

In many ways, the TS7680 virtual tape solution is similar to traditional tape libraries. It is designed for easy use by tape operators that do not have to learn a new paradigm from an operational perspective.

Once installed, the TS7680 behaves and is managed like a standard library. It has a nominal capacity that is available for use by the tape processing application. This capacity is represented by a certain number of cartridges, each with a given capacity. Just as with real tape libraries, with the TS7680, the storage administrator uses the number of tapes available and the capacity of each as the key indicators. Just as with real libraries, the TS7680 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 tape processing totals.

However, while there are many similarities between the TS7680 and traditional tape libraries with capacity management from the storage administrator's point of view, there are also some differences as well.

Cartridge capacities within the TS7680 fluctuate; as a result, the number of tapes used per day fluctuates, even when the payload stays constant. Space reclamation is also different. New data sets have a bigger impact on capacity, and capacity expansion requires additional steps that will need to be learned as part of the operational process.

How the TS7680 manages changes in nominal capacity

The initial factoring ratio is always an estimate and the actual data change rate of data that enters the TS7680 fluctuates each day and over time. The result is that the *nominal capacity* of the overall system fluctuates. The TS7680 manages these changes in 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 storage administrator to manage fluctuating capacity. The "results" of the learning algorithm are visible to the storage administrator through the usage of cartridges on a daily basis. If capacities decline, the Early Warning (EW) of an individual tape will arrive sooner, which in turn requires the tape processing application to request a new tape. The overall effect on a daily basis is that more tapes will be consumed. Inversely, if capacities increase, the EW of a cartridge arrives later, and less tapes are used overall. Just as with a traditional tape library, the administrator is able to track the tape usage statistics to manage

the system capacity. Thus, the paradigm of traditional tape library management is closely mimicked.

Managing capacity fluctuations

As mentioned, cartridge capacity changes to reflect the system-wide shift in nominal capacity. In Table 5-6, the factoring ratio at “equilibrium” is greater than the factoring ratio that was used when the TS7680 was first installed. There are 1000 tape cartridges, but since 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 TS7680 handles this through the “learning algorithm” mentioned earlier. The tape processing 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 TS7680, the increase in cartridge capacity is transparent to the tape processing application.

Table 5-6 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 5-7, the factoring ratio at “equilibrium” is less than the factoring ratio that was used when the TS7680 was first installed. As you can see in Table 5-7, there are 1000 tape cartridges, but since 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 5-7 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	100 TB	100 GB
Day 30	10 TB	1000	8:1	80 TB	80 GB

As cartridge size changes, the Early Warning (EW) signal will arrive sooner or later than originally. In the example shown in Table 5-6, the EW for each cartridge will occur 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 5-7, the EW for each cartridge will occur 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 less 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 for the TS7680

The objective of capacity management is to understand the total capacity requirement of the VTL, factoring in the compression that will be expected.

Capacity is a function of the daily and weekly workload, and the duration in which the data remains accessible on tape before being returned to scratch.

Capacity management begins from the moment the TS7680 initially starts to accept data.

Capacity management implications: initialization phase

During the first weeks of a TS7680 implementation, the daily fluctuations in capacity are more pronounced. This is normal system behavior as the TS7680 learns the true data change rate. You should not be alarmed during this phase when cartridge capacity oscillates, sometimes significantly. Once the system runs a full tape processing cycle (for all data sets that it will manage), the capacity changes should stabilize.

Capacity management implications: management of the TS7680

From an ongoing management perspective, the storage administrator must be aware of the following when running a TS7680 virtual tape solution within the corporate tape processing environment. As cartridge size changes, the Early Warning (EW) signal will arrive sooner or later than originally. In the example in Table 5-6 on page 87, the EW for each cartridge will occur 20 GB later on Day 30 than on Day 1, allowing more data to fit on a given cartridge. As a result, more or less tapes will be consumed during a given day's workload. Therefore, storage administrators for the TS7680 must keep track of the number of cartridges used as a key indicator to indicate capacity fluctuations.

Capacity management implications: adding new data sets to an existing TS7680

Often during the initial phase, or some time thereafter, you will decide to send more data to a TS7680. While this is a common occurrence, this situation creates new implications of which the storage administrator must be aware.

A new data stream will have a very high change rate (given that all of the data is new to the TS7680). This will cause an increase in the system-wide change rate, and a decrease in the nominal capacity since the factoring ratio is going to decrease. As the new data set runs through a full cycle, the nominal capacity may or may not return to what it was previously, depending on the data change rate of the new data set. Given the variability that is inherent in this situation, you must be aware of the phenomenon and understand the impact. The best way to add new data streams is to first sample the data to project the likely impact. In some cases this may create a need for more physical disks.

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 storage. This process is easy to manage and understand: When a cartridge is returned to the available pool, its full capacity is available for new storage.

From the tape processing application point of view, the process of tape recycling is exactly the same in a TS7680: Cartridges expire and are returned to the available pool for new data. However, the underlying process of space reclamation in the TS7680 is unique.

As soon as a cartridge begins to receive new data at the beginning of its media, the TS7680 knows that it has been recycled by the tape processing 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 TS7680 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%.

Performance planning best practices

These are some simple rules of thumb for high-level planning:

- ▶ Configure for future performance and capacity requirements.
- ▶ For optimal performance, configure a minimum of 24 file systems.
- ▶ One file system per RAID group.
- ▶ For selecting RAID types, more spindles per RAID group for User Data yield more usable capacity, but for performance you want to ensure that there are enough file systems.
- ▶ Compare the effect of different disk options.

Summary of TS7680 ProtecTIER systems capacity management

In summary, the benefit of disk savings enabled by the TS7680 also introduces some new capacity management challenges of which the administrator must be aware. The TS7680 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 TS7680 to maintain the “language of tape” to the storage administrator:

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

Archived

TS7680 setup

This chapter provides information about how to power up the 3958-DE2 (TS7680) in the right sequence and configure it at the initial installation.

We will describe the following items:

- ▶ Powering up the 3958-DE2 (TS7680)
- ▶ Setting up the TSSC for use with the 3958-DE2
- ▶ Configuring the Reliability, Availability and Serviceability (RAS) package
- ▶ Verify and create a complex and testing call home
- ▶ Installing the ProtectTIER Manager GUI on external TSSC
- ▶ Configuring ProtectTIER using ptconfig
- ▶ Adding a node and creating a repository
- ▶ Configuring Upper PT Server, creating a library and port configuration
- ▶ Testing a clustered system configuration
- ▶ ProtectTIER Library attachment to the enterprise controllers

6.1 PowerUP

This topic provides instructions for powering up the 3958-DE2 (TS7680) hardware components.

Important: Disk expansion modules and disk controllers that need to be attached to this TS7680 as Disk Repository must be already powered up and ready.

You must power up the components in the order shown here (refer to Figure 6-1 on page 93 for a view of the front and rear of the 3958-DE2 frame):

1. Frame breakers (your main wall power circuit breaker)
2. Power Control Assembly (PCAs) by turning the Emergency Unit Power Off (EUPO) switch on the front of the frame to the ON position.
3. Western Telematic Inc. (WTI) network power switch, located on the rear left frame
4. TS3000 System Console (TSSC), keyboard, video, mouse (KVM), and KVM switch
5. ProtecTIER server (lower), and wait until the Login panel appears
6. ProtecTIER server (upper) and wait until the Login panel appears
7. Verify that the Enhanced Routers are powered on, located on the rear middle frame
8. Enterprise Controller (lower controller), located on the front bottom frame
9. Enterprise Controller (upper controller), located just above the lower controller

Note: The 3958-DE2 has a keyboard monitor tray which pulls out for use. Use the PrtSc key to select either the upper or lower ProtecTIER server or the TSSC if it is in the frame.

The TSSC is used for connection to the Enterprise controller to open a terminal window and Telnet to 172.31.1.xx2 for the lower controller, or 172.31.1.xx7 for the upper controller.




Figure 6-1 TS7680 front and rear view

6.1.1 TS3000 System Console, keyboard, video, mouse; and keyboard, video, and mouse switch

Complete this task to power on the TS3000 System Console (TSSC), keyboard, video, mouse (KVM), and KVM switch.

Perform the following steps to power on the TSSC, KVM, and KVM switch:

1. If the TSSC is located in the 3598-DE2 (TS7680) frame, then press the white power control button on the front of the TSSC.
2. If the TSSC is external, then verify that the power switch is in the ON position.
3. Loosen the thumb screws on the front of the KVM.
4. Pull the KVM out and open the display.
5. Press the black button on the front of the KVM display panel.
6. Verify that the power switch on the KVM switch is in the ON position.

6.1.2 ProtecTIER servers

This topic defines what occurs when the ProtecTIER servers are powered on.

The back-end disk array subsystem must be powered up first when powering up the whole 3598-DE2 frame.

When the ProtecTIER server is connected to an AC power source but is not turned on, the operating system does not run, and all core logic except for the service processor is shut down. However, the ProtecTIER server can respond to requests from the service processor, such as a remote request to turn on the ProtecTIER server. The power-on light-emitting diode (LED) flashes to indicate that the ProtecTIER server is connected to AC power, but the ProtecTIER server is not powered on.

Important: Power on the lower ProtecTIER server first, then power on the upper ProtecTIER server.

Approximately 20 seconds after the ProtecTIER server is connected to power, the power-control button becomes active, and one or more fans might start running to provide cooling while the ProtecTIER server is connected to power.

- To power on the ProtecTIER server, press the white, recessed power-control button on the ProtecTIER server operator panel, shown in Figure 6-2.
 - When the panel displays Registered calypsa with major device 249, press Enter.
 - Login input should be displayed.




Figure 6-2 ProtecTIER server operator panel

6.1.3 Enterprise controllers

Complete this task to power on the Enterprise controllers.

Important: Power on the lower Enterprise controller first, then power on the upper Enterprise controller.

Perform the following steps to power on the Enterprise controllers:

1. Verify that the Enhanced Routers are powered on.
2. Verify that the two line cords are connected to the Enterprise controller's power supplies.
3. When the Enterprise controllers are in standby mode (blinking green light-emitting diode (LED) 2 on the operator panel), press the power ON button 1, which is located on the Enterprise controller's operator panel. The controller panel displays READY or the information shown in 3 when the controller has powered on successfully; see Figure 6-3 on page 95.

This could take up to 15 minutes.




Figure 6-3 Enterprise Controller operator panel

Note: If the numbers on the operator panel do not change for over 5 minutes and you are unable to log into the Enterprise controller, then call your IBM System Service Representative.

6.2 Setting up the TSSC for use with the 3958-DE2

Before the TSSC can be used to support the 3958-DE2, you must first configure the TSSC's Call Home, firewall, Time for UTC time, Enable NTP service, and network communications settings.

1. If it is not already running, power up the TSSC. If you see a blank panel, using the KVM switch, press the PrtSc key on the TSSC's keyboard. Select **TSSC** from the list of devices.

Note: If you need to Backspace when typing, use Cntl-H or Cntl -Backspace.

2. If prompted for login information, enter the username service and the password service.
3. Right-click the TSSC's blue desktop.

The IBM TS3000 System Console menu displays.

Note: Verify, on top of the menu, that the TSSC licensed internal code level is 5.5.x or higher.

4. Select **System Console Actions** and from here select **Console Configuration Utility**. See Figure 6-4 on page 96.




Figure 6-4 TSSC System Console Actions

The “Username and password” panel displays. See Figure 6-5.

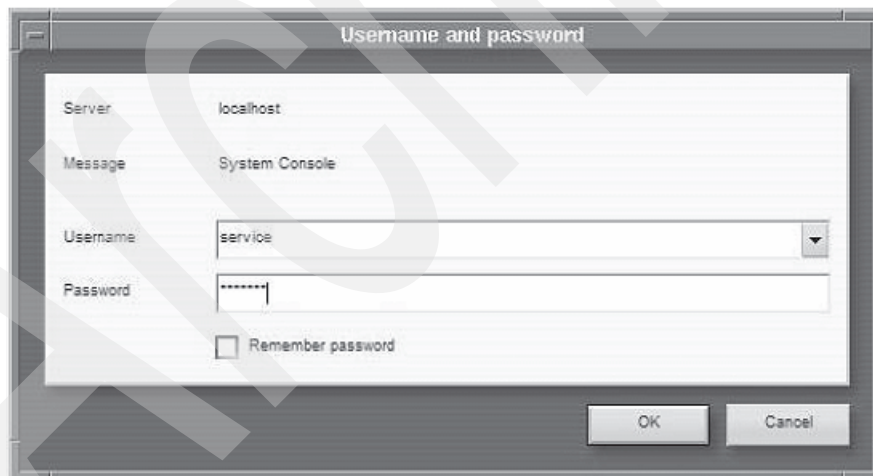


Figure 6-5 TSSC Username and password

5. Enter service in both the Username and Password fields.
6. Click **OK**.

The Console Configuration application starts and the Console Configuration utility panel displays. See Figure 6-6 on page 97.

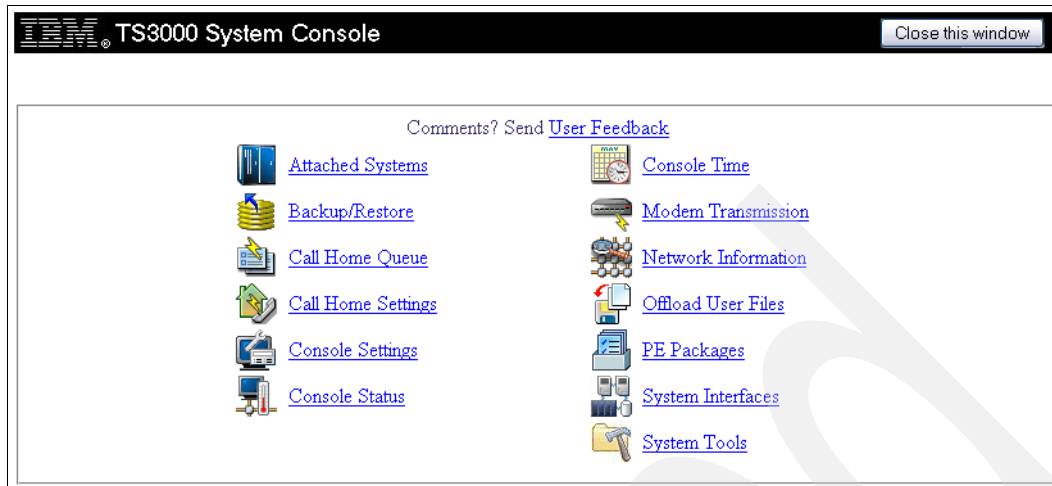


Figure 6-6 TSSC console configuration

7. Click **Console Settings**.

The Console Settings panel displays. See Figure 6-7.

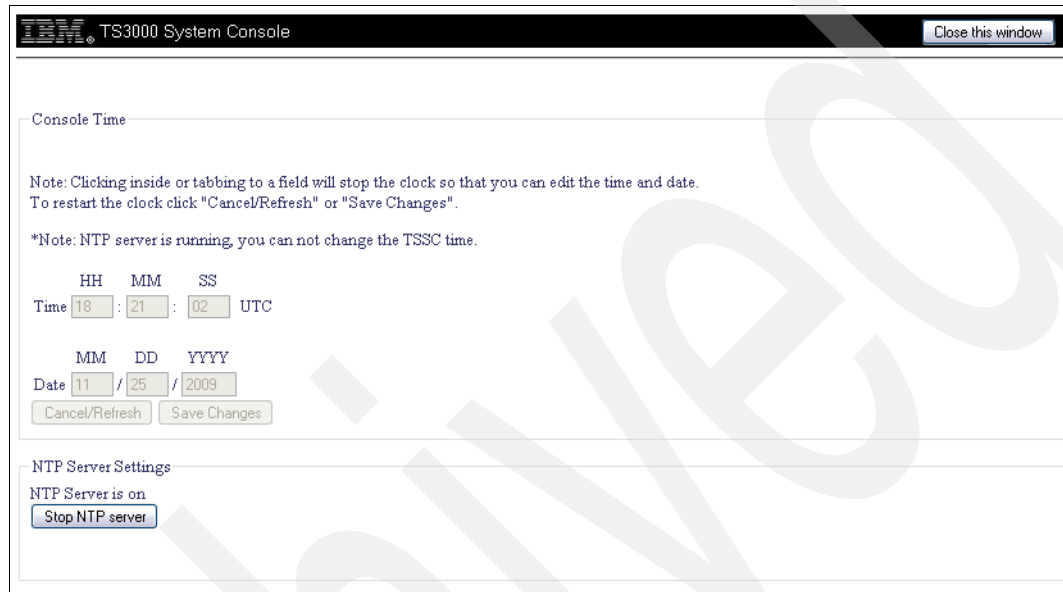
Figure 6-7 TSSC console settings

8. On the Internal Network Interface panel, enter IP Address: 172.31.1.1 and Subnet Mask: 255.255.255.0.
9. Fill in with the appropriate values for external address, which will be used to connect to this TSSC from your private network on the External Network and System Properties information fields.

10. Click **Save Changes** and then **Close the Window**. The main Console Configuration Utility panel will be displayed.

6.2.1 Setting Console Date, Time and Network Time Protocol (NTP) server

1. From the Console Configuration utility panel, select **Console Time**. The following panel is displayed; see Figure 6-8.



IBM TS3000 System Console Close this window

Console Time

Note: Clicking inside or tabbing to a field will stop the clock so that you can edit the time and date. To restart the clock click "Cancel/Refresh" or "Save Changes".

*Note: NTP server is running, you can not change the TSSC time.

Time : : UTC

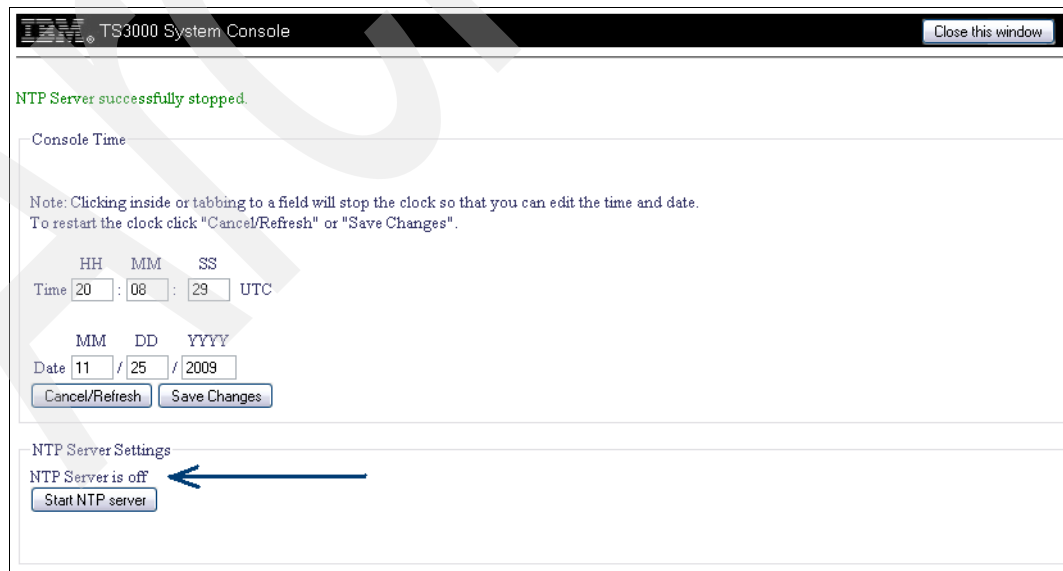
Date / /

NTP Server Settings

NTP Server is on

Figure 6-8 TSSC Console Time

2. To change the system clock on the TSSC, you must first make sure that the NTP server is stopped. To do this, click **Stop NTP server** and you will get an update panel saying that the NTP server is off; see Figure 6-9.



IBM TS3000 System Console Close this window

NTP Server successfully stopped.

Console Time

Note: Clicking inside or tabbing to a field will stop the clock so that you can edit the time and date. To restart the clock click "Cancel/Refresh" or "Save Changes".

Time : : UTC

Date / /

NTP Server Settings

NTP Server is off

Figure 6-9 TSSC NTP server stopped

3. Enter the values in the Time and Date fields.

4. Click **Save Changes** when you are finished.
5. To start the NTP server, click **Start NTP server**.

Note: If the sysplanar has just been replaced or if the time difference between the TSSC and the Enterprise Controller and/or the ProtecTIER server is off by more than 2 or 3 minutes, then follow the NTP instructions in the TS7680 Service Information Center before starting the NTP server.

The time zone on all systems must be the same: one of UTC, CUT, or GMT. Click **Save Changes**. It may be necessary to log out and log in again to refresh the date and time indicator on the desktop application bar. It is also possible that the screen saver may be invoked when resetting the time.

6. Verify that the NTP service is running by selecting **Network Information** on the main menu of the Console Configuration Utility window; see Figure 6-6 on page 97. Select **NTP** from the button on the left side of the panel; see Figure 6-10.

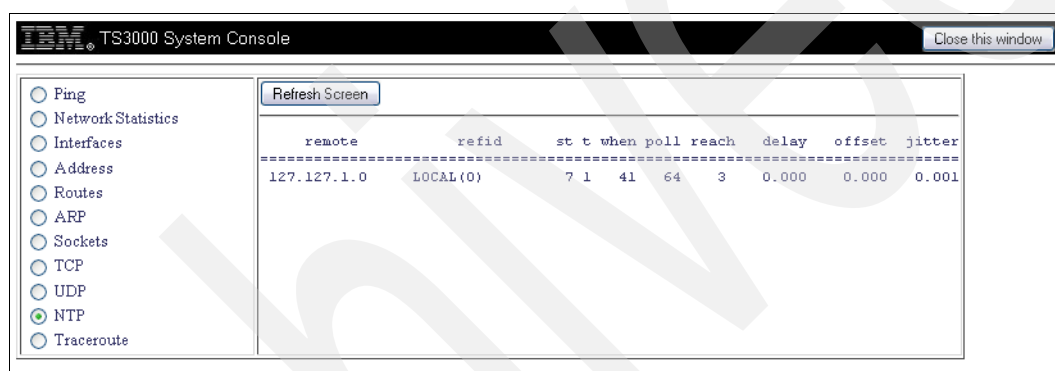


Figure 6-10 TSSC network information

6.3 Configuring the RAS package

Configuration of the Reliability, Availability, and Serviceability (RAS) package allows the ProtecTIER servers to interact with the TSSC for Call Home, log collection, and system health monitoring.

Verify that the following prerequisites are met before starting the RAS package configuration:

- The disk components are installed, configured, cabled, and powered on.
- The TSSC is installed, configured, cabled, and updated to code level 5.5.x or higher. To verify the TSSC's code level:
 - On the TSSC, right-click the blue desktop. The IBM TS3000 System Console menu displays.
 - Make note of the TSSC code level, displayed in the IBM TS3000 System Console title bar.
 - Verify that the TSSC is being used as the NTP server.

Note: For detailed information about TSSC code installation and cable connection, refer to *IBM System Storage TS3000 System Console (TSSC) Maintenance Information*.

- The TSSC's network IP addresses have been determined.
- If you are adding new 3958-DE2 systems to an existing TSSC, verify the IP addresses of any systems that are already attached to the TSSC. Doing so will allow you to avoid potential conflicts and adjust the IP addresses of the new systems to use alternate ranges. Each 3958-DE2 requires a block of ten IP addresses. Starting at xx0. Where xx is the frame location number.
- Verify that an existing TSSC meets the minimum hardware requirements. See *IBM System Storage TS3000 System Console (TSSC) Maintenance Information*.

The RAS package configuration procedure performs the following tasks:

Sets up the TSSC network (172.31.1.xxx) for the ProtecTIER servers and their Remote Service Adapter (RSA) cards.

6.3.1 ProtecTIER RAS package configuration

During the installation, you will be prompted to enter the system and customer contact information. You will need the information for each ProtecTIER server in the cluster. It is recommended that you collect the necessary information before you start the installation, so it is readily available when needed.

- ▶ Frame (location) number for this system.

Note: Each node in the frame is assigned a different IP address within the range of xx0 to xx9, where xx is the frame (location) number that can be 10 or 20... up to 240.

- ▶ Machine type and model number (3958-DD3)
- ▶ Business company name
- ▶ Machine location
- ▶ Callback number
- ▶ Voice phone number
- ▶ Disk array machine type and model number
- ▶ Disk array serial number
- ▶ Customer number
- ▶ Customer SMTP IP address
- ▶ Customer SMTP email address
- ▶ Country code

Note: This is the two-digit, IBM-assigned country code used to order software or to acquire software support. Do *not* confuse this with the three-digit RETAIN country code.

1. Log in to lower ProtecTIER Server. To do so:
 - a. If the TSSC resides inside the 3958-DE2 frame, go to Step 2.
 - b. If the TSSC resides outside the 3958-DE2 frame, verify a connection from its network to port 5 of the TSSC Network Switch in the 3958-DE2 frame.
2. Using the KVM switch, press the PrtSc key on the TSSC's keyboard.

- a. Select the lower ProtecTIER Server from the list of devices. If the panel goes blank, press Enter.
 - b. At the lower ProtecTIER Server command prompt, log in with the ID root and the password admin.
3. From the command line, enter the following command:
rsCerCfgA11 (you may need to enter the entire string **/opt/ras/bin/rsCerCfgA11.**)

Note: USB device errors might be displayed during the install.

4. Follow the on-screen instructions to complete the RAS package configuration process.

Note: If you receive an error that the frame range number entered is already in use, specify a different range in an increment of 10. For example, if frame range 20 is already in use, try frame range 30. Use the same frame number for each Enterprise controller. The default frame is 240. Change this to the lowest unused frame number for future growth. (TSSC IP address ranges xx0—xx9).

5. When the RAS package configuration is complete on the lower ProtecTIER server, repeat Steps 1 through 4 to configure the RAS package on the upper ProtecTIER server.

Note: Use the same frame range number for lower and upper ProtecTIER servers.

Type yes to the question *Are you setting up the second node in a cluster?*

When asked: *Do you want to change/update the DS or RSA configuration?* Answer yes.

6.3.2 Verifying the cluster's Ethernet connections

Before performing RAS verification, use the following procedures to ensure that the Ethernet connections for the clusters are configured correctly.

Start on the lower node.

Note: The MAC addresses will be different for each system that is installed. The following outputs are only examples. The hwaddrs are specific to the Ethernet cards in each system.

1. At the lower ProtecTIER server command prompt, log in with the ID root and the password admin.
2. Use kudzu to run the following command and view the output to ensure that the Intel Ethernet cards are assigned to eth0, eth1, eth2, and eth3, and the Broadcom Ethernet cards are assigned to eth4 and eth5:

```
[root@localhost ~]# kudzu -p -c network | grep -A3 ?device:? | more <enter>
```

```
device: eth0
driver: e1000e
desc: ?Intel Corporation 82571EB Gigabit Ethernet Controller? network.hwaddr:
00:15:17:94:37:aa
--
device: eth1
driver: e1000e
```

```

desc: ?Intel Corporation 82571EB Gigabit Ethernet Controller? network.hwaddr:
00:15:17:94:37:ab
--
device: eth2
driver: e1000e
desc: ?Intel Corporation 82571EB Gigabit Ethernet Controller? network.hwaddr:
00:15:17:94:37:ab
--
device: eth3
driver: e1000e
desc: ?Intel Corporation 82571EB Gigabit Ethernet Controller? network.hwaddr:
00:15:17:94:41:15
--
device: eth4
driver: bnx2
desc: ?Broadcom Corporation NetXtreme II BCM5709 Gigabit Ethernet?
network.hwaddr: 00:1a:64:db:36:24
--
device: eth5
driver: bnx2
desc: ?Broadcom Corporation NetXtreme II BCM5709 Gigabit Ethernet?
network.hwaddr: 00:1a:64:db:36:26

```

3. Use ethtool to verify the following:

For eth0 and eth5: Speed = 100 Mb/s or 1000 MB/s

For eth1 and eth2: Speed = 1000 Mb/s

For eth0 and eth5 (Standalone): Link detected = yes

For eth0, eth5, (eth1 eth2, when Clustered): Link detected = yes

For example:

```

# ethtool eth1
Settings for eth1:
Supported ports: [ TP ]
Supported link modes: 10baseT/Half 10baseT/Full
100baseT/Half 100baseT/Full 1000baseT/Full
Supports auto-negotiation: Yes
Advertised link modes: 10baseT/Half 10baseT/Full
100baseT/Half 100baseT/Full
1000baseT/Full
Advertised auto-negotiation: Yes
Speed: 1000Mb/s
Duplex: Full
Port: Twisted Pair
PHYAD: 1
Transceiver: internal
Auto-negotiation: on
Supports Wake-on: d
Wake-on: d
Current message level: 0x00000001 (1)
Link detected: yes

```

4. Repeat steps 1 to 3 for the upper node.

6.4 Verify or create a complex and testing Call Home

In this topic we describe how to verify or create a complex with the attached systems and test ProtecTIER and Enterprise controllers Call Home.

Note: The Enterprise controllers configuration as well as licensed internal code update, Cisco Router and TSSC Call Home configurations must be already done by your IBM System Service Representative (SSR).

1. If you are not already logged into the TSSC, do so now, with the username service and the password service.
2. Right-click the TSSC's blue desktop.
The IBM TS3000 System Console menu displays.
3. **System Console Actions** → **Console Configuration Utility**, see Figure 6-6 on page 97.
If prompted for a username and password, enter service in both fields.
The Console Configuration Utility starts and the Console Configuration window displays.
4. In the Console Configuration Utility window, select **Attached Systems**. See Figure 6-11 as a reference.

IBM TS3000 System Console Close this window

[Attached Systems](#) [Complex View](#)

<input type="checkbox"/>	Device-Model	Serial Number	ID	Host Name	IP address	Subnet Mask	Call Home Switched	Standby	Complex System
<input checked="" type="checkbox"/>	3592C06	78C6670	1	3590cu10	172.31.1.12	255.255.0.0	1	0	3958DE2 13FF007
<input checked="" type="checkbox"/>	3592C06	78C6668	2	3590cu20	172.31.1.17	255.255.0.0	1	0	3958DE2 13FF007
<input checked="" type="checkbox"/>	3958DD3	78KHKYZ	3	tssc_Marines	172.31.1.10	255.255.255.0	1	0	3958DE2 13FF007
<input checked="" type="checkbox"/>	3958DD3	78HZBNR	4	tssc_Navy	172.31.1.15	255.255.255.0	1	0	3958DE2 13FF007

172 . 31 . 1 .

Health Legend:

- No Problems Returned From System Health Check
- Warning(s) Found in System Health Check
- Failed Status(es) Found in System Health Check
- Communication Failure Between TSSC and Attached System

Figure 6-11 TSSC Attached Systems

5. Select **Complex View** in the Attached Systems panel. If the installation configuration is already done by the IBM System Service Representative, you will get a panel like the one

in Figure 6-13 on the right side If the Complex has not been created yet, you will receive a panel like the one in Figure 6-12 with nothing listed under the Complex View title.

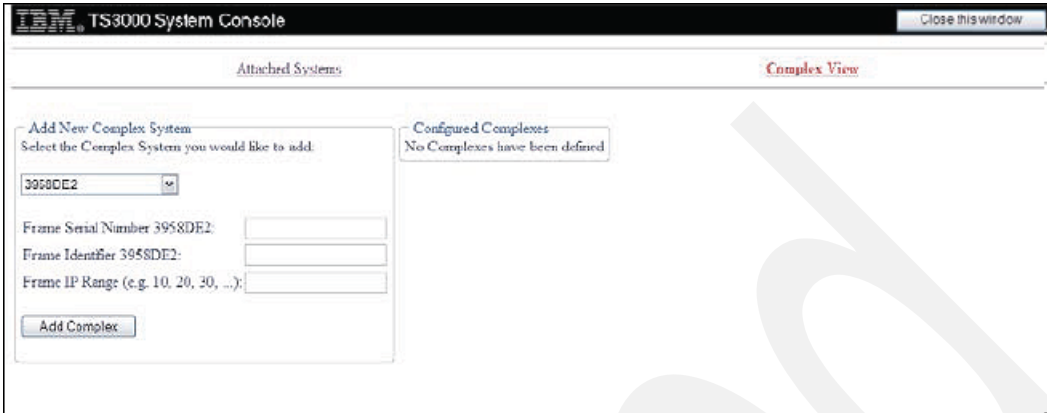


Figure 6-12 TSSC Complex view

6. Add a new complex system by selecting the Complex Type from the drop-down menu. Select **3958DE2** for this complex.
7. Enter the Frame Serial Number, Frame Identifier (for the Frame Identifier enter the Composite Library Sequence Number) and Frame IP Range (for example, 10, 20, 30) (frame location number). Press **Add Complex**. Figure 6-13 shows *all* the 3958DE2 options displayed. Each attached system within the frame number range is added to the frame cluster.

Note: The frame serial number is located in the bottom rear of the frame.




Figure 6-13 TSSC Add Complex

Note: All four systems must be attached to create the frame complex. The window will display any errors in your Complex. Correct the errors that are displayed and attempt the process again. Check it under Complex Pair Information.

8. Select **Close this window**.

6.4.1 Testing Call Home

This topic covers the procedure to test the Call Home service on ProtecTIER servers.

Testing Call Home on both Enterprise controllers is also possible, but should be done by your IBM System Service Representative (SSR).

The 3958-DE2 supports four types of Call Home activity:

- ▶ Error Initiated Call Home – Triggered by failures on a ProtecTIER server.
- ▶ Heartbeat/MRPD Call Home – Regularly scheduled calls to report system status, aliveness, and configuration of the ProtecTIER server.
- ▶ Test Call Home – Triggered from the RAS command line interface (CLI) to test the path to the IBM Customer Configuration Profile File (CCPF) system.
- ▶ User Initiated Call Home – Triggered from the TSSC graphical user interface (GUI) to collect a product engineering (PE) package.

Testing Call Home on both ProtecTIER servers

This will guide you in testing the ProtecTIER servers Call Home.

The lower ProtecTIER server will have an IP of 172.31.1.xx0; the upper ProtecTIER server will have an IP of 172.31.1.xx5 (xx will be the frame number).

Example: Frame 10 will create Lower ProtecTIER server as 172.31.1.10 and Upper ProtecTIER Server as 172.31.1.15.

1. Do the following to log on to a ProtecTIER server:
 - a. Right-click the TSSC blue panel.
 - b. Select **System Console Actions** → **Telnet to Tape Systems**, then select the system to log on to (Lower or Upper ProtecTIER server).
 - c. Enter the password ptadmin.
2. Complete the Call Home process on the server by either entering the command from the Command line or running the command using the rasMenu tool.
 - Command line method
 - i. To enable Call Home on the server from the command line, enter the following command from the command line: **rsCerCHFunction -e**.

Note: If you get a return message that Call Home is already enabled, continue to the next step.

- ii. To complete the Call Home test on the server from the command line, enter the following command: **rsCerCHTest** and press Enter.

After a few seconds, you should receive the message, Test Call Home sent successfully. Continue to step 3.
- rasMenu Method
 - i. To complete Call Home on the server using the rasMenu tool, at the server command line, enter the following command: **rasMenu** and press Enter.
 - ii. Enter the number for the option for Call Home Commands.

- iii. In the Call Home Commands submenu, enter the number for the option to Enable Call Home.

Note: If you get a return message that Call Home is already enabled, continue to the next step.

- iv. In the Call Home Commands submenu, enter the number for the option to Test Call Home.

After a few minutes, you should receive the message, Test Call Home sent successfully. Continue to step 3.

3. The TSSC Call Home Queue should have the test message listed or pending transmission. To check the Call Home Queue, right-click an empty area of the TSSC's desktop, and from the IBM TS3000 System Console menu, select **System Console Actions** → **Console Configuration Utility**.
4. From the menu, select the icon for **Call Home Queue**. See Figure 6-14.

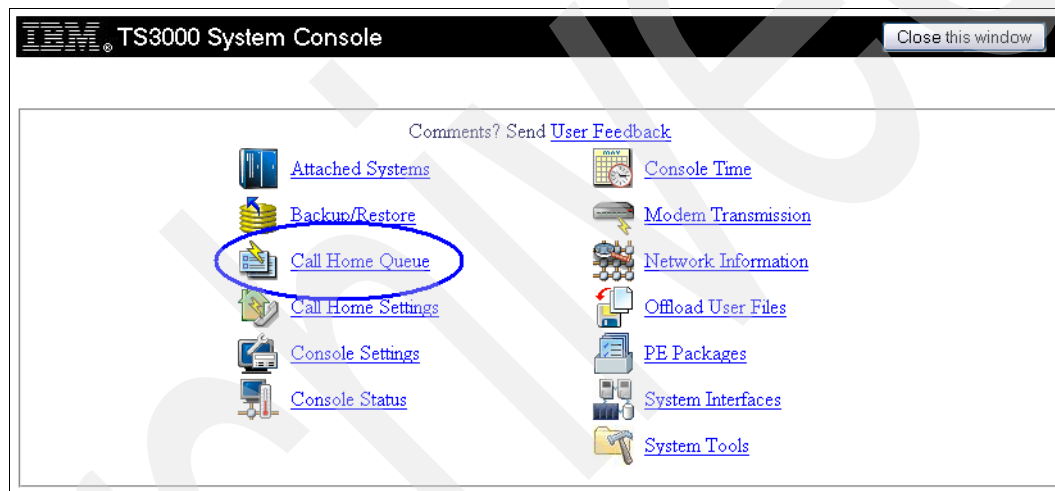


Figure 6-14 TSSC console configuration

5. Check the Call Home Queue for the existence of the test message.
A Call Home Queue window similar to Figure 6-15 on page 107 appears.

IBM TS3000 System Console Close this window

[Call Home Queue](#) [Call Home Event Log](#) [Failed Queue](#)

Selection Functions

- Delete Selected Entries
- Increase Selected Entries Priority
- Show Selected Entries Details
- Offload Selected Entries to USB ☐ Delete after successful offload
- Offload Selected Entries to CD ☐ Delete after successful offload

Last refreshed on: Nov 30 21:21:22 UTC 2009 Refresh Queue

<input type="checkbox"/>	Priority	Type	Hostname	Machine Type	Machine Model	Serial Number	Cluster	Name	Status	Size	Modified	Created
<input type="checkbox"/>	1	DC	usmc-fusion	3958	DE2	13FF007	1	v32322	Pending Transfer	204800	Nov 30 21:19:32 UTC 2009	Nov 26 02:05:07 UTC 2009
<input type="checkbox"/>	2	DC	tssc_Marines	3958	DD3	78KHKYZ	1	w2Modh	Pending Transfer	61440	Nov 30 21:19:33 UTC 2009	Nov 26 17:14:08 UTC 2009
<input type="checkbox"/>	3	DC	tssc_Navy	3958	DD3	78HZBNR	1	M49tj0	Pending Transfer	51200	Nov 30 21:19:34 UTC 2009	Nov 26 17:17:08 UTC 2009

Figure 6-15 TSSC Call Home Queue

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

Note: Each 3958-DE2 frame must be under valid warranty or Maintenance Agreement (MA) coverage or it will be rejected and no record or Problem Management Report (PMR) will be generated for the Call Home event.

6.5 Installing the ProtecTIER Manager GUI on external TSSC

Use this information to install the ProtecTIER Manager graphic user interface (GUI) on the TSSC.

1. Right-click anywhere on the desktop.



Figure 6-16 TSSC PT Manager functions

2. Select **Browser Functions** → **ProtectTIER Functions** → **Install GUI**; see Figure 6-16. You are prompted to insert the IBM System Storage ProtecTIER Enterprise Edition V2.4 CD into the TSSC CD drive. Press Enter. Installation of the GUI begins.
3. You see the Introduction window. Click **Next**.
4. Click **Next** in the Choose Install Folder panel (not shown) to accept the default.
5. Click **Next** in the Choose Link Folder panel (not shown) to accept the default.
6. Read and accept the terms that are depicted on the License panel (not shown). Select **I accept both the IBM and the non-IBM terms**, and click **Next**.
7. You see Figure 6-17 on page 109. This panel shows you where the ProtecTIER GUI Manager software will be installed.

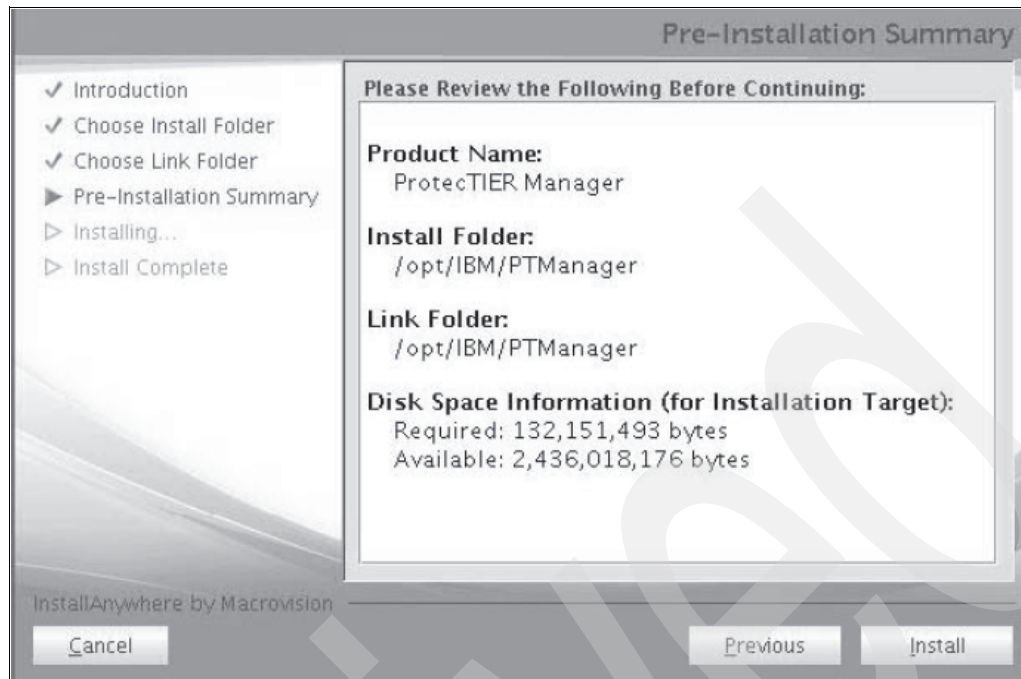


Figure 6-17 TSSC PT Manager GUI installation

8. Click **Install** to continue. You see an Installing ProtectTIER Manager panel (not shown).
9. When the installation successfully completes, you see an Install Complete panel (not shown). Click **Done**.
10. Press Enter when The ProtectTIER Manager GUI is now installed is displayed.
11. Remove the CD from the TSSC CD drive tray.

6.6 Configuring ProtectTIER using ptconfig

This section provides instructions for using the ptconfig utility to complete the ProtectTIER configuration process that was started in manufacturing.

Verify that the ProtectTIER Manager is installed on the TSSC.

To launch the ProtectTIER Manager on the TSSC, right-click the TSSC blue panel and select **Browser Functions** → **ProtectTIER Manager Functions** → **Launch GUI**.

If you need to install the ProtectTIER Manager on the TSSC, refer to 6.5, “Installing the ProtectTIER Manager GUI on external TSSC” on page 107 and return here.

Note: Throughout this chapter the terms server and node are used interchangeably, depending upon the task being performed.

Prerequisites

Before using ptconfig to configure ProtectTIER, verify that the following prerequisites and conditions have been met:

- ▶ Cabling and physical connections (including those to the client’s local area network (LAN) using assigned IP addresses) are complete.

- ▶ RAID groups have been created. This task occurs outside the scope of the gateway installation process. Verify that this has been done.
- ▶ Upper ProtecTIER server is powered off.
- ▶ You have the following server information:
 - External IP address
 - External, fully-qualified, host name (*hostname.domain.com*)
 - External default gateway address
 - External network subnet mask

Note: The host name is case sensitive. Be aware of the use of upper and lowercase characters when gathering the information and when entering it into the system.

6.6.1 Logging in to the server

1. Using the KVM switch, press the PrtSc key on the TSSC's keyboard, and select **LowerNode** from the list of devices.
2. Perform the following in this clustered configuration:
 - a. Verify that Upper Server is powered off:
 - If it is off, go on to step 3 on page 110.
 - If it is running, power it off now. To do so:
 - i. Using the KVM switch, press the PrtSc key on the TSSC's keyboard, and select **Upper Node** from the list of devices.
 - ii. At the server command line, enter the command **poweroff**.
When the server is completely powered off, the monitor goes blank and the green Power LED on the front panel of Upper Server flashes.
 - b. Verify that Lower Server is running:
 - If it is, go on to step 3.
 - If it is not running, power it on now, wait for the boot cycle to complete, and then go on to step 3.
3. At the localhost login: prompt, log in with the ID root and the password admin.
4. Go on to Configuring Lower Server.

6.6.2 Configuring Lower ProtecTIER Server

Perform the following steps on Lower Server in a clustered configuration:

1. At the server's command prompt, change to the `/opt/dtc/install` directory. To do so, enter the command **cd /opt/dtc/install**.
2. Update the network settings by typing the command **./ptconfig -updateNetwork**.
3. If asked if you would like to stop the VTFD service, type yes and press Enter.
4. You are then prompted, one at a time, to enter the following values. At each prompt, type the requested value and then press Enter.
 - Customer Network IP address
 - Customer Network netmask

- Customer Network default gateway
- Customer Network hostname

Note: The hostname should be free of extensions or .com suffixes, for example: server9000.

After you enter the hostname, the system automatically starts the network configuration process. See Figure 6-18 for the status messages display.

```
Configuring network [ Done ]
Updated network configuration successfully
update updateNetwork ended successfully
```

Figure 6-18 Configuring network

The system automatically restarts the VTFD service, and you are returned to the command prompt.

5. Go to the /opt/dtc/app/sbin directory by typing the command `cd /opt/dtc/app/sbin`.
6. Type the following command to create file systems: `./fsCreate -n`.

A message is displayed stating that any existing data on the back-end storage will be removed. At the prompt, type data loss to continue.

Note: The message Filesystems successfully mounted will appear upon a successful completion. A failed message will appear if there is a problem. Do *not* disrupt the script until a success or failure appears. The system may take up to an hour, depending on the size and configuration of the storage.

7. Continue to the next section to add Lower Server to the ProtecTIER Manager and to configure the repository.

6.7 Adding a node and creating a repository

Adding a node registers the node's IP address and port number with the instance of ProtecTIER Manager running on the ProtecTIER Manager workstation or TSSC.

You must add the node before you can create repositories and file systems.

If you need to install the ProtecTIER Manager on the TSSC, refer to 6.5, "Installing the ProtecTIER Manager GUI on external TSSC" on page 107 and return here.

1. Launch the ProtecTIER Manager.

To launch the ProtecTIER Manager on the TSSC, right-click the TSSC blue panel and select **Browser Functions** → **ProtecTIER Manager Functions** → **Launch GUI**.

On a Windows-based ProtecTIER Manager workstation, click **Start** → **Programs** → **IBM** → **ProtecTIER Manager 2.4.x** → **IBM ProtecTIER Manager**.

On a Linux-based ProtecTIER Manager workstation, double-click the **PT Manager** icon found on the Linux desktop, or in the shortcut location you specified during installation.

The ProtecTIER Manager panel opens. See Figure 6-19.

Any configured networked systems are listed along with their IP addresses in the Nodes section of the Navigation pane, and are available for login.

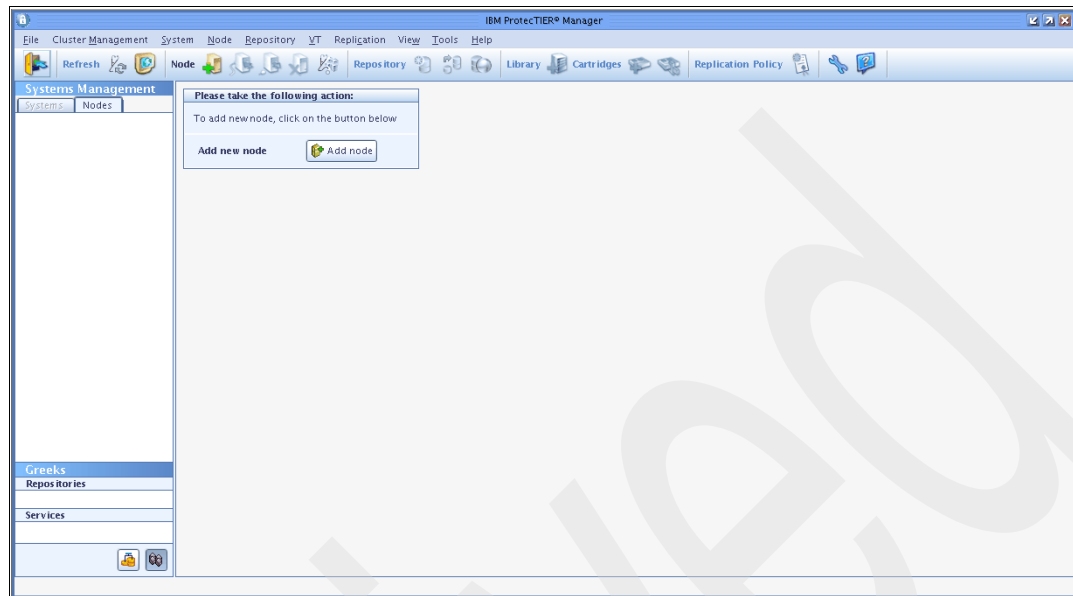


Figure 6-19 PT Manager

2. On the ProtecTIER Manager toolbar, click **Add node**; see Figure 6-20.

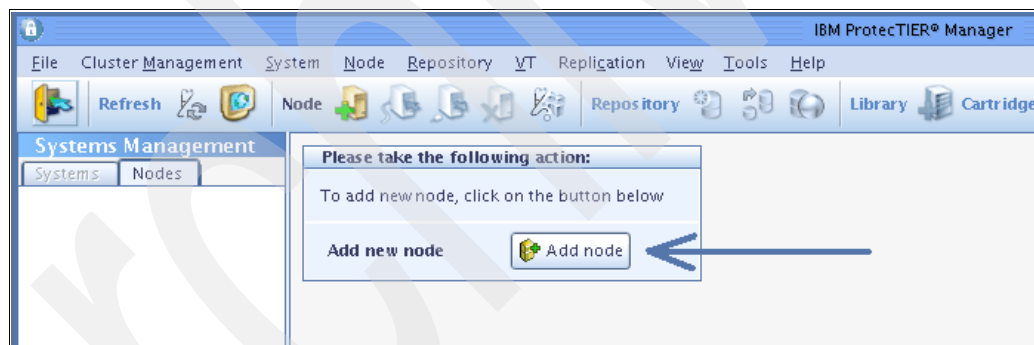


Figure 6-20 PT Manager Add node

The Add node dialog displays and prompts you for the IP address and port number of the node you want to add. See Figure 6-21 on page 112.




Figure 6-21 PT Manager_Add node_IP

3. Enter the external IP address of the node to be added, and then click **Ok**.

Note: Do not change the port number or the ping port number of the node unless directed to do so by IBM Support.

The node appears in the Nodes pane left side and the Login button displays in the View pane. See Figure 6-22.

The Success dialog displays and provides information about the repository, which is created in an upcoming procedure. See Figure 6-22.

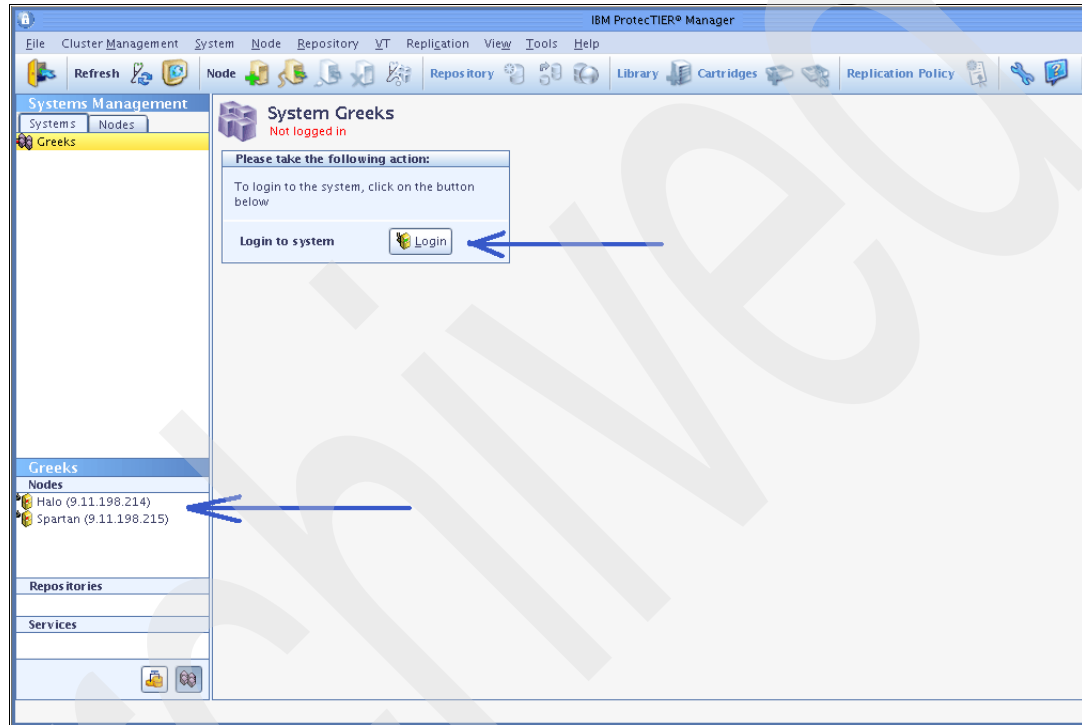


Figure 6-22 PT Manager_node added

4. Click **Login**. The login panel will prompt; see Figure 6-23 on page 113.




Figure 6-23 PT Manager_node login

Enter ptadmin for both the username and password, and then click **Ok**.

5. ProtecTIER Manager displays the information for the added node. If the node has an existing repository, the node's cluster displays on the Systems tab of the Navigation pane. If the cluster contains a second node, that node displays in the Nodes pane.

See Figure 6-24 for a reference of PT Manager system information just after login.

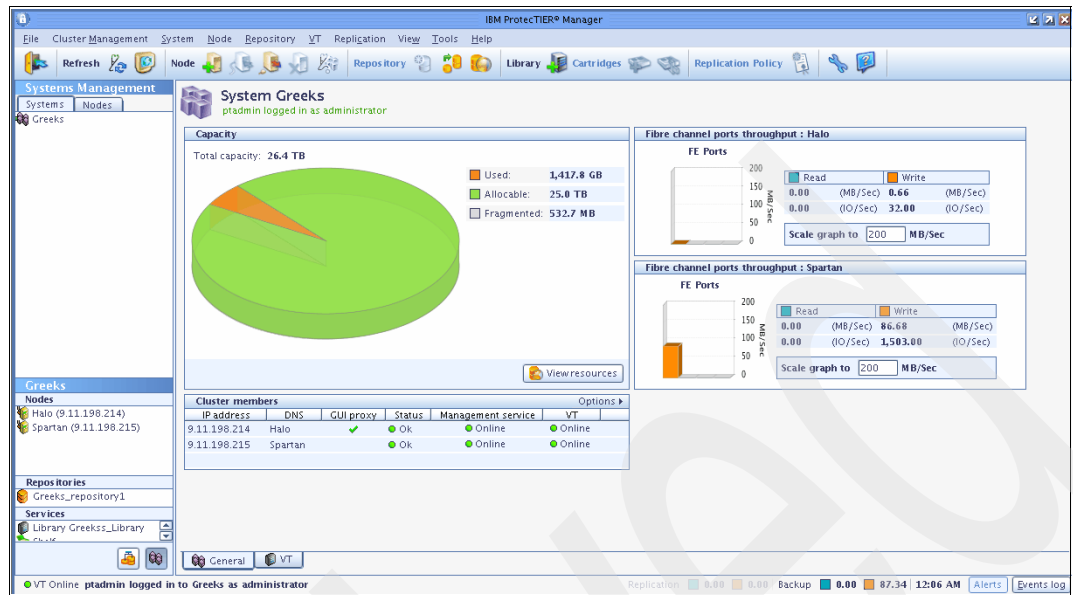


Figure 6-24 PT Manager _system info

6.7.1 Planning the repository

This topic describes how to create the Repository Planning from ProtecTIER Manager.

Use the Repository Planning wizard in conjunction with guidance from IBM Support to determine the optimum repository size and meta data file system arrangement for the repository.

1. On the ProtecTIER Manager workstation, run ProtecTIER Manager.
2. Select **Repository** → **Create repository planning**; see Figure 6-25 on page 114.




Figure 6-25 PT Manager _Create Repository

The “Create repository planning” wizard opens. See Figure 6-26.

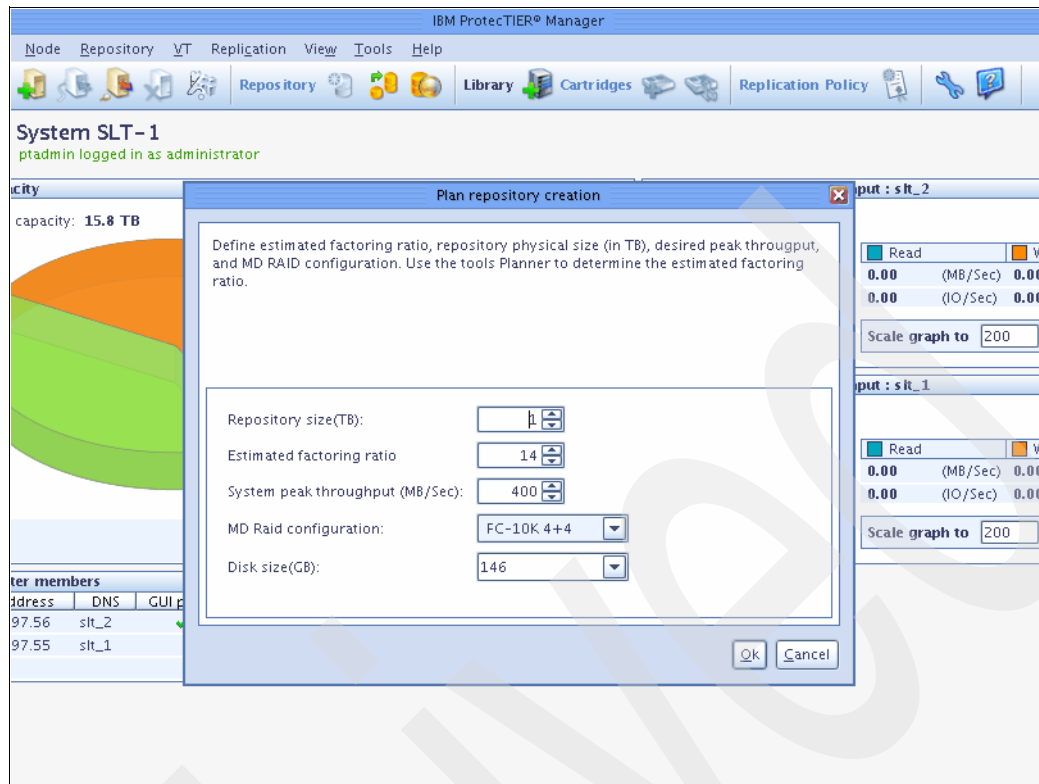


Figure 6-26 PT Manager_Repository wizard

3. In the “Estimated factoring ratio” field, enter the value estimated for the environment based on the data change rate, tape processing policies, and retention period.
4. In the Repository size field, select the size of the repository you want to create.

Note: The maximum possible repository physical size is 1 petabyte (1PB).

5. In the “System peak throughput” field, specify the rate of system peak throughput that the meta data file systems can support.
6. In the “MD Raid configuration” field, select the RAID configuration of the logical volumes on which the repository meta data file systems are to be created.

For example, select FC-10K 2+2 for a configuration of RAID 10 2+2 with fibre channel 10 KRPM disks.

7. Click **Ok**.

Once created, you can check the total size and usage of repository resources in every file system by selecting **View resources from the** toolbar; see Figure 6-27.

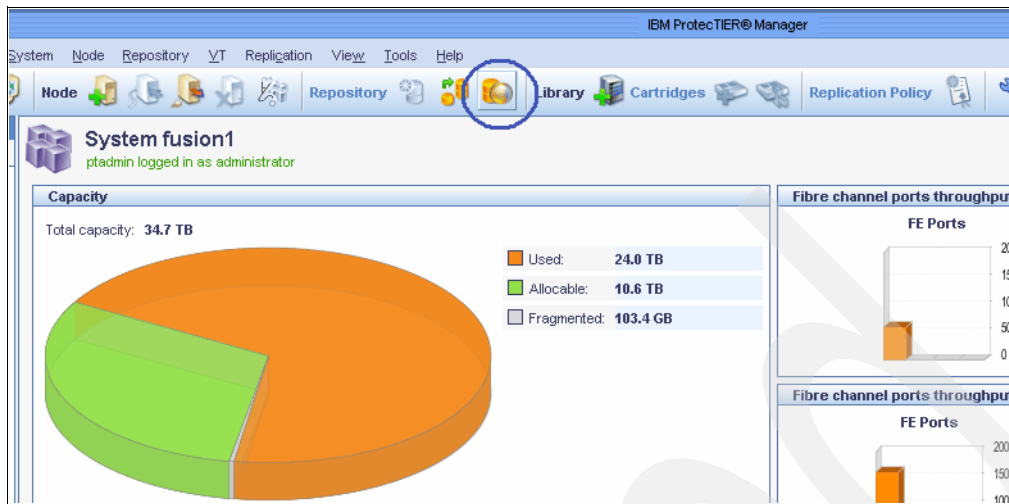


Figure 6-27 PT Manager_View resources

The output of repository resources will be posted, as shown in Figure 6-28.

Storage resources

The following table includes the total size and usage of resources in every file system.

Selected node: fusion2 (9.11.197.33)

File systems on fusion2 (9.11.197.33)

Mount point	Device	FS type	Size	Usage	RAID type	Disk size(GB):	Options
/mnt/vg0-lv_vg0	/dev/mapper/vg0-lv_vg0	gfs	1,115.1 GB	User data			
/mnt/vg1-lv_vg1	/dev/mapper/vg1-lv_vg1	gfs	2,230.6 GB	Metadata	FC-15K 8+8	300	
/mnt/vg10-lv_vg10	/dev/mapper/vg10-lv_vg10	gfs	1,115.1 GB	User data			
/mnt/vg11-lv_vg11	/dev/mapper/vg11-lv_vg11	gfs	1,115.1 GB	User data			
/mnt/vg12-lv_vg12	/dev/mapper/vg12-lv_vg12	gfs	1,115.1 GB	User data			
/mnt/vg13-lv_vg13	/dev/mapper/vg13-lv_vg13	gfs	1,115.1 GB	User data			
/mnt/vg14-lv_vg14	/dev/mapper/vg14-lv_vg14	gfs	1,115.1 GB	User data			
/mnt/vg15-lv_vg15	/dev/mapper/vg15-lv_vg15	gfs	1,115.1 GB	User data			
/mnt/vg16-lv_vg16	/dev/mapper/vg16-lv_vg16	gfs	1,115.1 GB	User data			
/mnt/vg17-lv_vg17	/dev/mapper/vg17-lv_vg17	gfs	1,115.1 GB	User data			
/mnt/vg18-lv_vg18	/dev/mapper/vg18-lv_vg18	gfs	1,115.1 GB	User data			
/mnt/vg19-lv_vg19	/dev/mapper/vg19-lv_vg19	gfs	1,115.1 GB	User data			
/mnt/vg2-lv_vg2	/dev/mapper/vg2-lv_vg2	gfs	1,115.1 GB	User data			
/mnt/vg20-lv_vg20	/dev/mapper/vg20-lv_vg20	gfs	1,115.1 GB	User data			
/mnt/vg21-lv_vg21	/dev/mapper/vg21-lv_vg21	gfs	1,115.1 GB	User data			
/mnt/vg22-lv_vg22	/dev/mapper/vg22-lv_vg22	gfs	1,115.1 GB	User data			
/mnt/vg23-lv_vg23	/dev/mapper/vg23-lv_vg23	gfs	1,115.1 GB	User data			
/mnt/vg24-lv_vg24	/dev/mapper/vg24-lv_vg24	gfs	1,115.1 GB	User data			
/mnt/vg25-lv_vg25	/dev/mapper/vg25-lv_vg25	gfs	1,115.1 GB	User data			
/mnt/vg26-lv_vg26	/dev/mapper/vg26-lv_vg26	gfs	1,115.1 GB	User data			
/mnt/vg27-lv_vg27	/dev/mapper/vg27-lv_vg27	gfs	1,115.1 GB	User data			
/mnt/vg28-lv_vg28	/dev/mapper/vg28-lv_vg28	gfs	1,115.1 GB	User data			
/mnt/vg29-lv_vg29	/dev/mapper/vg29-lv_vg29	gfs	1,115.1 GB	User data			

Ok

Figure 6-28 PT Manager_View Repository resources_output

- To print the information in the View Resources table, or to save the information as a .csv file, click **Options**; see Figure 6-29.

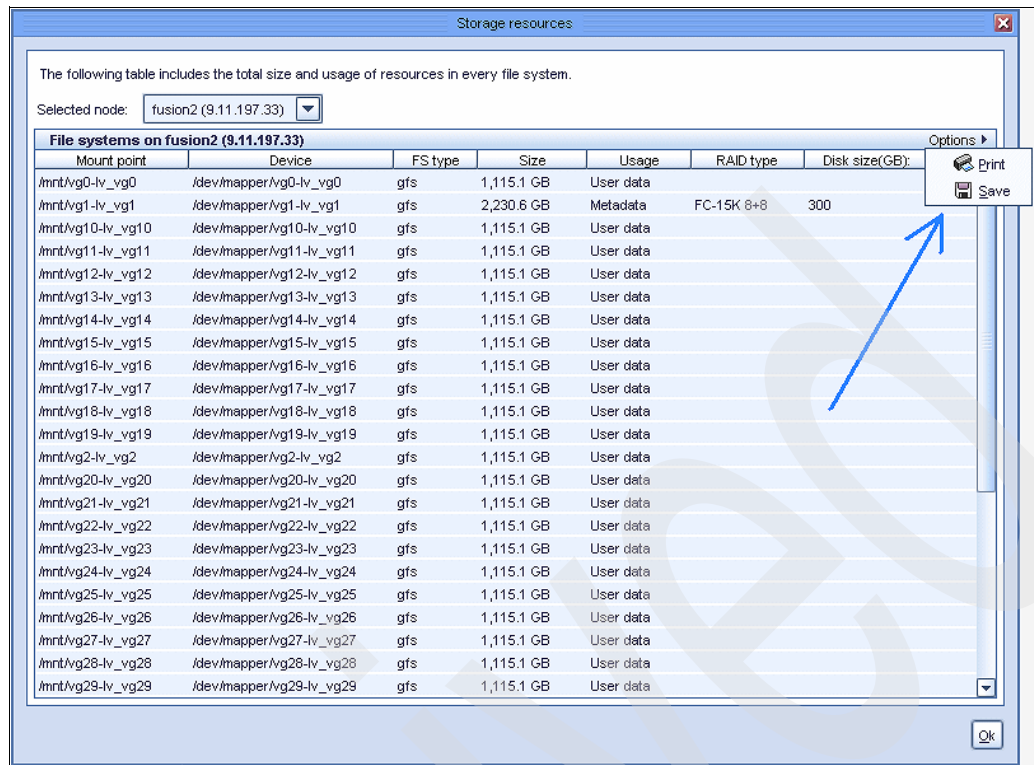


Figure 6-29 PT Manager_View Repository Resources_options

6.7.2 Creating the repository

After the necessary file systems have been created, use the information generated during the repository planning process to create the repository. 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, as repository creation *must* be performed on Lower Node.

1. On the ProtecTIER Manager workstation, if it is not already running, run the ProtecTIER Manager application. To do so:

Click **Start** → **Programs** → **IBM** → **ProtecTIER Manager 2.4.x** → **IBM ProtecTIER Manager 2.3.x**.

To launch the ProtecTIER Manager on the TSSC, right-click on the TSSC blue panel and select **Browser Functions** → **ProtecTIER Manager Functions** → **Launch GUI**.

The ProtecTIER Manager window opens. See Figure 6-30 on page 118.

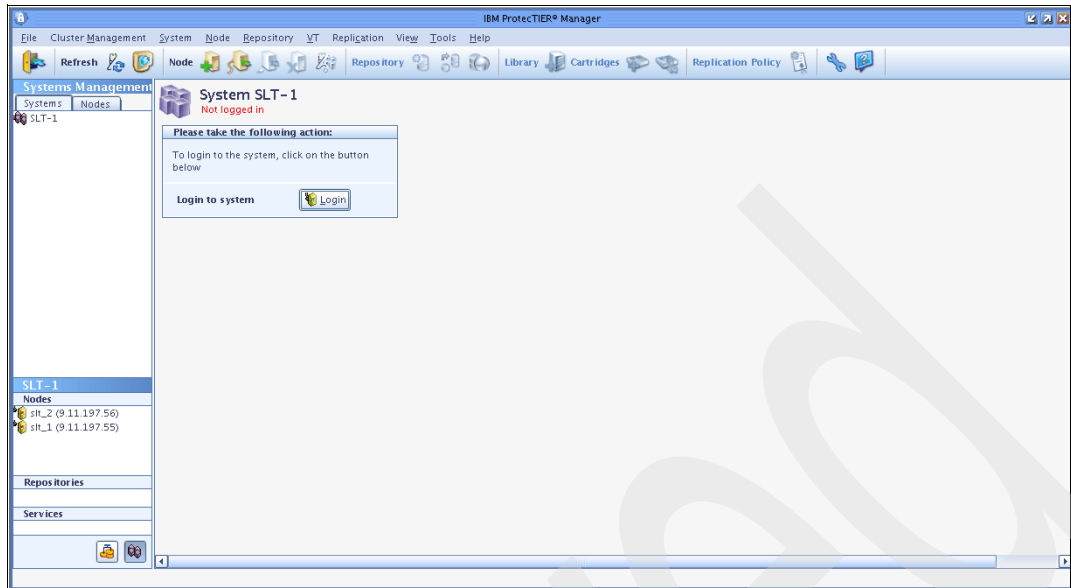


Figure 6-30 PT Manager_initial

2. In the Nodes pane, left side, select the node on which to create the repository.
3. Right-click on the node and select **Create repository**.

The Create repository wizard starts the data collection process. When data collection is complete, the Create repository wizard Welcome window opens. See Figure 6-31.

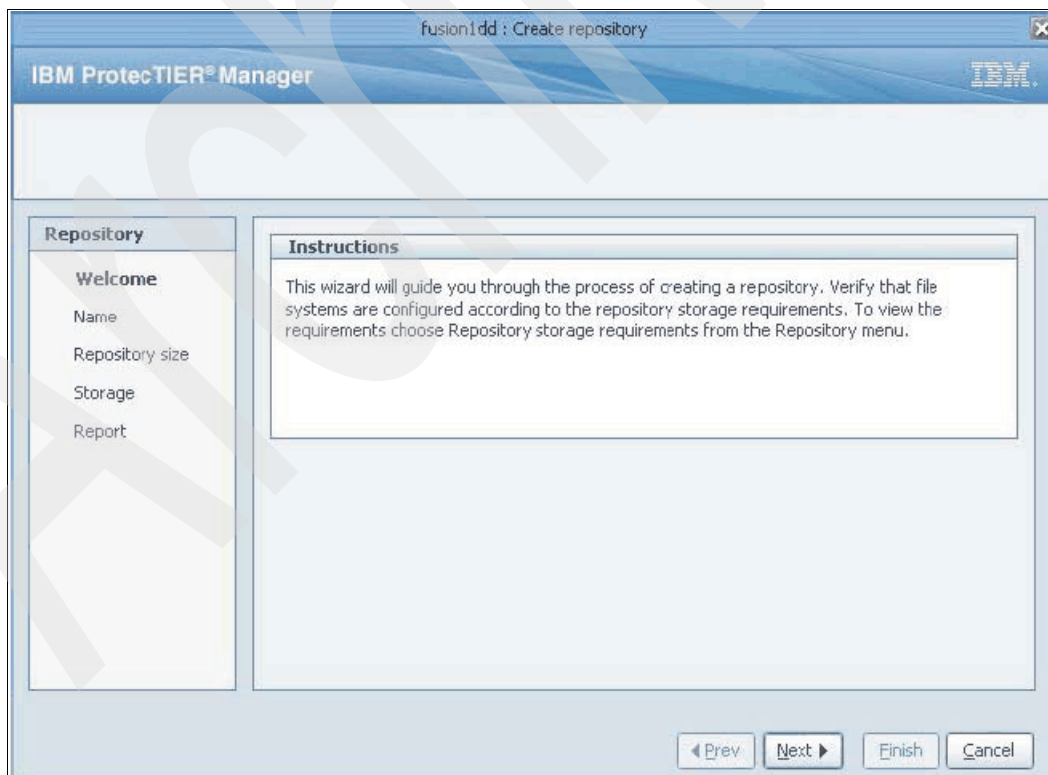


Figure 6-31 PT Manager_Create Repository_initial

4. Read the information on the Welcome panel, and then click **Next**. The Name panel opens. Enter the System name and click **Next**. See Figure 6-32.

fusion1dd : Create repository

IBM ProtecTIER® Manager

System and repository names should contain at least 4 characters.
System name is limited to 16 characters, and repository name is limited to 32 characters.

Repository

- ✓ Welcome
- Name**
- Repository size
- Storage
- Report

Name

System name: TS7680

Repository name: TS7680_repository1

◀ Prev Next ▶ Finish Cancel

Figure 6-32 PT Manager_Repository name

The Repository size panel opens. See Figure 6-33 on page 120.

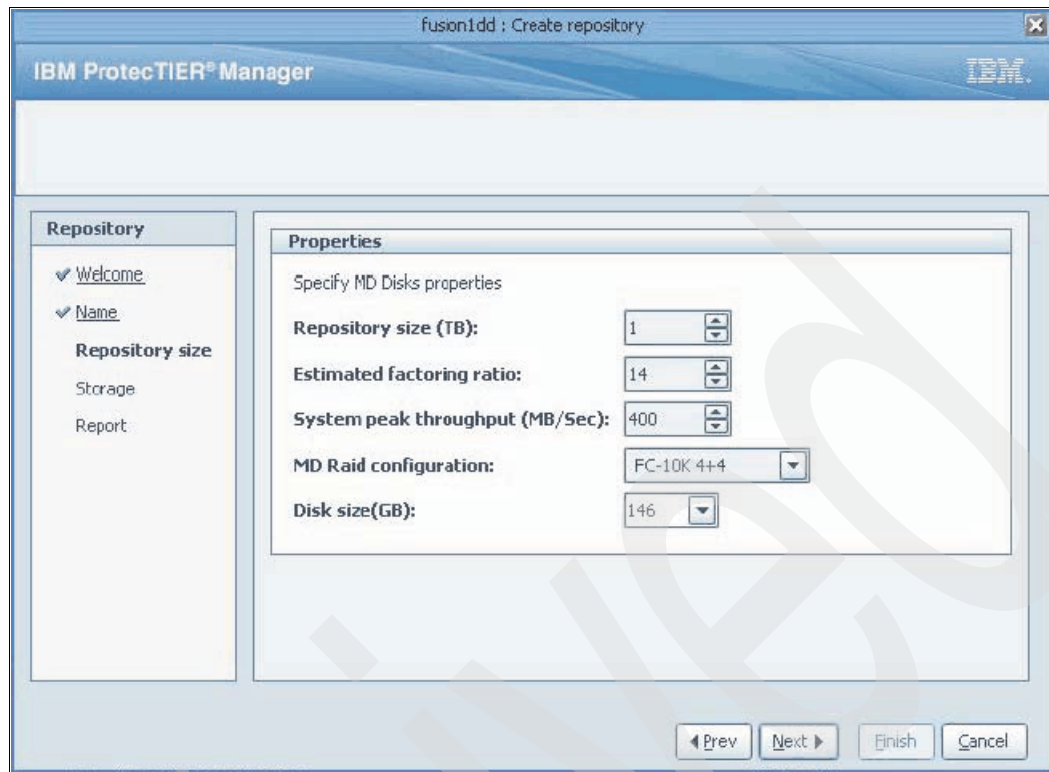


Figure 6-33 PT Manager_Repository size

5. In the Repository Size panel, provide the following information:
 - a. In the Repository size (TB) field, enter the repository size (in terabytes) that you determined using the Create repository planning wizard.
 - b. In the Estimated factoring ratio field, enter the estimated factoring ratio value that was determined with the assistance of IBM Support personnel.
 - c. In the System peak throughput (MB/Sec) field, specify the rate of system peak throughput that the meta data file systems can support.
 - d. In the MD Raid configuration field, select the RAID configuration of the logical volumes on which the repository meta data file systems are to be created.
For example, select FC-10K 2+2 for a configuration of RAID 10 2+2 with Fibre Channel 10K rpm disks.
 - e. Disk Size (GB) - Select disk size to match those in the array.
6. Click **Next**. The Repository resources dialog opens. See Figure 6-34 on page 121.

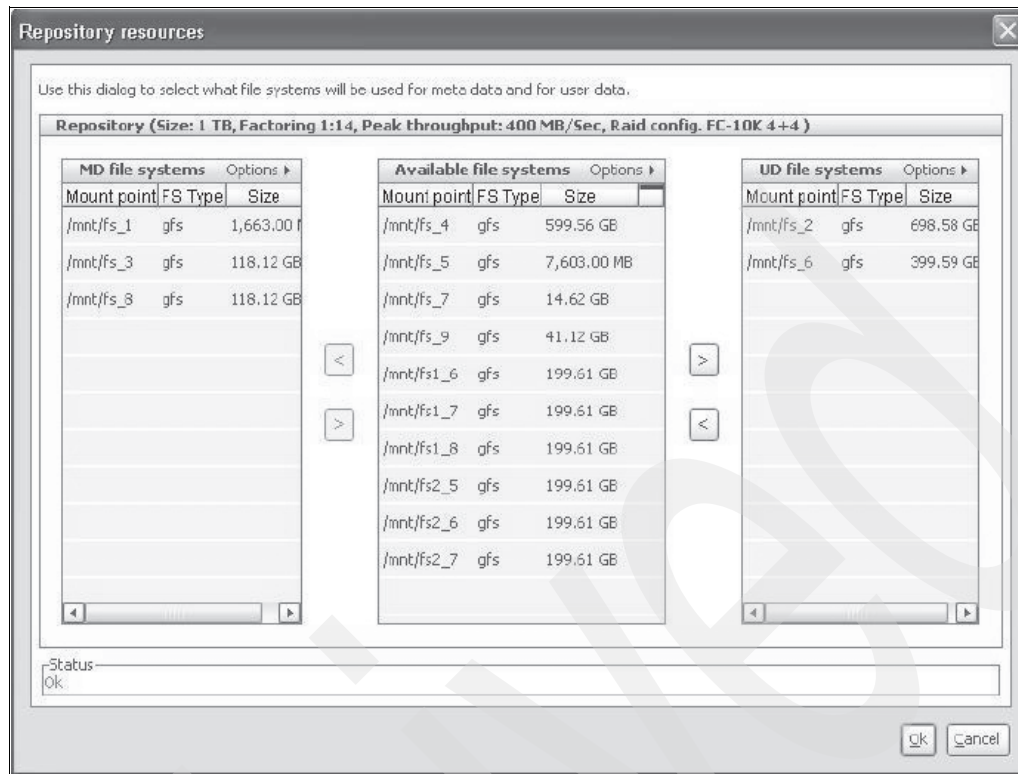


Figure 6-34 PT Manager_Repository resources

7. Verify that the correct file systems are selected for meta data and user data, based on the meta data file system sizes indicated by the repository planning process.

Note: By default, the ProtecTIER system generally selects the smallest available file systems for use as meta data file systems. The remaining file systems are available for user data. Storage space for user data cannot exceed the repository size defined in the Repository Size panel.

8. If the file systems selected by ProtecTIER for meta data and user data do not match the file systems created for those purposes, change the assignment. To do so:
 - a. From the Available file systems list, select **File Systems**.
 - b. Use the left and right arrows to move file systems to and from the MD file systems (meta data) and UD file systems (user data) lists.
9. Click **Ok**.

The Repository resources dialog closes.

10. The Meta data window opens. See Figure 6-35 on page 122. On this panel, the “Allocated meta data size” field displays the amount of disk space allocated for meta data, and the “Allocated user data size” field displays the amount of disk space allocated for user data, based on the estimated factoring ratio and the set of existing file systems.

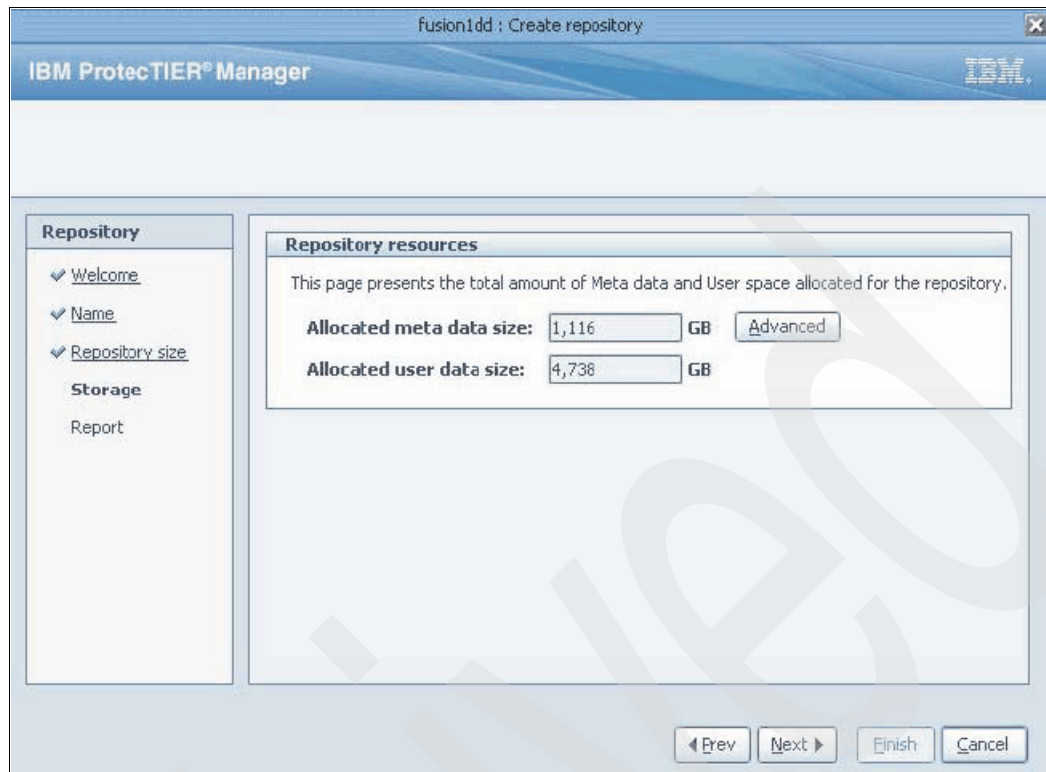


Figure 6-35 PT Manager_Meta data

11. Click **Next** to view the Report, then click **Finish**.

12. The Confirm operation panel appears. See Figure 6-36. Click **Yes** to continue.

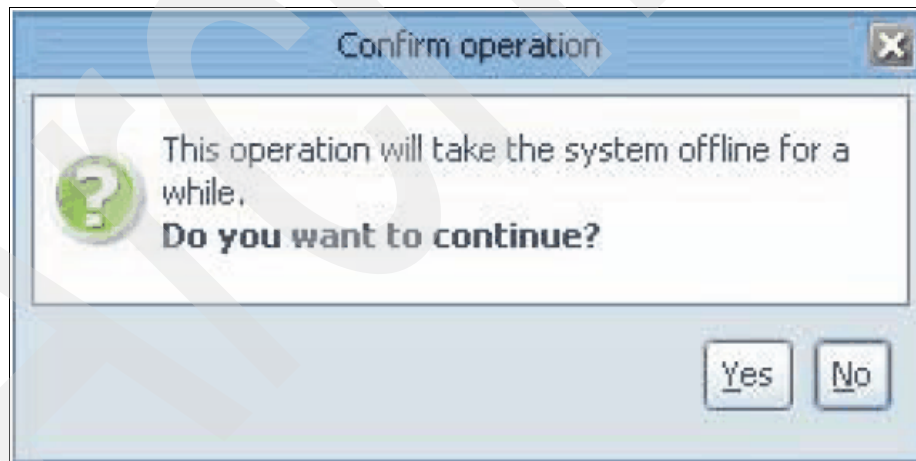


Figure 6-36 PT Manager_Meta data confirm

13. A Run in Background dialog panel appears. Do not run in Background.

Note: Do not continue until the Repository creation has completed. The time required for the Repository creation varies based on the amount of back-end disk storage.

6.8 Configuring Upper PT Server, creating a library and port configuration

This topic describes how to configure the Upper ProtecTIER server, which is the same procedure used for the Lower server. Also, we show you how to add the Upper server to the cluster (using the PT Manager), create a library and do port configuration.

Note: Ignore file system errors on boot-up.

1. Using the KVM switch, press the PrtSc key on the TSSC's keyboard, then select **Upper Node** from the list of devices.
2. Power on the Upper ProtecTIER server.
To power on the ProtecTIER servers, press the white, recessed power-control button on the ProtecTIER server operator panel. See Figure 6-2 on page 94.
 - When the panel displays registered calypsa with major device 249, press Enter.
 - Login input should be displayed.
3. When the localhost login: prompt appears, log in as **root** and type the password **admin**.
4. At the server's command prompt, change to the `/opt/dtc/install` directory. To do so, enter the command `cd /opt/dtc/install`.
5. Update the network settings by typing the command `./ptconfig -updateNetwork`.

Note: In case of any failure, enter the command **reboot**, wait until the login panel appears and reenter the above command.

6. If asked if you would like to stop the VTFD service, type yes and press Enter.
7. You are then prompted, one at a time, to enter the following values. At each prompt, type the requested value and then press Enter.
 - Customer Network IP address
 - Customer Network netmask
 - Customer Network default gateway
 - Customer Network hostname

Note: This is the name provided by the client to identify this server on their local network. The hostname should be free of extensions or .com suffixes, for example: server9000.

After you enter the hostname, the system automatically starts the network configuration process. See Figure 6-37 for the status messages displayed.

```
Configuring network [ Done ]
Updated network configuration successfully
update updateNetwork ended successfully
```

Figure 6-37 Configuring network

The system automatically restarts the VTFD service, and you are returned to the command prompt.

8. Go to the `/opt/dtc/app/sbin` directory by typing the command `cd /opt/dtc/app/sbin`.
9. Type the command `./fsCreate -t` to duplicate the mount points on Upper Server

6.8.1 Add Upper Server to the cluster (using the ProtecTIER Manager)

1. Launch the ProtecTIER Manager, if not already open.

To launch the ProtecTIER Manager on the TSSC, right-click on the TSSC blue panel and select **Browser Functions** → **ProtecTIER Manager Functions** → **Launch GUI**.

On a Windows-based ProtecTIER Manager workstation, click **Start** → **Programs** → **IBM** → **ProtecTIER Manager 2.4.x** → **IBM ProtecTIER Manager**.

On a Linux-based ProtecTIER Manager workstation, double-click the **PT Manager** icon found on the Linux desktop, or in the shortcut location you specified during installation.

2. Click **Login**.
3. When prompted for login information, enter `ptadmin` for both the username and password, and then click **Ok**.
4. On the ProtecTIER Manager toolbar, click **Cluster Management**, then select **Add cluster member**. See Figure 6-38.




Figure 6-38 PT Manager_add cluster

5. The Add Cluster member wizard window will appear. See Figure 6-39 on page 125.

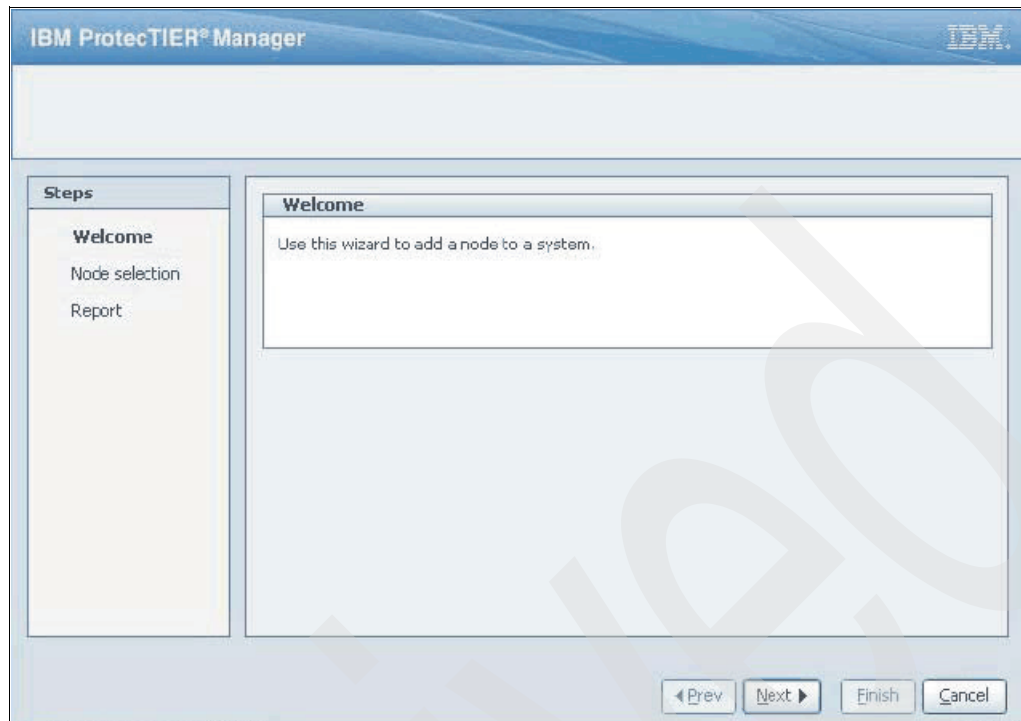


Figure 6-39 PT Manager_add cluster wizard

6. Select **Next** to start the process; the panel shown in Figure 6-40 will appear.




Figure 6-40 PT Manager_Node selection

7. In the Node Selection pane enter the Node IP of Upper Server, then select **Next**.
8. The report panel will appear. Select **Finish**.

Note: Wait 5 minutes for the node to come online.

6.8.2 Create a library using ProtecTIER Manager

In this topic we explain how to create a library on a configured TS7680 using the ProtecTIER Manager.

The 3958-DE2 (TS7680) is a two-node system. You can create only one library with a fixed configuration that contains the following values:

- ▶ 2 cluster members
- ▶ 4 FE ports on each node
- ▶ 128 tape drives assigned to each node (divided equally between the ports)
- ▶ 1,000,000 slots
- ▶ 0 Import and Export slots

No cartridges are created when the library is created. However, you can add up to one million cartridges at a later stage by performing the add cartridges task.

Perform the following steps to create the library:

1. Log into the ProtecTIER Manager.
2. In the Systems pane, select a cluster on which to add a library.
3. From the Toolbar, click **Create new library**; see Figure 6-41.



Figure 6-41 PT Manager_add library

The Create new library wizard Welcome panel is displayed.

4. Click **Next**. The Library details panel is displayed.
5. In the VT name field, enter a name for the library.
6. Click **Next** and **Finish**. The Create new library wizard closes and the ProtecTIER system temporarily goes offline to create the library. The library is displayed in the Services pane and the VT monitoring panel is displayed.

Note: This will take from fifteen minutes to about an hour to create, depending on the size of your repository. Wait for the operation to complete before continuing.

6.8.3 Port configuration

This topic guides you on how to verify and set up the ProtecTIER Servers port configuration.

1. Use the ProtecTIER Manager GUI to view or change the attributes.

- a. Select the **node** and right-click it. A menu will appear with the node options. Select **Port attributes** from this menu. See Figure 6-42.

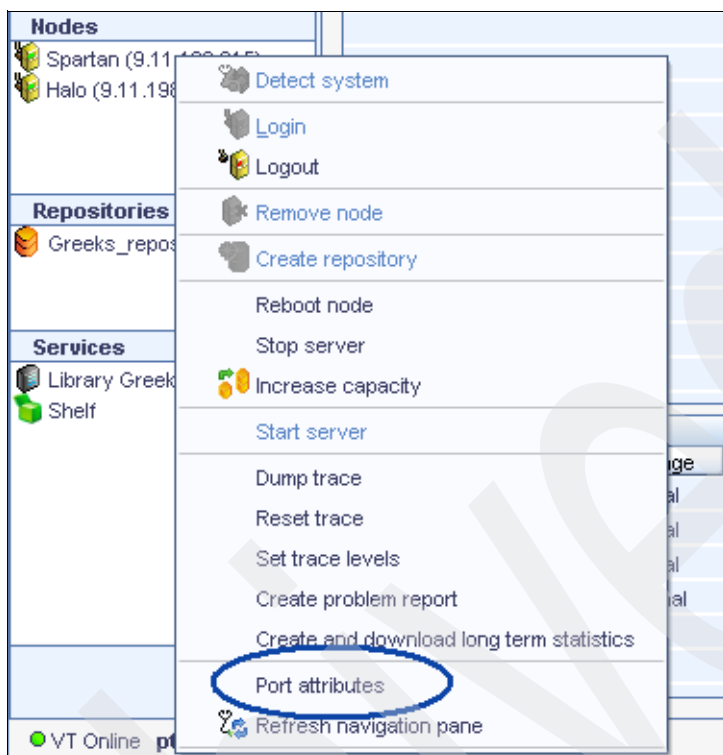


Figure 6-42 PT Manager_Port attributes

- b. A Port Attributes wizard window appears. See Figure 6-43 on page 128.

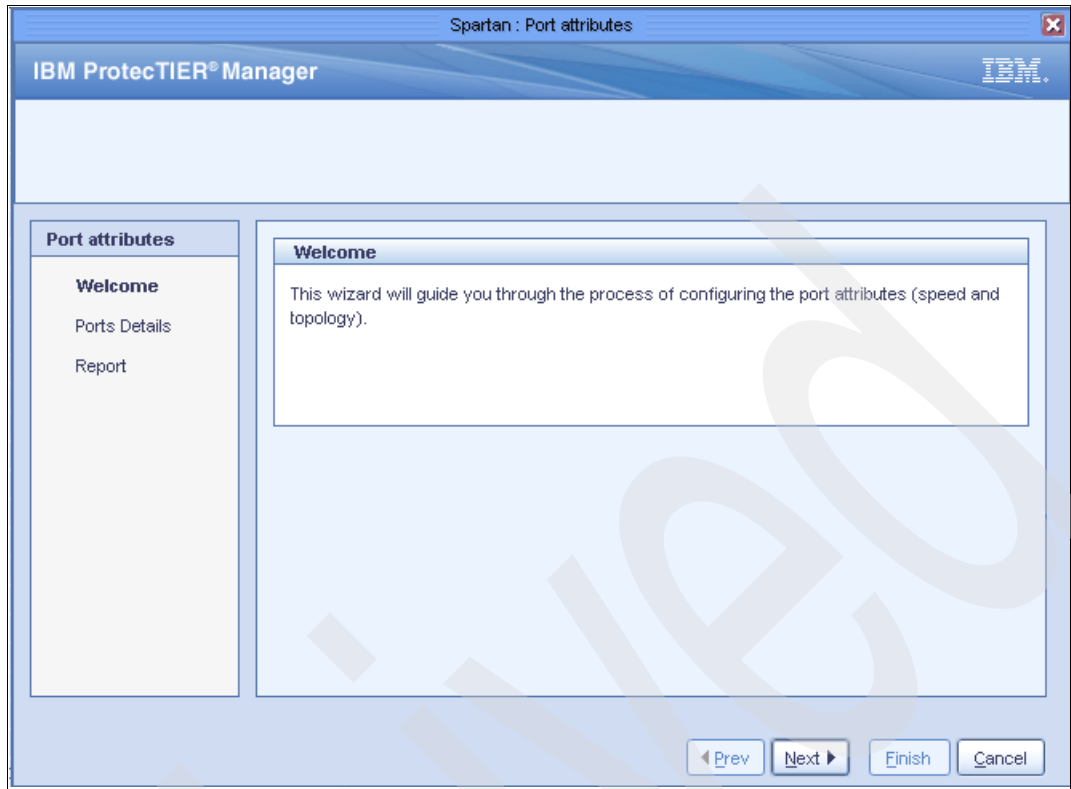


Figure 6-43 PT Manager_Port attribute wizard

- c. Select **Next** to proceed to the Port Details menu. Verify that the topology from the default setting is LOOP. See Figure 6-44 on page 129.

Note: The ports need to be configured as LOOP in order for the Enterprise controllers to configure the devices.

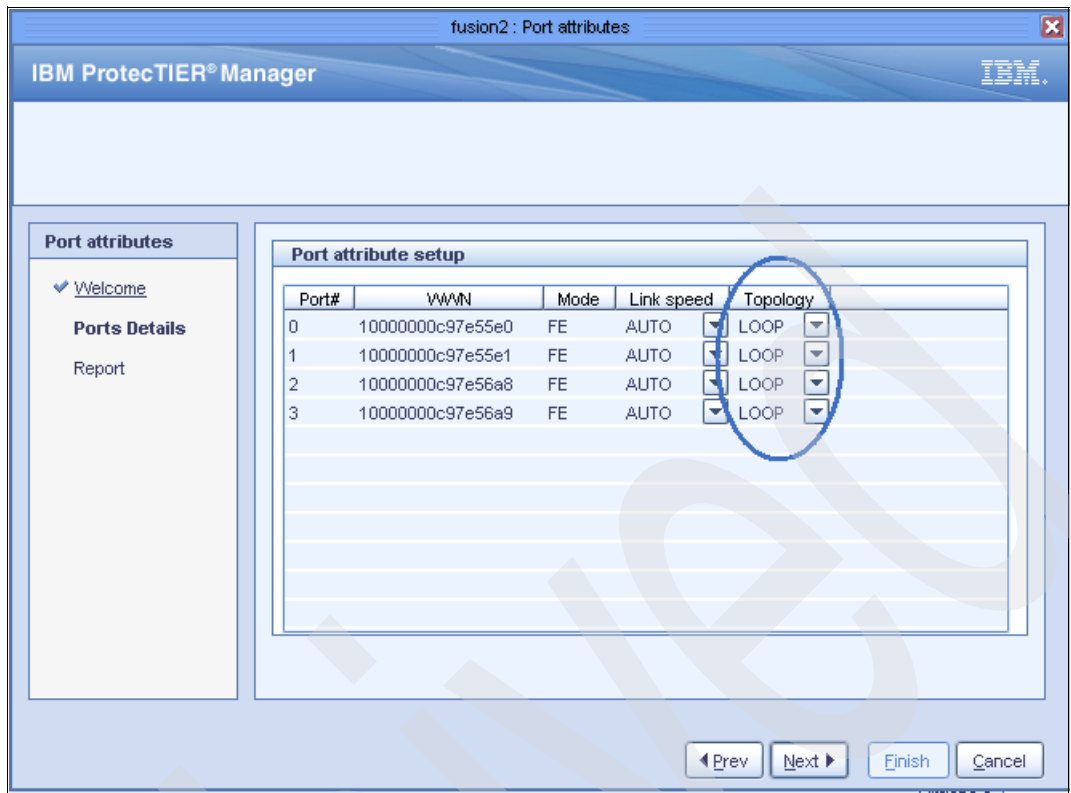


Figure 6-44 PT Manager_Port details

2. Press **Finish** to complete the topology change. The changes will take a few minutes to update.

Note: The ports will need to be changed on both Upper and Lower Server.

6.9 Testing a clustered system configuration

After using `ptconfig` to successfully configure both nodes in a clustered configuration, run the system verification tests. The tests perform checks on the Red Hat cluster, including fencing the other node, and wait for the system to reboot.

Important: Run the test on both nodes, starting with Lower Node. Verify that the DVD drive is empty.

1. Using the KVM switch, press the `PrtSc` key on the TSSC's keyboard, then select **Lower Node** from the list of devices.

Note: Upper and Lower Node names depend on what you called them during KVM setup.

2. Log in to Node A with the username `root` and password `admin`.
3. At the command prompt, enter the command `cd /opt/dtc/install`.
4. From the command prompt, enter the command `./ptconfig -validate`.

Note: During the validation process, you will perform a fence test. When prompted to continue with the fence test, type `fence test` and press `Enter`. When prompted to stop the VTFD service, type `yes` and press `Enter`. The fencing process will restart the remote node as the test progresses. This is to be expected because the clustering capabilities of both nodes are tested. Success is indicated by `Validation Ended`.

5. After the test finishes on Lower Node, using the KVM switch press the `PrtSc` key on the TSSC's keyboard, then select **Upper Node** from the list of devices.
6. Log into Upper Node with the username `root` and password `admin`.
7. Repeat steps 3 and 4 on Upper Node.

Note: Alerts may be generated as the validation tests are running. The ProtecTIER Manager Alerts Log allows you to monitor the alerts as they occur. To do so, on the ProtecTIER Manager workstation, click **Alerts** in the lower right corner of the ProtecTIER Manager window. The Alerts Log opens. When you are finished reviewing alerts, click **Clear Alerts** in the lower right corner of the Alerts Log panel.

8. If the test succeeds on both nodes, system validation is complete.
9. If the test fails on either node, verify that all cluster links except to the NPS are detected and operating at 1000 MB link speed. If it still fails, contact your next level of support to resolve the failure conditions.

6.10 ProtecTIER Library attachment to the Enterprise controllers

This topic explains how to verify and configure the Enterprise controllers to attach all 3958 ProtecTIER devices.

The 3958-DE2 has a keyboard monitor tray that pulls out for use. Use the PrtSc key to select either the Upper or Lower ProtectTIER server or the TSSC if it is in the frame.

The TSSC is used for connection to the Enterprise controller to open a terminal window and Telnet to 172.31.1.xx2 for the lower controller or 172.31.1.xx7 for the upper controller. (Factory default settings are 242 and 247.)

1. Perform the following to authenticate locally and Telnet to 172.31.1.xx2 for the lower controller and 172.31.1.xx7 for the upper controller.
 - a. Using the KVM switch, press the PrtSc key on the TSSC's keyboard. Select **TSSC** from the list of devices.
 - b. Right-click on the TSSC blue panel.
 - c. Select **Terminal**, or select **System Console Actions** → **Telnet to Tape System**, then select the system to log on to.
 - d. Telnet to 172.31.1.xx2 for the lower controller and 172.31.1.xx7 for the upper controller. See Figure 6-45 as example.

```
Base Level: 1.23.1.125
AIX Level: 5.3.9.0
SERIAL: 78-00C6733

IBM 3592-C06 CONTROL UNIT

Service Terminal

Available logins on [3590cu10]:

Enhanced
Support
Service

Authentication Id

Please login:
```

Figure 6-45 Enterprise Controller_login

- e. Type the userid used earlier or Service and press Enter.
- f. Use the password saved earlier if available or at the password request for Service, type in the case-sensitive password (ibm2serv). You now must authenticate locally. Skip to step n on page 133 if using an existing userid and password.
- g. Select option **1) TSSC Login (console launcher)**. See Figure 6-46 on page 132.

```

Select the method you are using to access this system:
1) TSSC Login (console launcher)
2) TSSC Login (serial connection)
3) EBTERM/NETTERM Login (serial connection)
4) User selectable terminal type (i.e. vt100, vt102, etc.)

NOTE: The user selectable terminal type may result in service menus
      and other displays that do not show up on your screen correctly

Enter your selection (1-4) and press <ENTER>:

```

Figure 6-46 Enterprise Controller_console launcher

- h. The following window will appear. See Figure 6-47 and just type Enter to continue.

```

Setting terminal type to: xterm
Pinging TSMC IP 172.31.1.1
/*****/
/*
/* Info:      Call Home will be disabled for 120 minutes
/*
/*****/
/*
/*          Press <enter> to continue
/*
/*****/

```

Figure 6-47 Enterprise Controller_console launcher_continue

- i. At the Authenticate User Login panel, type your Authentication ID in the Enter User ID box. See Figure 6-48 as an example.

Authenticate User Login	
Type or select values in entry fields. Press Enter AFTER making all desired changes.	
Enter User ID. This id should match your DFS login ID EXACTLY in order for remote authentication to run correctly.	[Entry Fields] [luciano]
Enter login type (Remote or Local)	Local
NOTICE: If this system does not have at least one ethernet connection running, you will be unable to login in when prompted following the authentication. To log in, you must exit out of the authentication completely and return to the login herald. From there enter the user id used above and your authenticated password in order to access the system.	

Figure 6-48 Enterprise Controller_Authenticate login

- j. An access code appears in the controller display panel.
- k. Type the access code into the space requesting the code. See Figure 6-49.

```

Type 3592-C06 SN: 78-00C5003
Mon 12 Day 04 Year 09
Creating user 1

User ID 1 expires on: 12/04/09 at 16:32

Please note: The access code is not case sensitive but the password is!

The access code is displayed on the Op Panel LED.
Type the access code and press Enter.

```

Figure 6-49 Enterprise Controller_access code

- l. A case-sensitive password will appear on top of the panel; see Figure 6-50 as an example. This can be manually noted or copied onto your clipboard. (Write it down. This password is good for 24 hours.)

```

Welcome luciano! Please login as luciano with password: zasn01fg

Do you have a locally generated password or a password from
an authorized decryption site, AND want to login now?

(y or n and <ENTER>)

```

Figure 6-50 Enterprise Controller_login password

- m. Type y and Enter, then type the password received in the above step and Enter. You will again receive the Console launcher window; see Figure 6-51.

```

Select the method you are using to access this system:
1) TSSC Login (console launcher)
2) TSSC Login (serial connection)
3) EBTERM/NETTERM Login (serial connection)
4) User selectable terminal type (i.e. vt100, vt102, etc.)

NOTE: The user selectable terminal type may result in service menus
      and other displays that do not show up on your screen correctly

Enter your selection (1-4) and press <ENTER>:

```

Figure 6-51 Enterprise Controller_console launcher

- n. Select option **1) TSSC Login (console launcher)** and you will get the IBM Control Unit Subsystem Maintenance menu; see Figure 6-52 on page 134.

```
IBM Control Unit Subsystem Maintenance

Move cursor to desired item and press Enter.

System Checkout Menus
SIM, Error Log, Diagnostics, Trace/Dump Menus
Utility Menus
Microcode Maintenance Menus
Subsystem Configuration Menus
Removal/Replacement Menus
Control Unit Online/Offline Control
Control Unit Shutdown and Restart Menu
PFE/Support Tools Menus
Using SMIT (information only)
```

Figure 6-52 Enterprise Controller_Maintenance menu

2. Take the controller offline. Select **Control Unit Online/Offline Control** → **Vary Control Unit Offline**. Once the offline is complete, press PF3.
3. Perform initial 3958 configuration on each Enterprise Controller to attach all 3958 ProtecTIER devices. Start the lower controller first and delay starting the upper controller by five minutes.
 - a. Select **Control Unit Service Utilities** → **Utilities Menu** → **3958 Utilities Menus** → **Perform Initial 3958 Configuration**.

Note: The time this takes depends on the cache. It may be as quick as 1/2 hour or as long as 4 1/3 hours. Do not interrupt the process or it will require help from support.

- b. Select **Yes** and press Enter when asked if you are sure.
 - c. Each Fibre Channel interface attaches to 32 virtual drives. Verify that all 4 channel interfaces are correctly configured with 32 devices.
4. Verify that the controllers are online.
 - a. Press PF3 to go back to the **IBM Control Unit Subsystem Maintenance** menu.
 - b. Select **Control Unit Online/Offline** → **Display Current Control Unit Offline/Online Status**.
5. Verify that the system checks out successfully.
 - a. Press PF3 to go back to the **IBM Control Unit Subsystem Maintenance** menu.
 - b. Select **System Checkout Menus** → **Refresh / Rerun ALL System Checks**.
 - c. When the checks are finished, press PF3 to back up to System Checkout Menus. This could take up to 5 minutes.
 - d. Select **Display Current Status of all running/completed checks**. Verify that all checks pass OK.
 - e. If the test fails, contact your IBM System Service Representative to resolve the failure conditions.
6. Verify that the FICON host connections are present.
 - a. Press PF3 to go back to the Utilities Menu.
 - b. Select **Host Adapter Utility Menus** → **Display Control Unit, FICON, and Tape Device Status**.
 - c. Verify that all FICON Host links are online.

- d. If links fail, verify the host connection and then call your IBM System Service Representative to resolve the failure conditions.

Archived

Archived

Host attachment

In this chapter we describe how to implement the IBM System Storage TS7680 ProtecTIER Deduplication Gateway for System z on the Series z host system. The tasks you have to complete include the definitions of the hardware to the hosts through the Hardware Configuration Definition (HCD) dialog boxes and other, operating system-specific tasks depending on the System z software you are using.

The TS7680 Tape Library includes the Tape Enterprise Controller Library Manager and it enables the Open Systems Small Computer System Interface (SCSI) virtual medium changer library to be connected and used by System z-attached hosts. The Library Manager uses the existing system-managed tape library support so the definitions and System Managed Storage (SMS) constructs required for implementation of the TS7680 in the system-managed environment are the same as those for the existing TS3500 or 3494 Tape Library.

In a z/OS environment, Data Facility Storage Management Subsystem (DFSMS) provides automated system management of tape libraries. The interface between the System z host and the library robotics is managed by a Library Manager, which is included in the Licensed Internal Code of the Enterprise Tape Controllers of the TS7680. The virtual library robotics are assigned to all attached ports so it can take advantage of the IBM Control Path Failover (CPF) and Data Path Failover (DPF) features.

System z hosts attach to the controllers, which in turn have a set of virtual drives attached, rather than each drive having a separate attachment to a host channel. The IBM System Storage TS7680 ProtecTIER Deduplication Gateway is designed to deliver a high performance inline data deduplication solution on System z. Like the IBM System Storage TS7720 Virtualization Engine, the TS7680 provides a disk-only virtual tape solution. The TS7680 emulates an IBM tape library and D/T3592 Model J1A tape drives. The TS7680 appears to the host as an automated tape library (AL) and not as a virtual tape library (VL) and is displayed as a 3958-DE2 by the DISPLAY SMS,LIBRARY command. The TS7680 is defined to the host as any other IBM automated tape library; composite and distributed libraries do not apply to the TS7680.

For high availability, the TS7680 supports two nodes, with each node consisting of an enterprise tape control unit and deduplication engine. Each node supports up to 128 virtual tape devices (8 logical control units with 16 devices per logical control unit) and emulates the IBM System Storage 3592 Model J1A tape drive and JA (MEDIA5) media for a maximum

configuration supporting up to 256 devices. Even though the library supports up to 256 virtual tape devices, based on the size of the back-end disk repository, only a subset of the drives may be supported. Devices in the first/lower tape control unit are defined using subsystems X'01' through X'08', and devices in the second/upper tape control unit are defined using subsystems X'11' through X'18'. With the TS7680, one tape controller or deduplication engine can be offline (for code upgrade or repair) and, with a shared disk cache, all logical volumes are accessible through the remaining node's virtual device addresses.

The capacity of a TS7680 logical volume is less than a traditional JA (MEDIA5) physical cartridge (100 GB versus 300 GB). Performance scaling and performance segmentation, which are options with physical JA (MEDIA5) media, are not supported with the TS7680 and will be ignored. The host's compression setting is also ignored by the TS7680. The TS7680 will attempt compression regardless of the setting. Logical volumes defined to the TS7680 (through the IBM ProtecTIER Manager GUI) do not take up any disk space until they are written by the host. The TS7680 supports up to a million logical volumes.

Note: Even though 7-digit slot-values are being supported in the DISPLAY SMS,LIBRARY command output (TOTAL SLOTS and EMPTY SLOTS), ISMF and IDCAMS will still only support 6-digit slot- values. The slot counts reflect the maximum number of logical volumes that can be defined to the TS7680 (TOTAL SLOTS) and the number of logical volumes that can still be defined (EMPTY SLOTS).

The Enterprise Tape Controllers come with eight FICON links as front-end ports of the TS7680 to the System z host, each assigned to 32 virtual tape drives. The single virtual library of the TS7680 provides a fixed layout with 256 virtual tape drives and 1,000,000 cartridge slots.

7.1 TS7680 is fully integrated with z/OS software support

Because the host systems are, in general, unaware of the type of system-managed tape library that they use, no additional software support is *required* to attach a TS7680 Tape Library configuration to a System z host. See 7.6, “Cache management” on page 156 for more information on host and software support. Seamless integration with z/OS by the TS7680 is achieved because of the following:

- ▶ It is managed using system-managed tape (SMStape).
- ▶ No JCL changes required.
- ▶ Appears to the host as an automated tape library with 3592 Model J1A devices supporting MEDIA5 cartridges.
- ▶ No host application or tape management changes needed.
- ▶ Integration with the TS7680 during return to scratch processing to free back-end cache.
- ▶ Host alerts when cache starts to run low (warning and critical state notification).
- ▶ Single frame implementation is easy to install.
- ▶ ProtecTIER software PID is ordered as 5639-FPA.

Various registration and renewal models

7.2 Planning for software implementation

This topic defines the Input/Output (I/O) of the z/OS operating system used by the System z host and how to configure virtual tape drives to achieve load balancing.

Within each TS7680, there are two Enterprise Tape Controllers. Each Enterprise Tape Controller supports up to eight logical control units (or subsystems), each of which represents 16 tape devices for a total of 128 devices. The devices supported by the first Enterprise Tape Controller are defined using subsystem identifications X'01' through X'08', while the devices supported by the second Enterprise Tape Controller are defined using subsystem identifications X'11' through X'18'.

7.2.1 Software requirements

This topic defines the hosts supported by the TS7680.

- ▶ The 3958-DE2 (TS7680) supports attachment to the IBM System z host.
- ▶ The 3958-DE2 (TS7680) supports
 - z/OS V1R9 and above; refer to APAR OA27796.

Other versions of z/OS may be supported on an RPQ basis.

Installation of the host support is recommended for usability of the TS7680. Refer to the 3958 and 3952 Preventive Service Planning (PSP) buckets for the latest information on Software Maintenance. PSP buckets are kept on the IBM RETAIN system and are available to clients at:

<http://www14.software.ibm.com/webapp/set2/psearch/search?domain=psp>

- z/VM 5.3 and above including DFSMS/VTM FL221 with PTFs.

Refer to APAR VM64773 for additional information on the support being provided and the z/VM customer information center at:

<http://publib.boulder.ibm.com/infocenter/zvm/v5r3/index.jsp>

The System z host connects to the TS7680 using the 4-Gb FICON adapters. Each installed adapter supports 256 logical paths. FICON attachments between a System z host and the TS7680 can be single-mode long wave laser or multi-mode short wave laser.

7.2.2 z/OS software environments

System-managed tape allows you to manage tape volumes and tape libraries according to a set of policies that determine the kind of service to be given to the data sets on the volume.

The automatic class selection (ACS) routines process every new tape allocation in the system-managed storage (SMS) address space. The production ACS routines are stored in the active control data set (ACDS). These routines allocate to each volume a set of classes that reflect your installation's policies for the data on that volume. The ACS routines are invoked for every new allocation. Tape allocations are passed to the object access method (OAM), which uses its library control system (LCS) component to communicate with the Library Manager through the Asynchronous Operation Manager (AOM).

The Storage Class ACS routine determines whether a request is SMS-managed. If no Storage Class is assigned, the request is not SMS-managed, and allocation for non-specific mounts is made outside the tape library.

For SMS-managed requests, the Storage Group routine assigns the request to a Storage Group.

7.2.3 Host configuration definition

You must define the hardware to the host using the Hardware Configuration Definition (HCD) dialog boxes. We recommend that you use LIBRARY-ID and LIBPORT-ID.

LIBRARY-ID

In the configuration used with the IBM Virtualization Engine TS7680, each virtual device that is attached to a System z host reports back the same library sequence number, known as the *LIBRARY-ID*.

Subsystem identification

Each logical control unit, or 16-device group, must present a unique subsystem identification (ID) to the System z host. This ID is a 1-byte field that uniquely identifies each device associated with the logical control unit within the Enterprise Tape Controller. The value of this ID cannot be 0. The subsystem ID must match the LIBPORT-ID defined in the host Hardware Configuration Definition (HCD). Table 7-1 shows the required definitions of subsystem IDs.

Table 7-1 Subsystem identification

Enterprise Tape Controller	Logical control unit	Subsystem ID / LIBPORT-ID
lower	0-7	0x01 - 0x08
upper	0-7	0x11 - 0x18

Note: Even though the library supports up to 256 virtual tape devices, based on the size of the back-end disk repository, only a subset of the devices may be supported.

Load balancing

In order to balance the workload across the two Enterprise Tape Controllers, the virtual tape devices should be divided equally across the controllers. For example, if the size of your disk repository supports fewer than 256 devices (for example, 128 devices), define 64 devices using subsystems X'01' → X'04' and 64 devices using subsystems X'11' → X'14'. If a given host only needs a subset of the devices, spread the online devices in equal proportions across the two Enterprise Tape Controllers.

If an imbalance across the two Enterprise Tape Controllers is observed with one Enterprise Tape Controller getting the majority of the mounts, consider using the MVST[™] Device Allocation TAPELIB_PREF randomization option BYDEVICES. This is especially important if the imbalance is impacting the overall throughput rate of the TS7680. This randomization option is available on z/OS V1R11, on z/OS V1R8 through V1R10 with APAR OA26414. This should not only help with device randomization across libraries, but also within libraries.

7.2.4 Sharing a TS7680 Virtualization Engine

A FICON-attached TS7680 supports eight physical channels, each of which is capable of supporting 256 logical paths. Each logical path can address any of the 256 virtual devices in the TS7680.

We recommend that you use a FICON Director when connecting the TS7680 to more than one system.

The TS7680 places no limitations on the number of hosts that can use those channel paths, the types of hosts, or their operating system environments (as with any tape technologies that are supported in IBM tape libraries). An operating environment, however, through its implementation, does impose limits. z/OS DFSMS SCDS can support up to 32 systems or groups of systems.

Basically, anything that can be done with native drives in a physical Tape Library can be done with the virtual drives in a TS7680.

The TS7680 attaches to the host system or systems through eight FICON channels. Each FICON channel provides 256 logical paths that will result in a total of 2048 logical paths per TS7680.

You can use the following formula to calculate the number of logical paths required in an installation:

Number of logical paths = number of hosts \times number of CU \times number of channels

This formula assumes all hosts access all control units in the TS7680 with all channel paths.

Note: You cannot partition the TS7680 to support both open systems and System z environments. The TS7680 is optimized for a System z environment. There are no plans to enable the TS7680 to address open systems and System z environments concurrently.

7.3 Host implementation considerations

The implementation of the library consists of three steps:

1. Hardware Configuration Definition (HCD) to I/O

You use the HCD to define the tape control unit and the tape drives that belong to an IBM TS7680 to the Input/Output Definition File (IODF).

These definitions are hardware-related, and you must create them independently of the host operating system.

We discuss the HCD definitions that are required for the definition of tape controllers and their attached tape drives in detail in 7.5, “Hardware I/O configuration definition” on page 146.

2. Definitions in the hardware

The TS7680 has a fixed configuration internally, so no specific settings are required here. The actual configuration is needed to define the virtual library to the System z using Interactive Storage Management Facility (ISMF).

3. Software definitions

Here you define the new tape library to the individual operating system. If you use z/OS DFSMS and SMS tape, you update SMS constructs and Automatic Class Selection (ACS) routines, Object Access Method (OAM) definitions, and your tape management system during this phase.

Management tools

From the z/OS point of view, the TS7680 will be managed as a standard SMS tape library. From the ProtecTIER point of view the TS7680 provides a management console called the ProtecTIER Manager. This management interface allows the user to perform all repository

operations, such as add capacity or add virtual volumes. It also allows for repository monitoring and statistical analysis through various embedded tools and utilities.

Perform these tasks during the installation:

1. Update the tape management system:
 - Update the tape management exits as appropriate.
 - Define the VOLSER ranges.
2. Update the storage management system software; see the OAM Tap Library PISA at:
http://publibz.boulder.ibm.com/cgi-bin/bookmgr_OS390/BOOKS/dgt2o361/CCONTENTS?SHELFEZ2ZBK0I&DN=SC35-0427-09&DT=20100112161502

7.3.1 Channels, adapters, and protocols

The protocol used by System z servers to communicate with the TS7680 is FICON, which is different from the Fibre Channel Protocol (FCP) or SCSI protocol used by Open Systems. Open Systems hosts with FCP host bus adapters (HBAs) cannot be connected to FICON channels even if they are attached through directors that support both protocols. System z hosts use a FICON controller between the drives and the host channels.

7.3.2 System z tape controllers

System z-attached virtual drives are addressed using a front-end controller¹. FICON channels from the host are directly connected, often using a director, to the IBM TS7680 Enterprise Tape Controllers (IBM System Storage TS1120 Model C06 Controller). The virtual IBM tape drives attach to the controllers using the internal interfaces between the ProtecTIER servers and the Enterprise Tape Controllers. You can connect up to 128 tape drives to one of the internal Enterprise Tape Controllers.

7.4 Remote installations and switch support

The TS7680 attaches to the System z host through FICON channel attachments. There are three basic types of switch connections that can be used between the host and TS7680:

- ▶ Direct connect
- ▶ Single switch
- ▶ Cascaded switches

You can also use Dense Wave Division Multiplexers (DWDMs) or FICON channel extenders between the System z host and the TS7680. See Figure 7-1 on page 144 for more details about the distances supported.

7.4.1 Factors that affect performance at a distance

Fibre Channel distances depend on many factors, including:

- ▶ Type of laser used: longwave or shortwave
- ▶ Type of fiber optic cable: multi-mode or single-mode

¹ This discussion excludes the special case of Linux logical partitions (LPARs) in System z machines to which FCP drives can be directly attached.

- Quality of the cabling infrastructure in terms of dB signal loss:
 - Connectors
 - Cables
 - Bends and loops in the cable

Native shortwave FC transmitters have a maximum distance of 500 m with 50 micron diameter, multi-mode, optical fiber. Although 62.5 micron, multi-mode fiber can be used, the larger core diameter has a greater dB loss and maximum distances are shortened to 300 m. Native longwave FC transmitters have a maximum distance of 10 km when used with 9 micron diameter single-mode optical fiber.

Link extenders provide a signal boost that can potentially extend distances to up to about 100 km. These link extenders simply act as a very big, fast pipe. Data transfer speeds over link extenders depend on the number of buffer credits and efficiency of buffer credit management in the FC nodes at either end of this fast pipe. Buffer credits are designed into the hardware for each FC port. FC provides flow control that protects against collisions.

This is extremely important for storage devices, which do not handle dropped or out-of-sequence records. When two FC ports begin a conversation, they exchange information about their buffer capacities. An FC port will send only the number of buffer frames for which the receiving port has given credit. This not only avoids overruns, but also provides a way to maintain performance over distance by filling the “pipe” with in-flight frames or buffers. The maximum distance that can be achieved at full performance depends on the capabilities of the FC node that is attached at either end of the link extenders.

This relationship is very vendor-specific. There should be a match between the buffer credit capability of the nodes at either end of the extenders. A host bus adapter (HBA) with a buffer credit of 64 communicating with a switch port with only eight buffer credits would be able to read at full performance over a greater distance than it would be able to write. This is because, on the writes, the HBA can send a maximum of only eight buffers to the switch port, while on the reads, the switch can send up to 64 buffers to the HBA. Until recently, a rule of thumb has been to allot one buffer credit for every 2 km in order to maintain full performance.

Buffer credits in the switches and directors have a large part to play in the distance equation. The buffer credits in the sending and receiving nodes heavily influence the throughput that is attained in the Fibre Channel. Fibre Channel architecture is based on a flow control that ensures a constant stream of data to fill the available pipe. A rule-of-thumb says that to maintain acceptable performance, one buffer credit is required for every 2 km of distance covered. Refer to *Introduction to SAN Distance Solutions*, SG24-6408.

7.4.2 FICON Director support

All FICON Directors are supported with 1 Gbps, 2 Gbps, or 4 Gbps links. The components will auto-negotiate to the highest speed allowed.

You cannot mix different vendors (McData (CNT & Inrange), CISCO, and Brocade) but you can mix models of one vendor. See the switch Web pages for specific intermix combinations supported.

Access the latest information about SAN switches and directors here:

<http://www.ibm.com/servers/storage/san>

See also the IBM Redbooks publication *FICON Planning and Implementation Guide*, which is available at:

<http://www.redbooks.ibm.com/redbooks/pdfs/sg246497.pdf>

7.4.3 FICON channel extenders

FICON channel extenders are available working in one of the following modes:

- ▶ Frame shuttle or tunnel mode
- ▶ Emulation mode

Using the *shuttle* or *tunnel* mode, the extender receives and forwards FICON frames without doing any special channel or control unit processing. The performance is limited to the distance between the sites and the normal round trip delays in FICON channel programs.

Emulation mode can go unlimited distances, and it monitors the I/O activity to devices. The channel extender interfaces emulate a control unit by presenting command responses and CE/DE status ahead of the controller and emulating the channel when running the pre-acknowledged write operations to the real remote tape device. Thus, data is accepted early and forwarded to the remote device to maintain a full pipe throughout the write channel program.

7.4.4 Supported distances

When directly attaching to the host, the TS7680 can be installed at a distance of up to 10 km from the host. With FICON Switches, also called FICON Directors, or Dense Wave Division Multiplexers (DWDMs), the TS7680 can be installed at extended distances from the host.

Supported FICON extended distances are summarized in Figure 7-1.

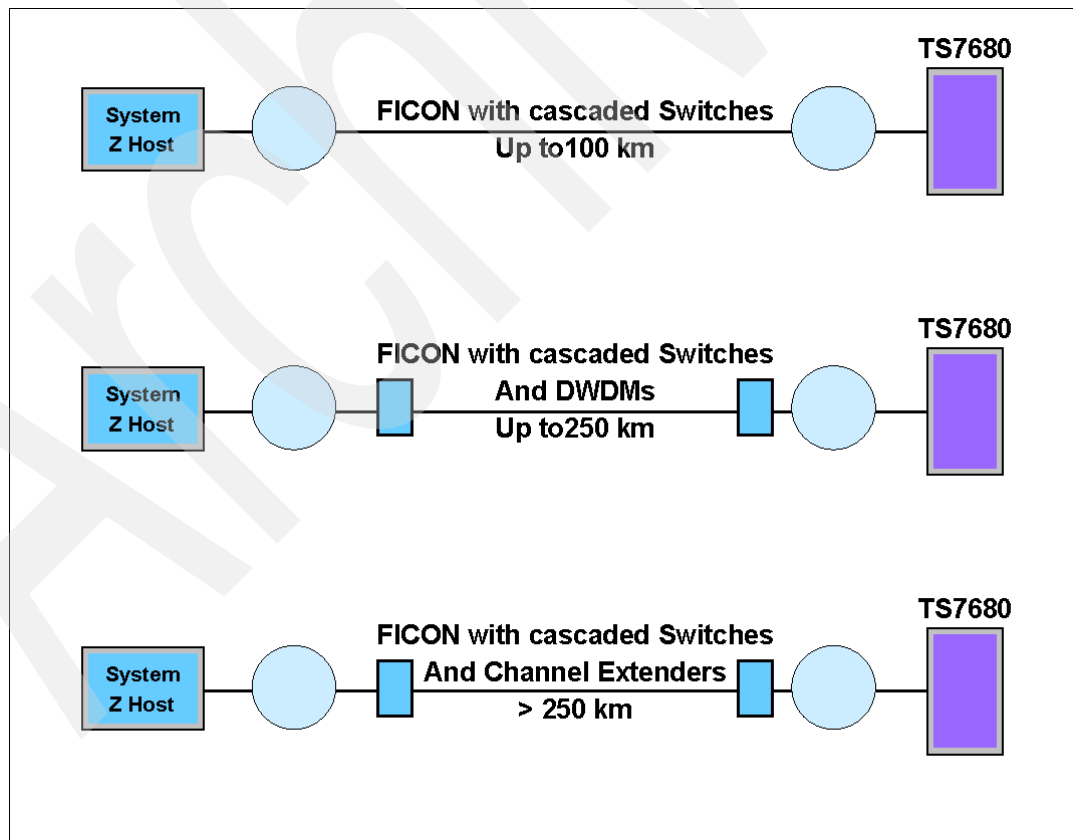


Figure 7-1 TS7680 Extended Distance Support

Figure 7-2 shows a complete diagram that includes the type and model of common DWDM and FICON Director products other than shortwave and longwave specifications. Although not shown in this diagram, FICON Directors and director cascading are supported.

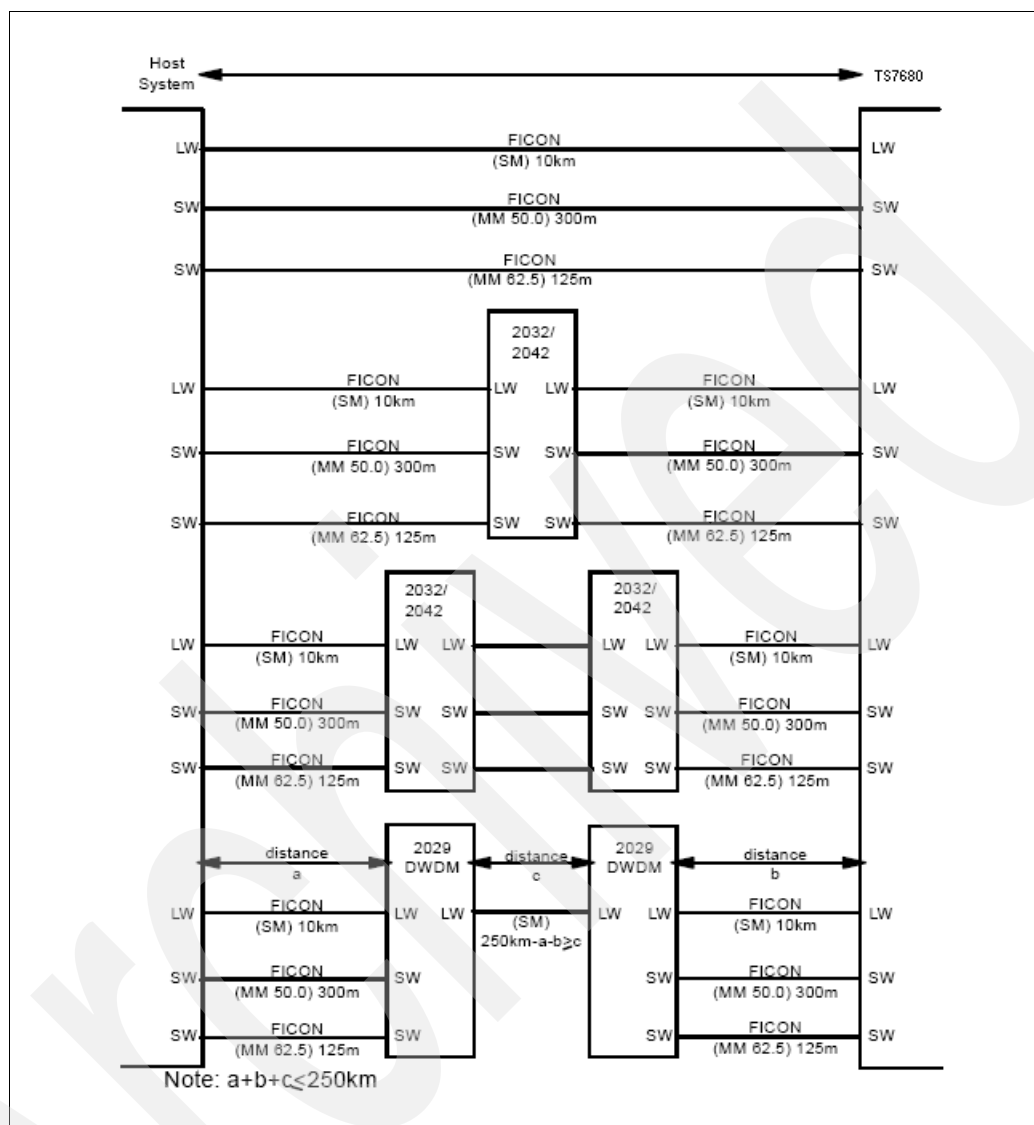


Figure 7-2 System z host attachment distances

7.4.5 Implementing cascaded switches

The following list summarizes the general configuration rules for configurations with cascaded switches:

- Director Switch ID

This is defined in the Director GUI setup.

The inboard Director Switch ID is used on the SWITCH= parameter in the CHPID definition. The Director Switch ID does not have to be the same as the Director Address. We recommend that you keep them the same to reduce configuration confusion and simplify problem determination work, although the example uses a different ID and address for clarity.

The following allowable Director Switch ID ranges have been established by the manufacturer:

- McDATA must be in the range x'61' to x'7F'
- CNT/Inrange must be in the range x'01' to x'EF'
- Brocade must be in the range x'01' to x'EF'

► Director Address

This is defined in the Director GUI setup.

The Director Domain ID is the same as the Director Address that is used on the LINK parameter in the CNTLUNIT definition. The Director Address does not have to be the same as the Director ID, but again, we recommend that you keep them the same to reduce configuration confusion and simplify PD work.

The following allowable Director Address ranges are established by the manufacturer:

- McDATA must be in the range x'61' to x'7F'.
- CNT/Inrange must be in the range x'01' to x'EF'.
- Brocade must be in the range x'01' to x'EF'.

► Director Ports

The Port Address might not be the same as the Port Number. The Port Number identifies the physical location of the port, and the Port Address is used to route packets.

The Inboard Director Port is the port to which the CPU is connected. The Outboard Director Port is the port to which the control unit is connected. It is combined with the Director Address on the LINK parameter of the CNTLUNIT definition:

- Director Address (hex) combined with Port Address (hex) (two bytes)
- Example: LINK=6106 would indicate a Director Address of x'61' and a Port Address of x'06'

► External Director connections

- Inter-Switch Links (ISLs) connect to E Ports.
- FICON Channels connect to F Ports.

► Internal Director connections

Port type and Port-to-Port connections are defined using the Director's GUI setup.

7.5 Hardware I/O configuration definition

Hardware Configuration Definition (HCD) is the required I/O configuration definition tool that defines which channel subsystems, controllers, and devices exist and can be accessed by the operating systems. All I/O configuration settings are stored in an I/O configuration repository called the *I/O Definition File (IODF)*. Before you can use a tape drive on a z/OS host, you must define it to the hosts through HCD.

The HCD definitions for an Enterprise Tape Controller is no different than the definitions when these devices are installed in an physical tape library. For completeness, we have included the HCD panels to define a Enterprise Tape Controller and its drives through HCD, as well as a discussion of the support for LIBRARY-ID and LIBPORT-ID in HCD.

7.5.1 Defining a TS7680

The definition of a TS7680 is almost the same as that for a native library, because from a host perspective, the TS7680 Virtualization Engine looks like sixteen IBM tape control units with 16 devices each attached through FICON. The most important points to observe are:

- HCD definitions are required.
- You must define 16 control units with 256 devices.
- Use CUADD = 0 through CUADD = 7 and LIBPORT-IDs of 01 through 08 for the first eight control units as shown in Table 7-2.

Table 7-2 CUADD and LIBPORT-ID for the first set of 128 virtual devices

CU lower	1	2	3	4	5	6	7	8
CUADD	0	1	2	3	4	5	6	7
Subsystem-ID/ LIBPORT-ID	01	02	03	04	05	06	07	08

For the ninth to sixteenth control units, use CUADD = 0 through CUADD = 7 and LIBPORT-IDs of 11 through 18. Refer to Table 7-3.

Table 7-3 CUADD and LIBPORT-ID for the second set of virtual devices

CU upper	1	2	3	4	5	6	7	8
CUADD	0	1	2	3	4	5	6	7
Subsystem-ID/ LIBPORT-ID	11	12	13	14	15	16	17	18

- Keep the link address blank when you do not use a FICON director.
- Specify LIBRARY=YES when you use system-managed tape.

Figure 7-3 shows the Add Control Unit HCD panel with the definition of one TS7680 Virtualization Engine control unit. Note that the control unit type differs from a native library, because only the emulated tape drives are defined in the HCD, not physical tape drives.

```

----- Add Control Unit -----
Specify or revise the following values.

Control unit number . : 0441          Type . . . . . : 3590
Processor ID . . . . . : A2097        A2097
Channel Subsystem ID . : 1            CSS1 ON A2097

Channel path IDs . . . . 3D    E8    _ _ _ _ _ _ _ _ _ _ +
Link address . . . . . 20    23    _ _ _ _ _ _ _ _ _ _ +

Unit address . . . . . 00          _ _ _ _ _ _ _ _ _ _ +
Number of units . . . . 016        _ _ _ _ _ _ _ _ _ _

Logical address . . . . 1 + (same as CUADD)

Protocol . . . . . _ + (D,S or S4)
I/O concurrency level . _ + (1, 2 or 3)

```

Figure 7-3 HCD: Add Control Unit for TS7680 Virtualization Engine

The frame in Figure 7-3 is for the second control unit using the Unit Address or CUADD = 01. *Control Unit Address* is used to allow access to all logical control units of a TS7680 Virtualization Engine through all channel paths.

The next two sections describe a step-by-step definition of control units and devices.

7.5.2 Control unit definition

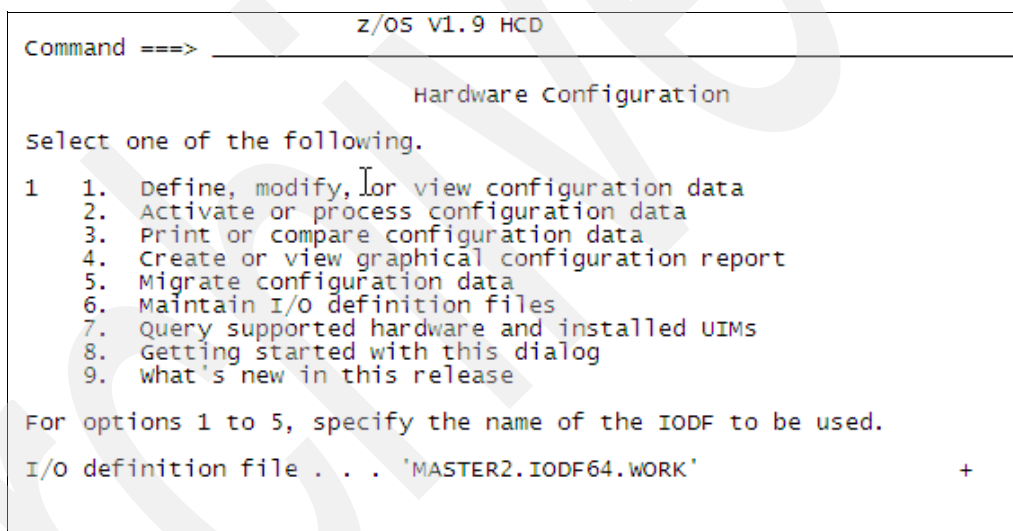
To define a control unit, follow these steps:

1. From the Hardware Configuration primary display (Figure 7-4), enter the name of the IODF file that you want to update and select option **1. Define, modify, or view configuration data**. Press Enter.

If this is not an IODF work file, the system prompts you to create one. The name of the IODF file is in the format *hlq.IODFcc.yyyyyyy*.

Note the following explanations:

- *hlq* is a high-level qualifier of up to eight characters.
- *cc* is for any two hexadecimal characters.
- *yyyyyyy* is up to eight optional characters.



```

Command ===> z/OS V1.9 HCD

Hardware Configuration

Select one of the following.

1  1. Define, modify, or view configuration data
   2. Activate or process configuration data
   3. Print or compare configuration data
   4. Create or view graphical configuration report
   5. Migrate configuration data
   6. Maintain I/O definition files
   7. Query supported hardware and installed UIMS
   8. Getting started with this dialog
   9. What's new in this release

For options 1 to 5, specify the name of the IODF to be used.
I/O definition file . . . 'MASTER2.IODF64.WORK'
  
```

Figure 7-4 Hardware Configuration definition primary panel

2. To define the control unit, select option **4. Control units** on the Define, Modify, or View Configuration Data panel (refer to Figure 7-5 on page 149). Press Enter.

```

                                z/OS V1.9 HCD
                        Define, Modify, or View Configuration Data

select type of objects to define, modify, or view data.

4_ 1. operating system configurations
    consoles
    system-defined generics
    EDTs
    esoterics
    user-modified generics
2. Switches
    ports
    switch configurations
    port matrix
3. Processors
    channel subsystems
    partitions
    channel paths
4. Control units
5. I/O devices

```

Figure 7-5 HCD: Define, Modify, or View Configuration Data panel

- Now, the Control Unit List panel appears. Refer to Figure 7-6.

```

Goto Filter Backup Query Help
-----
Control Unit List                                Row 560 of 765
Command ==>                                     Scroll ==> CSR
Select one or more control units, then press Enter.  To add, use F11.

/ CU  Type +   CUADD CSS MC  Serial-# + Description
- 1604 3490    3      1      FIG BARR31
- 1605 3490    4      1      FIG BARR31
- 1606 3490    5      1      FIG BARR31
- 1607 3490    6      1      FIG BARR31
- 1608 3490    7      1      FIG BARR31
- 1609 3490    8      1      FIG BARR31
- 160A 3490    9      1      FIG BARR31
- 160B 3490    A      1      FIG BARR31
- 160C 3490    B      1      FIG BARR31
- 160D 3490    C      1      FIG BARR31
- 160E 3490    D      1      FIG BARR31
- 160F 3490    E      1      FIG BARR31
- 1610 3490    F      1      ZORO BARR39
- 1611 3490    0      1      ZORO BARR39
- 1612 3490    1      1      ZORO BARR39
- 1613 3490    2      1      ZORO BARR39
- 1614 3490    3      1      ZORO BARR39
- 1615 3490    4      1      ZORO BARR39
- 1616 3490    5      1      ZORO BARR39
- 1617 3490    6      1      ZORO BARR39
- 1618 3490    7      1      ZORO BARR39
- 1619 3490    8      1      ZORO BARR39
- 161A 3490    9      1      ZORO BARR39

```

Figure 7-6 HCD: List Control Unit panel

- On the Control Unit List panel, press F11 to get to the Add Control Unit panel that Figure 7-7 on page 150 shows.

```

                                Add Control Unit

Specify or revise the following values.
Control unit number . . . . . 4000 +
Control unit type . . . . . 3590 +

Serial number . . . . . _____
Description . . . . . _____

Connected to switches . . . A9 _ _ _ _ _ +
Ports . . . . . 01 _ _ _ _ _ +

If connected to a switch:
Define more than eight ports . . 2 1. Yes
                                   2. No
Propose CHPID/link addresses and
unit addresses . . . . . 2 1. Yes
                           2. No

```

Figure 7-7 HCD: Add Control Unit panel

- When the Add Control Unit panel appears (Figure 7-7), enter the responses that match your environment. You might want to update this panel with the machine's serial number and a description. If you are uncertain about any of the required responses, press F1 for Help. Fields that have a plus sign to their right display a set of prompts when you press F4. We enter control unit number 4000 in our example as shown in Figure 7-7. The Enterprise Tape Controller runs in 3590 emulation mode; therefore, we define tape control units as Control unit type 3590. We have a FICON switch in our sample configuration, so we must enter the switches and ports.

Note: For load balancing, failover, and redundancy reasons we recommend to define more than one CHPD (port) or switch in the Add Control Unit panel, depending on your environment.

- After you press Enter, the Select Processor / CU panel appears. Select the processor and the Channel Subsystem (N2094.0 in our example) to which you will attach the control unit by typing an s to the left of N2094.0 as shown in Figure 7-8 on page 151. In addition, enter the Channel Path ID and the Link Address (60.01 in our example).

```

                                Select Processor / CU      Row 1 of 4 More:
Command ==> _____ Scroll ==> CSR
select processors to change CU/processor parameters, then press Enter.
Control unit number . . : 4000      Control unit type . . . : 3590

/ Proc.CSSID 1-----2-----3-----4-----5-----6-----7-----8-----
s N2094.0    60.01_  _____  _____  _____  _____  _____  _____
- N2094.1    _____  _____  _____  _____  _____  _____  _____
- N2094.2    _____  _____  _____  _____  _____  _____  _____
- N2094.3    _____  _____  _____  _____  _____  _____  _____
***** Bottom of data *****

```

Figure 7-8 HCD: Select Processor / Control Unit panel

Note: If you plan to attach the control unit to multiple processors, type a g to the left of each processor that you want to attach to the control unit.

7. Press Enter to continue and the Unit Address, and the number of units (Unit Range) for each processor to which you will attach the control unit, as shown in Figure 7-9.

```

                                Select Processor / CU      Row 1 of 4 More: < >
Command ==> _____ Scroll ==> CSR
select processors to change CU/processor parameters, then press Enter.
Control unit number . . : 4000      Control unit type . . . : 3590

/ Proc.CSSID Att CU -----Unit Address . Unit Range + -----
s N2094.0    00  ADD+ 1-----2-----3-----4-----5-----6-----7-----8-----
- N2094.1    _____  _____  _____  _____  _____  _____  _____
- N2094.2    _____  _____  _____  _____  _____  _____  _____
- N2094.3    _____  _____  _____  _____  _____  _____  _____
***** Bottom of data *****

```

Figure 7-9 HCD: Select Processor / Control Unit panel continued

8. After pressing Enter you have to add the Channel Path ID (CHPID). If you have more than one channel subsystem specified in your Central Processing Complex (CPC), also enter the Channel Subsystem ID (Refer to Figure 7-10 on page 152).

```

Select Processor / CU
Add Control Unit

Specify or revise the following values.
Control unit number . . . : 4000      Type . . . . . : 3590
Processor ID . . . . . : N2094      MVSC5,MVSC6,SY55, VMT24 on N2094
Channel subsystem ID . . : 0        CSS0 ON N2094

Channel path IDs . . . . 60      _ _ _ _ _ _ _ _ _ _ +
Link address . . . . . 01      _ _ _ _ _ _ _ _ _ _ +
Unit address . . . . . 00      _ _ _ _ _ _ _ _ _ _ +
Number of units . . . . 16      _ _ _ _ _ _ _ _ _ _ +
Logical address . . . . 0      + (same as CUADD)
Protocol . . . . . _      + (D, S or S4)
I/O concurrency level . _      + (1, 2 or 3)

```

Figure 7-10 HCD: Add Control Unit panel

9. Press Enter to complete and save the definition of the control unit. Press F3 to exit from the Control Unit Definition menu.

Next, you must define the attached devices for this controller.

7.5.3 Device definition

To define the attached devices for this controller:

1. From the HCD entry panel that is shown in Figure 7-5 on page 149, select option 5. From the next panel, press F11 to get the Add Device panel.
2. On the Add Device panel (refer to Figure 7-11), add the device number, the number of devices, the device type, and the connected control unit. Press Enter.

```

Add Device

Specify or revise the following values.
Device number . . . . . 5000      + (0000 - FFFF)
Number of devices . . . . . 16
Device type . . . . . 3590      _ _ _ _ _ +
Serial number . . . . . _ _ _ _ _ _
Description . . . . . _ _ _ _ _ _ _ _ _ _ _ _ _ _ _
Volume serial number . . . . . _ _ _ _ _ (for DASD)
Connected to Cus . . 4000      _ _ _ _ _ _ _ _ _ _

```

Figure 7-11 HCD: Add Device panel

3. The Device / Processor Definition panel appears (Figure 7-12 on page 153). Type an s to the left of the processor that will use the devices. Press Enter.


```

                                Device / Processor Definition
                                Row 1 of 1
Command ===> _____ Scroll ===> CSR

Select processors to change device/processor definitions, then press
Enter.

Device number . . . : 5000      Number of devices . : 16
Device type . . . : 3590

/ Proc.CSSID  SS+  UA+  Time-Out  STADET  CHPID +  Device Candidate List
s N2094.0      _   _   No         Yes      _       Explicit      Null
***** Bottom of data *****

```

Figure 7-12 HCD: Device / Processor Definition panel

Restriction: You cannot use Device number 0000. As documented in APAR OW56336, a restriction exists regarding the use of device address 0000 for all SMS-managed tape libraries, because software uses 0000 to indicate a null entry in library-related tables and control blocks.

4. The Define Device / Processor panel appears (Figure 7-13). This panel describes the processor's view of the device. Press Enter twice to get the Define Device to Operating System Configuration panel.

```

                                Define Device / Processor

Specify or revise the following values.

Device number . . . . : 5000      Number of devices . . . . : 16
Device type . . . . : 3590
Processor ID . . . . : N2094      MVSC5,MVSC6,SYS5,VMT24 on N2094
Channel Subsystem ID : 0         CSS0 ON N2094

Subchannel set ID . . . . . : _   +
Unit address . . . . . : 00      + (Only necessary when different from
                                the last 2 digits of device number)
Time-Out . . . . . : No         (Yes or No)
STADET . . . . . : Yes         (Yes or No)

Preferred CHPID . . . . . : _   +
Explicit device candidate list : No (Yes or No)

```

Figure 7-13 HCD Define Device / Processor panel

5. The Define Device to Operating System Configuration panel appears next; refer to Figure 7-14 on page 154. Type an s to the left of the operating system or systems that will use the TS7680 library. Press Enter.

```

Define Device to Operating System Configuration
Command ===> _____ Row 1 of 2
                                scroll ===> CSR
Select OSS to connect or disconnect devices, then press Enter.
Device number . : 5000          Number of devices : 16
Device type . . : 3590

/ Config. ID  Type  SS Description          Defined
s DFRMM3      MVS   MVS5,MVSCK AND SYS5
_ TPCR        MVS   SYSQ,SYSR
***** Bottom of data *****

```

Figure 7-14 HCD: Define Device to Operating System Configuration panel

6. The Define Device Parameters / Features panel appears. Refer to Figure 7-15. On this panel, you link the operating system to the tape library that you are installing.

```

Define Device Parameters / Features
Command ===> _____ Row 1 of
                                scroll ===> CSR
Specify or revise the values below.
Configuration ID . : DFRMM3      MVS5,MVSCK AND SYS5
Device number . . : 5000        Number of devices : 16
Device type . . . : 3590

Parameter/
Feature  Value +      R Description
OFFLINE  Yes         Device considered online or offline at IPL
DYNAMIC  Yes         Device supports dynamic configuration
LOCANY   YES         UCB can reside in 31 bit storage
LIBRARY  YES         Device supports auto tape library
AUTOSWITCH NO       Device is automatically switchable
LIBRARY-ID 00032     5 digit library serial number
LIBPORT-ID 01       2 digit library string ID (port number)
MTL      NO         Device supports manual tape library
SHARABLE NO         Device is sharable between systems
COMPACT  YES        Compaction
***** Bottom of data *****

```

Figure 7-15 HCD: Define Device Parameters / Features panel

In the panel that is shown in Figure 7-15, you define the only connection between the tape library and the devices. Note that you do not need to define the IBM TS7680 configuration through the HCD, you only define the control units and the devices. However, you must define the devices with the parameter LIBRARY Yes.

We recommend that you always specify YES to the parameters DYNAMIC and LOCANY.

Always specify COMPACT YES; otherwise, you turn off drive hardware compression. In general, we recommend that you use drive compression. If you do not want to use drive compression for specific applications, we recommend that you use the Data Class specification to temporarily turn off drive hardware compression.

You can accept all of the default parameter values, except for the following parameters:

- OFFLINE depends on your environment. OFFLINE specifies whether z/OS considers the device online or offline at IPL. If Yes, the device is considered offline at IPL. If No (the default), the device is considered online at IPL. We recommend that you specify Yes and use the COMMNDxx member of PARMLIB to vary drives online. This method

is also necessary if you want to use Automatic Tape Switching (ATS STAR or IEFAUTOS).

- DYNAMIC specifies whether you allow dynamic activation through an activate command. Always specify YES.
- LOCANY specifies whether the unit control block (UCB) can reside in 31-bit storage. Always specify YES.
- LIBRARY specifies whether to indicate that the device belongs to an automated tape library. Specify YES.
- AUTOSWITCH defines the devices as automatically switchable. Specify YES to indicate that the device will be treated as automatically switchable. For tape drives to be automatically switchable, they must be shared by systems in a Parallel Sysplex®.
- LIBRARY-ID and LIBPORT-ID are described in detail in “HCD support for LIBRARY-IDs and LIBPORT-IDs” on page 155.
- SHARABLE specifies whether you want to share the defined device between multiple processors. Specify YES. For tape drives, the OFFLINE parameter must be set to YES for using sharable tape devices.
- COMPACT specifies whether to indicate that compaction is available for tape devices. Compaction is standard on 3490, 3490E, and all 3590 and 3592 tape drives. Specify YES.

Press Enter.

7. The Assign/Unassign Device to Esoteric panel displays. Defining esoterics is site-dependent and optional. You do not have to define tape library resident devices to an esoteric in system-managed tape. Use esoteric device names when the number of installed physical drives is fewer than the number of devices defined in HCD in order to prevent allocations going to offline devices in the tape library. You must handle the esoteric device names in your System Managed Storage (SMS) Automatic Class Selection (ACS) routines and assign them to an appropriate tape Storage Group. Now you have defined your tape library and drives, and a production IODF is built, loaded, and then activated.

7.5.4 HCD support for LIBRARY-IDs and LIBPORT-IDs

The LIBRARY-ID is the unique identification number of a tape library. It specifies the hardware ID associated with the tape library that you define. LIBRARY-ID is defined by the IBM SSR at the time of the library installation. In terms of the z/OS operating system, LIBPORT-ID reflects the order in which the tape control units connect to the Library Manager and provides the tape drive pool ID, which is transparent and only used by allocation. LIBRARY-ID and LIBPORT-ID are optional z/OS HCD parameters. They allow the HCD to provide configuration information to the logical library that is normally obtained by the operating system at IPL time.

Note: You can find more information about the parameters of the installed TS7680 in Appendix B, “Checklists” on page 297.

When adding the virtual tape drives in the IBM TS7680 Tape Library through the HCD, you can optionally supply the information for LIBRARY-ID and LIBPORT-ID. If you do not provide the information, the operating system attempts to obtain the information from the tape subsystem at IPL or I/O activation time. If the operating system cannot get the information from the subsystem (which is the case if the subsystem is not powered on during IPL), the devices cannot be used. Therefore, we recommend that you provide the default information for LIBRARY-ID and LIBPORT-ID.

Note: If you do not use LIBRARY-ID and LIBPORT-ID, devices that are unavailable during IPL cannot be varied online without reactivating the IODF.

The IBM SSR assigns a unique five character *Library Sequence Number* to the logical library. For the TS7680 library, this number can be the last five digits of the serial number of the TS7680 base frame. This Sequence Number must match the LIBRARY-ID number used in the HCD definition for the partition, as well as the LIBRARY-ID coded in the library definition in Interactive Storage Management Facility (ISMF).

Note: We strongly recommend that you assign the last five digits from the serial number of the IBM TS7680 base frame as the LIBRARY-ID.

7.6 Cache management

When the TS7680 detects that available cache space has fallen below preset thresholds, the TS7680 sends CBR3792E (limited cache free space warning state reached) and CBR3794A (out of cache resource critical state reached) attention messages to the attached hosts (same messages as with the TS7720 Virtualization Engine). The CBR3792E message can be used to trigger return to scratch processing or the copying of data to another library. If the amount of available cache subsequently reaches the critical state, all fast ready (scratch) mounts are failed and any specific mount operations are allowed; however, any attempt to write to the volume will be failed. Mount operations that have been accepted before entering this state complete and volumes currently mounted can continue to perform host I/O operations. As appropriate, the VARY SMS,STORGRP operator command can also be used to steer scratch allocations to another library that is eligible for the scratch request. The DISPLAY SMS,LIBRARY command with DETAIL can also be used periodically to display the CACHE PERCENTAGE USED status line.

With software support installed, the host will notify the library when a volume is returned to scratch. The TS7680 then applies (by default) a nine-day grace period to the volume. The default grace period can be overridden at install time and set to a value from 0-9 days (0 indicates no grace period). After the grace period elapses (see Table B-1 on page 299 and the scratch delay option), the data associated with the scratch volume will be deleted from the library to free up back-end disk space, rendering the contents of the tape volume unusable. When the TS7680 reaches the critical Out of Cache Resources state (CBR3794A), it can also be configured (at install time) to automatically delete data associated with scratch volumes that are in the “grace period”. Volumes with the shortest time remaining in the grace period will be deleted first. The default behavior is to honor the grace period.

Upon reuse of the scratch volume, since the library no longer has the previous VOL1/HDR1/HDR2 information, the library will return dummy VOL1/HDR1/TM. The library will also return (in HDR1), “VirtualSCRTCH” at offset 61 (starting with 1); the system code field bytes 61-73. This may be needed by the tape management system to allow the reuse of a volume that had previously been used (especially since the TS7680 library, drives, and volumes are not reported to the host as virtual). The TS7680 will also return dummy VOL1/HDR1/TM for a newly inserted logical volume, and to be consistent with reuse processing will also return VirtualSCRTCH at offset 61 (HDR1). Support for VirtualSCRTCH has been provided in DFSMSrmm for several years (OA09171). If you are not using DFSMSrmm, check with your tape management system for any needed support in this area.

Note: The host and software support (APAR OA27786) provides the recommended updates for the TS7680. This support is available at z/OS VR9 and above. Usage of the TS7680 prior to z/OS V1R9 requires that a Request for Price Quotation (RPQ) be submitted.

Without host support

1. Volumes can still be returned to scratch, but the associated disk cache will not be made available until the volume is rewritten by the host.
2. Cache threshold messages introduced with the TS7720 Virtualization Engine (CBR3792E CBR3793I CBR3794A CBR3795I) may appear without this support (at a minimum OA24966 is needed). However, this support is needed for the cache percentage used line to appear in the DISPLAY SMS,LIBRARY command output of a TS7680 library.
3. This support is needed for the TS7680 to detect when a scratch mount is sent to the library. Without this support, when the library reaches its critical cache state, scratch mounts will be accepted and will fail on the first write (same as with specific mounts).

This support is needed for the library to track when a volume is returned to scratch and without this support, volumes returned to scratch cannot be ejected (purged) from the library.

7.7 Cartridge entry and eject

Volumes are defined to the TS7680 (through the IBM ProtectTIER Manager GUI) and go through the same cartridge entry processing at the host as IBM's other tape libraries. Also, since an eject request will purge the volume and its data from the library, as with other IBM Virtual Tape Libraries, a volume will only be ejected (purged) from the TS7680 if it is marked in their database as a scratch volume, otherwise the eject request will be failed by the library (function incompatible). Any request to eject a volume from the TS7680 should be initiated from the host (to keep databases in synch).

7.8 Tape initialization

For volumes in an automated tape library data server, you have the option to use DFSMSdfp OPEN processing as an alternative to using DFSMSrmm EDGINERS or IEHINITT to label scratch volumes.

Note: Initialization for logical volumes is not required with the TS7680. Stacked and logical volumes are initialized transparently to the user and host at the time of first use.

Upon reuse of the scratch volume, since the library no longer has the previous VOL1/HDR1/HDR2 information, the library will return dummy VOL1/HDR1/TM. The library will also return (in HDR1), VirtualSCRTCH at offset 61 (starting with 1); the system code field bytes 61-73. This may be needed by the tape management system to allow the reuse of a volume that had previously been used (especially since the TS7680 library, drives, and volumes are not reported to the host as virtual).

The TS7680 will also return dummy VOL1/HDR1/TM for a newly inserted logical volume and to be consistent with reuse processing will also return VirtualSCRTCH at offset 61 (HDR1). Support for VirtualSCRTCH has been provided in DFSMSrmm for several years (OA09171).

If you are not using DFSMSrmm, check with your tape management system for any needed support in this area.

7.8.1 DFSMSrmm for initialization

If you want to use DFSMSrmm instead of DFSMSdftp to initialize new tapes in a library, follow the steps in this procedure:

1. Perform one of the following actions:
 - Enter the undefined volumes into the TS7680 while DFSMSrmm is active.
 - Define the volumes as scratch to DFSMSrmm with LOCATION (*atlname*), and enter the volumes into the TS7680 with DFSMSrmm active.
2. Volumes must now be defined to DFSMSrmm with SCRATCH status. They must be known to be in the library.
3. Use the RMM CV VOLSER INIT(Y) command to set the initialize action for each volume. Use the following command to build the commands:

```
RMM CV VOL(*) STATUS(SCRATCH) LOC(atlname)
```
4. Run EDGINERS in automatic mode.

In the sample DFSMSrmm EDGINERS (Figure 7-16 on page 159), an automatic run of EDGINERS is scheduled to find and initialize up to 99 volumes residing in an automated tape library data server called MYATL. All tape cartridges are labeled as appropriate for the drive type on which they are mounted and for their current media characteristics.
5. DFSMSrmm temporarily sets the TCDB status to PRIVATE for the tapes to be initialized, because no specific mounts (as they are required for labeling a cartridge) are allowed for SCRATCH tapes inside a library.

Note: The automatic synchronization between DFSMSrmm and the TCDB works only if DFSMSrmm runs in PROTECT mode.

EDGINERS determines whether a volume in a system-managed tape library can be mounted on the current system. If the volume cannot be mounted, possibly because it is defined in a TCDB on another system, DFSMSrmm skips that volume.

The control statement description is:

- ▶ Tape DD and SYSIN DD are not required for a system-managed tape environment.
- ▶ PARM values request initialization of 99 cartridges in library MYATL. No verification is done. Verification causes each cartridge to be mounted twice: once for initialization and once for verification.

DFSMSrmm ensures that volumes in a system-managed tape library that are to be initialized or erased are in the private category, because the automated tape library data server does not support specific mounts of scratch volumes. You must define a volume in a system-managed tape library to DFSMSrmm before you can initialize or erase it. Any volume that is not defined to DFSMSrmm is requested to be mounted on the drive that is allocated by the TAPE DD statement in the JCL for EDGINERS as long as the drive is not in a system-managed library.

During demount processing, DFSMSrmm ensures that errors detected on volumes mounted in an automated tape library are reflected in the TCDB. For example, DFSMSrmm ensures that the TCDB contains information about write-protected, wrong volume, and wrong label type errors. DFSMSrmm skips the volume rather than having the operator correct the error.

```
//STEP1 EXEC PGM=EDGINERS,
//      PARM='COUNT(99),LOCATION(MYATL),INITIALIZE,NOVERIFY'
//SYSPRINT DD SYSOUT=A
```

Figure 7-16 Sample DFSMSrmm EDGINERS

7.8.2 Scratch pooling

System-managed tape does not support multiple scratch pools of a single media type. Refer to Table 7-4.

Table 7-4 Supported Media type for TS7680s

Media type	Name	Device type	Recording format	WORM or R/W	Cartridge capacity
MEDIA5	IBM 3592 Enterprise Tape Cartridge (ETC)	3592	EFMT1	R/W	100 GB

Note: The capacity of a TS7680 logical volume is less than a traditional JA (MEDIA5) physical cartridge (100 GB versus 300 GB).

Duplicate volume serial numbers

For system-managed tape, all VOLSERS in the same SMSplex must be unique across tape and DASD.

DFSMSrmm does not support duplicate VOLSERS and cannot manage volumes that are not defined to it. By defining a RACF® profile named STGADMIN.EDG.IGNORE.** with access of read, you allow specific users to bypass this check with the JCL parameter EXPDT=98000. If you even have a duplicate VOLSER in RMM that resides in an IBM robot, you must add the parameter STORCLAS=DUPT@SMS, as shown in Example 7-1.

Example 7-1 Duplicate VOLSER residing in IBM robot

```
//TAPE1 DD DSN=E40488.SASLDBE1,
//      DISP=OLD,UNIT=TAPEV,EXPDT=98000,
//      VOL=SER=676620,LABEL=(,BLP) ,STORCLAS=DUPT@SMS
```

In the IBM TS7680, all volumes must be unique. You have to use distinct volume serial number ranges for the three volume types.

7.8.3 Defining the library through ISMF

You define your TS7680 library to the system through the ISMF library application.

For details about defining your library, refer to *z/OS DFSMSdfp Storage Administration Reference*, SC26-7402.

When you define your library, you specify:

- LIBRARY-ID** This is the five-character ID associated with the TS7680 library.
- Console name** This is the optional z/OS console name if you defined an optional z/OS console name in SYS1.PARMLIB member CONSOLxx.

Entry default Data Class

This specifies the name of the Data Class that you want as the default for tape cartridges entered into the library that you define.

Entry default use attribute

This specifies the use attribute for cartridges that are entered into the library (SCRATCH or PRIVATE).

Eject default

This is the default action for the TCDB volume record when a tape cartridge is ejected from the library (PURGE or KEEP).

Scratch threshold

This specifies the threshold below which a message is issued to the operator requesting that the operator enter scratch volumes of the specified media type into the library.

Initial online status

This specifies whether the library is online, offline, or not connected to the systems or system groups in the SMSplex each time that the SCDS is activated. We recommend that you specify online to ensure that the library is accessible after the activation of an updated SCDS.

Note: When you connect a TS7680 Tape Library to a system group rather than to a system, you lose the ability to vary that library online or offline to the individual system in the group. We recommend strongly that you connect the TS7680 Tape Library to individual systems only.

7.8.4 Defining SMS constructs through ISMF

To direct allocations to system-managed tape, you have to define SMS constructs through ISMF. In the Data Classes, you specify the media type, the recording technology, and whether to use hardware compaction when allocating a system-managed tape data set.

You do not have to specify new Storage Classes. You can use existing classes. The Storage Class is used only to indicate that this is an allocation to a system-managed tape library. However, we recommend that you create new Storage Classes for tape, so that you can select Storage Groups on the basis of the Storage Class assignment and keep the automatic class selection (ACS) routines simple.

As for system-managed DASD allocations, the Management Class is optional. System-managed tape uses only the expiration attributes and retention limit parameters. If you use a tape management system, specify a retention limit of NOLIMIT.

You need to define a tape Storage Group and specify which libraries belong to that Storage Group. You also define the Storage Group status.

Although a blank Storage Group is allowed for system-managed tape volumes, we strongly recommend that you assign a Storage Group to private volumes when they are entered into the TS7680 Tape Library. The blank Storage Group is always enabled for all attached systems. You can specify the Storage Group during the definition of an existing private volume to DFSMSrmm or during cartridge insert processing.

Compared to the implementation of DFSMS for DASD, system-managed tape has the following differences:

- ▶ Tape data sets do not have to be cataloged. If they are to be cataloged, you catalog them at step termination time.
- ▶ System-managed tape is the management of tape cartridges, not tape data sets. No data set-related information is stored in the TCDB.

- ▶ A DASD (type POOL) Storage Group comprises one or more DASD volumes.
A tape (type TAPE) Storage Group comprises one or more tape libraries. Cartridge information is stored in the TCDB, not in the SMS active control data set (ACDS).
- ▶ Tape volumes are not preassigned to Storage Groups. They are assigned a Storage Group when their status changes to PRIVATE. Scratch volumes do not have a Storage Group assigned.

A blank Storage Group is allowed for system-managed tape.

7.8.5 Defining Data Classes

A Data Class provides the tape device selection information or tape data sets. The attributes that you can specify are:

- ▶ The type of media to use, which for the TS7680 only MEDIA5 is valid
- ▶ Whether the data is to be compacted
- ▶ Recording technology (18 track, 36 track, 128 track, 256 track, 384 track, EFMT1)
- ▶ The maximum volume count that your data set can span

Use ISMF panels to define your Data Classes:

1. Choose option **4 (Data Class)** on the ISMF PRIMARY OPTION MENU display. The DATA CLASS APPLICATION SELECTION panel appears.
2. On the DATA CLASS APPLICATION SELECTION panel, specify the SCDS name and the name of the Data Class that you are about to define.
3. Choose option **3 (Define)** to create a new Data Class or option **4 (Alter)** to change an existing Data Class on the panel. Figure 7-17 shows the panel where you specify the Data Class name.

```

Panel  Utilities  Help
-----
                        DATA CLASS APPLICATION SELECTION

To perform Data Class Operations, Specify:
CDS Name . . . . . 'DFSMS150.SCDs.SYS32'
                                     (1 to 44 character data set name or 'Active' )
Data Class Name . . HSMDC11M (For Data Class List, fully or partially
                               specified or * for all)

Select one of the following options :
4  1. List      - Generate a list of Data Classes
   2. Display   - Display a Data Class
   3. Define    - Define a Data Class
   4. Alter     - Alter a Data Class

If List Option is chosen,
Enter "/" to select option      Respecify View Criteria
                                Respecify Sort Criteria

Command ==>
F1=Help   F2=Split   F3=End   F4=Return  F7=Up     F8=Down   F9=Swap
F10=Left  F11=Right  F12=Cursor

```

Figure 7-17 Data Class Application Selection panel

4. Now, the first page of the DATA CLASS DEFINE or DATA CLASS ALTER panel appears as shown in Figure 7-18. The panels are the same for both Data Class Define and Data Class Alter; in the our example, we chose Data Class Alter.

Panel Utilities Scroll Help	
DGTDCDC1	DATA CLASS ALTER Page 1 of 5
SCDS Name . . . : DFSMS150.SCD.SYS32	
Data Class Name : HSMDC1M	
To ALTER Data Class, Specify:	
Description ==> DATA CLASS TAPE MEDIA 5, EE2 TRACK, COMPACTED, ENC, SCALED	
==>	
Recfm	(any valid RECFM combination or blank)
Lrecl	(1 to 32761 or blank)
Space Avgrec	(U, K, M or blank)
Avg Value	(0 to 65535 or blank)
Primary	(0 to 999999 or blank)
Secondary	(0 to 999999 or blank)
Directory	(0 to 999999 or blank)
Retpd or Expdt	(0 to 9999, YYYY/MM/DD or blank)
Volume Count 1	(1 to 255 or blank)
Add'l Volume Amount	(P=Primary, S=Secondary or blank)
Command ==>	
F1=Help	F2=Split F3=End F4=Return F7=Up F8=Down F9=Swap
F10=Left	F11=Right F12=Cursor

Figure 7-18 Data Class Alter panel (1 of 5)

5. Specify the following information for the Data Class definition in the current SCDS:
 - **Retpd or Expdt**- Specify how long the data sets in this Data Class remain valid (Figure 7-18).
 - **Volume Count**- Specify the maximum number of cartridges that you expect to use to store a data set in this Data Class.
6. Figure 7-19 on page 163 shows the second page of the Data Class definition process. You use this panel to specify compaction.

DGTDCDC2		DATA CLASS ALTER	Page 2 of 5
SCDS Name . . . : DFSMS150.SCD.SYS32			
Data Class Name : HSMDC1M			
To ALTER Data Class, Specify:			
Data Set Name Type	_____	(EXT, HFS, LIB, PDS, Large or blank)	
If Ext	_____	(P=Preferred, R=Required or blank)	
Extended Addressability . . .	N	(Y or N)	
Record Access Bias	_____	(S=System, U=User or blank)	
Space Constraint Relief . . .	N	(Y or N)	
Reduce Space Up To (%) . . .	_____	(0 to 99 or blank)	
Dynamic Volume Count	_____	(1 to 59 or blank)	
Compaction	Y	(Y, N, T, G or blank)	
Spanned / Nonspanned	_____	(S=Spanned, N=Nonspanned or blank)	
Command ==> _____			
F1=Help	F2=Split	F3=End	F4=Return
F10=Left	F11=Right	F12=Cursor	

Figure 7-19 Data Class Alter panel (2 of 5)

- **Compaction**- Specify whether to use data compaction for data sets assigned to this Data Class. A modified and more efficient Ziv-Lempel algorithm, Stream Lossless Data Compression (SLDC), is used by 3592 tape drives. We recommend that you always set the compaction to Y. The compaction attribute overrides the system default located in PARMLIB member DEVSUPxx, but it is overridden by JCL specification TRTCH. The valid Data Class values for the compaction attribute are Y, N, T (TCOM), G(GEN), and blank. TCOM and GEN do not apply to tape.

7. On the third page of the Data Class definition (Figure 7-20), you provide the Media Type and Recording Technology. The EFMT3 and EEFMT3 formats have been added here.

DATA CLASS ALTER		Page 3 of 5
Command ==> _____		
SCDS Name . . . : DFRMM1.SCD.S.TEST		
Data Class Name : DCFORE3		
To ALTER Data Class, Specify:		
Media Interchange		
Media Type	5	(1, 2, 3, 4, 5, 6, 7, 8, 9, 10 or blank)
Recording Technology . .	E3	(18,36,128,256,384,E1,E2,EE2,E3,EE3 or ' ')
Performance Scaling . .	Y	(Y, N or blank)
Performance Segmentation		(Y, N or blank)
Block Size Limit		(32760 to 2GB or blank)
Recorg		(KS, ES, RR, LS or blank)
Keylen		(0 to 255 or blank)
Keyoff		(0 to 32760 or blank)
CIsize Data		(1 to 32768 or blank)
% Freespace CI		(0 to 100 or blank)
CA		(0 to 100 or blank)
Use ENTER to Perform Verification; Use UP/DOWN Command to View other Panels;		

Figure 7-20 Data Class Alter panel (3 of 5) extended with the EFMT3 and EEFMT3 format

- **Media Type** - Specify the tape cartridge type that you use for data sets associated with this Data Class. If you do not enter anything for this field (the field is blank), the library that has the most SCRATCH cartridges is selected. This field is optional if you only use one valid media type within a TS7680 library, MEDIA5.

However, the definition is mandatory to allow selection of a media pool for non-specific mounts if multiple media types are present. For a description of the existing media types, refer to Table 7-4 on page 159.

- **Recording Technology**- This field is optional since only EFMT (E1) is supported for the TS7680.
- **Performance Scaling**- The 3592 tape drives allow you to record the data on the initial one-fifth (20% per default) of the media when performance is your main consideration. If you want fast access to the Media5 3592 data cartridge, the option Performance Scaling=y allows you to keep the data on the initial 60 GB (for EFMT1), and 100/128 GB (for EFMT2/EFMT3) with MEDIA5. This function is dynamic. If the tape is returned to scratch and later reused, the cartridge is reformatted to its scaled or full capacity as indicated through the Data Class assigned.
- **Performance Segmentation**- Performance segmentation, if selected, divides the tape into two segments: One segment is a fast access segment to be filled first, and the other segment is additional capacity to be filled after the fast access segment.

Note: Performance scaling and performance segmentation are not supported with the TS7680 and will be ignored. The host's compression setting is also ignored by the TS7680. The TS7680 will attempt compression regardless of the setting.

8. Press F8 after you have entered or updated the Recording Technology.

On the following panel, you also need to enter the Key Labels and the Encoding for both Key Labels, as shown in Figure 7-21.

DGTDGDC8
DATA CLASS ALTER
Page 4 of 5

SCDS Name . . . : DFSMS150.SCDS.SYS32
Data Class Name : HSMDC1M

To ALTER Data Class, Specify:

Encryption Management

Key Label 1 . . . (1 to 64 characters or blank)
tape sol tst shr pvt 1024 lbl 01

Key Label 2 . . .
tape sol tst shr pvt 1024 lbl 02

Encoding for Key Label 1 L (L, H or blank)
Encoding for Key Label 2 H (L, H or blank)

Command ==> _____

F1=Help F2=Split F3=End F4=Return F7=Up F8=Down F9=Swap
F10=Left F11=Right F12=Cursor

Figure 7-21 Data Class Alter panel (4 of 5)

If you change existing Data Classes, verify your Data Class ACS routine to make sure that you are assigning the correct constructs. If you create new Data Classes, update your Data

Class ACS routine to have the new constructs assigned to those tape data sets that you want to have encrypted.

To activate the new SMS definitions:

- ▶ Translate the Data Class ACS routine.
 - ▶ Validate the ACS routines.
 - ▶ Activate the SMS SCDS.
9. In the fifth page of the Data Class definition (Figure 7-22), there are no options that apply to tape data sets. We merely provide this panel for your convenience.

DGTD CDC6
DATA CLASS ALTER
Page 5 of 5

SCDS Name . . . : DFSMS150.SCD.SYS32

Data Class Name : HSMDC1M

To ALTER Data Class, Specify:

Shareoptions Xregion . . . -	-	(1 to 4 or blank)
Xsystem . . . -	-	(3, 4 or blank)
Reuse	N	(Y or N)
Initial Load	B	(S=Speed, R=Recovery or blank)
BWD	-	(TC=TYPECICS, TI=TYPEIMS, NO or blank)
Log	-	(N=NONE, U=UNDO, A=ALL or blank)
Logstream Id	-	
FRlog	-	(A=ALL, N=NONE, R=REDO, U=UNDO or blank)
RLS CF Cache Value	A	(A=ALL, N=NONE, U=UPDATESONLY)
RLS Above the 2-GB Bar	N	(Y or N)
Extent Constraint Removal	N	(Y or N)

Command ==>

F1=Help
F2=Split
F3=End
F4=Return
F7=Up
F8=Down
F9=Swap

F10=Left
F11=Right
F12=Cursor

Figure 7-22 Data Class Alter panel (5 of 5)

Note: Remember that the Data Class ACS routine is driven for both system-managed and non-system-managed data sets.

DFHSM Percent Full recommendations

Problems can arise when using DFHSM's Tape Copy Utility if the target volume signals that it is almost full with additional data still to be copied. The TS7680 (or any tape device) indicates that the tape volume is almost full by signaling a unit exception shortly before the tape drive has reached the end of the volume.

The part of a tape volume between when the unit exception begins to be reported and when a physical end of volume error is reported is called the Logical End of Volume (LEOV) region. For most tape drives, the LEOV region can store between 100s of KB to a few MB. For the TS7680, it is 8 MB. The amount that the region can store is more than enough for the host application to write its last few records and close the volume with necessary trailer labels. So, it is expected that when an application sees the unit exception, it finishes the writes that it needs and closes the volume. This method is the typical way end-of-volume handling is performed.

For HSM, when it is writing its original volume, it handles the LEOV indication as expected. This results in a volume that has just a little more data on it past the LEOV point. For HSM

tape copy, it treats the LEOV indication as a failure condition, thinking that it has run out of space and that the target volume is too small.

To avoid this issue, the Percent Full Utilization values that were recommended in the past need to be changed so that HSM will consider a volume full before it enters the LEOV region. Then when it copies the volume, there will be no issues with encountering the LEOV region.

To complicate this a bit, the Percent Full values need to take into account the media type and the logical volume size being used. Virtual volumes appear to HSM as either Media 1 or Media 2 (determined at insert time when the virtual volume was created).



Part 3

Managing the TS7680

This part covers the daily management tasks associated with the TS7680, including such activities as monitoring and reporting.

Archived

Managing and administering TS7680

The topics in this chapter provide information about how to manage the 3958-DE2 (TS7680), including:

- ▶ General managing
- ▶ Power managing
- ▶ Managing the Enterprise controller
- ▶ Managing the ProtecTIER server

The topics in this section define tasks used to manage the ProtecTIER server.

- ▶ Using the ProtecTIER Manager
- ▶ Managing the ProtecTIER Virtual Tape system
- ▶ Managing virtual libraries
- ▶ ProtecTier system (errors, logs, and recovery task)

The topics in this section provide tips, guidelines, and tasks for handling error situations that might arise in the ProtecTIER system.

8.1 General management

The topics in this section provide details for the general managing of the 3958-DE2 (TS7680).

8.1.1 About the TS3000 System Console and its related components

This topic describes the components of the TS3000 System Console (TSSC).

When a TSSC (FC 2722) is purchased with the 3958-DE2 (TS7680), the TSSC and its related components (keyboard, video, mouse [KVM], KVM switch, and 16-port network switch) are factory-installed in the frame. As a result, the KVM switch provides input and video

connections for each ProtecTIER server in the frame and no USB keyboard or graphics-capable monitor is needed for the 3958-DE2 (TS7680) configuration.

8.1.2 User login on IBM TS3000 System Console

Complete this task to log in to the TS3000 Console.

The keyboard, video, mouse (KVM) switch configuration is done in manufacturing.

1. On the TS3000 System Console (TSSC) keyboard, press **PrtSc** to launch the KVM switch's Main panel. The Main panel displays. You can select TSSC with the Up or Down arrow on the keyboard; see Figure 8-1.




Figure 8-1 Main panel 3958 (TSSC selected)

Note: At the main panel, click **Setup** to change the name.

The console login panel is displayed; see Figure 8-2 on page 171.

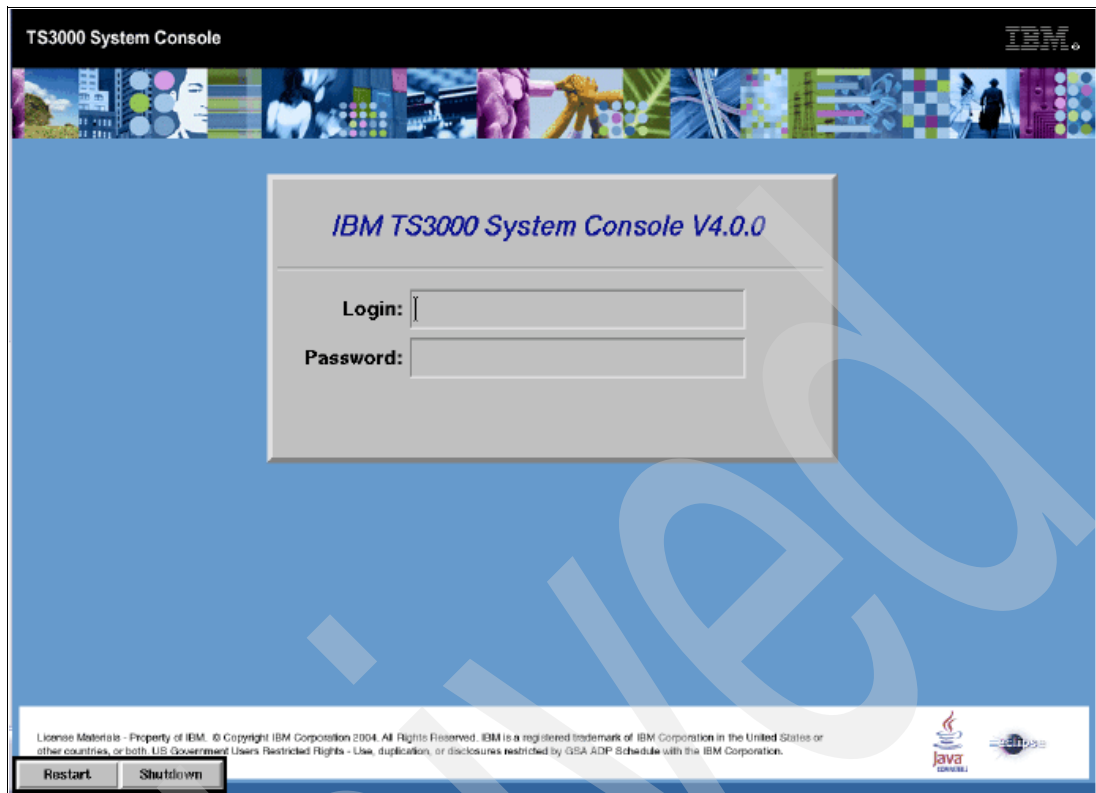


Figure 8-2 Console Login window, example. The version number (Vx.x.x) reflects your level of microcode

2. Type service. Press the Tab key to move your cursor to the Password field.
3. Enter the case-sensitive password service and press Enter.

From here right-click the TSSC's blue desktop for the TSSC System Console Action menu. See Figure 8-3 on page 172 for an example.



Figure 8-3 TSSC System Console

8.1.3 Logging into and out of the ProtectTIER Manager

Complete this task to log into and log out of the ProtectTIER Manager.

ProtectTIER Manager has default user accounts corresponding to three user permission levels: Administrator, Operator, and Monitor. The default usernames and passwords for each of these accounts are shown in Table 8-1.

Table 8-1 Default usernames and passwords (ProtectTIER Manager)

Permission Level	Default Username	Default Password
Administrator	ptadmin	ptadmin
Operator	ptoper	ptoper
Monitor	ptuser	ptuser

Log in to each ProtectTIER cluster that you want to manage using the ProtectTIER Manager.

1. Click **Login**. The Login dialog box is displayed.



Figure 8-4 ProtecTIER Login

2. Enter your username and password; see Figure 8-5.



Figure 8-5 ProtecTIER user, password

3. Click **Ok**. The Login dialog box closes and you are logged into the ProtecTIER Manager.

It is recommended that you change or replace these default user accounts.

Note: Only one Administrator can be logged into a ProtecTIER cluster at a time. It is, therefore, recommended that you log out at the end of each session by clicking **Logout**.

If you log in with Administrator level permission while another Administrator is already logged in, a message box is displayed; see Figure 8-6.

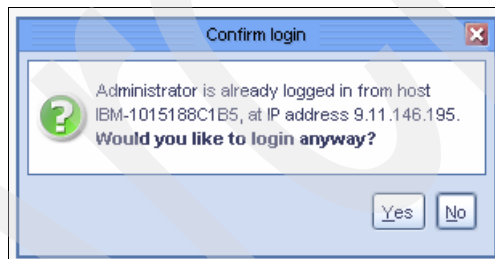


Figure 8-6 Administrator already logged

Click **Yes** to force the other Administrator to log out.

8.1.4 Changing user account passwords in the ProtecTIER Manager system

Complete this task to change the user account password in the ProtecTIER Manager system.

Ensure that you have your original password because you must use it to change to a new password. If the original password is not available, the administrator must delete the user account and create a new account.

Perform the following steps to change the user account password:

1. Log into the ProtecTIER Manager system and select **System** → **Users Management**. The Users Management dialog is displayed.
2. Select a user account and click **Change password**. The Change password dialog is displayed.
3. Type the user account's original password in the password field.
4. Type the new password in the New password field.
5. Type the new password in the Verify password field.
6. Click **Ok** and the new password is added to the ProtecTIER Manager system. Make a note of the new password and inform the user of what it is. Ensure that it is recorded for later use.

8.1.5 Adding a user to the ProtecTIER Manager system

Complete this task to add a user to the ProtecTIER Manager system.

Only a person assigned the administrator permission level can add a user to the ProtecTIER Manager system.

Perform the following steps to assign a new user account:

1. Log into the ProtecTIER Manager system and select **System** → **Users Management**. The User Management window is displayed.
2. Click **Add**. The Add account dialog is displayed.
3. Complete the input fields on the dialog window:
 - User name** - Type a user name that complies with your company policy for user names.
 - New password** - Type a password that complies with your company policy for passwords.
 - Verify password** - Enter the same password again to verify the password.
 - Permission** - Input the permission level (administrator, operator, or monitor) you want to assign the user.

Note: Before processing, record this information in a manner that allows you to notify the new user. Remind the new user to record the password because neither you nor the system stores the passwords for later use.

4. Click **OK**. The new account is added to the ProtecTIER Manager system.

You can delete a user by selecting the user account from the User Account list and clicking **Remove**.

8.1.6 Logging in to the ProtecTIER server

Complete this task to log in to the ProtecTIER server (example for Upper).

The keyboard, video, mouse (KVM) switch configuration is done in manufacturing.

1. On the TS3000 System Console (TSSC) keyboard, press **PrtSc** to launch the KVM switch's Main panel. The Main panel displays. You can select Upper 3958 server with the Up or Down arrow on the keyboard; see Figure 8-7.




Figure 8-7 Main panel 3958 (Upper server selected)

The Server command prompt, Log in, is displayed (as shown in Figure 8-8 on page 175).

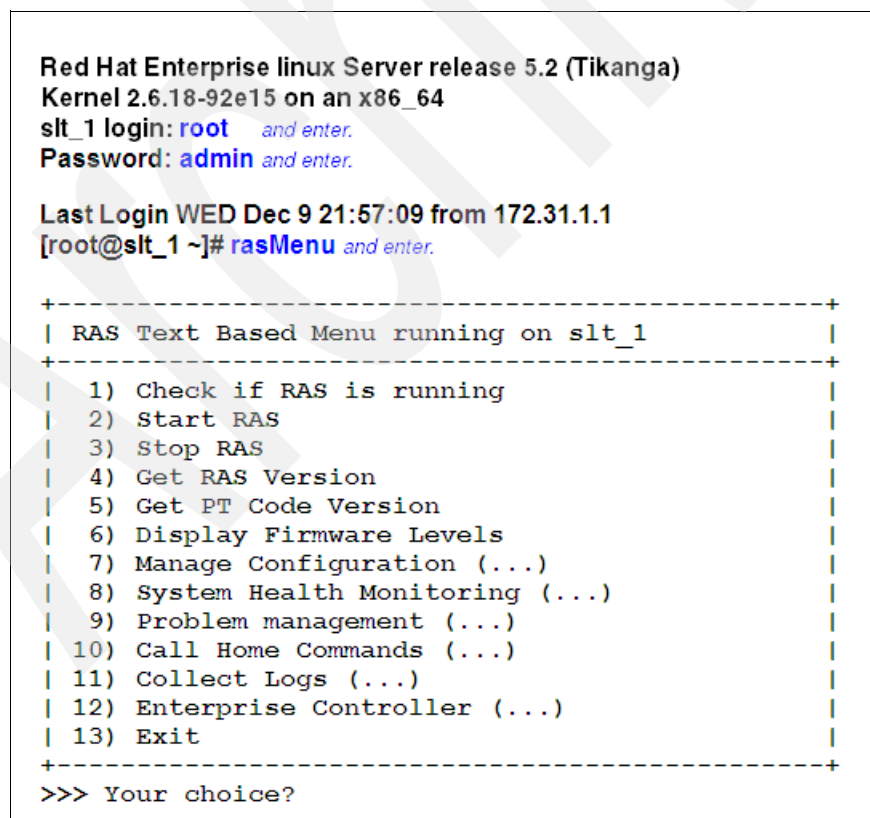


Figure 8-8 Red Hat Login

2. Type root and press Enter.
3. Type **admin** as password and press Enter.
4. Initiate the ProtecTIER service menu by typing `rasMenu` at the command line and pressing Enter. The RAS Text Based Menu is displayed.

8.1.7 Host reporting

The topics in this section describe the statistical analysis and reporting system messages that apply to the host.

Service information messages

This topic describes Service Information Message (SIM) functions as primary factors for improved product availability.

A SIM alerts you when an abnormal operational condition in the Enterprise controller requires service attention. Information in the SIM identifies the affected drive, the failing component and severity of its fault condition, and the expected operational impact of the pending service action. This information helps the user to initiate and expedite appropriate recovery and service procedures so that normal operation is restored with maximum efficiency and minimal disruption.

A SIM contains the machine type, machine serial number, and Field Replaceable Unit (FRU), which allows the dispatch of the appropriate service personnel and the replacement parts required to correct the machine fault. This procedure helps improve service response time and helps reduce the time required for machine repair.

You can select the severities you wish to see, depending on your software. For example, you may only want to see the acute SIM, or prefer to see all SIMs sent to the host. SIM filtering by severity can be done through configuration options. The four severity codes are:

- ▶ Severity 0 (FID4) means the device requires service, but normal drive function is not affected.
- ▶ Severity 1 (FID3) indicates moderate severity.
- ▶ Severity 2 (FID2) indicates serious severity.
- ▶ Severity 3 (FID1) indicates acute severity.

SIMs can be reported multiple times. A configuration option allows reporting the same SIM more than once. The time between repeat SIMs is eight hours. For example, if the configuration option is set to two, a SIM is reported when an error occurs, is repeated again eight hours later, and, then again, eight hours after that message. The default is to not repeat SIMs.

Service Information Message presentation

This topic lists the specific Service Information Message (SIM) presentation by the system.

SIMs reporting varies for different systems. The following SIM presentations apply to the 3598-DE2 (TS7680) system:

- ▶ System Presentation zSeries® (S/390®)
IEA480E and IEA486E messages, as well as Environmental Record Editing and Printing (EREP) reports
- ▶ AIX

SIM messages are logged to EREP reports.

- Linux

Via taped DAEMON, SIM messages are logged to /var/log with file names of the form <drive serial #>.<time stamp>.sim.

8.2 Power management

The topics in this section provide details for powering off and on the components of 3958-DE2 (TS7680).

8.2.1 Shutting down the 3958-DE2 (TS7680)

Complete this task to shut down the 3958-DE2 (TS7680) (entire frame).

Shutting down the 3958-DE2 (TS7680) involves shutting down each component before powering off the 3958-DE2 (TS7680) frame.

It is necessary to perform the following steps in order:

1. Vary all devices offline at the hosts (128 devices for each Enterprise controller); see 8.3.1, “Varying devices offline” on page 184.
2. Shut down the upper Enterprise controller; refer to 8.2.2, “Shutting down the Enterprise controller” on page 179.
3. Shut down the lower Enterprise controller. Use the same procedure as in Step 2.

Wait until the green lights on both of the Enterprise controllers start flashing. The shutdown of the Enterprise controllers is complete. Continue with the following set of steps to shut down the ProtecTIER servers next. Be sure to begin with the upper ProtecTIER server.

4. Enter the **poweroff** command at the upper ProtecTIER server; refer to 8.2.3, “Shutting down the ProtecTIER server” on page 181.
5. Enter the **poweroff** command at the lower ProtecTIER server. Use the same procedure as in Step 4.

Take the appropriate action depending on the location of the disk repository:

- If the disk repository is located in a separate frame, proceed to the following step.
- If the disk repository is located inside the 3958-DE2 (TS7680) frame, then follow the power off procedures from the disk provider to power off the repository now. Then return here and proceed with the following step.

The shutdown of the ProtecTIER servers is complete. Continue with the following set of steps to shut down the TSSC next.

6. Log out of the TSSC by performing the following substeps:
 - a. Right-click the desktop to get the Main Menu.
 - b. Select **Logout**. The Login window displays (Figure 8-9).

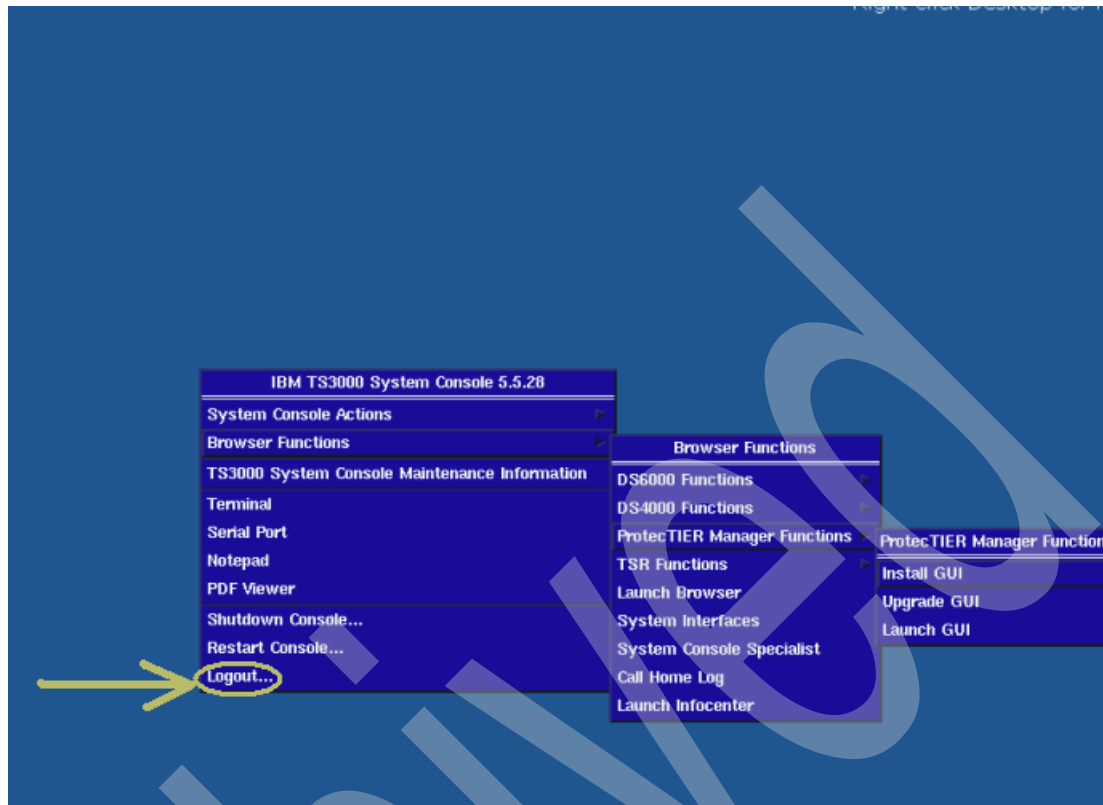


Figure 8-9 TSSC Logout

7. At the Login window, select **Shutdown** located at the bottom left portion of the window to begin the shutdown process (see Figure 8-10 on page 179).

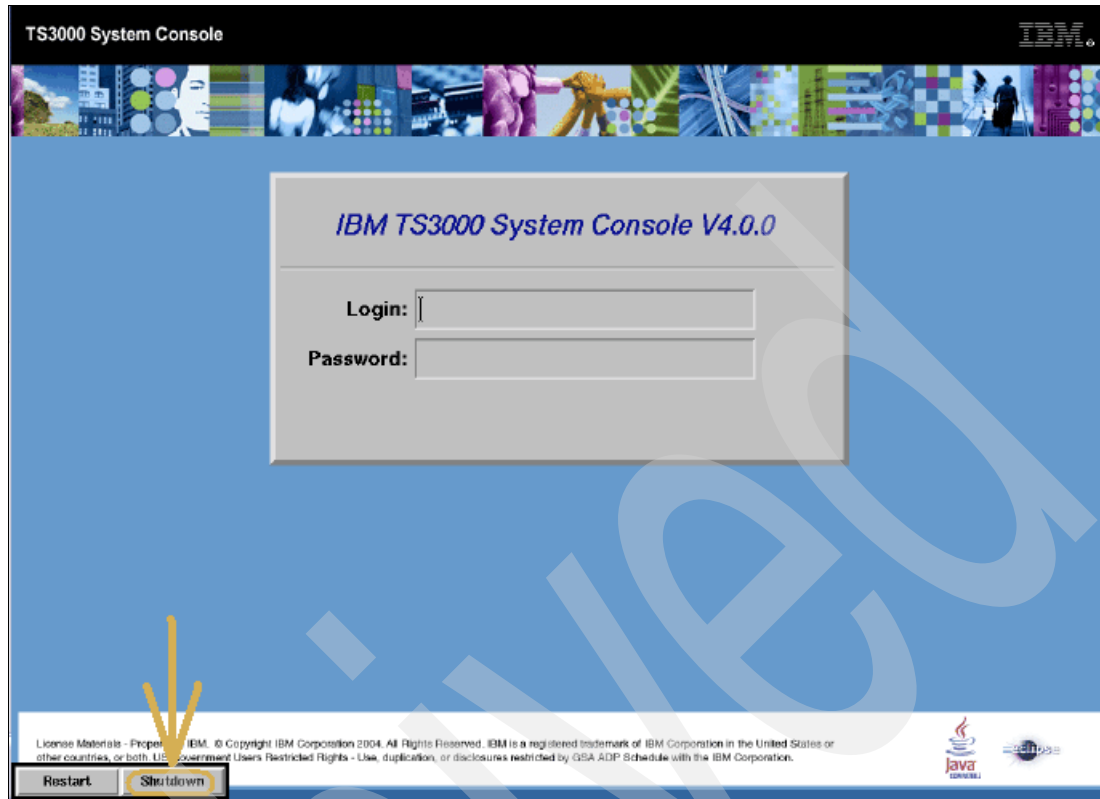


Figure 8-10 TSSC Shutdown

8. Press the TSSC power button to power off when prompted to do so on the window.

The shutdown of the TSSC is complete. Continue with the following step to shut down the Power Control Assembly (PCA) and 3958-DE2 (TS7680) frame.

9. Turn the EUPO switch on the front on the 3958-DE2 (TS7680) frame to the OFF position. This completes this shutdown of the 3958-DE2 (TS7680).

8.2.2 Shutting down the Enterprise controller

This topic defines the shutdown procedure for the Enterprise controller.

1. Ensure that all devices associated with the Enterprise controller are varied offline in preparation for shutdown. Refer to 8.3.1, "Varying devices offline" on page 184.
2. Telnet to both ProtecTIER servers using Telnet to Tape Systems on the TS3000 System Console (TSSC) by performing the following substeps:
 - a. Right-click the TSSC desktop to display the Main Menu.
 - b. Select **System Console Actions**, then select **Telnet to Tape Systems**.
 - c. Select the appropriate ProtecTIER server (Upper or Lower) and type the password: ptadmin.
3. From the ProtecTIER server, type rasMenu and press Enter. The RAS Text Based Menu displays (Example 8-1 on page 180).

Example 8-1 Ras Text Based Menu

```
+-----+
| RAS Text Based Menu running on slt_2 |
+-----+
| 1) Check if RAS is running           |
| 2) Start RAS                        |
| 3) Stop RAS                         |
| 3) Stop RAS                         |
| 4) Get RAS Version                  |
| 5) Get PT Code Version              |
| 6) Display Firmware Levels          |
| 7) Manage Configuration (...)       |
| 8) System Health Monitoring (...)   |
| 9) Problem management (...)         |
| 10) Call Home Commands (...)        |
| 11) Collect Logs (...)              |
| 12) Enterprise Controller (...)      |
| 13) Exit                            |
+-----+
>>> Your choice?
```

4. Enter the number for Enterprise Controller and press Enter. The Enterprise Controller submenu displays (see Example 8-2).

Example 8-2 Submenu of Enterprise Controller

```
+-----+
| > Sub-menu of 'Enterprise Controller'|
+-----+
| 1) Notify the Enterprise Controller  |
|    the PT server is going down      |
| 2) Notify Enterprise Controller the  |
|    PT server is up                  |
| 3) Shutdown Enterprise Controller    |
| 4) Execute timezone check           |
| 5) Back to parent menu              |
| 6) Exit                             |
+-----+
>>> Your choice?
```

Note: Remove all CDs and DVDs from the Enterprise controller by pressing the eject button on each tray for each controller before you proceed.

5. Enter the number for Shutdown the Enterprise Controller and press Enter. Then, type y to answer yes to both questions (see Example 8-3).

Example 8-3 Shutdown response

```
Are you sure you would like to continue (y/n)? y
Do you want to shut down both nodes (y/n)?
y
```

Note: This process will take about 5 minutes. Do not press CTRL-C or exit this process.

Wait until the green lights on the Enterprise controllers start flashing. The shutdown of the Enterprise controllers is complete.

If the above procedure is not working:

Press and hold down the Power ON/OFF button 6 for approximately 5 seconds.

Note: In the upper right corner of the Operator Panel, when the number decrements from 4 to 1 and then no number, release the button.

Be patient. If the power down process stops before reaching the standby mode, repeat Step 1.

The Enterprise controller is powered down to standby mode when the green LED 5 on the Operator Panel blinks and OK is displayed 1 (see Figure 8-11).




Figure 8-11 Enterprise controller operator panel

When the standby mode is reached, you can remove the power from the Enterprise controller.

8.2.3 Shutting down the ProtecTIER server

This topic defines the shutdown procedure for the ProtecTIER server.

1. Telnet to both ProtecTIER servers using Telnet to Tape Systems on the TS3000 System Console (TSSC) by performing the following substeps:
 - a. Right-click the TSSC desktop to display the Main Menu.
 - b. Select **System Console Actions**, then select **Telnet to Tape Systems**.
 - c. Select the appropriate ProtecTIER server (upper or lower) and type the password: ptadmin.
2. Enter the **poweroff** command at the lower ProtecTIER server; see Figure 8-12 on page 182.

```
Last login: Thu Dec 10 22:34:54 2009 from 9.11.144.151
slt_2: /opt/dtc/ptadmin> poweroff
```

Figure 8-12 Poweroff command

8.2.4 Powering on the 3958-DE2 (TS7680)

Complete these tasks in order to power on the 3958-DE2 (TS7680) (entire frame).

In order to power on the 3958-DE2 (TS7680), you must power up the component in the order shown here.

Note: Before you begin these tasks, verify that the disk subsystems (backend disk repository) are powered up and ready for use.

- ▶ Switch on the frame circuit breakers (main wall power switch).
- ▶ Power Control Assembly (PCA) (turn the EUPO switch on the front of the frame to the ON position.)
- ▶ Verify power on of the following items (from rear of the frame):
 - WTI network power switch
 - Enhanced Routers
 - TS3000 System Console (TSSC), KVM, and KVM switch
 - Backend disk repository system
- ▶ ProtecTIER server A (Lower); refer to 8.2.5, “Powering on the ProtecTIER servers”.
- ▶ ProtecTIER server B (Upper); use the same procedure.
- ▶ Enterprise controller (lower controller); refer to 8.2.6, “Powering on the Enterprise controllers” on page 183.
- ▶ Enterprise controller (upper controller); use the same procedure.
- ▶ Visually inspect the indicator and fault LEDs on all 3958-DE2 (TS7680) components.

If all link-up indicators and fault LEDs show normal operation, close the front and rear doors of the frames.

If an amber LED on any component is lighted, refer to the documentation for that component to diagnose and remedy the problem.

8.2.5 Powering on the ProtecTIER servers

Complete this task to power on the ProtecTIER servers.

When the ProtecTIER server is connected to an AC power source but is not turned on, the operating system does not run, and all core logic except for the service processor is shut down. However, the ProtecTIER server can respond to requests from the service processor, such as a remote request to turn on the ProtecTIER server. The power-on LED flashes to indicate that the ProtecTIER server is connected to AC power, but the ProtecTIER server is not powered on.

Important: Power on the lower ProtecTIER server (server A) first, then power on the upper ProtecTIER server (server B).

Approximately 20 seconds after the ProtecTIER server is connected to power, the power-control button becomes active, and one or more fans might start running to provide cooling while the ProtecTIER server is connected to power.

To power on the ProtecTIER server, complete the following steps:

1. Press the white, recessed power-control button on the ProtecTIER server operator panel; refer to Figure 8-13.




Figure 8-13 Server operator information panel

2. Connect to the server; refer to 8.1.6, “Logging in to the ProtecTIER server” on page 174.
3. When the screen displays registered calypsa with major device 249, press Enter.
4. Login input should be displayed.

8.2.6 Powering on the Enterprise controllers

This task is part of the 3958-DE2 (TS7680) power-on procedure. Ensure that you are familiar with the information in Powering on the 3958-DE2 (TS7680) before continuing with this task.

Important: Power on the lower Enterprise controller first, then power on the upper Enterprise controller.

Perform the following steps to power on the Enterprise controllers:

1. Verify that the Enhanced Routers are powered on.
2. Verify that the two line cords are connected to the Enterprise controllers' power supplies.
3. When the Enterprise controllers are in standby mode (blinking green LED 2 on operator panel), press the power ON button 1, which is located on the Enterprise controller operator panel. See Figure 8-14 on page 184. When the controller has powered on successfully, the controller panel displays READY on. This could take up to 15 minutes.




Figure 8-14 Enterprise Controller's operator panel

8.3 Managing the Enterprise controller

The topics in this section describe managing the Enterprise controller.

8.3.1 Varying devices offline

Complete this task to vary devices offline at the host.

Before the Enterprise controllers are varied offline for any procedure, it is first necessary to vary the associated devices offline at the host. This ensures that all jobs are stopped before the controller goes offline.

Perform the following steps at the host(s) in order to vary the devices offline:

1. Display the tape subsystems (Figure 8-4) and devices associated with each Enterprise controller with the MVS operator command:

```
DEVSERV QLIB,library-ID,
```

Example 8-4 Displaying the tape subsystems

```
DEVSERV QLIB,CA030, where CA030 is the 5-CHAR library-ID.
```

Note: Tape subsystems X'01' through X'08' are associated with the lower Enterprise controller and represent 128 devices. Tape subsystems X'11' through X'18' are associated with the upper Enterprise controller and also represent 128 devices.

2. Vary the devices offline at the host(s). It may only be necessary to vary offline a subset of devices, or only the devices associated with one of the Enterprise controllers, depending on what procedure you perform next.

Note: On each Enterprise controller, tape devices rmt00priA through rmt3FpriB are associated with the HBA card in slot 1. On each Enterprise controller, tape devices rmt40altA through rmt7FaltB are associated with the HBA card in slot 4.

- Interface priA(0A-08) correctly configured with 32 devices, 00 - 1F
- Interface priB(0A-09) correctly configured with 32 devices, 20 - 3F
- Interface altA(00-08) correctly configured with 32 devices, 40 - 5F
- Interface altB(00-09) correctly configured with 32 devices, 60 - 7F

8.3.2 Notifying the Enterprise controller of change in ProtecTIER server status

Complete this task to notify the Enterprise controller before varying the ProtecTIER server offline and after the ProtecTIER server is online again.

Certain procedures, such as performing a code upgrade or expanding the repository, require that one or both of the ProtecTIER servers are varied offline. If the Enterprise controller is not notified in advance, it automatically sends a Call Home when a ProtecTIER server is varied offline. In order to prevent this unnecessary Call Home, you must first notify the Enterprise controller. It is also necessary to notify the Enterprise controller after the ProtecTIER server is online again.

Complete the following steps to notify the Enterprise controller before the ProtecTIER server is varied offline for a code load or repository expansion. Also complete these steps to notify the Enterprise controller after the ProtecTIER server is back online after a code load or repository expansion is complete.

Notes:

- ▶ You must first ensure that you have a connection to the server. To establish a connection:
 - Activate the keyboard, video, mouse (KVM) switch and use the TS3000 System Console (TSSC).
 - Open a command prompt on the TS3000 System Console (TSSC). Type `tsys` and press Enter. At the selection menu, type the number of the ProtecTIER server to connect to and press Enter.
- ▶ The devices associated with the affected Enterprise controller should be varied offline from the host before beginning this procedure. Refer to 8.3.1, “Varying devices offline” on page 184.
- ▶ This task must be performed for each ProtecTIER server.

1. Log in to the ProtecTIER server at the local host login: prompt with the ID `root` and the password `admin`. If you are a service representative and the root user is blocked or the password does not work, contact the client.
2. Initiate the ProtecTIER service menu by typing `rasMenu` at the command line and pressing Enter. The RAS Text Based Menu displays.
3. At the `>>> Your choice?` prompt, type the number for the Enterprise Controller and press Enter. The Enterprise controller submenu displays.
4. Make the appropriate selection based on whether the ProtecTIER server is being varied offline or is back online.
 - Offline: If you are varying the ProtecTIER server offline, type the number for Notify Update PT code start at the `>>> Your choice?` prompt and press Enter.
 - Online: If the ProtecTIER server is back online, type the number for Notify Update PT code end at the `>>> Your choice?` prompt and press Enter.

Note: After the code load or repository expansion is complete, you must repeat these steps to notify the Enterprise controller after the ProtecTIER server is back online.

8.3.3 Performing a standalone mount

Complete the following task to mount a standalone logical volume. When the host cannot send mount commands to the 3958-DE2 (TS7680), it might need to perform standalone operations by using these standalone procedures.

Perform the following steps:

1. Determine the device to be used to perform the standalone mount and initial program load (IPL).
2. Vary the device offline on all hosts except the host that is going to use this device to IPL from. The drive that is being used in the standalone mode must be varied offline from all hosts except the host that is being used in this special mode. This prevents unwanted interaction from all hosts except the desired one.
3. Map the device selected in step 1 to a device on the ProtecTIER system using Table 8-2. The selected device has an associated subsystem ID and device number that maps to a drive number on the ProtecTIER system. You have to know how to map the drive address on the ProtecTIER graphical user interface (GUI) to a host device address. From the PT GUI perspective, drives 0 - 255 maps are shown in Table 8-2.

Table 8-2 Mapping the drive address to the host device address

PT GUI	Subsystem ID	Device number
0 - 15	0x01	0x00 - 0x0F
16 - 31	0x02	0x00 - 0x0F
32 - 47	0x03	0x00 - 0x0F
48 - 63	0x04	0x00 - 0x0F
64 - 71	0x05	0x00 - 0x0F
80 - 95	0x06	0x00 - 0x0F
96 - 111	0x07	0x00 - 0x0F
112 - 127	0x08	0x00 - 0x0F
128 - 143	0x11	0x00 - 0x0F
144 - 159	0x12	0x00 - 0x0F
160 - 175	0x13	0x00 - 0x0F
176 - 191	0x14	0x00 - 0x0F
192 - 207	0x15	0x00 - 0x0F
208 - 223	0x16	0x00 - 0x0F
224 - 239	0x17	0x00 - 0x0F
240 - 255	0x18	0x00 - 0x0F

4. From the ProtecTIER Manager GUI Library page, under the slots tab, select the slot that contains the VOLSER to be mounted. Right-click the line and select **Move Cartridges**. A dialog box is displayed.
5. In the dialog box, and in the destination type, select the drive and in the destination field select **Drive No.** (standalone drive number.) Press **OK** to continue. This causes ProtecTIER to mount the VOLSER on the specified drive. The drive is ready for the host, which enables the host to IPL from this device.

8.3.4 Performing a standalone demount

Perform the following steps to demount a tape on a standalone device.

1. Access the ProtecTIER Manager GUI Library page, under the Drives tab. It is presumed that you are already logged into the ProtecTIER Manager.
2. Scroll down the list of drives and select the drive number that contains the VOLSER to be demounted. Right-click the drive number line and select **unload drive**.

Note: If the drive is reserved, the unload operation is failed. To unload this drive and complete a demount process, you must first reset the drive.

3. Right-click the drive number selection again and select **move cartridges**. A dialog box is displayed.
4. In the dialog box, select **slot** for the destination type and **next available** for the destination field. Then click **OK** to complete the demount process.

8.4 Managing the ProtecTIER server

The topics in this section define tasks to manage the ProtecTIER server.

Important: The time setting on the ProtecTIER server is synchronized with all subsystems within the 3958-DE2 (TS7680) frame. Do not change the time or the time zone on the ProtecTIER server. Changing the time or time zone may result in data loss if not done properly. If it is necessary to change the time, call software support.

8.4.1 Adding and removing nodes from ProtecTIER Manager

The topics in this section describe how to add a node, a subnetwork node, and remove a node from the ProtecTIER system using the ProtecTIER Manager.

Adding a node registers the node's IP address and port number with the instance of ProtecTIER Manager at your workstation. Similarly, removing a node removes the node's registration from ProtecTIER Manager at that workstation.

Note: The 3958-DE2 (TS7680) is a dual-node system. It cannot be run in single-node mode.

Adding nodes

Complete this task to add a node to your ProtecTIER system using the ProtecTIER Manager.

Perform the following steps to add a node to your ProtecTIER system:

1. Run the ProtecTIER Manager application.
 - For a Windows-based ProtecTIER Manager workstation, run the ProtecTIER Manager application:
Click **Start** → **Programs** → **IBM** → **ProtecTIER Manager x.x** → **IBM ProtecTIER Manager**.
 - For a Linux-based ProtecTIER Manager workstation, click the icon for ProtecTIER Manager on the desktop or from the location of the shortcut that you selected during the installation.

The ProtecTIER Manager window is displayed (see Figure 8-15).

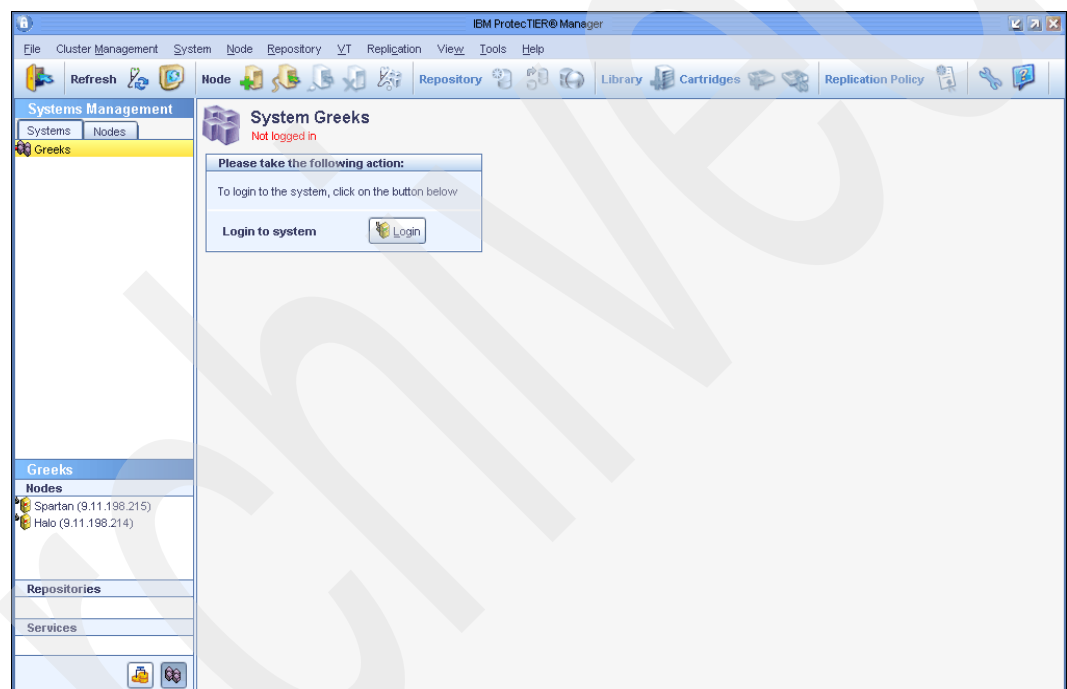


Figure 8-15 ProtecTIER Manager window

2. Click **Add node**. The Add node dialog box is displayed, prompting you for the IP address and Port number of the node that you want to add.
3. Enter the IP address of the node and click **OK**. The node is displayed in the Nodes pane and the Login button is displayed in the View pane.

Note: Do not change the port number of the node unless directed to do so by IBM Support.

4. Click **Login**. You are prompted for your username and password.
5. Enter your username and password and click **OK**. ProtecTIER Manager displays the information for that node. If the node has a repository, the node's cluster is displayed in the Systems tab of the Navigation pane. If that cluster already contains a second node, the second node is also displayed in the Nodes pane.

Adding node subnetworks

This topic describes how to add a node subnetwork to your ProtecTIER system using the ProtecTIER Manager.

You can use the ProtecTIER Manager to add addresses for subnetworks to which nodes are connected. When ProtecTIER Manager restarts, it automatically detects all nodes on the added subnetworks of the TCP/IP network.

Perform the following steps to add a subnetwork address to your ProtecTIER system:

1. Log into the ProtecTIER Manager.
2. Select **Tools** → **Preferences** from the toolbar. A confirmation message box is displayed.
3. Click **Yes**. The Preferences dialog box is displayed.
4. For each subnetwork you want to add, click a **Sub network** checkbox and enter the subnetwork address in the corresponding field.
5. Click **Ok**. The Preferences dialog box closes and the subnetwork address is added to the ProtecTIER Manager. When you restart the ProtecTIER Manager, all nodes on the defined subnetwork addresses are automatically added to the ProtecTIER Manager.

Removing nodes

This topic describes how to remove a node from your ProtecTIER system using the ProtecTIER Manager.

Removing a node stops the instance of ProtecTIER Manager at your workstation from registering the node and being able to manage it. The node itself, and the cluster with which the node is associated, is unaffected by removing a node in this way.

Note: In a two-node cluster, the second node is also removed.

Perform the following steps to remove a node:

1. Log into the ProtecTIER Manager.
2. Click the **Nodes** tab in the Navigation pane. The Node pane is displayed.
3. In the Node pane, select the node that you want to remove.
4. Click **Remove node**. A confirmation message box is displayed.
5. Click **Yes**. The node is removed.

8.4.2 Monitoring nodes using the ProtecTIER Manager

This topic defines how to monitor nodes using ProtecTIER Manager.

Open the IBM ProtecTIER Manager and click **Nodes**. The Node window is displayed with the following selections:

Port attributes

The Port attribute window displays the type of information shown in Table 8-3 on page 190 for each port in the node.

Table 8-3 Port attribute information

Column	Definition
Port	Indicates the port number and port mode. All ports in the ProtecTIER system are front-end ports that connect the node with the System z hosts. This node is labeled FE.
WWNN	The Worldwide Node Name of the port. Note: If you need to find the Worldwide Node Name, you can click Scan to open the Scan Port. The Scan Port window displays a numbered list of the Worldwide Node Names of the remote ports detected by the port.
Link Speed	Indicates the transmission speed of the port.
Topology	Fibre Channel topology of the port. The possible values are as follows: <ul style="list-style-type: none"> ▶ Point-to-point (P2P) ▶ Fibre channel-arbitrated loop ▶ Down There is no fibre channel connection.
User setup	User-assigned link speed and topology.

Version information

The Version information window displays the version numbers for the versions of ProtecTIER, the PT Linux RPM, and the DTC Emulex RPM installed and running on the node. In addition, click **Show fixes** to display the list of temporary fixes featured in the installed ProtecTIER version.

Fibre channel ports throughput

An instance of the Fibre Channel Ports Throughput window is displayed for each node. This window displays the rate of data movement and I/O operations for both reading and writing operations for the node. The data movement rate is also displayed graphically for each front-end fibre channel port on the node. You can change the scale of the graph by editing the value in the Scale graph field.

Network configuration

The Network configuration window displays information about the setup of the network interface cards (NIC) for the node (see Table 8-4).

Table 8-4 Network interface card information

Column	Definition
Device	The devices in the NIC. <ul style="list-style-type: none"> ▶ Eth0 is the node port that communicates with the ProtecTIER Manager workstation. ▶ Eth1 and Eth2 are the node ports used in the cluster-internal network. ▶ Bond0 is the virtual bond master device to which Eth1 and Eth2 are enslaved. Note: Bond devices are defined as part of the installation process.
IP Address	IP address of the device.
Usage	Indicates whether the device is used for the cluster-internal network or to communicate with the ProtecTIER Manager workstation and the external network.
Master device	The master device or bond device, if any, to which the device is enslaved.
Status	Indicates whether the device is functioning properly.

Column	Definition
Speed	The supported speed of data transfer across the device in megabits per second.
MTU	Configured maximum transmission unit for the device.

8.4.3 Upgrading the ProtecTIER software

The topics in this section explain how to upgrade the ProtecTIER software in the 3958-DE2 (TS7680) system.

The ProtecTIER software upgrade can be a concurrent or a nonconcurrent procedure. The concurrent procedure involves upgrading the software on one ProtecTIER server at a time, which keeps one server online throughout the process. The nonconcurrent procedure involves varying both ProtecTIER servers offline at the same time. It is recommended that you select the concurrent process, unless otherwise instructed, because the nonconcurrent process causes the library to fail.

In order to upgrade the ProtecTIER software, perform the following two tasks in the order shown.

ProtecTIER software upgrade procedure

Complete this task to upgrade the ProtecTIER software.

Before upgrading the ProtecTIER software, ensure that you have completed the following requirements:

- ▶ Inform the 3958-DE2 (TS7680) system users that they will lose access to some or all of the devices for a period of time. The clients lose access to half of the devices for a concurrent upgrade and to all devices for a nonconcurrent upgrade. The expected outage time must be planned for before starting the upgrade.
- ▶ Check for outstanding alerts in the alerts log. If any error conditions are found, contact your service support representative (SSR) before continuing with this procedure.

Complete the following steps to concurrently upgrade the ProtecTIER software. Each ProtecTIER server must be upgraded separately.

Important: For a nonconcurrent upgrade, refer to the note at the end of this procedure before beginning.

1. Quiesce tape activity at the host by performing the following substeps:
 - a. Direct scratch allocations to another library (as appropriate). This will happen automatically if multiple libraries are eligible for the scratch allocations when the devices in the 3958-DE2 (TS7680) are varied offline in the following substep.
 - b. Vary the appropriate devices offline at the hosts. Refer to 8.3.1, "Varying devices offline" on page 184 and then return here and proceed with the following step.
2. Once the devices are offline, notify the Enterprise controller that the ProtecTIER server is being varied offline from the ProtecTIER service menu. See 8.3.2, "Notifying the Enterprise controller of change in ProtecTIER server status" on page 185.
3. Proceed with upgrading the first node and then return here.
4. Vary the offline devices back online at the hosts.
5. Notify the Enterprise controller that the ProtecTIER server is back online.

6. Repeat Steps 1 and 2 for the second ProtecTIER server.
7. Proceed with upgrading the second node and then return here.
8. Vary the offline devices back online at the hosts.
9. Notify the Enterprise controller that the ProtecTIER server is back online.

For a nonconcurrent upgrade, follow the above procedure with the following exceptions:

- ▶ Vary all devices offline at Step 1.
- ▶ At Step 2, go to the rasMenu for each ProtecTIER server and notify the attached Enterprise controller that the ProtecTIER server is being varied offline.
- ▶ Skip Steps 4 through 6.
- ▶ At Step 9, go to the rasMenu for each ProtecTIER server and notify the attached Enterprise controller that the ProtecTIER server is being varied online.

Upgrading ProtecTIER on a two-node cluster

The following procedure describes how to upgrade ProtecTIER on a two-node cluster. In a two-node cluster, each node must be upgraded separately.

A repository must be configured on a node before upgrading ProtecTIER.

Note:

- ▶ Do not start the procedure to upgrade ProtecTIER software here. Start this procedure at the ProtecTIER software upgrade procedure.
- ▶ Some steps in this procedure vary depending on whether you are performing a concurrent or nonconcurrent software upgrade. Ensure that you follow the appropriate steps.

Upgrading the first node

To upgrade the first node:

1. Log in as root and stop the vtfd using the following steps:
 - a. For a concurrent upgrade, activate the keyboard, video, mouse (KVM), and verify a connection to Server A (the first server). For a nonconcurrent upgrade, activate the KVM and verify a connection to Server B (the second server).
 - b. When the localhost login: prompt appears, log in as root and type the password admin.
 - c. Monitor the vtfd service on Node A (for a concurrent upgrade) or Node B (for a nonconcurrent upgrade) by typing the following command:


```
service vtfd status      <Enter>
```
 - d. If the vtfd does not report as stopped, enter the following command to stop the process:


```
service vtfd stop      <Enter>
```

The process will discontinue and return you to a command prompt.
 - e. Activate the KVM and verify a connection to Server A (the first server).
 - f. When the localhost login: prompt appears, log in as root and type the password admin.
2. Mount and copy the code to the installation directory, using the following steps:
 - a. Insert the IBM System Storage ProtecTIER Enterprise Edition CD and wait for the CD to stop blinking.
 - b. Create the /mnt/cdrom directory by typing the following command:


```
mkdir /mnt/cdrom <Enter>
```

Note: If the /mnt/cdrom directory already exists, the following message will be displayed:

```
mkdir: cannot create directory '/mnt/cdrom': File exists
```

If this message appears, ignore it and proceed to the next step.

- c. Mount the CD-ROM drive:

```
mount /dev/cdrom /mnt/cdrom <Enter>
```

The following output is displayed:

```
mount: block device /dev/cdrom is write-protected. mounting read-only
```

- d. Type the following command to change the current directory to the local installation directory:

```
cd /mnt/cdrom
```

- e. From the CD, locate the following tar file, where <filename> indicates the version number and date:

- List the files on the CD from the command line: `ls` <Enter>
- Locate the .tar file and copy it to the /install directory on the hard drive. Type the following command to copy the file:

```
cp <filename>.tar /install <Enter>
```

- f. Access the directory to which you copied the tar file (for example, type `cd /install`).

- g. Type the following command to extract the installation files in the /install directory:

```
tar -xvf <tar filename>.tar <Enter>
```

The <filename> directory is created.

- h. Type the following command to change to the < tar filename> directory. Perform the rest of the installation from this directory:

```
cd /install/<filename> <Enter>
```

Example: <filename>=PT_TS7680_V1.1.0.0x86_64

- i. Type the following command to unmount the CD-ROM drive:

```
umount /dev/cdrom <Enter>
```

- j. Type the following command to eject the ProtecTIER installation CD from the CD-ROM drive:

```
eject /dev/cdrom
```

- k. Type the following command:

```
./autorun <Enter>
```

- l. Enter y if you are prompted to stop the vtfd service.

- m. The autorun utility will install the ProtecTIER application. The following message is displayed at the end of the installation:

```
The system will now reboot!\
```

Note: After the reboot, the vtfd will be ready only after the service vtfd status command reports that vtfd is running, not that the service is not finished running yet.

Press <Enter> at the prompt to reboot the system. The following message is displayed:

```
After boot, please set user to ptadmin by invoking 'su - ptadmin'
(default password is ptadmin).
Press <CR> to continue...
```

Upgrading the second node

To upgrade the second node:

1. Log in as root and stop the vtfd using the following steps:

- a. For a concurrent upgrade, activate the keyboard, video, mouse (KVM), and verify a connection to Server B (the second server). For a nonconcurrent upgrade, activate the KVM and verify a connection to Server A (the first server).
- b. When the localhost login: prompt appears, log in as root and type the password admin.
- c. From the command line, enter the following command to stop vtfd:

```
service vtfd stop      <Enter>
```
- d. Activate the KVM and verify a connection to Server B (the second server).
- e. When the localhost login: prompt appears, log in as root and type the password admin.

2. Mount and copy the code to the installation directory using the following steps:

- a. Insert the IBM System Storage ProtecTIER Enterprise Edition V2.4 (or later) CD and wait for the CD to stop blinking.
- b. Create the /mnt/cdrom directory by typing the following command:

```
mkdir /mnt/cdrom      <Enter>
```

Note: If the /mnt/cdrom directory already exists, the following message will be displayed:

```
mkdir: cannot create directory '/mnt/cdrom': File exists
```

If this message appears, ignore it and proceed to the next step.

c. Mount the CD-ROM drive:

```
mount /dev/cdrom /mnt/cdrom      <Enter>
```

The following output is displayed:

```
mount: block device /dev/cdrom is write-protected. mounting read-only
```

d. Type the following command to change the current directory to the local installation directory:

```
cd /mnt/cdrom
```

e. From the CD, locate the following tar file, where <filename> indicates the version number and date:

- List the files on the CD from the command line: **ls** <Enter>
- Locate the tar file and copy it to the /install directory on the hard drive. Type the following command to copy the file:

```
cp <filename>.tar /install      <Enter>
```

f. Access the directory to which you copied the tar file (for example, type **cd /install**).

g. Type the following command to extract the installation files in the /install directory:

```
tar -xvf <tar filename>.tar      <Enter>
```

The <filename> directory is created.

- h. Type the following command to change to the < tar filename> directory. Perform the rest of the installation from this directory:

```
cd /install/<filename>      <Enter>
```

Example: <filename>=PT_TS7680_V1.1.0.0x86_64

- i. Type the following command to unmount the CD-ROM drive:

```
umount /dev/cdrom      <Enter>
```

- j. Type the following command to eject the ProtecTIER installation CD from the CD-ROM drive:

```
eject /dev/cdrom
```

Run the autorun command only after the service on Server A has completed its startup (after the reboot).

- k. Type the following command:

```
./autorun      <Enter>
```

- l. Enter y if you are prompted to stop the vtfd service.

- m. The autorun utility will install the ProtecTIER application. The following message is displayed at the end of the installation:

```
The system will now reboot!\
```

Press <Enter> at the prompt to reboot the system. The following message is displayed:

```
After boot, please set user to ptadmin by invoking 'su - ptadmin'  
(default password is ptadmin).  
Press <CR> to continue... <Enter>
```

8.4.4 Setting up the ProtecTIER system

The topics in this section define the tasks used to set up the ProtecTIER system.

Prerequisites for completing the ProtecTIER system setup

Use the ptconfig utility to complete the ProtecTIER configuration process that was started in manufacturing.

Before executing ptconfig, verify that the following tasks have been completed:

- ▶ Cabling and physical connections (including those to the client's local area network (LAN) using assigned IP addresses) are complete.
- ▶ You have collected the following server information:
 - External IP address
 - External, fully-qualified, host name (hostname.domain.com)
 - External gateway address
 - External network subnet mask

Note: The cluster name and host name are case-sensitive. Be aware of the use of upper- and lowercase characters when gathering the information and when entering it into the system.

Logging in to the ProtecTIER server

Follow these steps to log in to the ProtecTIER server:

1. Verify that Servers A and B are running:
 - If they are, go on to Step 2.
 - If they are not running: power on any servers that are not running, wait for the boot cycle to complete, and then go on to Step 2.
2. At the localhost login: prompt, log in with the ID root and the password admin.
3. Go on to “Configuring the ProtecTIER server”.

Configuring the ProtecTIER server

If you have a code-level update package, execute the ptconfig that is resident on the server first, then apply the update from the CD.

Perform the following steps on Server A and Server B.

To configure server A:

1. At the server's command prompt, change to the /opt/dtc/install directory. To do so, enter the following command: `cd /opt/dtc/install <Enter>`
2. Change the server's IP address, netmask, default gateway, and host name from the values that were set in manufacturing, to the values specific to the environment. To do so:
 - a. Enter the following command:

```
./ptconfig -updateNetwork <Enter>
```

The following status messages display:

```
Starting Cluster, please wait
Starting cluster [Done]
Cluster Started
```

- b. If asked if you would like to stop the vtfd service, type yes <Enter>.

The following status message displays as the system initiates shutdown:

```
Stopping VTFD [/]
```

The shutdown process may take a few minutes to complete.

When shutdown completes, the following status message displays:

```
Stopping VTFD [Done]
```

- c. You are then prompted, one at a time, to enter the following values. At each prompt, type the new value and then press <Enter>.
 - Customer Network IP address
 - Customer Network netmask
 - Customer Network default gateway
 - Customer Network host name (this is the server's fully-qualified host name. For example: hostname.domain.com)

Note: The values that were assigned to the server during manufacturing appear in brackets after each prompt. For example: Customer Network netmask:

```
[255.255.255.0]
```

After you enter the hostname, the system automatically starts the network configuration process. The following status messages display:

```
Configuring network [Done]
Updated network configuration successfully
update updateNetwork ended successfully
```

The system automatically restarts the vtfd service, and you are returned to the command prompt.

3. Change the system name from the one assigned during manufacturing to one specific to the your environment. This command only needs to be run from one of the nodes because it affects the shared name of the system. To do so:

- a. Enter the following command:

```
./ptconfig -updateSystemName <Enter>
```

The following status messages display:

```
Starting Cluster, please wait
Starting Cluster [Done]
Cluster Started
```

- b. When prompted, type the new system name of the server and press Enter.

Note: The system name that was assigned to the server in manufacturing appears in brackets after the prompt. For example: [PORTLAND]. After you enter the system name, the system automatically starts the Update System Name process. The following status messages display:

```
Change system name [Done]
Updated system name successfully
update updateSystemName ended successfully
```

You are returned to the command prompt.

4. Repeat the steps in this section on Server B, then go on to 8.5.1, “Installing ProtecTIER Manager” on page 199.

Enabling Simple Network Management Protocol compatibility

This section details enabling Simple Network Management Protocol (SNMP) support and compatibility, and the IBM Management Information Base (MIB) definition file.

ProtecTIER software responds to SNMP, implementing MIB-2 and generating the appropriate traps.

The ProtecTIER server responds to SNMP discovery and queries, and to the standard MIB requests.

Enabling ProtecTIER Simple Network Management Protocol support

When the server is installed at the user site, the IP address of the SNMP management station must be made available to the ProtecTIER code.

To enable SNMP support:

1. Edit the configuration file. From the command line, use the vi editor:

```
vi /etc/snmp/snmpd.conf <Enter>
```

2. Scroll down the list to the bottom of the file and locate this line:

```
trapsink localhost
```

3. Use the arrow key on the keyboard to place the cursor under the letter l in the word localhost. Press the Delete key until localhost is removed. Press the a key (this is for add or insert mode). Enter the IP address with the TS3000 System Console (TSSC) SNMP management station IP address.

Example: Remove localhost and insert 172.31.1.1.

172.31.1.1 is the standard TSSC IP address. After the IP address has been entered, press the Esc key, then Shift+colon (:) and type wq! (write-quit) and press Enter. This will save the file.

IBM Management Information Base definition file

IBM supplies an MIB definition file called DILIGENT-MIB.txt that can be found in the /usr/share/snmp/mibs directory. This file describes each of the supported traps and should be imported to the Simple Network Management Protocol (SNMP) application used at the client site.

Simple Network Management Protocol compatibility

This topic defines the ProtecTIER Simple Network Management Protocol (SNMP) implementation.

ProtecTIER implements SNMP as follows:

- ▶ MIB-2 implementation

In the MIB-2 System Group, the following fields are implemented: sysDescr, sysObjectID, sysUpTime, sysContact, sysName, sysLocation, and sysServices.

All other parts of the MIB-2 respond in such a way that management tools understand that they are not implemented.

- ▶ Traps

The traps generated are: coldStart, warmStart, authenticationFailure, operatorInterventionRequired (proprietary), and recoverableErrorNotification (proprietary).

If authentication of a user fails more than five times in a row, an SNMP authenticationFailure trap is generated.

- ▶ Startup consistency checks

ProtecTIER checks its persistent data on startup, using a Power On Self Test (POST) procedure, in order to verify that the application can run. Initialization files are checked for consistency. Resource allocation—specifically memory—is checked.

Errors encountered may be recoverable or unrecoverable.

- Recoverable errors

A recoverable error is an error from which the ProtecTIER server can recover without losing the user's data. Recoverable errors are logged in the ProtecTIER logs and generate an SNMP warning notification trap.

- Unrecoverable errors

An unrecoverable error is an error that prevents the platform from booting correctly, an unrecoverable consistency check error during ProtecTIER startup, or any other error that could cause or has caused loss of user data. The server is left in offline state (booted, responding to TCP/IP SNMP inquiries, and responding to console, Telnet and modem logins). Unrecoverable errors are logged in the ProtecTIER logs and generate an SNMP error trap.

- ▶ Restarting on error

If ProtecTIER detects an error at runtime, it recovers by rebooting and restarting the ProtecTIER process.

If multiple restarts are detected within a short time period, ProtecTIER declares an unrecoverable error.

► Alerts

The server generates SNMP traps to higher-level management frameworks.

In particular, whenever the system enters the online state, it generates a coldStart trap if the platform has rebooted, and a warmStart trap if the system only returned to online state from the offline state.

A recoverable error generates a recoverableErrorNotification trap. An unrecoverable error generates an operatorInterventionRequired trap.

8.5 Managing the ProtecTIER Manager

This section discusses how to manage the ProtecTIER Manager.

8.5.1 Installing ProtecTIER Manager

Complete this task to install the ProtecTIER Manager.

The ProtecTIER Manager installer is provided on your IBM System Storage ProtecTIER Manager CD.

Different ProtecTIER Manager installers are used for Windows and Linux. Ensure that the installer that you are using is correct for the operating system running on your workstation.

- If you are installing ProtecTIER Manager on a workstation running Windows, see “Installing on Windows” on page 200.
- If you are installing ProtecTIER Manager on a workstation running Linux, see “Installing on Linux” on page 201.

If you are installing ProtecTIER Manager on a workstation on which an older version of ProtecTIER Manager is already installed, uninstall the older version first. For more information, see 8.5.2, “Uninstalling the ProtecTIER Manager” on page 202.

Prerequisites for the ProtecTIER Manager workstation

This topic describes the workstation prerequisites for installing and running the ProtecTIER Manager.

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

- One of the following operating systems:
 - Windows 32 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.

Installing on Windows

Complete this task to install ProtecTIER Manager on a Windows system.

1. Insert the IBM System Storage ProtecTIER Manager CD into the CD-ROM drive of the designated ProtecTIER Manager workstation.
 - If the ProtecTIER Manager autorun launches and starts the installation, go on to Step 2.
 - If the ProtecTIER Manager autorun process does not launch automatically, do the following:
 - a. On the Windows task bar, click **Start** → **Run**. The Run dialog box opens.
 - b. In the Open box, type D: (where D: is the server's CD-ROM drive).
 - c. Click **OK**.
The content of the IBM System Storage ProtecTIER Manager CD displays.
 - d. From the list of files, locate the ProtecTIER Manager for Windows installation file and double-click the file to start the installation.
2. Read the Introduction panel, and then click **Next**.
3. Read and accept the license agreement provided, and click **Next**.
The Choose Install Folder window is displayed.
4. Specify the folder where the ProtecTIER Manager program files will be installed, and click **Next**.
The Choose Shortcut Folder panel is displayed.
5. Select the location where the program icons will be created:
 - In a new Program Group - Creates a new program group in the Program list of the Start menu.
 - In an existing Program Group - Adds the shortcut to an existing program group in the Program list of the Start menu.
 - In the Start Menu
 - On the desktop
 - In the Quick Launch Bar
 - Other - Enables you to enter a path location for the shortcut, or to browse for a location by clicking **Choose**.
 - Do not create icons - No shortcuts are created.

Note: When relevant, you can select **Create icons for All Users** to create a shortcut in the defined location for all user accounts on the workstation.

6. Click **Next**.

The Pre-Installation Summary window displays the Install and Shortcut folder locations and the disk space information of the target for installation.

7. Review the Summary panel, and click Install to start the installation.

The Installing ProtecTIER Manager window is displayed.

8. When the installation is complete and ProtecTIER Manager has been successfully installed, click **Done**.

The Install Complete window is displayed. The ProtecTIER Installation wizard closes.

Installing on Linux

Perform the following steps to install ProtecTIER Manager on Linux:

1. Insert the IBM System Storage ProtecTIER Manager CD into the CD-ROM drive of the designated ProtecTIER Manager workstation.
2. Run the ProtecTIER Manager installer.

Note: The following presumes that the workstation has a Linux graphical interface, which is required for ProtecTIER Manager.

- a. From the Linux desktop, select and open the CD drive icon.
- b. Select and open the folder for your Linux version—Linux for version 64 or Linux32 for version 32.
- c. When the folder opens, drag the InstallLinuxXX.bin file from the folder to the desktop. (XX = either 32 or 64, depending on the Linux folder that you selected.)
- d. Close the open windows.
- e. Right-click an open area of the desktop, and from the menu options displayed, select **Open Terminal**.
- f. At the Terminal command prompt, change to the desktop directory using the following command (Note: Desktop is case sensitive. Type it using a capital D.):

```
cd Desktop      <Enter>
```

- g. From the desktop directory in the Terminal window, run the ProtecTIER Manager installer:

```
./InstallLinuxXX.bin      <Enter>
```

(XX = either 32 or 64, as noted above)

If the message Permission Denied displays, enter the following commands:

```
chmod +x InstallLinuxXX.bin      <Enter>
```

```
./InstallLinuxXX.bin      <Enter>
```

The IBM ProtecTIER Manager wizard Introduction window is displayed.

3. Click **Next**. Two separate Software License Agreement windows display.
4. Read the terms for each license agreement, select **I accept both the IBM and non-IBM terms of the License Agreement** and click **Next**. The Choose Install Folder window is displayed.
5. Enter the path to the location where the ProtecTIER Manager program files will be installed. Click **Choose** to browse for a location.

Note: Click **Restore Default Folder** to revert to the default installation path.

6. Click **Next**. The Choose Link Folder window is displayed.
7. Select the location where the program links will be created:
 - In your Home folder – Creates the links in the directory where your files are typically stored. For example, /home/bill.
 - Other – Creates the links in the default location (/opt/IBM/PTManager). To specify a different location, click **Choose** and select a directory on the workstation's hard drive.
 - Do not create links – No links are created.
8. Click **Next**. The Pre-Installation Summary window is displayed.
9. Click **Install**. The Installing ProtecTIER Manager window is displayed and ProtecTIER Manager is installed on your computer. When the installation is complete, the Install Complete panel is displayed.
10. Click **Done**. The ProtecTIER Manager wizard closes. When the command prompt returns in the Terminal window, type `exit` to close the window.

8.5.2 Uninstalling the ProtecTIER Manager

Perform the following steps to uninstall the ProtecTIER Manager:

1. From the ProtecTIER Manager directory, run the ProtecTIER Manager uninstaller. The Uninstall ProtecTIER Manager wizard Introduction window is displayed.
2. Click **Uninstall**. The Uninstalling window is displayed and ProtecTIER Manager is removed from your computer. The Uninstall Complete panel is displayed.
3. Click **Done**. The Uninstall ProtecTIER Manager wizard closes.

8.5.3 Managing users of the ProtecTIER Manager system

The topics in this section define how to manage users of the ProtecTIER Manager system.

The ProtecTIER Manager system allows a system administrator to manage the users of the system in the following areas:

- ▶ Levels of permission for each user
- ▶ Adding or deleting new users
- ▶ Assigning user passwords

Managing ProtecTIER Manager user permissions

The ProtecTIER system supports the following permission levels and associated activities:

- ▶ Administrator

Administrator has full access to the ProtecTIER Manager system. Only a person with administrator permission can add and delete users.

- ▶ Operator

Operator permission allows a user to access ProtecTIER Manager monitoring panels and perform the following tasks:

- Toggle cartridges between read/write and read-only modes.

- Set the HyperFactor mode for libraries.
 - Reset virtual tape drives and robots.
 - Unload and move cartridges from virtual tape drives.
- Monitor
- Monitor permission allows a user to only access the ProtecTIER Manager monitoring panels.

Adding a user to the ProtecTIER Manager system

Complete this task to add a user to the ProtecTIER Manager system.

Only a person assigned the administrator permission level can add a user to the ProtecTIER Manager system.

Perform the following steps to assign a new user account:

1. Log into the ProtecTIER Manager system as Administrator. Refer to 8.1.3, “Logging into and out of the ProtecTIER Manager” on page 172 for “Default username and password” and select **System** → **Users management** (Figure 8-16).




Figure 8-16 Select Users management

The Users management window is displayed (see Figure 8-17 on page 204).

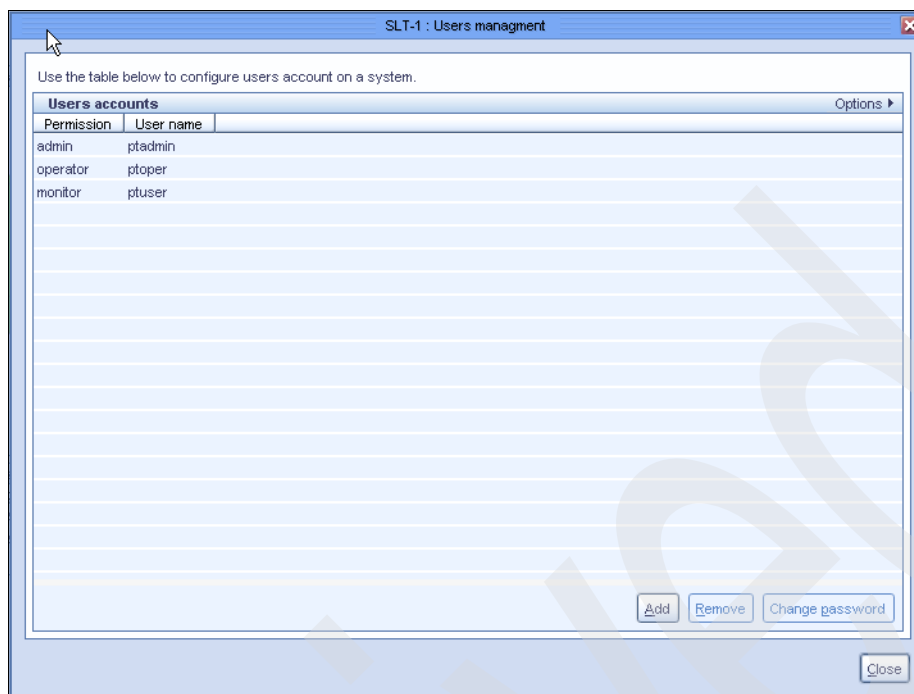


Figure 8-17 Users management

2. Click **Add**. The Add account dialog is displayed.
3. Complete the input fields on the dialog window.

The fields consist of the following information:

User name	Input a user name that complies with your company policy for user names.
New password	Input a password that complies with your company policy for passwords.
Verify password	Enter the same password again to verify the password.
Permission	Input the permission level (administrator, operator, or monitor) you want to assign the user.

Note: Before processing, record this information in a manner that allows you to notify the new user. Remind the new user to record the password because neither you nor the system stores the passwords for later use

4. Click **OK**. The new account is added to the ProtecTIER Manager system.
You can delete a user by selecting the user account from the User Account list and clicking **Remove**.

Changing user account passwords in the ProtecTIER Manager system

Ensure that you have your original password because you must use it to change to a new password. If the original password is not available, the administrator must delete the user account and create a new account.

Perform the following steps to change the user account password:

1. Log into the ProtecTIER Manager system as Administrator and select **System** → **Users management**. The Users management dialog is displayed. Refer to Step 1 on page 203.

2. Select a user account and click **Change password**. The Change password dialog is displayed (see Figure 8-18).

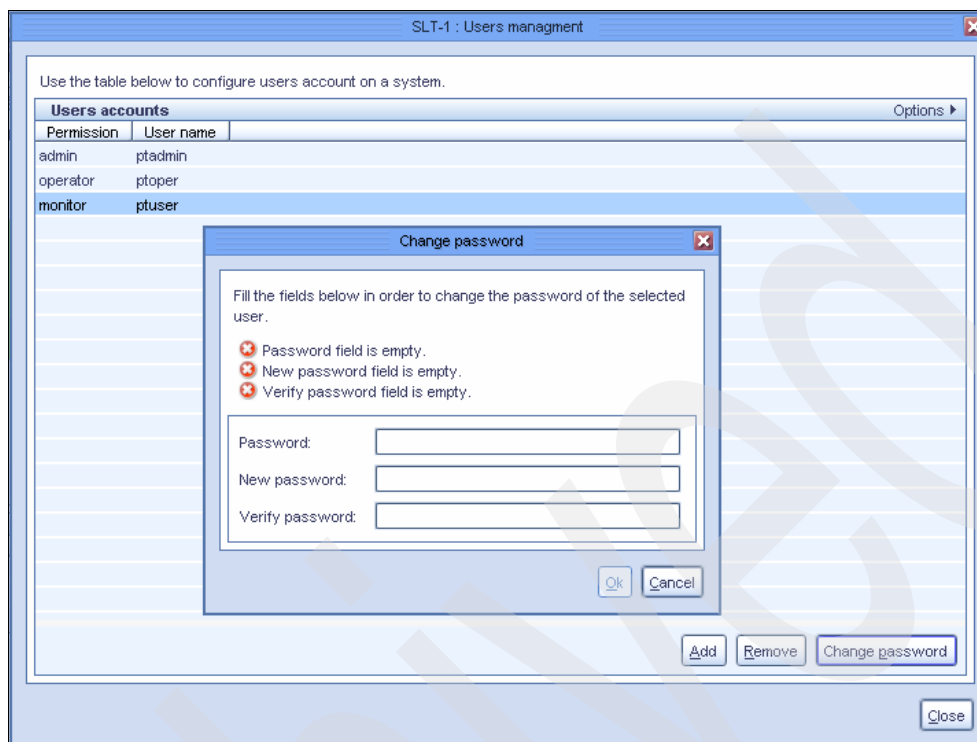


Figure 8-18 Change user password

3. Type the user account's original password in the password field.
4. Type the new password in the New password field.
5. Type the new password in the Verify password field.
6. Click **OK** and the new password is added to the ProtecTIER Manager system. Make a note of the new password and inform the user of what it is. Ensure that it is recorded for later use.

8.6 Managing the ProtecTIER Virtual Tape system

The topics in this section describe managing the ProtecTIER Virtual Tape Library (VTL) system.

8.6.1 Managing cartridges

The topics in this section define how to add and remove cartridges from your libraries in the ProtecTIER system.

Adding cartridges

The 3958-DE2 (TS7680) system is configured with 1,000,000 slots available. It is possible that this operation could cause input/output (I/O) activity to fail.

Perform the following steps to add cartridges:

1. Select **Library xxx** on the left pane of the IBM ProtecTIER manager window (see Figure 8-19).

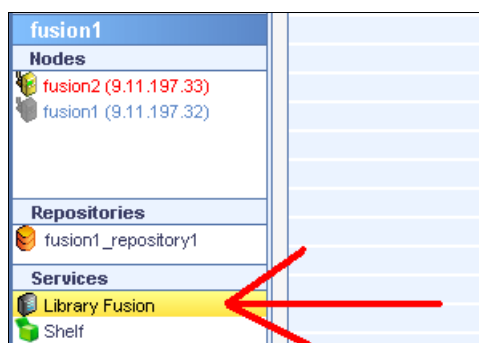


Figure 8-19 Select library

2. Click **Add cartridges**.
3. If Confirm operation panel appears, click Yes (see Figure 8-20).




Figure 8-20 Confirm operation "might cause I/O activity to fail"

4. The Add cartridges wizard Welcome panel is displayed (Figure 8-21 on page 207).

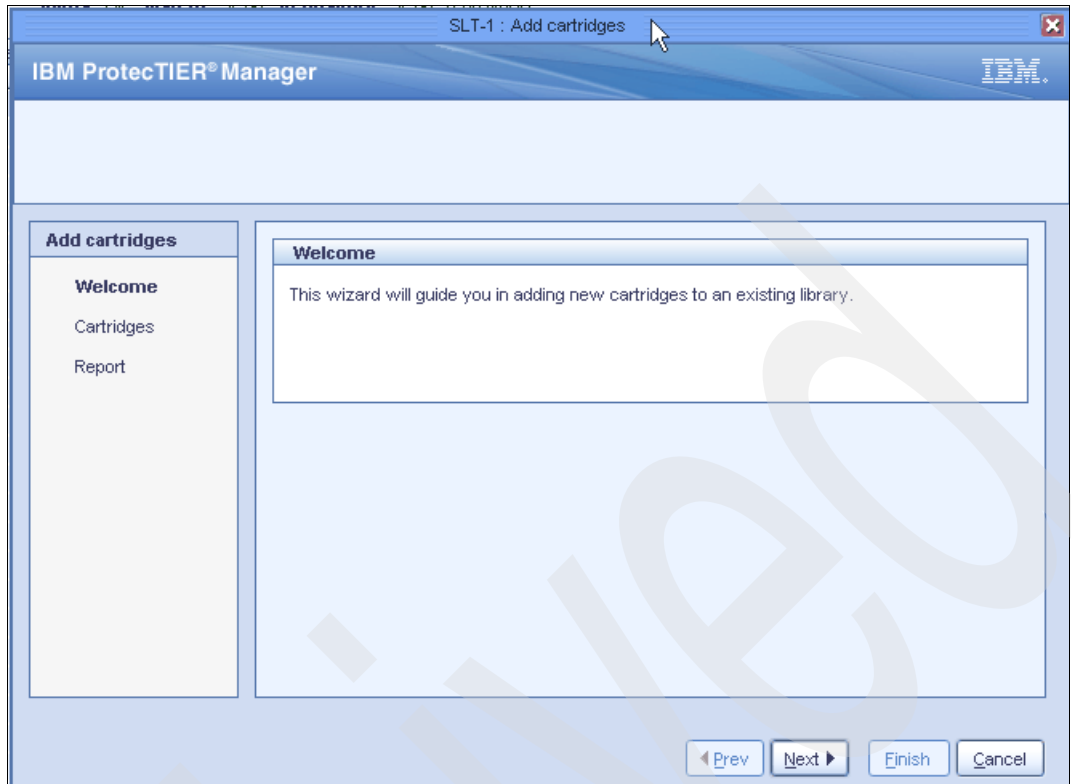


Figure 8-21 Add cartridges

- Click **Next**. The Cartridges panel is displayed (Figure 8-22).

Figure 8-22 Cartridges panel

- In the No. of cartridges field, enter the number of cartridges that you want to have in the library. The cartridge size is set at a value of 100 GB.
- In the Barcode seed field, enter a value for the barcode seed. The default barcode seed is the continuation of the initial barcode seed assigned when the library was created.

Note: The barcode seed must contain only numbers and capital letters.

- Click **Next** and **Finish**. The Add cartridges wizard closes and the cartridges are added to the library (see Figure 8-23 on page 208).

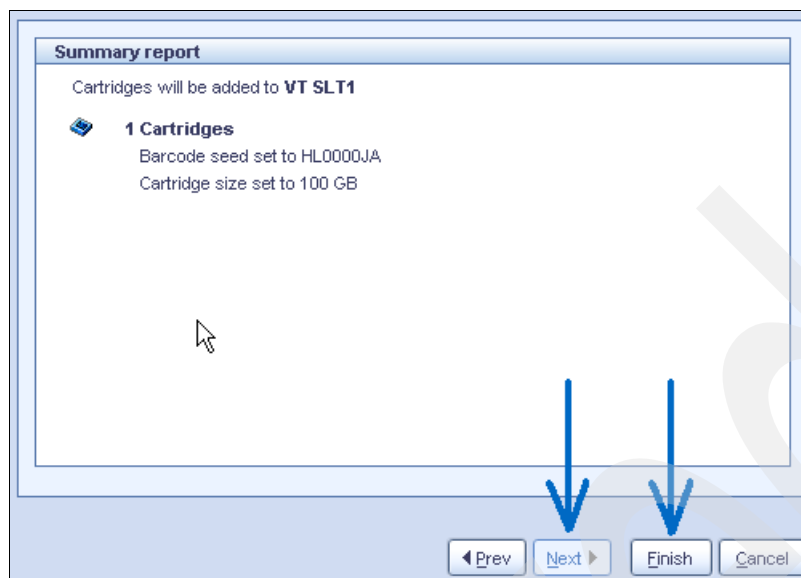


Figure 8-23 Add cartridges wizard

Recovering from database inconsistencies

We recommend that you only delete a virtual cartridge by initiating an eject request from the System z host. However, if you were instructed to delete a volume using the ProtecTIER Manager, you may need to complete this task in order to recover from database inconsistencies.

Complete the following steps to recover from database inconsistencies between the System z host software and the 3958-DE2 (TS7680).

1. From the System z host, enter the DISPLAY SMS,VOLUME command to determine whether the volume exists in the Tape Configuration Database (TCDB). If it does, continue with the following step (see Figure 8-24).

```
08.41.36      d sms,vol(fus002)
08.41.36 STC00015 CER1180I QAM tape volume status: 965
VOLUME MEDIA STORAGE LIBRARY USE W C SOFTWARE LIBRARY
TYPE GROUP NAME AIR P P ERR STAT CATEGORY
FUS002 MEDIA5 SGAIL ATILCA030 P N NOERROR NOTAVAIL

RECORDING TECH: EFMT1 COMPACTION: NO
SPECIAL ATTRIBUTE: NONE ENTER/EJECT DATE: 2009-07-31
CREATION DATE: 2009-05-13 EXPIRATION DATE:
LAST MOUNTED DATE: 2009-05-27 LAST WRITTEN DATE: 2009-05-27
SHELF LOCATION:
OWNER:
```

Figure 8-24 Display SMS, Volume

2. Use the IDCAMS utility to delete or alter the volume record in the TCDB. Alter can be used to retain information about the volume. If the intent is to reenter the same volumes into the library, no changes to the TCDB are required
3. Update the tape management system accordingly to remove or update a volume record in the inventory.

8.6.2 Monitoring system functions using ProtecTIER

The topics in this section define how to use ProtecTIER to monitor your clusters, repositories, nodes, and ProtecTIER Virtual Tape library (VTL) services.

Monitoring clusters with ProtecTIER

The topics in this section define how to use ProtecTIER to monitor clusters. The monitoring of clusters falls into two categories: General and VTL (Virtual Tape library).

Monitoring general clusters

Open the IBM ProtecTIER Manager, click **System** and select the cluster that you want to view. At the bottom of the System window, click the **General** tab. The General Clustering window is displayed with the following selections:

- ▶ Capacity
- ▶ Fibre channel ports throughput
- ▶ Cluster members

Refer to Chapter 9, “Monitoring the system” on page 231 for detailed procedures.

Monitoring Virtual Tape clusters

This topic defines the *drives* information provided by the ProtecTIER Manager when it is used for monitoring Virtual Tape Library (VTL) services.

Open the IBM ProtecTIER Manager, click **System** and select the cluster that you want to view. At the bottom of the System window click the **VT** tab. The VT tab allows information about the VTL service libraries associated with the cluster to be displayed. The VT clustering window is displayed with the following selections:

- ▶ Library front-end
- ▶ Libraries performance
- ▶ Libraries configuration

Refer to Chapter 9, “Monitoring the system” on page 231 for detailed procedures.

Monitoring the ProtecTIER Virtual Tape service

The type of monitoring information that is provided for VTL services consists of the following types of information:

- ▶ General
- ▶ Drives
- ▶ Cartridges
- ▶ Slots
- ▶ Import/Export

8.6.3 Setting Control Path Failover

This topic details the Control Path Failover (CPF), which is supported by the emulated IBM library models.

By default, CPF is enabled on the emulated IBM library models. When troubleshooting host connectivity issues, you might be asked to disable CPF.

To disable CPF, click a library from the Services window in the navigation pane.

1. Select Library xxx on the left pane of the IBM ProtecTIER manager window (see Figure 8-25).

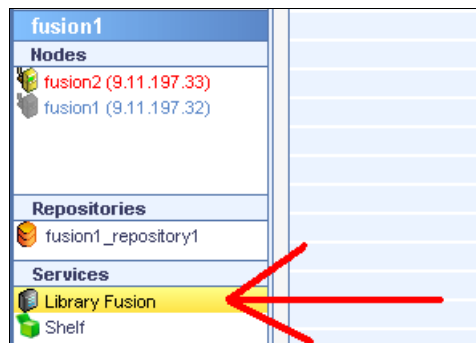


Figure 8-25 Select library

2. Select **VT** → **VT Library** → **Set control path failover mode**. The Set control path failover mode dialog is displayed (see Figure 8-26).

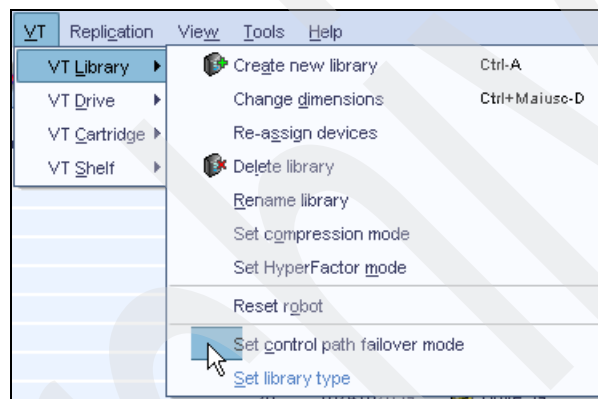


Figure 8-26 Select Control Path Failover

3. Set the CPF mode to Control path failover disabled (see Figure 8-27).




Figure 8-27 Set Control path failover disabled

8.7 Managing virtual libraries

The topics in this section describe managing virtual libraries.

8.7.1 Creating libraries

Complete this task to create a library for your ProtecTIER system.

The 3958-DE2 (TS7680) is a two-node system. You can create only one library with a fixed configuration that contains the following values:

- ▶ Two cluster members
- ▶ Four FE ports on each node
- ▶ 128 tape drives assigned to each node (divided equally between the ports)
- ▶ 1,000,000 slots
- ▶ No Import and Export slots

No cartridges are created when the library is created. However, you can add up to one million cartridges at a later stage by performing the add cartridges task. Refer to 8.6.1, “Managing cartridges” on page 205.

Perform the following steps to create the library:

1. Log into the ProtecTIER Manager.
2. In the Systems pane, select a cluster on which to add a library.
3. From the Toolbar, click **Create new library**. The Create new library wizard Welcome panel is displayed.
4. Click **Next**. The Library details panel is displayed.
5. In the VT name field, enter a name for the library.
6. Click **Next** and **Finish**. The Create new library wizard closes and the ProtecTIER system temporarily goes offline to create the library. The library is displayed in the Services pane and the VT monitoring window is displayed.

Note: To be done only at setup time.

Refer to Chapter 6, “TS7680 setup” on page 91 for details.

8.7.2 Renaming libraries

Perform the following steps to rename an existing library:

1. Log in to the ProtecTIER Manager.
2. In the Services pane, select a library.
3. Choose **VT → VT Library → Rename library** (see Figure 8-28 on page 212).




Figure 8-28 Renaming libraries

4. Enter a new name for the selected library and click **Ok**. The Rename library dialog closes and the library's name is changed.

8.7.3 Deleting libraries

Complete this task to delete a ProtecTIER library (see Figure 8-29).

Attention: Deleting a library results in the loss of all data contained in that library. Only a person with administrator authority can perform this task.

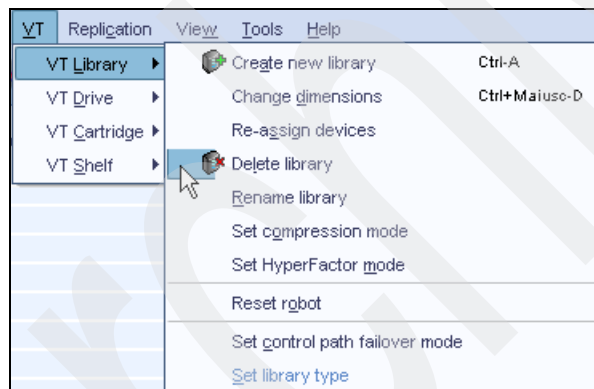


Figure 8-29 Deleting libraries

8.8 ProtecTIER system

The topics in this section provide tips, guidelines, and tasks for handling error situations that might arise within the ProtecTIER system.

8.8.1 Viewing the ProtecTIER system logs

The ProtecTIER system event logs consist of the following reports:

Alerts log Provides a list of current error events (200 viewable at a time) that have occurred in the ProtecTIER system. The detail is provided on a node by node basis.

Events log Provides a list of the current events (200 viewable at a time) that have occurred in the ProtecTIER system. The detail is provided on a node by node basis.

Perform the following steps to view the event and alert logs:

Note: The following steps presume you are already logged into the ProtecTIER Manager.

Select **View** on the tool bar. The View window is displayed.

1. Click **Alerts** on the bottom right corner of the View window (seeFigure 8-30).

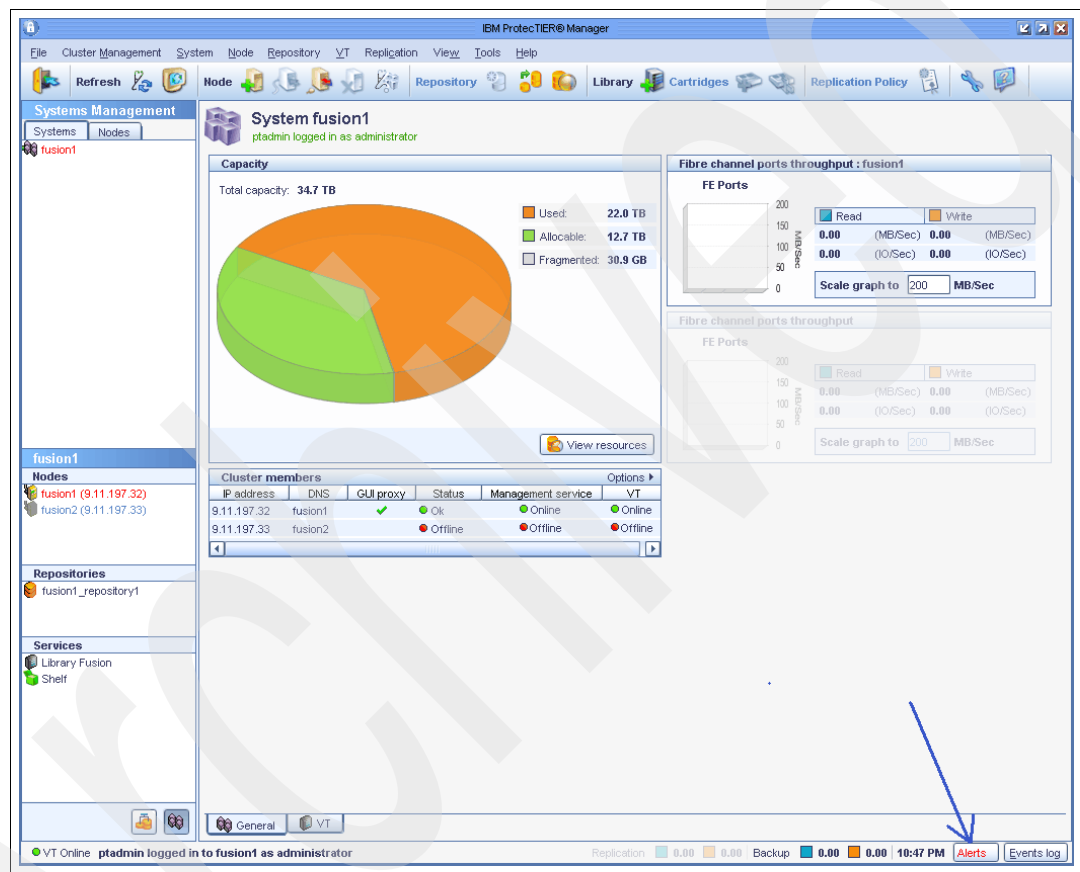


Figure 8-30 ProtecTIER Manager alert

Note:

- ▶ When a new error has been reported, the Alerts button turns red and features a blinking exclamation mark.
- ▶ You can view the Alerts log directly on the ProtecTIER server by opening the /pt_work/log/vtf_error.log file.
- ▶ You can clear selected Alerts by clicking **Clear Alerts** on the lower right corner of the Alerts Log window.

2. Select **Events Log** on the bottom right corner of the View window. The Events Log window is displayed.

Note: Refer to Chapter 6, “TS7680 setup” on page 91 for details.

8.9 Wizard error messages

ProtectTIER Manager wizards feature a Message area to inform you of issues that relate to each wizard panel.

There are two types of messages displayed in the ProtectTIER Manager wizard:

- ▶ **Critical Message** - Indicated by a red circle with a white X. A Critical Message must be resolved to continue with the wizard.
- ▶ **Warning Message** - Indicated by a yellow triangle with a white exclamation mark. The wizard can continue without resolving Warning Messages, but it is not recommended.

8.9.1 Generating a ProtecTIER system status report

IBM Service might request that you generate a ProtecTIER system status report. The generated report files can be attached to a support ticket and sent to IBM Service. Status reports can only be generated for one node at a time.

Perform the following steps to generate a status report:

Notes:

- ▶ It is presumed that you are already logged into the ProtecTIER Manager.
- ▶ You can generate a service report directly on a ProtecTIER server by entering the `opt/dtc/app/ sbin/report_problem` command.
- ▶ You can perform a system check on the server by typing `sosreport`. However, this should only be done if you are instructed to do so by IBM Service.

1. Click the folder **Nodes** on System Management (on the left side) to select a node where you will generate a status report (see Figure 8-31 on page 215).

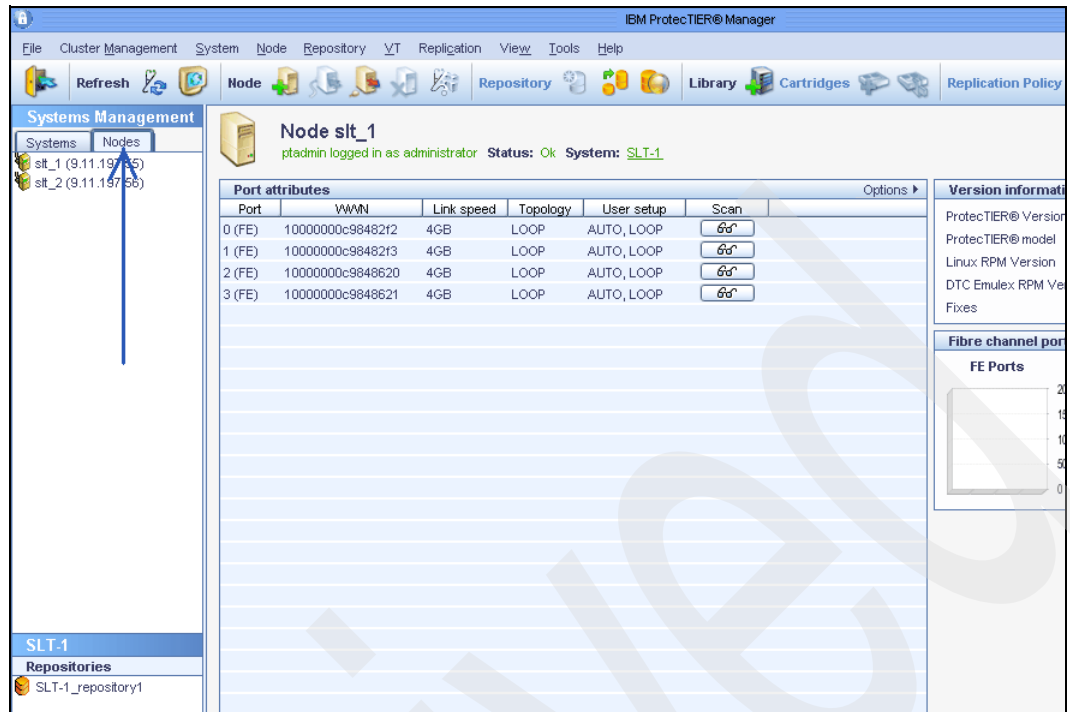


Figure 8-31 ProtecTIER Manager Nodes

- Click **Node** on the ProtecTIER Manager toolbar. The Node pane is displayed.
- Select a node and click **Create problem report** (see Figure 8-32).

slt_1 : Create problem report

Fill in the service report fields to continue.

Customer information

Customer site:

Backup application

Host operating system:

Backup application name:

Backup application version:

Problem description

Briefly describe the problem

Ok Cancel

Figure 8-32 ProtecTIER Manager "Create problem report"

- Complete each of the input fields and click **OK**. A confirmation message box is displayed.

5. Click **Yes**. The ProtecTIER system downloads the report files from the ProtecTIER server to the ProtecTIER Manager workstation and the Create problem report wizard “Add report files” pane is displayed.
6. Click **Ok**. A standard save dialog box is displayed, enabling you to save the report zip file.
7. Click **Save**. The report file is saved to the selected location.

8.9.2 Creating a long-term statistics report

You can generate a long-term statistics report for each node in your system when you want IBM to analyze your system's throughput. After the report is completed, you can save it to your workstation and also send it to IBM Service as a zip file.

Perform the following steps to generate a long-term statistics report for a node:

Note: It is presumed that you are logged on to the ProtecTIER Manager.

1. Click **Nodes**. The Nodes panel is displayed.
2. Select a node and click **Create and download long term statistics**. A confirmation dialog box is displayed.
3. Click **Yes**. The long-term statistics report file is created on the ProtecTIER server and a standard save dialog box is displayed.
4. Save the file on the ProtecTIER Manager workstation.
5. Zip the file and send it along with a support ticket to IBM Service for analysis when directed to do so.

8.9.3 Turning on the WTI power switch outlets

There might be a time when both cluster nodes have been powered off and their power switch outlets are off. Should this happen, you can easily bring the ports back online by using this task.

Perform the following steps to turn on the WTI power switch outlets and bring the ports back online:

1. Locate the Default button.
2. Press and hold the Default button down for three seconds.

8.9.4 Removing a cluster member

Complete this task to remove a cluster member.

You must be logged into ProtecTIER Manager.

There might be times when you are asked to delete a ProtecTIER repository. Before this can be done you must remove one of the cluster members in a two-node configuration.

Perform the following steps to remove the cluster member:

1. Select **Nodes**. The Nodes panel is displayed.
2. Select the node that you want to remove.
3. Select **Node** → **Stop ProtecTIER server** (Figure 8-33 on page 217).




Figure 8-33 ProtecTIER Manager “Select Node”

You are prompted to enter your username and password.

4. Enter your username and password and click **Ok**. The ProtecTIER service stops for the selected node and you are automatically logged out.
5. Click **Login**. You are prompted for your username and password.
6. Enter your username and password and click **Ok**. You are logged in.
7. In the Systems tab of the Navigation pane, select the cluster node that you are removing.
8. Verify that the selected node is offline.
9. Select **Cluster Management** → **Remove cluster member**. The Remove cluster member wizard Welcome panel is displayed.
10. Click **Next**. The Node selection panel is displayed.
11. In the Node IP field, select the node that you want to remove from the cluster.
12. Read the instructions and select the **I read the warnings above and I agree to continue** checkbox.
13. Click **Next** and **Finish**. The Remove cluster member wizard closes and the selected node is removed from the cluster.

8.9.5 Adding a cluster member

You are logged into the ProtecTIER Manager.

The need for this action becomes necessary when you are asked to remove a ProtecTIER repository. When this occurs you generally have to remove one of the nodes from a two-node cluster, and later add it back.

Perform the following steps to add a cluster member into a two-node cluster.

1. Select a cluster in the Systems panel that you want to work with.
2. Select **Cluster Management** and then **Add cluster member**. The Add cluster member wizard Welcome panel is displayed.
3. Click **Next**. The Node Selection panel is displayed.
4. Click the drop-down list in the Node IP field and select the IP address of the node that you want to associate with the cluster.

Note: If the IP address you need is not in the list, you can enter the IP address you need in the input field.

5. Click **Next** and **Finish**. The Add cluster member wizard closes and the node is added to the selected cluster.

8.9.6 Changing World Wide Node Name

Complete this task to change the World Wide Node Name (WWNN).

After replacing Host Bus Adapters (HBA) on a node, or replacing a whole node, you may want to change the World Wide Names (WWN) for the HBAs to match the names from the previous hardware. This minimizes the disruption to the tape processing application. ProtecTIER enables you to do this using the `wwnutil` utility from within the Linux shell environment.

The `wwnutil` utility operates with the following restrictions:

- ▶ Only works with Emulex HBAs.
- ▶ It does not perform a duplication check for name changes.
- ▶ It does not maintain a record of the changes made once the changes have been confirmed.
- ▶ You must have root permission to use the `wwnutil` utility.

Note: Because of these restrictions, it is recommended that you maintain records of the WWNs in order to ensure that WWNN numbers are not duplicated when they are changed.

Perform the following steps to change the WWNN number:

1. From the node server working directory, enter **`service vtfd stop`**. This command stops the ProtecTIER service on the node.
2. Enter `/opt/dtc/app/utls/wwnutil`. This opens the `wwnutil` utility to its main menu (see Figure 8-34 on page 219).

```

+-----+
|                                     |
|               MAIN MENU           |
|                                     |
+-----+
| [1] Show available FE ports names |
| [2] Edit WWN by port instance     |
| [3] View Changes                  |
| [4] Submit or discard changes     |
|                                     |
| [x] Exit                          |
|                                     |
+-----+
Select:

```

Figure 8-34 ProtecTIER Manager “World Wide Node Name”

3. Enter 2 to edit a WWN by port number.
4. Enter a port number.
5. Enter a node name.
6. Enter a port name.
7. Enter `m` to return to the `wwnutil` main menu.
8. Enter 4 and then `c` to accept the name change. The WWNN is changed.
9. Enter **`service vtfd start`** to restart the node server.

After the name change has been confirmed, the list of WWNs reflects the new name, and the name change is no longer registered as a change to the system.

8.9.7 Disabling defragmentation

Complete this task to disable defragmentation.

You must be logged in to the ProtecTIER Manager.

The ProtecTIER system automatically defragments fragmented repository disk space as a background task at a rate that does not cause the system to slow down. Stop defragmentation to free the resources used by the defragmentation process.

Note: Do not initiate this process unless you have been directed to do so by IBM Support.

1. Select **Repository** → **Defragmentation control** (see Figure 8-35 on page 220).




Figure 8-35 ProtecTIER Manager “Disable defragmentation”

2. Select **Disable defragmentation** and click **Ok**. The Defragmentation control pane closes and defragmentation is disabled.

Note: Selecting Enable defragmentation in the Defragmentation control panel resumes system defragmentation.

8.9.8 Disabling data compression

Complete this task to disable data compression.

You must be logged into the ProtecTIER Manager.

Under normal circumstances, the ProtecTIER system compresses data. Stop compression on a specific virtual library to free the resources usually demanded by the compression process.

Note: Do not initiate this stop process unless you have been directed to do so by IBM Support.

Perform the following steps to disable data compression for a specified Virtual Tape Library (VTL).

1. Select **VT**, **VT Library**, and finally **Set compression type**. The compression mode dialog box is displayed.
2. Select **Disable compression** and click **Ok**. The compression mode dialog closes and compression is stopped.

8.9.9 Enabling data compression

Complete this task to enable data compression on a Virtual Tape Library (VTL) that has had the compression process stopped.

You must be logged into the ProtecTIER Manager.

Under normal circumstances, the ProtecTIER system compresses data. You might need to enable the compression process on a specific virtual library after it was initially stopped to free resources used by compression.

Note: Do not initiate this enable process unless you have been directed to do so by IBM Support.

Perform the following steps to enable data compression.

1. Select **VT**, **VT Library**, and finally **Set compression type**. The compression mode dialog box is displayed.
2. Select **Enable compression** and click **OK**. The compression mode dialog box closes and data compression begins once more on the specified VTL library.

8.9.10 Changing the HyperFactor mode

Complete this task to change the HyperFactor mode for a specified library.

You must be logged into the ProtecTIER Manager.

By default, ProtecTIER factors all new incoming data. It detects recurring data and stores only the data segments that have not previously been written to the repository. You can change the default HyperFactor mode for each library, one library at a time.

Note: Do not initiate this change to the HyperFactor mode unless you have been directed to do so by IBM Support.

Perform the following steps to change the HyperFactor mode on a specified Virtual Tape (VTL) library:

1. Select **VT** and **VT Library**.

2. Select the VT library for which you want to stop the HyperFactor mode. Then select **Set HyperFactor mode** (see Figure 8-36). The ProtecTIER VT HyperFactor mode dialog box is displayed.

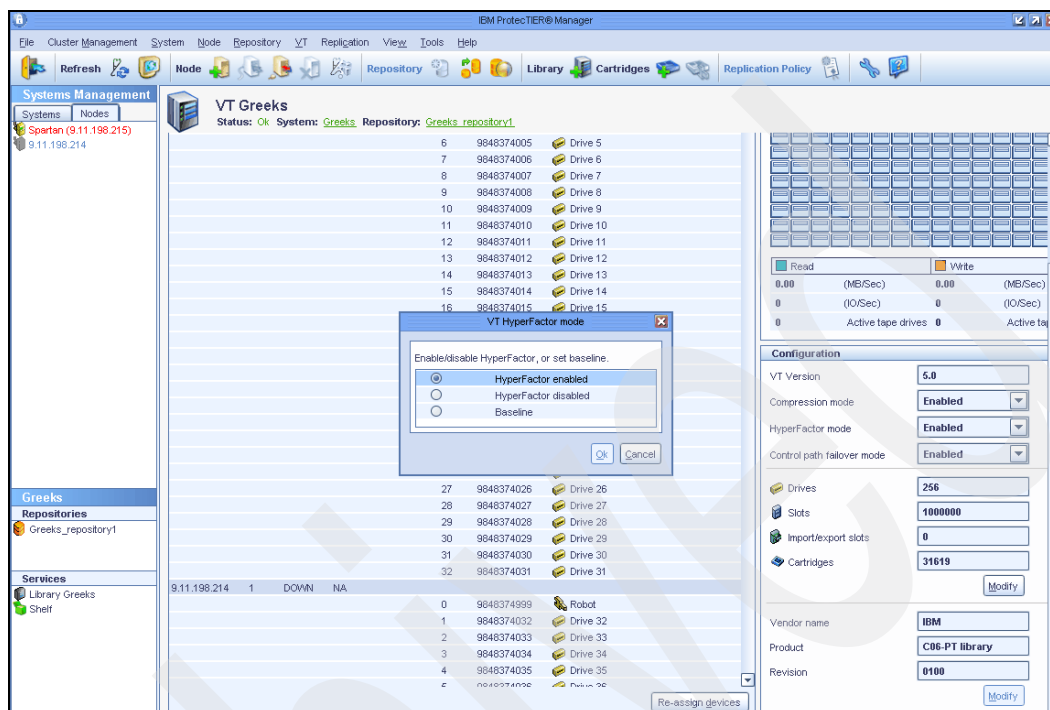


Figure 8-36 ProtecTIER Manager “HyperFactor mode”

3. 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 was stopped.
 - Baseline - HyperFactor stops factoring incoming data and uses the newly stored nonfactored data as the reference for factoring new data after Hyperfactor is resumed.
4. Click **OK**. The ProtecTIER VTL HyperFactor mode dialog box closes and HyperFactor stops for the specified library.

8.9.11 Dumping the trace buffer contents for a node

Complete this task to create a file of the trace buffer contents for a specified node and save it to the ProtecTIER server.

You must be logged in to the ProtecTIER Manager.

The ProtecTIER system stores run-time information in a cyclic memory buffer. IBM Support may direct you to dump the trace buffer for analysis.

Note: Do not initiate this action unless directed to do so by IBM support.

Perform the following steps to generate a file displaying the results of a dump of the trace buffer contents:

1. Select **Nodes**. The Nodes panel is displayed.
2. Select a node and then select **Dump trace**. A confirmation message dialog box is displayed.
3. Click **Yes**.

The trace buffer report is generated for the node and saved to the ProtecTIER server.

Note: Record the file name so that you can find it when IBM Support asks you to send it in with an attached support ticket.

8.9.12 Changing the trace levels

Complete this task to change the trace levels for a specified node.

You must be logged into the ProtecTIER Manager.

The ProtecTIER system traces and records many types of operation information at various levels of detail. These trace levels are initially set in manufacturing.

IBM Support might direct you to reduce the level of detail traced for certain components. This request is made so that system resources are freed up or to increase the level of detail for system components that are suspected to be problematic.

Note: Do not initiate this action unless directed to do so by IBM Support.

Perform the following steps to set the trace levels for a specified node:

1. Click **Node** and then select the node that you want to change the trace settings for (see Figure 8-37 on page 224).

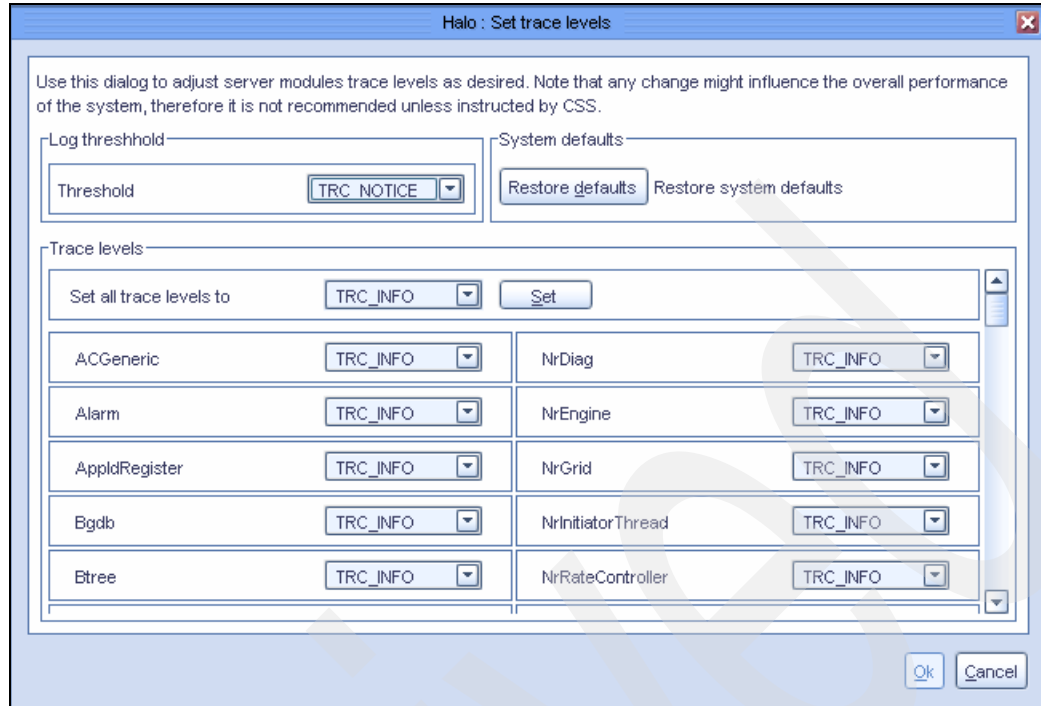


Figure 8-37 ProtecTIER Manager “Trace levels”

2. Select **Set trace levels**. The Set trace levels dialog box is displayed.
3. Change the trace level settings, as directed by IBM Support.
4. Click **Ok**. The Set trace levels dialog box closes and the new trace levels are set.

8.9.13 Resetting the trace buffer

Complete this task to reset the trace buffer on a specified node.

You must be logged in to ProtecTIER Manager.

Resetting the trace buffer empties the buffer. All information that has been gathered for the node is erased when you initiate a reset. If there is any reason to have a record of the node values before a reset, you might decide to do a trace buffer dump and then do the trace buffer reset.

Note: Do not initiate this action unless directed to do so by IBM Support.

Perform the following steps to reset the buffer:

1. Click **Nodes** and then select a node.
2. Click **Reset trace**. A confirmation message dialog box is displayed.
3. Click **Ok**. The trace buffer is reset for the designated node and new information begins to be gathered for that node.

8.9.14 Resetting virtual robots

Complete this task to reset a robot.

You must be logged into ProtecTIER Manager.

There might be times when a virtual robot is locked and you need to reset the robot to break any existing SCSI reservations on the robot.

Notes:

- ▶ Do not initiate this action unless directed to do so by IBM Support.
- ▶ Resetting a robot while the tape processing application is accessing the library can harm the tape processing operations.

Perform the following steps to reset a robot:

1. Navigate to the Services pane and select a library.
2. Select **VT** → **VT Library** → **Reset robot** (see Figure 8-38). A confirmation message dialog box is displayed:

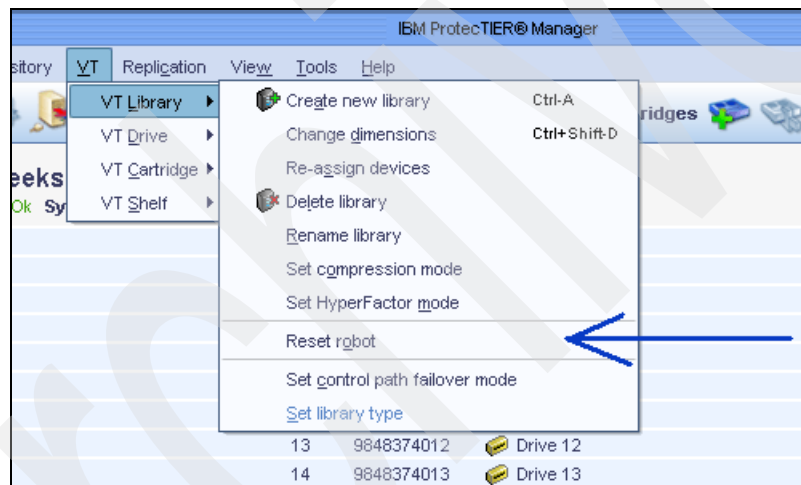


Figure 8-38 ProtecTIER Manager “Resetting virtual robots”

3. Click **Yes**. The robot is reset.

8.9.15 Resetting the virtual tape drives

Complete this task to reset a virtual tape drive.

You must be logged in to ProtecTIER Manager.

There might be times when a virtual tape drive is locked and you need to reset the drive to break any existing SCSI reservations on the robot.

Notes:

- ▶ Do not initiate this action unless directed to do so by IBM Support.
- ▶ Resetting a virtual tape drive while the tape processing application is accessing the library can harm the tape processing operations.

Perform the following steps to reset a virtual tape drive:

1. Navigate to the Services pane and select a library.
2. Click the **Drives** tab.
3. Select a drive.
4. Select **VT** → **VT Drive** → **Reset drive**. A confirmation message dialog box is displayed.
5. Click **Yes**. The virtual tape drive is reset.

8.9.16 Moving cartridges

Complete this task to move a cartridge after it has been disconnected (unloaded) from a virtual tape drive.

You must be logged in to ProtecTIER Manager.

You might need to move a cartridge to a new slot after you have disconnected (unloaded) it from a virtual tape drive slot.

When contemplating this action and only if you are advised to do so by IBM Support, consider the following conditions and proceed accordingly:

Notes:

- ▶ Do not initiate this action unless directed to do so by IBM Support.
- ▶ Moving a cartridge while the tape processing application is accessing the library can harm the tape processing operations. The move is not detected by your tape processing application and can result in the loss of synchronization between your tape processing application and ProtecTIER.

Perform the following steps to disconnect (unload) a cartridge from a drive:

1. Navigate to the Services pane and select a library.
2. Click the **Drives** tab.
3. Select a drive that has no cartridge loaded.
4. Select **VT** → **VT Drive** → **Move cartridge**. The Move cartridge dialog box is displayed.
5. Select one of the following destination types in the Destination Type field:

Drive

Slot

Import / Export

6. Select one of the following destinations in the Destination area:

Next Available - The cartridge is placed in the next available location of the selected type.

Slot/Drive No. - The cartridge is placed in the slot or drive with the number specified in the field. The name of this field depends on your selection in the Destination Type field.

7. Click **Ok**. The cartridge is moved to the specified location.

8.9.17 Checking and repairing errors

The topics in this section define how to bring consistency to your file system and then use ProtecTIER Manager to check and repair system errors.

Before attempting to check and recover errors using ProtecTIER Manager, ensure that the file systems are in a consistent state. If they are not consistent, you might have to run fsck. Do not report errors that require running fsck. Running ProtecTIER Manager's check and recovery wizard on top of inconsistent file systems can harm the consistency of the repository.

Checking for file system consistency

Complete this task to check for file system consistency. This task should be done before you run ProtecTIER Manager to check for and repair system errors.

The purpose of this task is to discover all error messages that indicate the need to run fsck. These error messages consist of wording that says something like the following:

- ▶ Fatal file system consistency error
- ▶ Needs recovery jjd 0 nodeid 2 status 1

Perform the following steps to check for file system consistency errors:

1. Open the `/var/log/messages` file.
2. Search for error messages of the following types:

```
GFS: fsid=<cluster name>:<file system name>: fatal filesystem consistency error
gfs_controlId[2905]: <file system name> finish: needs recovery jjd 0 nodeid 2
status 1
```

Note: If any error messages of these types are found, run fsck on the file systems referenced by the error messages using the file systems' volume names.

Determining a file system's volume name

Complete this task to find a file system's volume name.

You must use the volume names referenced by the error messages when you run fsck to make the file systems consistent.

Perform the following steps to determine a file system's volume name.

1. Run `gfs_tool df`. Summary information for all Global File System (GFS) file systems in your repository is displayed.
2. Locate the desired file system names and determine their mount points.

Figure 8-39 on page 228 gives an example of the file that is produced when you run `gfs_tool df`. The file system name is in bold and the mount point name is italicized, which you can use as a guide when reviewing your report.

```

/mnt/fs_4
SB lock proto = "lock_dlm"
SB lock table = "romeo_juliet:gfs_sda4_new"
SB ondisk format = 1309
SB multihost format = 1401|
Block size = 4096
Journals = 3
Resource Groups = 720
Mounted lock proto = "lock_dlm"
Mounted lock table == "romeo_juliet:gfs_sda4_new"
Mounted host data = "jid=0:id=589825:first=1"
Journal number = 0
Lock module flags = 0
Local flockes = FALSE
Local caching = FALSE
Oopses OK = FALSE

```

Type	Total	Used	Free	use%
inodes	6	6	0	100%
metadata	1296	17	1279	1%
data	47081258	0	47081258	0%

Figure 8-39 Example run "gfs_tool df

3. Open etc/fstab.
4. Locate the mount points determined using gfs_tool df, and determine the corresponding logical volume names.

The following is an example of the fstab entry format:

```

/dev/<volume group name>/<logical volume name> /mnt/<mount point
name> gfs defaults,noatime,nodiratime,noquota 0 0

```

Running fsck to provide file system consistency

Complete this task to run fsck to make the file systems consistent before you run ProtecTIER Manager to check and repair any ProtecTIER system errors.

To correct the consistency problem using fsck, do the following:

1. Shut down one of the active ProtecTIER nodes.
2. Stop the specified services in the remaining active node.
3. Run fsck on the problem file systems.
4. Restart the services.
5. Restart the node that you shut down.

Perform the following steps to run fsck.

1. Select a node and enter service vtfd stop to stop the ProtecTIER service.
2. Shut down the node that you just stopped service on.
3. Switch to the active node and enter service vtfd stop.
4. Enter service gfs stop on this node.
5. Run fsck on the problematic file systems. Enter gfs_fsck <logical volume name> for each file.
6. Enter service gfs start on the active node. This restarts the GFS service on the node.

7. Enter service `vtfd` start on the active node.
8. Restart the node that you turned off.

Checking the system using the ProtecTIER Manager

Complete this task to check the system for errors. This is a time-consuming process and you might want to use it only when there is a suspected severe problem.

Ensure that problematic file systems have been repaired to ensure consistency by running `fsck`. Running ProtecTIER Manager's check and recovery wizard on top of inconsistent file systems can harm the consistency of the repository.

1. You must be logged in to ProtecTIER Manager.

The check and recovery process using ProtecTIER Manager is time-consuming. The ProtecTIER system goes offline for the duration of the process. It is therefore recommended that the Check and recover wizard only be used for severe problems.

Perform the following steps to use the ProtecTIER Manager to check for errors:

1. Select **System** and then **Check and recover**. A confirmation dialog box is displayed.
2. Click **Yes**. The ProtecTIER system goes offline and scans itself. After the scan is completed, the Check and recover dialog window is displayed with the results of the scan.
3. Analyze the results of the scan and determine what your next actions should be. The following types of results are displayed:
 - Positive checks
There are no errors. No further action is needed.
 - ProtecTIER recoverable errors
These errors are repairable by using the repair feature in ProtecTIER Manager.
 - Support required errors
These errors cannot be repaired without the help of IBM Support. Contact IBM Support for further assistance.

Repairing the system using ProtecTIER Manager

Complete this task to use ProtecTIER Manager to repair system errors that were discovered during the system check.

You must be logged in to the ProtecTIER Manager.

The systems check process has completed and you have done your analysis of the diagnostic report.

You have allotted enough time for the repair process to complete. The recovery process is time-consuming and the ProtecTIER system goes offline for the duration of that process.

Perform the following steps to repair the ProtecTIER recoverable errors:

1. Select the checkbox in the Categories subpane for each category for which you want to repair errors (see Figure 8-40 on page 230).

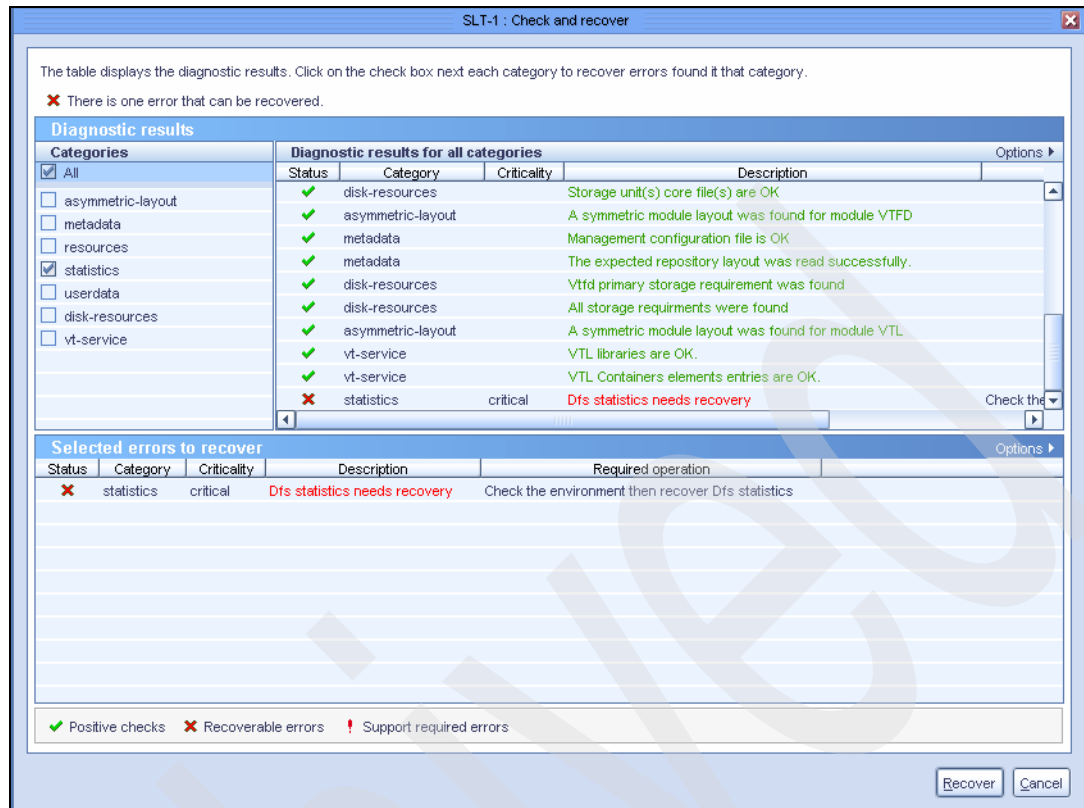


Figure 8-40 ProtecTIER Manager “Repairing the system”

2. Click **Recover**. ProtecTIER Manager attempts to repair the errors.
3. Check the list of transactions after the repair process has completed. If any transaction failed, call IBM Support.

Monitoring the system

The ProtecTIER Manager application is the main monitoring tool for the ProtecTIER software installed on the TS7680 systems. In this chapter, we describe what to look for in the ProtecTIER Manager user interface and what actions to take when problems are encountered in this window. We also describe other means of monitoring the system.

We cover the following topics:

- ▶ Monitoring the ProtecTIER software
- ▶ Monitoring the ProtecTIER VTL Service
- ▶ Reporting on ProtecTIER activity
- ▶ Monitoring the virtual library through the z/OS host

9.1 Monitoring ProtecTIER

In this section, we describe the information about the ProtecTIER systems that is available through the ProtecTIER Manager application. Figure 9-1 shows the first window you will see after starting the ProtecTIER Manager application on your workstation. We describe the different parts of this window in detail in the following sections.

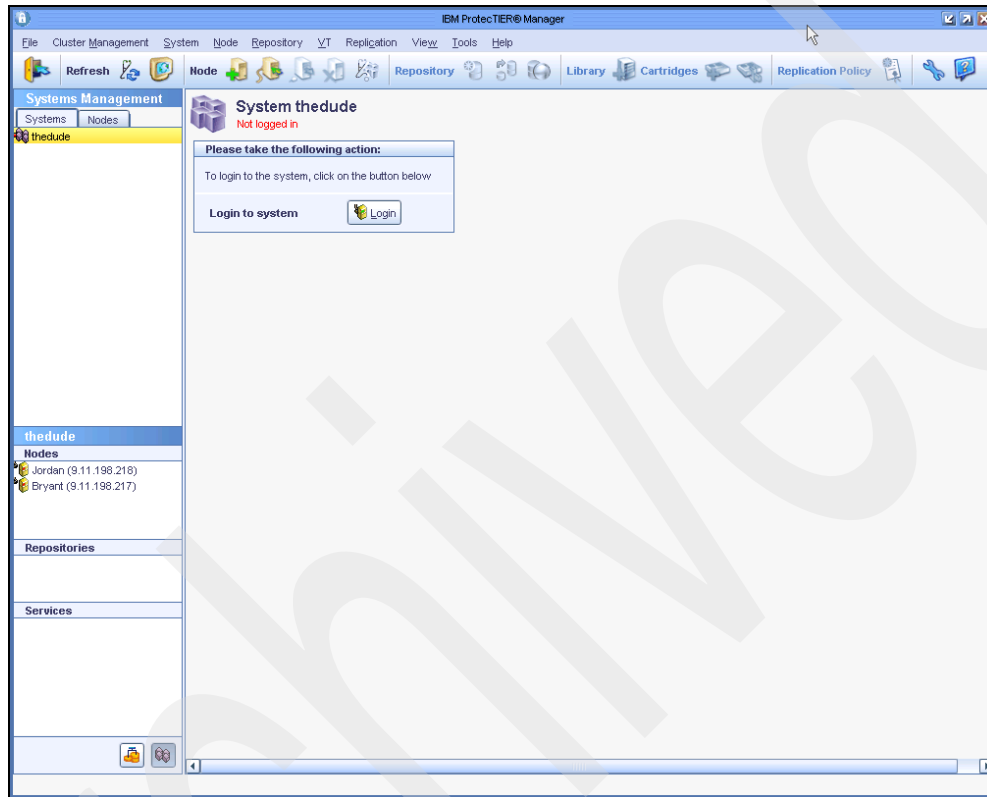


Figure 9-1 Initial login window in ProtecTIER Manager

Click **Login** to begin your ProtecTIER Manager session. You will be shown a dialog box, asking for your user ID and password to authenticate the session (see Figure 9-2).




Figure 9-2 User ID and password dialog box in ProtecTIER Manager

Enter your user ID and password details. If desired, you can have ProtecTIER Manager save your user ID and password details for future logins by checking the **Save password** check box. Once your user ID and password have been successfully authenticated, the Systems window in Figure 9-8 on page 236 will be displayed.

9.1.1 The Status line

At the bottom of every window in ProtecTIER Manager there is a brief system status displayed on a single line (see Figure 9-3.)




Figure 9-3 Status line in ProtecTIER Manager

It lists the following information, from left to right:

- ▶ The current state of the Virtual Tape task, as indicated by text and color:
 - Online/Green/Static for available
 - Offline/Red/Blinking for unavailable
- ▶ The user ID name you have logged in with, the system you are on and the user authority level you are using.
- ▶ The current read and write throughput rates in MB/sec.
- ▶ The current system time for the ProtecTIER (supplied from the TS7680 server running Linux).
- ▶ The Alerts and Events Log buttons.

The Alerts button

The Alerts button turns red and blinks if there are any alert or error messages to be viewed. Click **Alerts** to view these messages in the Alerts Log window (see Figure 9-4.)

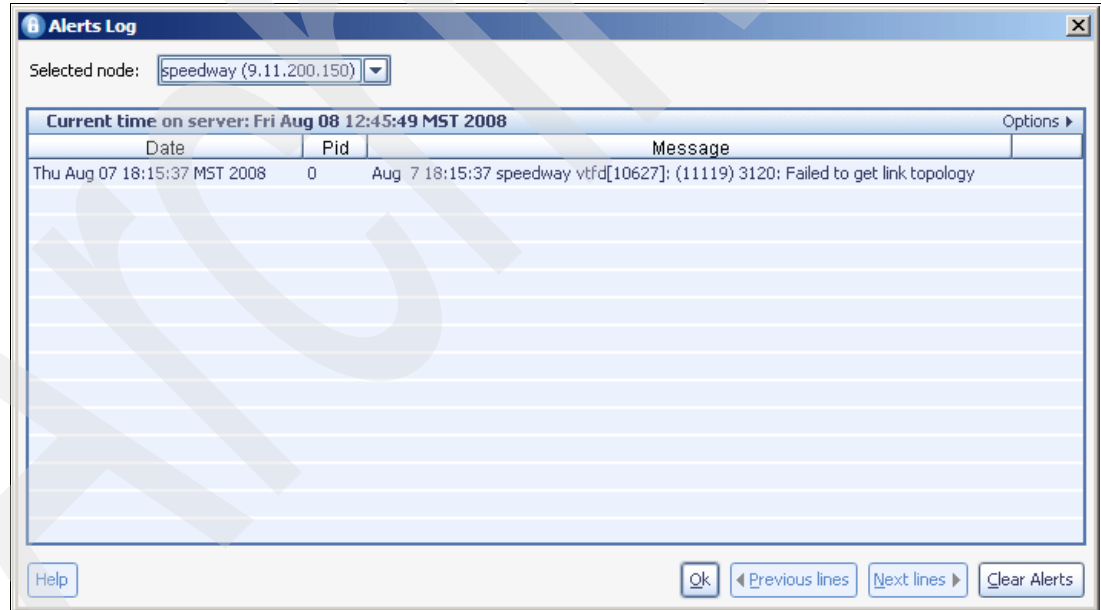


Figure 9-4 Alerts window in ProtecTIER Manager

Once you have finished reading the messages, you can click **Clear Alerts** to acknowledge and remove the messages from the view. They will still remain in the ProtecTIER Manager events log, along with all information messages. If you choose to exit without clearing them, click **Ok** to close the window. When you return to the main ProtecTIER Manager window, the Alerts button will have stopped blinking and be greyed out.

The Events Log button

Click **Events Log** to view the log of all actions you have performed during the current session, displayed in a window (see Figure 9-5.) Click **OK** to close the window.

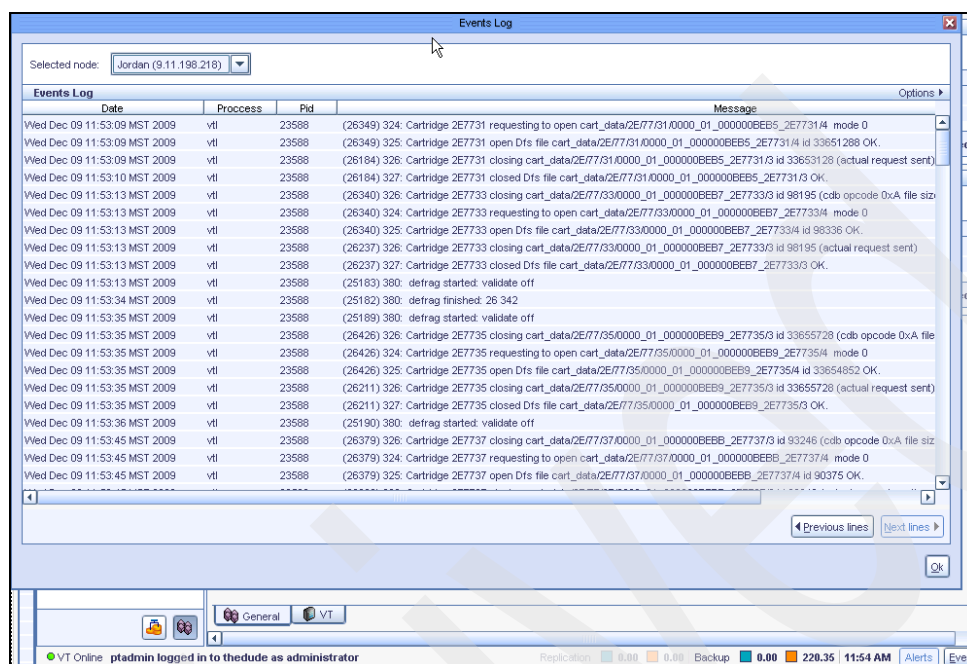


Figure 9-5 Events Log window in ProtectTIER Manager

Both the Alerts Log and the Events Log windows only display information for one node in a system at a time. The node whose messages are currently being displayed is shown in the Selected Node field. To change to another node in the same system, click the down arrow in the Selected Node field to open a menu, where you can select a new node by name.

In addition, the Alerts Log and Events Log windows display only up to 200 messages per page. You can navigate to other pages of messages by clicking **Previous Lines** or **Next Lines**.

9.1.2 The Navigation pane

The Navigation pane is located on the left side of the ProtectTIER Manager window (see Figure 9-6 on page 235) and is constant across all windows. It runs the entire height of the window and contains navigational features that can be used to move quickly to any desired view.

At the top of the Navigation pane are two tabs, Systems and Nodes. These lead to the Systems window (see 9.1.3, “The Systems window” on page 235) and the Nodes window (see 9.1.4, “The Nodes window” on page 240.) The other features are described now.

The Nodes pane

The Nodes pane is on the left side of the window below the Navigation pane (see Figure 9-6 on page 235). Its heading is the name of the system you are currently viewing (the system is called “thedude” in this example.) The Nodes pane contains the names and IP addresses of nodes that are defined to the system.




Figure 9-6 Nodes pane, General tab, Systems window in ProtectTIER Manager

The Repositories and Services panes

The Repositories pane and the Services pane are both on the left side of the window below the Nodes Pane (see Figure 9-7.) The Repositories pane contains the repository that is defined to the system you are currently viewing (there is only ever one repository per system.) It is explained in more detail in 9.1.5, “The Repository window” on page 244.

The Services pane contains the name of the Virtual Tape Library (VTL) defined to the system you are currently viewing. This library is explained in more detail in 9.2.1, “The Library window” on page 248.



Figure 9-7 Repositories and Services panes in ProtectTIER Manager

9.1.3 The Systems window

When you first log in to ProtectTIER Manager, you see the General tab of the Systems window (see Figure 9-8 on page 236).

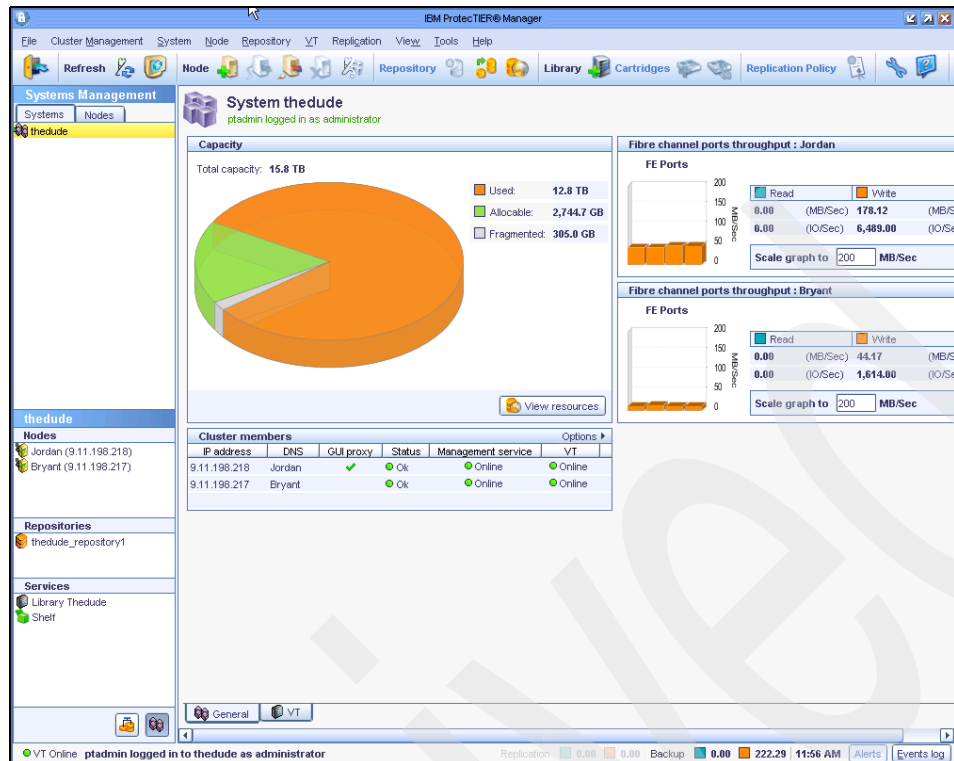


Figure 9-8 General tab, Systems window in ProtecTIER Manager

The General tab

In the General tab of the Systems window, the name of the system you have logged in to is automatically selected in the Systems tab of the Navigation pane on the left side. It is also displayed along the top line under the menu and tool bars.

Other systems ProtecTIER Manager has detected are also displayed in this tab. On the left side of the icon for each system is a small dot that is colored green if you are logged in to that system and is colored red if you are not.

You can change to another system by clicking a different system name. You must log in to each system separately.

When you select the **Nodes** tab in the Navigation pane, ProtecTIER Manager displays all the nodes it knows about (see Figure 9-16 on page 240). Again, there is a small part of the icon for each node that is colored red or green depending if you are logged into that system or not.

The General tab of the Systems window has a short summary of the system as a whole. More detailed information about each component can be found in the following window descriptions in this chapter.

The important items to note in the General tab of the Systems window are:

- ▶ The repository summary
- ▶ The Fibre Channel port status
- ▶ The cluster member and virtual tape library status

The Capacity section

The Capacity section (see Figure 9-9) is a graphical summary of the current state of the repository of the system you are currently viewing. It is exactly the same as the view in the Repository window, and is explained in more detail in “The Capacity section” on page 244.




Figure 9-9 Capacity section, General tab, Systems window in ProtecTIER Manager

The Fibre Channel Port Throughput section

The Fibre Channel Port Throughput section (Figure 9-10) displays the current Fibre Channel port throughput (both read and write) for each node. It is exactly the same as the view in the Nodes window, and is explained in more detail in “The Fibre Channel Ports Throughput section” on page 242.

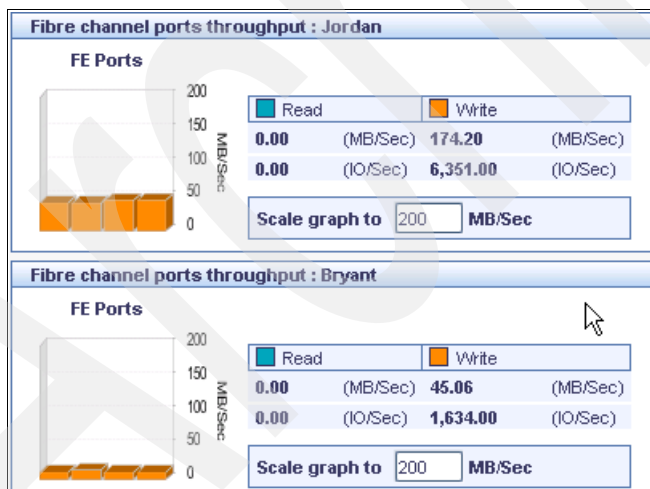


Figure 9-10 Fibre Channel Port Throughput section, General tab, Systems window in ProtecTIER Manager

The Cluster Members section

The Cluster Members section (Figure 9-11 on page 238) displays the Internet Protocol (IP) address and Domain Name Server (DNS) names for each node in the system you are currently viewing. The GUI Proxy column indicates which node is currently being used by ProtecTIER Manager to monitor the two-node cluster. The Status column indicates whether the ProtecTIER software is online or offline on the node and the VT column indicates whether the VTL service for that node is online or offline.

Cluster members						Options ▾
IP address	DNS	GUI proxy	Status	Management service	VT	
9.11.198.218	Jordan	✓	● Ok	● Online	● Online	
9.11.198.217	Bryant		● Ok	● Online	● Online	

Figure 9-11 Cluster Members section, General tab, Systems window in ProtecTIER Manager

The VT tab

In the Systems window, there is a second tab that can be selected for viewing, labelled VT. When you click the **VT** tab next to the General tab in the bottom section of the window, the window shown in Figure 9-12 will appear.

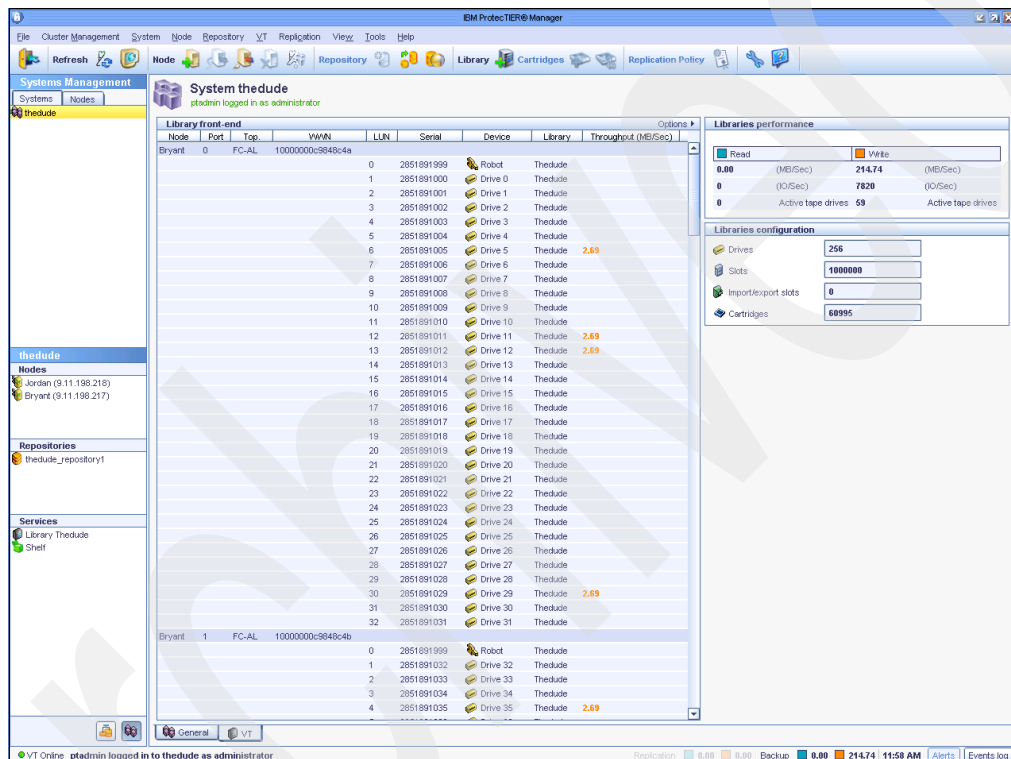


Figure 9-12 VT tab, Systems window in ProtecTIER Manager

In the VT tab, Systems window, the name of the system you have logged in to is automatically displayed in the Systems tab of the Navigation pane on the left side. It is also displayed along the top line under the menu and tool bars. You can change to another system by clicking a different system name in the Navigation pane at any time.

The important items to note in the VT tab of the Systems window are:

- ▶ All defined tape devices for this system
- ▶ Total current Front End Fibre Channel port throughput rates
- ▶ VTL configuration totals

The Library Front End section

The Library Front End section (Figure 9-13 on page 239) lists all the virtual tape devices (robots and drives) across all the VTLs defined to the system you are currently viewing, sorted in the order of node, Front End port number, LUN, device, and library name. It is the

same as the view in the General tab of the Library window (except that is across *all* VTLs.) The columns are explained in more detail in “The Library front-end section” on page 249.

Library front-end									Options
Node	Port	Top.	WWN	LUN	Serial	Device	Library	Throughput (MB/Sec)	
Bryant	0	FC-AL	10000000c9848c4a	0	2851891999	Robot	The dude		
				1	2851891000	Drive 0	The dude		
				2	2851891001	Drive 1	The dude		
				3	2851891002	Drive 2	The dude		
				4	2851891003	Drive 3	The dude		
				5	2851891004	Drive 4	The dude		
				6	2851891005	Drive 5	The dude	2.66	
				7	2851891006	Drive 6	The dude		
				8	2851891007	Drive 7	The dude		
				9	2851891008	Drive 8	The dude		
				10	2851891009	Drive 9	The dude		
				11	2851891010	Drive 10	The dude		
				12	2851891011	Drive 11	The dude	2.66	
				13	2851891012	Drive 12	The dude	2.66	
				14	2851891013	Drive 13	The dude		
				15	2851891014	Drive 14	The dude		
				16	2851891015	Drive 15	The dude		
				17	2851891016	Drive 16	The dude		
				18	2851891017	Drive 17	The dude		
				19	2851891018	Drive 18	The dude		
				20	2851891019	Drive 19	The dude		
				21	2851891020	Drive 20	The dude		
				22	2851891021	Drive 21	The dude		
				23	2851891022	Drive 22	The dude		
				24	2851891023	Drive 23	The dude		
				25	2851891024	Drive 24	The dude		
				26	2851891025	Drive 25	The dude		
				27	2851891026	Drive 26	The dude		
				28	2851891027	Drive 27	The dude		
				29	2851891028	Drive 28	The dude		
				30	2851891029	Drive 29	The dude	2.66	
				31	2851891030	Drive 30	The dude		
				32	2851891031	Drive 31	The dude		
Bryant	1	FC-AL	10000000c9848c4b	0	2851891999	Robot	The dude		
				1	2851891032	Drive 32	The dude		
				2	2851891033	Drive 33	The dude		
				3	2851891034	Drive 34	The dude		
				4	2851891035	Drive 35	The dude	2.66	

Figure 9-13 Library Front End section, VT tab, Systems window in ProtecTIER Manager

The Libraries Performance section

The Libraries Performance section (Figure 9-14) shows the cumulative usage and performance figures for all the VTLs defined to the system you are currently viewing. This display includes the Read and Write throughput rates in MB/sec and IO/sec and how many total tape drives are active for each type of operation (Read and Write).

Libraries performance			
Read		Write	
0.00	(MB/Sec)	216.28	(MB/Sec)
0	(IO/Sec)	7873	(IO/Sec)
0	Active tape drives	59	Active tape drives

Figure 9-14 Libraries Performance section, VT tab, Systems window in ProtecTIER Manager

The Libraries Configuration section

The Libraries Configuration section (Figure 9-15 on page 240) shows the cumulative totals for library elements, such as drives, slots, I/E slots, and cartridges, across all the VTLs defined to the system you are currently viewing.

Note: The TS7680 comes with a fixed configuration with 256 drives, 1,000,000 slots, no Import/Export slots, and as many cartridges as you defined for this library.

Libraries configuration	
Drives	256
Slots	1000000
Import/export slots	0
Cartridges	60995

Figure 9-15 Libraries Configuration section, VT tab, Systems window in ProtectTIER Manager

9.1.4 The Nodes window

You can access the Nodes window by clicking the **Nodes** tab in the Navigation pane on the left side of the window. This will list all the nodes currently defined to ProtectTIER Manager. Select the node you wish to display and click its name. If you are not yet logged in to that system, ProtectTIER Manager will display the node login window (see Figure 9-1 on page 232.) If you are logged in already, the window shown in Figure 9-16 will appear.

Node Jordan
ptadmin logged in as administrator Status: Ok System: thedude

Port	WWN	Link speed	Topology	User setup	Scan
0 (FE)	10000000c98485d6	4GB	LOOP	AUTO, LOOP	60"
1 (FE)	10000000c98485d7	4GB	LOOP	AUTO, LOOP	60"
2 (FE)	10000000c9848db6	4GB	LOOP	AUTO, LOOP	60"
3 (FE)	10000000c9848dbf	4GB	LOOP	AUTO, LOOP	60"

Version information

ProtectTIER® Version	1.0.0.5
ProtectTIER® model	TS7680
Linux RPM Version	7120.016-V1.0.0.5
DTC Emulex RPM Version	5213.003-1
Fixes	Not included

Fibre channel ports throughput

FE Ports

Read (MB/Sec)	Write (MB/Sec)
0.00	169.16
0.00	6,166.00

Scale graph to 200 MB/Sec

Network configuration

Device	IP address	Usage	Master device	Status	Speed (Mbit)	MTU (BYTES)
bond0	10.0.0.52	Internal				
eth1		Internal	bond0	up	1000	1500
eth2		Internal	bond0	up	1000	1500
eth0	9.11.198.218	External		up	100	1500

Figure 9-16 Nodes window in ProtectTIER Manager

In the Nodes window, the name of the node you have logged in to is displayed along the top line under the menu and tool bars.

The important items to note in the Nodes window are:

- ▶ The port attributes for the four Front End (FE) ports
- ▶ The ProtecTIER version information
- ▶ The Fibre Channel ports throughput display
- ▶ The network configuration

The Port attributes section

The current state of health for each of the four FE ports associated with the node is displayed in the Port attributes section. The worldwide name (WWN) for each port is displayed along with the Link speed, Topology (and state of the link), and the User setup (see Figure 9-17).

Port attributes						Options ▾
Port	WWN	Link speed	Topology	User setup	Scan	
0 (FE)	10000000c98485d6	4GB	LOOP	AUTO, LOOP		
1 (FE)	10000000c98485d7	4GB	LOOP	AUTO, LOOP		
2 (FE)	10000000c9848db0	4GB	LOOP	AUTO, LOOP		
3 (FE)	10000000c9848dbf	4GB	LOOP	AUTO, LOOP		

Figure 9-17 Port Attributes section, Nodes window in ProtecTIER Manager

The Link speed column shows the transmission speed of the port. Possible values are:

- ▶ AUTO - A transmission speed that is auto-negotiated between the two ports depending on the combined highest possible link speed.
- ▶ 1 Gb - A fixed transmission speed of 1 Gigabit per second.
- ▶ 2 Gb - A fixed transmission speed of 2 Gigabits per second.
- ▶ 4 Gb - A fixed transmission speed of 4 Gigabits per second.
- ▶ DOWN - There is no Fibre Channel connection.

The Topology column displays the Fibre Channel topology of the port. Possible values are:

- ▶ LOOP - Fibre Channel Arbitrated Loop connection.
- ▶ P2P - Peer to Peer connection.
- ▶ DOWN - There is no Fibre Channel connection.

The User setup column is the user-assigned link speed and topology. Possible values are a combination of the Link speed and Topology column values above, separated by a comma.

There is also the Scan button (marked with a pair of glasses icon), displayed in the far right column for each port. Clicking this icon will open the Scan Port dialog box. Scanning the port displays a numbered list of the Worldwide Names (WWNs) of the remote ports detected by the port. This is useful during the initial setup or when diagnosing problems.

Note: Scanning the port will cause a disruption to any active traffic using the link. A dialog box is displayed, asking you to confirm that you want to continue with the port scan (see Figure 9-18 on page 242.)

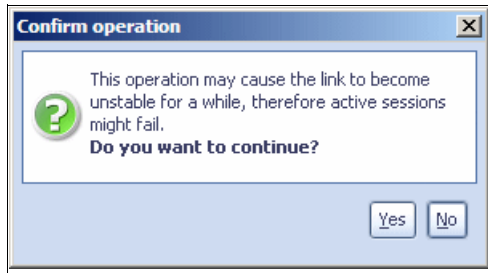


Figure 9-18 Scan Port confirmation dialog box in ProtecTIER Manager

Once the scan port operation has completed, the window shown in Figure 9-19 appears.

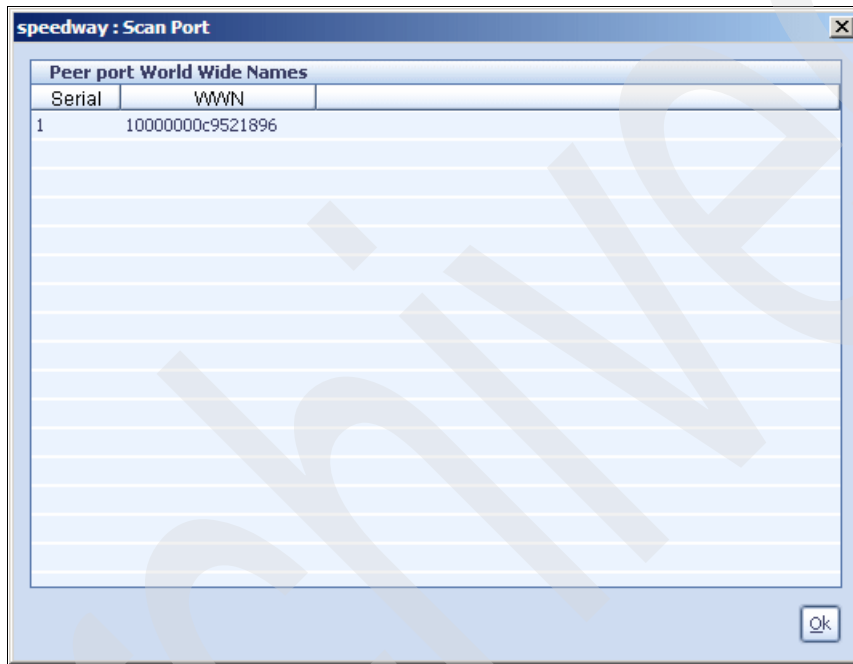


Figure 9-19 Scan Port result window in ProtecTIER Manager

The Version Information section

The Version Information pane (Figure 9-20) displays information about the version of the ProtecTIER, the model, in this case TS7680, the Linux RPM version, and the DTC Emulex RPM installed and running on the node.

Version information	
ProtecTIER® Version	1.0.0.5
ProtecTIER® model	TS7680
Linux RPM Version	7120.016-V1.0.0.5
DTC Emulex RPM Version	5213.003-1
Fixes	Not included

Figure 9-20 Version Information section, Nodes window in ProtecTIER Manager

The Fibre Channel Ports Throughput section

The Fibre Channel (FC) Ports Throughput pane displays the rate of data movement and I/O operations for both Read and Write operations for the node (see Figure 9-21 on page 243). The data movement rate is also displayed graphically for each Front End Fibre Channel port

on the node. The bars will be colored blue (for Read) or orange (for Write). There is enough space for four bars to be displayed at once in the bar graph, one bar for each FC port.

You can change the scale of the graph by editing the value in the Scale Graph To field, to see the throughput rates in finer or less details.




Figure 9-21 Fibre Channel Ports Throughput section, Nodes window in ProtecTIER Manager

The Network Configuration section

The Network Configuration pane displays information about the setup of the network interface cards (NIC) for the node (see Figure 9-22.)

Network configuration							Options ▾
Device	IP address	Usage	Master device	Status	Speed (MBit)	MTU (BYTES)	
bond0	10.0.0.52	Internal					
eth1		Internal	bond0	up	1000	1500	
eth2		Internal	bond0	up	1000	1500	
eth0	9.11.198.218	External		up	100	1500	

Figure 9-22 Network Configuration section, Nodes window in ProtecTIER Manager

The column values are explained in Table 9-1.

Table 9-1 Column definitions for Network Configuration values in ProtecTIER Manager

Column	Definition
Device	<p>The devices in the NIC.</p> <ul style="list-style-type: none"> ► Eth0 is the node port that communicates with the ProtecTIER Manager workstation. ► Eth1 and Eth2 are the node ports used in the two node cluster-internal network. ► Bond0 is the virtual bond master device to which Eth1 and Eth2 are enslaved. <p>Bond devices are defined as part of the installation process.</p>
IP Address	IP address of the device.
Usage	Indicates whether the device is used for the two node cluster-internal network or to communicate with the ProtecTIER Manager workstation and the external network.
Master Device	The master device or bond device, if any, to which the device is enslaved.
Status	Indicates whether the device is functioning properly.
Speed	The supported speed of data transfer across the device in Megabits per second.
MTU	Configured maximum transmission unit for the device.

9.1.5 The Repository window

You can access the Repository window by clicking a repository name in the Repositories pane on the left side of any window (see Figure 9-23.).

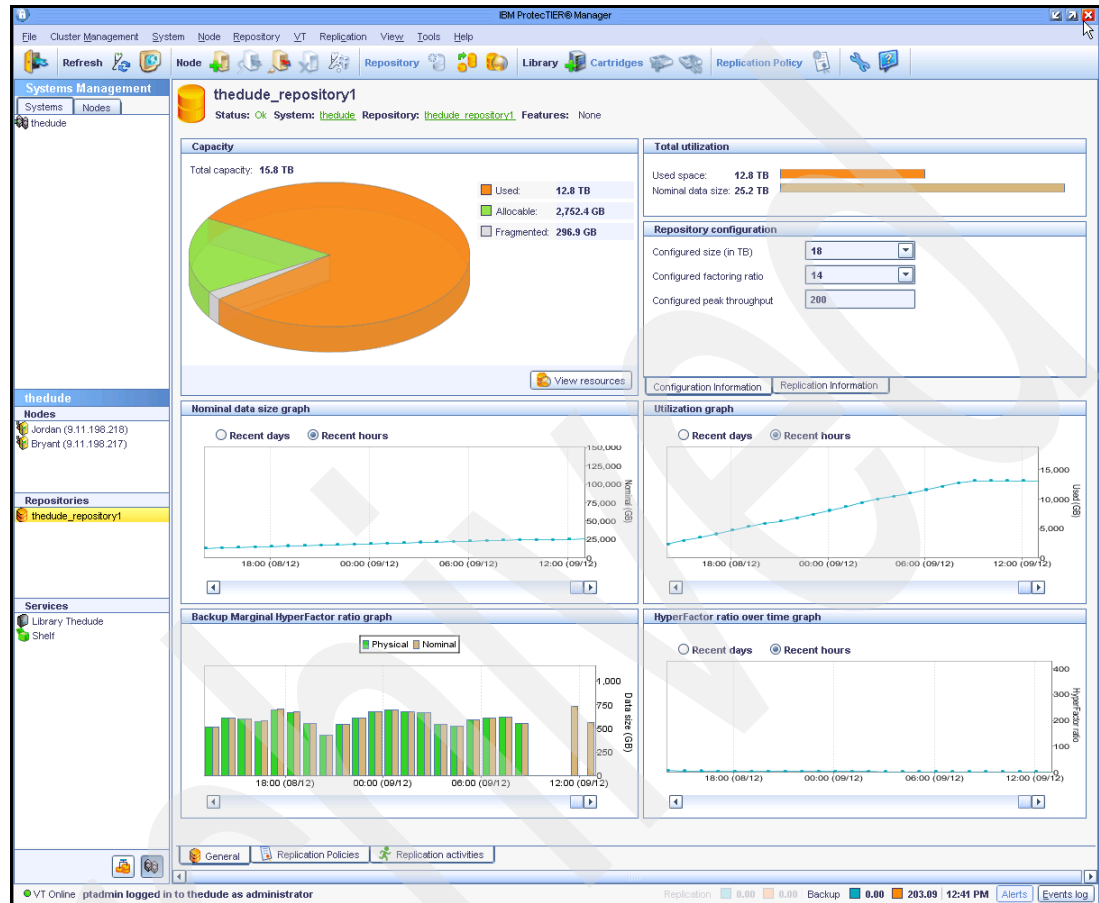


Figure 9-23 Repository window in ProtecTIER Manager

The Capacity section

The repository “state of health” is displayed in a pie chart in the Capacity section near the center of the window. This gives a visual indication of the current capacity status. It also shows three repository capacity totals: Used, Allocable and Fragmented.

Used

This total represents the amount of space presently occupied in the repository *after* any data deduplication has taken place. It is the real consumption level of the space allocated to the repository.

Allocable

This total represents the amount of unoccupied space in the repository.

Fragmented

This total represents the amount of data in the repository that is in pieces smaller than 1 MB. ProtecTIER runs housekeeping tasks in the background to constantly reassemble these small fragments into larger pieces of free space to make the best use of all available space.

Note: Repository Space Usage

The Repository view represents the space used on all the virtual cartridges defined, regardless of the logical state (“active” or “expired”) of the data in the tape processing application that created the data. Once a virtual cartridge has been filled, the amount of data it contains is included in the Used portion of the pie chart in the repository view. There is no reconciliation between the tape processing application and the ProtecTIER repository to delete expired data from the repository. This emulates the process used for physical cartridges; when a tape processing application logically expires data, physically the data is still there on the cartridge media surface. See “Steady state” on page 26 for more information.

The View Resources button

The View Resources button just below the pie chart graphic displays the layout of the repository at the Linux file system level in a window (see Figure 9-24). It shows the mount point names, the file system type (always GFS, meaning global file system), the file system size, and the current usage type of the file system (either User Data or Meta Data.) Nothing can be changed from this window; it is intended for informational purposes only.

Mount point	Device	FS type	Size	Usage	RAID type	Disk size(GB)
/mnt/vg0-lv_vg0	/dev/mapper/vg0-lv_vg0	gfs	1,674.2 GB	User data		
/mnt/vg1-lv_vg1	/dev/mapper/vg1-lv_vg1	gfs	2,091.9 GB	Metadata	FC-15K 8+8	450
/mnt/vg10-lv_vg10	/dev/mapper/vg10-lv_vg10	gfs	1,674.2 GB	User data		
/mnt/vg11-lv_vg11	/dev/mapper/vg11-lv_vg11	gfs	1,674.2 GB	User data		
/mnt/vg2-lv_vg2	/dev/mapper/vg2-lv_vg2	gfs	635.0 MB	Metadata		
/mnt/vg3-lv_vg3	/dev/mapper/vg3-lv_vg3	gfs	1,674.2 GB	User data		
/mnt/vg4-lv_vg4	/dev/mapper/vg4-lv_vg4	gfs	1,674.2 GB	User data		
/mnt/vg5-lv_vg5	/dev/mapper/vg5-lv_vg5	gfs	1,674.2 GB	User data		
/mnt/vg6-lv_vg6	/dev/mapper/vg6-lv_vg6	gfs	1,674.2 GB	User data		
/mnt/vg7-lv_vg7	/dev/mapper/vg7-lv_vg7	gfs	1,674.2 GB	User data		
/mnt/vg8-lv_vg8	/dev/mapper/vg8-lv_vg8	gfs	1,674.2 GB	User data		
/mnt/vg9-lv_vg9	/dev/mapper/vg9-lv_vg9	gfs	1,674.2 GB	User data		

Figure 9-24 View resources window in ProtecTIER Manager

The Total utilization section

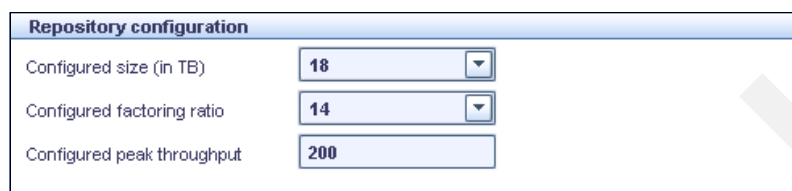
The Total utilization section of the Repository window is an at-a-glance view of the nominal data received compared to the physical data stored, with totals (see Figure 9-25).



Figure 9-25 Total utilization section, Repository window in ProtecTIER Manager

The Repository configuration section

The Repository configuration section of the Repository window lists important repository information in a single place (see Figure 9-26).



Repository configuration	
Configured size (in TB)	18
Configured factoring ratio	14
Configured peak throughput	200

Figure 9-26 Repository configuration section, Repository window in ProtecTIER Manager

It contains the information shown in Table 9-2.

Table 9-2 Field definitions for Repository configuration values in ProtecTIER Manager

Field	Definition
Configured size (in TB)	The physical repository size in terabytes (User Data only).
Configured factoring ratio	The estimated HyperFactor factoring ratio that was used to create the repository.
Configured peak throughput	The expected maximum peak throughput specified when the repository was created.

The ProtecTIER Manager data graph section

Depending on your window resolution size, the following graphs will appear in tabs or displayed directly, one above the other. The following paragraphs describe the tabbed display.

Four tabs showing two different graphical views of the data sent to and stored by the ProtecTIER are displayed beneath the Capacity section in the Repository window. The graphs are named:

1. The Nominal Data Size graph
2. The Utilization graph
3. The Marginal HyperFactor ratio graph
4. The HyperFactor ratio over time graph

The Nominal Data Size graph

By clicking the **Nominal Data Size graph** tab, you can display how much data has been sent to ProtecTIER over time (see Figure 9-27).




Figure 9-27 Nominal Data Size graph tab, Repository window in ProtecTIER Manager

This graph represents the amount of data written to the ProtecTIER systems by the tape processing application without any compression or deduplication. You can display data figures from recent hours or recent days using the radio buttons in the top left corner of the section. There is also a scroll bar to view earlier historical data, if available.

The Utilization graph

By clicking the **Utilization graph** tab, you can display how much actual data has been stored by ProtecTIER over time (see Figure 9-28).




Figure 9-28 Utilization graph tab, Repository window in ProtecTIER Manager

This graph represents the amount of data written to the disk repository by the ProtecTIER system after deduplication and compression has taken place. It is likely to be significantly less than the nominal data amount. You can display data figures from recent hours or recent days using the radio buttons in the top left corner of the section. There is also a scroll bar to view earlier historical data, if available.

The Marginal HyperFactor ratio graph

By clicking the **Marginal HyperFactor ratio graph** tab, you can see how well ProtecTIER is performing (see Figure 9-29).




Figure 9-29 Marginal HyperFactor ratio tab, Repository window in ProtecTIER Manager

This section gives a visual comparison through a bar graph of data written by ProtecTIER (physical, green) versus data received by ProtecTIER (nominal, grey) over the previous few hours, summarized for each hour. It is possible to immediately see whether the data sent recently has factored well or not.

The HyperFactor ratio over time graph

By clicking the **HyperFactor ratio over time graph** tab, you can see the recent performance of ProtecTIER (see Figure 9-30 on page 248).




Figure 9-30 HyperFactor ratio over time tab, Repository window in ProtecTIER Manager

This section shows the deduplication ratio through a line graph. The ratio is a measure of the amount of data received by the ProtecTIER system from the tape processing application to the amount of data written by the ProtecTIER system out to the disk repository over the previous few hours, summarized for each hour.

9.2 Monitoring the ProtecTIER VTL service

In this section, we describe monitoring of the ProtecTIER virtual tape library.

9.2.1 The Library window

You can access the Library window by clicking the library name in the Services pane on the left side of any window (see Figure 9-31 on page 249). The General tab of the Library window appears.

All of the tabs in the Library window display a status line across the top of the page beneath the menu and tool bars that shows the name of the library, its status, the system it is defined to, and the repository it is attached to.

All of the tabs except the General tab have an area beneath this status line containing page navigation options for when multiple pages are displayed. On this tab, you can also sort the display according to any of the columns by clicking the column heading. Note, though, that this only sorts the data contained on the current page, not the entire data set.

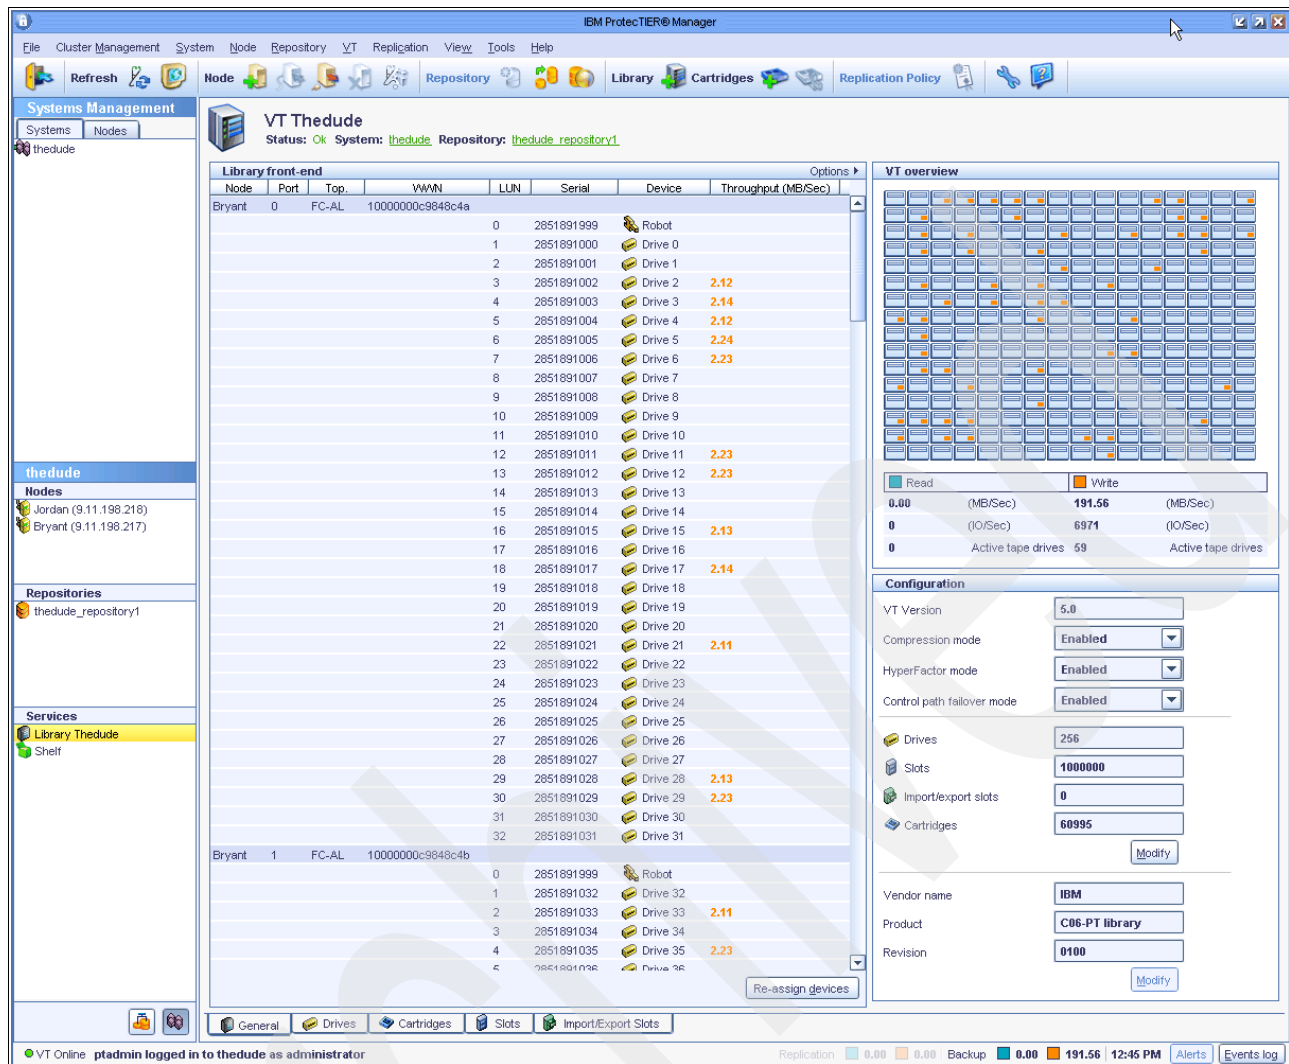


Figure 9-31 General tab, Library window in ProtecTIER Manager

There are several tabs along the bottom of the Library window. These are discussed in detail in the next few sections.

The General tab

The General tab gives a summary view of the selected library (see Figure 9-31). There are four sections in this window.

The Library front-end section

The Library front-end section shows the virtual robot and tape drives for the selected library, sorted by node, then by the Linux LUN number (see Figure 9-32 on page 250).

Node	Port	Top	WWN	LUN	Serial	Device	Throughput (MB/Sec)
Bryant	0	FC-AL	1000000c9848c4a				
				0	2851891999	Robot	
				1	2851891000	Drive 0	
				2	2851891001	Drive 1	
				3	2851891002	Drive 2	2.13
				4	2851891003	Drive 3	2.13
				5	2851891004	Drive 4	2.11
				6	2851891005	Drive 5	2.24
				7	2851891006	Drive 6	2.24
				8	2851891007	Drive 7	
				9	2851891008	Drive 8	
				10	2851891009	Drive 9	
				11	2851891010	Drive 10	
				12	2851891011	Drive 11	2.23
				13	2851891012	Drive 12	2.24
				14	2851891013	Drive 13	
				15	2851891014	Drive 14	
				16	2851891015	Drive 15	2.14
				17	2851891016	Drive 16	
				18	2851891017	Drive 17	2.14
				19	2851891018	Drive 18	
				20	2851891019	Drive 19	
				21	2851891020	Drive 20	
				22	2851891021	Drive 21	2.10
				23	2851891022	Drive 22	
				24	2851891023	Drive 23	
				25	2851891024	Drive 24	
				26	2851891025	Drive 25	
				27	2851891026	Drive 26	
				28	2851891027	Drive 27	
				29	2851891028	Drive 28	2.11
				30	2851891029	Drive 29	2.23
				31	2851891030	Drive 30	
				32	2851891031	Drive 31	
Bryant	1	FC-AL	1000000c9848c4b				
				0	2851891999	Robot	
				1	2851891032	Drive 32	
				2	2851891033	Drive 33	2.10
				3	2851891034	Drive 34	
				4	2851891035	Drive 35	2.24

Figure 9-32 Library front-end section, General tab, Library window in ProtecTIER Manager

There are several columns of information. Some columns apply to the node in general and others apply to individual drives. They are explained in Table 9-3.

Table 9-3 Column definitions for Library front-end section, General tab, Library window in ProtecTIER Manager

Column	Definition
Node	The node on which the virtual device is assigned.
Port	The port within the node on which the virtual device is assigned.
Top	The Fibre Channel topology of the port. Possible values are: <ul style="list-style-type: none"> ► P2P (Point-to-point) ► FC-AL (Fibre Channel-arbitrated loop) ► DOWN (There is no Fibre Channel connection.)
WWN	The World Wide Name of the port.
LUN	The logical unit number of the robot or tape drive relative to the port.
Device	The name of the robot or tape drive.
Throughput	The rate of data transfer across the device.

If a drive is currently active, the Throughput (MB/Sec) column shows a value in one of two colors:

- ▶ Orange for a WRITE operation
- ▶ Blue for a READ operation

This section also contains the Reassign Devices button. You cannot use this button to change the assignments of the devices to the Front End ports across both nodes because this is a fixed configuration.

The VT overview section

The VT overview section shows a summary of current drive activity and updates dynamically. The drives are represented graphically and display an orange or blue icon if they are being used (the colors are the same as above). Hovering your cursor over the graphic displays the drive number, the current read/write rate of the drive in megabytes (MBs) per second, and the percentage of time that the tape drive is idle during tape processing operations due to low tape processing application data transfer rates (see Figure 9-33).




Figure 9-33 VT overview section, General tab, Library window in ProtecTIER Manager

Summary totals are listed for both Read and Write operations and for total active tape drives in the selected library.

The Configuration section

The Configuration section (Figure 9-34 on page 252) displays the current number of drives, slots, import/export slots, and cartridges in the selected library.

Configuration	
VT Version	5.0
Compression mode	Enabled
HyperFactor mode	Enabled
Control path failover mode	Enabled
Drives	256
Slots	1000000
Import/export slots	0
Cartridges	60995
Modify	
Vendor name	IBM
Product	C06-PT library
Revision	0100
Modify	

Figure 9-34 Configuration section, General tab, Library window in ProtecTIER Manager

The section also contains the Modify button. You cannot use this button to change the dimensions of the selected library because of its fixed layout.

Note: Because the TS7680 has a fixed layout and is optimally adjusted to work with the System z hosts, you cannot modify these settings.

Also the Vendor name, the Product name, and the Revision for this virtual library are shown in this window (see Figure 9-34).

The Drives tab

The Drives tab displays detailed information about the virtual tape drives in the selected library (see Figure 9-35 on page 253).

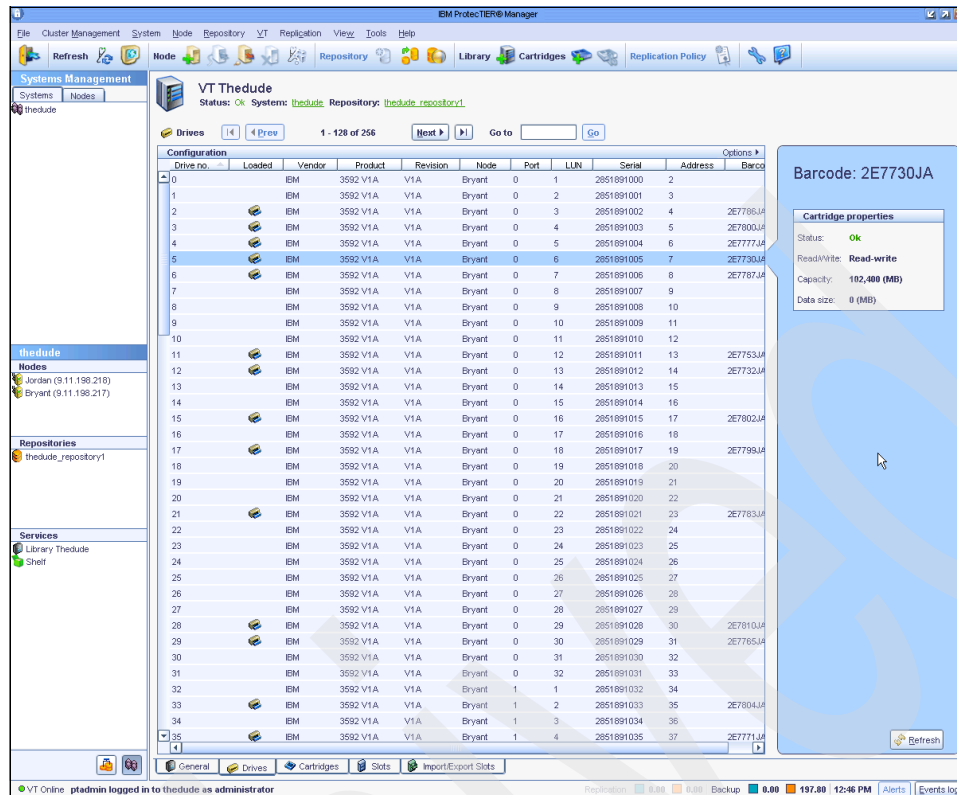


Figure 9-35 Drives tab, Library window in ProtecTIER Manager

The column values are explained in Table 9-4.

Table 9-4 Column definitions for Drives tab, Library window in ProtecTIER Manager

Column	Definition
Drive No.	The drive number in ProtecTIER.
Loaded	Whether the drive is loaded with a virtual cartridge. If the drive is loaded with a cartridge, an icon is displayed.
Vendor Name	The vendor whose product the virtual drive emulates.
Product	The product name for the product that the virtual drive emulates.
Revision	The revision number for the product that the virtual drive emulates.
Node	The node to which the drive is assigned.
Port	The port on the node to which the drive is assigned.
LUN	The drive's logical unit number relative to the port. Note this is for ProtecTIER only.
Serial	The drive's serial number.
Address	The drive's address within the library.
Barcode	If the drive is loaded with a cartridge, this column displays the cartridge's barcode.

The Cartridges tab

The Cartridges tab displays detailed information about the cartridges in the selected library (see Figure 9-36).

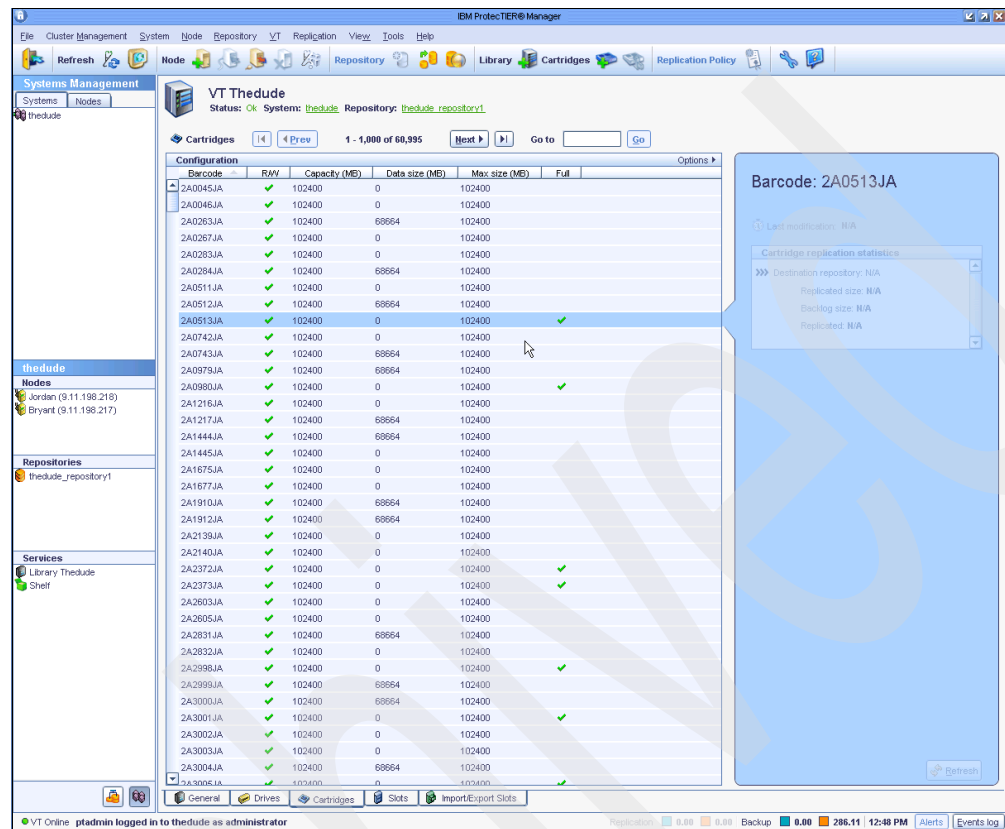


Figure 9-36 Cartridges tab, Library window in ProtecTIER Manager

The column values are explained in Table 9-5.

Table 9-5 Column definitions for Cartridges tab, Library window in ProtecTIER Manager

Column	Definition
Barcode	The cartridge's barcode.
Write enabled	Whether the cartridge is write enabled. If it is, a green tick icon is displayed; otherwise, a red cross icon is displayed.
Capacity (MB)	The cartridge's estimated data capacity in megabytes. This value varies over time depending on the HyperFactor ratio and the number of cartridges configured in the system.
Data size (MB)	The amount of nominal data, in megabytes, currently stored on the cartridge.
Max size (MB)	The maximum (fixed) amount of nominal data at which the ProtecTIER will announce Early Warning for this cartridge to the tape processing application.
Full	Indicates whether the cartridge has reached the Early Warning threshold. If so, the cartridge is regarded as full and a green tick icon is displayed.

The Slots tab

The Slots tab displays detailed information about the slots in the selected library (see Figure 9-37), some of which is repeated from the Drives tab. The information includes slot number, element address, and cartridge barcode, estimated capacity, and data size (amount of nominal data stored) if the slot contains a cartridge.

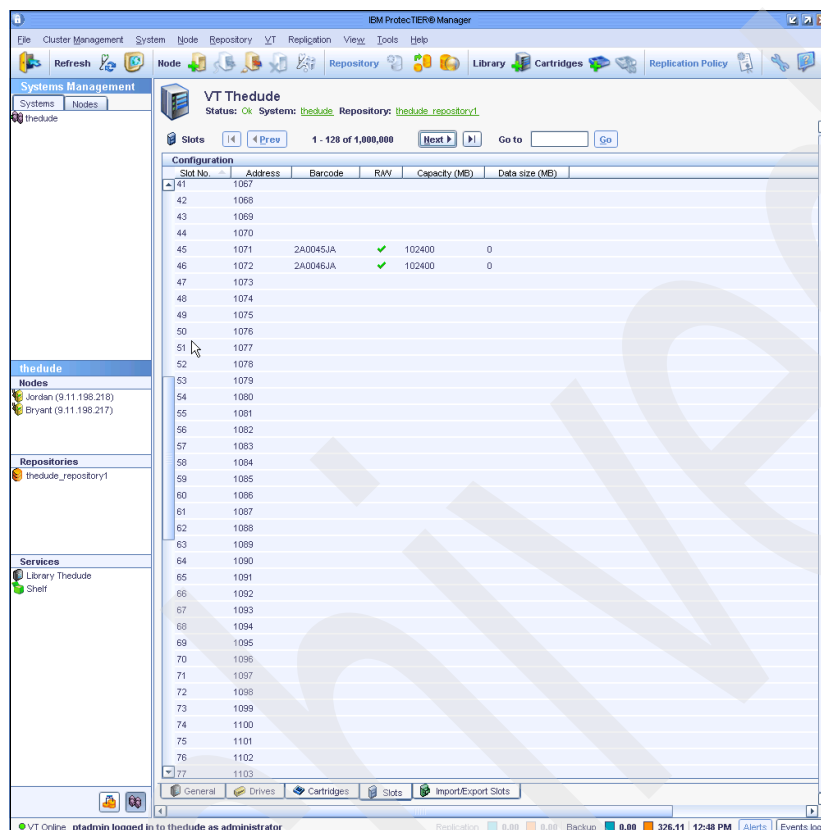


Figure 9-37 Slots tab, Library window in ProtecTIER Manager

9.3 Reporting on ProtecTIER activity

In this section, we discuss the two types of reports that can be generated by you for ProtecTIER activities. Both reports produce spreadsheet format files that can be viewed and analyzed.

9.3.1 The analyze_sessions utility

You can gather and view ProtecTIER tape processing statistics to see how the system is performing. These statistics show throughput, compression, and HyperFactor rates for tape processing data sent to the system. The ProtecTIER provides a utility called *analyze_sessions* to perform this data collection task.

The *analyze_sessions* utility analyzes and summarizes the contents of the current, cumulative compressor log. A compressor log is generated each time the Linux vtfd daemon is started and accumulates entries until the vtfd daemon is shut down. Repeated runs of the utility between restarts will report mostly identical data each time, plus any new session data

since the last execution of the utility. When the `analyze_sessions` utility is run after a restart of the ProtecTIER software, it only reports on new session data since the restart.

The recommended use of the `analyze_sessions` utility is for the deliberate testing of a specific workload or type of data. The best practice would be isolate a desired workload to be the only data being sent to the ProtecTIER for a certain period with a five minute break on either side of it, to separate it in a single session from any other data. You would then run the `analyze_sessions` utility and examine the output to determine how well the data might be suited to data deduplication processes.

Running the utility

To run this command and gather the data, you need to execute the following steps:

1. Log on to one of the TS7680 Linux servers as root.
2. Change to the `/opt/dtc/app/utils/` directory.
3. Run `./analyze_sessions`.

This will create an output file in comma-separated-values (CSV) format in the following directory:

```
/pt_work/<server name>-<date>-<time>.csv
```

See Figure 9-38 for an example of how to run the `analyze_sessions` utility on the Linux server and what output will be returned.

```
login as: root
root@9.11.200.150's password:
Last login: Wed Aug 06 13:00:28 2008 from 10.0.0.52
[root@speedway ~]# cd /opt/dtc/app/utils/
[root@speedway utils]# ./analyze_sessions

analyze_sessions: this program analyzes the logs on the current
ProtecTIER server.  Its output is the change rate between each
session and data that was available in the repository at the time.

This program and its output contain proprietary information of
International Business Machines Corporation.

(c) Copyright International Business Machines Corporation 2008. All rights
reserved.

Cutoff date: 2008-5-6 13:00:50
Processing file: 2/2
Read 1MB, throughput: 7.35MB/s
Output file is /pt_work/speedway-2008-08-06-13-00-50.csv
[root@speedway utils]#
```

Figure 9-38 Running the `analyze_sessions` utility on ProtecTIER Linux system

Examining the output

You can either view the generated CSV format file on the TS7680 Linux server or you can use **ftp** commands or similar utilities to download the file to a Windows workstation to open it. When you open the CSV file in Microsoft Excel or an equivalent product, you will see data similar to that shown in Figure 9-39.

Name	Total data (TB)	Total data (KB)	System change rate	Factoring ratio	start time	end time
Grand totals						
all	0.74	791260000	18.67%	5.35738	29/07/2008 18:22	5/08/2008 19:44
By session (summary)						
2008-7-29 18:22:54 to 2008-7-29 18:22:54	0.00	826810	100%	1	29/07/2008 18:22	29/07/2008 18:22
2008-7-29 18:29:41 to 2008-7-29 18:34:14	0.00	489220	35.34%	2.82926	29/07/2008 18:29	29/07/2008 18:34
2008-8-4 11:56:26 to 2008-8-4 12:13:09	0.01	8388610	99.82%	1.00183	4/08/2008 11:56	4/08/2008 12:13
2008-8-4 12:26:03 to 2008-8-4 12:41:38	0.01	9437180	10.43%	9.58874	4/08/2008 12:26	4/08/2008 12:41
2008-8-4 12:52:20 to 2008-8-4 13:07:06	0.01	9259260	0.13%	760.404	4/08/2008 12:52	4/08/2008 13:07
2008-8-4 13:14:42 to 2008-8-4 13:53:30	0.02	16777200	38.56%	2.59367	4/08/2008 13:14	4/08/2008 13:53
2008-8-4 14:02:13 to 2008-8-4 14:19:06	0.04	40894500	100%	1	4/08/2008 14:02	4/08/2008 14:19
2008-8-4 14:39:15 to 2008-8-4 14:39:15	0.00	340736	100%	1	4/08/2008 14:39	4/08/2008 14:39
2008-8-4 15:14:05 to 2008-8-4 16:17:16	0.21	227044000	25.99%	3.84737	4/08/2008 15:14	4/08/2008 16:17
2008-8-4 16:22:24 to 2008-8-4 18:04:55	0.12	132392000	13.30%	7.51867	4/08/2008 16:22	4/08/2008 18:04

Figure 9-39 Example of the analyze_sessions utility output

In ProtecTIER terms, a “session” is defined as a sustained period of data throughput (for example, data sent to the PT server from a tape processing application) preceded by and followed by an idle period of at least five minutes.

In the data, there is a Grand Totals row summarizing all the sessions in the output file, followed by detail rows for each individual session. There are several columns for each row, containing statistical information. See Table 9-6 for an explanation of the fields.

Table 9-6 Recent tape processing session statistics definitions

Statistic	Description
Total data (TB)	Total amount of data backed up during the session, in terabytes.
Total data (KB)	Total amount of data backed up during the session, in kilobytes.
System change rate	Percentage of data in the tape processing session recognized as changed relative to the previous tape processing session.
Factoring ratio	The ratio of the quantity of actual backed-up data over the total amount of physical data.
Start time	Start time of the tape processing session.
End time	End time of the tape processing session.

9.4 Monitoring ProtecTIER through ptmon

This section describes how to query the ProtecTIER system through the CLI using ptmon. The ptmon is loaded during the installation of ProtecTIER software and is located in the `/opt/dtc/app/sbin/` directory. Make sure you are in this directory when running ptmon commands from the command prompt.

Usage

Use the **ptmon** command to query the system via the command line interface (CLI), and to receive various statistics about the ProtecTIER system. This information can provide valuable insight to the administrator about the performance, capacity, configuration and operation of the system, and can be accessed by other management applications.

You can print the usage information message by running **ptmon** with no arguments. The usage information message will also be printed when the arguments list is malformed.

The ptmon command is issued from the command line as follows:

```
./ptmon <-parameter> <variable>
```

Parameters

-repository

Displays repository statistics.

-libs

Displays information about the unique ID of all configured libraries on the repository.

-libinfo -libid *library id*

Displays statistics on the configured tape drives of a specific library.

-ncarts -libid *library id*

Displays statistics on the number of configured cartridges of a specific library.

-cartsinfo -libid *library id* -from *first cartridge index* -count *number of cartridges*

Displays statistics for each of the cartridges in a specified range.

-node_vtl_statistics -hours *hours*

Displays the statistics history on the local host for the last <hours> hours.

-version

Displays the version of the build of the local ProtecTIER installation. *pt-version* is the major version release.

ptmon responses

ptmon responses always appear in a well structured XML document printed.

The XML always has the same root element: “response” with attribute “status” which may have the values “success” or “failed”. The response element displays the reply as a child element on success and, in this case, returns an exit code of zero to the shell. If status is “failed” the response element has a single child element: “error” with attributes “description” (a human readable string) and a non-zero “code” value. The same value is returned to the shell. The code value is also returned as the ptmon exit code and enables grouping of errors to general types (RPC error, argument error, and so on).

Example - on success

Example 9-1 ptmon on success response output

```
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<response status="success">
<libraries-list host-name="somehostname" time="1235485682"
time-string="24-Feb-2009 16:28">
<library name="somehostname_vt" unique-id="2199023256193"/>
</libraries-list>
</response>
```

Example - on error

Example 9-2 ptmon on error output

```
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<response status="failed">
<error code="2" description="Error: failed to execute
GetStatisticsSamples"/>
</response>
```

ptmon -repository

This command generates the repository statistics shown in Table 9-7.

Table 9-7 Repository command statistics

Field	Description
current-factoring-ratio	Current factoring ratio in the repository, which is the nominal data divided by the physical data.
repository-free-physical-storage-mb	Amount of unused space in the repository in MB.
repository-total-nominal-storage-mb	Amount of nominal data in the repository in MB.
repository-total-physical-storage-mb	Amount of physical storage configured for the repository in MB.
repository-used-physical-storage-mb	Amount of physical storage used by the repository in MB.
repository-allocable-physical-storage-mb	Amount of physical storage available for the repository in MB.

An example of this command is seen in Example 9-3.

Example 9-3 ptmon -repository output

```
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<response status="success">
<repository-statistics current-factoring-ratio="1.00" host-name="beast"
repositoryallocablephysical-
storage-mb="49152" repository-free-physical-storage-mb="49152"
repositorytotalnominal-
storage-mb="0" repository-total-physical-storage-mb="49152"
repository-usedphysicalstorage-
mb="0" time="1236210447" time-string="04-Mar-2009 18:47"/>
</response>
```

ptmon -libs

This command displays information about the unique ID of all configured libraries on the repository. This library-unique ID can later be used for receiving detailed information about each of the configured libraries. Execution of the command is shown in Example 9-4.

Example 9-4 ptmon -libs output

```
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<response status="success">
<libraries-list host-name="beast" time="1236210637"
time-string="04-Mar-2009 18:50">
<library name="2" unique-id="2199023255719"/>
<library name="20" unique-id="2199023255824"/>
</libraries-list>
</response>
```

ptmon -libinfo -libid <library id>

This command displays statistics about the configured tape drives of a specific library (see Example 9-5 and Example 9-8).

Table 9-8 Table: Statistics on the configured tape drives of a specific library

Field	Description
num-total-drives	Amount of configured tape drives
library-name	Name of the library
libid	Unique ID for the library

Example:

Example 9-5 ptmon -libinfo -libid <library id> output

```
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<response status="success">
<library-info host-name="beast" library-name="2"
library-unique-id="2199023255719" numtotaldrives="
12" time="1236210783" time-string="04-Mar-2009 18:53">
<active-drives/>
</library-info>
</response>
```

ptmon -ncarts -libid <library id>

This command displays statistics on the number of configured cartridges of a specific library (see Example 9-9 and Example 9-6):.

Table 9-9 Table: Statistics on the number of configured cartridges of a specific library

Header	Header
libid	Library unique ID taken from ptmon -libs output.

Example 9-6 ptmon -ncarts -libid <library id> output

```
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<response status="success">
```

```
<library host-name="beast" num-configured-cartridges="5"
time="1236210795" time-string="04-
Mar-2009 18:53" unique-id="2199023255719"/>
</response>
```

ptmon -cartinfo -libid <library id> -from <first cartridge index> -count <number of cartridges>

This command displays statistics for each of the cartridges in a specified range (see Table 9-10).

- ▶ Cartridge barcode
- ▶ Indication of whether the cartridge is full
- ▶ Total amount of nominal data currently on the cartridge

Table 9-10 Statistics of the cartridges in a specified range

Field	Description
libid	Library-unique ID taken from -libs output.
from	Number of cartridges before the first printed cartridges in the sorted list of cartridges (by barcode).
count	Maximum number of cartridges in the output. This list may be shorter if the sorted list of cartridges is exhausted before printing <number of cartridges> cartridges.

The following is the command syntax; Example 9-7 shows the execution example.

```
ptmon -cartinfo -libid <library id> -from <first cartridge index> -count
<number of cartridges> output
```

Example 9-7 Output

```
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<response status="success">
<cartridges-information host-name="beast" library-unique-id="2199023255719"
time="1236210879" time-string="04-Mar-2009 18:54">
<cartridge barcode="000000" early-warning-reported="false"
nominal-size-mb="0"/>
<cartridge barcode="000001" early-warning-reported="false"
nominal-size-mb="0"/>
<cartridge barcode="000002" early-warning-reported="false"
nominal-size-mb="0"/>
<cartridge barcode="000003" early-warning-reported="false"
nominal-size-mb="0"/>
<cartridge barcode="000004" early-warning-reported="false"
nominal-size-mb="0"/>
</cartridges-information>
</response>
```

ptmon -node_vtl_statistics -hours <hours>

This command (see Table 9-11) displays the statistics history on the local host for the last <hours> hours.

Table 9-11 Statistics history on the local host for the last hours

Field	Description
hours	Number of hours of statistics to be included in the output. By default, the output includes 4 chronicles (for example, statistic records) per hour (if the uptime is greater than or equal to the number of hours).

Example 9-8 shows execution of the command.

Example 9-8 *ptmon -node_vtl_statistics -hours <hours> output*

```
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<response status="success">
<statistics num-returned="4">
<chronicle n-samples-per-poll="30" sample-date="04-03-2009 18:00:32"
sample-timestamps="
30">
<vtl-overall num-active-tapes="0" num-reading-tapes="0"
num-writing-tapes="0" readio.
average-bytes="0" read-io.max-bytes="0" read-io.min-bytes="0"
valid="true" writeio.
average-bytes="0" write-io.max-bytes="0" write-io.min-bytes="0"/>
</chronicle>
<chronicle n-samples-per-poll="30" sample-date="04-03-2009 18:15:33"
sample-timestamps="
30">
<vtl-overall num-active-tapes="0" num-reading-tapes="0"
num-writing-tapes="0" readio.
average-bytes="0" read-io.max-bytes="0" read-io.min-bytes="0"
valid="true" writeio.
average-bytes="0" write-io.max-bytes="0" write-io.min-bytes="0"/>
</chronicle>
<chronicle n-samples-per-poll="30" sample-date="04-03-2009 18:30:34"
sample-timestamps="
30">
<vtl-overall num-active-tapes="0" num-reading-tapes="0"
num-writing-tapes="0" readio.
average-bytes="0" read-io.max-bytes="0" read-io.min-bytes="0"
valid="true" writeio.
average-bytes="0" write-io.max-bytes="0" write-io.min-bytes="0"/>
</chronicle>
<chronicle n-samples-per-poll="30" sample-date="04-03-2009 18:45:35"
sample-timestamps="
30">
<vtl-overall num-active-tapes="0" num-reading-tapes="0"
num-writing-tapes="0" readio.
average-bytes="0" read-io.max-bytes="0" read-io.min-bytes="0"
valid="true" writeio.
average-bytes="0" write-io.max-bytes="0" write-io.min-bytes="0"/>
</chronicle>
</statistics>
</response>
```

ptmon -version

This command (see Example 9-9) displays the version of the build of the local ProtecTIER installation. pt-version is the major version release.

Example 9-9 ptmon -version output

```
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<response status="success">
<version-info>
<version build="0.0.0" pt-version="2.2.0.TEST1"/>
</version-info>
</response>
```

9.5 Monitoring the virtual library through the z/OS host

In this section we describe some basic Data Facility Storage Management System (DFSMS) commands to gather the key information about the TS7680 virtual tape library from the host point of view. In addition, we provide operation-critical messages that can occur on your z/OS host system.

9.5.1 Display SMS Library

To display the status of the TS7680 virtual tape library use the DFSMS Display SMS Library command shown in Figure 9-40. The output of this command shows the Automated Library (AL) TS7680 as the Device Type 3958-DE2. You can also see the fixed configuration of 1,000,000 storage slots. On the bottom of this figure you see two messages occurring because of the field CACHE PERCENTAGE USED: 099 state for this example: *Limited Cache Free Space - Warning State* and *Out of Cache Resource - Critical State*. See 9.5.4, “z/OS host messages” on page 264 for more information about these messages.

13.45.24 d sms,lib(atlca030),detail										
13.45.24 STC00016 CBR1110I OAM library status: 869										
TAPE	LIB	DEVICE	TOT	ONL	AVL	TOTAL	EMPTY	SCRATCH	ON	OP
LIBRARY	TYP	TYPE	DRV	DRV	DRV	SLOTS	SLOTS	VOLS		
ATLCA030	AL	3958-DE2	16	1	1	1000000	988176	0	Y	Y

MEDIA	SCRATCH		SCRATCH		SCRATCH					
TYPE	COUNT		THRESHOLD		CATEGORY					
MEDIA5	3		2		0405					

LIBRARY ID: CA030										
CACHE PERCENTAGE USED: 099										
OPERATIONAL STATE: AUTOMATED										
ERROR CATEGORY SCRATCH COUNT: 0										

Limited Cache Free Space - Warning State.										
Out of Cache Resources - Critical State.										

Figure 9-40 DFSMS Display SMS Library command

9.5.2 Display SMS Volume

To get information about a specific logical volume, use the DFSMS command Display SMS Volume. Figure 9-41 shows this command for the logical volume fus002. The output shows information on the media type used and the recording technology that are always used for the

TS7680 virtual tape library volumes. The media type is MEDIA5 and the recording technology is EFMT1.

```

08.41.36          d sms,vol(fus002)
08.41.36 STC00015  CBR1180I OAM tape volume status: 965
VOLUME  MEDIA    STORAGE  LIBRARY  USE  W  C  SOFTWARE  LIBRARY
        TYPE      GROUP   NAME    ATR  P  P  ERR  STAT  CATEGORY
FUS002  MEDIA5    SGATL   ATLCA030 P    N    NOERROR  NOTAVAIL
-----
RECORDING TECH:    EFMT1                COMPACTION:      NO
SPECIAL ATTRIBUTE: NONE                ENTER/EJECT DATE: 2009-07-31
CREATION DATE:    2009-05-13           EXPIRATION DATE:
LAST MOUNTED DATE: 2009-05-27         LAST WRITTEN DATE: 2009-05-27
SHELF LOCATION:
OWNER:
-----

```

Figure 9-41 DFSMS Display SMS Volume command

9.5.3 Library Display Drive

If you want to show information about a specific virtual TS7680 tape drive, you can use the DFSMS command Library Display Drive. Figure 9-42 shows this command for drive number 1bd0. The TS7680 virtual tape drives are all Device Type 3592-J1A, which is shown here as 3592-J.

```

08.46.51          li dd,1bd0
08.46.51          CBR1220I Tape drive status: 968
DRIVE  DEVICE    LIBRARY  ON  OFFREASN  LM  ICL  ICL  MOUNT
NUM    TYPE      NAME    LI  OP  PT  AV  CATEGRY  LOAD  VOLUME
1BD0   3592-J    CA030   N   N   Y   Y   -  --N/A--  -

```

Figure 9-42 DFSMS Library Display Drive command

Note: For more information about how to gather information about the TS7680 from the z/OS host point of view, refer to *z/OS DFSMS Object Access Method Planning, Installation and Storage Administration Guide for Tape Libraries*.

9.5.4 z/OS host messages

In this section we give you four examples of messages that can occur in critical circumstances when the back-end storage of the TS7680 tape library runs out of space.

CBR3792E

CBR3792E Library *library-name* has entered the limited cache free space warning state.

Explanation: The available cache in library *library-name* has entered the limited cache resource warning state. When the TS7680 determines that the amount of available cache space is less than 3 terabytes, the library enters this state.

System Action: Mounts and host I/O transfers continue to be accepted.

This is a warning state and indicates that the TS7680 might soon run out of cache resources unless older data is removed from the TS7680. This state is left when the TS7680 determines that the amount of available cache space is greater than 3.5 terabytes. Message CBR3793I is

issued when the library has left this state. As appropriate, copy or remove (return to scratch) older data from the library.

CBR3793I

CBR3793I Library *library-name* has left the limited cache free space warning state.

Explanation: The available cache in library *library-name* has left the limited cache resource warning state.

When the TS7680 determines that the amount of available cache space is greater than 3.5 terabytes, the library exits this state.

CBR3794A

CBR3794A Library *library-name* has entered the out of cache resources critical state.

Explanation: The available cache in library *library-name* has entered the out-of-cache resource critical state.

When the TS7680 determines that the available cache space reaches a critical level, the library enters this state. This state takes into account the number of jobs (devices) that are writing data to ensure that currently running jobs can complete. At a minimum, this state can also be entered when the amount of free space falls below 500 gigabytes (may be entered sooner based on the number of jobs that are running and the projected cache usage). As the available cache space continues to decline, the TS7680 may also reach a point where it needs to throttle the write activity of currently running jobs, resulting in job delays.

System Action: Mount operations that have been accepted before entering this state complete and volumes currently mounted can continue to perform host I/O operations. Any scratch mount operations received in this state are failed by the library. Any specific mount operations received in this state are accepted by the library; however, any write operation to the volume is failed.

While in this state, the TS7680 can be configured (at install time) to automatically delete data associated with scratch volumes that are in the "grace period". Volumes with the shortest time remaining in the grace period will be deleted first. This state is left when the TS7680 determines that the amount of available cache space can accommodate the currently running write jobs plus an additional number of jobs. Message CBR3795I is issued when the library has left this state. When this state is left, CBR4196D provides an opportunity to retry failing mount requests. To make cache space available, copy or remove (return to scratch) older data from the library as appropriate, and respond to any outstanding CBR4196D messages. In addition to this, the TS7680 will also attempt to free up cache space by deleting data associated with scratch volumes (refer to the System Action for additional detail).

CBR3795I

CBR3795I Library *library-name* has left the out of cache resources critical state.

Explanation: The available cache in library *library-name* has left the out-of-cache resource critical state.

When the TS7680 determines that the amount of available cache space can accommodate the currently running write jobs plus an additional number of jobs, the library exits this state.

Note: For a complete list of available z/OS messages, refer to the *z/OS message help website* at:

<http://www-03.ibm.com/systems/z/os/zos/bkserv/lookat/index.html>

9.6 Other user notifications

Besides the z/OS console messages you may also get notifications in the following ways:

- ▶ PT Manager GUI “Used space” physical disk space consumed
- ▶ PT Manager GUI “Nominal data size” client data

9.7 Enterprise controller tracing, logs and tools

The Enterprise controller also has some utilities and functions available for monitoring and troubleshooting:

- ▶ AIX trace
- ▶ AIX Error Report
- ▶ Library Commands by Device
- ▶ LIE Traces and Logs
- ▶ RASUTIL
 - FICON Interface Status
 - FICON Read/Write Performance Monitoring
 - Virtual Tape Device Status
- ▶ IRMM Logs
- ▶ List ALL C06 control unit processes
- ▶ SCSI Configuration Expert (SCE)
 - Configure the tape and library devices (C06 must be offline)
 - Display the tape and library configuration

Disaster recovery and failover scenarios

Disaster recovery is the process of recovering production site data in a remote location which was the target for the replication operation from the primary site prior to the disaster occurrence.

In case of a disaster or a situation where the production (or primary) site has gone offline, the remote, disaster recovery site, can take the place of the production site until the primary site comes back online. When the primary site comes back online, previously replicated as well as newly created tapes can be moved to the main production site using the fallback process so it can once again become the production/primary site for the user.

Note: In the TS7680 environment the initial release will provide the native High Availability configuration as the built in Disaster Recovery option

This chapter summarizes the different stages in handling a TS7680 disaster scenario.

10.1 Disaster failover to remote site

For temporary or extended outage of production site All steps are run at the DR site Requires unique scratch pool configured at DR host Volumes replicated from the production site are read only at DR site

10.1.1 DR Failover procedure

This is a high level example of the Failover procedure (see Figure 10-1)

1. Enter DR mode using the ProtecTIER Replication Manager GUI
2. Retrieve list of inconsistent volumes
3. Insert additional scratch volumes as required
4. Vary on devices at the DR host

Figure 10-1 Disaster Recovery failover

10.1.2 Disaster Recovery failback

This high-level procedure returns to the production site when operations are restored (see Figure 10-2). Data created at the DR site may be replicated back to the production site through the failback replication policy. Data that was previously created at the production site before DR failover is read only at the DR site and therefore was not modified at the DR site.

DR failback procedure (once production site has been restored)

1. Create and execute the failback replication policy to copy any volumes that were created at the DR site.
2. Multiple failback replication policies may be executed. The bulk of data created during the outage is replicated first while production at the DR site continues. Repeat the failback replication process for volumes created during the first bulk copy.
3. DR host vary off devices.
4. Allow the final failback replication policy to complete.
5. Exit DR mode at the DR site.
6. Production host vary on devices and resume operations.

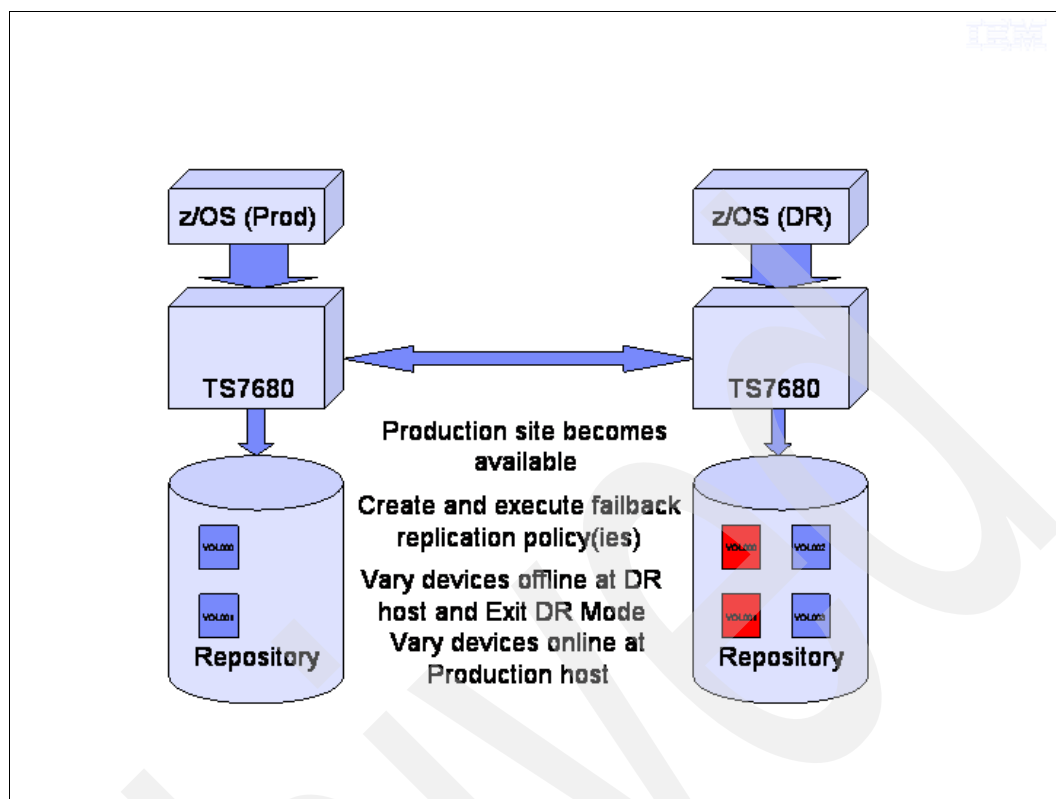


Figure 10-2 Failback process

10.1.3 Disaster Recovery takeover

This topic is for total loss of system at the production site. All steps are run at the DR site (see Figure 10-3). May require adding cartridges to the scratch pool at the DR host (inserted scratch cartridges are not replicated until written). Volumes replicated from the production site are changed to read/write at the DR site.

DR takeover procedure

1. Initiate a takeover while in DR mode using the ProtecTIER Replication Manager GUI.
2. Insert additional scratch volumes as required.
3. Vary on devices at the DR host.

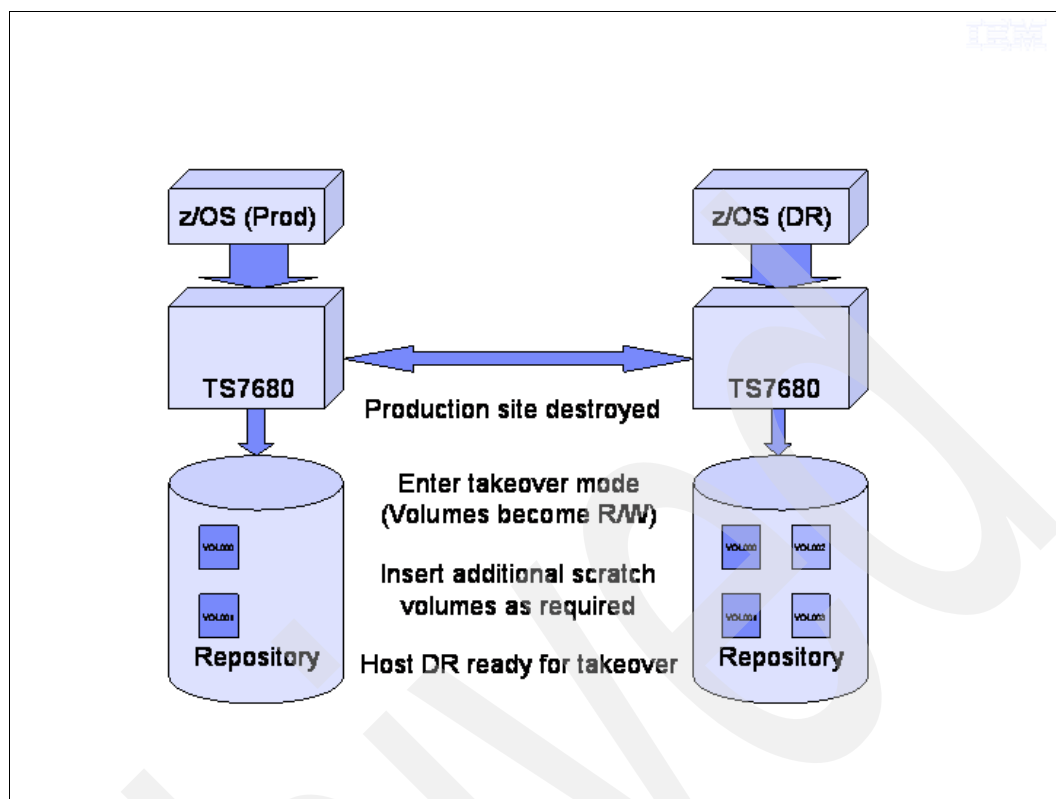


Figure 10-3 Disaster Recovery takeover

10.2 Enterprise controller (3592-C06) contribution to Disaster Recovery

The dual Enterprise controller standard configuration in a TS7680 environment provides disaster recovery in case one controller fails or will not come online.

Note: 128 virtual addresses owned by the failing Enterprise controller will not be available for the time the controller is offline.

The Enterprise controller is configured in manufacturing. All configuration settings are set through the System Management Interface Tool (SMIT) interface. There is no client management interface. A SMIT interface is provided to set a “grace period” in days during which scratched volumes will not be deleted. You may request the System Service Representative (SSR) to change the default value at the time of installation.

A SMIT panel is provided to set the Library ID. This is a unique value assigned by you that is used in z/OS to correlate the hardware library with a software-defined Automated Tape Library (ATL).

In case of a disaster that requires the SSR to rebuild and reconfigure the Enterprise controller in the field, the existing procedures for Enterprise controller are followed with the exception of setting the library type to 3958 (this is the indicator that the Enterprise controller is attached to a TS7680 VTL) and any other changes to the default configuration.

Note: The Enterprise controller does not maintain a backup of the configuration settings. In case of a disaster, all configuration settings must be restored by the SSR.

10.2.1 ProtecTIER Server (3958-DD3) contribution to disaster recovery

In the 3958-DE2 (TS7680), two ProtecTIER servers or nodes are arranged in a cluster.

A two-node system uses two servers in an active-active cluster and enables a sophisticated system with high availability and disaster recovery benefits.

Clustered ProtecTIER servers provide access to the virtual volumes on disk and hardware redundancy in the event of a disaster recovery for a node failure or maintenance.

- ▶ Both ProtecTIER Servers have access to all volumes in the library.
- ▶ Vary off a portion of the host devices to perform concurrent maintenance.
- ▶ Concurrent code load by updating one server at a time.

Note: In case one ProtecTIER server goes down (example: Lower), the 128 addresses corresponding to the Lower Enterprise controller will also not be available until the server is back online.

10.2.2 Failover scenarios

In this topic we show you the most important cases of a TS7680 disaster recovery scenario.

We assume the following configuration names as example; see Figure 10-4.

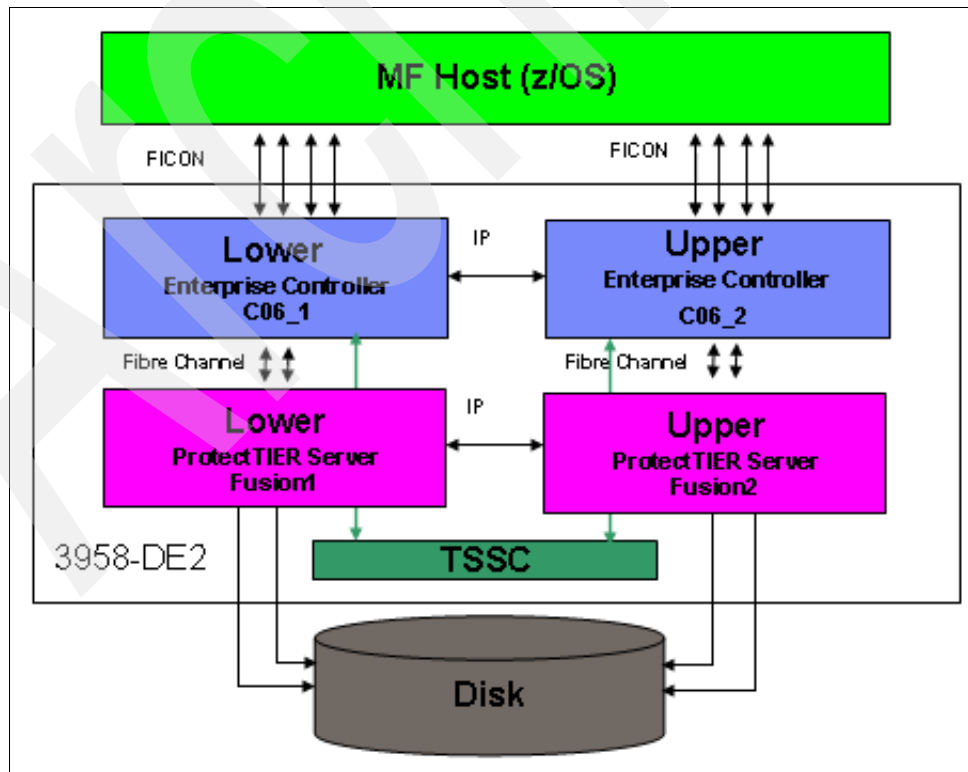


Figure 10-4 Failover scenario

One Enterprise controller down

Let us assume that one Enterprise controller goes either OFFLINE or POWER® DOWN; see Figure 10-5 on page 272.

Note: This scenario applies also in case Enterprise controller maintenance is needed (example: Licensed internal code upgrade).

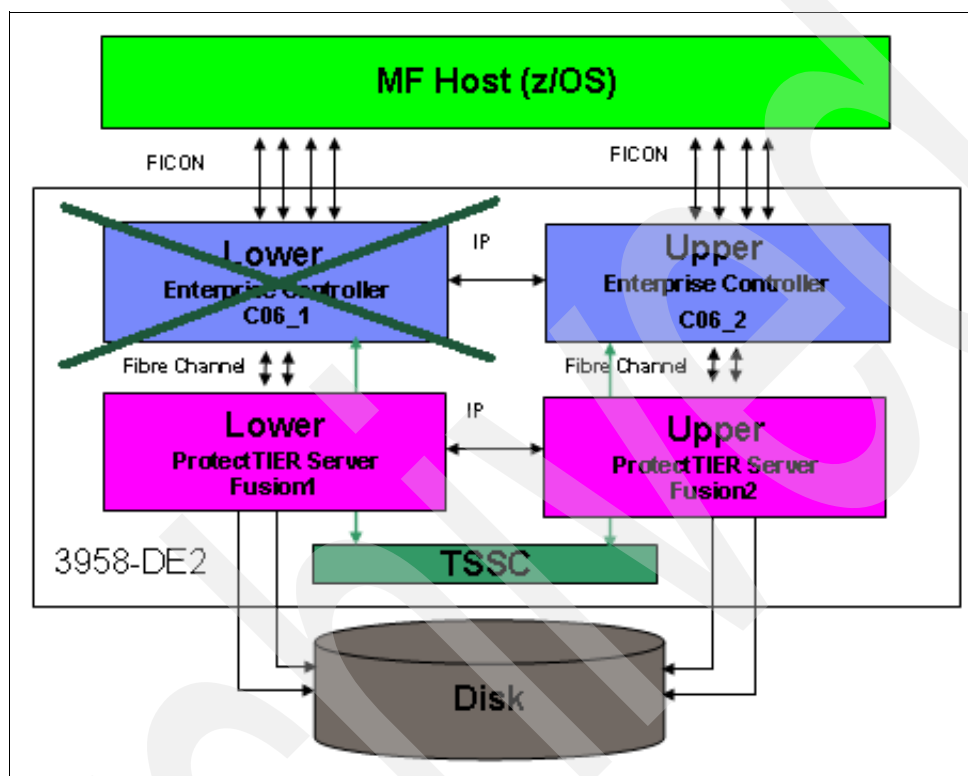


Figure 10-5 Lower Enterprise controller down

As stated before, you will lose 128 virtual addresses (4 virtual control units with 32 virtual tape drives each), but all jobs will continue to run because the volumes are stored on external cache (disk) that is accessed by both ProtecTIER servers. The total throughput will drop because only one Enterprise controller is used at that point.

You will not see any Alert from ProtecTIER Manager; see Figure 10-6 on page 273

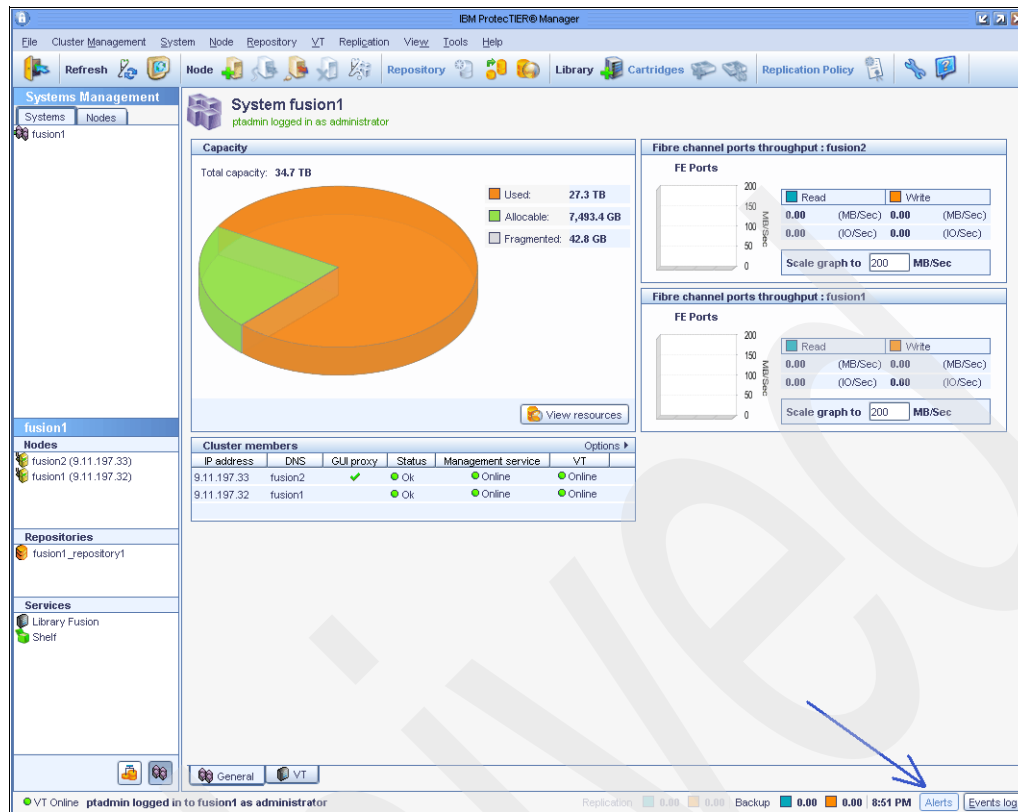


Figure 10-6 PT Manager_no Alert

You can only see Link Speed and Topology under Port Attributes by clicking the corresponding node (**Fusion1**) at the left side of the panel; see Figure 10-7 on page 274.

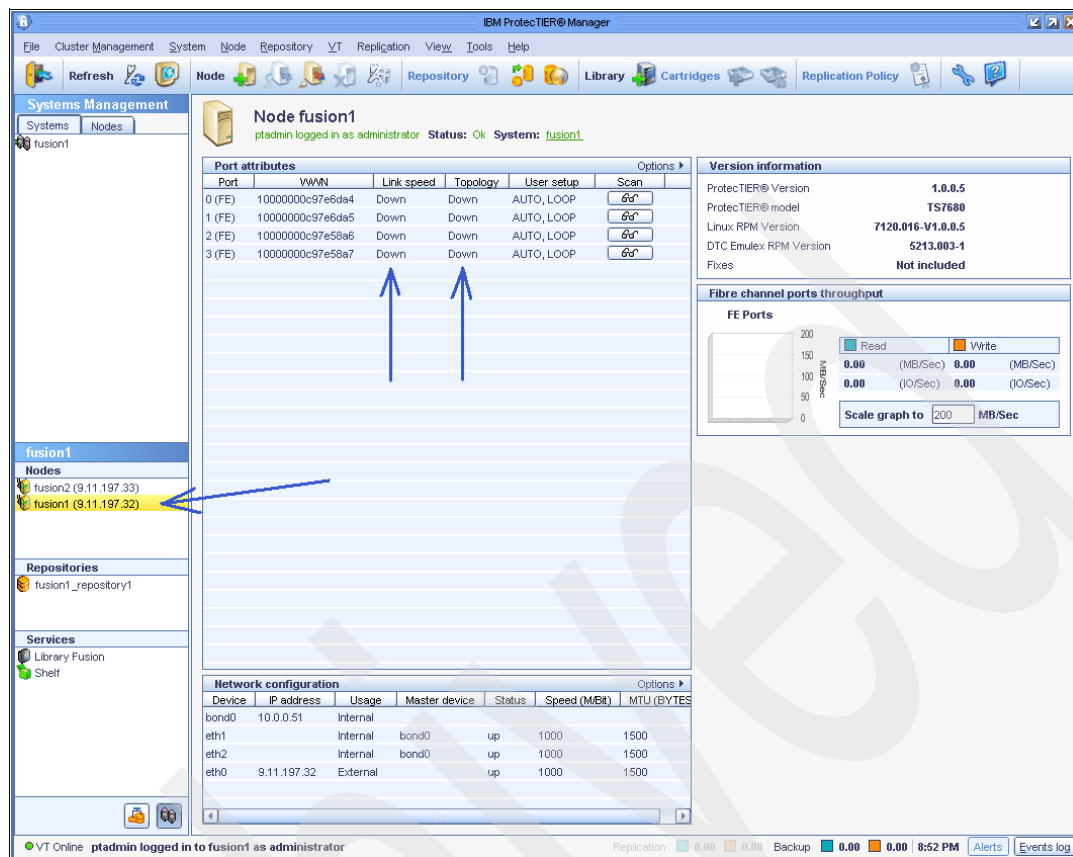


Figure 10-7 PT Manager_Link down

Link Down is also confirmed by clicking **Library Fusion** on the Services panel; see Figure 10-8 on page 275.

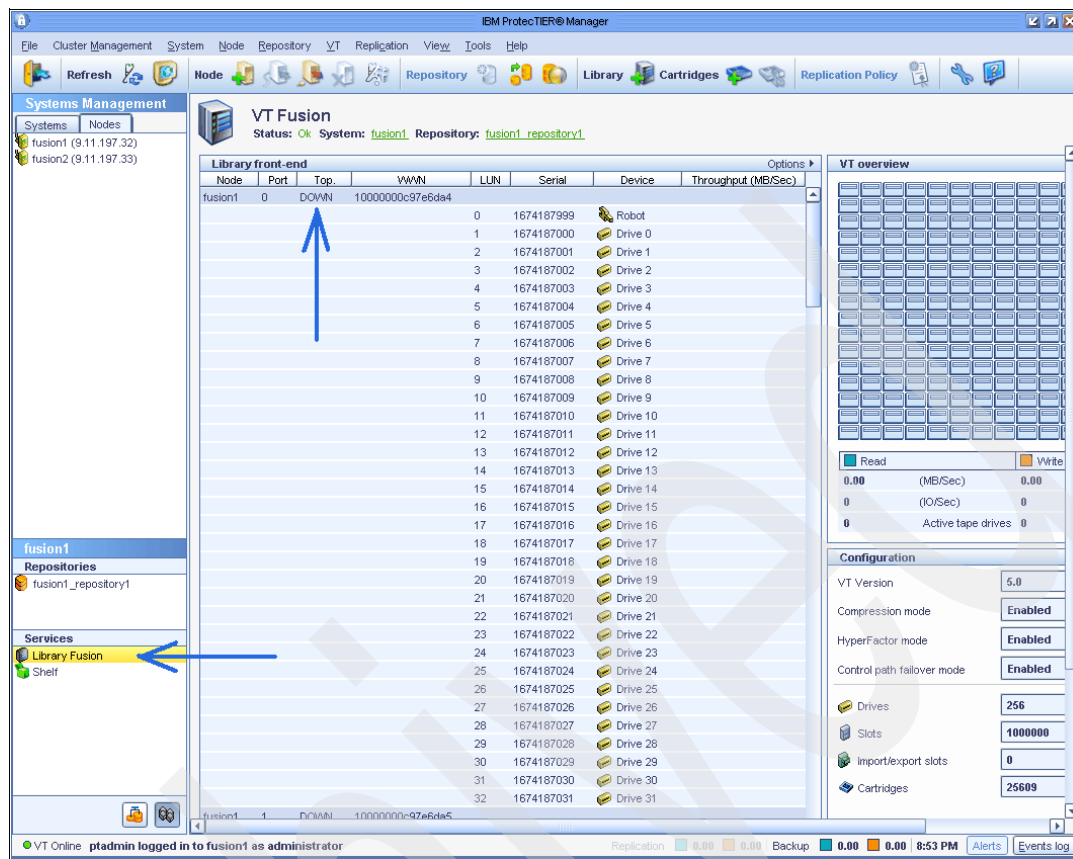


Figure 10-8 Library Fusion_link down

FIX

Call your IBM System Service Representative (SSR) and provide all necessary information.

One ProtecTIER server down

Let us assume that one ProtecTIER server goes either OFFLINE or POWER DOWN; see Figure 10-9 on page 276. In the event of an error condition detected by the tape daemon in the Enterprise controller, a signal is sent to the TS7680 to initiate a log snap shot. These logs will remain on the TS7680 to be gathered later by Service when required.

Note: This scenario also applies in case ProtecTIER server maintenance is needed (example: Licensed internal code upgrade).

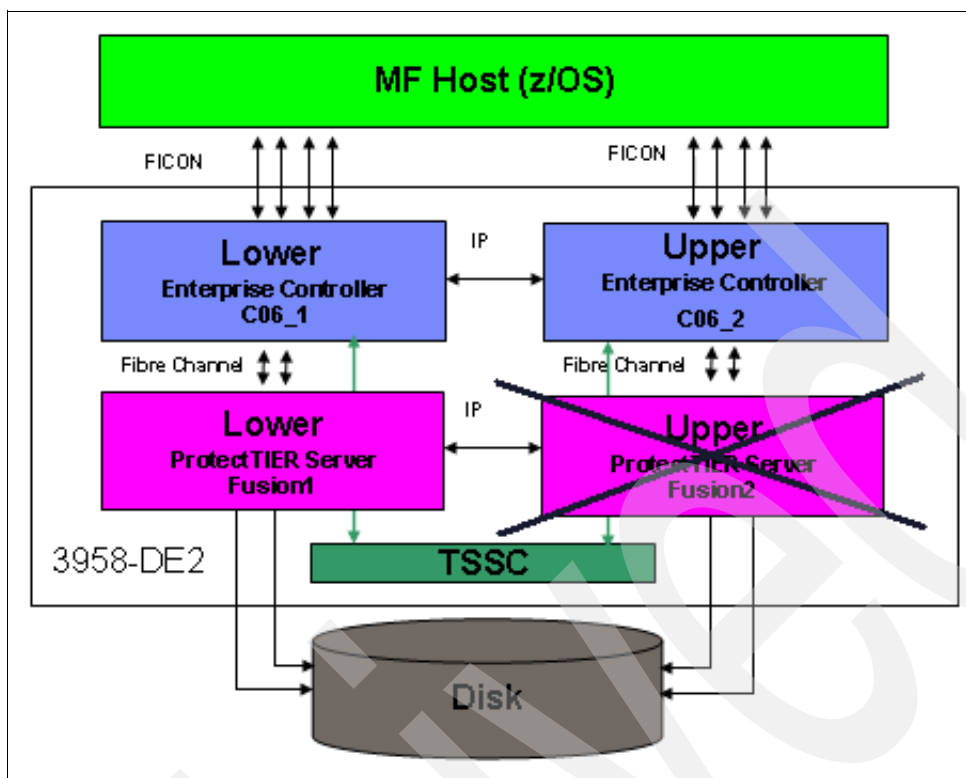


Figure 10-9 Upper PT server down

Even in this case you will lose 128 virtual addresses, because they are owned by either Enterprise controller and the ProtectTIER server.

From ProtectTIER Manager you will immediately see the server OFFLINE, as shown in Figure 10-10 on page 277.

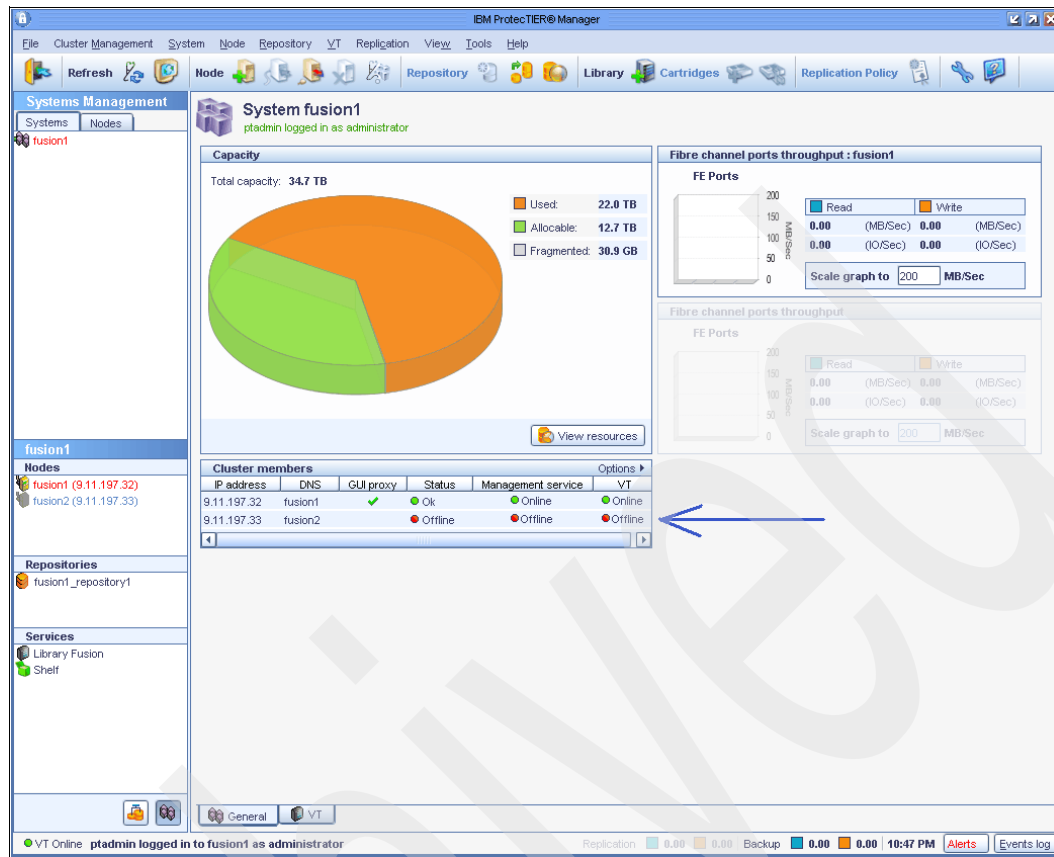


Figure 10-10 PT Server_Offline

Here you can check by clicking the left pane node name (**Fusion2**); you will get info about PT Server Status, shown in Figure 10-11 on page 278.

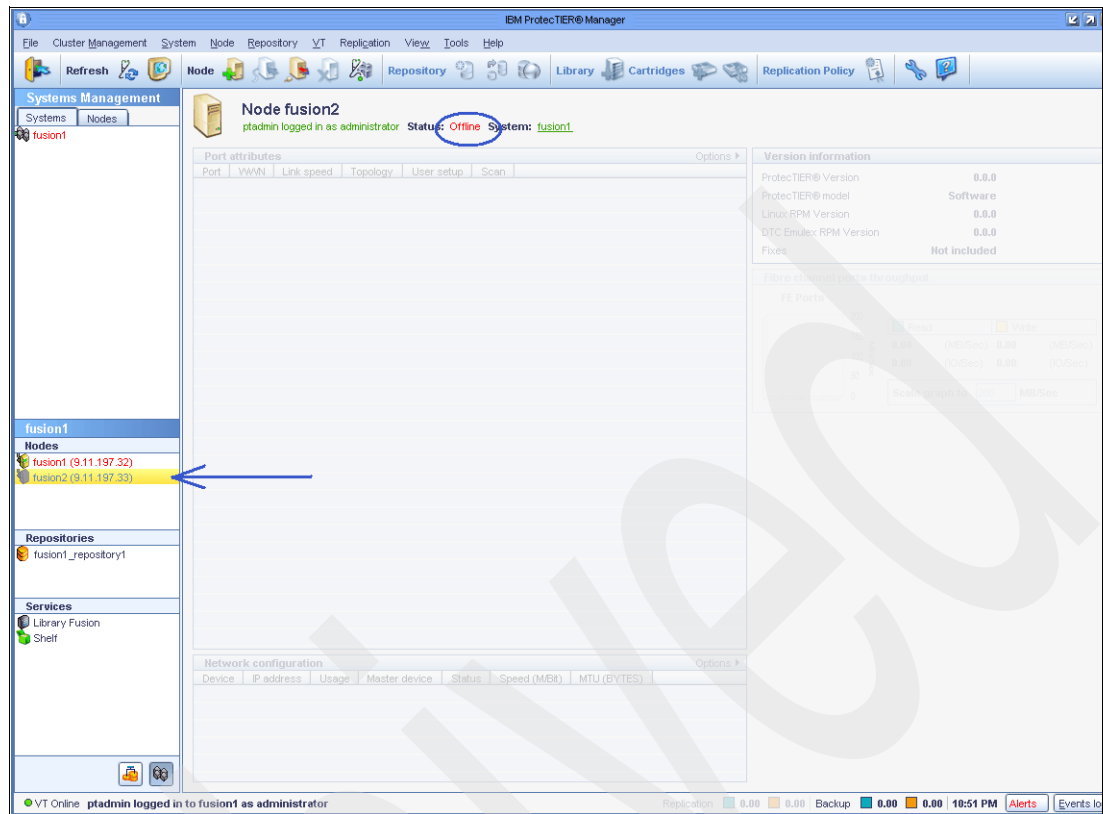


Figure 10-11 PT Server_Offline_chk

Also, as stated before, you will lose the 128 virtual addresses configured on the Enterprise controller connected to this ProtecTIER server (Upper); see Figure 10-12 on page 279.

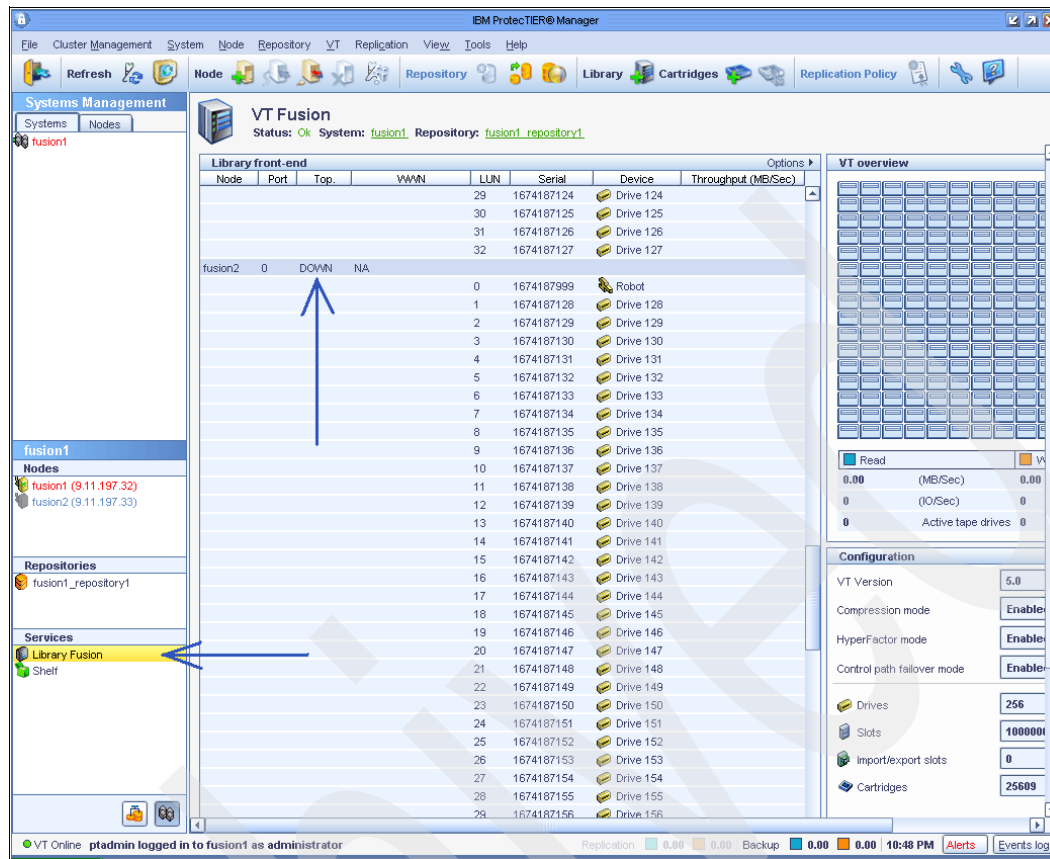


Figure 10-12 PT Server_128 addr_down

There is no problem on the Lower ProtecTIER server, which will continue to work.

Click the server name (**Fusion1**) on the left of the panel and you will get what is seen in Figure 10-13 on page 280.

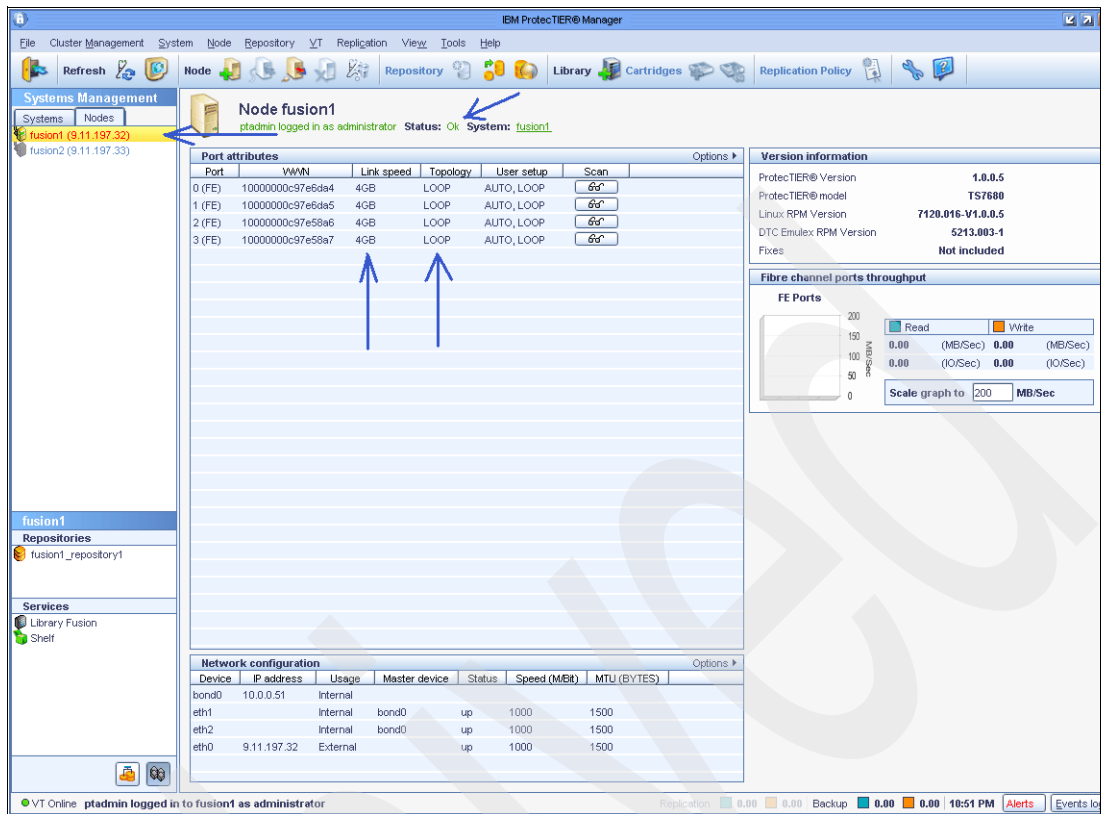


Figure 10-13 PT Server_Upper_OK

If you want to see more details about the Fusion2 problem, just click **Alert** in the lower right corner; see Figure 10-14 on page 281.

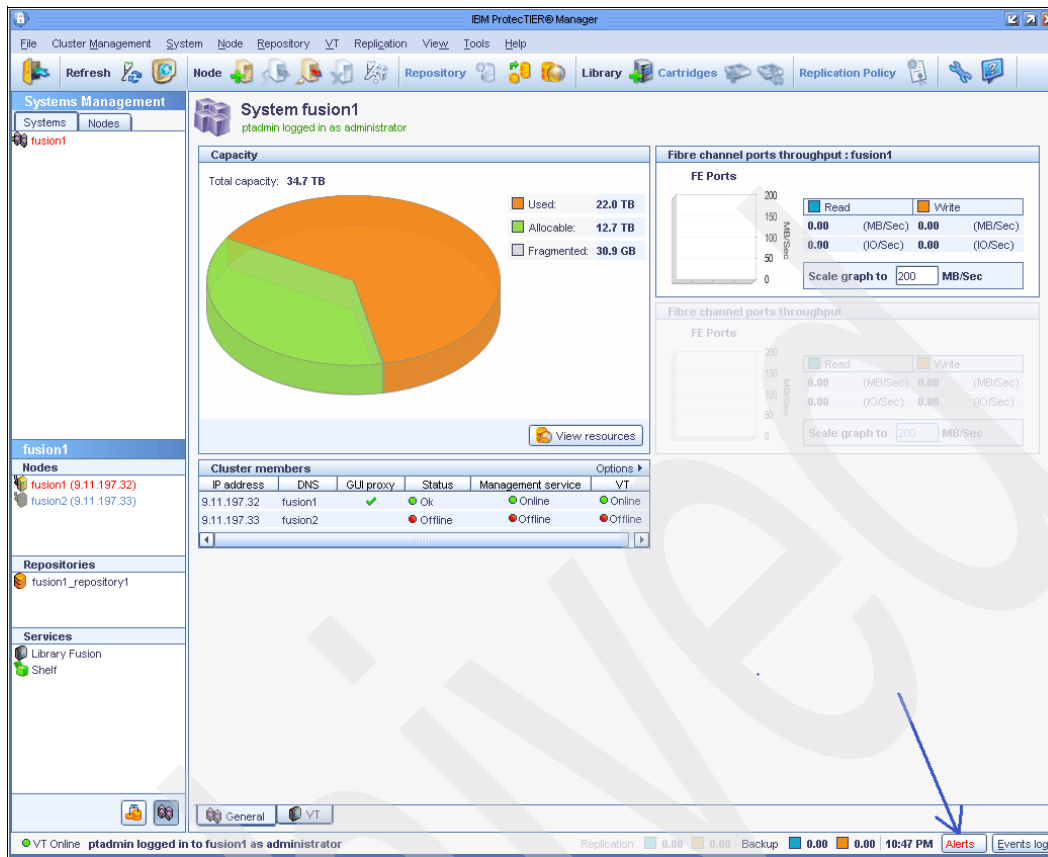


Figure 10-14 More details about the Fusion 2 problem

You will receive the Alert Log panel, shown in Figure 10-15.

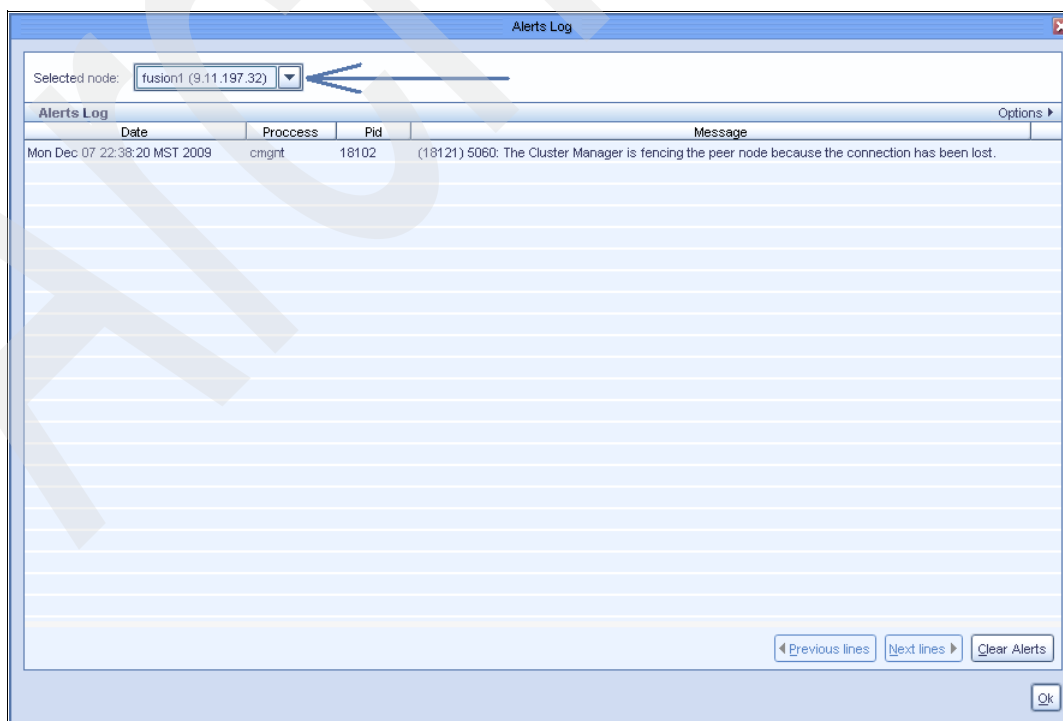


Figure 10-15 PT Server_Alerts log

This alert message is from the running PT server (Fusion1).

If you try to see the alert message from the down PT server (Fusion2) by clicking the **Selected node** menu (see Figure 10-16), you will receive a Communication error message.

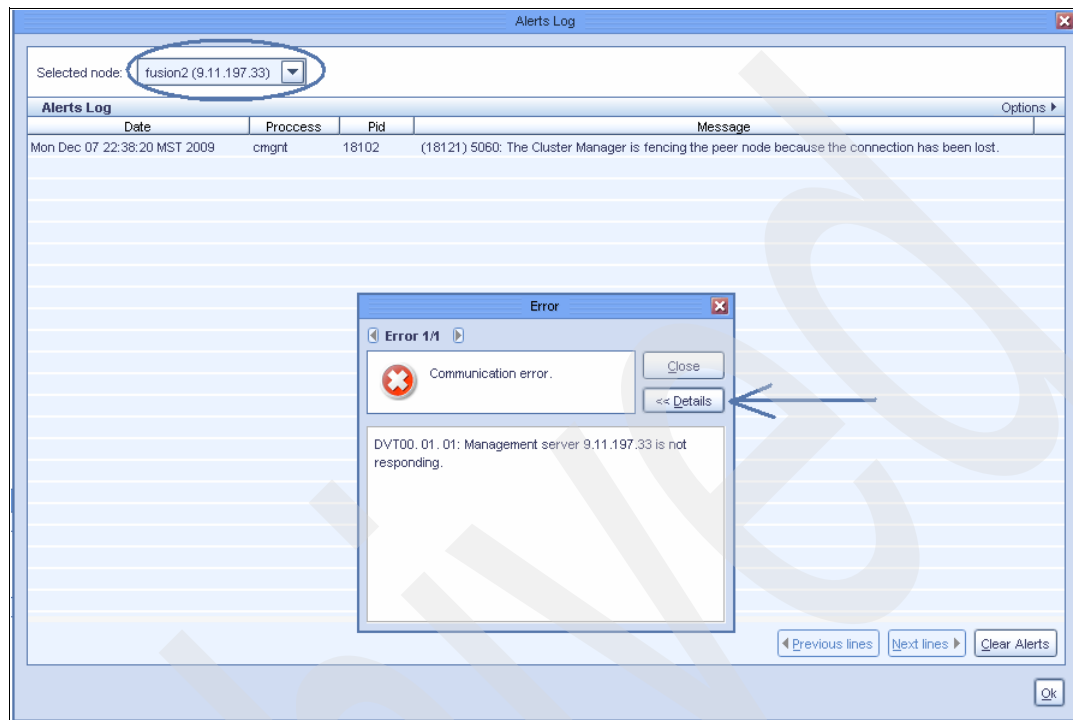


Figure 10-16 PT Server_Error

FIX

Call your IBM System Service Representative (SSR) and provide all necessary information.

After your SSR has performed the service action you can restore the TS7680 from the disk repository (Figure 10-17).

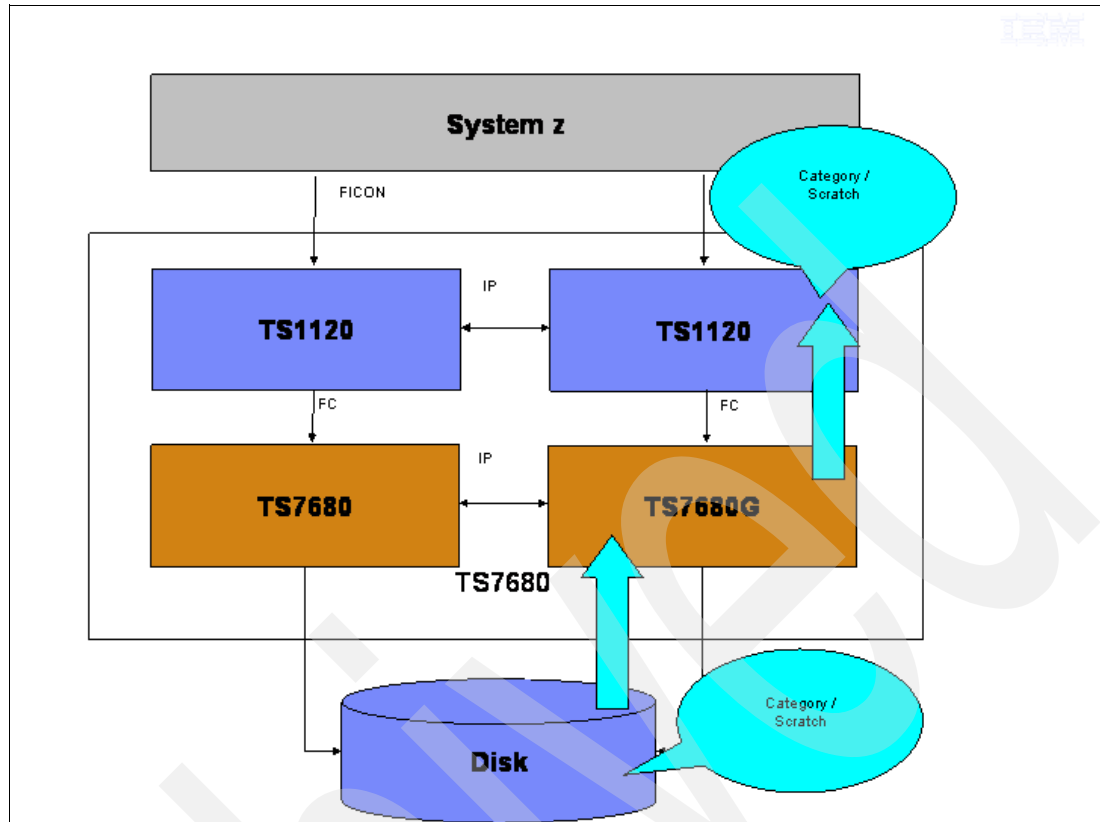


Figure 10-17 Restoring failed TS7680

Link down between Enterprise controllers

In case of link problems between Enterprise controllers (see Figure 10-18 on page 284), the Integrated Removable Media Manager (IRMM) will stop the synchronization process between primary and alternate databases for 2 minutes. Then, if the problem is on the alternate controller (checked by Heartbeat), the IRMM will continue on the primary controller and no takeover will occur.

Generally, if the Ethernet link between control units goes down, it would not immediately affect performance because I/O would still continue. However, new mounts on the standby would fail because that side cannot talk to IRMM to issue the mounts. So, eventually the drives on the standby CU would no longer be used. A Service Informational Message (SIM) is also generated in this case and it will be posted to the system console. However, there would be nothing in the information on the console that would specify that the problem is with the Ethernet link. Your System Service Representative (SSR) will be able to figure out the problem by looking at the detailed information generated at the time of the SIM.

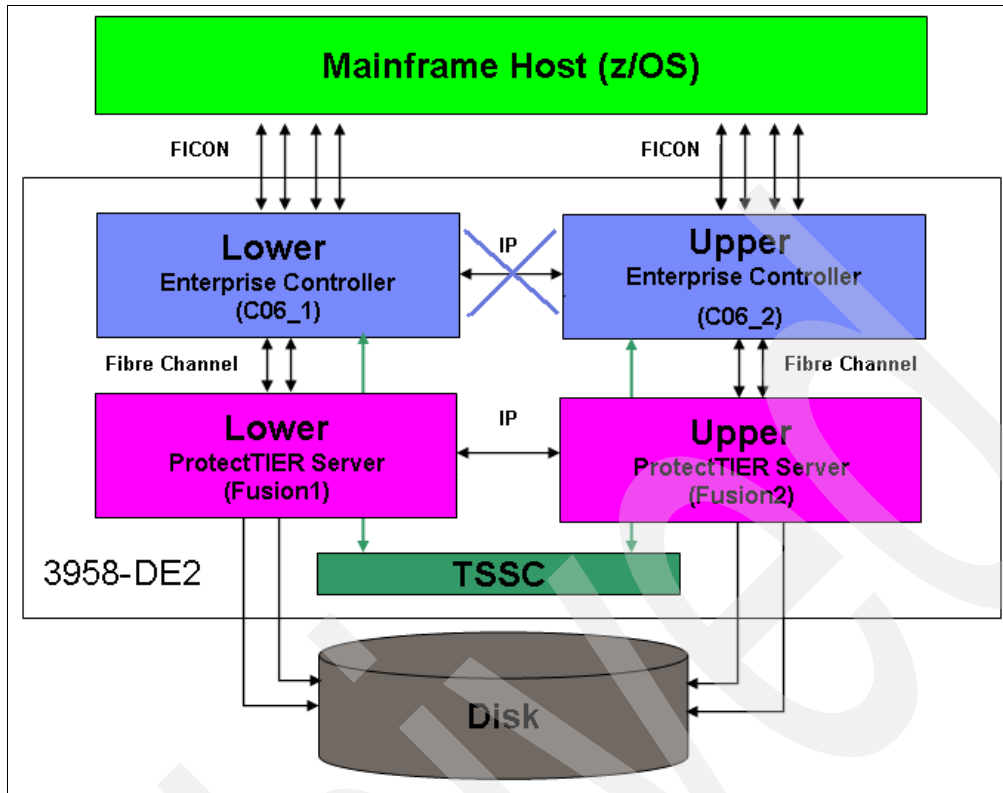


Figure 10-18 Failover Scenario_Link Enterprise Controller_down

Cabling between the two Enterprise controllers is not a main data route. This is for IRMM communication. If this link goes down, IRMM communication is gone and no library activities are possible (mounts, and so on), but running read, write tape activities are still possible.

The monitors are still up because this communication goes through the protectTIER servers. See Figure 10-19 on page 285.

A Service Informational Message (SIM) will be generated by the standby Upper Enterprise controller (Enterprise Controller_2), and the corresponding error log entry will report "IP is no longer functional." The standby Library Integration & Emulation (LIE) will lose communication to IRMM and notify the primary. The primary Lower Enterprise controller (Enterprise Controller_1) will continue to function as normal. All takeover processes will be disabled because the databases will get out of sync.

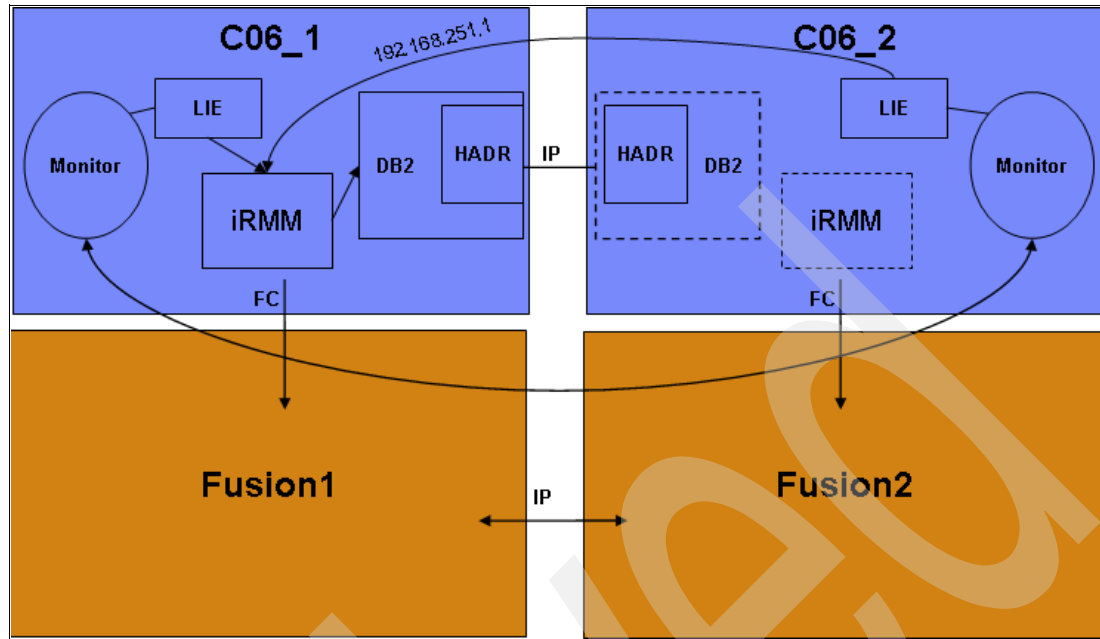


Figure 10-19 Failover Scenario_Monitor process

Enterprise Controller_1 communicates with Enterprise Controller_2 using a virtual IP address (192.168.251.1) so that can easily change: in case one Enterprise controller goes down, the other Enterprise controller takes over.

There is only one active DB2 at the time and only one IRMM running, so in case of switchover the standby Enterprise controller will start IRMM.

Monitor process

The monitor process running in each Enterprise controller monitors the health of its own system and communicates its status to the other Enterprise controller's monitor process through the ProtectTIER server.

Monitors:

- ▶ DB2/HADR status
- ▶ TCP/IP between Enterprise controllers
- ▶ Library function status
- ▶ ProtectTIER status

When a failure of the primary node occurs, the monitor process on the standby takes over the primary role. When service is initiated on the primary node, the monitor process on the primary requests standby to take over.

Lower Enterprise controller (Enterprise Controller_1) loses power

In case of a power problem on the primary Lower Enterprise controller (Enterprise Controller_1), see Figure 10-19, the standby Upper Enterprise controller (Enterprise Controller_2) will see a TCP/IP failure, and Enterprise Controller_1 will no longer send the heartbeat message. A Service Informational Message (SIM) is generated by the standby Upper (Enterprise Controller_2), and the corresponding error log entry reports "IP is no longer functional." Enterprise Controller_2 will initiate a takeover of the database, configure the 192.168.251.1 address, and start IRMM. All drives on the down Enterprise controller

(Enterprise Controller_1) will be demounted, and all takeovers will be disabled until Enterprise Controller_1 is back online.

Upper Enterprise Controller (Controller_2) loses its PT server (Fusion2)

The Upper Enterprise controller will lose communication to all fibre channel devices, a Service Informational Message (SIM) will be generated by the Upper Enterprise Controller_2, and the corresponding error log entry will report “FIBER is no longer valid.” The monitor process on Enterprise Controller_2 will no longer be able to send status and all drives on the standby CU will be demounted. All takeovers will be disabled because Enterprise Controller_2 cannot communicate to the Virtual Tape Library (VTL).

Maintenance done on Lower Enterprise Controller_1

When the Lower Enterprise controller (Enterprise Controller_1) is taken offline, it requests the standby Control Unit (CU) Enterprise Controller_2 (Upper Enterprise controller) to start a takeover.

If the system is in a good state, the takeover is allowed and should complete successfully. Once Enterprise Controller_1 is offline, maintenance can be done without interrupting the host on the other CU.

Check and Recover

At any time and mainly during a problem determination you can use the Check and Recover procedure by simply clicking the **System** menu and choosing **Check and Recover**, as shown in Figure 10-20.

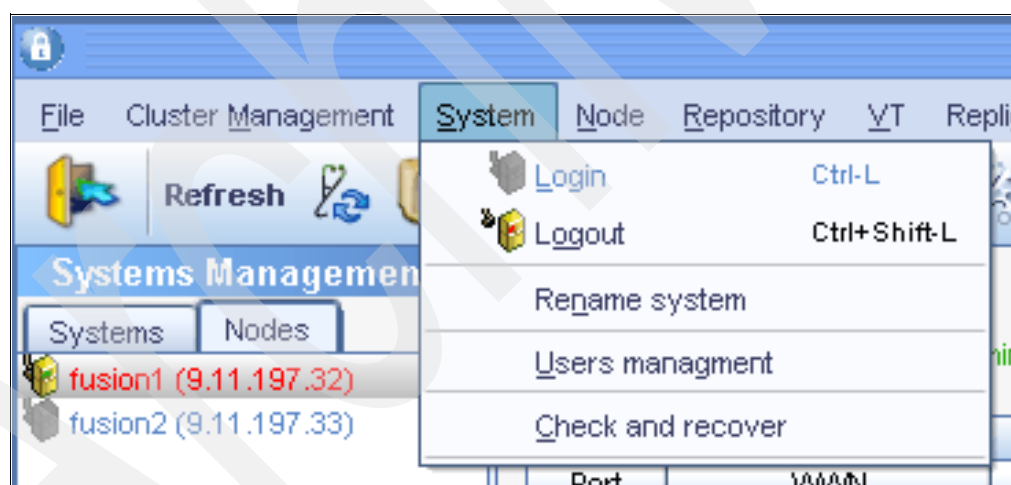


Figure 10-20 Check and Recover

Note: The complete system must be offline.

You will receive a prompt message for confirmation to put the complete system offline; see Figure 10-21 on page 287.

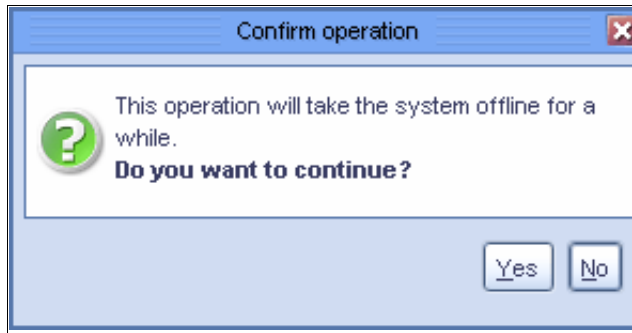


Figure 10-21 Check and Recover_confirm

ProtectTIER microcode will do all tests and give you a result; see Figure 10-22 as example.

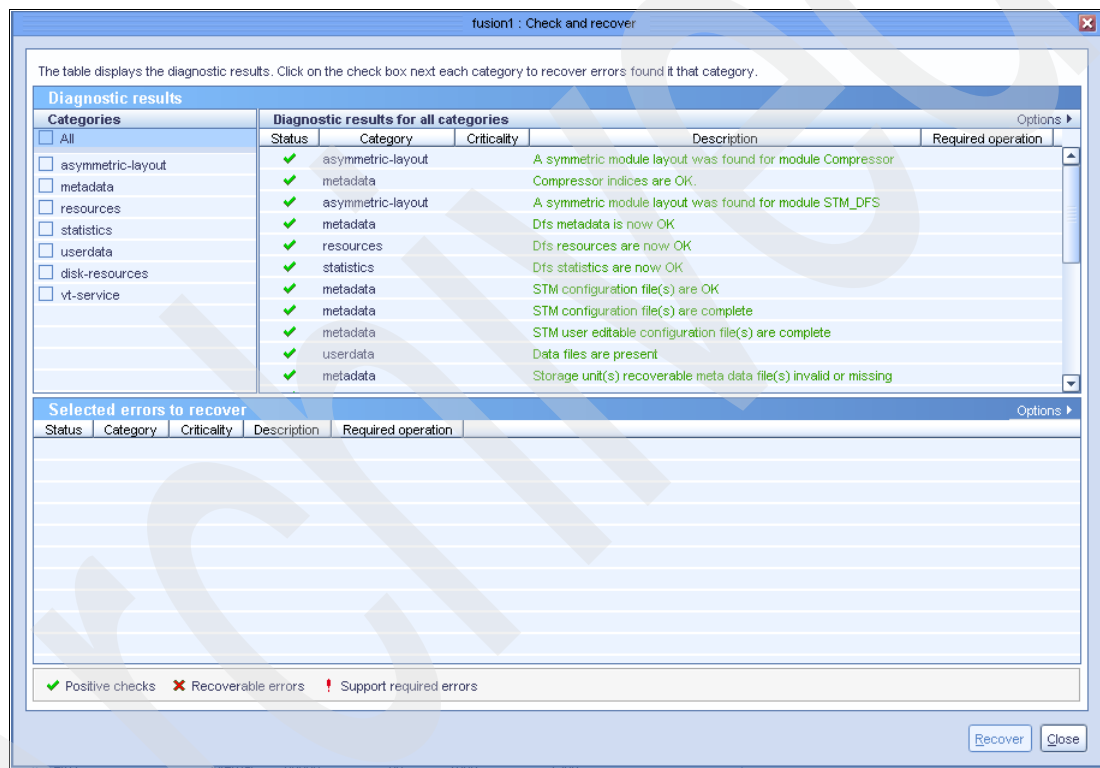


Figure 10-22 Check and Recover_result

Depending on what the message says in case of a problem, you can call your IBM System Service Representative (SSR) and provide him with all test message results.

Archived



Part 4

Appendixes

In this part we offer several appendixes dealing with checklists, power, and upgrades.

Archived

Reliability, availability, and serviceability update

This appendix provides a procedure to update ProtectTIER RAS licensed internal code.

The Reliability Availability and Serviceability (RAS) licensed internal code update procedure will perform the following steps:

1. Vary the Enterprise controller virtual devices offline from z/OS, corresponding to the related ProtectTIER server.
2. Notify the Enterprise controller of changes in the ProtectTIER server status.
3. RAS licensed internal code update option.
4. Vary the Enterprise controller virtual devices back online from z/OS.

A.1 Vary virtual devices offline

Before starting with RAS licensed internal code update the virtual devices of Enterprise controller related to the ProtecTIER server must be varied offline from z/OS.

A.2 Notify the Enterprise controller that the ProtecTIER server will be offline

Notify the Enterprise controller before the ProtecTIER server is varied offline by logging into the ProtecTIER server and from the command line type `rasMenu` [Enter] and you will get the (RAS) menu shown in Figure A-1.

```
+-----+
| RAS Text Based Menu running on XXX.XXX.XXX.XXX |
+-----+
| 1) Check if RAS is running |
| 2) Start RAS |
| 3) Stop RAS |
| 4) Get RAS Version |
| 5) Get PT Code Version |
| 6) Display Firmware Levels |
| 7) Manage Configuration (...) |
| 8) System Health Monitoring (...) |
| 9) Problem management (...) |
| 10) Call Home Commands (...) |
| 11) Collect Logs (...) |
| 12) Enterprise Controller (...) |
| 13) Exit |
+-----+
```

Figure A-1 RAS menu

Select **Enterprise Controller** and you will get Figure A-2.

```
+-----+
| > Sub-menu of 'Enterprise Controller' |
+-----+
| 1) Notify Enterprise Controller of Update PT code Start |
| 2) Notify Enterprise Controller of Update PT code End |
| 3) Shutdown Enterprise Controller |
| 4) Execute timezone check |
| 5) Back to parent menu |
| 6) Exit |
+-----+
```

Figure A-2 RAS Enterprise controller menu

From the this menu select **Notify Enterprise Controller of Update PT code Start** and wait for the following message shown in Figure A-3 on page 293.

```

Begin Procedure: Notify Enterprise Controller of Update PT code Start

>>> This operation will notify the Enterprise Controller that the PT
code is being updated.
This will cause the Enterprise Controller to not report errors for the
PT paths being down.
Are you sure you would like to continue (y/n)?

```

Figure A-3 RAS Enterprise Controller menu_01

Answer **yes** and a new question will appear, shown in Figure A-4.

```

>>> Are you upgrading both nodes concurrently (y/n)?

```

Figure A-4 RAS Enterprise Controller menu_01_1.gif

Answer **no** and wait for the message shown in Figure A-5.

```

>>> Please upgrade PT code, when done please return here and run the o
Enterprise Controller of Update PT code End".

End Procedure: Notify Enterprise Controller of Update PT code Start
Press any key to continue

```

Figure A-5 RAS Enterprise Controller menu_01_2

Just type any key and the menu shown in Figure A-6 will display.

```

+-----+
| > Sub-menu of 'Enterprise Controller' |
+-----+
| 1) Notify Enterprise Controller of Update PT code Start |
| 2) Notify Enterprise Controller of Update PT code End   |
| 3) Shutdown Enterprise Controller                       |
| 4) Execute timezone check                               |
| 5) Back to parent menu                                  |
| 6) Exit                                                  |
+-----+

```

Figure A-6 RAS Enterprise Controller menu

Select **Exit** to go back to the command line.

10.2.3 RAS licensed internal code update option

Get the latest RAS licensed internal code from your IBM Representative.

Upload the updateRAScode-4-3.x.xx.tar file to a temporary directory on the ProtecTIER server to be updated.

Stop the RAS licensed internal code by using the **rasMenu** tool.

At the ProtecTIER server command line type the following command:

```

rasMenu [enter]

```

You will get the RAS menu, shown in Figure A-7 on page 294.

```

+-----+
| RAS Text Based Menu running on XXX.XXX.XXX.XXX |
+-----+
| 1) Check if RAS is running |
| 2) Start RAS |
| 3) Stop RAS |
| 4) Get RAS Version |
| 5) Get PT Code Version |
| 6) Display Firmware Levels |
| 7) Manage Configuration (...) |
| 8) System Health Monitoring (...) |
| 9) Problem management (...) |
| 10) Call Home Commands (...) |
| 11) Collect Logs (...) |
| 12) Enterprise Controller (...) |
| 13) Exit |
+-----+

```

Figure A-7 RAS menu

From this menu select **Stop RAS** and a message will appear, as seen in Figure A-8.

```

Begin Procedure: Stop RAS

>>> Do you want to stop RAS on both nodes?

```

Figure A-8 RAS menu_03

Answer **no** and the message shown in Figure A-9 will appear.

```

RAS is down

End Procedure: Stop RAS
Press any key to continue

```

Figure A-9 RAS menu_03_1

Press any key to go back to the RAS menu (see Figure A-7) and select **Exit** to go back to the command line.

From the command line, enter the following command to expand the tar file:

```
tar -xvf updateRAScode-4-3.x.xx.tar [enter]
```

The following files are created:

- ▶ updateRAScode
- ▶ RASupdate.tar
- ▶ README

From the command line, enter the following command to update the RAS licensed internal code:

```
./updateRAScode
```

This command expands the RAS licensed internal code and runs the **rsCerCfgWebInstall** command to update the RAS licensed internal code. After the installation, the expanded RAS

licensed internal code is erased.

Wait until the message shown in Figure A-10 is displayed.

```
Cerberus RAS installation completed successfully
```

Figure A-10 RAS update complete

Delete both the updateRAScode and updateRAScode-4-3.x.xx.tar files on your temporary directory.

Note: RAScode will start automatically at the end of the update.

10.2.4 Notify the Enterprise controller that ProtecTIER server is back online

From the command line type `rasMenu` [enter] and the menu shown in Figure A-11 will display.

```
+-----+
| RAS Text Based Menu running on XXX.XXX.XXX.XXX |
+-----+
| 1) Check if RAS is running |
| 2) Start RAS |
| 3) Stop RAS |
| 4) Get RAS Version |
| 5) Get PT Code Version |
| 6) Display Firmware Levels |
| 7) Manage Configuration (...) |
| 8) System Health Monitoring (...) |
| 9) Problem management (...) |
| 10) Call Home Commands (...) |
| 11) Collect Logs (...) |
| 12) Enterprise Controller (...) |
| 13) Exit |
+-----+
```

Figure A-11 RAS menu

From this menu select **Enterprise Controller** and the menu shown in Figure A-12 will display.

```
+-----+
| > Sub-menu of 'Enterprise Controller' |
+-----+
| 1) Notify Enterprise Controller of Update PT code Start |
| 2) Notify Enterprise Controller of Update PT code End |
| 3) Shutdown Enterprise Controller |
| 4) Execute timezone check |
| 5) Back to parent menu |
| 6) Exit |
+-----+
```

Figure A-12 RAS Enterprise Controller menu

From the this menu select **Notify Enterprise Controller of Update PT code End** and the message shown in Figure A-13 on page 296 will appear.

```
Begin Procedure: Notify Enterprise Controller of Update PT code End

>>> This operation will notify the Enterprise Controller that the PT
code is done updating its code.
This will cause the Enterprise Controller to start reporting errors
for the PT paths being down.
Are you sure you would like to continue (y/n)?
```

Figure A-13 RAS Enterprise Controller menu_02

Answer **yes** and the following message will appear (Figure A-14).

```
>>> Did you upgrade both nodes concurrently (y/n)?
```

Figure A-14 RAS Enterprise Controller menu_02_1

Answer **no** and the following message will appear (Figure A-15).

```
Initiating takeover on lower Enterprise Control Unit
Command completed successfully on lower Enterprise Control Unit

End Procedure: Notify Enterprise Controller of Update PT code End
Press any key to continue
```

Figure A-15 RAS Enterprise Controller menu_02_2

Press any key and select **Exit** from the RAS menu to go back to the command line.

10.2.5 Vary virtual devices online

Vary the Enterprise controller virtual devices back online from z/OS.

Note: If not at the same RAS licensed internal code level, or in case of need, repeat this procedure for the other ProtectTIER server.

Checklists

In this appendix, we summarize some checklists and worksheets that will help you during planning and installation of IBM System Storage TS7680 (3958-DE2).

- ▶ Client prerequisites for preinstallation tasks

This topic defines the client prerequisites for the 3958-DE2 (TS7680) preinstallation tasks.

- ▶ Client preinstall responsibilities for the 3958-DE2 (TS7680) installation

This topic provides a list of client responsibilities that must be performed before the installation of the 3958-DE2 (TS7680).

- ▶ 3958-DE2 (TS7680) physical specifications.

This topic provides the physical specifications for the 3958-DE2 (TS7680).

- ▶ 3958-DE2 (TS7680) electrical power ratings.

This topic provides the electrical power ratings for the 3958-DE2 (TS7680).

- ▶ Planning worksheets.

The topics in this section detail the planning worksheets for the 3958-DE2 (TS7680).

- IP address worksheet

Use this worksheet to specify the IP addresses assigned to the 3958-DE2 (TS7680) components.

- Client information worksheet

Use this worksheet to record your company information.

Client installation responsibilities

The topics in this section list and describe the responsibilities of the client in order to ensure the successful installation and configuration of the 3958-DE2 (TS7680).

Client prerequisites for preinstallation tasks

This topic defines the client prerequisites for the 3958-DE2 (TS7680) preinstallation tasks.

You are responsible for ensuring that the following prerequisites are complete before the IBM Service Support Representative (SSR) begins the preinstallation tasks in preparation for the installation of the 3958-DE2 (TS7680).

1. Ensure that the disk repository is installed and configured (refer to Disk repository overview and Disk repository configuration guidelines for more information).

Note: During the pre-sales process, you will receive a spreadsheet containing information to determine the optimal disk repository configuration. It is your responsibility to use this information to configure the disk repository and verify that it is configured correctly before the SSR comes on site for installation. A preinstalled conference call with the Field Technical Support Specialist (FTSS) is also required before the disk repository is configured.

2. Ensure that the fibre channel cables are available (refer to 3958-DE2 (TS7680) feature codes for more information).
3. Ensure that power is available for the 3958-DE2 (TS7680) (refer to 3958-DE2 (TS7680) electrical power ratings for more information).
4. Ensure that the Ethernet IP addresses are assigned and the cables are pulled (one for each ProtecTIER server and one for the TS3000 System Console [TSSC]) (refer to the IP address worksheet).
5. If the internal TSSC (FC 2722) is not ordered with your 3958-DE2 (TS7680), ensure that a TSSC is available. If you are using an external TSSC (FC 2714), the total line of sight distance from the frame to the external TSSC cannot exceed 31 m (100 ft.). (Refer to 3958-DE2 (TS7680) feature codes for more information.)
6. Ensure that the System z host is configured for the 3958-DE2 (TS7680) (refer to Host compatibility for more information).
7. Ensure that FICON cables are available (refer to 3958-DE2 (TS7680) feature codes for more information).

Your responsibilities for the 3958-DE2 (TS7680) installation

This topic provides a list of your responsibilities that must be performed before the installation of the 3958-DE2 (TS7680).

The 3958-DE2 (TS7680) installation requires participation from you, as well as the coordinated efforts of IBM Service Support Representative (SSR), Lab-based Services, and Customer Service Center (CSS) personnel. Table B-1 on page 299 outlines your responsibilities and lists the required action, or a link to another topic that provides additional information.

Table B-1 Customer responsibilities for 3958-DE2 (TS7680) installation.

Customer responsibility	Required Action.
Complete the planning and preparation tasks described in the 3958-DE2 (TS7680) Customer Information Center.	Ensure that your site meets all of the requirements outlined in Physical specifications and site planning.
Meet the pre installation requirements outlined in the 3958-DE2 (TS7680) Customer Information Center.	Ensure that all preinstallation prerequisites are met prior to the start of preinstallation tasks. Refer to Customer prerequisites for preinstallation tasks.
Complete the Company Information and IP Address worksheets.	Refer to the Company information worksheet. Refer to the IP Address worksheet.
Select your scratch delay and scratch delay override settings. These settings are set by the SSR at the time of install. Note: If, at any point after installation, these settings need to be changed, the system must be taken offline and there is a charge for a Customer Engineer (CE) to make the changes.	Scratch Delay <ul style="list-style-type: none"> ▶ 0 - data on volumes put into the scratch category will be erased from the repository immediately. ▶ 9 - data on volumes put into the scratch category will be erased from the repository after a 9-day lapse. Scratch Delay Override <ul style="list-style-type: none"> ▶ No – data on volumes placed in the scratch category will not be erased from the repository until 9 days have elapsed. ▶ Yes – data on volumes placed in the scratch category may be erased early if repository limited space critical state reached.
If a TS3000 System Console (TSSC) is not installed in the frame, confirm that the TSSC to be used with the 3958-DE2 (TS7680) has FC 2719 or equivalent and that FC 2714 has been ordered.	Refer to 3958-DE2 (TS7680) feature codes.
It is recommended that the customer have one or more workstations designated to run the ProtecTIER Manager software.	Refer to Prerequisites for the ProtecTIER Manager workstation.
Purchase and have available additional Ethernet (Cat 5e or higher) and any fibre channel cables, if needed.	Refer to the topics in the Planning section for more information.

3958-DE2 (TS7680) physical specifications

Table B-2 displays the physical specifications of the 3958-DE2 (TS7680).

Table B-2 3958-DE2 (TS7680) physical specifications

Width	644 mm (25.35 in)
Depth	1102 mm (43.39 in)
Height	1804 mm (71.02 in)
Weight (Note 1)	442.25 kg (975 lb)

Note 1: This is the total weight of the frame with all components that comprise the 3958-DE2 (TS7680) including:

- ▶ Enterprise controllers (2)
- ▶ ProtecTIER servers (2)
- ▶ TS3000 System Console (TSSC)
- ▶ Routers
- ▶ Switches
- ▶ Power Distribution Unit

3958-DE2 (TS7680) electrical power ratings

The 3958-DE2 (TS7680) ships with two internal Power Control Assemblies (PCA).

Table B-3 provides the electrical power ratings per line cord for the 3958-DE2 (TS7680).

Table B-3 3958-DE2 (TS7680) electrical power ratings

Unit	Voltage AC	Frequency	Current (Amps)	Inrush current (A)	Power (Watts)	kVA	KBtu/hr
3958-DE2 (TS7680)	200-240 VAC	50-60 HZ	15	250	3000	3.0	10.24

Planning worksheets

The topics in this section detail the planning worksheets for the 3958-DE2 (TS7680).

IP address worksheet

Use this worksheet to specify the IP addresses assigned to the 3958-DE2 (TS7680) components.

IBM Service Representatives use the information provided to define the IP addresses of components supported by the TS3000 System Console (TSSC). When the TSSC sends Call Home information to IBM through VPN or modem, or sends you notices about serviceable events, these settings are included in the information to identify and provide important information about the TSSC that sent a service request.

Table B-4 on page 301 shows the default IP addresses for the ProtecTIER server. Table B-5 on page 301 shows the default IP addresses for the TSSC. Table B-6 on page 302 shows the IP addresses of the 3958-DE2 (TS7680) components.

Note: All components use subnet mask 255.255.255.0.

Table B-4 Factory default ProtecTIER server IP addresses for a clustered 3958-DE2 (TS7680)

3958-DE2 (TS7680) clustered	Component	Port	IP Address
Node A (the ProtecTIER server located in the lower part of the rack)	Server	eth0	Customer IP
	Note: By default, the ProtecTIER server uses the IP address range 10.0.0.50 through 10.0.0.59 for the power control network. The server IP addresses do not change from frame to frame.		
	Server	eth1	10.0.0.51
	Server	eth2	10.0.0.51
	Server	eth3	N/A
	Server	eth4	N/A
	Server	eth5	172.31.1.xx
	Network Power Switch	N/A	10.0.0.51
	Customers Gateway Address		
	WAN Subnet Mask		
	1st Customer DNS		
	2nd Customer DNS		
Node B (the ProtecTIER server located in the upper part of the rack)	Server	eth0	Customer IP
	Note: By default, the ProtecTIER server uses the IP address range 10.0.0.50 through 10.0.0.59, for the power control network. The server IP addresses do not change from frame to frame.		
	Server	eth1	10.0.0.52
	Server	eth2	10.0.0.52
	Server	eth3	N/A
	Server	eth4	N/A
	Server	eth5	172.31.1.xx
	Network Power Switch	N/A	10.0.0.51

TSSC IP addresses

Table B-5 TSSC IP addresses

TSSC	Ethernet Port	IP Address
TSSC	eth0	Customer IP
TSSC	eth1	172.31.1.1

IP addresses for components housed in the 3958-DE2 (TS7680)

Table B-6 ProtecTIER server and Enterprise controller IP addresses

Component	IP addresses: Node A in a cluster (the server located in the lower part of the rack)	IP addresses: Node B in a cluster (the server located in the upper part of the rack)
ProtecTIER server	.xx0 and .xx1	.xx5 and .xx6
Component	IP addresses: Lower controller	IP addresses: Lower controller
Enterprise controller	.xx2 and .xx3	.xx7 and .xx8

Customer information worksheet

Use this worksheet to record your company information.

IBM Service Representatives use the information that is provided on the company information worksheet to customize your IBM storage complex. When you use any of the remote support features, the TS3000 System Console (TSSC) sends this information to IBM so an IBM service representative can contact you.

Note: There may be situations where the information in this worksheet is needed. Keep this worksheet in a safe place for future reference.

Table B-7 Company information worksheet

Required information	Description	Your information
Business company name	The full name of your company. IBM Service Representatives use this information to identify your company when they receive Call Home reports from your IBM storage system. Ensure that the company name provided is consistent with all other machines that correspond to your IBM customer account.	
Customer number	The IBM-assigned customer number for your company.	
Country code	The two-digit number that must be used in order to reach your country by phone or fax, from another country. This is not the three-digit RETAIN country code. Example: it for Italy	

Required information	Description	Your information
Frame range number	<p>This value is specific for setting the IP address range for a TSSC connection (172.31.1.xxx).</p> <p>With two nodes, the frame range number entered should be the same number used for the first frame. For example: 10 for the original server's frame number (172.31.1.10), and when prompted, enter 10 for the second node. The system will automatically assign the frame range to 15 and the IP of the second node (in this case: 172.31.1.15).</p> <p>For each 3958-DE2 (TS7680) within the 10 IP address space, IP addresses are used as shown in the following example:</p> <ul style="list-style-type: none"> ▶ ProtecTIER server - 172.31.1.10 ▶ ProtecTIER server - 172.31.1.15 ▶ Enterprise controller - 172.31.1.12 ▶ Enterprise controller - 172.31.1.17 <p>To check the original server's IP address, view the Attached System list on the TSSC, where the last octet of the IP address is the frame number (172.31.1.10 = 10).</p>	<p>Node 0: 3958-DE2 (TS7680) Frame Range Number _____</p> <p>Node 1: 3958-DE2 (TS7680) Frame Range Number _____</p>
System administrator information Provide information about your storage system administrator in the following section.		
Administrator name	The name of the individual at your site whom IBM Service Representatives should contact about IBM storage system service matters.	
Administrator e-mail address	The storage system administrator's email address.	
Voice phone number	The primary telephone number that IBM Service Representatives should use to contact the storage system administrator. Include the area code and the country code, if appropriate.	
Fax number	The primary fax number that IBM Service Representatives should use to fax documents to the storage system administrator. Include the area code and the country code, if appropriate.	
Alternate fax number	An alternate fax number that IBM Service Representatives can use to fax documents to the storage system administrator. Include the area code and the country code, if appropriate.	

Required information	Description	Your information
ProtectTIER server and storage system information Provide basic information about your ProtectTIER servers and storage system in the following section.		
Machine type and model number	The machine type and model number for the ProtectTIER servers.	3958-DE2
Machine location	The address of the facility where the ProtectTIER servers reside. If different from the administrator mailing address above, provide the full street address, building (if appropriate), city or locality, state or province, and postal or zip code.	
Call back phone number	The phone number of the modem being used for Call Home. Include the area code and the country code, if appropriate.	
Disk storage machine type(s) and model number(s)	The machine type(s) and model number(s) for the attached disk storage subsystem(s). For non-IBM equipment, also provide vendor name(s).	
Disk storage serial number(s)	The serial number(s) for the attached disk storage subsystem(s).	
ProtectTIER server A (lower)	Customer IP Address _____ Customer Gateway _____ WAN Subnet Mask _____ 1st Customer DNS _____ - _____ 2nd Customer DNS _____ - _____	
ProtectTIER server B (upper)	Customer IP Address _____ Customer Gateway _____ WAN Subnet Mask _____ 1st Customer DNS _____ - _____ 2nd Customer DNS _____ - _____	
TSSC server information Provide basic information about your TSSC server in the following section.		

Required information	Description	Your information
TSSC	TSSC name _____ Customer IP Address _____ _____ Customer Gateway _____ WAN Subnet Mask _____ _____ 1st Customer DNS _____ - _____ 2nd Customer DNS _____ - _____	
Enterprise controllers server information Provide basic information about the Enterprise controller servers in the following section.		
Library sequence number	A 5-character (hexadecimal) identification number assigned by the customer to this library.	
Scratch delay	The amount of time that scratch volumes are held (The default value of 9 days is recommended.)	(0 - 9 days)
Scratch delay override	Allows scratch volumes being held for 9 days to be deleted immediately if disk storage is limited	(yes or no)

Archived

WTI Network Power Switch

The 3958-DE2(TS7680) uses the Red Hat Cluster Suite for clustering two servers together.

The Red Hat Cluster Suite checks the health of the nodes in the cluster and will prevent data corruption in a clustered environment by “fencing” the peer in the cluster using the WTI Network Power Switch.

For example, if there are two nodes in a cluster, Node 1 and Node 2, and Node 2 stops responding to Red Hat Cluster Suite heartbeat messages, the Red Hat Cluster Suite running on Node 1 will automatically Telnet to the WTI Network Power Switch and toggle the power supplies connected to Node 2. This results in Node 2 rebooting, preventing Node 2 from modifying data on the shared storage when it cannot synchronize the modifications with its peer; see Figure C-1.

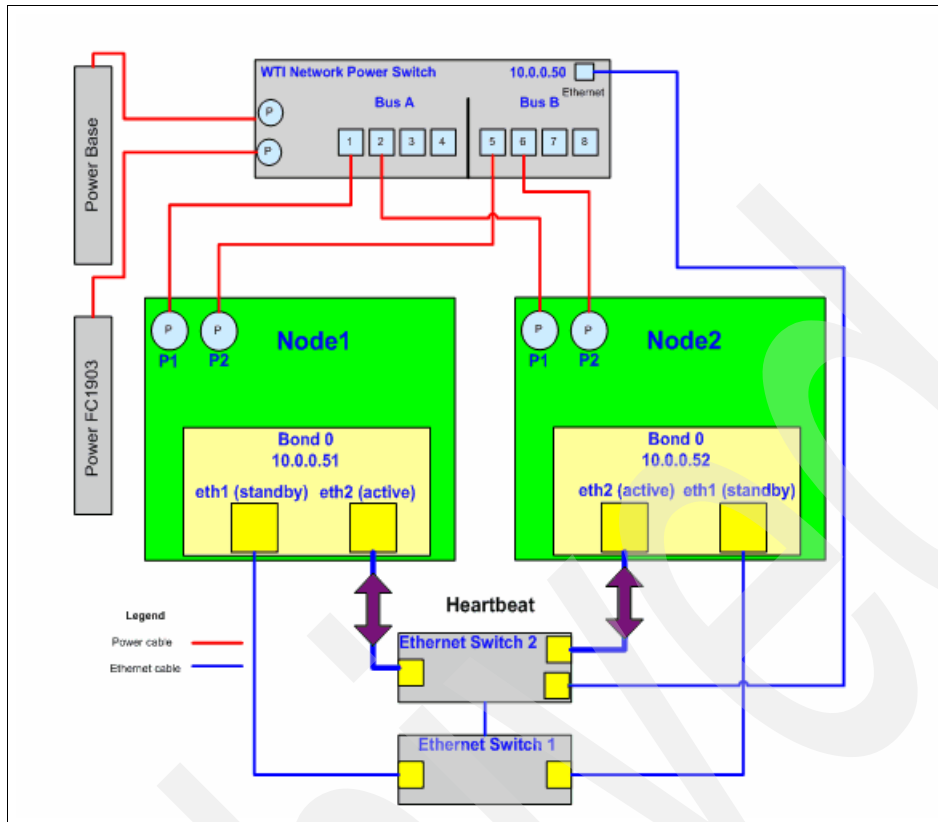


Figure C-1 Power fencing

If one of the nodes is fenced, the view panel from the ProtecTIER Manager also shows the fenced status; see Figure C-2.

Cluster members						Options ▾
IP address	DNS	GUI proxy	Status	Management service	VT	
9.11.201.19	Naples	✓	● Ok	● Ok	● Online	
9.11.200.233	Italy		● Fenced	● Offline	● Offline	

Figure C-2 Fenced status

In the rare case that the fenced status is not changing automatically from Fenced to Ok after the reboot of the node, you have to check the status of the ports of the WTI switch. Use a program, for example Putty if using Windows, or you can use SSH natively on UNIX systems to make an SSH connection to one of the ProtecTIER nodes. Once you are connected to your node, make a Telnet connection to the WTI switch. The default IP address is 10.0.0.50 and the default password is password. All ports should have the status ON; see Figure C-3 on page 309.

```
[root@roberto~]# telnet 10.0.0.50
Trying 10.0.0.50...
Connected to 10.0.0.50 (10.0.0.50).
Escape character is '^]'.

Enter Password: *****      Default password is "password"

Internet Power Switch v1.41h   Site ID: (undefined)

Plug | Name                | Password      | Status | Boot/Seq. Delay | Default |
-----+-----+-----+-----+-----+-----+
1 | internal_node1       | (undefined)   | ON     | 5 Secs          | ON      |
2 | internal_node2       | (undefined)   | ON     | 5 Secs          | ON      |
3 | (undefined)          | (undefined)   | ON     | 0.5 Secs        | ON      |
4 | (undefined)          | (undefined)   | ON     | 0.5 Secs        | ON      |
5 | internal_node1       | (undefined)   | ON     | 5 Secs          | ON      |
6 | internal_node2       | (undefined)   | ON     | 5 Secs          | ON      |
7 | (undefined)          | (undefined)   | ON     | 0.5 Secs        | ON      |
8 | (undefined)          | (undefined)   | ON     | 0.5 Secs        | ON      |
-----+-----+-----+-----+-----+

"/H" for help.

IPS>
```

Figure C-3 WTI Power Switch status

Figure C-3 shows the status of all the power ports. Port 1 and Port 5 are connected to Node 1, and Port 2 and Port 6 are connected to Node 2.

A useful function is the /H function. It will show you all the options available on the WTI switch.

```
IPS> /H and enter

Internet Power Switch v1.41h   Site ID: (undefined)

Display                                Configuration
/H      Display Help Screen        /G      View/Set General Parameters
/S      Display Plug Status         /P [n]   View/Set Plug Parameters
/SN     Display Network Status      /C      View/Set Serial Parameters
Control                                /N      View/Set Network Parameters
/D      Set Plugs to Default        /T      View/set Telnet Parameters
/Boot <n> Boot Plug n               /W      View/Set Web Server
/On <n>  Turn On Plug n             /E      Save Parameters
/Off <n> Turn Off Plug n            /R      Recall Parameters
/X      Exit/Disconnect             /DL     Download Parameters to File

Utilities
-----+-----+
| [n] = optional plug name or number |
| <n> = required plug name or number |
| n+n or n n = plug n and plug n    |
| n:n = plug n through plug n       |
| * = all plugs                      |
| ,Y = bypass "Sure? (y/n)"         |
+-----+-----+

IPS>
```

Figure C-4 Help function for the WTI switch

In the next example we show you how to set a port from Status Off to status On.

Internet Power Switch v1.41h Site ID: (undefined)					
Plug	Name	Password	Status	Boot/Seq. Delay	Default
1	internal_node1	(undefined)	OFF	5 Secs	ON
2	internal_node2	(undefined)	ON	5 Secs	ON
3	(undefined)	(undefined)	ON	0.5 Secs	ON
4	(undefined)	(undefined)	ON	0.5 Secs	ON
5	internal_node1	(undefined)	ON	5 Secs	ON
6	internal_node2	(undefined)	ON	5 Secs	ON
7	(undefined)	(undefined)	ON	0.5 Secs	ON
8	(undefined)	(undefined)	ON	0.5 Secs	ON

Figure C-5 WTI port 1 is OFF

Figure C-5 displays that from Node 1 port (plug) number 1 has the status OFF.

From the CLI type `/on 1` (see Example C-1) and Enter. A question will be Sure? (Y/N). Answer that question with Y and the port status will be changed from OFF to ON (see Figure C-6.)

Example C-1 Changing the port status

“/H” for help.

IPS> `/on 1 and enter`

Plugs to be turned on:

Plug 1: internal_node1

Sure? (Y/N): y

Internet Power Switch v1.41h Site ID: (undefined)					
Plug	Name	Password	Status	Boot/Seq. Delay	Default
1	internal_node1	(undefined)	ON	5 Secs	ON
2	internal_node2	(undefined)	ON	5 Secs	ON
3	(undefined)	(undefined)	ON	0.5 Secs	ON
4	(undefined)	(undefined)	ON	0.5 Secs	ON
5	internal_node1	(undefined)	ON	5 Secs	ON
6	internal_node2	(undefined)	ON	5 Secs	ON
7	(undefined)	(undefined)	ON	0.5 Secs	ON
8	(undefined)	(undefined)	ON	0.5 Secs	ON

Figure C-6 WTI status

Figure C-6 displays the status of the WTI after turning on Port 1.

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 312. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *IBM System Storage TS7600 with ProtecTIER*, SG24-7652-01
- ▶ *IBM TotalStorage Enterprise Tape 3592: Presentation Guide*, REDP-3749
- ▶ *z/OS V1R6 DFSMS Technical Guide*, SG24-6651
- ▶ *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring*, SG24-2229-07

Other publications

These publications are also relevant as further information sources:

- ▶ *z/OS DFSMS Object Access Method Planning, Installation and Storage Administration Guide for Tape Libraries*, SC35-0427
- ▶ *IBM System Storage TS7680 ProtecTIER Deduplication Gateway for System z Introduction and Planning Guide*, GA32-0650
- ▶ *License Information for IBM System Storage ProtecTIER Deduplication Gateway for System z*, GA32-0690
- ▶ *IBM System Storage TS7680 ProtecTIER Deduplication Gateway for System z User's Guide*, GA32-0651

Online resources

These Web sites are also relevant as further information sources:

- ▶ IBM System Storage TS7680 ProtecTIER Deduplication Gateway for System z Customer Information Center
<http://publib.boulder.ibm.com/infocenter/ts7680/cust/index.jsp>
- ▶ OA27786: NEW FUNCTION
<http://www-01.ibm.com/support/docview.wss?uid=isg10A27786>
- ▶ OA27787: NEW FUNCTION
<http://www-01.ibm.com/support/docview.wss?uid=isg10A27787>

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

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

Symbols

“visibility” features
42

Numerics

3590 20
3592-J1A 18
3958 17–18, 130
3958 configuration 134
3958 DD3 17–18
3958-DD3 100
3958-DE2 29, 92, 107
3958-DE2 frame 31, 93
5.2 Advanced Platform 64-bit 49

A

a la carte 21
ACS routines 160
AIX 17
AIX's EtherChannel support 31
algorithm 86
allocatable space 27
allocated meta data size 121
allocated user data size 121
appliance 21
Automated Tape Library function 28
available file systems 121

B

Backup 42
backup 4
backup administrator 86, 88
backup application 7, 22, 74, 88, 251
 changes 71
 read command 22
 tape cartridges 22
backup application standpoint 40
backup architecture 69
backup data 5, 27, 255
backup operators 86
backup policy 5, 24, 77
backup tiers 69
backup window 8–9
backup policies 20
bandwidth 11
bandwidth 41
baseline 75
beginning of tape (BOT) 22
bit-by-bit comparison 6
block groups 27
BM System Storage TS7680 ProtecTIER 48
BOT 22

Broadcom 101
buffer 78
byte-level 23

C

cache 67
Call Home commands 105–106
Call Home commands submenu 106
Call Home is already enabled 106
Call Home queue 106–107
capacity fluctuations 87
capacity management 87
capacity planning 71
cartridge capacity 87
cartridges 42
change rate 256–257
chunk 5–6, 23
chunks of data 23
Cisco router 103
cluster member 236
cluster member wizard 124
Clustered ProtectTIER Servers 29
clustering 16
clusters 101
command line 105
command queuing 81
Complex Pair Information 104
Complex View 103
compression
 3592 SLDC 163
compression rate 80
CompressionRate 80
conceptualization 27
Concurrent licenced internal code load 29
Concurrent licensed internal code load 29
Concurrent maintenance 29
concurrent maintenance 29
Console Configuration Utility 103, 106
Console Configuration Utility Screen 103
Console Configuration Utility screen 99
Console Configuration utility 96
content aware deduplication 6
Content aware deduplication methods 6
contiguous blocks 27
contiguous space 27
control unit 147–148
Control Unit Online 134
Control Unit Service Utilities 134
control units 29
Controller cache 71
critical cache full condition 21
CU 32–33

D

- daily backups 72
- data center topology 67
- data change
 - rate 4, 24
- Data Change rate 76–77
- data change rate 77, 88
- Data Class 160
- Data deduplication 4, 8, 11
- data de-duplication
 - process 256
 - processing 4
 - ratio 4
- data deduplication 3, 5, 7–8, 11, 25, 68
- data deduplication process 9
- data deduplication software 10
- data deduplication system 10
- Data deduplication systems 9
- data element 4
- data factoring technology 19
- data online 11
- data proliferation 68
- data protection environment, 75
- Data Protection Survey 71, 76–78
- data protection survey 76
- data restoration 42
- data retention 4
- Data Retention period 76–77
- data streams 88
- Data type 72
- DATACLAS
 - compaction 163
 - media type 164
- DB 31
- DB2 29–32, 77
- DB2 database 29, 31
- DD/compression 18
- de-duplicate, 23
- deduplicated disk repository 17
- deduplicating data 5
- Deduplication 4
- de-duplication 4–5
- deduplication 4–5, 15, 24, 41, 66
- Deduplication assumptions 78
- deduplication process 6
- Deduplication Ratio 72
- de-duplication ratio 4, 248
- deduplication ratio 4, 8, 67
- deduplication ratio. 25
- de-duplication ratios 5
- deduplication server 10
- deduplication software 10
- deduplication solution 69
- deduplication system 9
- de-duplication technology 17
- deduplication technology 68
- default gateway 110–111
- default gateway 123
- define
 - DFSMS constructs 159

- Define policies 42
- Defining Data Classes 161
- defragmentation 27
- delta 79
- DFSMS 16
- DFSMS constructs 159
- DFSMSrmm
 - tape initialization 157
- differential 72
- differential incremental backups 4
- differentials 71
- Disk Array 71
- Disk array 71–72
- disk array 10, 71, 73
- disk array manufacturer 73
- disk array subsystem 93
- disk cache 21
- Disk capacity 71
- disk mirroring 84
- disk repository 25
- disk subsystem 21
- disk-to-disk (D2D) 68
- disk-to-disk-to-tape 68
- Dual parity 83
- duplicate data 4, 7, 9, 23
- duplicate items 4
- duplicate volume serial numbers 159
- DVD 130

E

- Early Warning 86
- Early Warning (EW) 86, 88
- early warning (EW) 254
- Early Warning (EW) signal 87
- eject command 21
- EMC CX300 50
- EMC CX600 50
- EMC CX700 50
- Enable Call Home 106
- encryption 73
- encryption features 74
- End of Cartridge 87
- end time 257
- Enhanced Routers 92
- Enterprise Controller 17, 19–21, 28, 32, 92, 94–95
- Enterprise Controller control unit 28
- Enterprise Controller monitor process 32
- Enterprise Controller schematic blocks 28
- Enterprise controller's 94
- Enterprise controller's operator panel 94
- Enterprise Controllers 17, 29, 31, 103, 105
- Enterprise controllers 15, 17, 94
- Enterprise Controllers' 94
- enterprise-class data protection platform 16
- ermmMonitor process 31
- Estimated factoring ratio field 120
- EtherChannel adapter 31
- EtherChannel link 31
- Etherchannel link aggregation 31
- ethernet communication 31

- ethtool 102
- EW 86–88
- EW signal 86
- exact chunk 23
- expired status 23

F

- fabric switches 72
- Factoring Ratio 76
- Factoring ratio 76
- factoring ratio 24, 26, 77, 86–87, 115, 121
- factoring ratio, 67, 76
- failed disk 84
- FC-10K 115
- FC-10K 2+2 120
- fence test 130
- FIBER 30, 33
- Fibre Channel
 - arbitrated-ILoop 250
 - arbitrated-loop 241
- fibre channel communication 32
- Fibre Channel disks 81
- Fibre Channel 32
- Fibre Channel-arbitrated loop (FC-AL) 250
- FICON HBA 70
- FICON host 134
- FICON Host links 134
- Ficon switch 70
- FICON® attach 16
- FICON/FC 17
- file system 27, 245
 - current usage type 245
- File systems 89
- first stripe 84
- frame complex 104
- free blocks 27
- full backup cycle 88
- full backups 4, 71
- full usage condition 21
- FullCapacityVersions 78
- FullFrequency 78–79
- FullRetention 78–79

G

- General Tab 235–239, 248–249, 251–252
- Global File System (GFS) 26, 245
- government regulations 11
- grace period 21

H

- HADR 29, 31–32
- hash 5
- hash algorithm 5
- Hash based deduplication 5–6
- hash collision 6
- hash collisions 6
- hash Fh 5
- hash index 5

- Hash-based data deduplication 23
- hash-based deduplication 23
- hash-based techniques 7, 23
- hashing 5
- hashing algorithm 5
- HBA 17
- HDDs 71
- HDR1 19
- HDS AMS1000 50
- HDS AMS2300 50
- HDS USP-V 50
- Heartbeat communication 31
- heartbeat message 32
- High Availability (HA) 29
- High Availability benefit 29
- High Availability components 30
- Host Adapter Utility Menus 134
- HyperFactor 5, 7–8, 19, 23, 25, 66
- HyperFactor data deduplication 8
- HyperFactor deduplication 8
- HyperFactor Deduplication technology 16
- HyperFactor technology 23
- HyperFactor. 25
- HyperFactor® technology 14
- HyperFactored data 73

I

- I/O 27
- I/O operation 242
- I/O size 82
- I/O write 75
- IBM 48, 68, 100
- IBM Control Unit Subsystem Maintenance menu 134
- IBM Customer Configuration Profile File 105
- IBM disk storage 48
- IBM DS3400 50
- IBM DS4200 50
- IBM DS4300 50
- IBM DS4700 50
- IBM DS4800 50
- IBM DS5020 50
- IBM DS5100 50
- IBM DS5300 50
- IBM DS8100 50
- IBM DS8300 50
- IBM Storage ProtecTIER Deduplication Gateway for System z V1 software 15
- IBM SVC 50
- IBM System p5® 520 (model 9131-52A) 49
- IBM System Service Representative 103, 134–135
- IBM System Service Representative (SSR) 50
- IBM System Storage IBM TS1120 Tape Controller 49
- IBM System Storage IBM TS1120 Tape Controllers 49
- IBM System Storage ProtecTIER Deduplication Gateway servers 49
- IBM System Storage ProtecTIER Enterprise Edition V2.4 software 49
- IBM System Storage ProtecTIER Enterprise Edition V2.4 CD 108
- IBM System Storage ProtecTIER for System z Gateway

- solution 15
- IBM System Storage ProtecTIER for System z V1 version 16
- IBM System Storage TS3000 System Console (TSSC) Maintenance Information 100
- IBM System Storage TS3000 System Console (TSSC) 99
- IBM System Storage TS7650 ProtecTIER Deduplication Appliance 48
- IBM System Storage TS7650G ProtecTIER Deduplication 48
- IBM System Storage TS7680 ProtecTIER Deduplication Gateway for System z 16, 48–49
- IBM System x3850 M2 server (model 7233) 49
- IBM System x3850 M2 Type 7233 18
- IBM technical personnel 71
- IBM TS3000 System Console 50, 95, 103
- IBM TS3000 System Console menu 106
- IBM TS3000 System Console 99
- IBM XIV 50
- IBM XIV, 50
- import/export slot 251
- import/export slots 42
- incremental backup 71
- incremental backup capacity 79
- incremental backups 79
- IncrementalCapacity 79
- incrementalCapacityVersions 79
- IncrementalChangeRate 80
- IncrementalFrequency 79
- IncrementalPhysicalCapacity 80
- IncrementalRetention 79
- index save 25
- init 157
- in-line 9
- In-line data 16
- In-line data deduplication 8
- in-line data deduplication 8
- in-line factoring approach 72
- in-line method 9
- input/output (I/O) 9
- Integrated Removable Media Manager (iRMM) 30
- IOPS 27, 81
- iostat 74
- IP address 234
- iRMM 29
- iRMM 30–32

K

- KRPM 115
- kudzu 101
- KVM 93
- KVM setup 130
- KVM switch 18, 93, 100, 110, 123, 130

L

- LAN 109
- learning algorithm 86–87, 89
- Lempel-Ziv-Haruyasu (LZH) 25

- Library function 29
- Library function status 32
- Library Integration and Emulation (LIE) 29
- Library Manager 13
- Library Manager Integration 31
- Library Manager Integration components 29
- Library Manager Integration functions 28
- Library window 239, 248–250, 253–254
- library wizard 126
- Licensed Internal Code (LIC) 49
- LIE 29, 31
- LIE component 31
- LIE HA aspects 31
- LIE process 31
- Linux desktop, 124
- logical drives 82
- logical unit number (LUN) 238, 249–250, 253
- LOOP 128
- LUNs 73, 80

M

- MAC 101
- mainframe hosts 68
- Management Class 160
- matrix 69
- maximum transmission unit (MTU) 243
- MD Raid configuration field 115, 120
- MD file systems 121
- MEDIA 6 164
- MEDIA1 163
- MEDIA2 163
- MEDIA3 163
- MEDIA4 163
- MEDIA5 164
- MEDIA7 164
- MEDIA8 164
- Memory Resident 8
- Memory Resident Index 8, 25
- memory resident index 25
- memory-resident index 7
- memory-resident index. 23
- Meta Data 22, 25–27, 67, 73, 76
 - file 27
 - file system 25
 - LUNs 27
- Meta data 121
- Meta Data files 27
- meta data files 27
- Meta Data LUNs 73
- MetaData 73
- monitor process 32–33
- monitor processes 32
- Monitoring Process Mechanism 32
- MRPD 105
- multiplexing 73
- Multiplexing features 74

N

- native data 8

- Navigation 113
- Navigation pane 112
- navigation pane 234, 236, 238, 240
 - Nodes tab 236
 - Systems tab 238
- netmask 123
- network power switch 49
- network settings 123
- network subnet mask 110
- node 101
- node failure 29
- nominal capacity 87–88
- nominal cartridge size 88
- NominalCapacity 78
- NTP 99
- NTP server 99

O

- Offline Control 134
- operating system 71
- optimal performance 89
- original backup stream 9
- outboard library manager 29
- outboard library manager function 28

P

- Peer to Peer (P2P) 241, 250
- Performance 14
- performance stabilizes 27
- physical capacity 87
- physical library 26
- physical tape cartridge 22
- physical tape library 22
- PLF 29
- Port attributes 127
- Port Attributes wizard window 127
- Port Details menu 128
- post-processing data deduplication method 9
- post-processing data deduplication period 9
- post-processing method 9
- Power Control Assembly (PCAs) 92
- pre-sales engagement 77
- Primary CU 33
- primary CU 32–33
- primary node 32
- Problem Management Report (PMR) 107
- product engineering (PE) 105
- ProtectTIER 7, 19, 24, 26–27, 66–69, 72, 75, 87, 92, 130, 231–237, 239–246, 248, 250–257
- ProtectTIER deduplication technology 7–8
- ProtectTIER devices 134
- ProtectTIER for meta data 121
- ProtectTIER Functions 108
- ProtectTier hardware 42
- ProtectTIER I/O pattern 73
- ProtectTIER implementation 72
- ProtectTIER Manager 109, 111, 113, 117, 124, 126, 231–243, 245–250, 253–254
 - application 232

- Library window 249–255
- Network Configuration values 243
- Nodes window 241–243
- password dialog box 232
- window 233
- workstation 243
- ProtectTIER Manager 2.3.x 111
- ProtectTIER Manager 2.4.x 124
- ProtectTIER Manager Functions 109, 111, 117, 124
- ProtectTIER Manager graphic user interface (GUI) 107
- ProtectTIER Manager GUI 68
- ProtectTIER Manager screen 117, 130
- ProtectTIER Manager screen 109
- ProtectTIER Manager toolbar 112, 124
- ProtectTIER Manager workstation 111, 117, 124, 130
- ProtectTIER Native Replication 42
- ProtectTIER Native replication 41
- ProtectTIER native replication 40, 42
- ProtectTIER Native Replication Management Interface 42
- ProtectTIER nodes 41, 74
- ProtectTIER Repository 76–78
- ProtectTIER repository 72, 75
- ProtectTIER Server 100, 105
- ProtectTIER server 18, 30, 92, 94, 100, 105, 123, 131, 30
- ProtectTIER server hardware 48
- ProtectTIER server operator panel 94, 123
- ProtectTIER servers 15, 17, 48–50, 93, 99, 123
- ProtectTIER servers. 48
- ProtectTIER Software 23–24, 26–27, 231, 237, 242, 244, 246–248, 253–255
- ProtectTIER Software server 25
- ProtectTIER Software server, 25
- ProtectTIER solution 69
- ProtectTIER system 27, 67, 69, 73, 75–76, 80, 121, 126
- ProtectTIER system sizing 68
- ProtectTIER systems 73–75, 77–78
- ProtectTIER uncompressed 74
- ProtectTIER GUI 108
- ProtectTIER server 94
- ProtectTIER, HyperFactor 23
- ProtectTIER-HyperFactor 23
- ProtectTIER 91, 103
- ProtectTIER devices 130
- ProtectTIER Library 91
- ProtectTIER Manager 126
- ProtectTIER Manager GUI 91
- ProtectTIER Server 32, 101, 105, 110, 123
- ProtectTIER Server status 32
- ProtectTIER Server. 17
- ProtectTIER Servers 29, 101, 105
- ProtectTIER Servers Call Home. 105
- ProtectTIER Servers port configuration 126
- ProtectTIER Virtual Tape (VT) 17
- ProtectTIER VT 17
- pseudo label 19
- PT Manager 123
- PT Manager system 114
- PT Server 91

PT Manager 111
ptadmin 113, 124
ptconfig 130

Q

Qlogic host bus adapters 17

R

RACF profiles 159
ragmentation 8
RAID 67, 72, 89, 110
RAID 0 82, 84
RAID 1 82
RAID 10 84–85
RAID 5 82, 84
RAID 5 array 82
RAID 6 83
RAID groups 73, 83
RAID level 82
RAID levels 81–82
RAID types, 89
RAID0 82
RAID10 84
RAM 23
Random writes 84
RAS 17, 101
RAS command line interface (CLI) 105
RAS package 99, 101
RAS package configuration 101
RAS verification 101
rasMenu 105
rasMenu tool 105
recording technology
 SLDC 163
recovery throughput. 67
recycled cartridge 88
Red Hat cluster, 130
Red Hat Enterprise Linux® 17
Redbooks Web site 312
 Contact us xiv
Redundant Array of Independent Disk (RAID) 27
reference table 26
Reliability Availability and Serviceability (RAS) 91
Reliability, Availability, and Serviceability (RAS) 48, 99
Remote Service Adapter (RSA) 100
replicate data 41
replication window” 42
repository 23, 25–26
Repository Planning 114
Repository Planning wizard 114
Repository size (TB) field 120
restore activity. 8
RETAIN 100
retention period 24, 79
RPM 71, 81
RSA 101
rsCerCfgAll 101
rsCerCHFunction 105
rsCerCHTest 105

rule of thumb 78

S

SAN 72
SAN Back End disk connectivity 72
SAS 72
SATA 81
SATA drives 81
SCDS 160
scrambled workloads 5
scratch category 20
scratch pool 86
scratch pooling 159
Scratch processing delete function 16
scratch storage 29
scratch threshold
 MEDIA1 160
 MEDIA2 160
 MEDIA3 160
scratch write 22
scratching volumes 21
SCSI Enclosure Services (SES) 30
Secure Hash Algorithm 5
seismic data 73
Service Information Message (SIM) 30
Service Representative (SSR) 47
Service Representative (SSR). 103
Services pane 235, 248
 library name 248
SHA-1 6
shelf 40
SIM 31–33
Single disk I/O rate 81
single instance store 19
SLDC 163
 3592 compression algorithm 163
SMTP 100
SMTP IP 100
source control data set
 See SCDS
space reclamation 88
Spin speed 81
standard backup activity, 27
standby 29–30
Standby CU 33
standby CU 33
standby LIE 33
standby mode 28
steady state 26–27
storage block 27
Storage Class 160
Storage Group 160
Storage Interoperation Center 50
storage performance optimization 81
storage requirement 5
Subnet Mask 97
Subsystem Maintenance menu 133
sync 32
sysplanar 99
System Checkout Menus 134

- System Checkout Menus 134
- System Console Actions 131
- System peak 115
- System peak throughput (MB/Sec) field 120
- System Service Representative (SSR) 105
- System Storage TS7680 ProtecTIER Deduplication Gateway for System z (3958-DE2) 15
- System Storage TS7680 ProtecTIER Deduplication Gateway for System z. 47
- System z 70
- System z attachment 49
- System z host attachment 17
- System z hosts 15
- system-managed tape 139
- Systems window 232, 234–240
 - General tab 235–236
 - VT tab 238
- Systems z applications 14
- system-wide nominal capacity 86

T

- takeover 32–33
- takeovers 33
- tape cartridge 22
- tape commands 21
- tape drive 239, 249–252
- tape initialization 157
- tape library 22
- tape library management 87
- tape management paradigm 89
- tape systems 88
- TB 87
- TCP/IP 32
- TCP/IP Communication 32
- TCP/IP failure 33
- TCP/IP Failure 33
- The IBM System Storage ProtecTIER for System z® Gateway (TS7680) 15
- The IBM System Storage TS7680 ProtecTIER Deduplication Gateway for System z 48
- The ProtecTIER Manager GUI 109
- throughput performance 72
- topology 71
- Total Storage System Console (TSCC) 50
- Total Storage System Console (TSSC) 18
- TS3000 System Console 50, 93
- TS3000 System Console (TSSC) 92–93
- TS3000 System Console, 50
- TS7650 ProtecTIER De-duplication appliance 14
- TS7650G 48, 75
- TS7650G Back End 72
- TS7650G ProtecTIER
 - De-duplication Gateway 13
- TS7650G ProtecTIER De-duplication Gateway 13
- TS7650G system 88
- TS7680 14–15, 18–20, 28, 33, 48–50, 65, 68, 70–72, 85–89, 92, 126
- TS7680 handles 87
- TS7680 implementation 72
- TS7680 installation 48

- TS7680 learns 88
- TS7680 ProtecTIER Deduplication Gateway 17
- TS7680 ProtecTIER Deduplication Gateway for System z 15
- TS7680 ProtecTIER servers 49
- TS7680 ProtecTIER Deduplication Gateway 17
- TS7680 repository 77
- TS7680 virtual tape solution 88
- TS7700 20
- TSSC 18, 31, 92–93, 95, 97–100, 103, 105, 131
- TSSC blue screen. 131
- TSSC C 109
- TSSC Call Home Queue 106
- TSSC graphical user interface (GUI) 105
- TSSC Login 133
- TSSC Network Switch 100
- TSSC's 95, 100, 103, 131
- TSSC's keyboard 95, 130
- two-node system 29

U

- UD file systems 121
- uncompressible 4
- uncompressible data 4
- unique data 88
- USB 101
- User Data 24, 26–27, 245–246
- UTC 99
- UTC time 95
- Utilities Menus 134

V

- virtual address 31
- virtual cartridge 22, 26–27, 245, 253
 - prior use 27
- virtual cartridges 42
- Virtual Device 147
- virtual device
 - first set 147
 - second set 147
- virtual IP address 31
- virtual library 22
 - space utilization 22
- virtual robots 17
- virtual scratch tape 27
- virtual tape cartridge 22
- virtual tape cartridges 22
- Virtual Tape Library 29
- virtual tape library 10, 16, 20–22, 68
- Virtual Tape Library with HyperFactor 13
- virtual volume 21–22
- virtual volumes 19
- Visibility Control 40
- vmstat 74
- volume label 19
- volume serial numbers 159
- VT monitoring screen 126
- VT name 126
- VTFD 110–111

VTFD service 123, 130
VTL 22
VTL (Virtual Tape Library) 29
VTL data storage 21
VTL export/import 42
VTL software 21
VTL solution 68
VTLs 21
VTS 147

W

Western Telematic Inc. (WTI) 92
write cache 82, 84
write policy 82

Z

z/OS 20



TS7680 Deduplication ProtecTIER Gateway for System z

(0.5" spine)
0.475" <-> 0.873"
250 <-> 459 pages



Redbooks®

TS7680 Deduplication ProtecTIER Gateway for System z

**Understand the
concepts of data
deduplication**

**Realize bandwidth
savings with the
TS7680**

**Reduce your storage
hardware
requirements**

This IBM® Redbooks publication introduces the IBM System Storage TS7680 ProtecTIER Deduplication Gateway for System z (3958-DE2) hardware and the IBM System Storage ProtecTIER Deduplication Gateway for System z V1.1 software. These are designed to help address the tape processing needs of System z data centers by improving the data protection infrastructure and more cost effectively managing and protecting critical client data.

Managing this growth has become the primary source of pain for storage professionals, who are grappling with the following challenges:

- ▶ Growing storage acquisition and management costs
- ▶ Data processing administration
- ▶ Shrinking batch windows
- ▶ Demanding service levels

The TS7680 helps alleviate these challenges.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-7796-00

ISBN 0738434558