


IBM Virtualization Engine TS7700 with R1.7

New cache options for capacity and performance

Improved disaster recover testing capabilities

Enhanced hybrid multi-cluster grid management



Alex Osuna
Babette Haeusser
Aderson J. C. Pacini
Jorma Salokari

Redbooks



International Technical Support Organization

IBM Virtualization Engine TS7700 with R1.7

March 2011

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

Third Edition (March 2011)

This edition applies to Release 1.7 of IBM Virtualization Engine TS7700

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

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

Contents

Notices	xiii
Trademarks	xiv
Preface	xv
The team who wrote this book	xv
Now you can become a published author, too!	xviii
Comments welcome	xviii
Stay connected to IBM Redbooks	xviii
Summary of changes	xix
March 2011, Third Edition	xix
Part 1. Architecture and planning	1
Chapter 1. Introducing the IBM Virtualization Engine TS7700	3
1.1 Overview	4
1.2 Concepts of storage virtualization	10
1.3 Benefits of tape virtualization	13
Chapter 2. Architecture, components, and functional characteristics	15
2.1 TS7700 Virtualization Engine architecture and terminology	16
2.1.1 TS7700 Virtualization Engine specific terminology	16
2.1.2 Multi-cluster grid terms	23
2.1.3 Tape Virtualization general terms	27
2.2 Architectural capabilities	32
2.3 TS7700 Virtualization Engine hardware components	36
2.3.1 Common components for the TS7720 Virtualization Engine and TS7740 Virtualization Engine models	36
2.3.2 TS7720 Virtualization Engine components	37
2.3.3 TS7740 Virtualization Engine components	44
2.3.4 TS3500 Tape Library	49
2.3.5 IBM System Storage TS1130 Tape Drive	53
2.3.6 IBM System Storage TS1120 Tape Drive	53
2.3.7 IBM TotalStorage 3592-J1A Tape Drive	54
2.4 Integrated Library Manager	54
2.5 TS7700 Virtualization Engine Concepts	56
2.5.1 Buffering data in the Tape Volume Cache	57
2.5.2 Tape Volume Cache Management	57
2.5.3 Logical and stacked volume management	63
2.5.4 Copy Export	70
2.5.5 Encryption	73
2.5.6 Logical WORM (LWORM) support and characteristics	75
2.5.7 Workflow Management Enhancements	76
2.6 TS7700 Virtualization Engine multi-cluster grid	77
2.6.1 Data integrity by volume ownership	77
2.6.2 Copy policy management	81
2.6.3 High availability and disaster recovery: Two-cluster grid	95
2.6.4 High availability and disaster recovery: three-cluster grid	98
2.6.5 Removal of a cluster from a Grid and cluster clean-up	100

2.6.6	Joining of clusters at different code release levels	103
2.7	Expanded Memory	105
Chapter 3.	Preinstallation planning and sizing	107
3.1	Hardware configurations	108
3.1.1	TS7740 Virtualization Engine configuration details.	109
3.1.2	TS7720 Virtualization Engine configuration details.	112
3.1.3	TS3500 Tape Library High Density frame for a TS7740 Virtualization Engine configuration	115
3.1.4	Cables.	116
3.1.5	Limitations.	117
3.1.6	TS7700 Virtualization Engine upgrades	120
3.2	Hardware installation and infrastructure planning	125
3.2.1	System requirements	125
3.2.2	TS7700 Virtualization Engine Release 1.7 Licensed Internal Code levels	127
3.2.3	TCP/IP configuration considerations.	128
3.2.4	LAN/WAN requirements	131
3.3	Remote installations and FICON switch support.	132
3.3.1	Factors that affect performance at a distance.	133
3.3.2	Host attachments	133
3.3.3	FICON Director support	135
3.3.4	FICON channel extenders	135
3.3.5	Implementing cascaded switches	136
3.4	High availability Grid considerations	137
3.5	Planning for software implementation	139
3.5.1	Host configuration definition	139
3.5.2	Software requirements	140
3.5.3	z/OS software environments.	140
3.5.4	Sharing and partitioning considerations	141
3.6	Planning for logical and physical volumes	142
3.6.1	Logical volumes	142
3.6.2	Logical WORM (LWORM).	145
3.6.3	Physical volumes for TS7740 Virtualization Engine	145
3.6.4	Data compression	148
3.6.5	Secure Data Erase	149
3.6.6	Copy Policy Override Settings	151
3.7	Planning for encryption in the TS7740 Virtualization Engine	153
3.7.1	Pool encryption settings for TS7740 Virtualization Engine	154
3.7.2	Encryption Key Manager (EKM)	156
3.7.3	Tivoli Key Lifecycle Manager	156
3.7.4	Tape drives	156
3.7.5	TS7740 Virtualization Engine	157
3.7.6	Encryption Key Manager IP addresses.	157
3.7.7	Encryption key management	159
3.7.8	Installation.	160
3.8	Tape analysis and sizing the TS7700 Virtualization Engine	160
3.8.1	IBM tape tools	160
3.8.2	BatchMagic	163
3.8.3	Workload considerations.	163
3.9	Education and training	167
3.9.1	Adding a TS7740 Virtualization Engine to an existing TS3500 Tape Library . . .	167
3.9.2	Sharing the TS7740 Virtualization Engine with other systems within a TS3500 Tape Library.	167

3.9.3 Implementation services	168
3.10 Planning considerations	169
Part 2. Implementation and migration.	173
Chapter 4. Hardware implementation	175
4.1 TS7700 Virtualization Engine implementation and installation considerations	176
4.1.1 Implementation tasks	176
4.1.2 Installation tasks	177
4.2 TS3500 Tape Library definitions (TS7740 Virtualization Engine only)	178
4.2.1 Defining a logical library	179
4.2.2 Adding drives to the logical library	186
4.2.3 Defining the Encryption Method for the new Logical Library.	188
4.2.4 Defining control path drives	190
4.2.5 Defining Cartridge Assignment Policies (CAP)	191
4.2.6 Inserting TS7740 Virtualization Engine physical volumes.	192
4.2.7 Assigning cartridges in TS3500 Tape Library to logical library partition	196
4.3 Setting up the TS7700 Virtualization Engine.	198
4.3.1 TS7700 Virtualization Engine definition with the management interface	198
4.3.2 Verifying the composite and distributed library sequence numbers	202
4.3.3 Defining VOLSER ranges for physical (stacked) volumes	203
4.3.4 Defining physical volume pools.	206
4.3.5 Defining Fast Ready categories	220
4.3.6 Defining the logical volume expiration time.	222
4.3.7 Operator Interventions	225
4.3.8 Defining TS7700 Constructs	225
4.3.9 TS7700 licensing	234
4.3.10 Defining Encryption Key Manager addresses.	236
4.3.11 Defining SNMP	239
4.3.12 Inserting logical volumes.	240
4.4 Virtualization Engine multi-cluster grid definitions.	243
4.4.1 TTS7700 Grid configuration considerations: two-cluster grid	243
4.4.2 Data access and availability characteristics	244
4.4.3 Defining Grid copy mode control.	245
4.4.4 Defining cluster families	250
4.4.5 Retain Copy Mode	253
4.4.6 Automatic and Enhanced Removal Policies for grids	254
4.4.7 Temporary Removal Threshold.	255
4.4.8 Backup and restore of construct definitions	258
4.4.9 Data management settings (TS7740 Virtualization Engine only)	258
4.5 Implementing Outboard Policy Management for non-z/OS hosts.	260
Chapter 5. Software implementation	263
5.1 Host implementation considerations	264
5.1.1 Sharing the TS7700 Virtualization Engine by multiple hosts.	264
5.1.2 Partitioning the TS7700 Virtualization Engine between multiple hosts	265
5.1.3 Logical path considerations	265
5.1.4 Library names, Library IDs, and port IDs	266
5.2 Hardware configuration definition	269
5.2.1 Defining devices through HCD for Cluster 0	270
5.2.2 Activate the I/O configuration	274
5.2.3 HCD considerations for multi-cluster grid operation	276
5.2.4 More HCD and IOCP examples: two-cluster grid	278
5.2.5 Display and control your settings	279

5.2.6 Set values for the Missing Interrupt Handler	284
5.3 TS7700 Virtualization Engine software definitions for z/OS	285
5.3.1 z/OS and DFSMS/MVS System-managed tape	286
5.3.2 Implementing Copy Export	288
5.4 Software implementation in z/VM and z/VSE	303
5.4.1 General support information	303
5.4.2 z/VM native support using DFSMS/VM	304
5.4.3 Native z/VSE	305
5.4.4 VM/ESA and z/VM guest support	307
5.4.5 z/VSE as a z/VM guest using a VSE Guest Server (VGS)	308
5.5 Software implementation in Transaction Processing Facility	310
5.5.1 Usage considerations for TS7700 with TPF	311
5.5.2 Performance considerations for TS7700 Multi-Cluster Grids with TPF	312
Chapter 6. Upgrade considerations	315
6.1 TS7700 Virtualization Engine component upgrades	316
6.1.1 TS7700 Cache Upgrade options	316
6.1.2 TS7700 Virtualization Engine concurrent system component upgrades	317
6.1.3 TS7700 non-concurrent component upgrades	318
6.2 TS7700 Virtualization Engine upgrade to Release 1.7	319
6.2.1 Planning for the upgrade	319
6.2.2 Upgrading from TS7700 R1.5 or later	320
6.2.3 Upgrading from TS7700 R1.4	320
6.3 Adding clusters to a grid	326
6.3.1 TS7700 Virtualization Engine multi-cluster grid considerations	326
6.3.2 Merge two TS7700 Single Clusters into a TS7700 two-cluster grid	327
6.3.3 Three-cluster grid	330
6.3.4 Four-cluster grid configurations	340
Chapter 7. Migration scenarios	349
7.1 Migration to TS7700 Virtualization Engine	350
7.2 Migration planning from VTS	351
7.2.1 VTS migration procedures and times	351
7.2.2 IBM 3494 Tape Library attachment	352
7.2.3 Licensed Internal Code levels	353
7.2.4 DFSMS definitions	353
7.2.5 HCD definitions	354
7.3 Hardware processes	354
7.3.1 Data format conversion	354
7.3.2 Back up VTS data	355
7.3.3 Restore VTS data	356
7.4 TS7700 Virtualization Engine hardware migration scenarios	357
7.4.1 Migrating a stand-alone VTS to a TS7740 single-cluster grid	357
7.4.2 Merge multiple VTSs to a TS7700 single-cluster grid	364
7.4.3 PTP VTS to TS7700 two-cluster grid	370
7.5 Moving data in and out of the TS7700 Virtualization Engine	376
7.5.1 Moving data into the TS7700 Virtualization Engine: Phased method	377
7.5.2 Moving data into the TS7700 Virtualization Engine: Quick method	378
7.5.3 Products to simplify the task	381
7.5.4 Considerations for static VOLSERS	382
7.5.5 Combining methods to move data into the TS7700 Virtualization Engine	383
7.5.6 Moving data out of the TS7700 Virtualization Engine	383
7.6 Migration of DFSMSHsm-managed data	386

7.6.1	Volume and data set sizes	387
7.6.2	TS7700 Virtualization Engine implementation considerations	390
7.6.3	DFSMSHsm task-related considerations	391
7.7	DFSMSrmm and other tape management systems	395
7.8	IBM Tivoli Storage Manager	396
7.8.1	Guidance for TS7700 Virtualization Engine usage	396
7.8.2	Native drives	398
7.8.3	IBM Tivoli Storage Manager parameter settings	398
7.9	DFSMSdss	398
7.9.1	Full volume dumps	399
7.9.2	Stand-Alone Services	399
7.10	Object Access Method	401
7.11	Database backups	402
7.11.1	DB2 data	402
7.11.2	CICS and IMS	404
7.11.3	Batch data	405
Part 3.	Operation	407
Chapter 8.	Operation	409
8.1	User interfaces	410
8.1.1	TS3500 Tape Library Specialist (TS7740 Virtualization Engine only).	411
8.1.2	Call Home (TS7740 Virtualization Engine with library attachment only).	412
8.1.3	Electronic Customer Care	412
8.2	TS7740 Virtualization Engine management interface	415
8.2.1	Connecting to the MI	415
8.2.2	Using the management interface	418
8.2.3	Health & Monitoring (TS7740 Virtualization Engine with Library attachment).	419
8.2.4	Performance and Statistics windows	435
8.2.5	Logical volumes	435
8.2.6	Physical volumes	460
8.2.7	Constructs	479
8.2.8	Configuration	486
8.2.9	User management	500
8.2.10	Service and troubleshooting	520
8.2.11	Drive cleaning	537
8.3	System-managed tape	540
8.3.1	DFSMS operator commands	540
8.3.2	Library LMPOLICY command	542
8.3.3	Host Console Request function	544
8.3.4	MVS System commands	565
8.4	Basic operations	567
8.4.1	Clock and time setting	567
8.4.2	Library online/offline to host	567
8.4.3	Library in Pause mode	568
8.4.4	Preparing for service	568
8.4.5	TS3500 Tape Library inventory	569
8.4.6	Inventory upload	570
8.5	Tape cartridge management	570
8.5.1	3592 tape cartridges and labels	570
8.5.2	Manual insertion of stacked cartridges	572
8.5.3	Exception conditions	573
8.6	Managing logical volumes	575

8.6.1	Scratch volume recovery for logical volumes	575
8.6.2	Ejecting logical volumes	575
8.7	Messages and displays	577
8.7.1	Console name message routing	577
8.7.2	Alert setting messages	577
8.7.3	Grid messages	578
8.7.4	Display Grid status	579
8.7.5	Warning link status degraded messages	581
8.7.6	Warning VTS operation degraded messages	582
8.7.7	Warning cache use capacity (TS7720 Virtualization Engine)	582
8.8	Recovery scenarios	582
8.8.1	Hardware conditions	582
8.8.2	TS7700 Virtualization Engine software failure	585
8.9	TS7720 Virtualization Engine considerations	586
8.9.1	Management interface (MI)	586
Chapter 9.	Performance and monitoring	589
9.1	TS7700 Virtualization Engine performance characteristics	590
9.1.1	Performance overview	591
9.1.2	Single cluster performance	594
9.1.3	Multi-cluster grid performance	596
9.1.4	Distance and performance	600
9.1.5	TS7740 tape drives and overall cluster performance	602
9.2	TS7700 components and tasks distribution	604
9.2.1	Tasks performed by CPU (processor cycles)	605
9.2.2	Tasks performed by the TS7700 cache	606
9.2.3	Tasks performed by the Grid	606
9.2.4	Tasks performed by the TS7740 Tape Drives	607
9.3	Throttling, tasks, and knobs	607
9.3.1	Throttling in the TS7700	609
9.3.2	Host Write Throttle	609
9.3.3	Copy Throttle	610
9.3.4	Deferred Copy Throttle	612
9.3.5	Managing tasks within TS7700	613
9.4	Data movement through cache	616
9.4.1	Cache Data Flow in a TS7700 two-cluster grid	617
9.4.2	Cache data flow in a TS7700 three-cluster grid	619
9.4.3	TS7700 Four-cluster grid considerations	621
9.4.4	TS7700 hybrid grid considerations	623
9.4.5	Cluster families and cooperative replication	623
9.4.6	Retain Copy Mode and Device Allocation Assist	625
9.5	Monitoring TS7700 Virtualization Engine performance	628
9.6	Using the TS3500 Tape Library Specialist for monitoring	629
9.7	Using the TS7700 Virtualization Engine management interface for monitoring	631
9.7.1	Performance and statistics	632
9.8	Tuning the TS7700	641
9.8.1	Tuning the TS7700	644
9.8.2	Grid links	648
9.9	TS7700 Virtualization Engine statistical data	652
9.9.1	Types of statistical records	652
9.9.2	Collecting and analyzing TS7700 Virtualization Engine statistical records	656
9.10	Bulk Volume Information Retrieval	657
9.10.1	Overview of the BVIR function	657

9.10.2	Prerequisites	660
9.10.3	Request data format	660
9.10.4	Response data format.	662
9.10.5	Interpreting the BVIR response data.	665
9.11	IBM Tape Tools	706
9.11.1	Introduction to IBM Tape Tools	706
9.11.2	Tools download and installation	707
9.11.3	IBM Tape Tools for TS7700 monitoring	709
9.12	Using VEHSTATS for monitoring and reporting	710
9.12.1	VEHSTATS tool overview	711
9.12.2	Running the VEHSTATS jobs	712
9.12.3	VEHSTATS reports.	713
9.13	z/OS commands for monitoring.	718
9.13.1	Display SMS command.	719
9.13.2	Library command	721
9.14	What to look for and where	722
Chapter 10.	Failover and disaster recovery scenarios	725
10.1	TS7700 Virtualization Engine Grid failover principles	726
10.2	Failover scenarios	729
10.2.1	Test configuration	730
10.2.2	Failover scenario 1	731
10.2.3	Failover scenario 2	732
10.2.4	Failover scenario 3	733
10.2.5	Failover scenario 4	734
10.2.6	Failover scenario 5	735
10.2.7	Failover scenario 6	736
10.2.8	Failover scenario 7	737
10.2.9	Failure scenario 9	738
10.3	Planning for disaster recovery.	739
10.3.1	Grid configuration	739
10.3.2	Planning guidelines.	740
10.3.3	Disaster recovery considerations with TS7740 and TS7720	741
10.4	Disaster recovery using Copy Export	742
10.4.1	Planning and considerations for testing Copy Export Recovery	742
10.4.2	Grid considerations for Copy Export.	744
10.4.3	Performing Copy Export Recovery	745
10.4.4	Restoring the host and library environments.	751
10.5	Geographically Dispersed Parallel Sysplex	752
10.5.1	GDPS functions	752
10.5.2	GDPS considerations in a TS7740 Grid configuration	753
10.6	Disaster recovery testing considerations	755
10.6.1	The test environment represents a point in time.	755
10.6.2	Breaking the interconnects between the TS7700 Virtualization Engines	755
10.6.3	Writing data during the test.	756
10.6.4	Protecting production volumes with TS7700 R1.7	757
10.6.5	Protecting production volumes with DFSMSrmm	758
10.6.6	Control of copies.	760
10.6.7	Return-to-scratch processing and test use of older volumes	760
10.6.8	Copies flushed or kept as LRU	762
10.6.9	Ownership takeover	762
10.7	Disaster recovery testing detailed procedures	763
10.7.1	Scenario 1 for a TS7700 two-cluster grid at R1.7	764

10.7.2 Scenario 2 for a TS7700 two-cluster grid Pre-R1.7	767
10.7.3 Scenario 3 for a TS7700 two-cluster grid Pre R1.7.	770
10.7.4 Scenario 4 for a three-cluster grid.	773
10.8 A real disaster	777
Part 4. Appendixes	779
Appendix A. Feature codes	781
Feature code (FC) lists	782
3952 F05 features	782
3494 B10 and B20 VTS, 3957-V06, and 3957-VEA Server features	782
3956-CC6, 3956-CC7, 3956-CC8, and 3956-CS8 Cache Controller features	784
TS7700 Cache Drawer features 3956-CX6, 3956-CX7, and 3956-XS7	784
TS7700 Virtualization Engine feature code descriptions	785
Appendix B. IBM Virtualization Engine TS7700 implementation steps	797
Appendix C. TS3500 and TS7700 checklists	803
TS7700 Virtualization Engine installation worksheets	804
Grid cluster descriptions	805
TS3500 Tape Library configuration information	805
TS3500 Tape Library drive information worksheet	808
Media volume serial range	810
TS7700 Virtualization Engine configuration information	812
TSSC remote support information worksheet	815
Grid local address	816
TSSC Grid configuration information	818
Appendix D. JES3 examples and information	821
JES3 support for system-managed tape	822
First configuration example	826
LDG definitions necessary for the first example	827
Device statements needed for this configuration	827
SETNAME statements needed for this configuration	827
HWSNAME statement needed for this configuration	828
Second configuration example	828
LDG definitions needed for the second configuration example	830
Device statement needed for the second configuration example	831
SETNAME statements needed for the second configuration example	831
High watermark setup name statements for the second example	832
Processing changes	832
JES3/DFSMS processing	833
Selecting UNITNAMEs	834
New or modified data sets	834
Old data sets	834
DFSMS catalog processing	834
DFSMS VOLREF processing	835
Fetch messages	835
JES3 allocation and mounting	835
Multi-cluster grid considerations	836
Appendix E. IBM Virtualization Engine TS7700 statistical record format	839
Introduction	840
Overview	841

Record types and sizes	841
vNode Virtual Device Point-in-Time (PIT) record	842
vNode Adapter Point-In-Time (PIT) record	847
vNode Virtual Device Historical record	851
vNode Adapter Historical record	854
hNode HSM Point-In-Time (PIT) record	858
hNode Grid Point-In-Time (PIT) record	862
hNode HSM Historical record	865
hNode Library Historical record	872
hNode Grid Historical record	885
Appendix F. Sample JCLs	889
BVIR jobs to obtain historical data	890
BVIRHSTS	890
BVIRHSTU	892
BVIRHSTV	893
Additional BVIR reporting	895
Volume Map report	895
Cache Contents report	897
Copy Audit report	897
Volume Status report	897
Physical volume status	899
Physical Volume Status report	900
Physical volume pool status report	901
Physical Volume and Pool Status Report Writer	902
VEHSTATS Reports.	904
Export List Volume sample JCL	909
JCL for TS7700 Virtualization Engine migration scenarios	910
Using EDGUTIL to validate TCDB inconsistencies.	910
IDCAMS example to delete a library definition in TCDB.	910
IDCAMS example to list all entries in TCDB.	910
IDCAMS example for change the TCDB.	911
JCL to change volumes in RMM.	911
REXX EXEC to update the library name.	911
Appendix G. Library Manager volume categories	913
Appendix H. DEVSERV QLIB command	923
Related publications	933
IBM Redbooks publications	933
Other publications	933
Online resources	935
How to get Redbooks.	935
Help from IBM	935
Index	937

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®	MVS™	System x®
AS/400®	Netfinity®	System z10®
CICS®	OS/400®	System z9®
DB2®	Parallel Sysplex®	System z®
DS4000®	POWER®	Tivoli®
DS6000™	RACF®	TotalStorage®
DS8000®	RDN®	VM/ESA®
ESCON®	Redbooks®	WebSphere®
eServer™	Redbooks (logo)  ®	z/OS®
FICON®	RETAIN®	z/VM®
GDPS®	RS/6000®	z/VSE™
Geographically Dispersed Parallel Sysplex™	S/390®	z10™
IBM®	System i®	z9®
IMS™	System p®	zSeries®
	System Storage®	

The following terms are trademarks of other companies:

Java, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

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

Preface

The IBM® Virtualization Engine TS7700 provides tape virtualization for the IBM System z® environment. The IBM Virtualization Engine TS7700 offers improved performance and capacity to help lower the total cost of ownership for tape processing. It introduces a modular, scalable, and high-performing architecture for mainframe tape virtualization. The TS7700 Virtualization Engine is available as the IBM Virtualization Engine TS7740 R1.7, which brings significant enhancements over its previous release.

The TS7700 Virtualization Engine integrates the advanced performance, capacity, and data integrity design of the 3592 Tape Drives, the IBM industry-leading tape technology, with high-performance disks and an advanced IBM System p® server to form a storage hierarchy, managed by robust storage management firmware with extensive self-management capability. It includes functions, such as advanced policy management to control physical volume pooling, cache management, dual copy, dual copy across a Grid network, and copy mode control. The TS7700 Virtualization Engine offers enhanced statistical reporting and a single, standards-based management interface for TS7700 Virtualization Engine management.

This IBM Redbooks® publication covers the TS7740 Virtualization Engine Release 1.7 and the IBM Virtualization Engine TS7720, and includes information about supported tape drives and libraries. IBM Virtualization Engine TS7700 Release 1.7 is the newest level, and it can support the following capabilities:

- ▶ New disk cache controller and expansion frame
- ▶ Expanded memory
- ▶ Selective write protect for disaster recover (DR) testing
- ▶ Enhanced removal policies
- ▶ Cluster joinings of R1.7 into R1.6 Grid

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 (ITSO), San Jose Center.

Alex Osuna is a Project Leader at the International Technical Support Organization, Tucson Center. He writes extensively and teaches IBM classes worldwide on all areas of storage. Before joining the ITSO 3 years ago, Alex worked as a Principal Systems Engineer for the IBM Tivoli® Western Region. Alex has over 32 years in the IT industry, focused mainly on storage hardware and software, related to managing storage. He holds certifications from IBM, Microsoft®, Red Hat, and the Open Systems Group.

Babette Haeusser is a consulting IBM Certified IT Specialist at the Customer Satisfaction Project Office Northwest Europe. She supports customers in all areas of storage, but primarily focuses on IBM tape and virtualization products. Babette joined IBM in 1973 as an Application Programmer. In 1987, she became an MVS™ Systems Engineer and specialized in IBM Storage hardware and software, which she has supported as a Technical Sales Specialist in various job roles since then. Before joining the project office in 2009, Babette was a Project Leader at the ITSO, managing all tape related projects.

Aderson J. C. Pacini works in the Tape Support Group in the IBM Brazil Hardware Resolution Center. He is responsible for providing second-level support for tape products in Brazil. Aderson has an extensive experience servicing a broad range of IBM products. While at the Country Support Center, he has worked with the Large System Group and Storage Group, mostly with direct access storage devices (DASDs) and tapes. He has specialized in tape products since the launch of IBM Virtual Tape Server B16. He has installed, implemented, and supported all the IBM Tape Virtualization Servers from the Virtual Tape Server (VTS) B16 to the TS7700 Virtualization Engine. Aderson joined IBM in 1976 as a Service Representative; his entire IBM career has been in IBM Services.

Jorma Salokari is an IBM Advisory Client Advocate System Services Representative in Canada. He joined IBM in 1985 and provides technical support to IBM System z clients. Jorma has participated in manifold disaster recovery, technology implementation, and data center relocation projects. He is a Guest Instructor at the IBM Technical Training Center, teaching virtualization concepts, implementation, and service strategies as they apply to the TS7700 Virtualization Engine. As a Team Leader, he mentors people who work with complex systems environments. Jorma currently supports a leading edge solutions provider that is an early adopter of the TS7700 Virtualization Engine.



Figure 1 Aderson, Alex, Babette, Jorma

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

Carl Bauske

Jim Fisher

Advanced Technical Skills (ATS), IBM Americas

David Reich

IBM Systems & Technology Group

Wayne Carlson

Dennis Haight

Norie Iwasaki

Khanh M. Ly

Kerri R. Shotwell

Joseph Swingler

Bill Travis

IBM Systems & Technology Group, Systems Hardware Development

Erika Dawson

Antoinette Kaschner

IBM Systems & Technology Group, Systems Software Development

Becky Swingler

IBM Systems & Technology Group, Growth Solutions

John M. Tarby

IBM Software Group, Application and Integration Middleware Software

Antoinette Kaschner

IBM Deutschland Research and Development

Snehal Bhakta

IBM DFSMSHsm Development

Les Geer

PMP Certified z/VM® Project Management

Mike Noel

Qwest

Roman Yusuf

Hydra FICON® Microcode Team Lead

Gary Anna

Virtual Tape Development

Takeshi Sohda

Virtual Tape Server Development, LTO & Enterprise Tape Systems, Systems Development Lab. Yamato IBM Japan. Qwest

Now you can become a published author, too!

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

ibm.com/redbooks

- Send your comments in an email to:

redbooks@us.ibm.com

- Mail your comments to:

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

Stay connected to IBM Redbooks

- Find us on Facebook:

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

- Follow us on Twitter:

<http://twitter.com/ibmredbooks>

- Look for us on LinkedIn:

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

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

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

- Stay current on recent Redbooks publications with RSS Feeds:

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

Summary of changes

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

Summary of Changes
for SG24-7712-02
for IBM Virtualization Engine TS7700 with R1.7
as created or updated on April 12, 2011.

March 2011, Third Edition

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

New information

The following new information applies to R1.7:

- ▶ New disk cache controller and expansion frame
- ▶ Selective write protect for DR testing
- ▶ Cluster joinings of R 1.7 into R 1.6 Grid

Changed information

The following changes apply to R1.7:

- ▶ Enhanced Removal Policies
- ▶ Expanded Memory



Part 1

Architecture and planning

In this part, we introduce the IBM Virtualization Engine TS7700 Release 1.7 family, consisting of the IBM Virtualization Engine TS7740 R1.7, and the disk-only virtualization solution, the IBM Virtualization Engine TS7720. In addition to a detailed technical description of the architecture and components of the TS7700 Virtualization Engine, we provide you with all required information to plan for the implementation of TS7720 Virtualization Engine and TS7740 Virtualization Engine in your environment.

This part contains the following chapters:

- ▶ Chapter 1, “Introducing the IBM Virtualization Engine TS7700” on page 3
- ▶ Chapter 2, “Architecture, components, and functional characteristics” on page 15
- ▶ Chapter 3, “Preinstallation planning and sizing” on page 107



Introducing the IBM Virtualization Engine TS7700

The TS7700 Virtualization Engine, introduced in 2006, is the fourth generation of IBM Tape Virtualization products for mainframes and replaces the highly successful IBM TotalStorage® Virtual Tape Server (VTS). In June 2010, IBM made available the TS7700 Virtualization Engine Release 1.7, which adds new disk controllers. Release 1.7 also introduces larger disk drives and a disk cache expansion frame, and memory expansion of the 8 GB main processor RAM to 16 GB.

Release 1.7 also introduces several usability and management features:

- ▶ A new Release 1.7 cluster may join a Release 1.6 gird
- ▶ Selective write-protect for disaster recovery
- ▶ Enhanced cache removal policy
- ▶ Encryption key management now has default key support
- ▶ Improved management interface
- ▶ Reliability and serviceability enhancements

1.1 Overview

Figure 1-1 shows the TS7700 Virtualization Engine.



Figure 1-1 TS7700 Virtualization Engine

The purpose of this publication is to explain, in detail, the new features introduced with the TS7700 Virtualization Engine Release 1.7 and all the concepts associated with it.

The TS7700 Virtualization Engine is a modular, scalable, and high performing architecture for mainframe tape virtualization and incorporates extensive self-management capabilities consistent with IBM Information Infrastructure initiatives. These capabilities can improve performance and capacity to help lower the total cost of ownership for tape processing and avoid the potential impact of human errors. A TS7700 Virtualization Engine can help improve the efficiency of mainframe tape operations by efficiently using disk storage, tape capacity, and tape speed, and by providing a large number of tape addresses.

The TS7700 Virtualization Engine uses outboard policy management functions to manage cache and volume pools, control selective dual copy, dual copy across a Grid network, copy mode control, encryption, and Copy Export. It also includes a new standards-based management interface and enhanced statistical reporting.

TS7700 Virtualization Engine provides tape virtualization for IBM System z servers, and can help you satisfy several major needs in your environment, such as reliability, the reduction of the window needed for the backup and restore process, the reduction of services downtime that are caused by physical tape drives and library outages, the reduction of the impact to the operational staff that manages daily backup and restore processing, and infrastructure simplification through reduction of the number of physical tape libraries, drives, and media.

TS7700 Virtualization Engine Release 1.7 is the newest level, and can deliver support for the following capabilities:

► New disk controllers

New 3956 CC8 (FC) and CS8 (SATA) models

► Larger disk drives

- Fibre Channel (TS7740 with tape back-end) 600 GB DDMs replace current 300 GB drives.
- SATA (TS7720 disk-only) 2 TB DDMs replace current 1 TB DDMs

► Disk Cache Expansion Frame TS7720 disk-only configurations; maximum raw capacity increases are as follows:

- From 70 TB to 444 TB for a single node
- Up to 888 TB for a two-cluster grid
- Up to 1332 TB for a three-cluster grid
- Up to 1776 TB for a four-cluster grid

► Selective write-protect for disaster recovery (DR) testing

To prevent the DR host test applications from overwriting any production data, the TS7700 can be put into a Write Protect Mode (Figure 1-2 on page 6), which prevents *any* data in the TS7700 from being created or changed. However, customers must be able to create new test data too, ideally targeting the same TS7700. This situation creates a conflict because some volumes must only be read, others must be read/write. Therefore, R1.7 introduces a function that is referred to as *selective write protect for disaster recovery (DR)*, which offers the following features:

- Management interface extension of cluster write protect
Certain categories can be excluded from cluster scope write protection.
- Can be used by customers for DR testing:
 - Customers can simulate a DR scenario.
 - DR test site can read production data while still being able to write DR test volumes within preconfigured categories.
 - DR test is prevented from modifying production volumes.
 - Fast-ready characteristics can be ignored by the DR host.
 - Cache attributes for a scratch category provide performance and capacity information.




Figure 1-2 Write-protect mode

The TS7740 Virtualization Engine and TS7720 Virtualization Engine are members of the TS7700 Virtualization Engine family. Most functionality of the family applies to both models, but several specific functions exist for each model. In this book, when we refer to a TS7700 Virtualization Engine, we are referring to functions that are available for both TS7740 Virtualization Engine and TS7720 Virtualization Engine. If we refer to a specific model, TS7740 Virtualization Engine or TS7720 Virtualization Engine, we do so because the functionality is specific for that model. Table 1-1 lists this nomenclature detail: what we refer to and what that engine applies to.

Table 1-1 TS7700 Virtualization Engine nomenclature for this book

Referring to	Applies to
TS7700 Virtualization Engine	TS7740 Virtualization Engine and TS7720 Virtualization Engine
TS7740 Virtualization Engine	TS7740 Virtualization Engine only
TS7720 Virtualization Engine	TS7720 Virtualization Engine only

The TS7720 Virtualization Engine further broadens the IBM storage solutions portfolio. With the TS7740 Virtualization Engine improvements, IBM continues delivering solutions to satisfy customer requirements in an environment where data is increasing in an exponential way, and data management is needed with the lowest total cost of ownership in the hardware devices.

The TS7700 Virtualization Engine creates a storage hierarchy through the integration of the following components:

- For the TS7740 Virtualization Engine (see Table 1-2):
 - One TS7740 Virtualization Engine Server Model V06
 - One TS7740 Virtualization Engine Cache Controller Model CC8
600GB/15K RPM FC 4Gbps disk drives go into CC8
 - Zero, one, or three TS7740 Virtualization Engine Cache Drawers Model CX7s
 - Connection to one IBM System Storage® TS3500 or IBM System Storage 3494 Tape Library

Note: Attachment to a 3494 Tape Library is not formally supported by code levels of 8.6.0.xx or later. Support for 3494 Tape Library attachment is available only through a request for price quotation (RPQ).

- Connection to high-capacity, high-performance 3592 Tape Drives, including the new IBM System Storage TS1130 Tape Drive

Table 1-2 TS7740 Cache configurations and capacity

Number of TS7740 Cache Controllers 3956-CC8	Number of TS7740 Cache Drawers 3956-CX7	Usable capacity
0	0	3.44 TB
		6.87 TB
		13.74 TB
		9 TB
	1	10.48 TB
	3	24.57 TB
	2	20.96 TB
1	0	7.04 TB
1	1	14.09 TB
1	3	28.17 TB

- For the TS7720 Virtualization Engine (Table 1-3):
 - One TS7720 Virtualization Engine Server Model VEA
 - One TS7720 Virtualization Engine SATA Cache Controller Model CS7

Table 1-3 Base TS7720 Cache configurations and capacities

Number of TS7720 Cache Controllers (3956-CS8)	Number of TS7720 Cache Drawers (3956-XS7)	Usable capacity in TB
1	0	19.84 TB
	1	43.68
	2	67.52
	3	91.35
	4	115.19
	5	139.03
	6	162.87

Important characteristics of these components are summarized as follows:

- The TS7700 Virtualization Engine Server provides host connections for up to four FICON channels for attachment to IBM System z servers with the appropriate levels of System z software.
- In a TS3500 Tape Library, the TS7740 Virtualization Engine supports 4 - 16 IBM 3592 tape drives.
- In a 3494 Tape Library, the TS7740 Virtualization Engine supports 4 - 12 IBM 3592 tape drives.
- TS7720 Storage Expansion frame: Release 1.7 introduces the TS7720 Storage Expansion frame, which significantly increases the total capacity of the TS7720 solution. The maximum usable capacity of a TS7720 base frame with an attached expansion frame is approximately 441 TB using Model CS8 and Model XS7 cache modules containing 2 TB SATA HDDs. The previous maximum capacity was approximately 70 TB.

The TS7720 Storage Expansion frame must contain two TS7720 Cache Controllers (3956 model CS8), see Table 1-4 on page 9. The expansion frame may optionally contain up to ten TS7720 Cache Drawers (3956 Model XS7). The optional cache drawers may be installed in the manufacturing location (plant) or in the field.

Users who purchased a TS7720 base frame prior to Release 1.7 might want to increase its storage capacity. These users may add the TS7720 Storage Expansion frame in the field to increase capacity to 376 TB or 346 TB, depending on whether the existing TS7720 is configured with 40 TB or 70 TB.

Table 1-4 Capacity options

Cache configuration in a new TS7720 Virtualization Engine	Additional TS7720 Storage Expansion frame cache controllers (3956-CS8)†	Total TS7720 cache units (including TS7720 Base Frame)‡	Usable capacity (when TS7720 Base Frame contains only 2 TB drives)
1 TS7720 Cache Controller (3956-CS8) 6 TS7720 Cache Drawers (3956-XS7)	2	9	202.55 TB
		10	226.39 TB
		11	250.22 TB
		12	274.06 TB
		13	297.90 TB
		14	321.74 TB
		15	345.57 TB
		16	369.41 TB
		17	393.25 TB
		18	417.08 TB
		19	440.92 TB

Note: Attachment to a 3494 Tape Library is not formally supported by code levels of 8.6.0.xx or later. Support for 3494 Tape Library attachment is available only by RPQ.

- ▶ The TS7720 Virtualization Engine does not attach to any tape library or drives.
- ▶ A TS7700 Virtualization Engine with Grid Enablement features must be installed on all TS7700 Clusters that operate in a grid configuration.
- ▶ Each TS7700 Virtualization Engine supports up to a maximum of 256 3490E virtual tape drives.
- ▶ A Grid (TS7700 Virtualization Engine single or multi-cluster) supports up to 1,000,000 logical volumes, each logical volume having a maximum capacity of 1.2 GB to 12 GB (assuming 3:1 compression ratio and using the 400 - 4000 MB volume sizes).

For more than 1,000,000 logical volumes up to 2,000,000 you can submit an RPQ. An RPQ is also available for 6 GB volumes.

Figure 1-3 shows the main components of the TS7700 Virtualization Engine for each model.

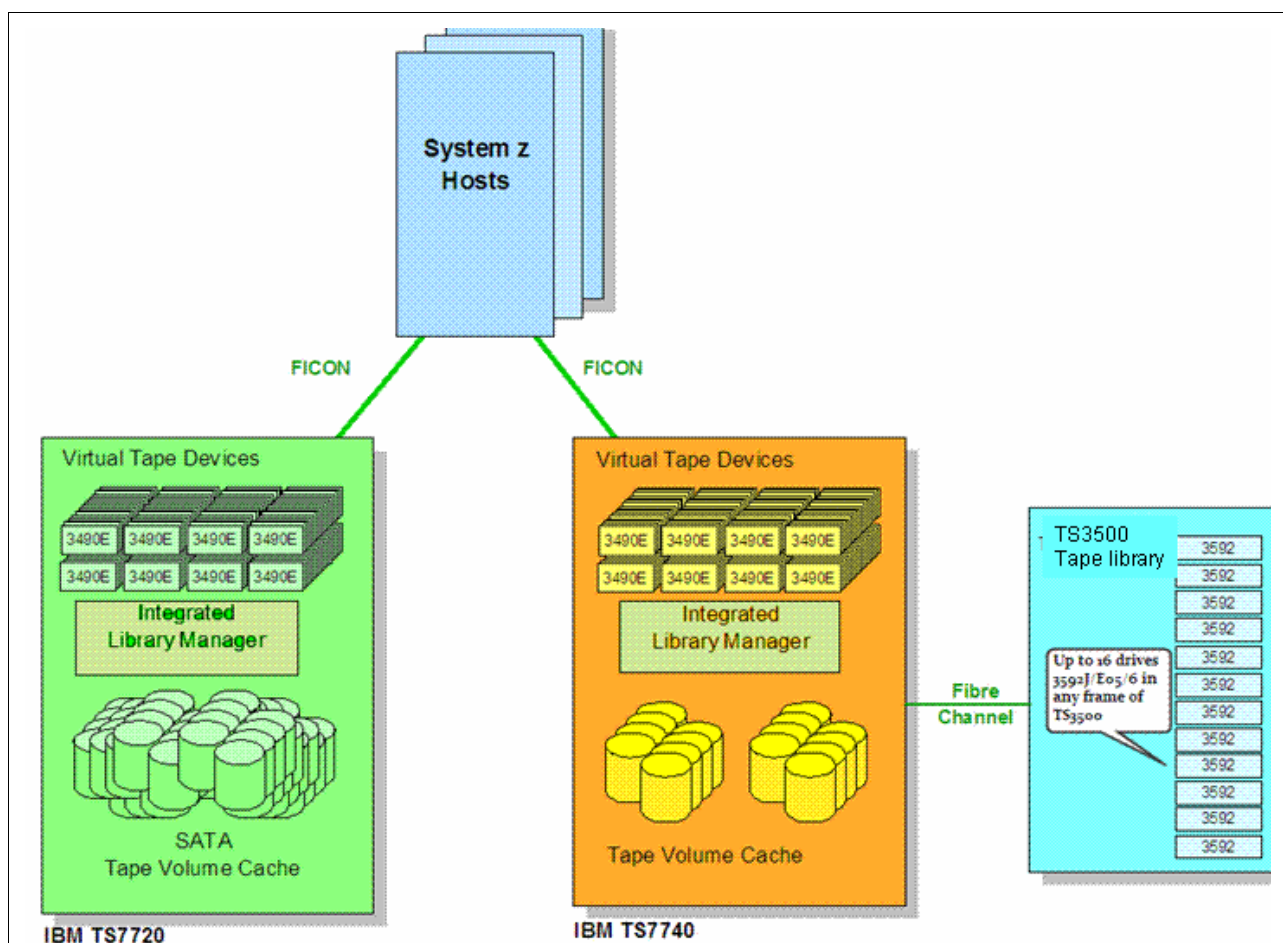


Figure 1-3 Main components of the TS7700 Virtualization Engine

1.2 Concepts of storage virtualization

A virtual tape subsystem presents emulated tape drives to the host and stores tape data on emulated tape volumes in a disk-based cache rather than on physical tape media. The TS7740 Virtualization Engine emulates the function and operation of IBM 3490 Enhanced Capacity (3490E) tape drives and uses a RAID 5 disk subsystem to store volumes written by the host. The disk space provided is called a *Tape Volume Cache (TVC)*.

Emulated tape drives are also called *virtual drives*. To the host, virtual 3490E tape drives look exactly the same as physical 3490E tape drives. Emulation is transparent to the host and to applications. The host always writes to and reads from virtual tape drives; it never accesses the physical tape drives in the back end in TS7740 Virtualization Engine configurations. In fact, it does not need to know that these tape drives exist. As a consequence, even an application that only supports 3490E tape technology can use the TS7700 Virtualization Engine without any changes to the application and still benefit from high capacity and high performance tape drives in the back end. For TS7720 Virtualization Engine configurations, no physical tape attachment exists, but virtual tape drives work the same for the host.

Tape virtualization provides you with a large number of virtual devices, up to 256 per TS7700 Virtualization Engine Cluster. When jobs are contending for native attached physical drives,

tape virtualization can solve this issue by providing a large number of virtual devices for concurrent use.

Because the host exclusively accesses the virtual tape drives, all data must be written to or read from emulated volumes in the disk-based Tape Volume Cache (TVC). We call these emulated tape volumes in the Tape Volume Cache *virtual volumes*.

When the host requests a volume that is still in cache, the volume will be virtually mounted. No physical mount is required. After the virtual mount is complete, the host can access the data at disk speed. Mounting of scratch tapes is also virtual and does not require a physical mount.

Although you define maximum sizes for your volumes, a virtual volume takes up just the space in cache that the data on the volume actually requires. For this reason, tape virtualization makes efficient use of disk capacity. In TS7740 Virtualization Engine configurations, the virtual volumes are copied from disk to tape, and they will also need only the amount of tape capacity occupied by the data, also making efficient use of disk and tape capacity.

Another benefit you get from tape virtualization is the large number of drives available to applications. Each TS7700 Virtualization Engine provides you with a maximum of 256 virtual tape devices. If you have a situation where applications are contending for tape drives and jobs have to wait because no physical tape drive is available, tape virtualization can efficiently address these issues by providing a large number of virtual tape drives.

The TS7740 Virtualization Engine manages the physical tape drives and physical volumes in the tape library and controls the movement of data between physical and virtual volumes.

In the TS7740 Virtualization Engine, data that is written from the host into the tape volume cache is scheduled for copying to tape at a later point in time. The process of copying data to tape that still exists in cache is called *premigration*. When a volume has been copied from cache to tape, the volume on the tape is called a *logical volume*. A physical volume can contain a large number of logical volumes. The process of putting several logical volumes on one physical tape is called stacking, and a physical tape containing logical volumes is therefore referred to as a *stacked volume*. This concept does not apply to TS7720 Virtualization Engine, because no physical tape is attached to it.

Because many applications are not able to fill the high capacity media of today's tape technology, you can end up with a large number of underutilized cartridges, wasting a lot of space on your physical media and requiring an excessive number of cartridge slots in your tape automation system. Tape virtualization reduces the space required by volumes and fully utilizes the capacity of current tape technology. Tape virtualization enables you to exploit the full potential of modern tape drive and tape media technology without making any changes to your applications or JCL.

When space is required in the TVC for new data, volumes that already have been copied to tape will be removed from the cache. By default, removal is based on a *least recently used* (LRU) algorithm to ensure that no new data or recently accessed data is removed from cache. The process of copying volumes from cache to tape and afterwards deleting them is called *migration*. Accordingly, volumes that have been deleted in the cache, and now only exist on tape, are called *migrated volumes*. In a TS7720 Virtualization Engine configuration, there are no migrated volumes, because there is no physical tape attachment. Instead, logical volumes are maintained in disk until they expire; for this reason, cache capacity for the TS7720 Virtualization Engine is bigger than the TS7740 Virtualization Engine cache capacity.

With the four-cluster hybrid grid configuration implemented by R1.6, in a TS7720 Virtualization Engine, which is a member of a hybrid grid, virtual volumes in the TS7720

Virtualization Engine cache can be automatically removed from cache by a Volume Removal Policy, if another valid copy exists elsewhere in the grid (TS7740 Virtualization Engine or another TS7720 Virtualization Engine); see 9.4.4, “TS7700 hybrid grid considerations” on page 623. Volume removal from a TS7720 Virtualization Engine policy will remove volumes based on the *least recently used* (LRU) algorithm at the R1.6 and later Licensed Internal Code level.

Going back to the TS7740 Virtualization Engine, when the host needs to access a volume that has previously been migrated, the volume has to be copied back from tape into the TVC, because the host has no direct access to the physical tapes. When the complete volume has been copied back into the cache, the host is able to access the data on this volume. The process of copying data back from tape to the Tape Volume Cache is called *recall*.

Figure 1-4 shows Tape Volume Cache processing.




Figure 1-4 Tape Volume Cache processing

Another benefit of tape virtualization is the data replication functionality. Two, three or four TS7700 Virtualization Engines can be interconnected through two 1 Gb Ethernet links to form a *multi-cluster grid configuration*. Logical volume attributes and data are replicated across the clusters in a Grid. Any data replicated between the clusters is accessible through any other cluster in a Grid configuration. Through remote volume access, you can reach any virtual volume through any virtual device, even if a replication has not been made.

By setting up policies on the TS7700 Virtualization Engines, you define where and when you want to have multiple copies of your data. You may also specify, for certain kinds of data, for example, test data, that you do not need a secondary or tertiary copy.

Starting with R1.6, you can group clusters within a grid into families, to make improved decisions for tasks, such as replication or tape volume cache selection.

Depending on the configuration, multiple TS7700 Virtualization Engines forming a grid provide you with a high availability, disaster recovery, or high availability and disaster recovery solution.

A multi-cluster grid configuration presents itself to the attached hosts as one large library with a maximum of 512 virtual devices for a two-cluster grid, 768 virtual tape devices for a three-cluster grid, or 1024 virtual devices for a four-cluster grid. The copying of the volumes in a grid configuration is handled by the clusters and is completely transparent to the host. From R1.6 on, the possibility exists to configure a two-, three-, or four-cluster grid, intermixing both TS7720 Virtualization Engine and TS7740 Virtualization Engine models of TS7700 Virtualization Engine.

For TS7740 Virtualization Engine Grid configurations, each Virtualization Engine TS7740 Virtualization Engine manages its own set of physical volumes and maintains the relationship between logical volumes and the physical volumes on which they reside.

In a multi-cluster grid configuration, it is now possible to remove one cluster from the grid in case you no longer require one of the TS7700 Virtualization Engines to participate in the grid. When removal of a cluster from a grid is required, the removed cluster will be disabled and cannot be returned to active use until a cluster cleanup process has been performed.

1.3 Benefits of tape virtualization

Today's global marketplace is increasingly information-oriented with far-reaching implications for businesses. The ability to securely use information can create a competitive advantage. Conversely, the inability to securely use information creates business risk. Four information-related business trends are causing an explosion of information and complexity in data centers, stressing already overtaxed infrastructures and staffs. These four information trends are as follows:

- ▶ Information availability requirements are increasing.
- ▶ Information security threats and privacy regulations are increasing.
- ▶ Information compliance is more complex and penalties are more severe.
- ▶ Information retention periods are longer, often exceeding the life of the storage media.

IBM, aligned with customer requirements, offers an extraordinary range of systems, storage, software, and services that are based on decades of innovation, all designed to help customers to get the right information to the right person at the right time while managing challenges such as exploding data growth, new applications, dynamic workloads, and new regulations.

IBM Information Infrastructure is a means of intelligently storing, retrieving, protecting, and distributing information to help clients achieve competitive advantage. Our clients are converting data centers to service-oriented architectures, for example, which means multiple service levels will be identified and supported, including information services. Certain information services must be high speed to support websites and databases. Sometimes, information services must be multiplexed to multiple locations, or require extra encryption and overwrite protection. IBM Information Infrastructure helps clients apply the right services and service levels, so vital information can be delivered where and when it is needed.

IBM Information Infrastructure solutions are designed to help customers manage the information explosion and address challenges regarding information compliance, availability, retention, and security (CARS). This approach can help your company move toward improved productivity and reduced risk without driving up costs.

The TS7700 Virtualization Engine is part of the IBM Information Infrastructure, a strategy to deliver information availability supporting continuous and reliable access to data, and information retention supporting responses to legal, regulatory, or investigatory inquiries for information.

The TS7700 Virtualization Engine can be the answer to growing storage requirements, shrinking backup windows and continuous access to data.

The following list is a short summary of the main benefits you can expect from tape virtualization:

- ▶ Brings efficiency to the tape operation environment.
- ▶ Reduces batch window.
- ▶ Provides high availability and disaster recovery configurations.
- ▶ Provides fast access to data through caching on disk.
- ▶ Provides utilization of current tape drive, tape media, and tape automation technology.
- ▶ Provides the capability of filling high capacity media to 100%.
- ▶ Provides a large number of tape drives or concurrent use.
- ▶ Provides data consolidation, protection, and sharing.
- ▶ Requires no additional software.
- ▶ Reduces total cost of ownership (TCO).



Architecture, components, and functional characteristics

In this chapter, we provide a detailed description of the architecture of the IBM Virtualization Engine TS7700, general virtualization concepts, new concepts and functionality introduced with TS7700 Virtualization Engine R1.7, and hardware components and configuration scenarios for high availability solutions.

We cover the following topics:

- ▶ Terms and expressions that are used throughout this book to describe the TS7700 Virtualization Engine
- ▶ Architecture of the TS7700 Virtualization Engine and the architecture's potential for future development
- ▶ Hardware components for TS7700 Virtualization Engine Release 1.7
- ▶ Attachment of the IBM Virtualization Engine TS7740 to an IBM System Storage TS3500 Tape Library
- ▶ Logical WORM support
- ▶ Four-cluster grid configuration
- ▶ User security and user access improvements
- ▶ Grid network load balancing improvements
- ▶ IBM Virtualization Engine TS7720/IBM TS7740 Virtualization Engine hybrid grid
- ▶ Immediate copy failure reporting on rewind/unload response
- ▶ Underlying concepts of tape virtualization within the TS7700 Virtualization Engine
- ▶ Functional characteristics of the TS7700 Virtualization Engine
- ▶ Host override for cache residency
- ▶ Device allocation assist
- ▶ Expanded memory

2.1 TS7700 Virtualization Engine architecture and terminology

The TS7700 Virtualization Engine consists of several primary components that can be described in terms of their system role or according to their technical functions. In this section, we review these components to help you better understand the TS7700 Virtualization Engine architecture, and to refresh the concepts of tape virtualization.

If you are familiar with the IBM Virtual Tape Server, the predecessor of the TS7700 Virtualization Engine, many of the terms in this chapter will be familiar to you. Some of the terms have not changed, some will have a new definition, while others were introduced with the TS7700 Virtualization Engine. To avoid confusion, we point out similarities and differences to the previous virtualization solution wherever it seems appropriate.

2.1.1 TS7700 Virtualization Engine specific terminology

The TS7700 Virtualization Engine is built on a distributed node architecture, and new terminology has been introduced. In this section, we explain the new architectural terminology.

Nodes

Nodes are the most basic components in the TS7700 Virtualization Engine architecture. A node has a different name depending on the role associated with it. There are three types of nodes: Virtualization Node, Hierarchical data storage management Node, and General Node.

Virtualization Node (vNode)

The *vNode* refers to a code stack that presents the virtual image of a library and drives to a host system. When the TS7700 Virtualization Engine is attached as a virtual tape library, the vNode receives the tape drive and library requests from the host and processes them as real devices would. The vNode then translates the tape requests through a virtual drive and uses a file in the cache subsystem to represent the virtual tape image. After it is created or altered by the host system through a vNode, a logical volume resides in disk cache.

Hierarchical data storage management (hNode)

The *hNode* refers to a code stack that performs management of all logical volumes residing in disk cache or physical tape after it has been created or altered by the host system through a vNode. The hNode is the only node that is aware of physical tape resources and the relationships between the logical volumes and physical volumes. It is also responsible for any replication of logical volumes and their attributes across site boundaries. An hNode uses standardized interfaces (TCP/IP) to communicate with external components.

General Node (gNode)

A *gNode* can be considered a vNode and an hNode sharing the same physical controller.

Based on the node architecture, a vNode or hNode can run on separate virtualization hardware or be combined to run on the same hardware. When a vNode and hNode are combined to run on the same virtualization hardware, they are referred to collectively as a General Node, or gNode. The TS7700 Virtualization Engine runs on a gNode.

Figure 2-1 shows a relationship between nodes.




Figure 2-1 Node architecture

Cluster

The TS7700 Virtualization Engine Cluster combines the TS7700 Virtualization Engine server with a disk subsystem, the TS7700 Virtualization Engine Cache Controller. This architecture permits expansion of cache capacity and the addition of nodes in future offerings to enhance the capabilities of the system.

Figure 2-2 shows the TS7700 Virtualization Engine configured as a cluster.




Figure 2-2 TS7700 Virtualization Engine cluster

A TS7700 Virtualization Engine cluster provides FICON host attachment and 256 virtual tape devices. The TS7740 Virtualization Engine cluster also includes the assigned TS3500 Tape Library partition, fiber switches, and tape drives. With TS7720 Virtualization Engine R1.7 and later, the cluster can include an optional 3592 cache expansion frame.

Figure 2-3 shows the components of a TS7740 Virtualization Engine cluster.

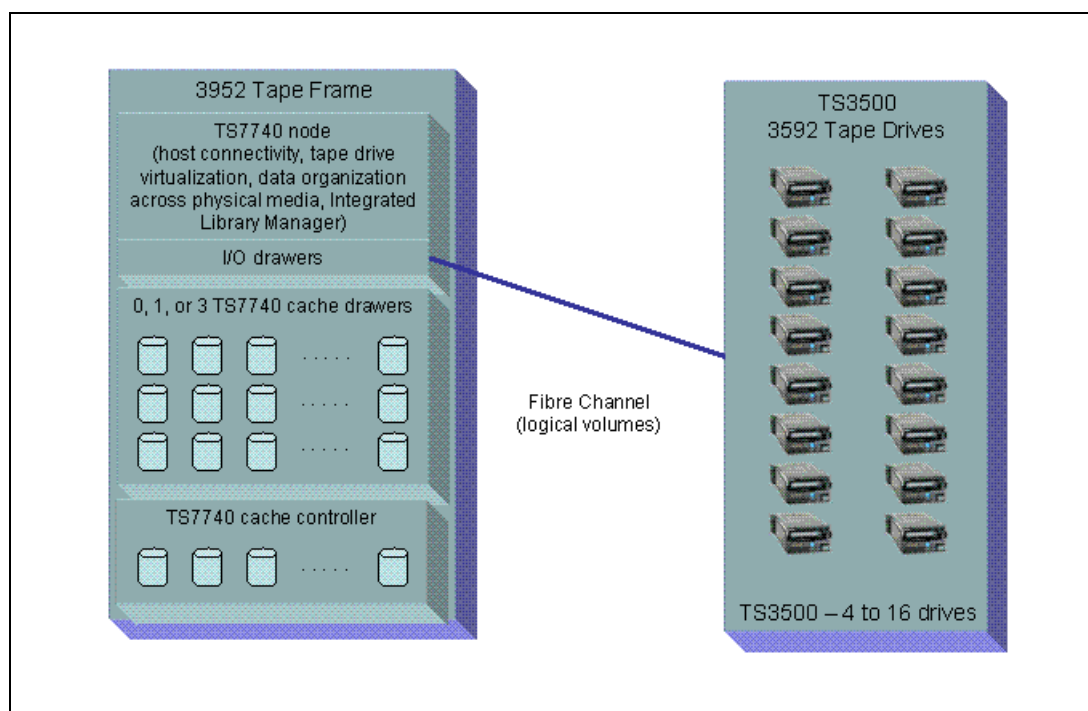


Figure 2-3 TS7740 Virtualization Engine cluster components

You may expand the TS7740 Virtualization Engine cache subsystems with CC8/CX7s to a maximum of four cache modules. This maximum configuration consists of one Model CC8 Cache Controller with three Model CX7 Cache Drawers.

With TS7720 Virtualization Engine R1.7 and later, cache subsystems can be configured with a maximum of three CC8/XS7 elements with the optional Storage Expansion Frame. The TS7720 Virtualization Engine Base Frame must be configured with one Model CC8 Cache Controller and six Model XS7 Cache Drawers before the TS7720 Storage Expansion Frame can be implemented. The maximum configuration of the TS7720 Storage Expansion Frame will have two Model CC8 Cache Controllers and ten Model XS7 Cache Drawers. The minimum expansion frame will have two CS8 cache controllers.

Grid

The TS7700 Virtualization Engine R1.7 Grid, or *Grid configuration*, is a series of two, three, or four clusters that are connected by a network to form a disaster recovery or high availability solution. For high data availability, the logical volume attributes and data are replicated across the clusters, joined by the *Grid network*. This data replication ensures the continuation of production work should a cluster become unavailable. Any data that is replicated between the clusters is accessible through any of the available clusters in a Grid configuration. Using the remote file access capabilities of the Grid network, a logical volume can be mounted through any virtual device.

A Grid configuration looks like a single storage subsystem to the hosts. Whether a single- or multi-cluster configuration, the entire subsystem appears as a single tape library to the attached hosts. This can be described as a composite library with underlying distributed libraries. The distributed libraries are not seen from the host and are similar in concept to the IBM TotalStorage Virtual Tape Server (VTS) Peer-to-Peer (PTP) VTS. Distributed and

Composite Libraries are explained in more detail in 2.1.2, “Multi-cluster grid terms” on page 23.

Several types of grids can be configured:

Single-cluster grid	A single-cluster grid is a single TS7700 Virtualization Engine. The TS7700 Virtualization Engine architecture has the Grid functionality built in. Although the single-cluster grid is not actively participating in a Grid, it is prepared to participate in a multi-cluster grid.
Two-cluster grid	A two-cluster grid consists of two TS7700 Virtualization Engine Clusters, attached across a LAN or WAN through the 1-Gbps Ethernet links. This configuration is also known as a <i>dual-cluster grid</i> .
Three-cluster grid	A three-cluster grid consists of three TS7700 Virtualization Engine Clusters, attached to each other across a LAN or WAN through the 1-Gbps Ethernet links. Each cluster can access the other two clusters.
Four-cluster grid	A four-cluster grid consists of four TS7700 Virtualization Engine Clusters, attached to each other across a LAN or WAN through the 1-Gbps Ethernet links. Each cluster can access the other three clusters.
Multi-cluster grid	A multi-cluster grid describes a Grid with two or more clusters.
Hybrid grid	A hybrid grid describes a multi-cluster grid with an intermix of TS7720 Virtualization Engine and TS7740 Virtualization Engine clusters.

Multiple TS7700 Virtualization Engine Grid configurations can be attached to host systems and operate independently of one another. Currently, a TS7700 Virtualization Engine R1.7 multi-cluster grid consists of two, three, or four clusters of various models. The architecture is designed to allow for more than four clusters to participate in a multi-cluster grid configuration. Five and six-way grids can be configured through RPQ.

Figure 2-4 shows a two-cluster grid configuration. Each cluster contains the TS7740 Virtualization Engine and a tape library with tape drives. The tape library in each cluster is a TS3500 Tape Library.




Figure 2-4 TS7740 Virtualization Engine two-cluster grid

Figure 2-5 shows a three-cluster grid. As with the two-cluster grid, each cluster contains a TS7720 Virtualization Engine.

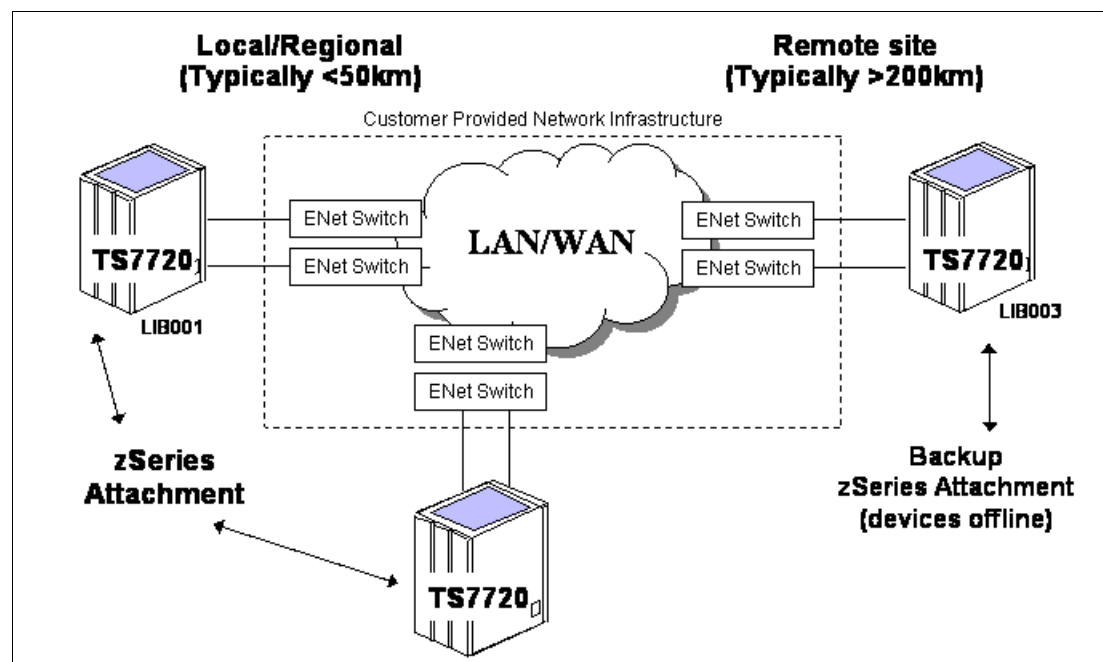


Figure 2-5 TS7720 Virtualization Engine three-cluster grid

Figure 2-6 shows a four-cluster hybrid grid. The configuration consists of two TS7720 Virtualization Engines and two TS7740 Virtualization Engines.

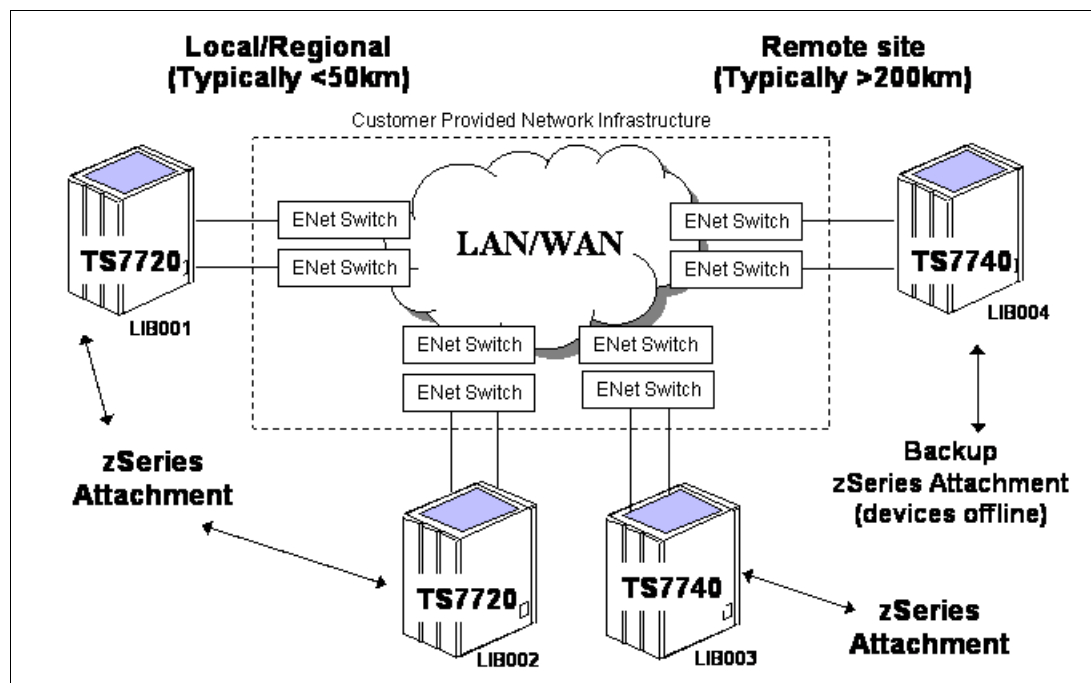


Figure 2-6 TS7720 Virtualization Engine four-cluster grid

Ownership

The concept of ownership has been introduced with the TS7700 Virtualization Engine to ensure data integrity. At any point in time, a logical volume is owned by one cluster. The owning cluster controls access to the volume and the attributes of a volume. Ownership can change dynamically. If a cluster needs to mount a logical volume on one of its virtual devices and it is not the owner of that volume, it has to obtain ownership for this specific volume first. We describe these processes in detail in 2.6.1, “Data integrity by volume ownership” on page 77.

I/O Tape Volume Cache

All vNodes in a Grid have direct access to all logical volumes in the Grid, either from the local cluster’s Tape Volume Cache (TVC) through Fibre Channel or from a remote cluster’s TVC through a Grid WAN connection. During non-Fast Ready mount processing, Device Allocation Assist (DAA) attempts to select the best available cluster for the mount point and a TVC will be designated as the I/O TVC. All I/O operations that are associated with the virtual tape drive on which the mount was issued are routed to and from its vNode to the I/O TVC. For details, see “I/O TVC selection” on page 87.

Management interface

The management interface (MI) is web-based interface and was introduced with the TS7700 Virtualization Engine. It is used to monitor and configure the system, to manage access, and to manage logical and physical volumes. All tasks that were previously done through the IBM System Storage 3494 Library Manager or the Enhanced Library Controller (ELC) console have been consolidated into this interface. We cover the management interface in more detail in Chapter 4, “Hardware implementation” on page 175 and Chapter 8, “Operation” on page 409.

User roles

Users of the TS7700 Virtualization Engine can be assigned one or more roles associated with the performance of certain system functions.

User roles are levels of access, assigned by the administrator, that allow users to perform certain functions. User roles are created by using the TS7700 Virtualization Engine management interface. When an administrator creates a new user account, the administrator must specify an initial password for the account. Multiple roles cannot be assigned to one user.

Administrators can assign the following roles when defining new users:

- ▶ Operator

The operator has access to monitoring information, but is restricted from changing settings for performance, network configuration, feature licenses, user accounts, or custom roles. The operator is also restricted from inserting and deleting logical volumes.

- ▶ Lead operator

The lead operator has almost all of the same permissions as the administrator, but may not change network configuration, feature licenses, user accounts, or custom roles.

- ▶ Administrator

The administrator has the highest level of authority, including the authority to add or remove user accounts. The administrator has access to all service functions and TS7700 Virtualization Engine resources.

- ▶ Manager

The manager has access to health and monitoring information, jobs and processes information, and performance data and functions, but is restricted from changing most settings, including those for logical volume management, network configuration, and feature licenses.

- ▶ Custom roles

The administrator can name and define two custom roles by selecting the individual tasks permitted for each custom role. All available tasks are selectable for a custom role, with the exception of creating, modifying, and deleting a user. Figure 2-7 shows the Roles & Permissions panel, including the custom roles.

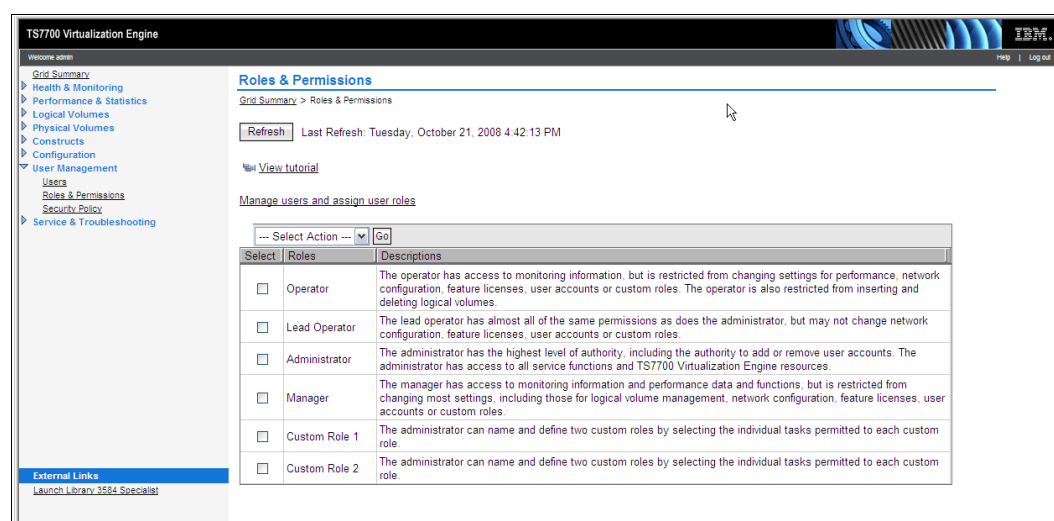


Figure 2-7 TS7700 Virtualization Engine MI Roles & Permissions panel

Table 8-4 on page 549 lists the tasks and which role is authorized to view a window or perform a certain task.

2.1.2 Multi-cluster grid terms

In this section, we explain terms that are specific to a multi-cluster grid configuration.

Composite Library

The Composite Library is the logical image of the Grid, which is presented to the host. Distinctly different from the IBM Virtual Tape Server, a *single cluster* TS7700 Virtualization Engine also has a five-character hexadecimal Composite Library ID defined. From an architectural perspective, a stand-alone TS7700 Virtualization Engine is considered a Grid that consists of only one cluster. We refer to such a cluster as a *single-cluster grid*.

With both TS7700 Virtualization Engine single-cluster grid configurations and TS7700 Virtualization Engine multi-cluster grid configurations, a Composite Library is presented to the host. In the case of a stand-alone TS7700 Virtualization Engine, the host sees a logical tape library with sixteen 3490E tape control units, each with sixteen IBM 3490E tape drives, attached through two or four FICON channel attachments.

Figure 2-8 illustrates the host view of a single-cluster grid configuration.

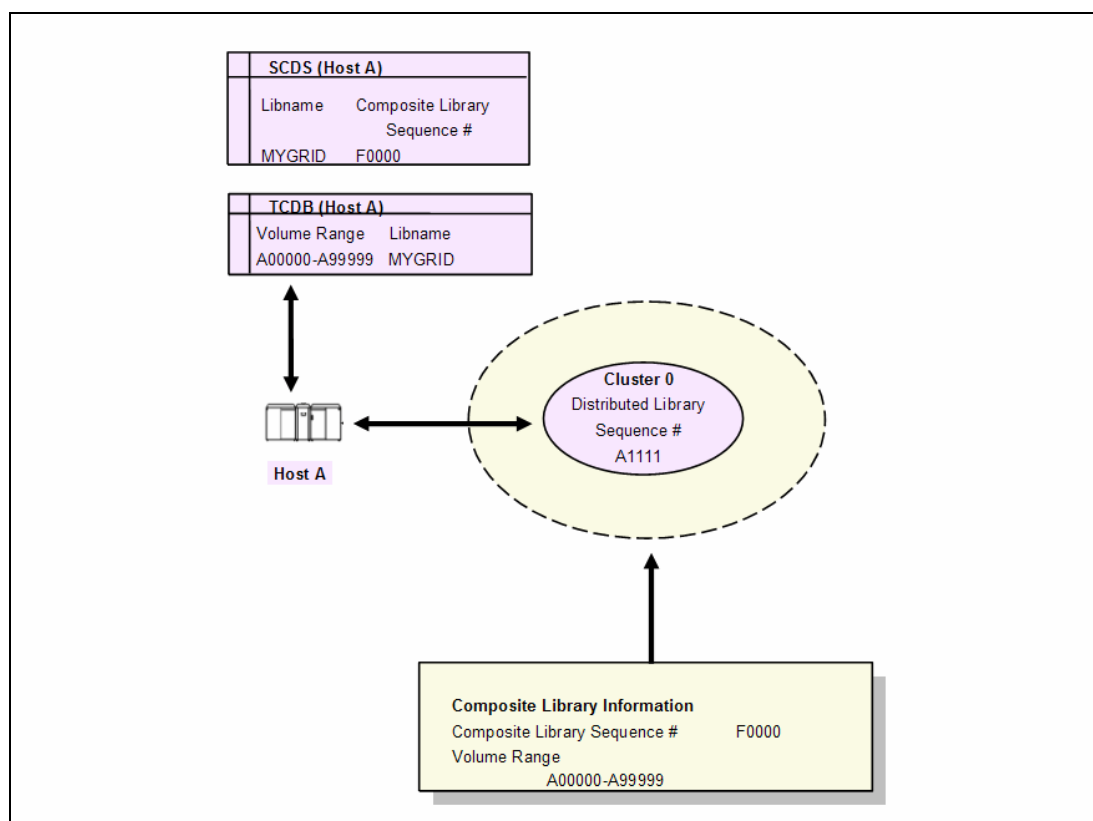


Figure 2-8 TS7700 Virtualization Engine single-cluster grid configuration

In the case of a multi-cluster grid, the host sees a logical tape library with sixteen 3490E tape control units per cluster, each with sixteen IBM 3490E tape drives, attached through four or more FICON channel attachments.

Figure 2-9 illustrates the host view of a two-cluster grid configuration.

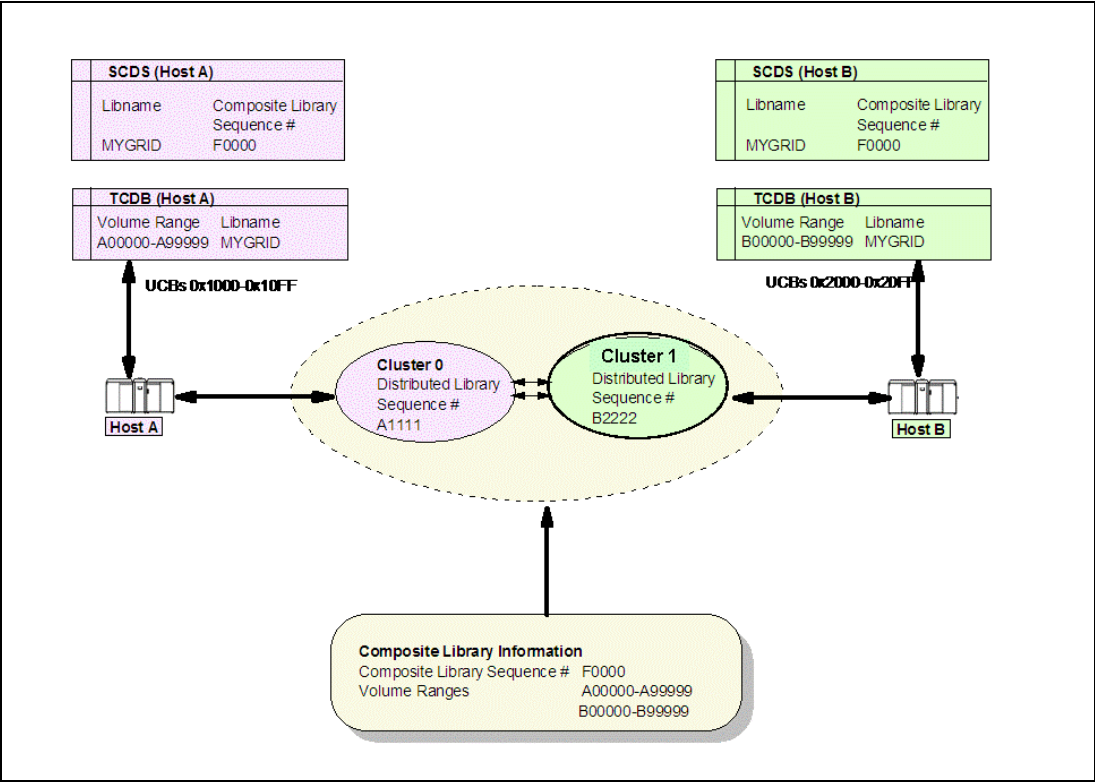


Figure 2-9 TS7700 Virtualization Engine two-cluster grid configuration

Figure 2-10 illustrates the host view of a three-cluster grid configuration.

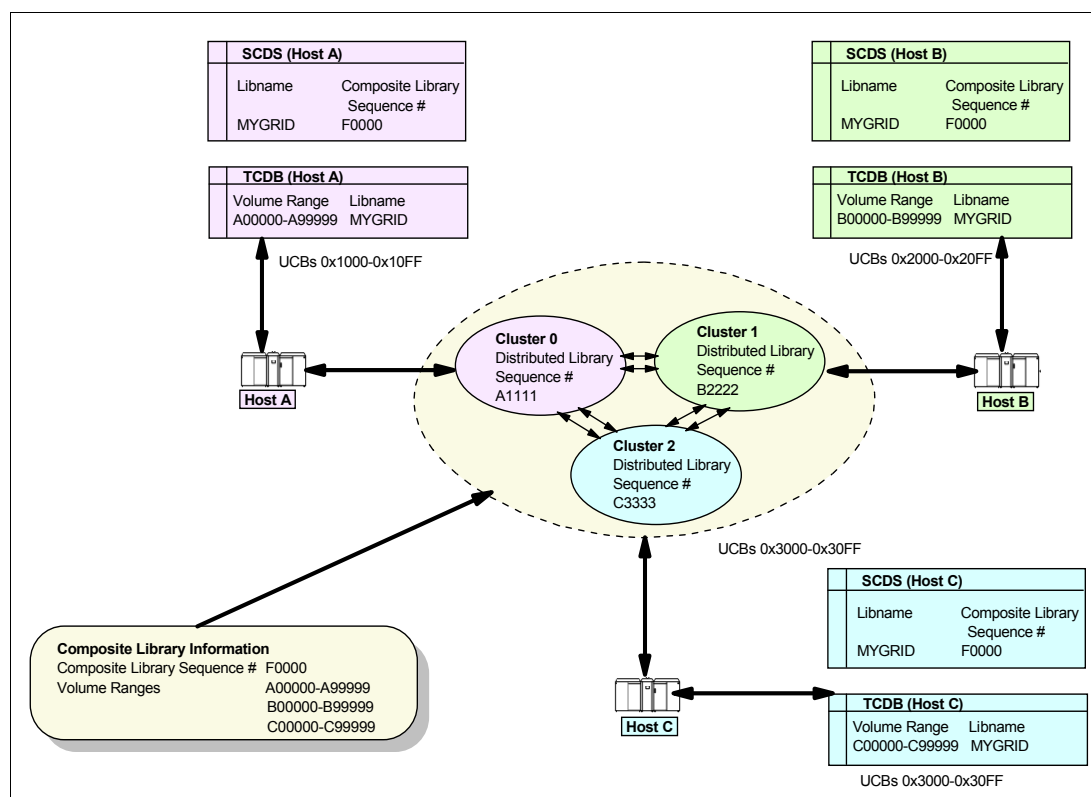


Figure 2-10 TS7700 Virtualization Engine three-cluster grid configuration

Important: A Composite Library ID must be defined both for a multi-cluster grid and also for a single-cluster grid. For a single-cluster grid, the Composite Library ID must not be the same as the Distributed Library ID. For a multiple grid configuration, the Composite Library ID must differ from any of the unique Distributed Library IDs. Both the Composite Library ID and Distributed Library ID are five-digit hexadecimal strings.

Distributed library

Each cluster in a Grid is a distributed library, which consists of a TS7700 Virtualization Engine. In the case of a TS7740 Virtualization Engine, it is also attached to a physical TS3500 Tape Library.

A distributed library consists of cluster hardware components:

- ▶ A Virtualization Engine
- ▶ A Disk Cache controller
- ▶ Disk expansion drawers
- ▶ A frame
- ▶ Attachment to a physical library (TS7740 Virtualization Engine only)
- ▶ A number of physical tape drives (TS7740 Virtualization Engine only)

In a Grid configuration that includes multiple TS7700 Virtualization Engines, each of which has the Grid Enablement Feature installed, a distributed library must be defined at the host for each cluster. The host has sufficient knowledge about the distributed libraries to allow appropriate console message handling from the Integrated Library Manager. At the host, the distributed library is only defined to SMS. It is defined using the existing ISMF panels and has no tape devices defined. The virtual tape devices are defined for the Composite Library only.

Copy consistency point

In a Grid environment, you can specify where and when you want to have a copy of your data. Currently, the three settings consist of two copy consistency points and an option to have no copy at all:

RUN	The copy will occur as part of the rewind/unload (RUN) operation and completes before the rewind/unload operation at the host finishes. This mode is comparable to the immediate copy mode of the Peer-to-Peer VTS.
Deferred	The copy will occur some time after the rewind/unload operation at the host. This mode is comparable to the deferred copy mode of the Peer-to-Peer VTS.
No Copy	No copy will be made.

On each cluster in a multi-cluster grid, a copy consistency point setting is specified for the local cluster and one for each of the other clusters. The settings need not be the same on all clusters in the Grid. When a volume is mounted on a virtual tape device, the copy consistency point policy of the cluster to which the virtual device belongs will be honored.

For a detailed description, see 2.6.2, “Copy policy management” on page 81.

Dynamic Grid Load Balancing

Dynamic Grid Load Balancing is an algorithm that is used within the TS7700 Virtualization Engine and that continually monitors the rate at which data is processed by each network link. It attempts to ensure that each link is used according to its capabilities.

Figure 2-11 on page 27 shows a two-cluster grid with two links that are connected through two paths that have different lengths. In the example, the primary path is 600 km in length; the secondary path is 800 km in length. Dynamic Grid Load Balancing can help to compensate unbalanced network performance in these types of situations.




Figure 2-11 Dynamic Grid Load Balancing in a two-cluster grid

With Dynamic Grid Load Balancing, the link capabilities are evaluated every five minutes. A host warning message is sent to the system console if the Grid links are severely degraded. The threshold for the warning is reached when a difference of 60% or more exists between links.

2.1.3 Tape Virtualization general terms

In this section, we explain the general terms that are related to Tape Virtualization. Some of the concepts are different or do not apply in the same way to the TS7720 Virtualization Engine because no physical tape is attached to it. Most of the concepts apply to Tape Virtualization in general and if you have previous experience with the predecessor to the TS7700 Virtualization Engine, the IBM TotalStorage Virtual Tape Server (VTS), you will be familiar with some of these terms.

Tape Volume Cache

The Tape Volume Cache (TVC) is disk space where virtual volumes are stored and retrieved. The TVC is under complete and exclusive control of the TS7700 Virtualization Engine.

When a TS7740 Virtualization Engine virtual volume in the TVC is closed and demounted, it is scheduled to be copied to a stacked volume. Volumes that have previously been copied to physical tape can be removed from cache to provide space for new virtual volume data or for the recall of existing logical volumes. By default, candidates for removal from cache are selected by using a least recently used (LRU) algorithm. The only external method with which to influence this is through tape volume cache management.

Virtual volumes in a TS7720 Virtualization Engine configuration always remain in the TVC because no physical tape drives are attached to the TS7720 Virtualization Engine. However, cache management policies can be configured to manage virtual volumes in a multi-cluster configuration. For further details, see 2.5.2, "Tape Volume Cache Management" on page 57.

Virtual drives

From a host perspective, a TS7700 Virtualization Engine appears as sixteen logical IBM 3490E tape control units, each with sixteen drives attached through FICON channels. Virtual tape drives and control units are defined just like physical IBM 3490s through hardware configuration definition (HCD). Defining a preferred path for the virtual drives gives you no benefit because the IBM 3490 control unit functions inside the TS7700 Virtualization Engine are emulated to the host.

Each virtual drive has the following characteristics of real tape drives:

- ▶ Uses host device addressing.
- ▶ Is included in the I/O generation for the system.
- ▶ Is varied online or offline to the host.
- ▶ Signals are ready when a virtual volume is loaded.
- ▶ Responds and processes all IBM 3490E I/O commands.
- ▶ Becomes *not* ready when a virtual volume is rewound and unloaded.

For software transparency reasons, the functionality of the 3490E Integrated Cartridge Loader (ICL) is also included in the virtual drive's capability. All virtual drives indicate that they have an ICL. For scratch mounts, there is no benefit in using the emulated ICL in the TS7700 Virtualization Engine to preload virtual cartridges. When the Fast Ready attribute is set, the virtual volume is created directly in the tape volume cache without the need to copy any data from a stacked volume. No mechanical operation is required to mount a logical scratch volume.

Physical drives

The physical tape drives used by a TS7740 Virtualization Engine are installed in a TS3500 Tape Library. The physical tape drives are not addressable by any attached host system and are completely under the control of the TS7740 Virtualization Engine. The TS7740 Virtualization Engine supports IBM 3592-J1A, TS1120, or TS1130 physical tape drives. See 2.3.3, “TS7740 Virtualization Engine components” on page 44 for more information about physical drives and libraries.

Important: Do not use Advanced Library Management System (ALMS) to reassign physical tape drives attached to a TS7740 Virtualization Engine in a TS3500 Tape Library. This function is used for Open Systems hosts and can cause unpredictable complications when used to reconfigure TS7740 Virtualization Engine attached tape drives without the assistance of trained IBM service personnel.

Virtual volumes

A virtual volume is created in the Tape Volume Cache when the host writes data to the TS7700 Virtualization Engine subsystem. All host interaction with tape data in a TS7700 Virtualization Engine is through virtual volumes and virtual tape drives.

Each virtual volume, like a real volume, has the following characteristics:

- ▶ It has a unique volume serial (VOLSER) number known to the host.
- ▶ It is loaded and unloaded on a virtual device.
- ▶ It supports all tape write modes, including Tape Write Immediate.
- ▶ It contains all standard tape marks and data blocks.
- ▶ It supports an IBM standard label (and non-labeled tape also).
- ▶ It can be appended to after it is initially written from the beginning of tape.

- The application is notified that the write operation is complete when the data has been written in the TVC.
- Each host-written record has a logical block ID.
- The end of volume is signaled when the total number of bytes written into the TVC after compression has reached 400 MB for an emulated cartridge system tape (CST), 800 MB for an emulated enhanced capacity cartridge system tape (ECCST) volume, or 1000, 2000, or 4000 MB using the larger volume size options.

Note: With the use of large logical volume sizes of 2000 or 4000 MB, recalls in a TS7740 Virtualization Engine can take considerably longer. This recall can occur when a volume is not resident in the Tape Volume Cache and has to be copied from physical tape to the cache.

Copy time between clusters is also affected by larger logical volumes because more data must be copied when the amount of data is larger.

Also, consider the consequences of large logical volumes and the size of your cache. The TS7700 Virtualization Engine does not restrict the logical volume size based on the cache configuration. With a 1 TB configuration and 4,000 MB logical volumes, where the logical volumes are filled, opening 256 virtual devices with this size of volume fills the cache and causes throttling. With smaller cache configurations, be sure not to use filled larger logical volume sizes for all devices at once.

With data compression based on the IBMLZ1 algorithm by the FICON channel card in a TS7700 Virtualization Engine node, the actual host data stored on a virtual CST or ECCST volume can vary from 1.2 GB up to 12 GB (assuming a 3:1 compression ratio). The default logical volume sizes of 400 MB or 800 MB, defined at insert time, can be overwritten at every individual scratch mount by the use of a Data Class construct.

Virtual volumes can exist only in a TS7700 Virtualization Engine. You can direct data to a virtual tape drive by directing it to a system-managed storage (SMS) tape Storage Group of the TS7700 Virtualization Engine and using the automatic class selection (ACS) routines of a system-managed tape environment. You can also pass Data Class name, Management Class name, Storage Class name, and Storage Group name. The TS7700 Virtualization Engine then uses these names to determine how to manage a virtual volume. You can define up to 1,000,000 virtual volumes per single or multi-cluster grid. The VOLSERS for the logical volumes are defined through the management interface.

To speed up the scratch mount process, you should associate a “Fast Ready” attribute with a scratch category of TS7700 Virtualization Engine virtual volumes. Although you might want to use logical volume sizes larger than 400 or 800 MB, you still define CST or ECCST emulated cartridges to the TS7700 Virtualization Engine. These simulate MEDIA1 with 400 MB capacity or MEDIA2 with 800 MB capacity. Virtual volumes go through the same cartridge entry processing as native cartridges inserted in a library.

Logical volumes

When a TS7740 Virtualization Engine virtual volume is copied from the TVC to a physical tape cartridge, it becomes a logical volume. This process is called *premigration*. When the volume is removed from the TVC, the process is called *migration*. When a logical volume is moved from a physical cartridge to the TVC, the process is called *recall* and the volume becomes a virtual volume again.

Although the TS7700 Virtualization Engine emulates a 3490E tape of a specific size (400, 800, 1000, 2000, or 4000 MB), the space used in the TVC is no more than that needed for the

number of bytes of data written to the virtual volume. When the TS7740 Virtualization Engine virtual volume is written to the physical tape, it uses only the space occupied by the data. The TVC or the physical tapes do not preallocate space in 400 MB, 800 MB, 1000 MB, 2000 MB, or 4000 MB segments.

As virtual volumes are copied from the TVC to a physical cartridge, they are stacked on the cartridge end to end, taking up only the space written by the host application. This arrangement maximizes utilization of a cartridge's storage capacity. The storage management software within the TS7740 Virtualization Engine node manages the location of the logical volumes on the physical cartridges. The only control the user has over the location of the data is when volume pooling is used. For details, see “Physical volume pooling” on page 63. If a logical volume cannot fit on the stacked volume, it will not span across two or more physical cartridges. Instead, the stacked volume is marked full and the logical volume is written on another stacked volume from the assigned pool.

Through the use of the Data Class construct that is applied to a volume during a scratch mount request, you have the choice of additional logical volume sizes. The default logical volume sizes of 400 and 800 MB are extended to 1000, 2000, and 4000 MB. The default logical volume sizes, which are still used at insert time, can be overwritten during a scratch mount according to the size that is specified for the volume's currently assigned Data Class. All definitions are outboard and can be defined through the TS7700 Virtualization Engine management interface.

Note: One million logical volumes are supported in both the TS7720 Virtualization Engine and TS7740 Virtualization Engine. However, in the TS7720 Virtualization Engine, the number of logical volumes is limited by the size of the cache, the average size of the logical volumes, and the host bus adapter (HBA) compression ratio. For example, with an average logical volume size of 800 MB and the cache size of 70 TB, the number of logical volumes is limited to 87,500 logical volumes. With a compression ratio of 2.5:1, that results in up to 175 TB of user data that can be stored in the TS7720 Virtualization Engine.

Stacked volume

Physical cartridges used by the TS7740 Virtualization Engine to store logical volumes are completely under the control of the TS7740 Virtualization Engine node and are not known to the hosts. Physical volumes are called stacked volumes. Stacked volumes must have unique, system-readable VOLSERS and external labels like any other cartridges in a tape library.

Through the TS3500 Tape Library Specialist, you define which physical cartridges are to be used by the TS7740 Virtualization Engine. Logical volumes stored on those cartridges are mapped by the TS7740 Virtualization Engine internal storage management software. When you use pooling, your stacked volumes can be assigned to individual pools. Logical volumes can then be assigned to specific stacked volume pools. In the case of out-of-scratch scenarios, pools can be set up to enable “borrowing” from other pools. The methodology and configuration of volume pooling is covered in “Physical volume pooling” on page 63.

Tokens

Tokens are used to track changes to the ownership, data, or properties of a logical volume. The tokens are mirrored at each cluster that participates in a Grid and represent the current state and attributes of the logical volume. Note the following information about tokens:

- ▶ Every logical volume has a corresponding token.
- ▶ The Grid component manages updates to the tokens.
- ▶ Tokens are maintained in an IBM DB2® database coordinated by the local hNodes.
- ▶ Each cluster's DB2 database has a token for every logical volume in the Grid.

Tokens are internal data structures that are not directly visible to you. However, they can be retrieved through reports generated with the Bulk Volume Information Retrieval (BVir) facility.

Hot tokens

In a multi-cluster configuration, Hot tokens are special tokens that are used to track changes against sites that are not operational or are in service mode. In a two-cluster grid, there is only one non-operational cluster and one remaining operational cluster. In a three or four-cluster grid, more than one cluster can be in service mode at one time. To avoid single points of failure, ensure that only one cluster is in service mode at any given time.

Hot tokens are used as follows:

- ▶ The remaining sites track logical volume updates in their local DB2 database.
- ▶ After the down site returns back into the Grid, all sites share their hot tokens.
- ▶ The returning site builds a hot token list of outstanding updates it needs to make and processes them in the background.
- ▶ A site can come online while hot tokens are being processed. Unlike a Peer-to-Peer VTS, any unresolved tokens are now dynamically resolved if a host issues a command and specifies a volume that has hot tokens associated with it.

Service Prep

The transition of a cluster into Service mode is called Service Prep. Service Prep allows for a cluster to be gracefully and temporarily removed as an active member of the Grid. The remaining sites are given the option to acquire ownership of the volumes owned by the site going into service while it is away from the Grid. If the option is not exercised, then ownership is retained by the original cluster. Operations that target the Distributed Library entering service are surrendered by the site going into service before transition to service completes. Other Distributed Libraries within the Composite Library will remain available. The host device addresses associated with the site in service send up Device State Change attentions to the host. This allows the host to remove them from device selection for new mounts.

When Service Prep completes and the cluster enters Service mode, nodes at the site in Service mode remain online, but are prevented from communicating with other sites. This allows service personnel to perform maintenance tasks on the site's nodes, run hardware diagnostics, and so on, without impacting other sites.

Only one Service Prep can occur within a Composite Library at any one time. If a second Service Prep is attempted at the same time, it will fail. After Service Prep for one cluster has completed and the cluster is in Service mode, then another cluster can be placed in Service Prep.

A site in Service Prep automatically cancels and reverts back to an ONLINE state if any ONLINE peer in the Grid experiences an unexpected outage while the Service Prep process is still active. The last ONLINE cluster in a multi-cluster configuration cannot enter the Service Prep state, which includes a stand-alone cluster. The operator or IBM System Services Representative (SSR) can cancel Service Prep processing to return the subsystem to a normal state.

After a cluster has completed Service Prep and has entered service mode, it will remain in this state until explicitly taken out of service mode by the operator or by the IBM SSR.

A best practice is to place only one cluster into service mode at a time. This eliminates the creation of a single point of failure within the Grid. Each cluster in a Grid can be upgraded in a serial manner. Two or more clusters can be placed in service mode at the same time, but this action should only be done during emergencies.

Time coordination

All nodes in the Grid subsystem *must* coordinate their time with one another. All nodes in the system keep track of time in relation to Coordinated Universal Time (UTC), also known as Greenwich Mean Time (GMT). Statistics are also reported in relation to UTC.

The preferred method to keep nodes in sync is by using a Network Time Protocol (NTP) server. The NTP server can be a part of the Grid WAN infrastructure, customer intranet, or a public server on the Internet (Figure 2-12).

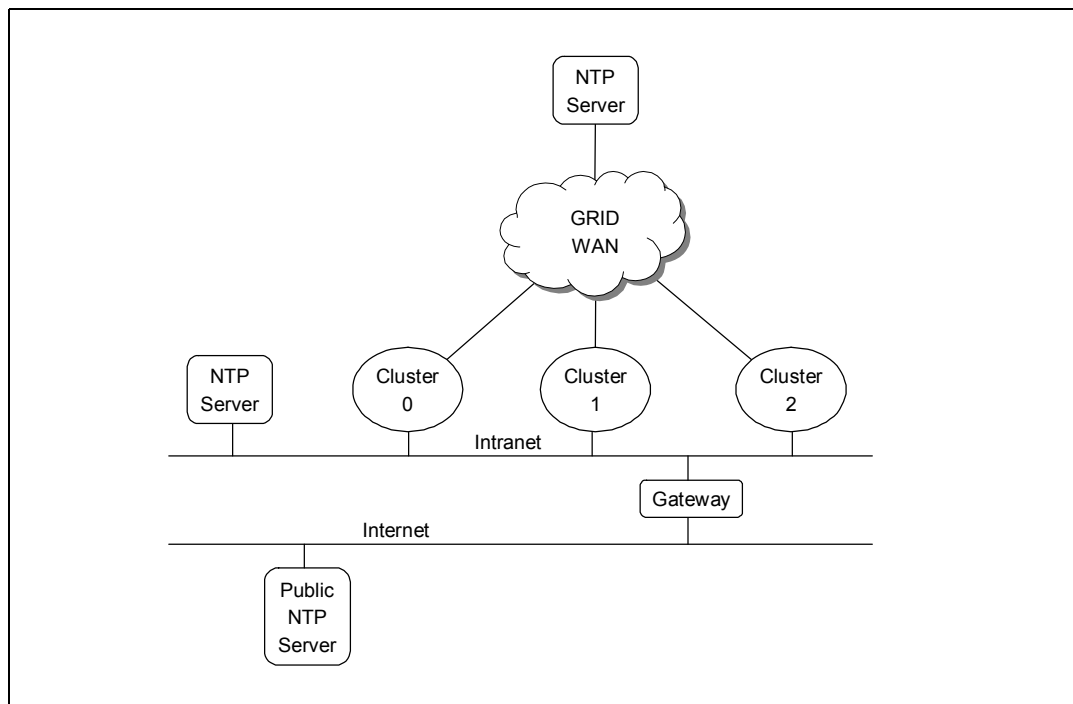


Figure 2-12 Time coordination with NTP servers

The NTP server address is configured into system VPD on a system-wide scope, such that all nodes will access the same NTP server. All clusters in a Grid need to be able to communicate with the same NTP server defined in VPD. In the absence of a NTP server, all nodes will coordinate time with Node 0 or the lowest cluster index designation (Cluster 0, if Cluster 0 exists; if not, then Cluster 1, and so on).

2.2 Architectural capabilities

When you compare a Peer-to-Peer VTS with a TS7700 Virtualization Engine Multi Grid Cluster, the differences are quite obvious. With the TS7700 Virtualization Engine, you do not need external Virtual Tape Controller (VTC) hardware. Additionally, the peers in a Grid are connected through Ethernet, rather than through ESCON® or FICON, as in the Peer-to-Peer VTS. When you compare a stand-alone VTS and a single-cluster grid, the differences might not be apparent at first. They look much the same from a functional perspective, although under the covers, the architectural design of a TS7700 Virtualization Engine is different from that of a VTS.

In this section, we discuss the architectural design of the TS7700 Virtualization Engine and its potential capabilities. To help you understand the changes in architecture, we start with a short description of the VTS architecture.

Monolithic design of a VTS

A VTS performs all functions within a single IBM System p server. The VTS also serves as the RAID disk controller. The RAID system is tightly integrated into the system. When you want to expand the capabilities and functionality, these expansions must fit within the capabilities of the System p server. There are a number of drawbacks to this architectural concept:

- ▶ A limited number of expansion slots are available.
- ▶ The amount of RAM that is supported by the platform is limited.
- ▶ The CPU capabilities are limited.
- ▶ Increasing the VTS capabilities requires more powerful System p servers.

The *law of diminishing returns* comes into play for upgrades. When you reach a certain point, a small incremental increase in capabilities will require a large investment in a more powerful server. As code components are tightly integrated with one another, implementing new functions impacts large amounts of code. Finally, the system must always be upgraded or extended as a whole, because of the tight integration of its components.

All these concerns have been addressed in the architectural design of the TS7700 Virtualization Engine, which we now describe.

Modular design of the TS7700 Virtualization Engine

The modular design of the TS7700 Virtualization Engine separates the functionality of the system into smaller components with well-defined, open interfaces. It provides a platform where the smaller components can be scaled up from a small configuration into an extremely large one. This provides the capability to grow the solution to meet business objectives.

In the future, you might be able to plug in separate components of similar functionality to provide a solution for specific environments.

The TS7700 Virtualization Engine is built on a distributed node architecture, as we previously mentioned, which is composed of nodes, clusters and Grid configurations. In addition, node architecture is based on standards interfaces.

Standards-based interface

The standards-based interface between the elements is Fibre Channel Protocol (FCP) for data movement and TCP/IP for command and control. This includes the management interface. The functions of the VTCs in a PTP VTS have been integrated into the hNode into something we call the *grid layer*. The grid layer interconnects all of the nodes whether they are in the same cluster or in another cluster.

Comparison

A comparison of the modular design of a TS7700 Virtualization Engine and the monolithic design of a VTS is shown in Figure 2-13.

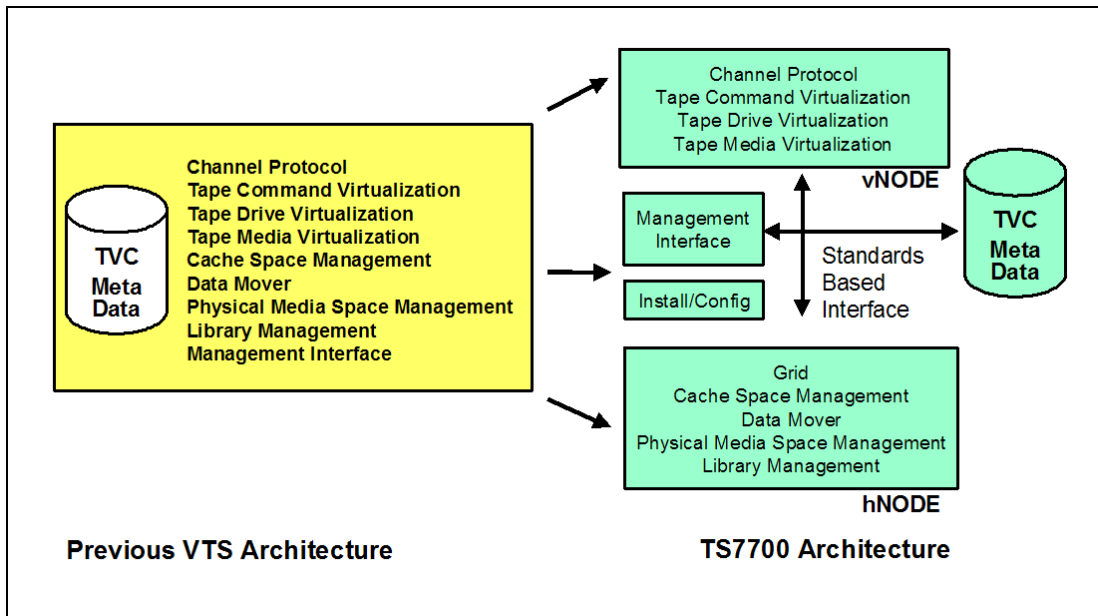


Figure 2-13 Monolithic design versus modular design

With the modular design of the TS7700 Virtualization Engine, it will be possible in the future to add gNodes, or add vNodes and hNodes separately to an existing cluster configuration. You might want to add vNodes to increase the number of virtual device addresses or add vNodes with some special functionality to perform additional tasks. More hNodes give you better performance and increased availability. Instead of replacing one monolithic system with a larger one, this new architecture allows horizontally scaling of the components to meet performance and functional goals.

Serviceability also benefits from having more than one node in a cluster. A nondisruptive code update for a single-cluster grid requires two gNodes. Redundant nodes eliminate or minimize the risk of planned and unplanned outages.

As the cache is a modular, highly or continuously available component of the TS7700 Virtualization Engine, you can upgrade the cache with minimal impact to the whole system.

Multi-cluster grid

When referencing a multi-cluster grid in this book, we are referring to a two-, three- or four-cluster grid. These are the only Grid configurations currently supported. From an architectural perspective, there is no need to limit the number of clusters to four.

In a multi-cluster grid with more than two clusters, copy consistency point policies, especially in combination with Override settings, become increasingly important. You will have to plan carefully where you want a copy of your data to reside and when to initiate it.

A logical volume can be accessed from any virtual device in the system, even if the logical volume has only a single replica. Should a failure occur at the source cluster, replicated logical volumes will still be available, because any data replicated between the clusters is accessible through any other cluster in a Grid configuration. Each copy is a valid source for the logical volume. A Grid configuration looks like a single storage subsystem to the hosts

attached to the clusters. It can be described as a Composite Library with underlying Distributed Libraries, similar to the prior generation's Peer-to-Peer (PTP) VTS.

Integration of the Library Manager into TS7700 Virtualization Engine

One of the benefits of the TS7700 Virtualization Engine architecture, and an enhancement included in TS7700 Virtualization Engine Release 1.5 and later, is the Integrated Library Manager. With the Integrated Library Manager, all functions previously performed through an external 3494 Library Manager or 3953 Enhanced Library Controller (ELC) are now accessed through the TS7700 Virtualization Engine management interface. This simplifies the environment and provides a single user interface for TS7700 Virtualization Engine configuration and operation.

See 2.4, "Integrated Library Manager" on page 54 for more details about the integration of Library Manager functions into the TS7700 Virtualization Engine.

TS7700 Virtualization Engine statistics for performance

The TS7700 Virtualization Engine statistics content and processes differ significantly from its VTS predecessor.

In addition to completely new statistical record formats, major changes have been made to the frequency of collection, data storage, host data retrieval, and reporting tools.

The TS7700 Virtualization Engine provides two types of statistical information that can be useful for performance tuning and capacity planning. They assist in evaluating how the TS7700 Virtualization Engine subsystem is performing in your environment:

- ▶ *Point-in-Time Statistics* are performance related. The TS7700 Virtualization Engine updates these statistics every 15 seconds, but does not retain them. Each 15-second update overlays the prior data. The user can retrieve the Point-in-Time statistics from the TS7700 Virtualization Engine at any time using the BVIR facility. A subset of Point-In-Time statistics is also available through the TS7700 Virtualization Engine management interface.
- ▶ *Historical Statistics* encompass a wide selection of performance and capacity planning information. Historical Statistics are collected by the TS7700 Virtualization Engine every 15 minutes, and stored for 90 days in a TS7700 Virtualization Engine database. For long term retention, the user can also retrieve these records with the Bulk Volume Information Retrieval (BVIR) function.

See 9.10, "Bulk Volume Information Retrieval" on page 657 for more information about BVIR function and use.

Call Home support

The Call Home function generates a service alert automatically when a problem is detected with a 3953 Tape System, TS3500 Tape Library, 3494 Tape Library, 3494 B10 or B20 VTS, VTC, TS7700 Virtualization Engine, 3590 A60 Controller, or 3592 Controller. Status information is transmitted to the IBM Support Center for problem evaluation. An IBM SSR can be dispatched to the installation site if maintenance is required.

2.3 TS7700 Virtualization Engine hardware components

TS7700 Virtualization Engine Release 1.7 provides capacity options that significantly enhance the performance capabilities of the subsystem. In this section, we describe the hardware components that are part of the TS7700 Virtualization Engine. We address the TS7720 Virtualization Engine disk-only solution and the TS7740 Virtualization Engine with its TS3500 Tape Library components.

TS7720 Virtualization Engine is available in one or two frame configurations:

- ▶ An IBM 3952 Tape Base Frame, which houses the following components:
 - One TS7720 Virtualization Engine server
 - One TS7720 Virtualization Engine Cache Model CS8 Controller with up to six optional TS7720 Virtualization Engine Cache Model XS7 Drawers
 - Two Ethernet routers
 - One TS3000 Service Console
- ▶ An optional IBM 3952 Storage Expansion Frame, which houses the following components:
 - Two TS7720 Virtualization Engine Model CS8 Cache Controllers
 - Up to ten optional TS7720 Virtualization Engine Model XS7 Cache Drawers

The TS7740 Virtualization Engine consists of the following configurations:

- ▶ An IBM 3952 Tape Frame, which houses the following components:
 - One TS7740 Virtualization Engine server
 - One TS7740 Virtualization Engine Model CC8 Cache Controller with zero, one, or three TS7740 Virtualization Engine Model CX7 Cache Drawers
 - Two Ethernet routers
- ▶ One TS3000 Service Console
- ▶ One TS3500 Tape Library and its four to 16 IBM 3592 Tape Drives and two FC switches

2.3.1 Common components for the TS7720 Virtualization Engine and TS7740 Virtualization Engine models

Several components are common for both models of the TS7700 Virtualization Engine. We describe these components in this section.

IBM 3952 Tape Base Frame

The IBM 3952 Tape Base Frame Model F05 is a frame that provides up to 36 EIA units of usable space to contain components of the defined tape solution. The 3952 Tape Base Frame is not a general purpose frame, but is designed to contain only the components of specific tape offerings, such as the TS7740 Virtualization Engine or TS7720 Virtualization Engine. Only components of one solution family may be installed in a single 3952 Tape Base Frame. The 3952 Tape Frame integrates a power supply to support these tape solution components and offers an optional Dual AC Power supply feature.

In a TS7700 Virtualization Engine configuration, the 3952 Tape Base Frame is used for the installation of the TS7700 Virtualization Engine node, the TS7700 Virtualization Engine Cache Controller, the TS7700 Virtualization Engine Cache Modules, two Ethernet routers, and optionally, the TS3000 Service Console. These components are discussed in detail for the TS7700 Virtualization Engine specific models in the following sections.

Ethernet routers

Two Ethernet routers are installed in the IBM 3952 Tape Base Frame. Each has eight internal LAN ports that are used for accessing the TS7700 Virtualization Engine management interface, and two WAN ports that are used for attachment to the customer LAN and to the TS3000 System Console. Figure 2-14 shows the Ethernet switches in detail.

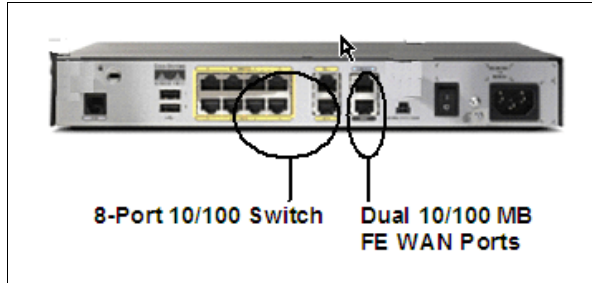


Figure 2-14 Ethernet router

TS7700 Virtualization Engine Grid adapters

The connection path between multiple TS7700 Virtualization Engines in a Grid configuration are the two Grid adapters in the TS7700 Virtualization Engine server. TS7700 Virtualization Engine Release 1.5 introduced adapters to support expanded Grid interconnections. The dual-ported 1 Gbps Ethernet adapters can be copper RJ45 or optical fiber (short wave). These optical adapters have an LC duplex connector. Only the top port on each dual-port adapter card will be activated.

The Ethernet adapters may not be intermixed. Both Grid adapters in a TS7700 Virtualization Engine must be the same feature code.

TS3000 System Console

The TS3000 System Console is a required component for the TS7700 Virtualization Engine. It can be a new console or, if you already have one, an existing TS3000 System Console can be used.

If this is a new TS3000 System Console, it can be installed in the TS7700 Virtualization Engine 3952-F05 Base Frame or a customer-provided rack.

When a TS3000 System Console is ordered with a TS7700 Virtualization Engine, it can be preinstalled in the 3952-F05 frame at the plant. The TS3000 System Console is a 1U server and includes a keyboard, display, mouse, and one 16 port Ethernet switch.

2.3.2 TS7720 Virtualization Engine components

The TS7720 Virtualization Engine is the latest member of the TS7700 Virtualization Engine family. The TS7720 Virtualization Engine is a disk-only Virtualization Engine that provides most of the benefits of the TS7740 Virtualization Engine, but without physical tape attachment. The TS7720 Virtualization Engine can be used to write tape data that does not need to be copied to physical tape. This allows access to the data from the Virtualization Engine Cache until the data expires.

Software configurations and support is transparent for IBM z/OS®, IBM z/VM, IBM z/VSE™, and IBM z/TPF, and it is also transparent to IBM or third-party tape management software.

The TS7720 Virtualization Engine is configured with different server models and different disk cache models than the TS7740 Virtualization Engine. A TS7720 Virtualization Engine

consists of 3952 Model F05 Tape Base Frame and an optional 3952 Model F05 Storage Expansion Frame.

The 3952 Model F05 Tape Base Frame houses the following components:

- ▶ One TS7720 Virtualization Engine Server, 3957 Model VEA
- ▶ One TS7720 Virtualization Engine SATA Cache Controller, 3956 Model CS8 with zero to six TS7720 Virtualization Engine SATA Cache Drawers, 3956 Model XS7
- ▶ Two Ethernet routers

The 3952 Model F05 Storage Expansion Frame houses the following component:

- ▶ Two TS7720 Virtualization Engine SATA Cache Controllers, 3956 Model CS8 with zero to ten TS7720 Virtualization Engine SATA Cache Drawers, 3956 Model XS7

Each TS7720 Virtualization Engine SATA Cache Controller, 3956 Model CS8 using 2-TB drives provides approximately 19.8 TB of capacity after RAID 6 formatting. The base frame cache controller provides a different capacity than the two in the expansion frame. The TS7720 Virtualization Engine SATA Cache Drawers, 3956 Model XS7 using 2 TB drives provides approximately 23.8 TB of capacity after RAID 6 formatting. The 3956 Model CS8 Cache Controller RAID 6 has two spare disks. There are no spare disks in the TS7720 Virtualization Engine SATA Cache Drawers, 3956 Model XS7 when attached to a TS7720 Virtualization Engine SATA Cache Controller, 3956 Model CS8. There are no spare disks within a 3956 Model XS7 cache drawer when attached to a TS7720 virtualization engine. The arrays that are configured within the expansion drawers will use the spares within the CS8 cache drawer.

Using 2-TB disk drives, the maximum configurable capacity of the TS7720 Virtualization Engine at R1.7 (PGA1 and later) with the 3952 Model F05 Storage Expansion Frame is 441 TB of data before compression. This works out to 1.32 PB of host data, assuming a 3:1 compression rate.

A single cluster TS7720 Virtualization Engine generates out-of-space warning messages to its attached hosts as available cache space falls below preset thresholds. Then, the storage administrator must decide what to delete from the appliance's cache to provide space for continued workload.

Homogeneous multi-cluster TS7720 Virtualization Engine R1.7 configurations also support automatic logical volume removal. With TS7720 Virtualization Engine R1.7 configurations now supporting up to 441 TB of usable capacity, the differences between cluster cache sizes can be significant.

The TS7720 Virtualization Engine also includes the Integrated Library Manager, internal Workflow Management controls, Grid capabilities, Cache Controller and Tape Expansion Frame adapters, along with the existing TS7700 Virtualization Engine functionality that is not related to physical tape library attachment or back-end tape drive support.

TS7720 Virtualization Engine Grids are supported in two-, three-, and four-site configurations. The intermix of TS7740 Virtualization Engines with TS7720 Virtualization Engines is supported as of TS7700 Virtualization Engine Release 1.6. See Figure 2-15 for an example of a three-cluster configuration of TS7720 Virtualization Engines. Notice that there is no tape library or Enhanced Library Controller attachment.

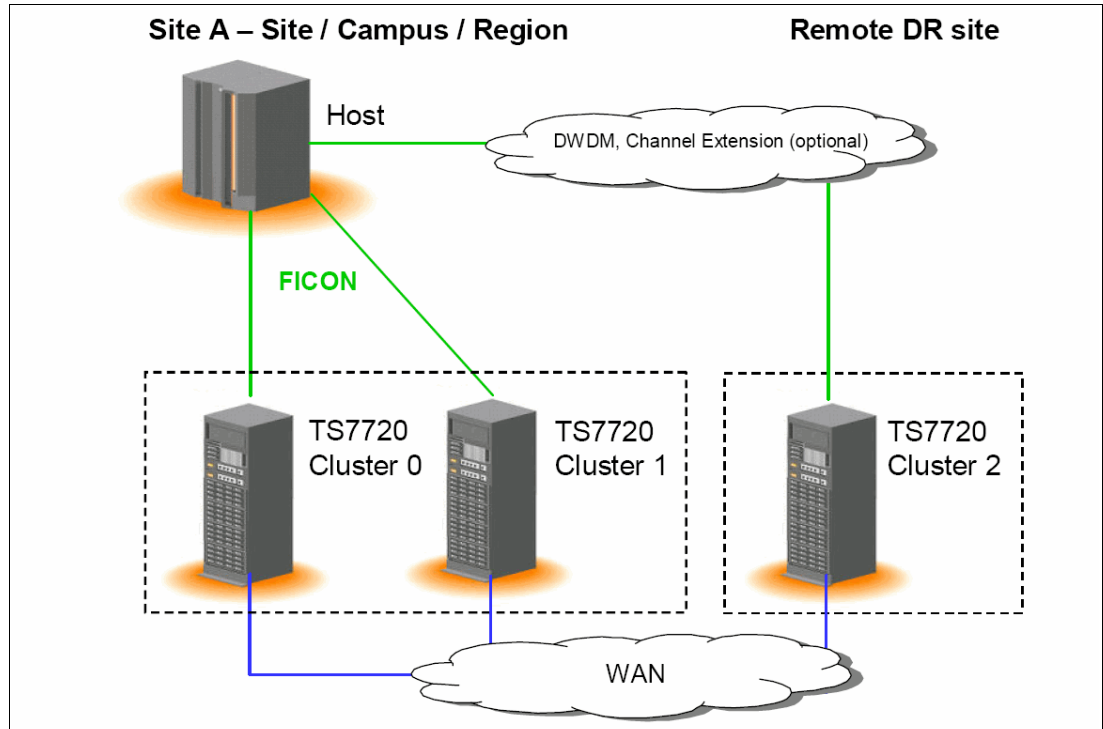


Figure 2-15 TS7720 Virtualization Engine three-cluster grid configuration

Figure 2-16 shows the components of a TS7720 Virtualization Engine Base Frame fully configured, including the optional TS3000 System Console.




Figure 2-16 TS7720 Virtualization Engine components

Figure 2-17 shows Storage Expansion Frame Cache Controller and Cache Drawer locations for a fully configured TS7720 Virtualization Engine.




Figure 2-17 TS7720 Virtualization Engine Expansion Frame components

TS7720 Virtualization Engine Server (3957-VEA)

The TS7720 Virtualization Engine Server consists of an IBM System POWER® 5++ server and two expansion drawers for I/O adapters. The TS7700 Virtualization Engine Server controls virtualization processes such as host connectivity and device virtualization, and hierarchical storage management (HSM) functions for logical volumes and replication.

Figure 2-18 shows the front view of the TS7720 Virtualization Engine Server.




Figure 2-18 TS7720 Virtualization Engine Server (front view)

The TS7720 Virtualization Engine Server offers the following features:

- ▶ Two dual-core, 64-bit, 2.1 GHz processors with 2 MB of L2 cache memory
- ▶ 16 GB of 266 MHz DDR2 memory
- ▶ Two hot-pluggable power supplies with embedded enclosure cooling units
- ▶ The following integrated features:
 - Two SCSI controllers
 - Two USB ports
 - Two serial ports
 - Six hot-pluggable PCI adapter slots (to house Ethernet adapters for Grid communication)
 - Eight hot-pluggable DASD (2 × 4-pack)
- ▶ Two dual port 1 Gbps Grid Ethernet adapters (copper or optical)
Only the top port on each adapter is used. The bottom port is reserved for future use.
- ▶ One hot-pluggable DVD-ROM drive
- ▶ One hot-pluggable DVD-RAM drive

Each Expansion Unit I/O adapter drawer offers six additional hot-pluggable PCI adapter slots. These are used to house FICON adapters for host attachment, Fibre Channel adapters for Cache Controller attachment, and Ethernet adapters for communication.

TS7720 Virtualization Engine Cache Controller (3956-CS8)

The TS7720 Virtualization Engine Cache Controller, 3956 Model CS8 is a self-contained 3U enclosure that mounts in the 3952 Tape Base Frame and the optional 3952 Storage Expansion Frame. Figure 2-19 shows the TS7720 Virtualization Engine Cache Controller from the front (left side) and from the rear (right side).




Figure 2-19 TS7720 Virtualization Engine Cache Controller, 3956-CS8 (front view and rear view)

The TS7720 Virtualization Engine Cache Controller provides RAID-6-protected virtual volume disk storage for fast retrieval of data from cache. The TS7720 Virtualization Engine Cache Controller offers the following features:

- ▶ Two Fibre Channel processor cards
- ▶ Two battery backup units (one for each processor card)
- ▶ Two power supplies with embedded enclosure cooling units
- ▶ 16 DDMs, each with a storage capacity of 2 TB, for a usable storage capacity of 19.84 TB
- ▶ Configurations support one or three TS7720 Virtualization Engine Cache Controllers
 - All configurations provide one TS7720 Virtualization Engine Cache Controller in the 3952 Tape Base Frame

The 3952 Tape Base Frame can have zero to six TS7720 Virtualization Engine SATA Cache Drawers, 3956 Model XS7

- All configurations with the optional 3952 Storage Expansion Frame provide two additional TS7720 Virtualization Engine Cache Controllers, 3956 Model CS8s.

The 3952 Storage Expansion Frame can have zero to ten TS7720 Virtualization Engine SATA Cache Drawers, 3956 Model XS7

TS7720 Virtualization Engine Cache Drawer (3956-XS7)

The TS7720 Virtualization Engine Cache Drawer is a self-contained 3U enclosure that mounts in the 3952 Tape Base Frame and in the optional 3952 Storage Expansion Frame. Figure 2-20 shows the TS7720 Virtualization Engine Cache Drawer from the front (left side) and from the rear (right side). It offers attachment to one of the TS7720 Virtualization Engine Cache Controllers.




Figure 2-20 TS7720 Virtualization Engine Cache Drawer (front view and rear view)

The TS7720 Virtualization Engine Cache Drawer expands the capacity of the TS7720 Virtualization Engine Cache Controller by providing additional RAID-6-protected disk storage. Each TS7720 Virtualization Engine Cache Drawer offers the following features:

- ▶ Two Fibre Channel processor cards
- ▶ Two power supplies with embedded enclosure cooling units
- ▶ Sixteen DDMs, each with a storage capacity of 2 TB, for a usable capacity of 23.84 TB per drawer

2.3.3 TS7740 Virtualization Engine components

TS7700 Virtualization Engine Release 1.7 uses the TS7740 Virtualization Engine Cache Controller, 3956 Model CC8. Each of these disk cache models includes sixteen 600 GB Fibre Channel HDDs, providing approximately 9.6 TB of usable capacity after RAID 5 formatting.

New TS7740 Virtualization Engine Release 1.7 plant-built configurations include:

- ▶ One TS7740 Virtualization Engine Server, 3957 Model V06
- ▶ One TS7740 Virtualization Engine Cache Controller, 3956 Model CC8
- ▶ Zero, one, or three TS7740 Virtualization Engine Cache Drawers, 3956 Model CX7

The total usable capacity of a TS7740 Virtualization Engine R1.7, with one 3956 Model CC8 and three 3956 model CX7s is approximately 28.17 TB before compression.

The Model CX7s may be installed in the plant or in the field. Field upgrades are available.

Figure 2-21 shows a summary of TS7740 Virtualization Engine components.




Figure 2-21 Virtualization Engine TS7740 components

TS7740 Virtualization Engine Server (3957-V06)

The TS7740 Virtualization Engine Server consists of an IBM POWER 5++ System and two expansion drawers for I/O adapters. The TS7740 Virtualization Engine Server controls virtualization processes such as host connectivity and device virtualization. Additionally, it controls hierarchical storage management (HSM) functions for logical volumes, such as replication, and organization of data across physical media and libraries.

Figure 2-22 shows the front view of the TS7740 Virtualization Engine Server.




Figure 2-22 TS7740 Virtualization Engine Server (front view)

The TS7740 Virtualization Engine Server offers the following features:

- ▶ Two dual-core, 64-bit, 2.1 GHz processors with 2 MB of L2 cache memory
- ▶ 16 GB of 266 MHz DDR2 memory
- ▶ Two hot-pluggable power supplies with embedded enclosure cooling units
- ▶ The following integrated features:
 - Two SCSI controllers
 - Two USB ports
 - Two serial ports
 - Six hot-pluggable PCI adapter slots (used to house Ethernet adapters for Grid communication)
 - Eight hot-pluggable DASDs (2 × 4-pack)
- ▶ Two dual-port 1 Gbps Ethernet adapters (copper or optical)
Only the top port of each adapter is used. The bottom port is reserved for future use.
- ▶ One hot-pluggable DVD-ROM drive
- ▶ One hot-pluggable DVD-RAM drive

Each Expansion Unit I/O adapter drawer offers six additional hot-pluggable PCI adapter slots, which are used to house FICON adapters for host attachment, Fibre Channel adapters for Cache Drawer attachment, Fibre Channel adapters for tape drive I/O, and Ethernet adapters for communication.

TS7740 Virtualization Engine Cache Controller (3956-CC8)

The TS7740 Virtualization Engine Cache Controller is a self-contained 3U enclosure that mounts in the 3952 Tape Frame.

Figure 2-23 shows the front and rear view of the TS7740 Virtualization Engine Model CC8 Cache Controller.




Figure 2-23 TS7740 Virtualization Engine Cache Controller (front and rear view)

Figure 2-24 shows a diagram of the rear view.




Figure 2-24 TS7740 Virtualization Engine Cache Controller (rear view)

The TS7740 Virtualization Engine Cache Controller provides RAID-6-protected virtual volume disk storage to temporarily hold data from the host before writing it to physical tape, then caches the data to allow fast retrieval from the disk. The TS7740 Virtualization Engine Cache Controller offers the following features:

- ▶ Two Fibre Channel processor cards
- ▶ Two battery backup units (one for each processor card)
- ▶ Two power supplies with embedded enclosure cooling units
- ▶ 16 DDMs, each possessing 600 GB of storage capacity, for a usable capacity of 7.04 TB
- ▶ Optional attachment to one or three TS7740 Virtualization Engine Cache Drawers

TS7740 Virtualization Engine Cache Drawers (3956-CX7)

The TS7740 Virtualization Engine Cache Drawer is a self-contained 3U enclosure that mounts in the 3952 Tape Frame.

Figure 2-25 shows the front view and the rear view of the TS7740 Virtualization Engine Model CX7 Cache Drawer.




Figure 2-25 TS7740 Virtualization Engine Cache Drawer (front and rear view)

The TS7740 Virtualization Engine Cache Drawer expands the capacity of the TS7740 Virtualization Engine Cache Controller by providing additional RAID 6 disk storage. Each TS7740 Virtualization Engine Cache Drawer offers the following features:

- ▶ Two Fibre Channel processor cards
- ▶ Two power supplies with embedded enclosure cooling units
- ▶ 16 DDMs, each with 600 GB of storage capacity, for a total usable capacity of 7.04 TB per drawer
- ▶ Attachment to the TS7740 Virtualization Engine Cache Controller

TS7740 Virtualization Engine Tape Library attachments, drives, and media

In a TS7740 Virtualization Engine configuration, the TS7740 Virtualization Engine is used in conjunction with an attached tape library. The tape library will have its own logical partition with dedicated tape drives and media for the TS7740 Virtualization Engine.

Tape libraries

The TS3500 Tape Library is the only library that is supported with TS7740 Virtualization Engine Release 1.6 and later. The TS7740 Virtualization Engine includes the functionality of the previously required IBM 3953 Tape System (below R1.5), consisting of one or more frames and the ELC.

Tape drives

The TS7740 Virtualization Engine supports the following tape drives inside a TS3500 Tape Library:

- ▶ **TS1130 Tape Drive**

TS7740 Virtualization Engine Release 1.6 and later, include support for TS1130 Tape Drives. When a TS1130 Tape Drive is attached to a TS7740 Virtualization Engine, all attached drives must be TS1130 Tape Drives. Intermixing is not supported with 3592-J1A or TS1120 Tape Drives. TS1130 Tape Drives can read data that has been written by either of the prior generation 3592 Tape Drives. Tapes written in E05 format are appended to in E05 format. The first write to supported tapes are written in the E06 format.

If TS1130 Tape Drives are detected and other generation 3592 Tape Drives are also detected, the TS1130 Tape Drives will not be configured. TS1130 Tape Drives read data that has been written by either of the prior generation 3592 Tape Drives.

- ▶ **TS1120 Tape Drive (native mode or emulating 3592-J1A Tape Drives)**

Tape Drive types that are used with the TS7740 Virtualization Engine cannot be intermixed, except for 3592-J1A Tape Drives and TS1120 Tape Drives operating in 3592-J1A emulation mode.

- ▶ **IBM 3592 Model J1A Tape Drive**

TS7740 Virtualization Engine supports 3592 Model J1A Tape Drives, but a better practice is to use more recent generations of the 3592 Tape Drives to obtain the maximum benefit of the TS7740 Virtualization Engine R1.7. The 3592 Model J1A Tape Drives cannot be intermixed with TS1130 Tape Drives. The 3592 Model J1A Tape Drives can be intermixed with TS1120 Tape Drives only when the TS1120 Tape Drives are set to J1A emulation mode.

All TS1120 or TS1130 Tape Drives associated with the TS7740 Virtualization Engine must be capable of encryption and must have encryption enabled. Encryption is not supported on 3592 Model J1A Tape Drives. If you intend to use TS1120 Tape Drives in encryption mode, you must enable encryption on all TS1120 Tape Drives attached to the TS7740 Virtualization Engine.

For more information regarding encryption, see 2.5.5, “Encryption” on page 73 and 3.7, “Planning for encryption in the TS7740 Virtualization Engine” on page 153.

Tape media

The TS7740 Virtualization Engine supports 3592 Media Types JJ, JA, and JB. Capacity for each media depends on the type of tape drive. Table 2-1 on page 49 shows the media capacity and performance for each model of 3592 Tape Drive (media format) and the various Media Types.

Table 2-1 Summary of capacity and performance for media type and format

Media type	E06 format capacity and datarate (min - max)	E05 format capacity and datarate (min - max)	J1A format capacity and datarate (min - max)
JB	1 TB 40 MBps - 160 MBps	700 GB 40 MBps - 150 MBps	N/A
JA	640 GB 40 MBps - 140 MBps	500 GB 40 MBps - 140 MBps	300 GB 30 MBps - 70 MBps
JJ	128 GB 40 MBps - 140 MBps	100 GB 40 MBps - 140 MBps	60 GB 30 MBps - 70 MBps

Note the following information:

- Use of JB tape cartridges with TS1120 Tape Drives is supported only when operating in native capacity mode. Use of JB tape cartridges is supported with TS1130 Tape Drives.
- When TS1120 Tape Drives operating in native mode are replaced with TS1130 Tape Drives, additional data written on the 3592-E05 formatted tapes will be appended to until they are full. As the active data on the E05 formatted tapes gets reclaimed or expired, the tape goes back to the scratch pool. On the next write, the tape will be reformatted to the 3592-E06 data format.
- TS1120 Tape Drives can operate in native E05 mode or in 3592-J1A emulation mode. However, all 3592 Tape Drives that are associated with the TS7700 Virtualization Engine must be TS1120 Tape Drives to operate in native E05 mode. To use TS1120 Tape Drives in native E05 mode, all attached drives must be set to E05 native mode. To use TS1120 Tape Drives in J1A emulation mode, all attached drives must be set to J1A emulation mode. The capacity of the tape media and performance will be as specified for a 3592-J1A.
- When 3592-J1A drives (or TS1120 Tape Drives in J1A emulation mode) are replaced with TS1130 Tape Drives, the TS7740 Virtualization Engine marks the J1A formatted tapes with a status of active data FULL. By marking these tapes full, the TS7740 Virtualization Engine does not append more data because the TS1130 Tape Drive cannot append data to a J1A formatted tape. As the active data on the J1A formatted tape gets reclaimed or expired, the tape goes back to the scratch pool where it will eventually get reformatted to the E06 data format.

See *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789 for more information about IBM 3592 Tape Drives and Media.

In the following sections, we describe the characteristics of the TS3500 Tape Library and TS1130 Tape Drive.

2.3.4 TS3500 Tape Library

The TS3500 Tape Library is part of a family of tape libraries designed for large automated tape storage and backup solutions.

Originally delivered in year 2000 with Linear Tape Open (LTO) Ultrium technology, the TS3500 Tape Library offers a robust enterprise library solution for mid-range and high-end open systems. Since its introduction, the library has been enhanced to accommodate different drive types and operating platforms, including the attachment of System z hosts and tape drive controllers.

Currently, the TS3500 Tape Library (Figure 2-26) is capable of connecting drives to host systems with FICON or ESCON attachments, and any combination of Fibre Channel and Ultra2/Wide High and Low Voltage Differential (LVD) SCSI. The combination of proven reliable tape handling and functional enhancements have resulted in an optimal design for a robust enterprise solution with outstanding retrieval performance. Typical cartridge move time is less than three seconds. Using TS1120, TS1130, or LTO Ultrium high density cartridge technology, the TS3500 Tape Library provides a powerful and robust tape storage solution for the whole enterprise with a minimal footprint.



Figure 2-26 TS3500 Tape Library

In summary, the TS3500 Tape Library provides:

- ▶ A modular, scalable, and automated Tape Library combining IBM tape and automation for open systems and mainframe hosts, using a variety of IBM drive types
- ▶ Attachment to IBM System z, IBM System i®, IBM System p, IBM RS/6000®, IBM System x®, Netfinity®, Oracle, Hewlett-Packard, and other servers that are not IBM
- ▶ Connectivity using FICON, ESCON, Fibre Channel, Low Voltage Differential (LVD) SCSI, and High Voltage Differential (HVD) SCSI
- ▶ IBM Multi-Path Architecture to support redundant control paths, mixed drive configurations, and library sharing between multiple applications

TS3500 Tape Library frame models

The TS3500 Tape Library frames come in several different models to provide a broad and flexible range of configuration options. The TS3500 Tape Library frame models are used to differentiate the functions they provide and their placement within a solution, such as the accessor frame, tape drive type, and storage cells supported. The following list describes the currently available frame models:

- ▶ TS3500 Tape Library Model L23: A base frame that contains up to 12 IBM 3592 Tape Drives and up to 260 IBM TotalStorage 3592 tape cartridges.
- ▶ TS3500 Tape Library Model D23: An expansion frame that can house up to 12 IBM 3592 Tape Drives and up to 400 IBM TotalStorage 3592 tape cartridges. Up to 15 expansion frames can be installed with a base frame.
- ▶ TS3500 Tape Library Model L53: A base frame that contains up to 12 Ultrium tape drives and up to 287 LTO Ultrium tape cartridges.

- ▶ TS3500 Tape Library Model D53: An expansion frame that contains up to 12 Ultrium tape drives and up to 440 LTO Ultrium tape cartridges. Up to 15 expansion frames can be installed with a base frame.
- ▶ IBM 3584 Model HA1: An optional high availability frame. In conjunction with service bay features on the TS3500 Tape Library Models D23 and D53, this model supports the installation of a second accessor in the TS3500 Tape Library.
- ▶ TS3500 Tape Library Model S24: A driveless storage-only high density expansion frame for up to 1000 3592 tape cartridges. Up to 15 expansion frames can be installed with a base frame. Advanced Library Management System (ALMS) is required for any library with a Model S24 frame. This frame can optionally be configured as service bay B. High Density (HD) slots contain the tape cartridges in a tiered architecture. The cartridge immediately accessible in the HD slot is a Tier1 cartridge, behind this is Tier2, and so on. The maximum number of tiers in an S24 frame is four.
- ▶ TS3500 Tape Library Model S54: A driveless storage-only high density expansion frame for up to 1320 Ultrium (LTO) tape cartridges. Up to 15 expansion frames can be installed with a base frame. Advanced Library Management System (ALMS) is required for any library with a Model S54 frame. This frame can optionally be configured as service bay B. High Density (HD) slots contain the tape cartridges in tiered architecture. The cartridge immediately accessible in the HD slot is a Tier1 cartridge, behind this is Tier2, and so on. The maximum number of tiers in a S54 frame is five.

Figure 2-27 shows the TS3500 Tape Library High Density frame.



Figure 2-27 TS3500 Tape Library High Density frame

The Lxx frames also contain an I/O station for 16 cartridges. If both LTO and IBM 3592 Tape Drives are installed inside the TS3500 Tape Library, the optional second I/O station is required for the second media format. The second I/O station is installed below the first I/O station. The drive type in the Lxx frame determines which I/O station is in the top position. In an L53 frame, this position is the I/O station for LTO cartridges, and in an L23 frame, the I/O station for 3592 cartridges are in the top position.

All currently available frame models can be intermixed in the same TS3500 Tape Library as previously installed frame models. Previous frame models include the L22, L32, L52, D22, D32, and D52.

TS3500 Tape Library attached to TS7740 Virtualization Engine

The TS3500 Tape Library houses the physical tape drives for a TS7740 Virtualization Engine. A dedicated logical library is required in the TS3500 Tape Library for each TS7740

Virtualization Engine. The back end drives will be assigned to this partition for exclusive use by the attached TS7740 Virtualization Engine. Defining a Cartridge Assignment Policy for the TS7740 Virtualization Engine partition is highly recommended to assist in insert processing of physical cartridges into the TS3500 Tape Library.

The tape drive types that are currently supported for System z attachment in the TS3500 Tape Library are the TS1130, TS1120, and the 3592 Model J1A Tape Drives. Although these tape drives also support Open Systems hosts attachment, some differences exist between System z and Open Systems attachment.

With the Library Emulation (LE) capability in the TS7700 Virtualization Engine R1.5 and later, all library management functions are performed by the TS7700 Virtualization Engine and are accessed through the management interface.

Note: Advanced Library Management System (ALMS) must be enabled in a TS3500 Tape Library to support attachment of tape subsystems to System z hosts.

The TS7740 Virtualization Engine requires a minimum of four dedicated physical tape drives and supports a maximum of 16 drives. These drives can reside in the TS3500 Tape Library Model L23/L22 Base Frame and the Model D23/D22 Drive Frames. Up to twelve drives can be installed in one frame. TS7740 Virtualization Engine attached drives do not have to be installed in contiguous positions. For availability and performance reasons, a good practice is to spread the drives evenly across two frames in close proximity to each other. TS7740 Virtualization Engine attached drives cannot be shared with other systems, but they can share a frame with tape drives attached to other TS7740 Virtualization Engines, tape controllers, or Open Systems hosts.

The TS7740 Virtualization Engine R1.7 supports TS1130, TS1120, and IBM 3592 Model J1A Tape Drives. The IBM 3592-J1A Tape Drive has been withdrawn from marketing, but existing drives can be used for TS7740 Virtualization Engine attachment. When TS1120 drives are intermixed with 3592-J1A drives on the same TS7740 Virtualization Engine, the TS1120 drives run in J1A Emulation mode. When only TS1120 drives are attached to a TS7740 Virtualization Engine, set them to run in native E05 mode. TS1130 drives cannot be intermixed with either TS1120 drives or J1A tape drives.

Logical libraries in the TS3500 Tape Library

The patented IBM Multi-Path Architecture enables SCSI medium changer libraries to be shared by heterogeneous host platforms, without middleware or a dedicated server acting as a library controller. This facility is implemented on a number of IBM libraries, including the TS3500 Tape Library.

With the Multi-Path Architecture, sharing is possible by allowing the allocation of library storage slots and tape drives to logical libraries. Servers can then run separate applications for each logical library. Heterogeneous hosts share the library robotics, but have the physical library drives and slots divided between them. Each host sees the same accessor, but has a different view of the physical attributes of the library. This existing TS3500 Tape Library concept of logical libraries is used to segregate the physical devices associated with TS7740 Virtualization Engines from the 3592-J70 and TS1120 Tape Controllers, and from the Open Systems drives. The TS7740 Virtualization Engine is allocated its own logical library partition in the TS3500 Tape Library. Multiple logical libraries can exist in the same TS3500 Tape Library connected to other TS7740 Virtualization Engines, connected to an IBM 3953 ELC or dedicated to Open System hosts.

For details of the TS3500 Tape Library, see *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789.

2.3.5 IBM System Storage TS1130 Tape Drive

The IBM System Storage TS1130 Tape Drive is the third generation in the IBM 3592 Tape Drive family. The TS1130 Tape Drive provides higher capacity and performance compared to previous generations. It provides a native data rate of up to 160 MBps and uncompressed capacity of 1 TB with the Extended Data Cartridge (JB) versus a 100 MBps data rate and 700 GB uncompressed capacity of the TS1120 Tape Drive.

The TS1130 Tape Drive is available in two 3592 models, E06 and EU6. The 3592 Model EU6 Tape Drive is only available as a field upgrade from the TS1120 Tape Drive Model E05.

The TS1130 Tape Drive also supports data encryption and key management and is downward read compatible (n-2) to the 3592 Model J1A and is downward write compatible (n-1) to the 3592 Model E05 formats. The TS1130 Tape Drive uses the same IBM 3592 Cartridges available for TS1120 and 3592-J1A. Attachments to System z and Open Systems platforms are maintained.

Other capabilities that are included in the TS1130 Tape Drive are as follows:

- ▶ A giant magnetoresistance (GMR) read/write head provides fewer data errors and improved reliability.
- ▶ Head overcoat technology provides improved corrosion protection and wear characteristics.
- ▶ Improved standby power management delivers energy savings through lower power consumption.

TS7740 Virtualization Engine support for TS1130 was introduced with TS7700 Virtualization Engine Release 1.5. See *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789 for more information regarding IBM TS1130 Tape Drives.

Note: TS1130 Tape Drives are supported at TS7740 Virtualization Engine Release 1.5 and later.

2.3.6 IBM System Storage TS1120 Tape Drive

The IBM System Storage TS1120 Tape Drive (3592-E05) is the second generation in the 3592 family with dramatically enhanced performance and capacity characteristics when compared to the 3592 Model J1A Tape Drive. When compared to previous tape drive families and generations, the TS1120 is unique in two ways:

- ▶ The E05 models can read cartridges written by J1A drives and write in J1A Emulation mode. The drive can dynamically change its mode of operation per physical cartridge.
- ▶ The 3592-E05 Tape Drives can be intermixed with 3592-J1A Tape Drives:
 - In the same tape library frame
 - Behind the same TS7740 Virtualization Engine
 - Behind the same Model TS1120 or 3592-J70 Controller
 - Behind the same VTS

Notes:

- ▶ When intermixed with 3592-J1A Tape Drives behind the same TS7740 Virtualization Engine or other controller, the TS1120 Model E05 Tape Drive always operates in J1A Emulation mode.
- ▶ To use the full capacity and functionality of the TS1120 Model E05, do not intermix it with J1A Tape Drives behind the same TS7740 Virtualization Engine or other controller.

2.3.7 IBM TotalStorage 3592-J1A Tape Drive

The IBM TotalStorage 3592-J1A was the first generation of the 3592 family. The 3592 Tape Drive family replaces the prior generation of 3590 Tape Drives. The 3592-J1A provides:

- ▶ Up to 40 MBps native data rate, which is over 2.5 times the 14 MBps for the 3590 E or H models
- ▶ Up to 300 GB native cartridge capacity (900 GB at 3:1 compression), which is a five times increase over the 60 GB for the 3590 H models
- ▶ A 60 GB Economy tape cartridge with fast read/write access
- ▶ 300 GB and 60 GB WORM (Write Once, Read Many) cartridges for increased security of data

Note: The IBM TotalStorage 3592-J1A has been withdrawn. Its replacements are the TS1120 Tape Drive and the TS1130 Tape Drive.

The IBM TotalStorage 3592-J1A does not support Media Type JB.

2.4 Integrated Library Manager

TS7700 Virtualization Engine Release 1.5 introduced the Integrated Library Manager. This functionality is available for both TS7720 Virtualization Engine and TS7740 Virtualization Engine models. In this section, we explain the Integrated Library Manager support, the aspects of the integration, and the differences for TS7720 Virtualization Engine and TS7740 Virtualization Engine.

With the Integrated Library Manager support, the TS7700 Virtualization Engine is now Library Emulation (LE) capable. Library Emulation is the integration of the virtual volume processing into the TS7700 Virtualization Engine. Prior to TS7700 Virtualization Engine Release 1.5, virtual volume processing for TS7740 Virtualization Engines was done by the Enhanced Library Controller (ELC) in the 3953-F05 frame.

The key aspects of the Integrated Library Manager are as follows:

- ▶ Logical volume management is now completely controlled by the TS7700 Virtualization Engine instead of being split between the ELC and the TS7700 Virtualization Engine. This means that all library related host commands are now executed by the TS7700 Virtualization Engine directly.
- ▶ The setup, policy management, and monitoring functions that relate to the logical library (logical volumes and drives), which were previously performed through the ELC, are now part of the TS7700 Virtualization Engine management interface.

- ▶ The ELC and its associated components in the 3953-F05 frame are no longer needed for the TS7740 Virtualization Engine at R1.5 and later.
- ▶ The TS7740 Virtualization Engine directly issues physical volume requests to the underlying automation. The TS7740 Virtualization Engine issues SCSI Medium Changer (SMC) commands through the tape drive interfaces to the TS3500 Tape Library.
- ▶ Other System z tape drives, such as TS1120 tape controller attached drives, still require the 3953 when connected to a TS3500 Tape Library. These other System z drives and their associated media are in a separate logical library from the drives and media of the TS7740 Virtualization Engine.

Library management integration into the TS7700 Virtualization Engine has many benefits:

- ▶ Increased performance of the library operations (for example, mount time and category updates)
- ▶ Removal of database duplication
- ▶ Removal of an external ELC and its operating system
- ▶ Removal of complex communications interface with the ELC
- ▶ Removal of ELC for virtual volume processing
- ▶ Use of the TS7700 Virtualization Engine management interface for ELC functions

TS7700 Virtualization Engine management interface

The TS7700 Virtualization Engine management interface (MI) is a web-based user interface to the TS7700 Virtualization Engine for monitoring and control of the subsystem. Using the management interface, you can:

- ▶ Monitor the status of the TS7700 Virtualization Engine functions and hardware components, including the back-end tape drives.
- ▶ Monitor the performance of the TS7700 Virtualization Engine and Grid
- ▶ Manage the TS7700 Virtualization Engine logical volumes
- ▶ Configure the TS7700 Virtualization Engine and Grid
- ▶ Manage the operation of the TS7700 Virtualization Engine and Grid
- ▶ Manage user access to the TS7700 Virtualization Engine
- ▶ Manage the TS7700 Virtualization Engine Policy Construct definitions
- ▶ Access the TS7700 Virtualization Engine service and customer information centers
- ▶ Set up back-end tape encryption (TS7740 Virtualization Engine only)
- ▶ Set up Copy Export pools (TS7740 Virtualization Engine only)
- ▶ Initiate Service Prep and shutdown of the TS7700 Virtualization Engine cluster
- ▶ Activate Ownership Takeover Mode for a given TS7700 Virtualization Engine cluster

The TS7700 Virtualization Engine can be connected to your LAN using the supplied Ethernet connections. During the installation process, the IBM SSR will configure the Ethernet of the TS7700 Virtualization Engine with customer assigned TCP/IP parameters.

With the TS7700 Virtualization Engine Integrated Library Manager, you can perform additional functions from the management interface. The following functions that were performed from the external ELC prior to TS7700 Virtualization Engine Release 1.5 are now integrated into the management interface for operational support:

- ▶ Fast ready category configuration
- ▶ Constructs
 - Storage groups
 - Management classes
 - Storage classes
 - Data classes
- ▶ Inhibit reclaim schedules
- ▶ Physical volume pools
- ▶ Move/Eject physical volumes
- ▶ Move logical volumes
- ▶ Operator interventions
- ▶ Physical library status
- ▶ Volume search
- ▶ Modify logical volumes
- ▶ Physical volume ranges
- ▶ Shutdown
- ▶ Backup and recovery

See 8.2, “TS7740 Virtualization Engine management interface” on page 415 for more details about the management interface.

TS7720 Virtualization Engine support

The Integrated Library Manager also applies to TS7720 Virtualization Engine, which means that all management functions are included in the TS7720 Virtualization Engine. The TS7720 Virtualization Engine handles the virtualized library interface to the attached System z hosts. The management interface disables or removes all physical tape and library related functions that are not applicable to the TS7720 Virtualization Engine.

2.5 TS7700 Virtualization Engine Concepts

In this section, we explain the main management concepts of the TS7700 Virtualization Engine. Although most of these concepts apply to both TS7720 Virtualization Engine and TS7740 Virtualization Engine, certain functions apply only to the TS7740 Virtualization Engine, because the TS7720 Virtualization Engine uses a disk-only configuration.

TS7720 Virtualization Engine accepts all storage constructs that are provided by the host to make configuration and support easier. TS7700 Virtualization Engine R1.7 provides enhanced cache management options for performance tuning in TS7720 Virtualization Engine hybrid and multi-cluster grid configurations.

2.5.1 Buffering data in the Tape Volume Cache

The Tape Volume Cache (TVC) of the TS7700 Virtualization Engine is the key element that enables both high-speed tape storage processing and the use of high-capacity tape technology. Buffering of host-created volumes on disk significantly improves the host tape processing response time to near disk speeds in both the TS7740 Virtualization Engine and the TS7720 Virtualization Engine. TS7740 Virtualization Engine subsequently stacks the logical volumes on high-capacity tape, which allows use of the full cartridge capacity that is provided by current tape technology.

In multi-cluster configurations, the TS7700 Virtualization cache resources are accessible by all participating clusters. The architecture allows for any logical volume in cache to be accessed by any cluster through the common Grid network. This capability results in the creation of a Composite Library effective cache size that is close to the sum of all Grid cluster cache capacities.

The TS7700 Virtualization Engine Tape Volume Cache is a disk buffer to which emulated tape volumes are written in 3490E format. For TS7740 Virtualization Engine configurations, the emulated tape volumes are then copied to physical tape cartridges. The host operating system sees virtual IBM 3490E Tape Drives, and the 3490 tape volumes are represented by storage space in a fault-tolerant disk subsystem. All host interaction is through the virtual control unit presented by the vNode. The host never writes directly to the physical tape drives attached to a TS7740 Virtualization Engine.

While resident in the tape volume cache, the user data is protected by RAID 5 (Fibre Channel drives) for TS7740 Virtualization Engine configurations, and RAID 6 (SATA drives) for TS7720 Virtualization Engine configurations. These RAID configurations provide continuous data availability to users. If one data disk in a RAID group becomes unavailable, the user data can be recreated dynamically from the remaining disks using parity data provided by the RAID implementation. The RAID groups contain hot spare disks to take the place of a failed hard disk. Using parity, the RAID controller rebuilds the data from the failed disk onto the hot spare as a background task. This allows the TS7700 Virtualization Engine to continue working while the IBM SSR replaces the failed hard disk in the TS7700 Virtualization Engine Cache Controller or Cache Drawer.

In addition to enabling the full use of high-capacity tape cartridges in a TS7740 Virtualization Engine, the following benefits are offered:

- ▶ Emulated 3490E volumes are accessed at disk speed. Tape commands such as space, locate, rewind, and unload are mapped into disk commands that are completed in tens of milliseconds rather than the tens of seconds required for traditional tape commands.
- ▶ Multiple, different, emulated 3490E volumes can be accessed in parallel because they physically reside in the tape volume cache. To ensure data integrity, a single virtual volume cannot be shared by different jobs or systems at the same time.

2.5.2 Tape Volume Cache Management

Tape Volume Cache Management allows management of the contents of the Tape Volume Cache (TVC) to improve average mount response time. You can use Data Facility Storage Management Subsystem (DFSMS) policy management to assign virtual volumes to preference groups. When space is needed in cache, the preference groups are used to determine which volumes in the TVC are preferred for removal after they have been copied to tape or other participating Grid clusters. Cache management is also used to improve mount response time by keeping required volumes in cache for longer periods.

TVC management processes

Four processes manage the Tape Volume Cache of the TS7700 Virtualization Engine:

- Premigration Management

This process becomes effective when the amount of TS7740 Virtualization Engine TVC data that is not copied to tape reaches a predefined threshold. It is intended to make sure that the TVC does not become completely full of data that has not been backed up to physical tape. It is the mechanism that takes the TS7740 Virtualization Engine from peak mode to sustained mode.

- Free-space Management

This process becomes effective when the amount of unused (free) TVC space reaches another predetermined threshold. It is intended to make sure that the TVC does not become completely full of data, copied to physical tape or not. It is the mechanism that keeps the input to the TVC from overrunning the available free space.

- Copy Management

This process applies only to a Grid configuration and becomes effective when the amount of un-copied data in the TVC reaches a predefined threshold. It applies in particular to deferred copy mode, and when invoked reduces the incoming host data rate independently of premigration or free-space management. Its purpose is to prevent logical volumes from being migrated to physical tape prior to being copied to one or more other TS7700 Virtualization Engine Clusters. This is done to avoid a possible recall operation from being initiated by remote clusters in the Grid. This process is also called *Copy Throttling*.

- Copy Time Management

This process applies to grid configurations where the rewind/unload (RUN) copy consistency point is used. When enabled, it limits the host input rate. It is intended to prevent any RUN copies from exceeding the Missing Interrupt Handler (MIH) timeout value for host jobs. In the event that limiting the host input does help, the TS7700 Virtualization Engine will allow the job to succeed before the MIH timer expires. In the event that limiting the host input does not help, the job will downgrade to deferred mode and an alert is posted to the host console that the TS7700 Virtualization Engine has entered the immediate-deferred state. You can modify this setting through the Host Console Request function to customize the level of throttling applied to the host input when this conditions is detected.

Preference level 0

Preference level 0 (Preference Group 0) is assigned to volumes that are unlikely to be accessed after being created, for example, volumes holding DASD image copies. There is no need to keep them in cache any longer than is necessary to copy them to physical tape. Informal studies suggest that the proportion of data that is unlikely to be accessed can be as high as 80%.

When a volume is assigned preference level 0, the TS7740 Virtualization Engine gives it preference to be copied to physical tape. When space is needed in the TVC, the TS7740 Virtualization Engine will first select a preference level 0 volume that has been copied to a physical volume and delete it from cache. Preference level 0 volumes are selected by largest size first, independent of how long they have been resident in cache. If no preference level 0 volumes have been copied to physical volumes to remove, the TS7740 Virtualization Engine will select preference level 1 (PG1) volumes.

In addition to removing preference level 0 volumes from cache when space is needed, the TS7740 Virtualization Engine will also remove them if the subsystem is relatively idle. There is a small amount of internal processing impact to removing a volume from cache, so there is

some benefit in removing them when extra processing capacity is available. In the case where the TS7740 Virtualization Engine removes PG0 volumes during idle times, it selects them by smallest size first.

Preference level 1

Preference level 1 is assigned to volumes that are likely to be accessed after being created. An example of this are volumes that contain master files created as part of the nightly batch run. Because the master files are likely to be used as input for the next night's batch run, it is beneficial for these volumes to stay in the TVC for as long as possible.

When a volume is assigned preference level 1, the TS7740 Virtualization Engine adds it to the queue of volumes to be copied to physical tape after a four minute time delay and after any volumes assigned to preference level 0. The four minute time delay is to prevent unnecessary copies from being performed when a volume is created, then quickly remounted and appended to again. When space is needed in cache, the TS7740 Virtualization Engine will first determine if there are any preference level 0 volumes that can be removed. If not, the TS7740 Virtualization Engine selects preference level 1 volumes to be removed based on a "least recently used" algorithm. This results in volumes that have been copied to physical tape and have been in cache the longest without access to be removed first.

Figure 2-28 shows cache utilization with policy-based cache management.

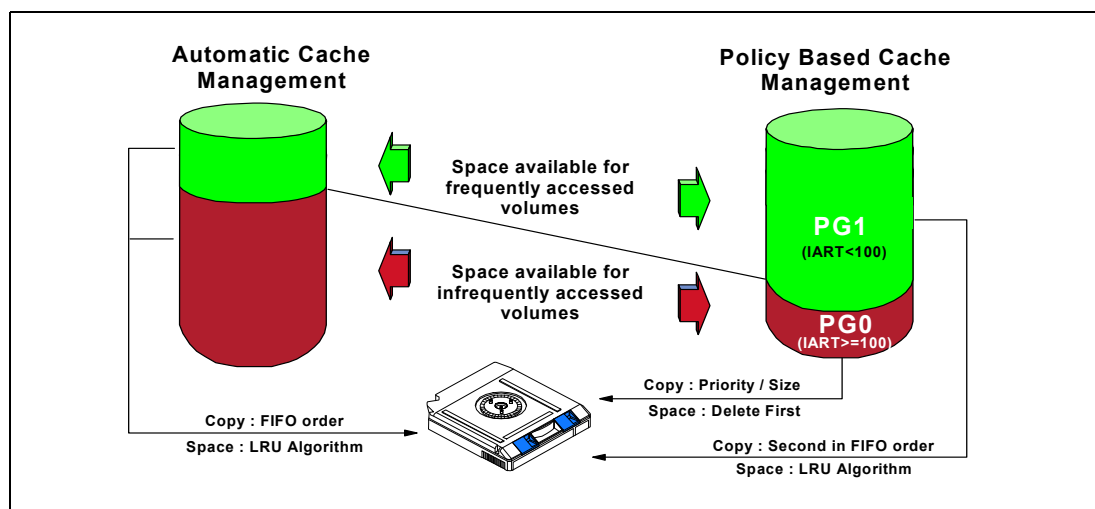


Figure 2-28 Cache utilization with policy-based cache management

When a preference level has been assigned to a volume, that assignment is persistent until the volume is reused for scratch and a new preference level is assigned. This means that a volume that is assigned a preference level 0 will maintain that preference level when it is subsequently recalled into cache.

The Storage Class name assigned to a volume in the ACS routine is directly passed to the TS7700 Virtualization Engine. Figure 2-29 on page 60 shows this process. For non-DFSMS environments, using the management interface, a Storage Class can be assigned to a range of logical volumes during insert processing. The Storage Class can also be updated through the management interface.

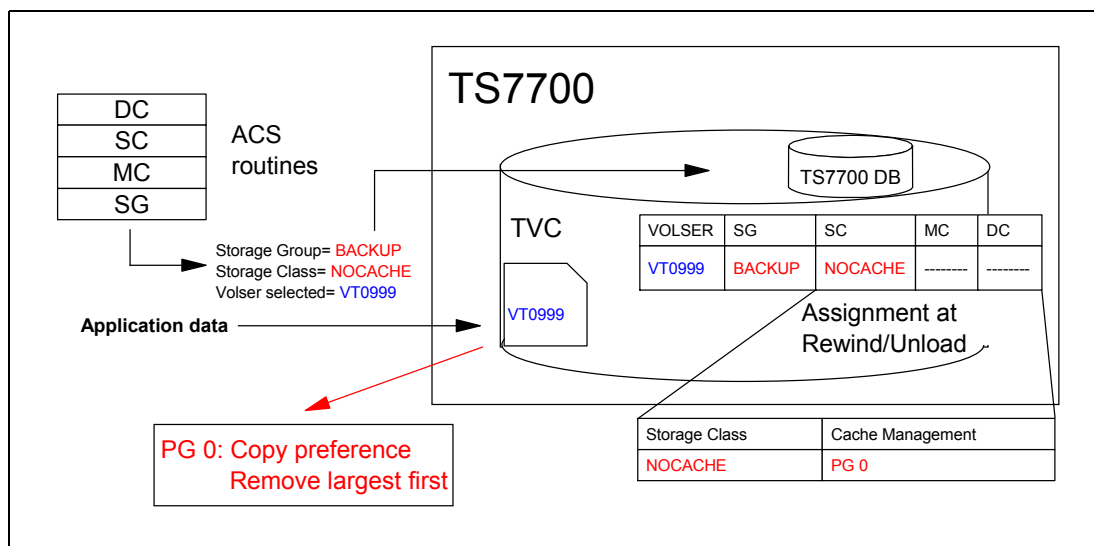


Figure 2-29 TVC Management through Storage Class

Through the management interface, you can define one or more Storage Class names and assign preference level 0 or 1 to them. To be compatible with the IART method of setting the preference level, the Storage Class definition also allows a *Use IART* selection to be assigned.

Even before Outboard Policy Management was made available for the VTS, you had the ability to assign a preference level to virtual volumes by using the Initial Access Response Time Seconds (IART) attribute of the Storage Class. The IART is a Storage Class attribute that was originally added to specify the desired response time (in seconds) for an object using the Object Access Method (OAM). If you wanted a virtual volume to remain in cache, you assign a Storage Class to the volume whose IART value is 99 seconds or less. Conversely, if you wanted a virtual volume preferenced to be out of cache, you assign a Storage Class to the volume whose IART value was 100 seconds or more.

Assuming that the *Use IART* selection has not been specified, the TS7700 Virtualization Engine sets the preference level for the volume based on the preference level 0 or 1 of the Storage Class assigned to the volume.

If the host passes a previously undefined Storage Class name to the TS7700 Virtualization Engine during a scratch mount request, the TS7700 Virtualization Engine will add the name using the definitions for the default Storage Class.

TS7720 Volume Removal

The TS7720 Logical Volume Removal Policy is enhanced with TS7700 Virtualization Engine R1.7. Initially used to support hybrid configurations, the new function provides additional tuning capabilities in hybrid and homogenous TS7720 Virtualization Engine configurations. The TS7720 Virtualization Engine cache configurations can significantly vary between participating clusters in a Grid. The enhancements increase the flexibility of the subsystem effective cache in responding to changes in z/OS workload. The Logical Volume Removal Policy settings can be coupled with Copy Consistency Policy settings to provide an automatic data migration function. To ensure data availability, a key criteria for a execution of the removal policy requires that two available copies of the logical volume must remain in the Grid.

With the Logical Volume Removal Policy additional controls are provided to manage if, how, and when data is removed. These new options are defined through the TS7720 Virtualization Engine management interface and are associated with the Storage Class construct. These options are as follows:

- **Pinned**

The copy of the volume is never removed from this TS7720 cluster. This feature must be used cautiously to prevent a TS7720 cache from filling.

- **Prefer Remove - When Space is Needed**

A copy of a Preference Group 0 private volume is removed in least-recently-used order if the available free space on the cluster has fallen below the removal threshold and the minimum retention time has elapsed since the last access.

- **Prefer Keep - When Space is Needed**

A copy of a Preference Group 1 private volume is removed in least-recently-used order if the available free space on the cluster has fallen below the removal threshold and the minimum retention time has elapsed since last access.

A requirement of the Logical Volume Removal Policy capability is that all clusters participating in the Grid are at TS7700 Virtualization Engine R1.7.

In addition to the policies, volumes that are assigned to a Fast Ready category, and that have not been previously delete-expired, are also removed from cache when the free space on a cluster has fallen below the removal threshold.

Host Command Line Query Capabilities are supported that help override automatic removal behavior and the ability to disable automatic removal within a TS7720 cluster.

TS7720 Cache Full Redirection

TVC selection in a hybrid configuration remains the same as for a homogenous TS7720 Virtualization Engine configuration until a cluster in the Grid enters the *Out of Cache Resources* state. TS7700 Virtualization Engine R1.6 Grid configurations can fail fast-ready mounts to any cluster in the domain if any TS7720 Virtualization Engine enters the Out of Cache Resources state. In this state, all non-fast ready mounts in the Grid had read-only access until the Out of Cache Resources condition was resolved.

The TS7700 Virtualization Engine R1.7 handling of a mount into a Grid configuration where one or more TS7720 Virtualization Engine clusters have crossed their Out of Cache Resources thresholds has been enhanced. During TVC selection in a Grid domain with a cluster in the Out of Cache Resources state, all TVC candidates are considered, even when the mount point cluster is in the Out of Cache Resources state. The Grid function will choose an alternate TS7700 cluster with a valid consistency point and, if dealing with a TS7720, available cache space. Fast-ready mounts involving a TVC candidate that is Out of Cache Resource will fail only if no other TS7700 cluster is eligible to be a TVC candidate. Non-Fast Ready mounts choose the TVC candidate that is in the Out of Cache Resources state only when no other TS7700 cluster is eligible to be a TVC candidate. As with R1.6, when all participating TVCs are in the Out of Cache Resources state, non-Fast Ready mounts will be mounted with read only access.

When the only TVC candidate is either in the *Paused* or *Out of Physical Scratch* resource state, just as is done for R1.6, the mount process waits for either the host to cancel the mount or for the TVC candidate to exit the undesirable state. If the mount process is to enter the wait state for these reasons, and one of the other TVC candidates was disqualified because it was Out of Cache Resources, the cluster that is Out of Cache Resources is included in the cluster mask of TVC candidates to notify the waiting mount process of a state change. The mount

process is resumed when a TVC candidate exits the Paused, Out of Physical Scratch, or Out of Cache Resources state.

The inherent risk is that if the TS7720 Virtualization Engine in the Out of Cache Resources state is not remedied expeditiously, the number of pending copies targeting a TS7720 Virtualization Engine in the Out of Cache Resources state can get quite large and might affect overall Grid performance.

Any mount, issued to a cluster that is in the Out of Cache Resources state and also has Copy Policy Override set to Force Local Copy, will be failed. The Force Local Copy setting excludes all other candidates from TVC selection.

Business continuance settings

For customers who have installed or are considering the installation of a multi-cluster grid for business continuance, the standard management algorithms of the subsystem might not meet all of their business or recovery objectives. To better meet those objectives, the following aspects of cache management have been integrated into the TS7700 Virtualization Engine workflow management functions and can be modified through the z/OS Host Console Request function.

Copy files preferred to reside in cache

Normally all caches in a multi-cluster grid are managed as one composite cache. This way increases the likelihood that a needed volume will be in a TVC by increasing the overall effective cache capacity. By default, the volume on the TVC selected for I/O operations is preferred to reside in the cache on that cluster, whereas the copy made on the other clusters is preferred to be removed from cache. *Preferred to reside in cache* means that when space is needed for new volumes, the oldest volumes are removed first (least recently used algorithm). *Preferred to remove from cache* means that when space is needed for new volumes, the largest volumes are removed first, regardless of when they were written to the cache.

For example, in a two-cluster grid, when you have set up a copy consistency point policy of *RUN, RUN* and the host has access to all virtual devices in the Grid, the selection of virtual devices combined with I/O TVC selection criteria automatically balances the distribution of original volumes and copied volumes across the TVCs. The original volumes (newly created or modified) will be preferred to reside in cache, while the copies will be preferred to be removed from cache. The result is that each TVC is filled with unique newly created or modified volumes, thereby roughly doubling the effective amount of cache the host sees.

For a multi-cluster grid that is used for remote business continuation, particularly when the local clusters are used for all I/O (remote virtual devices varied offline), the default cache management method might not be desired. In the case where the remote cluster of the Grid is used for recovery, the recovery time is minimized by having most of the needed volumes already in cache. What is needed is to have the most recently copied volumes remain in cache, not be preferred out of cache.

Based on your requirements, you can set or modify this control through the z/OS Host Console Request function for the remote cluster:

- ▶ When off, which is the default, logical volumes copied into the cache from a peer TS7700 Virtualization Engine are managed as PG0 (preferred to be removed from cache).
- ▶ When on, logical volumes copied into the cache from a peer TS7700 Virtualization Engine are managed using the actions defined for the Storage Class construct associated with the volume as defined at the TS7700 Virtualization Engine receiving the copy.

If you define the Storage Class construct names as the same on all TS7740 Virtualization Engines in the Grid configuration, the copied volumes will then be managed in the cache, the same as the original volumes. For example, the storage group constructs can be SCBACKUP=Pref Group 1, SCARCHIV=Pref Group 0. All logical volumes written that specify SCARCHIV are treated as PG0 in both the local and remote (copy) caches. All logical volumes written that specify SCBACKUP are treated as PG1 in both the local and remote caches.

Volumes that are written to an I/O TVC that is configured for PG0 have priority with respect to peer TS7700 Virtualization Engine replication priority. Therefore, copy queues within TS7700 Virtualization Engine Clusters will handle volumes with I/O TVC PG0 assignments prior to volumes configured as PG1 within the I/O TVC. This behavior is designed to allow those volumes marked as PG0 to be flushed from cache as quickly as possible and therefore not left resident for replication purposes. This behavior overrides a pure first-in first-out (FIFO) ordered queue.

Recalls preferred for cache removal

Normally, a volume recalled into cache is managed as though it were newly created or modified because it resides in the TVC selected for I/O operations on the volume. A recalled volume displaces other volumes in cache. If the remote cluster of a Grid is used for recovery, the recovery time is minimized by having most of the needed volumes in cache.

However, an unlikely situation is that all of the volumes that are used to restore will be resident in the cache and that some number of recalls will be required. Unless the customer can explicitly control the sequence of volumes used for restored, recalled volumes will likely displace cached volumes that have not yet been restored from, resulting in further recalls at a later time in the recovery process. After a restore process is completed from a recalled volume, that volume is no longer needed. A method is needed with which to remove the recalled volumes from the cache after they have been accessed so that they minimally displace other volumes in the cache.

Based on your requirements, you can set or modify this control through the z/OS Host Console Request function on the remote cluster:

- ▶ When off, which is the default, logical volumes that are recalled into cache are managed by using the actions defined for the Storage Class construct associated with the volume as defined at the TS7700 Virtualization Engine.
- ▶ When on, logical volumes that are recalled into cache are managed as PG0 (prefer to be removed from cache). This control overrides the actions that are defined for the Storage Class associated with the recalled volume.

This control is independent of and not affected by cache management controlled through the Storage Class construct. Storage Class cache management affects only how the volume in the I/O TVC is managed.

2.5.3 Logical and stacked volume management

Logical volumes and physical cartridges are managed by the TS7740 Virtualization Engine because it is Library Emulation Capable.

Physical volume pooling

Logical volumes can be assigned to selected Storage Groups. These Storage Groups point to primary storage pools. The pool assignments are stored in the TS7740 Virtualization Engine database. When a logical volume is copied to tape, it is written to a stacked volume that is

assigned to a Storage Pool. The Storage Pool is assigned to the Storage Group construct through the management interface.

Physical pooling of stacked volumes is identified by defining a pool number, as shown in Figure 2-30.

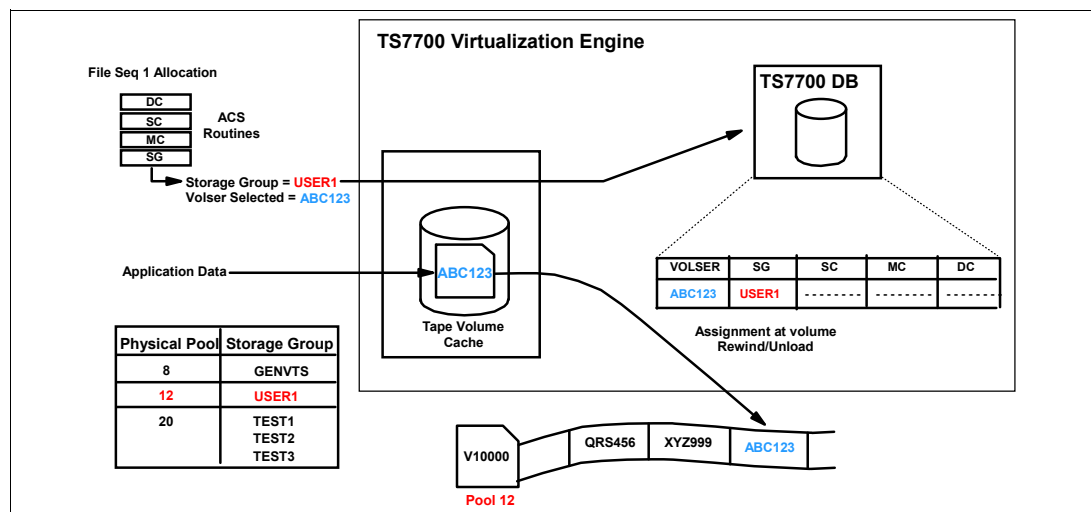


Figure 2-30 Logical volume allocation to specific physical volume pool flow

This definition can be set through the management interface. Each TS7740 Virtualization Engine that is attached to a TS3500 Tape Library has its own set of pools. A Common Scratch Pool (Pool 00) is a reserved pool containing only scratch stacked volumes. There are also 32 general purpose pools (Pools 01-32).

By default, there is one pool, Pool 01, and the TS7740 Virtualization Engine stores virtual volumes on any stacked volume available to it. This creates an intermix of logical volumes from differing sources, for example, LPAR, customer, and applications on a physical cartridge. The user cannot influence the physical location of the logical volume. Physical volume pooling addresses the following concerns:

- ▶ Data from separate customers on the same physical volume can contravene some outsourcing contracts.
- ▶ Customers prefer to be able to “see, feel, and touch” their data by having only their data on dedicated media.
- ▶ Charging for tape is complicated. Traditionally, users are charged by the number of volumes they have in the tape library. With Physical Volume Pooling, users can create and consolidate multiple logical volumes on a smaller number of stacked volumes and therefore reduce their media charges.
- ▶ The TS7740 Virtualization Engine will use any media type available to it. Small logical volumes on the Enterprise Tape Cartridges (JA) or the Enterprise Extended Tape Cartridges (JB) will take a longer time to recall than volumes on the Economy Cartridge (JJ).
- ▶ Protecting data through encryption can be set on a per pool basis. This enables you to encrypt all or some of your data when it is written to the back-end tapes.

Volume pooling allows the administrator to define pools of stacked volumes within the TS7740 Virtualization Engine and the user can direct virtual volumes to these pools through the use of SMS constructs. There can be up to 32 general purpose pools (01-32) and one common pool (00). Pool 00 is a source of scratch stacked volumes for the other pools, which can be

configured to borrow from pool 00. The pool can then return the “borrowed” volume when it becomes scratch or keep it indefinitely. Each pool can be defined to borrow single media (JA, JB, or JJ), mixed media, or have a “first choice” and a “second choice”.

VOLSER ranges can be defined with a home pool at insert time. Changing the home pool of a range has no effect on existing volumes in the library.

Common Scratch Pool (Pool 00)

The Common Scratch Pool (CSP) is a pool that only contains scratch stacked volumes and serves as a reserve pool. As the primary pools run out of scratch stacked volumes, they can borrow scratch stacked cartridges from the CSP (Pool 00) either on a temporary or permanent basis. The borrowing options can be set at the management interface when defining Stacked Volume Pool properties.

General purpose pools (Pools 01-32)

There are 32 general purpose pools available per TS7740 Virtualization Engine cluster. These pools can contain both physical scratch and private stacked volumes. All volumes, private and scratch, are assigned category FF04. When initially creating these pools, it is important to ensure that the correct borrowing properties are defined to each pool. See “Stacked volume pool properties” on page 65 for more information.

Using this facility, you can also perform the following tasks:

- ▶ Move stacked volumes to separate pools.
- ▶ Set a reclamation threshold at the pool level.
- ▶ Set Forced Reclamation policies for stacked volumes.
- ▶ Eject stacked volumes from specific pools.
- ▶ Intermix or segregate media types.
- ▶ Map separate Storage Groups to the same primary pools.

Note: Primary Pool 01 is the default private pool for TS7740 Virtualization Engine stacked volumes.

Borrowing and returning

Using the concept of borrowing and returning, out of scratch scenarios can be automatically addressed. With borrowing, stacked volumes can move from pool to pool and back again to the original pool. In this way, the TS7740 Virtualization Engine can manage out of scratch and low scratch scenarios, which can occur within any TS7740 Virtualization Engine from time to time.

Stacked volume pool properties

Logical volume pooling supports cartridge type selection. This can be used to create separate pools of 3592 tape cartridges with a variety of capacities from 128 GB up to 1 TB, depending upon the type of media and tape drive technology used.

Lower capacity JJ cartridges can be designated to a pool to provide fast access to applications such as HSM or Content Manager. Higher capacity JA or JB cartridges assigned to a pool can address archival requirements such as full volume dumps.

Selective Dual Copy

In a single-cluster grid, a logical volume and its internal data usually exists as a single entity that is copied to a single stacked volume. If the stacked volume is damaged, you can lose access to the data within the logical volume. Without the Selective Dual Copy Function, the only way to ensure data availability is to use host software to duplex the logical volume, or to set up a Grid environment.

With the Selective Dual Copy function, storage administrators have the option to selectively create two copies of logical volumes within any TS7740 Virtualization Engine. In a Grid environment, logical volumes that are copied to peer TS7740 Virtualization Engine clusters can also be duplexed to back-end tape.

The Selective Dual Copy function can be used along with the Copy Export function to provide multiple physical copies of TS7740 Virtualization Engine data for disaster recovery purposes. See 2.5.4, “Copy Export” on page 70 for more details concerning Copy Export.

The copy of the logical volume is created in a second physical pool ensuring physical cartridge separation. Control of Dual Copy is through the Management Class construct (see “Creating Management Classes” on page 227). The second copy is created when the original volume is premigrated.

The second copy created through the Selective Dual Copy function is only available when the primary volume cannot be recalled. It cannot be accessed separately and cannot be used if the primary volume is being used by another application, or if it has been corrupted by an application. In other words, the second copy prevents data loss because of media failures, but it cannot be used as a backup in case of operational or application errors.

Selective Dual Copy is defined to the TS7740 Virtualization Engine and has the following characteristics:

- ▶ The copy feature is enabled by the Management Class setting through the management interface, where you define the secondary pool.
- ▶ Secondary and primary pools can be intermixed:
 - A primary pool for one logical volume can be the secondary pool for another logical volume unless the secondary pool is used as a Copy Export pool.
 - Multiple primary pools can use the same secondary pool.
- ▶ At rewind unload time, the secondary pool assignment is determined and the copy of the logical volume is scheduled. The scheduling of the backup is determined by the premigration activity occurring in the TS7740 Virtualization Engine.
- ▶ The copy is created prior to the primary volume being migrated out of cache.

Mounting a scratch virtual volume

The management interface uses categories to group volumes. After virtual volumes are inserted through the management interface, they are placed in the insert category and handled exactly as native cartridges. When the TS7700 Virtualization Engine is varied online or during insert processing, the host operating system assigns newly inserted volumes to a particular scratch category.

When a request for a scratch is issued to the TS7700 Virtualization Engine, the request specifies a mount category. The TS7700 Virtualization Engine selects a virtual VOLSER from the candidate list of scratch volumes in the category.

Starting with release R1.5, scratch volumes chosen at the mounting cluster are chosen using the following priority:

1. All volumes in the source or alternate source category that are owned by the local cluster, not currently mounted, and not hot against peer clusters.
2. All volumes in the source or alternate source category that are owned by any available cluster, not currently mounted, and not hot against any peer cluster.

3. All volumes in the source or alternate source category that are owned by any available cluster and not currently mounted.
4. All volumes in the source or alternate source category that can be taken over from an unavailable cluster that has an explicit or implied takeover mode enabled.

For all scratch mounts, the volume is temporarily initialized as though the volume had been initialized using the EDGINERS or IEHINITT program and will have an IBM standard label consisting of a VOL1 record, an HDR1 record, and a tape mark. If the volume is modified, the temporary header information is applied to a file in the TVC. If the volume is not modified, the temporary header information is discarded and any previously written content (if it exists) is not modified. Besides choosing a volume, TVC selection processing, as described in “I/O TVC selection” on page 87, is used to choose which TVC will act as the I/O TVC.

Copying the virtual volume to tape

After the host closes and unloads a virtual volume, the storage management software inside the TS7740 Virtualization Engine schedules the virtual volume to be copied (also known as premigration) onto a physical tape cartridge. The TS7740 Virtualization Engine attempts to maintain a mounted stacked volume to which virtual volumes are copied. Therefore, mount activity is reduced because only one physical cartridge is mounted to service multiple virtual volume premigration requests.

Mounting a specific virtual volume

The process followed in the TS7700 Virtualization Engine to mount a specific tape is similar to the scratch mount scenario in that they both choose an I/O TVC (see “I/O TVC selection” on page 87). If the chosen I/O TVC contains a copy of the logical volume in cache, a physical tape mount is not required. In this case, the mount is signaled as complete and the host can access the data immediately.

If the volume has already been removed from the TVC and migrated to a stacked volume, recalling the logical volume from the stacked volume into the TVC is required before the host can directly access the volume. A recall operation typically requires a physical mount unless the stacked volume is already mounted from a previous volume recall. Mount completion is signaled to the host system only after the entire volume is available in the TVC. The virtual volume remains in the TVC until it becomes the least-recently used (LRU) volume, unless the volume was assigned a Preference Group of 0. If modification of the virtual volume did not occur when it was mounted, the VTS does not schedule another copy operation and the current copy of the logical volume on the original stacked volume remains active. Furthermore, copies to remote TS7700 Virtualization Engine Clusters are not required if modifications were not made.

In a z/OS environment, to mount a specific volume in the TS7700 Virtualization Engine, that volume must reside in a private category within the library. The tape management system prevents a scratch volume from being mounted in response to a specific mount request. Alternatively, the TS7700 Virtualization Engine treats any specific mount that targets a volume that is currently assigned to a category, which is also configured through the management interface as Fast Ready, as a host scratch mount. If this occurs, the temporary tape header is created and no recalls will take place, as described in “Mounting a scratch virtual volume” on page 66. In this case, DFSMS OAM usually fails the mount operation because the expected last written data set for the private volume was not found. Because no write operation occurs, the original volume’s contents should be left intact, which accounts for categories being incorrectly configured as Fast Ready within the management interface.

TS7700 Virtualization Engine R1.5, and later, supports Device Allocation Assist (DAA), which in conjunction with z/OS host software capabilities, attempts to determine the cluster that will provide the best performance for a given non-fast ready volume mount request

Expired virtual volumes

When a virtual volume is returned to a Fast Ready category during host software return-to-scratch processing, it can optionally be placed on an expiration list. This is an optionally configured Fast Ready attribute within the management interface. When the delete expired volume attribute is enabled, it allows virtual volume data still residing in any TVC or stacked volume within the Composite Library to be automatically deleted after the time elapses from the point a volume is moved to a Fast Ready category. Prior to the expiration of a volume, the data associated with the logical volume remains available in a TVC or on one or more stacked volumes until the logical VOLSER is reused as a scratch volume or if the expiration with deletion elapses, whichever occurs first. A reconfiguration of the Fast Ready category or the explicit movement of a volume out of the delete expired configured category can occur prior to the expiration of a volume if deletion is not the end goal.

The disadvantage of not having this option enabled is that scratched volumes needlessly consume TVC and physical stacked volume resources, therefore demanding more TVC active space while also requiring more physical stacked volumes in a TS7740 Virtualization Engine. The time it takes a physical volume to fall below the reclamation threshold is also increased because the data is still considered active. This delay in data deletion also causes scratched stale volumes to be moved from one stacked volume to another during reclamation.

With expired volume management, you may set a “grace period” for expired volumes ranging from 24 hours to approximately 144 weeks. After that period has elapsed, expired volumes become candidates for deletion. The deletion of expired logical volume data eliminates the need for the TS7740 Virtualization Engine to manage logical volume data that has been expired at the host. For details about expired volume management, see 4.3.6, “Defining the logical volume expiration time” on page 222.

An additional option referred to as “expire hold” can also be enabled if delete expired is enabled. When this option is also enabled in addition to delete expire, the volume cannot be accessed using any host initiated command until the grace period has elapsed. This additional option is made available to prevent any malicious overwriting or unintended overwriting of scratched data prior to the duration elapsing. After the grace period expires, the volume will simultaneously be removed from a held state and made a deletion candidate.

Expired data on a physical volume remains readable until the volume has been completely overwritten with new data. For example, some customers are concerned that a legal court order might expose them to liability and cost while trying to *find* an older version of a data volume.

Secure Data Erase

Another concern is security of old data. The TS7740 Virtualization Engine provides physical volume erasure on a physical volume pool basis that is controlled by an additional reclamation policy. With Advanced Policy Management, which is standard on a TS7740 Virtualization Engine, all reclaimed physical volumes assigned to that pool are erased with a random pattern prior to being reused. When Secure Data Erase is enabled, a physical cartridge is not available as a scratch cartridge as long as its data is not erased.

The Secure Data Erase function supports the erasure of a physical volume as part of the reclamation process. The erasure is performed by writing a random data pattern on the physical volume at the end of the reclamation task. A random data pattern is written on the physical volume being reclaimed so that the logical volumes that had been written on it are no longer readable. As part of this data erase function, an additional reclaim policy is added. The policy specifies the number of days a physical volume can contain invalid logical volume data before the physical volume becomes eligible to be reclaimed.

The data that is associated with a logical volume is considered invalidated if any of the following statements are true:

- ▶ A host has assigned the logical volume to a scratch category. The volume is subsequently selected for a scratch mount and data is written to the volume. The older version of the volume is now invalid.
- ▶ A host has assigned the logical volume to a scratch category that has the Fast Ready attribute set, the category has a nonzero delete expired data parameter value, the parameter value has been exceeded, and the TS7740 Virtualization Engine has deleted the logical volume.
- ▶ A host has modified the contents of the volume. This could be a complete rewrite of the volume or an append to it. The new version of the logical volume will be migrated to a different physical location and the older version is now invalid.

When a physical volume contains encrypted data, the TS7740 Virtualization Engine is able to perform a fast erase of the data under certain circumstances by “shredding” the encryption keys on the cartridge. This shredding of the encryption keys method is used when a stacked volume was encrypted on its previous use. This reduces the erasure time down to a couple of minutes instead of several hours. The first erasure of an encrypted volume causes the keys to be shredded and a pattern to be written to the tape.

The TS7740 Virtualization Engine keeps track of the amount of active data on a physical volume. It starts at 100% when a volume becomes full. Although the granularity of the percentage of full TS7740 Virtualization Engine tracks is one tenth of 100%, it rounds down, so even *one byte* of inactive data drops the percent to 99.9%. TS7740 Virtualization Engine keeps track of the time that the physical volume went from 100% full to less than 100% full by performing the following tasks:

- ▶ Checking on an hourly basis for volumes in a pool with a non-zero setting
- ▶ Comparing this time against the current time to determine if the volume is eligible for reclamation

This data erase function is enabled on a volume pool basis. It is enabled when a non-zero value is specified for the data erase reclamation policy. When enabled, all physical volumes in the pool are erased as part of the reclamation process, independent of the reclamation policy under which the volume became eligible for reclamation.

Any physical volume that has a status of read-only is not subject to this function and is not designated for erasure as part of read-only recovery.

If a customer uses the eject stacked volume function, no attempt is made to erase the data on the volume prior to ejecting the cartridge. The control of expired data on an ejected volume becomes a customer responsibility.

Volumes tagged for erasure cannot be moved to another pool until erased, but they can be ejected from the library, because such a volume is usually removed for recovery actions.

Using the *Move function* will also cause a physical volume to be erased, even though the number of days specified has not yet elapsed. This includes returning borrowed volumes.

Virtual volume reconciliation and reclamation

Every time a logical volume is modified, the data from the previous use of this logical volume, which is on a stacked volume, becomes obsolete. When you modify a logical volume, you have a new virtual volume in the TVC, and the copy on the stacked volume is invalidated, but still exists in its current state on the physical volume.

The reconciliation process checks for invalidated volumes. A reconciliation is that period of activity by the TS7740 Virtualization Engine when the most recent instance of a logical volume is determined as the active one and all other instances of that volume are deleted from the active volume list. This process automatically adjusts the active data amount for any stacked volumes that hold invalidated logical volumes.

A reconciliation normally occurs sometime during a 24-hour period just before every reclamation process. This makes the reclamation process more effective because it deletes more invalidated logical volumes from the active volume list and therefore reduces the percentage of active data on a greater number of stacked cartridges.

Reclamation, which consolidates active data and frees stacked volumes for return-to-scratch use, is part of the internal management functions of a TS7740 Virtualization Engine.

2.5.4 Copy Export

One of the key reasons to use tape is for recovery of critical operations in the event of a disaster. The TS7740 Virtualization Engine in a Grid configuration provides for automatic remote replication of data, which supports recovery time and recovery point objectives measured in seconds. For customers that do not require the recovery times that can be obtained in a Grid configuration, a new function called Copy Export is being introduced for the TS7740 Virtualization Engine. With Copy Export, logical volumes written to a TS7740 Virtualization Engine can be removed from the TS7740 Virtualization Engine and taken to an off-site location for disaster recovery purposes.

The Copy Export function is supported on all configurations of the TS7740 Virtualization Engine, including Grid configurations. In a Grid configuration, each TS7740 Virtualization Engine is considered a separate source TS7740 Virtualization Engine. This means that only the physical volumes exported from a source TS7740 Virtualization Engine can be used for the recovery of a source TS7740 Virtualization Engine. Physical volumes from more than one source TS7740 Virtualization Engine in a Grid configuration cannot be combined for recovery use. Recovery is only to a single cluster configuration. After recovery, the Grid MES offering can be applied to recreate a Grid configuration.

To perform a Copy Export operation, the TS7740 Virtualization Engine must have a minimum of four available physical tape drives.

The Copy Export function allows a copy of selected logical volumes written to the TS7740 Virtualization Engine to be removed and taken off site for disaster recovery purposes. The benefits of volume stacking, which places many logical volumes on a physical volume, are retained with this function. Because the physical volumes being exported contain copies of logical volumes, the primary logical volume remains accessible to the production host systems.

During a Copy Export operation, all of the physical volumes with active data on them in a specified secondary pool are removed from the library associated with the TS7740 Virtualization Engine. Only the logical volumes that are valid on that TS7740 Virtualization Engine are considered during the execution of the operation. If the virtual volumes are in the TVC, but have not yet been migrated to the secondary pool, copies are made as part of the Copy Export operation. If the TS7740 Virtualization Engine is in a Grid configuration, copies that have not been completed to the TS7740 Virtualization Engine performing the Copy Export are not considered during the execution of the operation. It is expected that Copy Export operations will be run on a periodic basis, resulting in multiple groups of physical volumes that contain copies of logical volumes from the TS7740 Virtualization Engine.

During the Copy Export operation, a copy of the current TS7740 Virtualization Engine's database is written to the exported physical volumes. To restore access to the data on the removed physical volumes, all exported physical volumes for a source TS7740 Virtualization Engine are placed into a library that is attached to an empty TS7740 Virtualization Engine. A disaster recovery procedure is then performed that restores access using the latest copy of the database.

The physical volumes exported during a Copy Export operation continue to be managed by the source TS7740 Virtualization Engine with regard to space management. As logical volumes that are resident on the exported physical volumes expire, are rewritten, or otherwise invalidated, the amount of valid data on a physical volume will decrease until the physical volume becomes eligible for reclamation based on the customer's provided criteria for its pool. Exported physical volumes that are to be reclaimed are not brought back to the source TS7740 Virtualization Engine for processing. Instead, a new secondary copy of the remaining valid logical volumes is made using the primary logical volume copy as a source.

The next time the Copy Export operation is performed, the physical volumes with the new copies are also exported. The physical volumes that were reclaimed (that are off site) no longer are considered to have valid data and can be returned to the source TS7740 Virtualization Engine to be used as new scratch volumes.

The host that initiates the Copy Export operation first creates a dedicated *export list volume* on the TS7740 Virtualization Engine that will perform the operation. The export list volume contains instructions regarding the execution of the operation and a reserved file that the TS7740 Virtualization Engine will use to provide completion status and export operation information. As part of the Copy Export operation, the TS7740 Virtualization Engine creates response records in the reserved file that lists the logical volumes exported and the physical volume on which they reside. This information might be used by the customer as a record for what data is off site. The TS7740 Virtualization Engine also writes records in the reserved file on the export list volume that provides the current status for all physical volumes with a state of Copy Exported.

The Bulk Volume Information Retrieval (BVIR) function can also be used to obtain a current list of exported physical volumes for a secondary pool. For each exported physical volume, information is available on the amount of active data that each cartridge contains.

We now describe the basic flow of the Copy Export function. Figure 2-31 on page 72 shows the flow of the Copy Export volumes and how they are used in a disaster recovery. We discuss a disaster recovery scenario where the production site is lost and the latest set of Copy Export volumes is used to restore the TS7740 Virtualization Engine at a remote site.

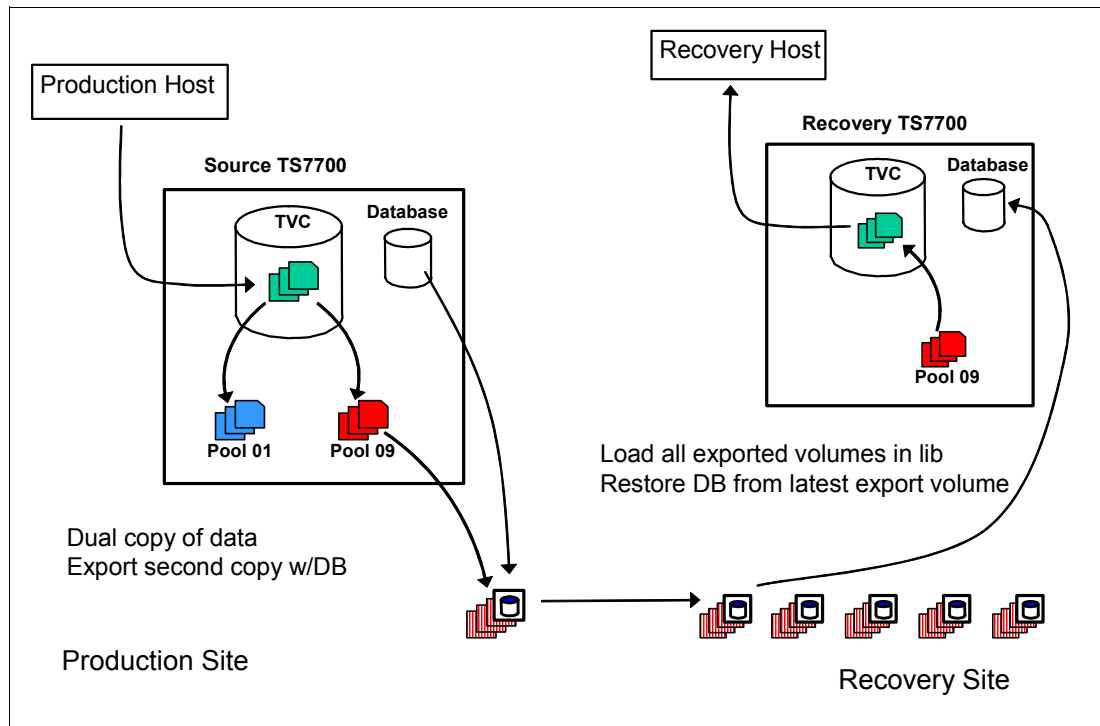


Figure 2-31 Copy Export overview

Storage Group and Management Class constructs are defined to use separate pools for the primary (pool 01) and secondary (pool 09) copies of the logical volume. The existing Management Class construct, which is part of Advanced Policy Management (APM), is used to create a second copy of the data to be Copy Exported. The Management Class actions are configured through the TS7700 Virtualization Engine management interface. An option on the management interface panel allows designation of a secondary pool as a Copy Export pool. As logical volumes are written, the secondary copy of the data is premigrated to stacked volumes in the Copy Export pool. Our example uses pool 09 as the Copy Export pool.

When the time comes to initiate a Copy Export, a Copy Export job is run from the production host. The TS7740 Virtualization Engine will premigrate any logical volumes in the Copy Export pool that have not been premigrated. Any new logical volumes written after the Copy Export operation is initiated will not be included in the Copy Export set of physical volumes. The TS7740 Virtualization Engine then writes a complete TS7740 Virtualization Engine database to each of the physical volumes in the Copy Export set.

The Copy Export job can specify whether the stacked volumes in the Copy Export set should be ejected immediately or placed into the export-hold category. When Copy Export is used with the export-hold category, the operator will need to manually request that the export-hold volumes be ejected. The choice to eject as part of the Copy Export job or to eject them later from the export-hold category will be based on customer operational procedures. The ejected Copy Export set will then be transported to a disaster recovery site or vault. The customer's Recovery Point Objective (RPO) will determine the frequency of the Copy Export operation.

A schedule is implemented for rotation of the Copy Export sets at the disaster recovery site. This will periodically return sets of stacked volumes to the production system TS7740 Virtualization Engine, where they are reintegrated into the library. The TS7740 Virtualization Engine will reconcile the logical volumes on the returning stacked volumes based on activity in the Composite Library while they were exported to the disaster recovery site.

In the event of a disaster, a recovery process is performed through the TS7740 Virtualization Engine management interface on an empty TS7740 Virtualization Engine. As part of that recovery process, all of the source TS7740 Virtualization Engine Copy Exported stacked volumes are inserted into the target library. If there are multiple pools that have been exported, one pool is recovered at a time. The stacked volume with the latest copy of the source TS7740 Virtualization Engine's database needs to be identified to the target TS7740 Virtualization Engine. The target TS7740 Virtualization Engine then restores the database from this stacked volume. At this point, the disaster recovery host can start operations.

2.5.5 Encryption

The importance of data protection has become increasingly apparent with news reports of security breaches, loss and theft of personal and financial information, and government regulation. Encryption of backstore tapes minimizes the risk of unauthorized data access without excessive security management burdens or subsystem performance issues.

The encryption solution for tape virtualization consists of several components.

IBM tape encryption solutions use the Tivoli Key Lifecycle Manager server as a central point from which all encryption key information is managed and served to the various subsystems. The Tivoli Key Lifecycle Manager server communicates with the Virtualization Engine TS7740 Virtualization Engine and tape libraries, control units, and Open Systems device drivers.

The TS1120 and TS1130 encryption-enabled tape drives provide hardware that performs the cryptography function without reducing the data transfer rate.

The TS7740 Virtualization Engine provides the means to manage the use of encryption and what keys are used on a storage pool basis. It also acts as a proxy between the tape drives and the Tivoli Key Lifecycle Manager servers, using Ethernet to communicate with the Tivoli Key Lifecycle Manager servers and Fibre Channel connections to communicate with the drives. Encryption must be enabled in each of the tape drives.

Encryption on the TS7740 Virtualization Engine is controlled on a storage pool basis. The Storage Group DFSMS construct that is specified for a logical tape volume determines which storage pool is used for the primary and optional secondary copies in the TS7740 Virtualization Engine. The storage pools were originally created for management of physical media and have been enhanced to include encryption characteristics. Storage pool encryption parameters are configured through the TS7740 Virtualization Engine management interface under Physical Volume Pools.

For encryption support, all drives that are attached to the TS7740 Virtualization Engine must be encryption-capable and encryption must be enabled. If TS7740 Virtualization Engine uses TS1120 Tape Drives, they must also be enabled to run in their native E05 format. The management of encryption is performed on a physical volume pool basis. Through the management interface, one or more of the 32 pools can be enabled for encryption.

With TS7700 Virtualization Engine R1.7, each pool can be defined to use *specific encryption keys* or the *default encryption keys* defined at the TKLM server:

- Specific encryption keys

Each pool that is defined in the TS7740 Virtualization Engine can have its own unique encryption key. As part of enabling a pool for encryption, enter two key labels for the pool and an associated key mode. The two keys may or may not be the same. Two keys are required by the Tivoli Key Lifecycle Manager servers during a key exchange with the drive. A key label can be up to 64 characters. Key labels do not have to be unique per pool and the management interface provides the capability to assign the same key label to multiple

pools. For each key, a key mode can be specified. The supported key modes are Label and Hash. As part of the encryption configuration through the management interface, you provide s IP addresses for a primary and an optional secondary key manager.

► Default encryption keys

As an extension of TS7740 Virtualization Engine encryption capabilities, TS7740 Virtualization Engine R1.7 supports the use of a default key. This support simplifies the management of the encryption infrastructure because no future changes are required at the TS7740 Virtualization Engine. After a pool is defined to use the default key, the management of encryption parameters is performed at the key manager:

- Creation and management of encryption certificates
- Device authorization for key manager services
- Global Default key definitions
- Drive level Default key definitions
- Default key changes as required by security policies

For logical volumes that contain data that is to be encrypted, host applications direct them to a specific pool that has been enabled for encryption using the Storage Group construct name. All data directed to a pool enabled for encryption will be encrypted when it is premigrated to the physical stacked volumes. The storage group construct name is bound to a logical volume when it is mounted as a scratch volume. Through the management interface, the storage group name is associated with a specific pool number. When the data for a logical volume is copied from the TVC to a physical volume in an encryption enabled pool, the TS7740 Virtualization Engine determines whether a new physical volume needs to be mounted. If a new cartridge is required, the TS7740 Virtualization Engine directs the drive to use encryption during the mount process. The TS7740 Virtualization Engine also provides the drive with the key labels specified by the customer for that pool. When the first write data is received by the drive, a connection is made to a key manager and the key needed to perform the encryption is obtained. Physical scratch volumes are encrypted with the keys in effect at the time of first write to beginning of tape (BOT). Any partially-filled physical volumes continue to use the encryption settings in effect at the time the tape was initially written from BOT. The encryption settings are static until the volumes are reclaimed and rewritten again from BOT.

Figure 2-32 on page 75 illustrates that the method for communicating with a key manager is through the same Ethernet interface that is used to connect the TS7740 Virtualization Engine to a customer's network for access to the management interface. The request for an encryption key is directed to the IP address of the primary key manager. Responses are passed through the TS7740 Virtualization Engine to the drive. If the primary key manager did not respond to the key management request, the optional secondary key manager IP address is used. After the TS1120 or TS1130 drive has completed the key management communication with the key manager, it accepts data from the TVC.

When a logical volume needs to be read from a physical volume in a pool enabled for encryption, either as part of a recall or reclamation operation, the TS7740 Virtualization Engine uses the key manager to obtain the necessary information to decrypt the data.

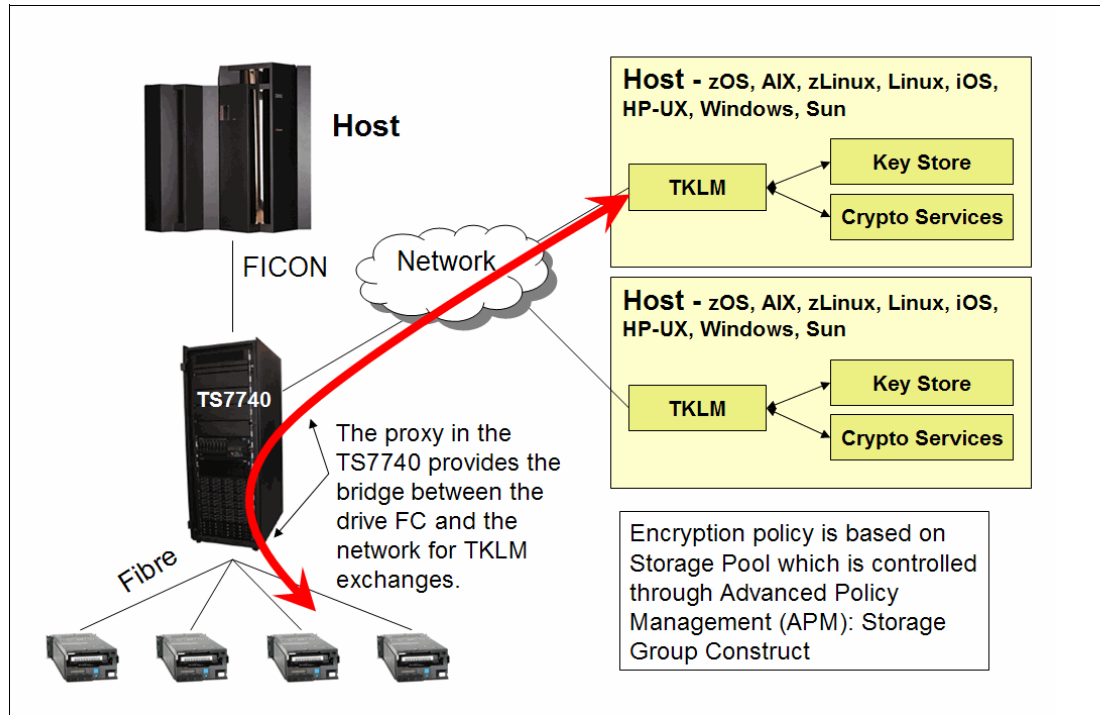


Figure 2-32 TS7740 Virtualization Engine encryption

The affinity of the logical volume to a specific encryption key or the default key can also be used as part of the search criteria through the TS7700 Virtualization Engine management interface.

2.5.6 Logical WORM (LWORM) support and characteristics

TS7700 Virtualization Engine, Release 1.6 and later, supports the logical write once, read many (WORM) function through TS7700 Virtualization Engine software emulation. The z/OS host software supports physical WORM media prior to z/OS Version 1.

As a result of supporting 3592 WORM compatible drives and media, the logical WORM enhancement can duplicate most of the 3592 WORM behavior, allowing ease of host integration. The host views the TS7700 Virtualization Engine as an LWORM-compliant library that contains WORM-compliant 3490E logical drives and media.

The LWORM implementation of the TS7700 Virtualization Engine emulates physical 3490E WORM tape drives and media, and uses the existing host software support of physical WORM media. TS7700 Virtualization Engine provides the following functions:

- ▶ Indicates to any attached host systems that the TS7700 Virtualization Engine Composite Library supports LWORM volumes only when all clusters in a grid are at R1.6 or later.
- ▶ Provides an advanced function Data Class construct property that allows volumes to be assigned as LWORM-compliant during the volume's first mount, where a write operation from BOT is required, or during a volume's reuse from scratch, where a write from BOT is required.
- ▶ Generates, during the assignment of LWORM to a volume's characteristics, a temporary World Wide Identifier that is surfaced to host software during host software open and close processing, and then bound to the volume during first write from BOT.

- ▶ Generates and maintains a persistent Write-Mount Count for each LWORM volume and keeps the value synchronized with host software.
- ▶ Allows only appends to LWORM volumes using physical WORM append guidelines.
- ▶ Provides a mechanism through which host software commands can discover LWORM attributes for a given mounted volume.

No method is available to convert previously written volumes to LWORM volumes without having to read the contents and rewrite them to a new logical volume that has been bound as an LWORM volume.

Note: Cohasset Associates, Inc. has assessed the logical WORM capability of the TS7700 Virtualization Engine. The conclusion is that the TS7700 Virtualization Engine meets all SEC requirements in Rule 17a-4(f), which expressly allows records to be retained on electronic storage media.

2.5.7 Workflow Management Enhancements

Workflow Management Enhancements were introduced with TS7700 Virtualization Engine Release 1.5. They provide the user with the ability to adjust the internal management and alert levels of the TS7700 Virtualization Engine.

These new controls are set and retrieved by using the TS7700 Virtualization Engine z/OS Host Console Request function. Settings are persistent, which means they are maintained even if the system restarts or has code upgrades.

Workflow Management Enhancements include alert controls and messages, cache controls, and throttling controls:

- ▶ Alert messages

Provides a mechanism to present messages to the host console when various metrics have crossed a customer set boundary. Alerts are defined at two levels configured by the customer: warning and critical. Messages are sent when the configured metric crosses the warning or critical levels, and when the system drops below the threshold level.
- ▶ Cache controls

Allows the customer to tailor how the TS7700 Virtualization Engine manages the data flow through the disk cache within the subsystem to better meet their objectives.
- ▶ Throttling controls

Allows the customer to tailor when and to what levels the TS7700 Virtualization Engine paces host I/O or replication through the throttling mechanism. See “Host Write Throttle” on page 609 for more details about throttling.

Supported Workflow Management Enhancements are as follows:

- ▶ Enable/Disable copies to follow storage class preference
- ▶ Set the premigration priority threshold
- ▶ Set the premigration throttling threshold
- ▶ Enable/Disable recalls preferred to be removed from cache
- ▶ Enable/Disable full cache copy throttling
- ▶ Set the deferred copy throttling value
- ▶ Enable/Disable immediate mode copy throttling

For each managed TS7700 Virtualization Engine resource, there are two alert thresholds that can be set. For each threshold, a message is sent to all attached hosts when the threshold is

exceeded and when the resource falls back into the acceptable range. The resources that are monitored are:

- ▶ The amount of data in cache that needs to be copied to a peer cluster
- ▶ The amount of data resident in the cache (for a TS7740 Virtualization Engine, this is also the amount of data that needs to be copied to physical tape)
- ▶ Number of logical scratch volumes
- ▶ Number of physical scratch volumes
- ▶ Number of available physical tape drives

Host Console Request function

The Host Console Request function can help you request information about a subsystem's status and its resources. This command also allows the user to perform an operation in the TS7700 Virtualization Engine. The Host Console Request function can be used with automation software to obtain and analyze operational data that can then be used to alert a storage administrator that further investigation might be warranted.

This command allows z/OS console operators to perform basic problem determination without having to access the web-based management interface. This command is also intended to allow you to automate monitoring of the physical resources of the TS7700 Virtualization Engine using z/OS facilities.

An additional function, the Host Copy Control, which is an extension of the Host Console Request function, allows host control of logical volume copying on a cluster basis.

See 8.3.3, "Host Console Request function" on page 544 for more details concerning the Host Console Request function.

2.6 TS7700 Virtualization Engine multi-cluster grid

When multiple clusters are connected to a Grid, logical volumes and their copies, if configured for copies, reside on some or all of the other clusters. In this section, we explain logical volume management and multi-cluster grid configuration options.

To configure a Multiclust er Grid, the Grid enablement feature must be installed on all TS7700 Virtualization Engines in the Grid, and an Ethernet connection must be established between the clusters. Each cluster has two dual-port copper RJ45 1 GBps Ethernet adapters or two optical fiber (short wave) adapters, as shown in 2.3.1, "Common components for the TS7720 Virtualization Engine and TS7740 Virtualization Engine models" on page 36. Earlier TS7740 Virtualization Engines that have been upgraded to R1.7 might still have the single-port adapters installed.

2.6.1 Data integrity by volume ownership

Any logical volume, or any copies of it, can be accessed by a host from any virtual device participating in a common Grid, even if the cluster associated with the virtual device does not have a local copy. The access is subject to *volume ownership*. At any point in time, a logical volume is "owned" by only one cluster. The owning cluster has control over access to the volume and over changes to the attributes associated with the volume (such as volume category or storage constructs).

The cluster that has ownership of a specific logical volume can change dynamically. When required, the TS7700 Virtualization Engine node transfers the ownership of the logical volume as part of mount processing. This action ensures that the cluster with the virtual device associated with the mount has ownership. When a mount request is received on a virtual device address, the TS7700 Virtualization Engine Cluster for that virtual device must have ownership of the volume to be mounted or must obtain the ownership from the cluster that currently owns it.

If the TS7700 Virtualization Engine Clusters in a Grid and the communication paths between them are operational, the change of ownership and the processing of logical volume-related commands are transparent to the host. If a TS7700 Virtualization Engine Cluster has a host request for a logical volume that it does not own and it cannot communicate with the owning cluster, the operation against that volume will fail unless some additional direction is given. In other words, clusters will not automatically assume or take ownership of a logical volume without being directed.

The volume ownership affinity protects the volume from being accessed or modified by different clusters simultaneously. If more than one cluster has ownership of a volume, it can result in the volume's data or attributes being changed differently on each cluster, resulting in a data integrity issue with the volume. If a TS7700 Virtualization Engine Cluster has failed or is known to be unavailable (for example, it is being serviced), its ownership of logical volumes is transferred to another TS7700 Virtualization Engine Cluster, as follows:

- **Read-only Ownership Takeover**

When Read-only Ownership Takeover (ROT) is enabled for a failed cluster, ownership of a volume is taken from the failed TS7700 Virtualization Engine Cluster. Only read access to the volume is allowed through the other TS7700 Virtualization Engine Clusters in the Grid. Scratch mounts that target volumes that are owned by the failed cluster are failed. Scratch mounts that target pre-owned volumes will succeed. After ownership for a volume has been taken in this mode, any operation attempting to modify data on that volume or change its attributes fails. The mode for the failed cluster remains in place until a different mode is selected or the failed cluster has been restored.

- **Write Ownership Takeover**

When Write Ownership Takeover (WOT) is enabled for a failed cluster, ownership of a volume is taken from the failed TS7700 Virtualization Engine Cluster. Full access is allowed through the requesting TS7700 Virtualization Engine Cluster in the Grid, and all other available TS7700 Virtualization Engine Clusters in the Grid. The mode for the failed cluster remains in place until a different mode is selected or the failed cluster has been restored.

- **Service preparation or service mode**

When a TS7700 Virtualization Engine Cluster is placed in service preparation mode or is in service mode, ownership of its volumes is allowed to be taken by the other TS7700 Virtualization Engine Clusters. Full access is allowed. The mode for the cluster in service remains in place until it has been taken out of service mode.

You can set the level of ownership takeover, Read-only or Write Ownership, through the TS7700 Virtualization Engine management interface. Note that you cannot set a cluster in service preparation after it has already failed.

Autonomic Ownership Takeover Manager

Autonomic Ownership Takeover Manager (AOTM) is an optional function by which, following a TS7700 Virtualization Engine Cluster failure, one of the methods for ownership takeover is automatically enabled without operator intervention. Enabling AOTM improves data availability levels within the Composite Library.

AOTM utilizes the TS3000 System Console (TSSC) associated with each TS7700 Virtualization Engine in a Grid, to provide an alternate path to check the status of a peer TS7700 Virtualization Engine. This is why each TS7700 Virtualization Engine in a Grid must be connected to a TSSC. Through a network that is separate from the Grid, the TSSCs communicate with each other to determine if a cluster is down or if the Grid network between the clusters is down. The TS7700 Virtualization Engine at each site can be configured to automatically enable Volume Ownership takeover through AOTM.

Without AOTM, an operator must determine if one of the TS7700 Virtualization Engine Clusters has failed and then enable one of the ownership takeover modes. This is required to access the logical volumes owned by the failed cluster. It is important that write ownership takeover be enabled only when a cluster has failed and not when there is only a problem with communication between the TS7700 Virtualization Engine Clusters. If it is enabled and the cluster in question continues to operate, data might be modified independently on other clusters, resulting in corruption of data. Although there is no data corruption issue with the read ownership takeover mode, it is possible that the remaining clusters will not have the latest version of the logical volume and present down-level data.

Even if AOTM is not enabled, it is a best practice to configure it. This provides protection from a manual takeover mode being selected when the cluster is actually functional.

With AOTM, one of the takeover modes is enabled if normal communication between the clusters is disrupted and the cluster performing the takeover can verify that the other cluster has failed or is otherwise not operational. If a mount request on a TS7700 Virtualization Engine requires the transfer of ownership from a cluster that it detects as failed, it will ask the TS3000 System Console to which it is attached to confirm that the other cluster has failed. If the TSSC validates that the other cluster has indeed failed, it will reply back and the requesting TS7700 Virtualization Engine will enter the AOTM configured ownership takeover mode. If it cannot validate the failure or if the system consoles cannot communicate with each other, AOTM will not be invoked. In this scenario, ownership takeover mode can only be enabled by an operator through the management interface.

To take advantage of AOTM, you must provide an IP communication path between the TS3000 System Consoles at the cluster sites. For AOTM to function properly, it must not share the same paths as the Grid interconnection between the TS7700 Virtualization Engines.

Operational example of AOTM

If a site failure is confirmed, then the logical volume ownership takeover can be automatically enabled by the system without requiring an intervention by an operator or by service personnel. Autonomic Ownership Takeover will be enabled with either R/O or R/W capability depending on the customer-configured policy (Figure 2-33).

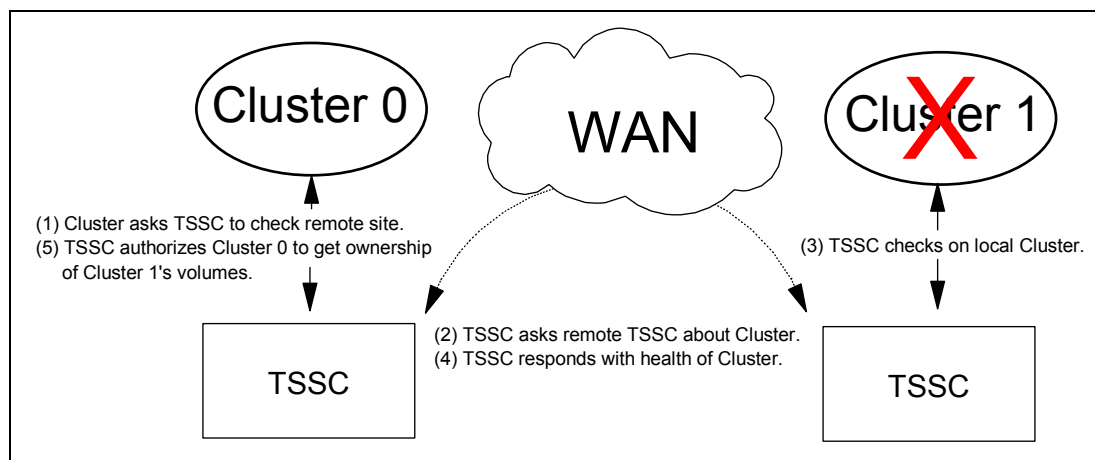


Figure 2-33 Successful Autonomic Ownership Takeover

If the TSSC is unable to ascertain the health of the remote site, then Autonomic Ownership Takeover is not performed (Figure 2-34).

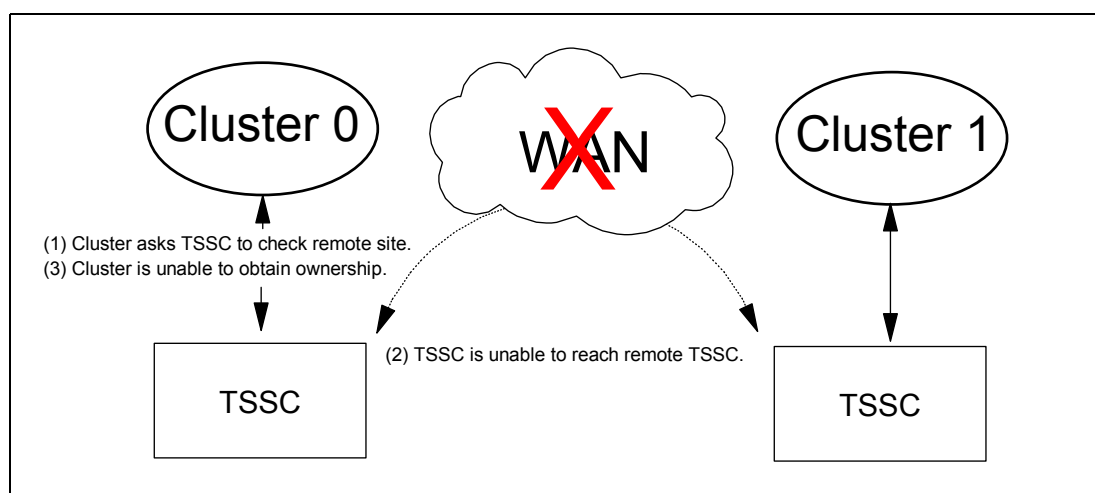


Figure 2-34 Unsuccessful Autonomic Ownership Takeover

Through options provided through SMIT menus, the IBM SSR can enable or disable this function and select which ownership takeover mode is to be entered when another TS7700 Virtualization Engine is determined not to be operational.

Note: In Figure 2-33 and Figure 2-34, the two local clusters can be attached to the same TSSC if the two clusters are in close proximity of each other.

With a three-cluster grid, an additional vote is available for peer outage detection (Figure 2-35 on page 81). For example, if Cluster 0 detects that Cluster 2 is unavailable, but Cluster 1 disagrees, then a network issue must be present. In this case, the ownership takeover will fail.

If the TSSCs of the available clusters agree that the third cluster is unavailable, then ownership takeover can occur.

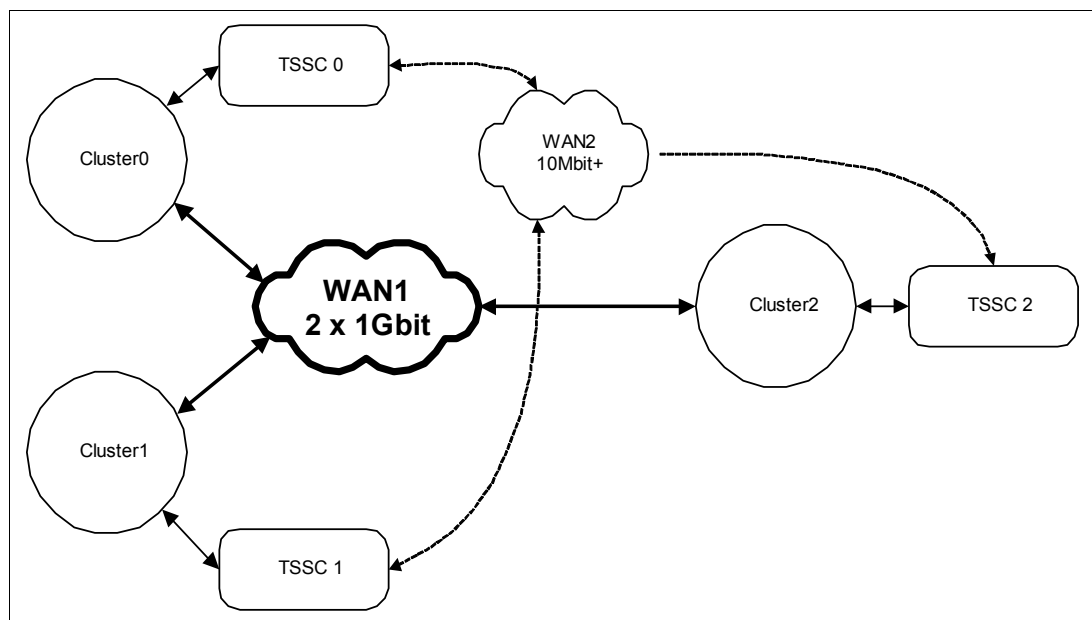


Figure 2-35 Autonomic Ownership Takeover for a three-cluster grid

Based on customer requirements, an IBM SSR enables or disables AOTM and also sets the default ownership takeover mode that is to be invoked.

Note: In Figure 2-35 on page 81, the two local clusters can be attached to the same TSSC if the two clusters are in close proximity to each other.

2.6.2 Copy policy management

When a TS7700 Virtualization Engine is part of a multi-cluster grid configuration, there are several customer policies and settings that can be used to influence the location of data copies, when the copies are to be made, and the access to data performance. This section describes the customer controls available. The implementation of multiple copies and their management differs from that of the prior generation PTP VTS. We note the differences throughout this section.

To put the following concepts in context, we first compare specific aspects of the architecture between the prior generation PTP VTS and the TS7700 Virtualization Engine. A configuration of the prior generation PTP VTS included two sets of storage server hardware, Virtual Tape Controllers (VTCs), and Virtual Tape Servers (VTSs). The VTCs served two purposes:

- ▶ To provide an access point from which a host wrote or read data that was resident in the Virtual Tape Servers
- ▶ To replicate data between the Virtual Tape Servers

With the TS7700 Virtualization Engine architecture, the hardware for the VTCs has been eliminated and the function of the VTCs has been embedded into the Virtualization Engine Controller. The access point function is now part of the vNodes and the replication function is part of the hNodes. Associated with each cluster of vNodes and hNodes is the Tape Volume Cache (TVC) and, like the prior generation VTS, this is where the data resides when being accessed.

A significant difference with the TS7700 Virtualization Engine architecture is that any vNode in the Grid can directly access the data within the TVC of any cluster. A copy of the data does not have to be in the TVC of the same cluster as the vNode used by the host for data access. Because of this enhancement, the implementation of volume copy policy management, and where data is written to and read from, is somewhat different from prior generation VTSs.

With the prior generation of PTP VTS, a VTC selected one of the two VTSs to act as the I/O VTS for a mounted volume based upon explicit direction in the Management Class, internal algorithms, and IBM SSR settings. With the TS7700 Virtualization Engine, the determination of which cluster's TVC will act as I/O TVC is based upon the copy requirements of the Management Class, internal algorithms, and customer settings. The differences are somewhat subtle and are described in more detail below.

Copy management

With the TS7700 Virtualization Engine's architectural capability of having more than two clusters peered together, copy policy management requirements are different than those of the prior generation PTP VTS. In a TS7700 Virtualization Engine Grid, you might want to have multiple copies of a virtual volume on different clusters. You might also want to specify when the copies are performed relative to the job that has written to a virtual volume and have that be unique for each cluster.

As with the prior generation PTP VTS, copy management is controlled through the Management Class storage construct. Using the management interface, you can create Management Classes and define where copies reside and when they will be synchronized to the host job that created them. Depending on your business needs for more than one copy of a logical volume, multiple Management Classes, each with a different set of definitions, can be created. The following key questions help to determine copy management in the TS7700 Virtualization Engine:

- ▶ Where do I want my copies to reside?
- ▶ When do I want my copies to become consistent with the originating data?
- ▶ Do I want logical volume copy mode retained across all Grid mount points?

Where copies will reside

With the prior generation PTP VTS, you had the ability, through a Management Class, to specify the VTS that is to be the I/O VTS and to indicate whether a copy resides on the other VTS. With the TS7700 Virtualization Engine, this concept is expanded to include multiple sites. What is now specified for a Management Class is where copies are to reside after all replication has been completed.

When a TS7700 Virtualization Engine is included in a Grid configuration, the Management Class definition window lists each cluster by its Distributed Library name and allows a copy policy for each. For example, assume three clusters are in a Grid:

- ▶ LIBRARY1
- ▶ LIBRARY2
- ▶ LIBRARY3

A portion of the Management Class definition window will include the cluster name and allow a copy consistency point to be specified for each cluster. If a copy is to reside on a cluster's TVC, then you indicate a copy consistency point. If you do not want a cluster to have a copy of the data, then you specify the No Copy option.

Note: The architecture of the TS7700 Virtualization Engine allows for multiple clusters in a Grid. The current version, Release 1.7 supports up to four clusters in a multi-cluster grid configuration.

When copies become consistent with the originating data

As part of the volume mount process within a Grid, the TVC of one of the clusters is selected as the I/O TVC (for details on the selection process and the factors considered as part of the process, see “I/O TVC selection” on page 87).

All host read and write operations are routed to the I/O TVC, which can be in a different cluster than the mount vNode. The mount vNode is in the cluster providing the virtual tape drive address against which the I/O operations from the host are being issued. To the application writing data to a volume, the data on the I/O TVC is always consistent with the application. Based on the policies defined for the Management Class assigned to the volume, the data from the I/O TVC is replicated to other TVCs associated with the other clusters in the Grid.

Two copy consistency points are provided:

- ▶ At volume *Rewind/Unload* time
- ▶ *Deferred* after *Rewind/Unload* time

Another option is *not* to have a copy of the data reside on a cluster.

If a copy consistency point of *Rewind/Unload* is defined for a cluster in the Management Class assigned to the volume, a consistent copy of the data must reside in that cluster’s TVC before command completion is indicated for the *Rewind/Unload* command.

If multiple clusters have a copy consistency point of *Rewind/Unload*, all of their associated TVCs must have a copy of the data before command completion is indicated for the *Rewind/Unload* command. Options are available to override this requirement for performance tuning purposes. See “Override settings” on page 89 for a description of the option that indicates command completion of the *Rewind/Unload* command when multiple clusters specify the *Rewind/Unload* copy consistency point.

If a copy consistency point of *Deferred* is defined, the copy to that cluster’s TVC can occur any time after the *Rewind/Unload* command has been processed for the I/O TVC. A mixture of copy consistency points can be defined for a Management Class.

Retain Copy Mode across Grid mount points

As the number of clusters participating in the Grid increases, the complexity of managing the total copy count becomes more complex. Depending on host connectivity and temporary cluster access inhibitors, such as upgrade events, mount scenarios can be invoked that result in the creation of additional logical volume copies than originally intended. The TS7700 Virtualization Engine R1.7 provides an enhanced option for managing Copy Consistency Policies (CCP) in multi-cluster grid configurations. The Retain Copy Mode setting will maintain the CCP of a non-Fast Ready volume that was in effect at the time of its initial creation. This prevents the mounting cluster from applying its CCP during non-Fast Ready mounts for reads or write appends.

Using a four-cluster TS7700 Virtualization Engine configuration, we can consider the application of the Retain Copy Mode setting. We use the Grid configuration depicted in Figure 2-36.

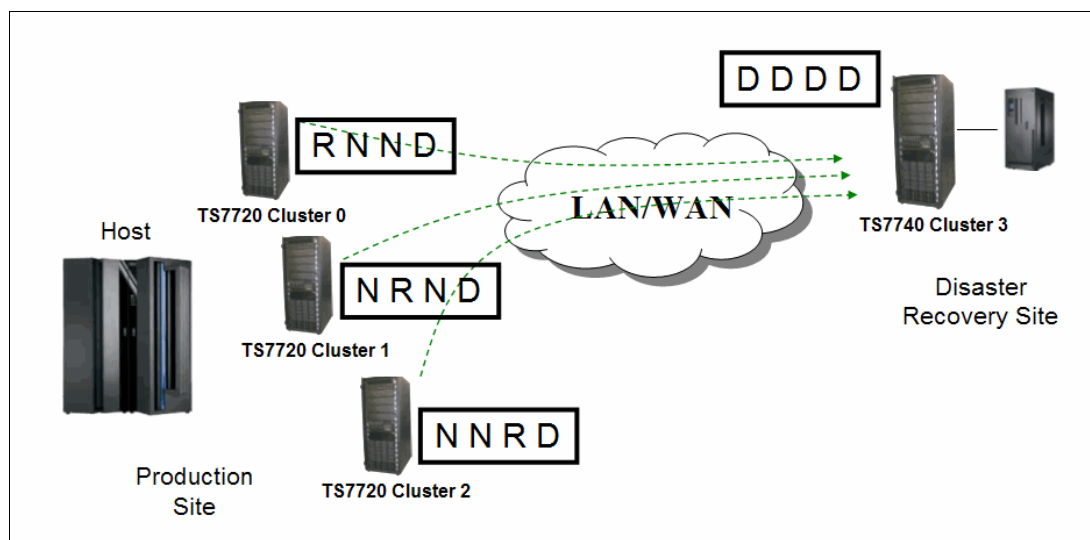


Figure 2-36 Retain Copy Mode across grid mount points

The Grid for this example spans two sites. The production site has three TS7720 Virtualization Engines that are attached to the same host. The fourth cluster in the Grid is a TS7740 Virtualization Engine and it is located at a disaster recovery site with no direct access to the production host. The goal is to maintain only two copies of a logical volume in the Grid at any given time. The CCP array is set for each TS7720 Virtualization Engine to create an immediate copy at the mounting cluster and to send a deferred copy to the TS7740 Virtualization Engine at the disaster recovery site. No copies are made between the TS7720 Virtualization Engine clusters. The CCP policies of the clusters would be as follows, where R is Run (immediate mode), N is No Copy, and D is Deferred:

- ▶ Production site TS7720 Cluster 0 - R N N D
- ▶ Production site TS7720 Cluster 1 - N R N D
- ▶ Production site TS7720 Cluster 2 - N N R D
- ▶ Disaster recovery site TS7740 Cluster 3 - D D D D

Under optimum operational conditions in the Grid, and using JES 2 Device Allocation Assist (DAA), the intended goal of producing two copies of a logical volume is met.

During allocation processing of non-Fast Ready logical volumes, DAA is used to determine which cluster is best for the specific mount request. Usually this is the cluster with the logical volume resident in cache. The host then attempts to allocate a device from the best cluster.

We next examine a change in the availability of Cluster 0 to the Grid and its impact on total logical volume copy count. In this scenario, Service Prep processing is activated on Cluster 0 so that a hardware upgrade can be performed. All subsequent mount allocations will be directed to Cluster 1 or Cluster 2.

Retain Copy Mode, not enabled or enabled, works as follows:

- Without Retain Copy Mode enabled

The host initiates mount processing for a non-fast ready logical volume and allocates a device on Cluster 1 as the mount point. The Grid selects Cluster 3, the disaster recovery TS7740 Virtualization Engine, as the I/O TVC. A cross-cluster mount is performed and a read operation followed by rewind-unload is executed on the volume. The cluster CCP is applied against the volume and an immediate copy is created on Cluster 1. At this point, copies of the logical volume are in Clusters 0, 1, and 2 of the Grid, which exceeds the total copy count target.

- With Retain Copy Mode enabled

The host initiates mount processing for a non-Fast Ready logical volume and allocates a device on Cluster 1 as the mount point. The Grid selects Cluster 3, the disaster recovery TS7740 Virtualization Engine, as the I/O TVC. A cross-cluster mount is performed and a read operation followed by rewind-unload is executed on the logical volume. The CCP in effect at the time the volume was created is enforced and no copy is created on Cluster 1. The Grid total copy count for this logical volume remains at two.

During allocation processing of non-Fast Ready logical volumes, DAA is used to determine which cluster is best for the specific mount request. Usually this is the cluster with the logical volume resident in cache. The host will then attempt to allocate a device from the best cluster.

JES3 does not support Device Allocation Assist, hence, using the configuration example, the host will allocate to clusters that do not have a copy in cache. Without Retain Copy Mode, three or four copies of a volume will exist in the grid after the dismount instead of the desired two.

Consistency point example

Figure 2-37 on page 86 shows the TS7700 Virtualization Engine management interface window where you can set the copy consistency points for a Management Class. In our example, we set the consistency points for Management Class COPY03 on Cluster 0 of the Grid. Every time a logical volume is mounted on Cluster 0, which has this Management Class assigned, we want a consistent copy of the data to reside in the TVC of Cluster 0 before command completion of Rewind/Unload is indicated to the host and a deferred copy on Clusters 1 and 2. Therefore, we specify RUN for the local Cluster 0 and specify Deferred for Clusters 1 and 2. Keep in mind that the consistency point settings for Cluster 1 and Cluster 2 can be different for the same Management Class.

"Juno[1]" (#BA39B): Add Management Class

Name:	*COPY03
Secondary Pool:	0
Description:	Immediate-Deferred-Deferred-No Copy
Retain Copy Mode:	<input checked="" type="checkbox"/>

Clusters	Copy Mode
"Zoro[0]" (#BA39A)	Rewind Unload (RUN)
"Juno[1]" (#BA39B)	Deferred
"Diana[2]" (#BA39C)	Deferred
"Herc[3]" (#BA39D)	No Copy

OK Cancel

Figure 2-37 Setting cluster copy data consistency points

You can specify the following consistency points:

No Copy

If a data consistency point of No Copy is specified, the data created on one TS7700 Virtualization Engine is not copied to the other TS7700 Virtualization Engine. If the TS7700 Virtualization Engine Cluster the data was written to fails, the data for that logical volume is inaccessible until that TS7700 Virtualization Engine is fully operational again.

Rewind Unload (RUN)

If a data copy consistency point of RUN is specified, the data created on one TS7700 Virtualization Engine is copied to the other TS7700 Virtualization Engine as part of successful Rewind/Unload command processing. This means that for completed jobs, a copy of the volume will exist on both TS7700 Virtualization Engines. Access to data written by completed jobs (successful Rewind/Unload) prior to the failure is maintained through the other TS7700 Virtualization Engine Cluster. Access to data of incomplete jobs that were in process at the time of the failure is not provided.

Deferred

If a data consistency point of Deferred is specified, the data created on one TS7700 Virtualization Engine is copied to the other TS7700 Virtualization Engine after successful Rewind/Unload command processing. Access to the data through the other TS7700 Virtualization Engine Cluster is available before the copy completes, as long as the source TS7700 Virtualization Engine is still available. Because there will be some delay in performing the copy, access might not be available when a failure occurs.

The default Management Class action is to have a Rewind/Unload consistent copy reside only on the local cluster. The other clusters have a No Copy option specified. The default actions are applied whenever a new construct is defined through the management interface or specified in a **mount** command where it was not previously defined.

Management Class locality to a cluster

Management Classes for the TS7700 Virtualization Engine are created at the management interface associated with a cluster. The same Management Class can be defined differently at each cluster and there are valid reasons for doing so. For example, one of the functions controlled through Management Class is to have a logical volume copied to two different physical volume pools.

You might want to have two separate physical copies of your logical volumes on one of the clusters and not on the others. Through the management interface associated with the cluster where you want the second copy, you specify a secondary pool when defining the Management Class. For the Management Class definition on the other clusters, you do not specify a secondary pool. For example, you might want to use the Copy Export function to extract a copy of data from the cluster to be taken to a disaster recovery site.

It is important to note that during mount processing, the copy consistency point information that is used for a volume is taken from the Management Class definition for the cluster with which the mount vNode is associated. A best practice is to define the copy consistency point definitions of a Management Class to be the same on each cluster so as to avoid confusion about where copies will reside. You can devise a scenario in which you define separate copy consistency points for the same Management Class on each of the clusters. In this scenario, the location of copies and when the copies are consistent with the host that created the data will differ, depending upon which cluster a mount is processed.

I/O TVC selection

The TVC associated with one of the clusters in the Grid is selected as the I/O TVC for a given tape mount request. All I/O operations associated with the virtual tape drive the mount was issued on are routed from its vNode to the selected I/O TVC. The vNode is referred to as the mount vNode. The I/O TVC is selected as part of the processing for the mount request. Similar to the prior generation's PTP VTS's I/O VTS selection process, the TS7700 Virtualization Engine takes into consideration such factors as data validity, cluster availability, type of mount request (Fast Ready (scratch) or non-Fast Ready), performance, local preferences, family affinity, and override settings. The many factors are evaluated and an I/O TVC is selected. The TS7700 Virtualization Engine also takes into consideration the copy consistency points defined in the Management Class associated with the volume being mounted. A cluster with a Rewind/Unload copy consistency point requirement is weighted heavier than one with a Deferred copy consistency point.

For example, Management Class MCPROD01 has been defined with the following copy data consistency points:

- ▶ LIBRARY1 Rewind/Unload
- ▶ LIBRARY2 Deferred

Assume a scratch mount request is received specifying a Management Class of MCPROD01 on a virtual drive address associated with cluster LIBRARY2 and no customer overrides have been established for that cluster. The selection process for picking the I/O TVC will select the TVC associated with cluster LIBRARY1 as long as it is available because it has the highest copy consistency point requirement. The reason for weighting the copy consistency point in this way is so that when the Rewind/Unload is received, the data will already be consistent and no delay will be introduced in performing a copy to another cluster's TVC prior to indicating that the Rewind/Unload command has completed.

If a Management Class had been defined that specified more than one Rewind/Unload cluster copy consistency point, those clusters' TVCs are given equal weight and other factors,

such as performance between the mount vNode and the TVCs, to steer the decision. For example, a Management Class MCPROD02 can be defined as follows:

- ▶ LIBRARY1 Rewind/Unload
- ▶ LIBRARY2 Rewind/Unload

If a scratch mount specifying MCPROD02 is received for a vNode in cluster LIBRARY2, the TVC associated with LIBRARY2 is selected because it meets the cluster consistency point requirement and is expected to have the best performance because the TVC is local. The TS7700 Virtualization Engine estimates the “best” performing TVC based on many factors, including whether a TVC is full and throttling. Because this is a scratch mount, other factors like consistency and cache residency are not considered. It is possible that the TVC in LIBRARY2 is full and throttling. In that case, the TVC in LIBRARY1 would likely be selected to avoid job throttling.

With the prior generation of PTP VTS, you could force a specific VTS to be the I/O VTS by selecting that option in the Management Class definition. With the TS7700 Virtualization Engine, you can also direct the TVC of a specific cluster to be used as the I/O TVC. The method of defining this capability is different than for the prior generation of PTP VTS and allows for additional flexibility.

Suppose that you want the data to reside only on the cluster that is associated with LIBRARY2. You define a Management Class that has a copy consistency point only for LIBRARY2. LIBRARY1 is specified as No Copy. Regardless of which vNode a mount using that Management Class was received on (assuming that you have defined the Management Class the same on all clusters), the TVC that is associated with cluster LIBRARY2 is selected. If the LIBRARY2 cluster is not available, the mount fails. It might be that failing the mount is not desired. If that is the case, you could define the copy point for LIBRARY2 as Rewind/Unload and a third cluster's (LIBRARY3) copy consistency point as Deferred. With that Management Class definition, if LIBRARY2 cluster is not available, the TVC associated with LIBRARY3 will be chosen and the mount does not failed. A deferred copy will still be made to LIBRARY2 when it becomes available.

Note: The copy consistency point is considered for both scratch and specific mounts. This is different from the PTP VTS, which only looked at the copy policy for scratch mounts.

Cross cluster mounts

The TS7700 Virtualization Engine offers maximum flexibility in implementation of individual cluster hardware and Grid configurations. This flexibility creates scenarios in multi-cluster configurations where mount request algorithms in the Grid component result in a remote cache selected as the I/O TVC. A cross cluster mount is created when the I/O TVC selected is not in the cluster which owns the allocated virtual device. The logical volume is accessed through the Grid network using TCP/IP. Each I/O request from the host will result in parts of the logical volume moving across the network. Logical volume data movement through the Grid is bidirectional and depends on whether the operation is a read or a write. The amount of data transferred depends on many factors, one of which is the data compression ratio provided by the host FICON adapters. In essence, to minimize Grid bandwidth requirements only compressed data used or provided by the host is transfer across the network.

Note: Cross cluster mounts use only one of the Grid links for data transfer throughout the entire time the volume is mounted. If the Grid link that is performing the remote mount fails the host job will fail. Resubmission of the job allows the use of an alternate Grid link.

Override settings

With the prior generation of PTP VTS, there were several optional override settings that influenced how an individual VTC selected a VTS to perform the I/O operations for a mounted tape volume. The override settings were determined by the customer but set by an IBM SSR. With the TS7700 Virtualization Engine, the customer defines and sets the optional override settings that influence the selection of the I/O TVC and replication responses through the management interface.

TS7700 Virtualization Engine overrides I/O TVC selection and replication response

The settings are specific to a cluster, meaning that each cluster can have different settings if desired. The settings take effect for any mount requests received after the settings were saved. Mounts already in progress are not affected by a change in the settings. The following override settings are supported:

► Prefer Local Cache for Fast Ready Mount Requests

This override selects the local TVC to the mount vNode cluster as the I/O TVC if it is available and a copy consistency point, other than No Copy, is specified for that cluster in the Management Class that is specified with the mount. The cluster must obtain volume ownership if it does not have it.

The default behavior of the TS7700 Virtualization Engine is to select an I/O TVC based on the highest copy consistency point defined for the Management Class associated with the mount request and availability considerations. If the cluster of the mount vNode has a low cluster copy consistency point, a remote TVC from the mount VNode cluster is selected. This might have undesired performance impacts on the application.

Note: The copy consistency rules are not overridden with this override. For example, if the cluster associated with the I/O TVC has a copy consistency point of Deferred and another cluster has a copy consistency point of Rewind/Unload, completion of the Rewind/Unload command is not indicated until the copy has been completed to the other cluster. Therefore, the local cluster is not preferred if it is configured as *D* and a remote cluster is still configured as *R* (that is, the preference does not override copy consistency points).

► Prefer Local Cache for non-Fast Ready Mount Requests

This override selects the TVC local to the mount VNode cluster as the I/O TVC as long as it is available and the cluster has a valid copy of the data, even if the data is only resident on a physical tape. Having an available, valid copy of the data overrides all other selection criteria. If the local cluster does not have a valid copy of the data, then the default selection criteria applies.

The default behavior of the TS7700 Virtualization Engine is to select an I/O TVC based upon the following criteria:

- Availability of a valid copy
- Cache residency
- The copy consistency point defined for the volume (highest is best)
- Performance characteristics between clusters
- Cluster availability

Note: If the volume is modified, the copy consistency rules are not overridden with this override.

► Force Local TVC to have a copy of the data

The default behavior of the TS7700 Virtualization Engine is to only make a copy of the data based on the definitions of the Management Class associated with the volume mounted and to select an I/O TVC that was defined to have a copy. If the mount vNode is associated with a cluster for which the specified Management Class defined a copy consistency point of No Copy, a copy is not made locally and all data access is to a remote TVC.

In addition, if the mount vNode has a specified defined copy consistency point of Deferred, remote Rewind/Unload clusters will be preferred. This override has the effect of overriding the specified Management Class with a copy consistency point of Rewind/Unload for the local cluster independent of its currently configured copy consistency point. In addition, it will require that the local cluster always be chosen as the I/O TVC. If the mount type is non-Fast Ready and a consistent copy is unavailable in the local TVC, a copy is performed to the local TVC prior to mount completion. The copy source can be any participating TS7700 Virtualization Engine in the Grid. In the case of a TS7740 Virtualization engine, the logical volume may have to be recalled from a stacked cartridge. If, for any reason, the vNode cluster is not able to act as the I/O TVC, a mount operation will fail even if remote TVC choices are still available when this override is enabled.

The override does not change the definition of the Management Class. It serves only to influence the selection of the I/O TVC or force a local copy.

Note: If the volume is modified, the copy consistency rules are not overridden with this override.

► Copy Count Override

This override limits the number of RUN consistency points in a two-, three-, or four-cluster grid that must be consistent prior to the surfacing of Initial Status or device end to a Rewind Unload command. Only copy consistency points of RUN are counted. For example, in a three-cluster grid, if the Management Class specifies copy consistency points of RUN, RUN, RUN, and the override is set to two, then initial status or device end is presented after at least two clusters configured with a RUN consistency point are consistent. This includes the original I/O TVC if that site is also configured with a Rewind/Unload consistency point. The third RUN consistency point is downgraded to a Deferred copy after at least two of the three RUN consistency points are proven to be consistent. The third site that has its copy consistency point downgraded to Deferred is called the *floating deferred site*. A floating deferred site is one that has not completed its copy when the Copy Count value has been reached.

This override has limited use in a two-cluster grid. The only limitation available is to limit the number of RUN consistency points to one. The override will only come into play for a Management Class with copy consistency points of RUN, RUN.

Overrides for Geographically Dispersed Parallel Sysplex (GDPS)

The default behavior of the TS7700 Virtualization Engine is to follow the Management Class definitions and configuration characteristics to provide the best overall job performance. In a Geographically Dispersed Parallel Sysplex™ (GDPS®), all I/O must be local to the mount vNode. There can be other customer requirements, such as disaster recovery testing, where all I/O must only go to the local TVC to ensure that the correct copy policies have been implemented and data is available where required.

In a GDPS environment, you have to set the Force Local TVC override to ensure that the local TVC is selected for all I/O. This setting includes:

- ▶ Prefer Local for Fast Ready Mounts
- ▶ Prefer Local for non-Fast Ready Mounts

Consistency configuration example 1

In the scenario shown in Figure 2-38, we have two sites with a System z host and a TS7700 Virtualization Engine installed in each host. The two clusters are part of a multi-cluster grid configuration. The hosts attach to both TS7700 Virtualization Engines. In our example, we assume that the FICON connections between the two sites only provides limited bandwidth. For this reason, we vary off the virtual devices on the remote cluster for each host during normal operation. Only in the case of disaster recovery do we vary these devices online. This means that during normal operation, Host A always accesses virtual devices on Cluster 0 and Host B accesses virtual devices on Cluster 1.

We want to have two consistent copies of the data to provide continued data access when a volume with a certain Management Class assigned to it is closed by the application. Because we want to have the same behavior at both sites, for each cluster we specify copy data consistency points of RUN, RUN (indicated by RR in Figure 2-38).

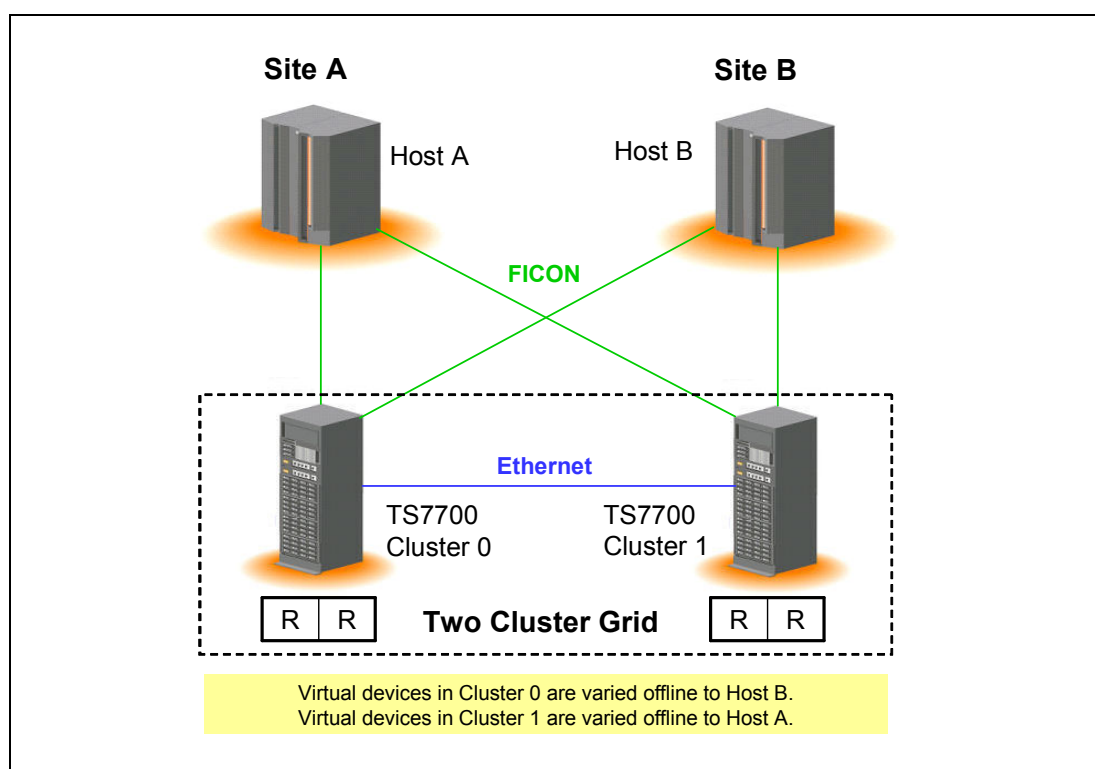


Figure 2-38 Example 1: Copy consistency settings RUN, RUN

Next, we look at what happens when Host A mounts a logical volume on a virtual device in Cluster 0. The mount vNode of Cluster 0 has access to both its local TVC and to the TVC of remote Cluster 1, so the local TVC will not necessarily be selected as I/O TVC. The cluster copy consistency point settings are the same for both clusters, so there is no preference as far as copy consistency is concerned. In this case, locality will probably tip the scales in favor of the local TVC, as the Fibre Channel attached drives of the local TVC give you a better performance than the Ethernet connected remote TVC under normal circumstances. Be aware that selection of the remote TVC as I/O TVC is still possible because numerous other

factors, such as cache residency or validity of data, are also taken into account and can overrule the locality factor. The same considerations apply to Host B and Cluster 1.

If you want to ensure that the local TVC is always selected as the I/O TVC, you have to fine-tune the system by using the override settings described in “TS7700 Virtualization Engine overrides I/O TVC selection and replication response” on page 89.

Although there is no one-to-one correspondence of TS7700 Virtualization Engine copy consistency points to the Peer-to-Peer Copy Modes of the previous generation subsystem, this TS7700 Virtualization Engine’s settings (RUN, RUN on both clusters) roughly compare to the Immediate Copy/Balanced Mode of the Peer-to-Peer VTS.

Consistency configuration example 2

In this scenario (Figure 2-39), we have a System z host and a TS7700 Virtualization Engine at Site A, and a remote TS7700 Virtualization Engine at Site B. The two Virtualization Engines form a multi-cluster grid. The host has access to all virtual devices of Cluster 0 and Cluster 1.

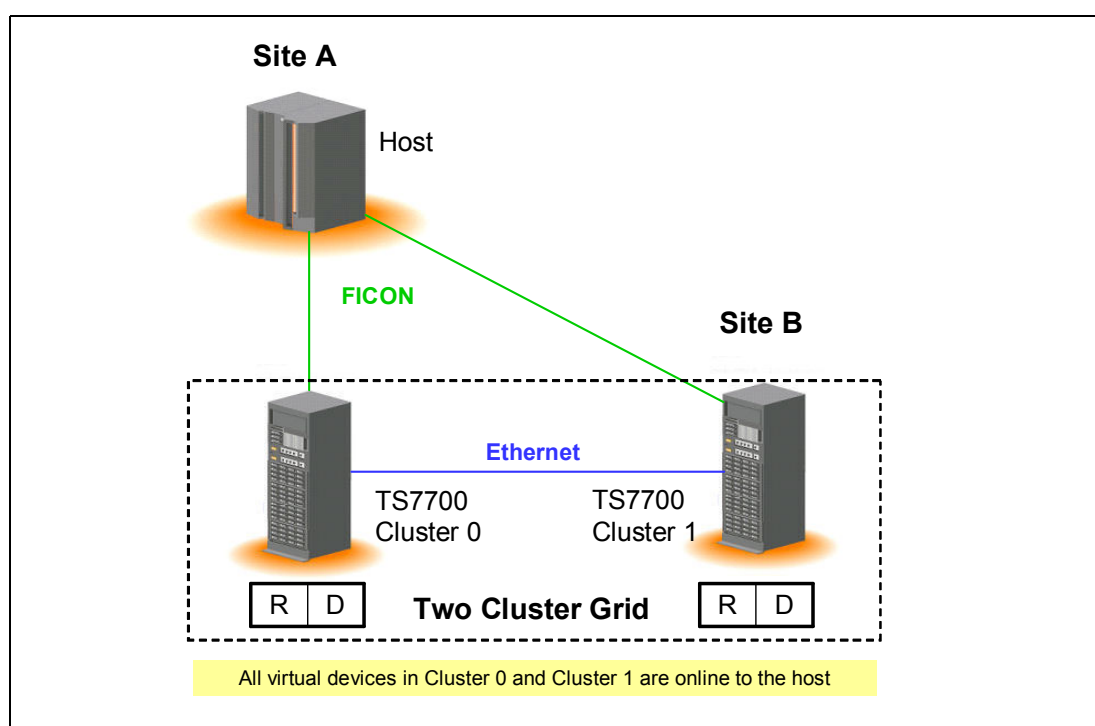


Figure 2-39 Example 2: Copy consistency point settings RUN/Deferred

We assume that a copy of the data is required on both clusters. Contrary to the previous example, we need a valid copy of the data at Rewind/Unload only on Cluster 0 in the local site. The data can be copied to Cluster 1 at a later time. Therefore, we define on both clusters the copy consistency points as RUN for Cluster 0 and Deferred for Cluster 1. When we look at a Fast Ready Mount, for example, we now have to differentiate between two cases:

- The host allocates a virtual drive on Cluster 0.

During the mount process, the vNode of Cluster 0 selects the I/O TVC. One of the factors it takes into consideration for this decision is the copy consistency points defined in the Management Class associated with the volume being mounted. The RUN setting for Cluster 0 weighs more than the Deferred setting for Cluster 1. The vNode of Cluster 0 will

therefore select its local TVC as I/O TVC. Locality also affects the TVC of Cluster 0, but the decisive factors here are the copy consistency points.

- The host allocates a virtual drive on Cluster 1.

During the mount process, the mount vNode in Cluster 1 selects the I/O TVC. It evaluates the copy consistency points for the decision, and again the TVC of Cluster 0 is selected as I/O TVC. The difference is that now the remote TVC is selected (remote from the perspective of the selecting vNode). Locality affects the local TVC as I/O TVC, but copy consistency point settings weigh heavier than locality. So data is written to the TVC of Cluster 0 first, using the Ethernet connection between the two clusters. The copy to the TVC of Cluster 1 happens at a later time.

You can modify the behavior of the clusters by using the override settings explained in “Override settings” on page 89.

Although there is no one-to-one correspondence of TS7700 Virtualization Engine copy consistency points and the Peer-to-Peer Copy Modes of the previous generation subsystem, these TS7700 Virtualization Engine settings (RUN/Deferred on both clusters) roughly compare to Deferred Copy/Preferred Mode of the Peer-to-Peer VTS.

Consistency configuration example 3

As we mentioned before, the copy consistency point settings can be different for the two clusters in a multi-cluster grid configuration. Although it is a best practice to have the same settings on both clusters, there might be scenarios with special requirements, where different settings are appropriate.

This scenario (Figure 2-40) is much the same as the scenario in example 1 (Figure 2-38 on page 91). The scenarios differ only in the copy consistency point settings. Here we have different copy consistency point settings defined on Cluster 0 and Cluster 1 of the Grid configuration.

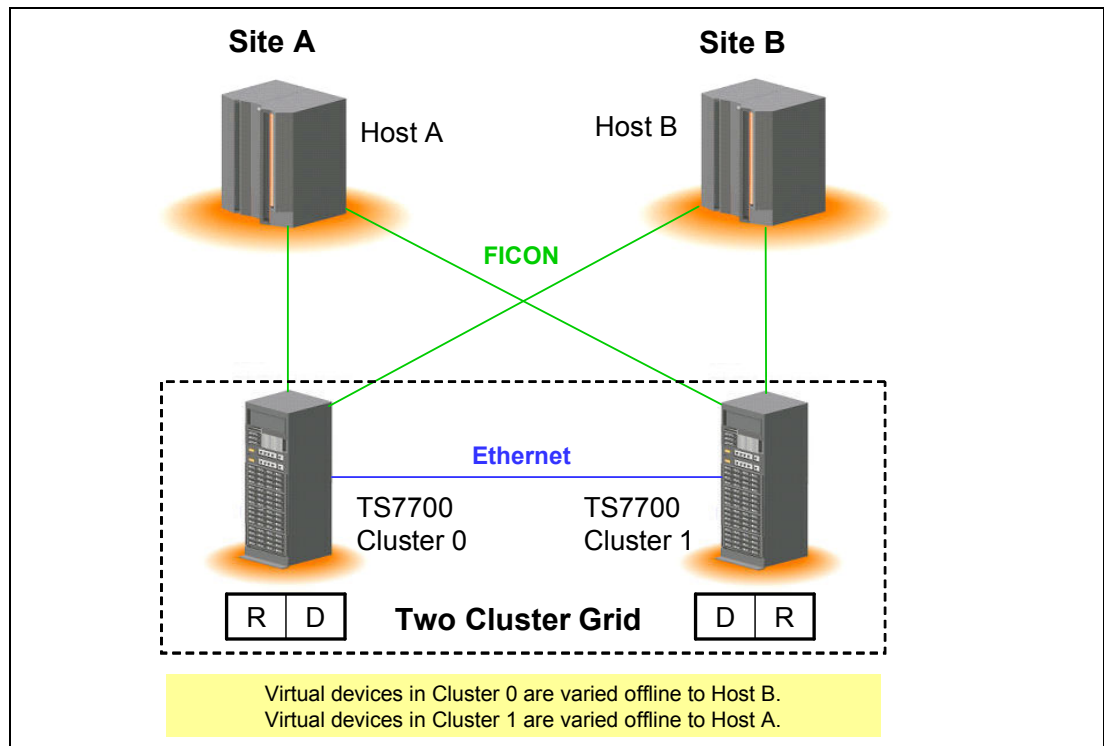


Figure 2-40 Example 3: Different consistency point settings in a Grid configuration

The hosts in this scenario have access to the virtual devices of their local cluster only, as the remote virtual devices are varied offline to the hosts. When a host mounts a logical volume on one of these virtual devices, we like to see that the respective local TVC is selected as the I/O TVC, that a copy exists in the local TVC before Rewind/Unload is indicated to the host, and that a deferred copy will be made to the TVC of the remote cluster at a later time.

To achieve this goal, we set the copy consistency points on Cluster 0 to RUN and to Deferred for Cluster 1 (RD in Figure 2-40 on page 93). On Cluster 1, we reverse these settings and define Deferred for Cluster 0 and RUN for Cluster 1 (DR in Figure 2-40 on page 93). Regardless of the site where a virtual volume is mounted, the local TVC is preferred to the remote TVC for I/O TVC and a copy must exist in the local TVC prior to Rewind/Unload, because the copy consistency point of RUN is set for the local cluster. The copy consistency point of Deferred is for the remote cluster and causes a deferred copy to be made to the remote TVC.

Note: This setting of RD and DR can cause strange behavior when you read data from the opposite Cluster 1 immediately after writing the volume to Cluster 0. If your workload requires immediate access from the other cluster, you should set the copy consistency points on both clusters to DD.

Although no R values exist, they are not required. The mounting vNode simply chooses one of the Ds as an I/O TVC and internally views it as a RUN site. In addition to DD-DD, we also suggest that you set “Prefer local Fast Ready”, so that Fast Ready mounts stay local. For our example, it then work as follows:

- ▶ Cluster 0 Fast Ready mount: Because “prefer local” is set, Cluster 0 is selected as TVC and is internally treated like a RUN site.
- ▶ Cluster 1 mount for read: Because Cluster 1 has not yet completed its copy, Cluster 0 is still chosen as I/O TVC (remote mount).
- ▶ When the volume is demounted, device end will not be held up, because no R values exists. The copy will remain queued to Cluster 1 and it will eventually complete.

In the same scenario, if you wanted to always have two valid copies at Rewind/Unload for mounts on Cluster 0, and just one valid copy in the local TVC for mounts on Cluster 1, you set the copy consistency points as follows:

- ▶ Cluster 0: RUN for Cluster 0, RUN for Cluster 1
- ▶ Cluster 1: No Copy for Cluster 0, RUN for Cluster 1

Device Allocation Assist

Support for Device Allocation Assist (DAA) is available in TS7700 Virtualization Engine R1.5 and later. DAA applies only to multi-cluster Grid configurations where devices from more than one cluster are online to the same host. The DAA function is a cooperative capability between the z/OS host and the TS7700 Virtualization Engine. Utilizing DAA the host will attempt to select the ‘best’ locally attached TS7700 Virtualization Engine cluster, either a TS7720 or TS7740, to satisfy a non-fast ready mount request.

With DAA, the host is now aware of which distributed library a virtual device is associated with. Prior to allocating a virtual device for a non-fast ready mount the host queries the TS7700 Grid for the allocation data. The logical volume identifier is received by the TS7700 Grid, analyzed and an ordered list of cluster identifiers is returned to the host with the ‘best’ candidate listed first. Subsequently, host processing will attempt to allocate a device from the first cluster returned in the list. If an online device is not available within that cluster, it will move to the next cluster in the list and try again until a device is chosen. This allows the host

to direct the mount request to the cluster that would result in the fastest mount, typically the cluster that has the logical volume resident in cache.

DAA improves a Grid's performance by reducing the number of cross-cluster mounts. This feature is especially important when copied volumes are treated as Preference Group 0 (removed from cache first) and when copies are not made between locally attached clusters of a common Grid. In conjunction with DAA, using the copy policy overrides to prefer local TVC for fast ready mounts provides the best overall performance. Configurations which include the TS7720's deep cache will dramatically increase their cache hit ratio.

Without DAA, configuring the cache management of replicated data as PG1 (prefer to be kept in cache with a LRU algorithm) is the best way to improve non-fast ready mount performance by minimizing cross cluster mounts. However, this performance gain comes with a reduction in the effective Grid cache size because multiple clusters are maintaining a copy of a logical volume. To regain the same level of effective Grid cache size an increase in physical cache capacity is required.

DAA includes updates to the TS7700 Virtualization Engine (R1.5 and later) and in host software (APAR OA24966 for z/OS V1R8, V1R9 and V1R10). DAA functionality is included in z/OS V1R11 and later.

Note: JES3 does not use Device Allocation Assist.

2.6.3 High availability and disaster recovery: Two-cluster grid

With a two-cluster grid, you can configure the Grid for disaster recovery, high availability, or both. In the following sections, we describe configuration considerations for two-cluster grids. The scenarios presented are typical configurations. Other configurations are possible and might be better suited for your environment.

Disaster recovery configuration

This section provides that information needed to plan for a TS7700 Virtualization Engine two-cluster grid configuration to be used specifically for disaster recovery purposes.

We assume that some natural or human-caused event has made the local site's TS7700 Virtualization Engine Cluster unavailable. The two TS7700 Virtualization Engine Clusters reside in different locations, separated by a distance dictated by your company's requirements for disaster recovery. The only connection between the local site and the disaster recovery site are the Grid interconnections. There is no host connectivity between the local hosts and the FICON channels on the disaster recovery site's TS7700 Virtualization Engine.

Figure 2-41 summarizes this configuration.

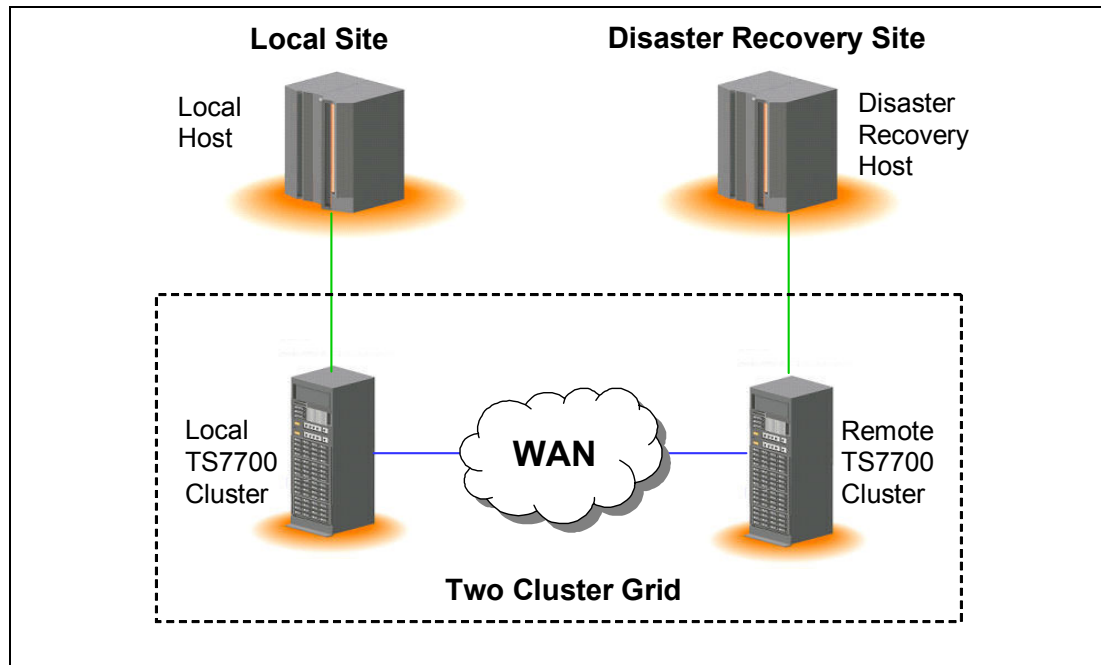


Figure 2-41 Disaster recovery configuration

As part of planning a TS7700 Virtualization Engine Grid configuration to address this solution, you need to consider the following items:

- ▶ Plan for the necessary wide area network infrastructure and bandwidth to meet the copy requirements that you need. You generally need more bandwidth if you are primarily using a copy consistency point of RUN, because any delays in copy time caused by bandwidth limitations will result in an elongation of job run times. If you have limited bandwidth available between sites, have data that is critical copied with a consistency point of RUN, with the rest of the data using the Deferred copy consistency point.
- ▶ Plan for host connectivity at your disaster recovery site with sufficient resources to perform your critical workloads.
- ▶ Design and code the DFSMS Automatic Class Selection routines to control what data gets copied and by which copy consistency point.
- ▶ Prepare procedures that your operators execute in the event the local site becomes unusable. The procedures include such tasks as bringing up the disaster recovery host, varying the virtual drives online, and placing the disaster recovery TS7700 Virtualization Engine Cluster in one of the ownership takeover modes.
- ▶ Keep in mind that if the disaster recovery site's local TS7700 Virtualization Engine Cluster becomes unavailable, there is no local host access to the data in the TS7700 Virtualization Engine Cluster.

Configuring for high availability

This section provides the information needed to plan for a TS7700 Virtualization Engine two-cluster grid configuration to be used specifically for high availability. The assumption is that continued access to data is critical and no single point of failure, repair, or upgrade can impact the availability of data.

In a high availability configuration, both TS7700 Virtualization Engine Clusters are located in close proximity to one another at the same site. These clusters are connected through a local

area network. If one of them becomes unavailable because it has failed, or is undergoing service or being updated, data can be accessed through the other TS7700 Virtualization Engine Cluster until the unavailable cluster is made available. As part of planning a TS7700 Virtualization Engine Grid configuration to address this solution, you will need to consider the following items:

- ▶ Plan for the virtual device addresses in both clusters to be configured to the local hosts. In this way, a total of 512 virtual tape devices are available for use (256 from each TS7700 Virtualization Engine Cluster).
- ▶ Set up a copy consistency point of RUN for both clusters for all data to be made highly available. With this copy consistency point, as each logical volume is closed, it is copied to the other TS7700 Virtualization Engine Cluster.
- ▶ Design and code the DFSMS Automatic Class Selection routines to set the necessary copy consistency point.
- ▶ Prepare procedures that your operators execute if one of the TS7700 Virtualization Engine Clusters becomes unavailable. The primary task is to place the remaining cluster in the mode so it can take over ownership of logical volumes that were owned by the unavailable TS7700 Virtualization Engine Cluster.

Figure 2-42 summarizes this configuration.

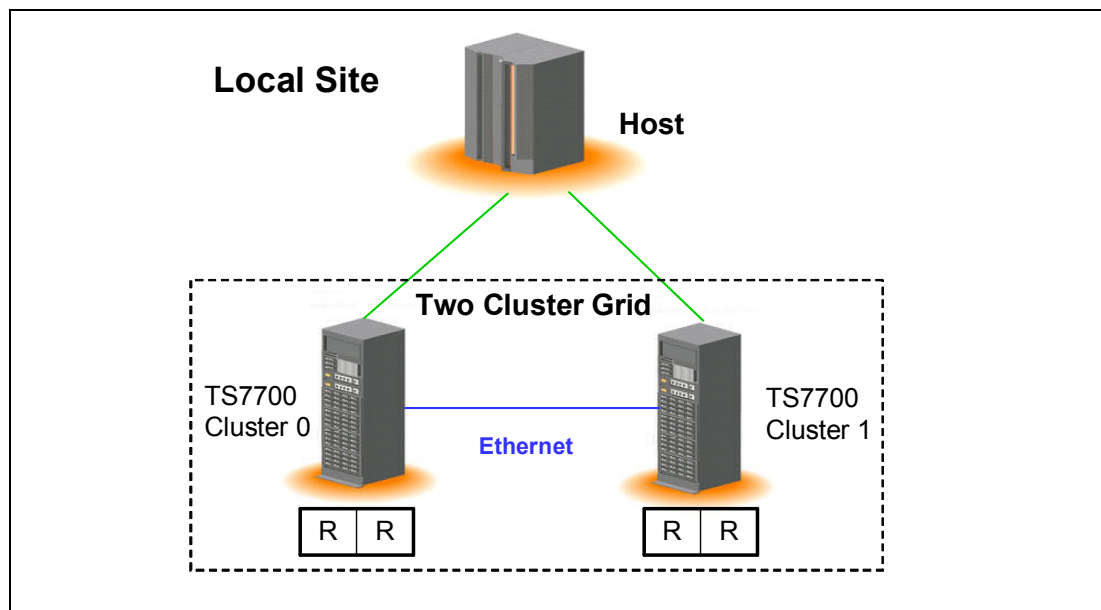


Figure 2-42 Availability configuration

Configuring for disaster recovery and high availability

A possibility is to configure a TS7700 Virtualization Engine two-cluster grid configuration to provide both disaster recovery and high availability solutions.

The assumption is that the two TS7700 Virtualization Engine Clusters will reside in separate locations, separated by a distance dictated by your company's requirements for disaster recovery. In addition to the configuration considerations for disaster recovery, you need to plan for the following items:

- ▶ Access to the FICON channels on the TS7700 Virtualization Engine Cluster located at the disaster recovery site from your local site's hosts. This can involve connections using DWDM or channel extension equipment, depending on the distance separating the two sites. If the local TS7700 Virtualization Engine Cluster becomes unavailable, you use this

remote access to continue your operations using the remote TS7700 Virtualization Engine Cluster.

- ▶ Because the virtual devices on the remote TS7700 Virtualization Engine Cluster are connected to the host through a channel extension, there can be a difference in read or write performance when compared to the virtual devices on the local TS7700 Virtualization Engine Cluster. If performance differences are a concern, consider only using the virtual device addresses in the remote TS7700 Virtualization Engine Cluster when the local TS7700 Virtualization Engine is unavailable. If that is an important consideration, in addition to the ownership takeover procedure, you need to provide for operator procedures to vary online and offline the virtual devices in the remote TS7700 Virtualization Engine.
- ▶ You might want to have separate copy consistency policies for your disaster recovery data versus your data that requires high availability.

Figure 2-43 summarizes this configuration.




Figure 2-43 Availability and disaster recovery configuration

2.6.4 High availability and disaster recovery: three-cluster grid

With a three-cluster grid, you can configure the Grid for disaster recovery and high availability or use dual production sites that share a common disaster recovery site. In the following sections, we describe configuration considerations for three-cluster grids. The scenarios presented are typical configurations. Other configurations are possible and might be better suited for your environment.

The planning considerations for a two-cluster grid also apply to a three-cluster grid.

High availability and disaster recovery

Figure 2-44 on page 99 illustrates a combined high availability and disaster recovery solution for a three-cluster grid. In this example, Cluster 0 and 1 are the high availability clusters and are local to each other (less than 50 kilometers apart). Cluster 2 is at a remote site that is away from the production site. The virtual devices in Clusters 0 and 1 are online to the host

and the virtual devices in Cluster 2 are offline to the host. The host accesses the 512 virtual devices provided by Cluster 0 and 1. Host data written to Cluster 0 is copied to Cluster 1 at run time. Host data written to Cluster 1 is written to Cluster 0 at run time. Host data written to Cluster 0 or 1 is copied to Cluster 2 on a Deferred basis.

The copy consistency points at the disaster recovery site (NNR) are set to only create a copy of host data at Cluster 2. Copies of data are not made to Clusters 0 and 1. This allows for disaster recovery testing at Cluster 2 without touching the production site clusters.

Figure 2-44 shows an optional host connection that can be established to remote Cluster 2 using DWDM or channel extenders. With this configuration, you need to define an additional 256 virtual devices at the host for a total of 768 devices.

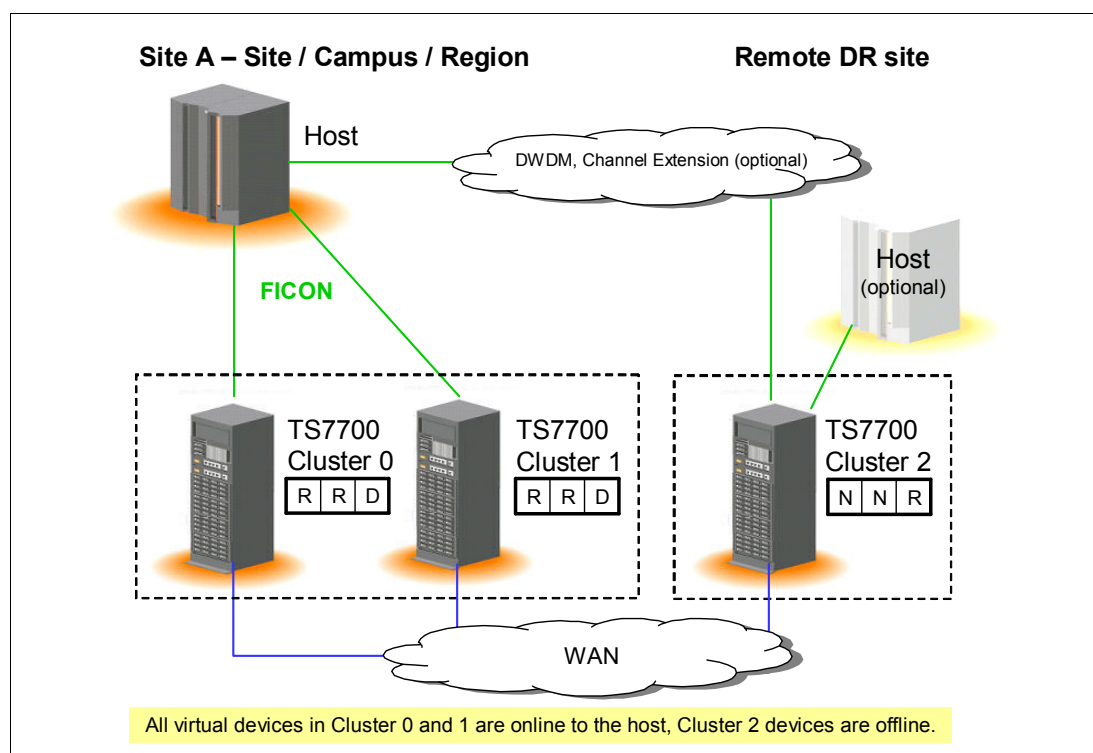


Figure 2-44 High availability and disaster recovery configuration

Dual production site and disaster recovery

Figure 2-45 on page 100 illustrates dual production sites that are sharing a disaster recovery site in a three-cluster grid. In this example, Cluster 0 and 1 are separate production systems that can be local to each other or might be distant from each other. The disaster recovery cluster, Cluster 2, is at a remote site at a distance away from the production sites. The virtual devices in Cluster 0 are online to host A and the virtual devices in Cluster 1 are online to host B. The virtual devices in Cluster 2 are offline to both hosts. Host A and host B access their own set of 256 virtual devices provided by their respective clusters. Host data written to Cluster 0 is not copied to Cluster 1. Host data written to Cluster 1 is not written to Cluster 0. Host data written to Cluster 0 or 1 is copied to Cluster 2 on a Deferred basis.

The copy consistency points at the disaster recovery site (NNR) are set to only create a copy of host data at Cluster 2. Copies of data are not made to Clusters 0 and 1. This allows for disaster recovery testing at Cluster 2 without touching the production site clusters.

Figure 2-45 shows an optional host connection that can be established to remote Cluster 2 using DWDM or channel extenders.

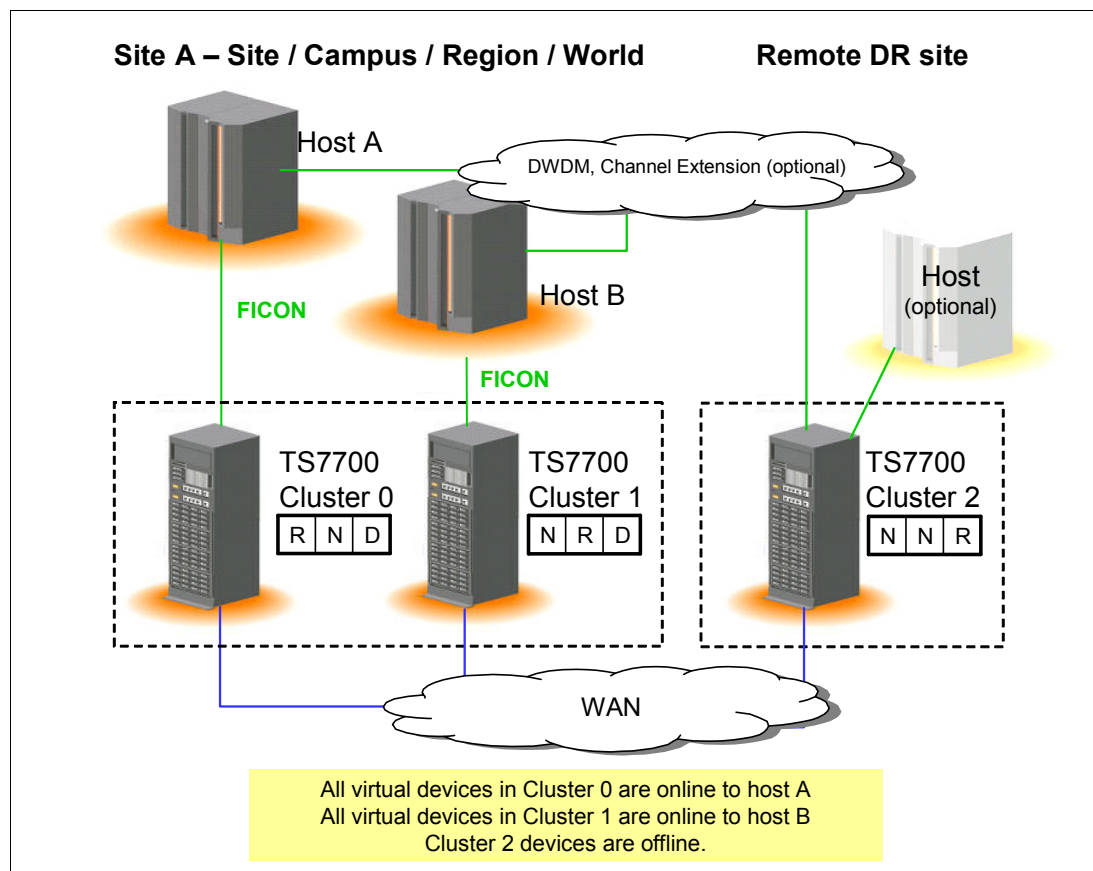


Figure 2-45 Dual production site with disaster recovery

2.6.5 Removal of a cluster from a Grid and cluster clean-up

It is possible to remove one cluster from a Grid if you have a multi-cluster Grid configuration and one of the TS7700 Virtualization Engines is no longer required to participate in the Grid.

The following features can be used if you are planning a data center consolidation or if you need to “ungrid” a TS7700 Virtualization Engine to meet a business objective:

- ▶ Removal of a cluster from a grid
- ▶ Cluster cleanup

Removal of a cluster from a Grid is a feature that delivers instructions for a one time process to remove or unjoin a cluster from a Grid configuration. This function can be used for removing one cluster from a two-, three-, or four-cluster grid configuration. Subsequent invocations can be executed to remove additional clusters as required.

Removal of a cluster from a Grid requires that all data in the cluster that is going to be removed be copied, removed, or expired before the cluster removal. After the removal, the cluster is disabled and cannot be returned to active until a cluster cleanup has been performed on this cluster.

Cluster cleanup cleans the database, logically deletes volumes from the Tape Volume Cache, and removes the configuration data for host connections from a TS7700 Virtualization Engine

Cluster. However, cluster cleanup does not physically delete (“shred”) user data from the Tape Volume Cache or from the back-end stacked tape cartridges of a TS7740 Virtualization Engine.

The cluster cleanup feature is required so that the removed cluster can be reused. This feature is a single use feature, and returns the removed cluster to a state similar to one received from manufacturing.

Note: These procedures are performed by IBM as a service offering. Contact your System Service Representative (SSR) about this offering and how to request the service.

Data center consolidation scenario

We use this scenario to consolidate data centers by collecting the data from remote locations and using the TS7700 Virtualization Engine Grid to move the data to the centralized data center.

In this scenario, you potentially have two clusters at the primary data center for high availability. The third cluster is located at a remote data center from which you want to move data. Figure 2-46 shows this initial status.

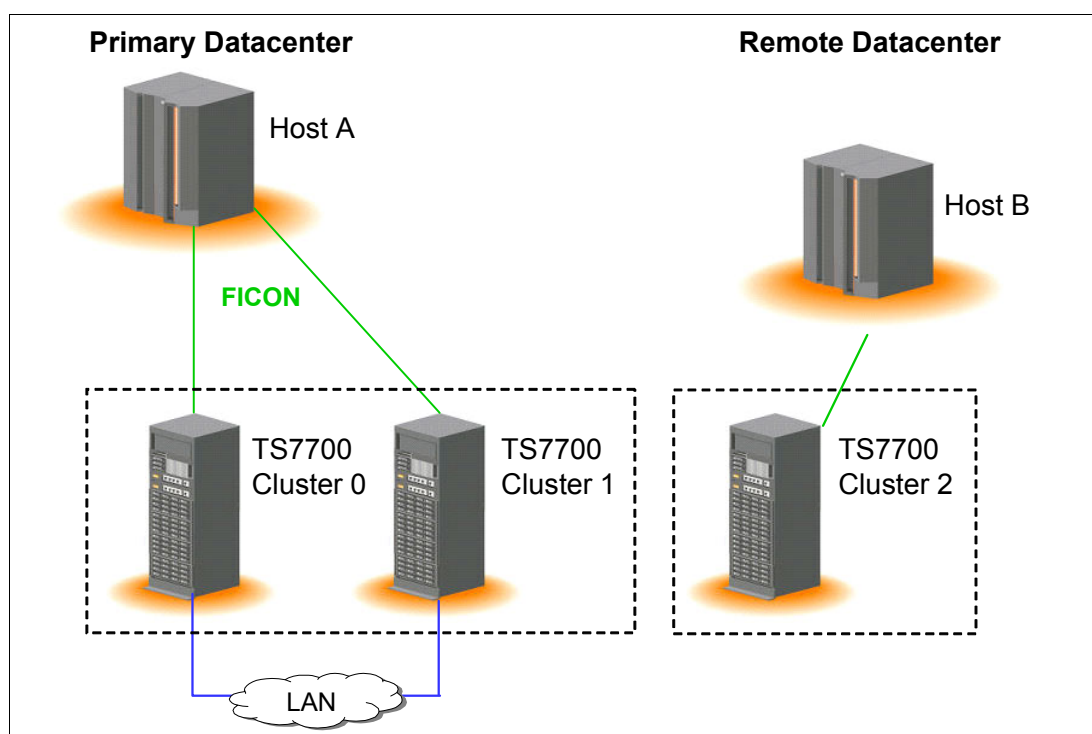


Figure 2-46 Data center consolidation initial status

Cluster 2 is joined with the two existing clusters. You use the remote hosts to copy data from their existing infrastructure into Cluster 2. The data is then moved to the primary data center using Grid replication. Figure 2-47 on page 102 shows this phase.

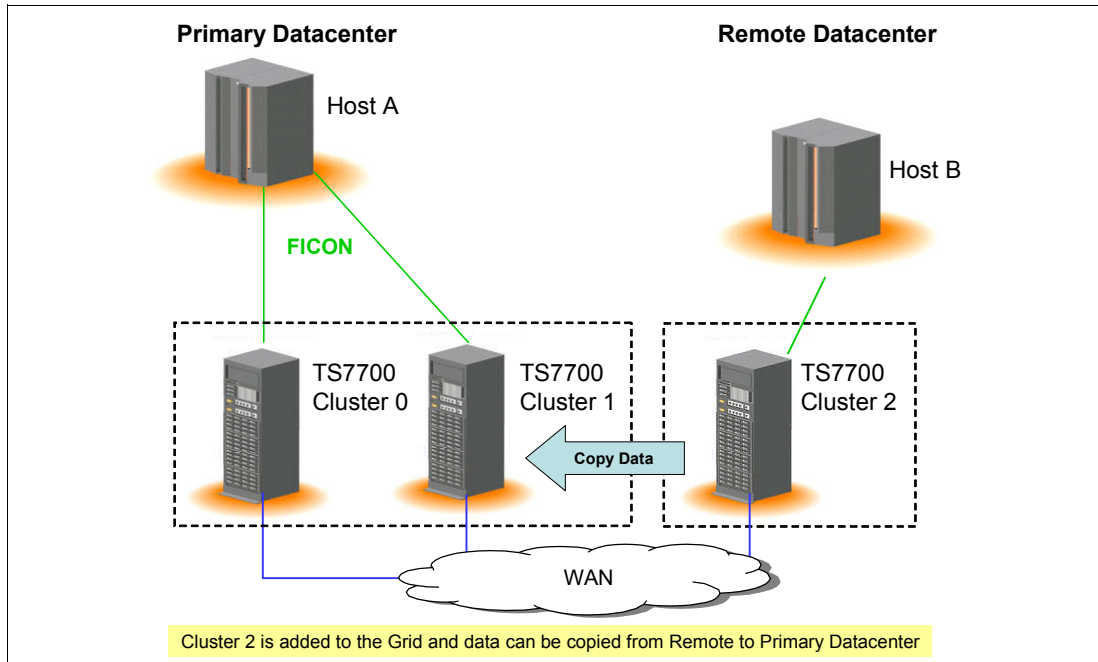


Figure 2-47 Data center consolidation join to Grid

After all of your data has been copied to the primary data center TS7700 Virtualization Engines, you can remove the third cluster from the remote data center and clean up the data from it. This TS7700 Virtualization Engine can now be relocated and the process can be repeated. Figure 2-48 shows this final status.

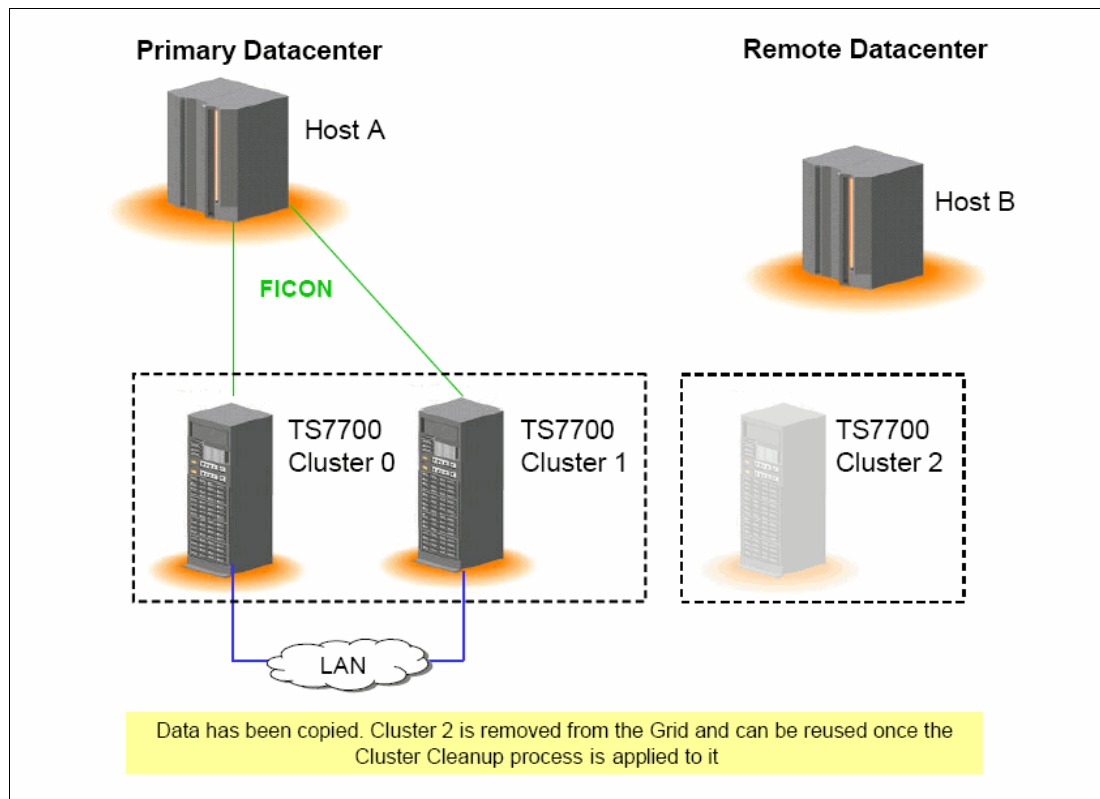


Figure 2-48 Data center consolidation final

TS7700 Virtualization Engine reuse

For a TS7700 Virtualization Engine reuse scenario, we assume that you have a multi-site Grid configuration and one TS7700 Virtualization Engine Cluster is no longer required at one site. This TS7700 Virtualization Engine Cluster can be removed (after all required data has been copied, removed, or expired) and you can redeploy this resource. However, before it can be reused, it needs to be removed from the Grid and cleaned up.

An important aspect to mention is that the cleanup process does not include physical cluster discontinuance or reinstallation. If you are planning to reuse a TS7700 Virtualization Engine after it is removed from a Grid, you must have a service contract for physical discontinuance services and the subsequent reinstall of the cluster.

During planning you must determine how to handle the volumes that have a copy consistency point only at the cluster that is being removed. You can eject them, move them to the scratch category, or activate a Management Class change on a mount/demount operation to get a copy on another cluster. This step must be done prior to starting the remove process. A new Bulk Volume Information Retrieval (BVIR) option, *Copy Audit*, is provided for generating the list of inconsistent volumes to help the customer plan their activities. See Chapter 8, “Operation” on page 409 for details about BVIR.

The removal of the cluster from the Grid is concurrent with customer operations on the remaining clusters, except for certain steps where inserts, ejects, and exports will be inhibited. Complete the removal of a cluster from the Grid during a period of minimized utilization of the remaining clusters.

No data, on cache or on tapes (if applicable), on the removed cluster will be available after the cluster has been removed from the Grid. The cluster cannot be rejoined with the existing data.

No secure erase or low level format will be done on the tapes or the cache as part of the removal of a cluster from a Grid and cluster cleanup process.

2.6.6 Joining of clusters at different code release levels

Prior to TS7700 Virtualization Engine Release 1.7, a requirement was for all clusters, which were to be joined, to be at the same release level. Enhancements in TS7700 Virtualization Engine R1.7 provide the capability to join a TS7700 Virtualization Engine R1.7 with a stand-alone or Grid configuration having Virtualization Engine R1.6 and R1.7 members. This provides additional installation planning flexibility for customers expanding their Grid capacity without the prerequisite of all clusters having to be at the same release level.

Purchasing an additional or redeploying an existing TS7700 Virtualization Engine into a Grid configuration had the requirement of the joining cluster having to be at the same release level as the existing cluster (or clusters). In some cases, the need to upgrade Grid capacity created a conflict in environments with fixed code upgrade dates and specific change management policies. This requirement has been removed and allows the joining of a TS7700 Virtualization Engine R1.7 with a stand-alone or multi-cluster configuration with Virtualization Engine R1.6 members. Generally speaking, the functional characteristics available to the composite library will be defined by the Grid cluster with the lowest release level. After all TS7700 Virtualization Engines in a multi-cluster configuration are upgraded to the same level, all licensed functional characteristics provided in the release will be enabled. In a Grid configuration that has clusters with different release levels, the TS7700 Virtualization Engine

R1.7 clusters will be able use their full cache capacity but will not be able to provide new capabilities such as:

- ▶ Selective write protect for DR testing
- ▶ Enhanced removal policies

All licensed functional characteristics provided in the release will be available after all of the TS7700 Virtualization Engines in a multi-cluster configuration are at the same level.

Expanding an existing Grid implementation resulting in a configuration with different code levels must adhere to three explicit conditions:

- ▶ Only one cluster can be joined at a time. The single cluster must be empty and not contain existing data.
- ▶ R1.7 systems can join existing configurations running all R1.6 or a mixture of R1.6 and R1.7.
- ▶ The join is performed to the existing cluster that has highest release level.

To present the capability to join clusters with different code levels, we use an upgrade scenario in which we have an existing two-cluster Grid at TS7700 Virtualization Release 1.6 and we want to expand it with two TS7700 Virtualization Engines at R1.7. We also apply the explicit rules to our upgrade path. See Figure 2-49.

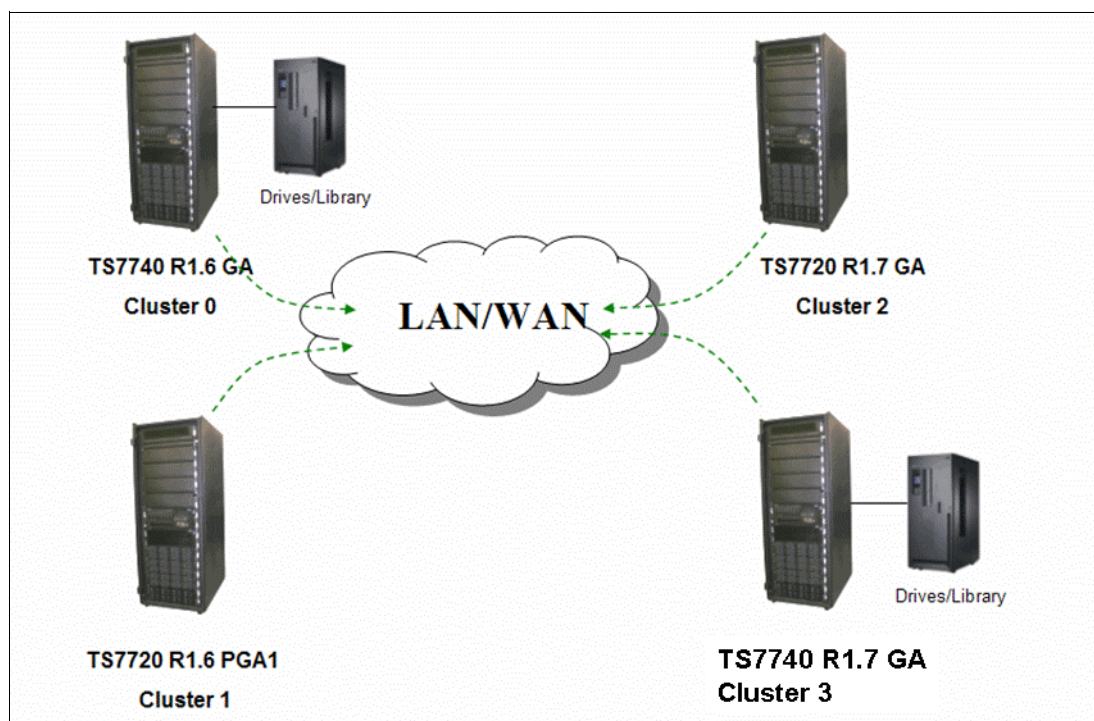


Figure 2-49 New Cluster Join With different Code Levels

The existing Grid configuration consists of Cluster 0, a TS7740 Virtualization Engine R1.6 GA, and Cluster 1, a TS7720 Virtualization Engine R1.6 PGA1. We will be joining a TS7720 Virtualization Engine R1.7 as Cluster 2 and a TS7740 Virtualization Engine R1.7 as Cluster 3. The identified code levels of the clusters that are shown in Figure 2-49 support the Grid join capability provided in R1.7.

We initially join Cluster 2 to Cluster 1 to satisfy the requirements to join one cluster at a time and target the existing cluster with the highest release level. We have now established a three

cluster Grid configuration. The cluster with the highest release level in the Grid is now Cluster 2. We then initiate a join of Cluster 3 to Cluster 2 to complete our goal of a four-cluster Grid configuration. We have now completed the formation of our four-cluster Grid configuration in the required order.

Note: TS7700 Virtualization Engine R1.7 supports joining to a TS7700 Virtualization Engine R1.6. TS7700 Virtualization Engine R1.7 cannot join a Grid with clusters at R1.5 or earlier. Future releases might provide N-1 support for joining various release levels.

2.7 Expanded Memory

This new feature provides capabilities to support additional physical memory for R1.7 systems. All new built systems starting with Hydra 1.7 will be shipped with a total of 16 GB of physical memory installed.

Existing TS7700 systems remain at 8 GB of physical memory, there is no plan to add memory to all existing systems. However, an upgrade from 8 GB to 16 GB of physical memory will be offered for existing systems. This update is targeted for systems that have more than 500,000 volumes and experience heavy host I/O throughput. Those systems that have the Expanded Memory, TS7700 R1.7 code will be performance tuned to take advantage of the additional memory.

The main objective of adding Expanded Memory is to provide an increase in performance. This is achieved by reduced system paging and allowing more data to be cached, thereby increasing overall performance.



Preinstallation planning and sizing

This chapter provides information to help you plan the installation and implementation of the IBM Virtualization Engine TS7700.

We cover the following topics:

- ▶ Hardware configurations
- ▶ Hardware installation and infrastructure planning
- ▶ Remote installations and switch support
- ▶ High availability Grid considerations
- ▶ Tape Library attachment
- ▶ Planning for software implementation
- ▶ Planning for logical and physical volumes
- ▶ Planning for encryption in the IBM Virtualization Engine TS7740
- ▶ Tape analysis and sizing the TS7740 Virtualization Engine
- ▶ Education and training
- ▶ Planning considerations

Use the preinstallation checklists in Appendix C, “TS3500 and TS7700 checklists” on page 803 to help plan your installation.

Note: For this chapter, we use the term “Tape Library” to refer to the IBM System Storage TS3500 Tape Library.

3.1 Hardware configurations

In this section, we describe the minimum configurations and the optional enhancements for the TS7700 Virtualization Engine. We cover both the IBM TS7740 Virtualization Engine and the IBM TS7720 Virtualization Engine.

Tape Library attachment

TS7700 Virtualization Engine Release 1.5 and later no longer implement a physical Library Manager. All virtual volume processing is now done in the TS7700 Virtualization Engine subsystem itself. The TS7720 Virtualization Engine is a disk-only solution and therefore requires no physical tape library. Consider the following details:

- ▶ Enterprise Library Controller (ELC) functions and associated components are integrated into the TS7700 Virtualization Engine frame, and includes the TS3000 Service Console.
- ▶ When migrating from levels earlier than R1.5 to R1.7, the migration of the 3953 ELC functions is automatically included.
- ▶ The IBM 3953 Tape System, including the 3953 Tape Frame and the 3953 ELC is still required for native drive applications.

See 2.3.4, “TS3500 Tape Library” on page 49 for more information about the TS3500 Tape Library and “Replacing the 3953-L05 Library Manager” on page 321 for information about the Licensed Internal Code upgrade to TS7740 Virtualization Engine Release 1.7.

Note: At the time of writing, TS7740 Virtualization Engine configurations with an IBM 3494 Tape Library will require initial request and approval of RPQ 8B3575 before ordering Release 1.7. Contact your IBM representative for availability of TS7740 Virtualization Engine Release 1.7 with the IBM 3494 Tape Library.

TS7740 Virtualization Engine configuration summary

This section summarizes the hardware components and configuration requirements for a single-cluster grid and for a multi-cluster grid configuration.

A TS7740 Virtualization Engine is composed of the following components:

- ▶ Frame (3952-F05)
- ▶ TS7740 Virtualization Engine Server
 - An IBM System p server as the node (3957 Model V06) with two 64-bit CPUs, 2.1 GHz with 16 GB of 266 MHz DDR2 memory and 36 MB of L3 cache memory
 - Integrated Enterprise Library Controller
- ▶ Disk cache: 3956 Model CC8 with zero to three 3956 Model CX7 providing up to 28.17 TB
- ▶ I/O Expansion drawers
- ▶ Redundant network routers
- ▶ Redundant Fibre Channel switches located in TS3500 Tape Library providing connectivity to TS1130/TS1120/3592-J1A drives
- ▶ TS3000 System Console & Keyboard/Display: Can be considered optional if you already have an external TS3000 Service Console.

TS7720 Virtualization Engine disk only configuration summary

A TS7720 Virtualization Engine is composed of the following components:

- ▶ Frame (3952-F05)
- ▶ TS7720 Virtualization Engine Disk Only Server
 - A IBM System p server as the node (3957 Model VEA) with two 64-bit CPUs, 2.1 GHz with 16 GB of 266 MHz DDR2 memory and 36 MB of L3 cache memory
 - Integrated Enterprise Library Controller
- ▶ I/O Expansion drawers.
- ▶ Redundant network routers.
- ▶ TS3000 Service Console & Keyboard/Display:
Can be considered optional if you already have an external TS3000 Service Console.
- ▶ Disk cache: A 3956 Model CS8 with zero to six 3956 Model XS7 for a maximum cache size of 161 TB
- ▶ Optional Storage Expansion Frame (3952-F05): Two 3956 Model CS8 with zero to ten 3956 Model XS7 for a maximum cache size of 441 TB

All components are installed in IBM 3952 Tape Frames. The Virtualization Engine is connected to the host through FICON channels. In a multi-cluster grid configuration, where we can have two, three, or four TS7700 Virtualization Engines, two 1 Gb Ethernet links per TS7700 Virtualization Engine are used to interconnect the clusters.

3.1.1 TS7740 Virtualization Engine configuration details

In this section, we describe the specific features of the TS7740 Virtualization Engine in a stand-alone and Grid configuration.

TS7740 Virtualization Engine minimum configuration

The minimum configuration of a TS7740 Virtualization Engine requires the following components, the feature code (FC) number is in parenthesis:

- ▶ One 3952 Tape Frame Model F05 with the following features:
 - TS7700 Virtualization Engine Base Frame (FC7312)
 - Plant Install 3957 V06 (FC5628)
 - Plant Install 3956 CC8 (FC7123)
 - Integrated Control Path (FC5759)
 - Dual AC Power (FC1903)

A power cord appropriate for the country of installation must be selected from features FC9954 through FC9959, or FC9966.
- ▶ One TS7740 Virtualization Engine Server (3957 Model V06) with the following required features:
 - One 1-TB Cache Enablement (FC5267)
 - 8 GB Memory Upgrade - Plant (FC9461)
 - One 100 MBps Increment (FC5268)
 - 3592 Tape Drives (FC5240)
 - Mainframe Attach (FC9000)

- One attach to TS3500 Tape Library (FC9219)
- Plant Install V06 in 3952 F05 (FC9350)
- Two of either, 1 Gb Grid Dual Port Copper Connection (FC1032), or 1 Gb Grid Dual Optical SW Connection (FC1033) to provide for Grid communication between TS7700 Virtualization Engine servers.
- Either two of host to Virtualization Engine FICON cables (FC0201 through FC0206) or one No Factory Cables (FC9700).
- Two of either FICON Shortwave Attachment (FC3441) or FICON Longwave Attachment (FC3442), or FICON 10 km Longwave Attachment (FC3443).
- Either Console Expansion (FC2714) or Console Attachment (FC2715).

3957 Model VEA features FC2714 and FC2715 may not be attached to a TotalStorage Master Console for Service (FC2713).

If the configuration has native tape drives, the TS3000 System Console (FC2721) may be installed in a 3953 Model F05, connecting to the TS7740 Virtualization Engine Server using Console Expansion (FC2714) or Console Attachment (FC2715).

The TS3000 System Console (FC2730 and FC2719) may be installed in the 3952 Tape Frame Model F05, connecting to the TS7740 Virtualization Engine Server using Console Expansion (FC2714) or Console Attachment (FC2715).

When FC2714 or FC2715 are installed on 3957 Model V06, the Console Upgrade (FC2719) is required on the machine type model where FC2718, FC2720, FC2721, or FC2730 is installed, as follows:

- ▶ One TS7740 Virtualization Engine Cache Controller (3956 Model CC8) with the following required features:
 - 9.6 TB Fibre Storage (FC7123)
 - Plant Install in 3952 F05 (FC9352)
 - Attach to 3957 V06 (FC9230)
 - Four of Intraframe Fibre Channel cable (FC6003)
- ▶ From 4 to 16, 3592 Tape Drives (Model J1A, E05, E06, or EU6) must be installed in a TS3500 Tape Library (3584 Tape Library) and connected to the TS7700 Virtualization Engine through the 4 Gb FC Switches. All attached tape drives must operate in the same mode. Therefore, 3592 Model E05 Tape Drives operating in native mode may not be intermixed with 3592 Model J1A or Model E06 / EU6 Tape Drives. 3592 Model E06 and EU6 Tape Drives may not be intermixed with other 3592 Tape Drives.

Note: FC9217 is no longer supported beginning with Release 1.5. FC9217 refers to the switches being installed in the 3953 F05 frame. The switches can remain in the 3953 F05 frame after upgrading to R1.5 and later.

- ▶ When FC9219 is installed on the TS7740 Virtualization Engine Server (3957 Model V06), the following items are required:
 - Two 4 Gb Fibre Channel switches are required in either the L23 or D23 frame:
 - If the switches remain in the 3953 Model F05, two of feature FC3488, 4 Gb Fibre Channel Switches, or FC4897, Reinstall 4 Gb Fibre Channel Switches, provide the switches.

- To remove the switches from the 3953 Model F05, order feature FC4748. Remove the 4 Gb switch, for each switch to be removed.
- The switches removed from the 3953 Model F05 may be installed in the 3584 Model L23 or D23 frame using feature number FC4873. Reinstall TS7740 Virtualization Engine back-end switches, which requires FC4871, TS7740 Virtualization Engine back-end switch mounting hardware.
- Two new 4 Gb Fibre Channel switches may be ordered for the 3584 frames using FC4872, TS7740 Virtualization Engine back-end 4 Gb switches.
- To install new switches in a 3584 frame or reinstall switches removed from a separate subsystem, the following features are required:
 - FC4871, TS7700 Virtualization Engine back-end switch mounting hardware
 - FC1950, Power Distribution Units
 - One power cord feature FC9954 through FC9959 or FC9966
- One or more 3584 Model L23 or D23 frames with the following components:
 - From 4 to 16,3592 tape drives. All attached tape drives must operate in the same mode. Therefore, 3592 Model E05 Tape Drives operating in native mode may not be intermixed with 3592 Model J1A or Model E06 / EU6 Tape Drives. 3592 Model E06 and EU6 Tape Drives may not be intermixed with older generation 3592 Tape Drives.
 - Up to 16 FC4874, Adjacent Frame Support for Back-End Drives. This provides the Fibre Channel cables to allow cabling from the back-end Fibre Channel switch to the drives.

Note: If the TS7740 Virtualization Engine was previously attached to the TS3500 Tape Library utilizing a 3953 Model L05 Library Manager, you might choose to leave the FC switches in the 3953 Model F05 Tape Frame. Optionally, the switches may be removed and reinstalled in the TS3500 Tape Library.

TS7740 Virtualization Engine configuration upgrades

The following options may be installed to modify the minimum TS7740 Virtualization Engine configuration:

- ▶ 3952 Tape Frame Model F05 with the following upgrade:
 - IBM System Storage TS3000 System Console (FC2730) with Console Upgrade (FC2719).
- ▶ TS7740 Virtualization Engine Server (3957 Model V06) with the following upgrades:
 - Grid Enablement (FC4015) allows this TS7700 Virtualization Engine Server to communicate to other TS7700 Virtualization Engine Servers through the Grid. Grid Enablement must be installed on each TS7700 Virtualization Engine Server that participates in the communication Grid. With R1.7, we support both TS7740 Virtualization Engine and TS7720 Virtualization Engine participation in hybrid and homogenous TS7720 Virtualization Engine Grid configurations.

The customer must provide appropriate Ethernet cables to attach the Grid Connection adapters (FC1030 through FC1033) to the communication Grid when Grid Enablement (FC4015) is installed.

- Two additional FICON attachments may be installed to provide a total of four FICON attachments on the TS7740 Virtualization Engine Server. Valid total quantities of FICON Shortwave Attachment (FC3441), FICON Longwave Attachment (FC3442), and FICON 10 km Longwave Attachment (FC3443) are two or four.

- Additional 1 TB Cache Enablements (FC5267) may be installed, up to a maximum quantity of 28.
- Additional 100 MBps Throughput Increments (FC5268) may be installed, up to a maximum quantity of 6.
- Encryption Configuration (FC9900) allows the TS7740 Virtualization Engine Server to be configured to support encryption for designated storage pools. To use encryption with the TS7700 Virtualization Engine, all attached tape drives must be TS1120 encryption capable tape drives, or TS1130 Tape Drives. TS1120 drives must be operating in 3592 Model E05 native mode with encryption enabled.
- ▶ TS7740 Virtualization Engine Cache Controller with the following upgrades:
 - One or three additional 3956 Model CX7 Cache Drawers may be installed on an existing 3956 Model CC7 Cache Controller. Valid total quantities of TS7740 Virtualization Engine Cache Drawer Model CX7 are zero, one, and three. Each 3956 Model CX7 must contain FC7121.
 - Enable 1st Expansion Drawer, FC7403, must be installed on the TS7740 Virtualization Engine Cache Controller Model CC7 when the first 3956 Cache Drawer Model CX7 is installed.

3.1.2 TS7720 Virtualization Engine configuration details

In this section, we describe the specific features of the TS7720 Virtualization Engine in the stand-alone and Grid configurations.

TS7720 Virtualization Engine minimum configuration

The minimum configuration of a TS7720 Virtualization Engine requires the components described in this section.

IBM 3592 Tape Frame Model F05

One 3952 Tape Frame Model F05 with the following required features is required:

- ▶ TS7720 Virtualization Engine Frame (FC7322)
- ▶ Plant Install 3957 VEA (FC5626)
- ▶ Plant Install 3956 CS8 (FC5635)
- ▶ Integrated Control Path (FC5759)
- ▶ Dual AC Power (FC1903)

A power cord appropriate for the country of installation must be selected from FC9954 through FC9959, or FC9966.

TS7720 Virtualization Engine Server (3957 Model VEA)

One TS7720 Virtualization Engine Server (3957 Model VEA) with the following required features is required:

- ▶ 100 MBps Throughput - Plant (FC9268)
- ▶ 8 GB Memory Upgrade - Plant (FC9461)
- ▶ Mainframe Attach (FC9000)
- ▶ Plant Install in 3952 F05 (FC9350)
- ▶ Two of either 1 Gb Grid Dual Port Copper Connection (FC1032) or 1 Gb Grid Dual Optical SW Connection (FC1033) to provide for Grid Communication between TS7700 Virtualization Engine servers

- ▶ Either two of host to Virtualization Engine FICON cables (FC0201 through FC0206) or one No Factory Cables (FC9700)
- ▶ Two of either FICON Shortwave Attachment (FC3441), FICON Longwave Attachment (FC3442), or FICON 10 km Longwave Attachment (FC3443)
- ▶ Either Console Expansion (FC2714) or Console Attachment (FC2715). Note the following information:
 - 3957 Model VEA features FC2714 and FC2715 may not be attached to an IBM TotalStorage Master Console for Service (FC2713).
 - The TS3000 System Console (FC2721) may be installed on 3953 Model F05, connecting to the TS7720 Virtualization Engine Server using Console Expansion (FC2714) or Console Attachment (FC2715).
 - The TS3000 System Console (FC2730 and FC2719) may be installed on 3952 Tape Frame Model F05, connecting to the TS7720 Virtualization Engine Server using Console Expansion (FC2714) or Console Attachment (FC2715).
 - When FC2714 or FC2715 is installed on 3957 Model VEA, the Console Upgrade (FC2719) is required on the machine type model where FC2718, FC2720, FC2721, or FC2730 is installed.

TS7720 Virtualization Engine SATA Cache Controller (3956 Model CS8)

The TS7720 Virtualization Engine minimum configuration requires one Cache Controller, 3956 Model CS8. With the optional TS7720 Virtualization Engine Expansion frame two additional Cache Controller, 3956 Model CS8 are provided.

The 3952 Model F05 base frame houses the following components:

- ▶ One TS7720 Virtualization Engine SATA Cache Controller, 3956 Model CS8 providing 19.84 TB of usable storage
- ▶ The 2 TB disks are formatted with RAID 6 for maximum data integrity as follows:
 - Two arrays of five Data Disks + two Parity Disks
 - Two hot spares

One TS7720 Virtualization Engine SATA Cache Controller, 3956 Model CS8 with the following features is required:

- ▶ 32 TB SATA Storage (FC7114)
- ▶ Plant Install in 3952 F05 (FC9352)

TS7720 Virtualization Engine configuration upgrades

The following options may be installed to modify the minimum TS7720 Virtualization Engine configuration:

- ▶ 3952 Tape Base Frame Model F05 with the following upgrade:
 - TS3000 System Console (FC2730) with Console Upgrade 3 (FC2719).
 - 3952 Storage Expansion Frame attachment (FC9323)
- ▶ 3952 Tape Expansion Frame Model F05 with the following features:
 - TS7720 Expansion Frame identifier (FC7323)
 - Dual AC Power (FC1903)
 - Install 3956 Cache Controller Model CS8- Plant (FC5635)
 - Optional are zero to ten 3956 Cache Drawer Model XS7- Plant (FC5646)

- ▶ TS7720 Virtualization Engine Server (3957 Model VEA) with the following upgrades:
 - Grid Enablement (FC4015) allows this TS7700 Virtualization Engine Server to communicate to other TS7700 Virtualization Engine Servers through the Grid. Grid Enablement must be installed on each TS7700 Virtualization Engine Server that participates in the Communication Grid. The customer must provide appropriate Ethernet cables to attach the Grid Connection adapters (FC1032 or FC1033) to the Communication Grid when Grid Enablement (FC4015) is installed.
 - Two additional FICON attachments may be installed to provide a total of four FICON attachments on the TS7700 Virtualization Engine Server. Valid total quantities of FICON Shortwave Attachment (FC3441), FICON Longwave Attachment (FC3442), and FICON 10 km Longwave Attachment (FC3443) are two or four.
 - Up to five 100 MBps Increment (FC5268) may be installed.
 - Dual Port FC HBA for Storage Expansion Frame - Plant (FC5240)
- ▶ TS7720 Virtualization Engine SATA Cache Upgrades

The TS7720 Virtualization Engine is available in single and two frame configurations.

- A single frame TS7720 Virtualization Engine provides one Cache Controller, 3956 Model CS8 and up to six TS7720 Virtualization Engine SATA Cache Drawers, 3956 Model XS7s.
- With the optional TS7720 Virtualization Engine Expansion frame two additional Cache Controller, 3956 Model CS8 are provided and up to ten TS7720 Virtualization Engine SATA Cache Drawers, 3956 Model XS7s can be configured.

All TS7720 Virtualization Engine Cache Drawers, 3956 Model XS7s are optional when configured with TS7720 Virtualization Engine Cache Controller, 3956 Model CS8. When ordering TS7720 Virtualization Engine Cache Drawers, 3956 Model XS7 each will require the following features:

- ▶ 32 TB SATA Storage (FC7114)
- ▶ Plant Install in 3952 F05 (FC9354)

The 3952 Model F05 base frame must be configured with one TS7720 Virtualization Engine Cache Controller, 3956 Model CS8 and six TS7720 Virtualization Engine Cache Drawers, 3956 Model XS7 before the optional 3952 Model F05 Storage Expansion Frame can be added.

The optional 3952 Model F05 Storage Expansion Frame is always configured with two TS7720 Virtualization Engine Cache Controller, 3956 Model CS8s. Addition of TS7720 Virtualization Engine Cache Drawers, 3956 Model XS7 to the expansion frame is performed in single drawer increments. The installation is balanced across the two TS7720 Virtualization Engine Cache Controller, 3956 Model CS8s starting with the lower string first.

Note: When attached to a TS7720 Virtualization Engine Cache Controller, 3956 Model CS8 there are no spare drives in the TS7720 Virtualization Engine Cache Drawer, 3956 Model XS7. The two spare drives in the TS7720 Virtualization Engine Cache Controller, 3956 Model CS8 can be used anywhere within its string configuration.

TS7720 Virtualization Engine: Grid configurations

Support for a hybrid Grid configuration with a TS7720 Virtualization Engine and TS7740 Virtualization Engine subsystems is provided in R1.6 and later. Figure 3-1 shows a homogenous TS7720 Virtualization Engine Grid with its maximum configuration.




Figure 3-1 TS7720 Virtualization Engine Grid configuration

The configuration provides the following maximums:

- ▶ Up to 1024 virtual tape drives
- ▶ Up to 1.76 PB of usable storage space before compression
- ▶ Up to 1600 MBps

3.1.3 TS3500 Tape Library High Density frame for a TS7740 Virtualization Engine configuration

A TS3500 Tape Library High Density frame is supported by a TS7740 Virtualization Engine and holds the stacked volumes.

If you are planning to size TS7740 Virtualization Engine physical volume cartridges, this high density (HD) frame can be a solution in terms of floor space reduction.

The TS3500 Tape Library offers high-density, storage-only frame models (HD frames) designed to greatly increase storage capacity without increasing frame size or required floor space. The new HD frame (Model S24 for 3592 tape cartridges) contains high density storage slots, as shown in Figure 3-2 on page 116.




Figure 3-2 TS3500 Tape Library HD frame (left) and top-down view

HD slots contain tape cartridges in a tiered architecture. The cartridge immediately accessible in the HD slot is a Tier 1 cartridge. Behind that is Tier 2 and so on. The maximum tier in a 3592 (Model S24) HD slot is Tier 4.

HD frame model S24 provides storage for up to 1,000 3592 tape cartridges.

The base capacity of Model S24 is 600 cartridges, which are stored in Tiers 0, 1, and 2. To increase capacity to the maximum for each frame, it is necessary to purchase the High Density Capacity on Demand (HD CoD) feature. This feature provides a license key that enables you to utilize the storage space available in the remaining tiers.

3.1.4 Cables

This section lists cable feature codes for attachment to TS7700 Virtualization Engine and additional cables, fabric components, and cabling solutions.

Required cable feature codes

The following cable feature codes are needed for attachment to TS7700 Virtualization Engine.

A TS7740 Virtualization Engine Server or TS7720 Virtualization Engine Server with the FICON Attachment features (FC3441, FC3442, or FC3443) can attach to the FICON channels of a IBM System z9®, IBM System z10®, or IBM S/390® host using Fibre Channel cable features ordered on the TS7700 Virtualization Engine Server. A maximum of four Fibre Channel cables, each 31 meters in length, can be ordered.

Order one cable for each Host System Attachment by using the following cable features:

► 4 GBps FICON Long-Wavelength Attachment Feature (FC3442 and FC3443)

The FICON Long-Wavelength adapter included with FC3442 (4 Gbps FICON Long-Wavelength Attachment) or FC3443 (4 Gbps FICON 10 km Long-Wavelength Attachment) has an LC Duplex connector, and can connect to FICON Long-Wavelength channels of System z9, System z10, or S/390 servers using a 9 micron single-mode Fibre Channel cable.

If host attachment cables are desired, they can be specified with one of the following feature numbers:

- FC0201: 9 Micron LC/LC 31 meter Fibre Channel cable
- FC0202: 9 Micron LC/SC 31 meter Fibre Channel cable

► 4 GBps FICON Short-Wavelength Attachment Feature (FC3441):

The FICON Short-Wavelength adapter shipped with FC3441 has an LC Duplex connector, and can connect to FICON Short-Wavelength channels of System z9, System z10, or S/390 servers using a 50 micron or 62.5 micron multimode Fibre Channel cable. If host attachment cables are desired, they can be specified with one of the following feature code numbers:

- FC0203: 50 Micron LC/LC 31 meter Fibre Channel cable
- FC0204: 50 Micron LC/SC 31 meter Fibre Channel cable
- FC0205: 62.5 Micron LC/LC 31 meter Fibre Channel cable
- FC0206: 62.5 Micron LC/SC 31 meter Fibre Channel cable

Additional cables, fabric components, and cabling solutions

Conversion cables from SC Duplex to LC Duplex are available as features on the System z servers for customers currently using cables with SC Duplex connectors that now require attachment to Fibre Channel components with LC Duplex connections. See the *IBM 2064 Sales Manual* for more information about these conversion kits. It is located at the following address:

<http://www.ibm.com/support/docview.wss?uid=pos1R1001750>

Additional cable options, along with product support services such as installation, are offered by IBM Global Services' Networking Services. See the *IBM Virtualization Engine TS7700 Introduction and Planning Guide*, GA32-0567 for Fibre Channel cable planning information.

3.1.5 Limitations

Consider the following limitations when performing the TS7700 Virtualization Engine preinstallation and planning:

- Each TS7700 Virtualization Engine node (3957 Model V06 or VEA) must be located within 100 feet of the TS3000 System Console (TSSC).
- Intermix of 3592 Model E06 / EU6 Tape Drives with other 3592 Tape Drives is not supported.
- Cisco routers do not perform IPv4 to IPv6 translation.
- Only one port of the dual port Grid adapters, FC1032 and FC1033, will be activated on each adapter.
- Only one port of the dual port FICON adapters, FC3441, FC3442, and FC3443, will be activated on each adapter.

TS7700 Virtualization Engine Grid interconnect LAN/WAN requirements

This section gives the LAN/WAN requirements for the TS7700 Virtualization Engine cross-site Grid TCP/IP network infrastructure:

- ▶ The TS7700 Virtualization Engine uses your LAN/WAN to replicate logical volumes, access logical volumes remotely, and perform cross-site messaging. The LAN/WAN should have adequate bandwidth to deliver the throughput necessary to meet your data storage objectives.
- ▶ The cross-site Grid Network is 1 Gb Ethernet with either shortwave Fibre Channel or copper (RJ-45). For copper networks, Cat 5E or 6 Ethernet cabling may be used, but Cat 6 cabling can achieve the highest throughput.
- ▶ The TS7700 Virtualization Engine does not encrypt the data it sends over the LAN/WAN.

Note: To avoid any network conflicts, the following subnets must *not* be used for LAN/WAN IP addresses or management interface primary, secondary, or virtual IP addresses:

- ▶ 192.168.251.xxx
- ▶ 192.168.250.xxx
- ▶ 172.31.1.xxx

Network redundancy

The TS7700 Virtualization Engine provides two independent 1 Gb copper (RJ-45) or shortwave Fibre Channel Ethernet links for Grid Network connectivity. Connect each link through an independent WAN interconnection to avoid a single point of failure that would disrupt service to both WAN paths from a node.

Note: For performance reasons and utilization of the full 1 Gb speed, use Cat 6 cables. Also, use two 1 Gb Ethernet links between clusters in a Grid to benefit from Dynamic Grid Network Balancing.

TS1120 and TS1130 Tape Drives (TS7740 Virtualization Engine only)

Support for the third generation of the 3592 drive family was included in TS7700 Virtualization Engine Release 1.5. The third generation of the 3592 drive family is the TS1130 Tape Drive, Models E06 and EU6. The 3592 Model EU6 Tape Drive is only available as a field upgrade from the TS1120 Tape Drive Model E05. When a TS1130 Tape Drive is attached to a TS7740 Virtualization Engine, all attached drives must be TS1130 Tape Drives; no intermix is supported with 3592 J1A or 3592 E05 drives. If TS1130 drives are detected by the TS7700 Virtualization Engine and there are also 3592 J1A or TS1120s detected, the TS1130 drives will not be configured. TS1130 drives will read data that has been written by either of the prior generations of 3592 drives. The TS1130 drives will append to filling E05-formatted tapes and will mark filling J1A-formatted tapes read-only. All new write from beginning of tape will be in E06 format.

The TS7740 Virtualization Engine uses the native capacity of TS1120 (3592-E05) and TS1130 (3592-E06) Tape Drives with the following restrictions:

- ▶ If any TS1130 Tape Drives are installed, then no 3592-J1A drives or TS1120 drives can exist in the TS7740 Virtualization Engine configuration.
- ▶ If any TS1120 Tape Drives are installed and configured in their native mode, then no 3592-J1A drives can exist in the TS7740 Virtualization Engine configuration.
- ▶ For encryption, all tape drives attached to the TS7740 Virtualization Engine must be TS1130 or TS1120 Tape Drives. The combination of encryption-capable and non-encryption-capable drives is not supported.

Note: If the TS7740 Virtualization Engine is returned to 3592 J1A emulation mode after physical volumes have been written in native 3592 E05 or E06 mode, the TS7740 Virtualization Engine does not return to an online status.

To address compatibility issues, the TS1120 Tape Drives can operate in native 3592 E05 mode or in 3592-J1A emulation mode. This way allows established configurations to add more back-end TS1120 drives when 3592-J1A drives already exist. However, all TS7740 Virtualization Engine back-end drives must operate in the same mode. Therefore, in a mixed configuration of TS1120 and 3592-J1A drives, the TS1120 Tape Drives must run in 3592-J1A emulation mode to provide a homogenous back-end tape format. For operation in TS1120 native mode, all back-end drives must be TS1120 Tape Drives. To use the TS1120 Tape Drives in emulation mode, the back-end drives must be configured by the IBM System Services Representative (SSR).

Note: An intermix of TS1130 E06 Tape Drives and prior 3592 drive types is not supported.

Supported media type

The TS1120, TS1130, and the 3592-J1A drives use the following rewritable media:

- ▶ JA cartridge (MEDIA5) with a native capacity of 300 GB in J1A Emulation mode, 500 GB in E05 native mode, and 640 GB in TS1130 native mode
- ▶ JJ cartridge (MEDIA7) with a native capacity of 60 GB in J1A Emulation mode, 100 GB in E05 native mode, and 128 GB in TS1130 native mode
- ▶ JB cartridge (MEDIA9) with a native capacity of 700 GB (TS1120 only) and 1 TB in TS1130 native mode

Note: Physical Write Once Read Many (WORM) cartridges are not supported on TS7700 Virtualization Engine attached tape drives.

Performance Scaling cannot be used on stacked volumes in a TS7700 Virtualization Engine.

TS3000 System Console

The TS3000 System Console (TSSC) is *mandatory* in a TS7700 Virtualization Engine environment. The TS7700 Virtualization Engine has a function to call the IBM Support center and provide summary information when a problem is detected. In addition, the TS7700 Virtualization Engine can be configured to automatically execute a health check and send the results to the TSSC for transmission to IBM for review and preventive actions.

The TSSC integrates service monitoring of up to forty-three controllers. It centralizes maintenance and service terminals, and enhances remote support capability. It is the only method for electronic call home, problem determination, and repair of the 3956 Cache.

The TSSC provides a service interface to the TS7700 Virtualization Engine through the following ways:

- ▶ Ethernet using Console Launcher
- ▶ Serial connection to the TS7700 Virtualization Engine
- ▶ Graphical User Interface (GUI) for servicing the 3956-CC6/CX6 or 3956-CS7/XS7

You can have two types of system consoles:

- ▶ Tower server (IBM System x 206m)

This stand-alone TSSC can be attached to 3494 Tape Library models Lxx, Bxx, CX1, 3590 Model A60, 3592 Model J70, TS1120 Model C06, TS3500 Tape Library, 3957 Model V06, and 3956 Model CC6.

The tower server is shown at the top of Figure 3-3.

- ▶ Rack-optimized server

This TSSC can be installed inside an IBM 3953-F05 frame or 3952-F05 frame for TS7700 Virtualization Engine configurations.

The rack-optimized server is shown at the bottom of Figure 3-3.



Figure 3-3 Tower server and rack-optimized TSSC

3.1.6 TS7700 Virtualization Engine upgrades

In this section, we review the existing TS7700 Virtualization Engine component upgrades for functions you might want to enable if you do not have the maximum TS7700 Virtualization Engine configuration.

Licensed Internal Code upgrades from levels earlier than Release 1.5 require a hardware reconfiguration scenario. See 6.2, “TS7700 Virtualization Engine upgrade to Release 1.7” on page 319 if you currently have a TS7740 Virtualization Engine with a microcode release prior to R1.5. In scenarios where your TS7720 Virtualization Engine is at R1.5 or if your TS7740 Virtualization Engine is at R1.5 and is configured with a TS3500 Tape Library, hardware migration is not required.

Currently, the upgrade options for the TS7700 Virtualization Engine include changing the number of tape drives, the grid connection, the amount of cache, and the performance increments.

TS7700 Virtualization Engine components upgrades

Several field-installable upgrades are available to enable an existing TS7700 Virtualization Engine for additional functions or capacities. In this section, we describe these upgrades, which can be both concurrent and disruptive.

Table 3-1 shows upgrades that are available for each model of the TS7700 Virtualization Engine. In the table, features marked with an asterisk (*) are removable features.

Table 3-1 TS7700 Virtualization Engine upgrades overview

Feature	TS7740 Virtualization Engine	TS7720 Virtualization Engine
FC1030 1 Gb Grid Copper*	Yes	No
FC1031 1 Gb Grid Optical*	Yes	No
FC1032 1 Gb Grid Dual Port Copper Connection*	Yes	Yes
FC1033 1 Gb Grid Dual Port Optical Connection*	Yes	Yes
FC3441 FICON SW*	Yes	Yes
FC3442 FICON LW*	Yes	Yes
FC3443 FICON 10 km*	Yes	Yes
FC4015 Grid Enablement	Yes	Yes
FC4016 Remove Cluster from Grid	Yes	Yes
FC4017 Cluster Cleanup	Yes	Yes
FC5240 Attach 3592 Tape Drives	Yes	No
FC5267 1 TB Cache Enablement	Yes	No
FC5268 100 MBps * (One FC9268 + Max of five FC5268s for TS7720 Virtualization Engines)	Yes	Yes
FC9219 Attach to a TS3500 Tape Library	Yes	No
FC9900 Encryption Enablement	Yes	No

Concurrent system component upgrades

Concurrent system upgrades can be installed when the TS7700 Virtualization Engine is online and operating. The following component upgrades can be made concurrently to an existing on-site TS7700 Virtualization Engine:

► FICON adapters

You can install FICON adapters to convert a two-FICON configuration to a four-FICON configuration, or to replace one pair of FICON adapters of a certain type with a pair of another type (shortwave/longwave or 4 km/10 km).

Replacement of an existing FICON adapter will require removal of the original feature and the addition of the new feature. Use FC3441, FICON Short-Wavelength Attachment, FC3442, FICON Long-Wavelength Attachment, and FC3443, FICON 10 km Long-Wavelength Attachment, to achieve these upgrades.

► Incremental disk cache capacity

This function is only available for TS7740 Virtualization Engine hardware.

You can add a 1 TB increment of disk cache to store logical volumes, up to 28 TB. Use FC5267, 1 TB Cache Enablement, to achieve this upgrade:

- For the 3956-CC6/CX6 cache, the number of usable features is 9.
- For the 3956-CC7/CX7 cache, the number of usable features is 14.
- For the 3956-CC8/CX7 cache, the number of usable features is 28.

► Incremental data throughput

Performance throughput capacity is provided in terms of nominal 100 MBps increments measured as data transferred through the TS7700 Virtualization Engine Server in a cluster, prior to any effect of compression. Based on the number of increments installed, the TS7700 Virtualization Engine will limit the host data transfer rate evenly across all running jobs. All host data transfers through the cluster with the virtual mount are taken into consideration for the data transfer limit. If a remote cluster is accessed through the Grid to satisfy a virtual mount, the incremental data throughput limit of the local cluster performing the mount is applied.

The number of features installed dictates when the system will limit the total throughput. Throughput might also be limited by other factors, such as record block sizes, compression ratio, and host channel capabilities.

You can add a 100 MBps increments of peak data throughput (FC5268) in any quantity as follows:

- From one to six for the TS7740 Virtualization Engine (3957-V06)
- From one to five for the TS7720 Virtualization Engine (3957-VEA)

This function provides you the flexibility to tailor a solution to your specific performance needs. With the TS7720 Virtualization Engine, the first performance increment is included with the 3957-VEA, and is specified by FC9268. When the maximum number of six performance increments are installed in a TS7700 Virtualization Engine, the system will no longer restrict performance.

The number of performance increments does not have to be the same on every cluster in a TS7700 Virtualization Engine Grid configuration.

Note: See Appendix A, “Feature codes” on page 781 for specific information about Feature Codes.

Non-concurrent system component upgrades

Non-concurrent upgrades require the TS7700 Virtualization Engine Cluster to be brought offline prior to installation. In some instances, the targeted component must be reconfigured before the upgrade will take effect. The following component upgrades must be made non-concurrently to an existing, on-site TS7700 Virtualization Engine.

TS7720 Virtualization Engine cache

The TS7720 Virtualization Engine R1.7 provides a broad range of supported cache sizes. Single- and two-frame configurations are available for the TS7720 Virtualization Engine.

- ▶ TS7720 Virtualization Engine Base Frame SATA Cache Upgrades allow for up to seven Cache Drawer, 3956 Model XS7s when installed with Cache Controller, 3956 Model CS8. Each Cache Drawer, 3956 Model XS7 provides 23.84 TB of usable storage.
- ▶ TS7720 Virtualization Engine Expansion Frame SATA Cache Upgrades when added to the Base Frame configuration is maximized. Zero to ten additional Cache Drawer, 3956 Model XS7s when installed with Cache Controller, 3956 Model CS8. Each Cache Drawer, 3956 Model XS7 provides 23.84 TB of usable storage.

TS7740 Virtualization Engine cache size

The TS7740 Virtualization Engine supports from 1 TB to 28 TB of tape volume cache in 1 TB increments using FC5267. A TS7740 Virtualization Engine supports physical cache sizes up to 28 TB in a 3956-CC8/CX7 configuration. See Table 3-2, which lists the physical cache options, whether the cache can be upgraded, and the possible cache enablement features. For a complete list, see Appendix A, “Feature codes” on page 781.

Table 3-2 Physical cache size and cache enablement options

Starting usable cache size	Can be field cache upgraded?	Additional cache enablement features required	Final usable cache size
7.04 TB (1x 3956-CC8, with 7x FC5267)	Yes. one 3956-CX7s can be added by FC5642.	7x FC5267	14.08 TB
14.08 TB (1x 3956-CC8, 1x 3956-CX7), with 14x FC5267	Yes. Two 3956-CX7s can be added by 2x FC5642.	14x FC5267	28.17 TB

Cache increments are activated using feature licenses. One license key is required for each 1 TB of enablement. See 4.3.9, “TS7700 licensing” on page 234 for details concerning how feature licenses are installed. Cache increments can be installed either by the manufacturing plant or in the field.

As shown in Table 3-2 (for the TS7740 Virtualization Engine), you can have the following upgrades:

- ▶ One to two TS7740 Virtualization Engine Cache Drawers
You can add one TS7740 Virtualization Engine Cache Drawer to an existing one-drawer TS7740 Virtualization Engine Cache subsystem. Use FC5642, Field install 3956–CX7, to achieve this upgrade. This addition permits an increase in capacity from 7.04 to 14.08 TB for 3956-CC8 and 3956-CX7 configurations.
- ▶ Two to four TS7740 Virtualization Engine Cache Drawers
You can add two TS7740 Virtualization Engine Cache Drawers to an existing two-drawer TS7740 Virtualization Engine cache subsystem. Use FC5642, Field install 3956-CX7, to

achieve this upgrade. This addition permits an increase in capacity from 14.08 to 28.17 TB for 3956-CC8 and 3956-CX7 configurations.

Note: No MES is available to upgrade an existing one-drawer TS7740 Virtualization Engine cache subsystem directly to a four-drawer TS7740 Virtualization Engine cache subsystem. A one-drawer cache subsystem must be upgraded to a two-drawer cache subsystem before an upgrade to a four-drawer cache subsystem can occur. However, they can be performed sequentially within the same upgrade window.

Number of tape drives

For a TS3500 Tape Library, if less than 16 tape drives are included in the initial TS7740 Virtualization Engine configuration, you can add tape drives, up to a maximum of 16. An intermix of 3592 models J1A and E05 attached to the same TS7740 Virtualization Engine is supported as long as the TS1120 is running in J1A Emulation mode. See the restrictions listed in “TS1120 and TS1130 Tape Drives (TS7740 Virtualization Engine only)” on page 118 and “Supported media type” on page 119.

Grid connection

In a Grid configuration, you can change your Grid connection from two 1 Gbps Copper connections (two FC1030s) to two 1 Gbps optical SW fiber connections by ordering FC1031. FC1031 provides two shortwave optical adapters and instructions to replace the copper adapters. You can also change your Grid connection from two 1 Gbps Optical SW connections (two FC1031s) to two 1 Gbps Copper connections by ordering FC1030. You can also upgrade to the dual-ported copper and optical connections using FC1032 and FC1033, respectively. These are supported only with R1.5 and later.

You cannot intermix copper and optical fiber adapters in a TS7720 Virtualization Engine or TS7740 Virtualization Engine. All adapters must be of the same type in each TS7700 Virtualization Engine subsystem. However, with customer-provided network switching equipment for a Grid, the implementation can have some clusters with copper Grid interfaces and others with optical fiber adapters.

Note: Upgrading the Grid connections requires the TS7700 Virtualization Engine to be offline.

Existing TS7700 Virtualization Engine stand-alone or TS7700 Virtualization Engine Grid upgrades

The system configuration upgrades that are listed in this section may be made to an existing stand-alone TS7700 Virtualization Engine Cluster or TS7700 Virtualization Engine Grid.

Notes:

- ▶ To achieve any of the following upgrades, you must order FC4015, Grid enablement, for all participating TS7700 Virtualization Engine Servers.
- ▶ Joining or merging of TS7700 Virtualization Engine Servers may have release level dependencies. Review your upgrade plans with an IBM System Services Representative.
- ▶ Hybrid configurations with a mix of TS7720 Virtualization Engine and TS7740 Virtualization Engine subsystems is supported.

The upgrades are as follows:

- ▶ Two-cluster grid
 - You may add a new, empty TS7700 Virtualization Engine Cluster to an existing TS7700 Virtualization Engine Cluster to create a two-cluster grid.
 - You may combine two existing TS7700 Virtualization Engine Clusters to create a two-cluster grid.
- ▶ Three-cluster grid
 - You may add a new, empty TS7700 Virtualization Engine Cluster to an existing Two-Cluster TS7700 Virtualization Engine Grid to create a three-cluster grid.
 - You may add an existing TS7700 Virtualization Engine Cluster to an existing Two-Cluster TS7700 Virtualization Engine Grid to create a three-cluster grid.
- ▶ Four-cluster grid
 - You may add a new, empty TS7700 Virtualization Engine Cluster to an existing Three-Cluster TS7700 Virtualization Engine Grid to create a four-cluster grid.
 - You may add an existing TS7700 Virtualization Engine Cluster to an existing Three-Cluster TS7700 Virtualization Engine Grid to create a four-cluster grid.

Note: Currently, you may not add an existing two-cluster TS7700 Virtualization Engine Grid to an existing two-cluster TS7700 Virtualization Engine Grid.

Two TS7740 Virtualization Engines connected to a single tape library must not be combined to form a TS7740 Virtualization Engine Grid. Each TS7740 Virtualization Engine cluster in a TS7740 Virtualization Engine Grid configuration must possess its own physical tape library. For additional information about configuring a TS7700 Virtualization Engine Grid, see 3.4, “High availability Grid considerations” on page 137.

3.2 Hardware installation and infrastructure planning

The topics in this section provide planning information related to your TS7700 Virtualization Engine. Topics covered include system requirements, infrastructure requirements, host compatibility, and available feature codes.

3.2.1 System requirements

You must ensure that your facility meets the system requirements for the TS7700 Virtualization Engine when planning for installation. System requirements for installation include such things as requirements for power, cooling, floor leveling, loading, and distribution, clearance, environmental conditions, and acoustics.

IBM 3952 Tape Frame specifications

The 3952 Tape Frame houses the components of the TS7700 Virtualization Engine. The specifications are listed in the following tables:

- ▶ Dimensions and clearance are listed in Table 3-3 on page 126.
- ▶ Weight, source power, and heat output are listed in Table 3-4 on page 126.

Table 3-3 IBM 3952 Tape Frame specifications: Dimensions and clearance

Description	Side in mm (in.)	Front in mm (in.)	Height in mm (in.)
Dimensions	1102 (43.4)	644 (25.4)	1804 (71.0)
Operator/Service clearance	0 (0)	690 (27.2)	660 (26.0)

Table 3-4 IBM 3952 Tape Frame specifications: Weight, source power, and heat output

Description	Value
Weight	565.6 kg (1247 lb.)
Source power	240 V AC 15 amp
Heat output	4.0 kVA

Environmental operating requirements

Your facility must meet specified temperature and humidity requirements before installing the TS7700 Virtualization Engine. Table 3-5 lists the environmental conditions for the TS7700 Virtualization Engine.

Table 3-5 Environmental specifications

Condition	Air temperature	Altitude	Relative humidity	Wet bulb temperature
Operating	10-32°C (50-80°F)	Up to 1.523m (5.000 ft.) AMSL	20-80%	23°C (73°F)
Operating	10-28°C (50-82.4°F)	1.524-2.133m (5001-700 ft.) AMSL	20-80%	23°C (73°F)
Power Off	N/A	N/A	8-80%	27°C (80°F)
Storage	1-60°C (33.8°F-140°C)	N/A	5-80%	29°C (84°F)

For a complete list of system requirements, see *IBM Virtualization Engine TS7700 Series Introduction and Planning Guide*, GA32-0567.

Power considerations

Your facility must have an available power supply to meet the input voltage requirements for the TS7700 Virtualization Engine.

The 3952 Tape Frame houses the components of the TS7700 Virtualization Engine. The standard 3952 Tape Frame ships with one internal power distribution unit. However, FC1903, Dual AC power, is required to provide two power distribution units to support the high availability characteristics of the TS7700 Virtualization Engine. The 3952 Tape Expansion Frame has two power distribution units and requires two customer provided power feeds.

Table 3-6 on page 127 lists the input power requirements for the 3952 Tape Frame, which houses the components of the TS7700 Virtualization Engine.

Table 3-6 3952 Tape Frame input power requirements

Power requirement	Value
Nominal input voltage	200-240 V AC, single phase
Input frequencies	50 ± 3.0 Hz 60 ± 3.0 Hz

TS7720 Virtualization Engine power requirements

Your facility must have an available power supply to meet the input voltage requirements for the TS7720 Virtualization Engine. The 3952 Tape Expansion Frame has two power distribution units and requires two customer provided power feeds.

Table 3-7 lists the correct values.

Table 3-7 TS7720 Virtualization Engine input power requirements

Power requirement	Value
Voltage	200-240 V AC
Frequency	50-60 Hz
Current	20 A
Inrush current	250 A
Power (W)	4000 W
kVa (single phase)	4.0 kVa
Thermal units	13.7 KBtu/hr

3.2.2 TS7700 Virtualization Engine Release 1.7 Licensed Internal Code levels

To support a TS7740 Virtualization Engine R1.7 attached to a TS3500 Tape Library, minimum Licensed Internal Code levels are required on all components of the cluster.

The minimum Licensed Internal Code levels for TS7740 Virtualization Engine R1.7 include support for:

- ▶ Four-cluster grid configurations
- ▶ Cluster families
- ▶ TS7740 Virtualization Engine and TS7720 Virtualization Engine Hybrid Cluster configurations
- ▶ Logical WORM
- ▶ User security enhancements

The minimum microcode levels for the TS7740 Virtualization Engine R1.7 cluster at the time of writing are listed in Table 3-8 on page 128. Contact your IBM System Services Representative (SSR) for currently supported levels.

Table 3-8 Minimum microcode levels for R1.7

Component	Microcode level
TS7700 Virtualization Engine	8.7.0.xx
TS3500 Tape Library	8890/9590
IBM 3952 Model J1A Tape Drive	D3I0_C90
TS1120 Model E05 Tape Drive	D3I1_E4C
TS1130 Model E06/Eu6 Tape Drive	D3I2_611

3.2.3 TCP/IP configuration considerations

In this section, we discuss configuration considerations and LAN/WAN requirements for the TS7700 Virtualization Engine. We cover Single and multi-cluster grid configurations.

Single node configuration

There are two NAT routers embedded in the TS7700 Virtualization Engine and used for its internal network, for connection to the Management Console, and for users to access the management interface (MI) from their intranet.

To provide access to the MI, you need to supply three TCP/IP addresses and two network cables for the physical connection from your intranet to the TS7700 Virtualization Engine. Two of those IP addresses are static, assigned to each of the two physical connections within the two NAT routers. The third IP is used as a virtual IP address, which will automatically move between the two static ones in case an outage should occur to one of the routers. Therefore, only use this third, so-called *virtual IP address*, for your access to the TS7700 Virtualization Engine management interface, although you can technically use any of the three IP addresses.

Important: All three provided IP addresses that will be assigned to one TS7700 Virtualization Engine Cluster for the management interface access must be in the same subnet, for example:

- ▶ 192.168.1.1
- ▶ 192.168.1.2
- ▶ 192.168.1.3 (*virtual IP address*)

You must use subnet mask 255.255.255.0. Also, be aware of any gateway requirements of your intranet.

Grid configuration

In a Grid environment, all network components should be duplicated from the perspective of the TS7700 Virtualization Engine, as shown in Figure 3-4.

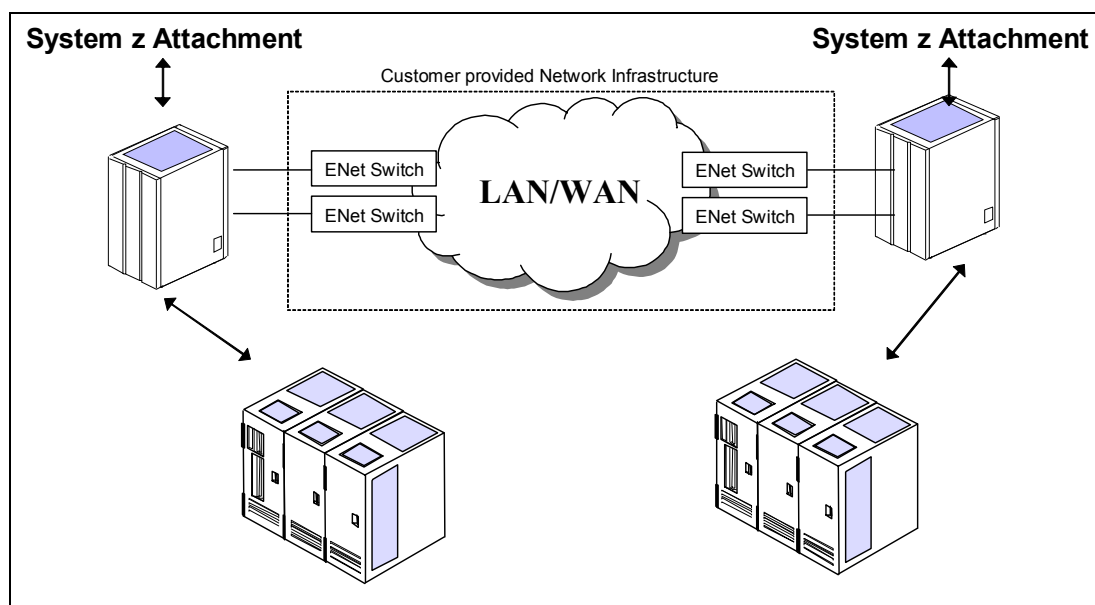


Figure 3-4 TS7740 Virtualization Engine two-cluster grid

The connections between the two TS7740 Virtualization Engine peers shown in Figure 3-4 are two separate 1 gigabit Ethernet (1 GigE) connections. The TS7700 Virtualization Engines are connected to the Grid network with customer-provided cables and switches. Standard TCP/IP protocol is used.

Notes:

- ▶ For performance reasons and utilization of the full 1 Gb bandwidth, use Cat 6 cables.
- ▶ Both links should provide 1 Gigabit Ethernet (1 GigE) between clusters to obtain the maximum benefit of the Dynamic Grid Network Load Balancing enhancements.

IP addresses for the management interfaces in a multi-cluster grid

In a TS7700 Virtualization Engine multi-cluster grid environment, you need to supply three IP addresses and two network cables per cluster. This provides the physical connections from your intranet to the TS7700 Virtualization Engines for accessing their management interfaces.

Similar to a Single Cluster configuration, three IP addresses are used for each cluster and two physical cable connections per cluster. Again, two of those IP addresses per cluster are static, assigned to each of the physical connections within the two TS7700 Virtualization Engine NAT routers. The third IP is used as a *virtual IP address* that automatically switches between the two routers if an outage occurs on either one.

Therefore, only use this third, *virtual IP address* per cluster, for access to each TS7700 Virtualization Engine management interface. Technically, you can use any of the three IP addresses per cluster.

Important: All three provided IP addresses that will be assigned to one TS7700 Virtualization Engine Cluster for the management interface access must be in the same subnet. As shown in Figure 3-5 on page 130, the local cluster is using 192.168.1.1, 192.168.1.2, and 192.168.1.3 with subnet mask 255.255.255.0.

In a TS7700 Virtualization Engine multi-cluster grid, it is possible to use a separate subnet to access each cluster's management interface, for each cluster, for redundancy. The remote TS7700 Virtualization Engine Cluster's management interface can also be accessed through the internal GigE Link, should the outside network to one of the clusters be unavailable.

This function is handled automatically by each cluster's management interface in the background. Figure 3-5 shows a sample setup for a two-cluster grid. For a three-cluster or four-cluster grid, simply increase the third octet by one so that a three-cluster grid uses addresses 192.168.3.1, 192.168.3.2, and 192.168.3.3, and a four-cluster grid uses 192.168.4.1, 192.168.4.2, and 192.168.4.3.

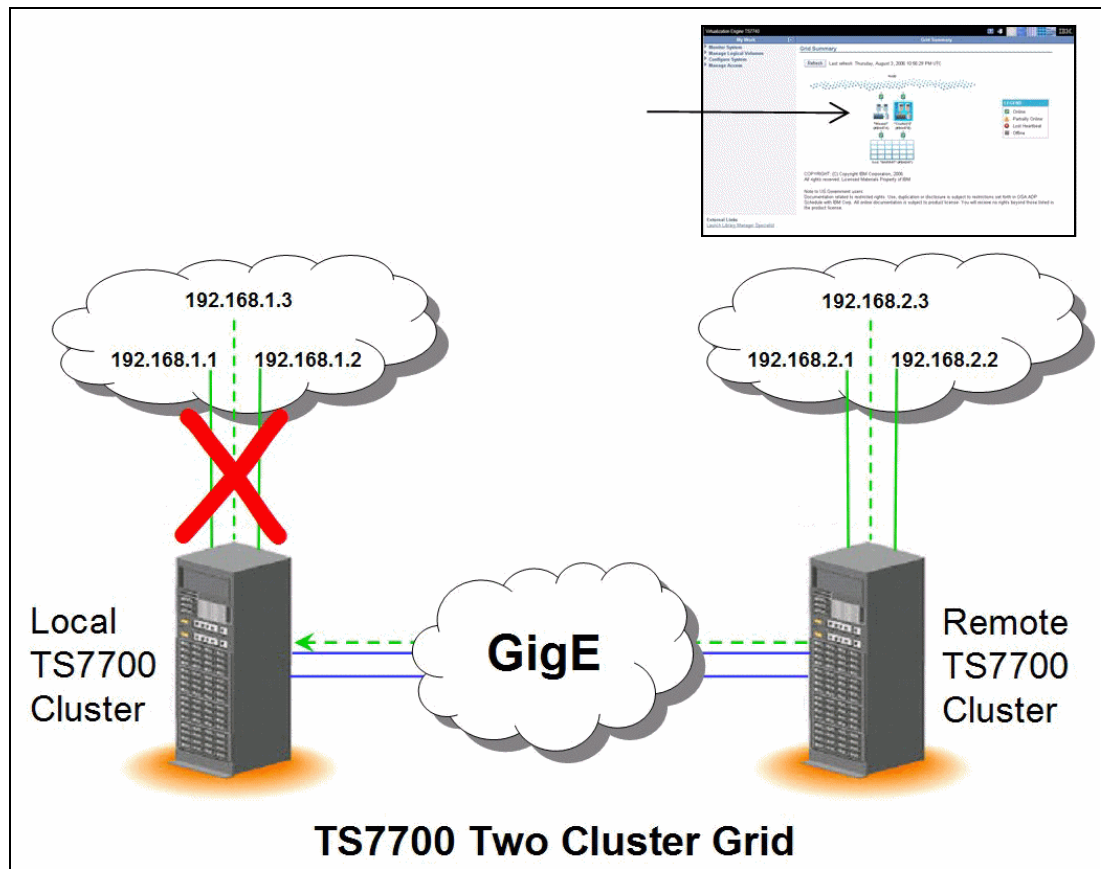


Figure 3-5 TS7700 Virtualization Engine MI access from a remote cluster

IP addresses for cross-site replication within a multi-cluster grid

In a TS7700 Virtualization Engine multi-cluster grid environment, you need to supply two additional IP addresses and two additional GigE Ethernet network connections per cluster for the physical links required by the TS7700 Virtualization Engine Grid for cross-site replication.

The physical Gb Ethernet links should be configured with IP addresses acting in separate and independent logical subnets for proper link failure and error detection. In other words, each

GigE Ethernet link operates with two TCP/IP addresses and a subnet mask different from the second GigE Ethernet link.

Important: Configure each GigE Ethernet link between the TS7700 Virtualization Engine Clusters as a separate subnet. For example, in a Two-Cluster configuration, use:

- ▶ 172.1.1.1 and 172.1.2.1 for the first TS7700 Virtualization Engine's GigE Ethernet links
- ▶ 172.1.1.2 and 172.1.2.2 for the second TS7700 Virtualization Engine's GigE Ethernet links
- ▶ 255.255.255.0 as the subnet mask for both

See Figure 3-6 for a sample configuration of a two-cluster grid.

Note: A three-cluster grid simply adds another cluster and uses IP addresses 172.1.1.3 and 172.1.2.3. A four-cluster grid can use 172.1.1.4 and 172.1.2.4 addresses.

These IP addresses are simply suggestions for illustration purposes. Work with your network management team to determine the appropriate IP addresses.




Figure 3-6 TS7700 Virtualization Engine GigE Ethernet link subnets

IP addresses for the TSSC

The TS7700 Virtualization Engine reserves a block of ten IP addresses on the TSSC network. Seven of these IP addresses are currently used. Three are reserved for future use.

The seven IP addresses are as follows:

- ▶ Two IP addresses for the routers
- ▶ Two IP addresses for the IBM 3956 Controller
- ▶ One IP address for the IBM Virtualization Engine 3957
- ▶ One Advanced Systems Management Interface (ASMI) IP address
- ▶ One IP address for the management interface

Note: The IP addresses that are used for the TSSC network must be entered as an increment of 10 in the range of 10 - 240. The router configuration uses this address and the next nine sequential addresses for TSSC configuration. To prevent network problems, do not configure these addresses on this TSSC for another system.

3.2.4 LAN/WAN requirements

The TS7700 Virtualization Engine uses your LAN/WAN to replicate logical volumes, access logical volumes remotely, and perform cross-site messaging. The LAN/WAN should have

adequate bandwidth to deliver the throughput necessary for customer requirements. The cross-site grid network can be shortwave optical fiber or 1 Gb copper Ethernet, which may be Cat 5E or 6 cabling. Cat 6 cabling will achieve the highest throughput when using copper connections. The TS7700 Virtualization Engine does not encrypt the data it sends over the LAN/WAN.

Cross-site dedication

Your LAN/WAN must be dedicated to cross-site grid communication and not shared with other intranet traffic or high bandwidth applications. The TS7700 Virtualization Engine connection to the cross-site Grid network can be shortwave optical fiber or 1 Gb copper Ethernet. With copper connections, use Cat 6 cabling.

The system uses all available bandwidth up to 1 Gb per link for its use. Sharing the WAN infrastructure with other high bandwidth applications might cause contention and adversely affect TS7700 Virtualization Engine performance.

Network protection

Consider the security of the data exchanged between the sites as part of the WAN topology and implementation, for example, as follows:

- ▶ Isolating your LAN/WAN from the general employee network
- ▶ Encrypting or otherwise securing data in flight

Bandwidth requirements

Your LAN/WAN must be able to support the bandwidth requirements of your Grid. These requirements are determined by the amount of data requiring replication and the expected time in which replication should be accomplished. The recommended bandwidth capacity is 1 Gbps per redundant paths. Work with your sales team to determine a recommended minimum bandwidth based on your requirements.

Network redundancy

The TS7700 Virtualization Engine provides two independent 1 Gb Ethernet connections. Connect each link through an independent WAN connection for redundancy to avoid a single point of failure that disrupts service to both cluster WAN paths.

3.3 Remote installations and FICON switch support

The TS7700 Virtualization Engine 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 TS7700 Virtualization Engine:

- ▶ 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 TS7700 Virtualization Engine. See Figure 3-7 on page 134 for more details about the distances supported.

3.3.1 Factors that affect performance at a distance

Fibre Channel distances depend on many factors:

- ▶ Type of laser used: Longwave or shortwave
- ▶ Type of fiber optic cable: Multi-mode or single-mode
- ▶ Quality of the cabling infrastructure in terms of decibel (dB) signal loss:
 - Connectors
 - Cables
 - Bends and loops in the cable

Native shortwave Fibre Channel 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 Fibre Channel 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 Fibre Channel nodes at either end. Buffer credits are designed into the hardware for each Fibre Channel port. Fibre Channel provides flow control that protects against collisions.

This configuration is extremely important for storage devices, which do not handle dropped or out-of-sequence records. When two Fibre Channel ports begin a conversation, they exchange information about their number of supported buffer credits. A Fibre Channel port will send only the number of buffer frames for which the receiving port has given credit. This approach both avoids overruns and 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 Fibre Channel node that is attached at either end of the link extenders.

This relationship is 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 is able to read at full performance over a greater distance than it is able to write. This is because, on the writes, the HBA can send a maximum of only eight buffers to the switch port; on the reads, the switch can send up to 64 buffers to the HBA. Until recently, a rule of thumb has been to allocate one buffer credit for every 2 km to maintain full performance.

Buffer credits within 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 distance covered. See *Introduction to SAN Distance Solutions*, SG24-6408 for more information.

3.3.2 Host attachments

The TS7700 Virtualization Engine attaches to System z hosts through either FICON Longwave or Shortwave channels. ESCON channel attachment is not supported. FICON channel extension and DWDM transport connections are supported.

Supported distances

When directly attaching to the host, the TS7700 Virtualization Engine 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 TS7700 Virtualization Engine can be installed at extended distances from the host.

Supported FICON extended distances are summarized in Figure 3-7.

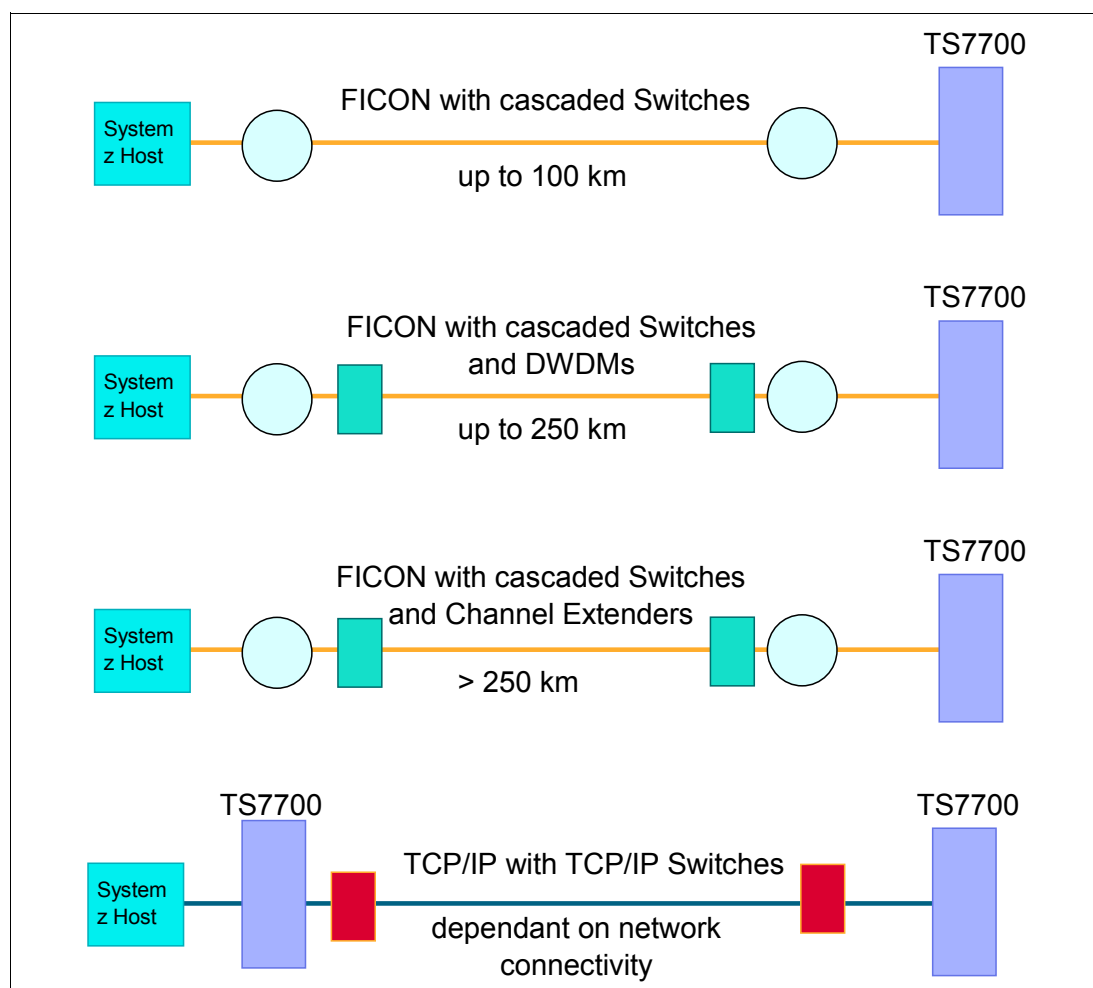


Figure 3-7 TS7700 Virtualization Engine Extended Distance Support

In a multi-cluster grid configuration, the TS7700 Virtualization Engines are also connected through TCP/IP connections. These network connections are not as sensitive as FICON to long distances when sufficient bandwidth is available.

Figure 3-8 on page 135 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 (see 3.3.3, “FICON Director support” on page 135).

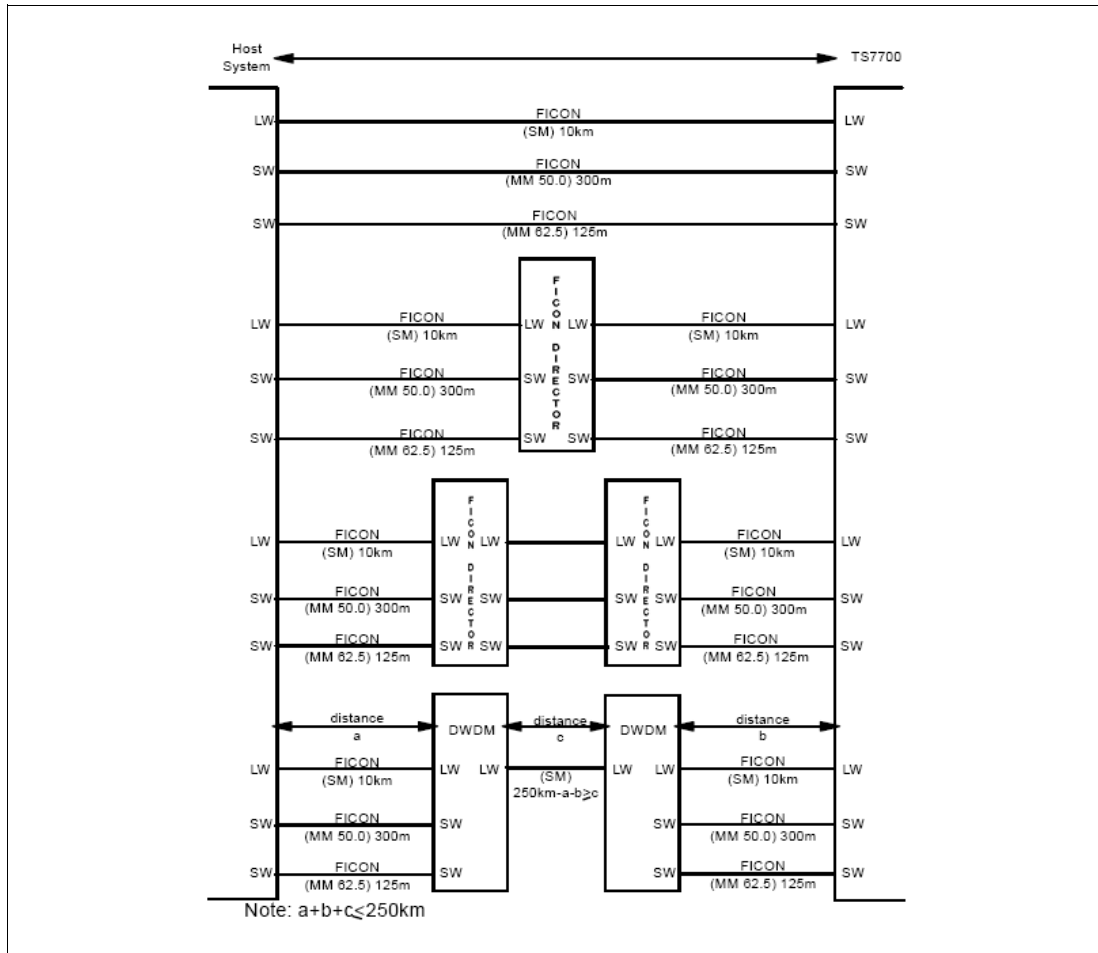


Figure 3-8 System z host attachments distances

3.3.3 FICON Director support

All FICON Directors are supported for Single and multi-cluster grid configurations with 1 Gbps, 2 Gbps, or 4 Gbps links. The components will auto-negotiate to the highest speed allowed.

You cannot mix different vendors (Brocade (McData (CNT & Inrange)) and CISCO), but you may mix models of one vendor. See the System Storage Interoperation Center (SSIC) for specific intermix combinations supported. You can find the SSIC at the following address:

http://www.ibm.com/systems/support/storage/config/ssic/displayesssearchwithoutjs.wss?start_over=yes

The FICON switch support matrix is at the following address:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/FQ116133>

3.3.4 FICON channel extenders

FICON channel extenders can operate in one of the following modes:

- ▶ Frame shuttle or tunnel mode
- ▶ Emulation mode

Using the frame shuttle or tunnel mode, the extender receives and forwards FICON frames without doing any special channel or control unit emulation 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.

The channel extenders listed in Table 3-9 are supported between the System z host and the TS7700 Virtualization Engine.

Table 3-9 FICON channel extender support

Channel extender vendor and model	Channel extender level of code supported	TS7740 Virtualization Engine level of microcode supported
Ciena CN2000	V4.1.2 ^a	8.0.0.x - 8.4.1.x
Brocade (McDATA) Edge	3.1.5 or V3.2.0.4.3 or V3.2.0.5	8.0.0.x - 8.4.1.x
Brocade (McDATA) USD and USD-x	3.4.2.0.6 or V4.1.0.7.3 or V4.1.0.11	8.0.0.x - 8.4.1.x

a. For optimal performance, flow control must be enabled.

3.3.5 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. Although the example uses a different ID and address for clarity, keep them the same to reduce configuration confusion and simplify problem determination work.

Allowable Director Switch ID ranges have been established by the manufacturer:

- McDATA range: x'61' to x'7F'
- CNT/Inrange range: x'01' to x'EF'
- Brocade 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, keep them the same to reduce configuration confusion and simplify PD work.

The following allowable Director Address ranges are established by the manufacturer:

- McDATA range: x'61' to x'7F'
- CNT/Inrange range: x'01' to x'EF'
- Brocade range: x'01' to x'EF'

- ▶ Director Ports

The Port Address may 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 indicates 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.

3.4 High availability Grid considerations

The TS7700 Virtualization Engine Grid provides configuration flexibility to meet a variety of customer needs. Those needs are both customer and application dependent. Here we specifically address planning a two-cluster grid configuration to meet high availability needs. The discussion easily translates to the two production clusters of a high availability, disaster recovery, Three-cluster grid configuration where the third cluster is strictly a disaster recovery site.

High availability means being able to provide continuous access to logical volumes through planned and unplanned outages with as little customer impact or intervention as possible. It does not mean that all potential for customer impact or action is eliminated. The basic guidelines for establishing a Grid configuration for high availability are as follows:

- ▶ The production systems (sysplexes and LPARs) have FICON channel connectivity to both clusters in the Grid. This means that DFSMS library definitions and IODF have been established and the appropriate FICON Directors, DWDM attachments, and fiber are in place. Virtual tape devices in both clusters in the Grid configuration are varied online to the production systems. If virtual tape device addresses are not normally varied on to both clusters, the virtual tape devices to the standby cluster need to be varied on in the event of a planned or unplanned outage to allow production to continue.
- ▶ The workload placed on the Grid configuration should be such that when using only one of the clusters, performance throughput is sufficient to meet customer service level agreements. If both clusters are normally used by the production systems (the virtual devices in both clusters are varied online to production), then in the case where one of the clusters is unavailable, the available performance capacity of the Grid configuration is reduced to half.
- ▶ For all data that is critical for high availability, assign a Management Class whose copy consistency point definition has both clusters having a copy consistency point of RUN. This means that each cluster is to have a copy of the data when the volume is closed and unloaded from the source cluster.
- ▶ The distance of grid links between the clusters should be no more than 50-100 km using low-latency directors, switches, or DWDMs. This minimizes the impact to job execution times because of volume copying between clusters at volume close time. Network Quality of Service (QoS) or other network sharing methods should be avoided, because it can

introduce packet loss, which directly reduces the effective replication bandwidth between the clusters.

- ▶ For data that you want to have two copies of, but the copy operation can be deferred, the copy consistency points for the two production clusters that have the online virtual devices should be set to deferred (DD) for both clusters. This, in concert with the *Prefer Local Cluster for Fast Ready Mounts* override, will provide the best performance.
- ▶ To prevent remote tape volume cache accesses during scratch mounts, the *Prefer Local Cluster for Fast Ready Mounts* copy override setting must be configured on both TS7700 Virtualization Engine Clusters within the Grid.
- ▶ To improve performance and take advantage of cached versions of logical volumes, the *Prefer Local Cluster for Non-Fast Ready Mounts* and *Force Local Copy override* settings must not be configured in either cluster.
- ▶ To minimize operator actions when a failure has occurred in one of the clusters that makes it unavailable, the Autonomic Ownership Takeover Manager (AOTM) must be set up to automatically place the remaining cluster in at least the read ownership takeover mode. Use read/write ownership takeover mode if you want to write data to the remaining cluster. If AOTM is not used, or it cannot positively determine if a cluster has failed, an operator must make that determination and, through the management interface on the remaining cluster, manually select one of the ownership takeover modes.
- ▶ If multiple Grid configuration is available for use by the production systems, on detection of a failure of a cluster in one of the Grid configurations, the state of the Storage Group (or groups) that is associated with that Grid configuration: The composite library must be changed to disallow scratch mounts. By doing so, all future write workloads are directed to the other fully operational Grid configurations. If this is the case, using AOTM with read ownership takeover is preferred and the impact to job performance might be minimized.

By following these guidelines, the TS7700 Virtualization Engine Grid configuration supports the availability and performance goals of customer workloads by minimizing the impact of the following outages:

- ▶ Planned outages in a Grid configuration, such as microcode or hardware updates to a cluster. While one cluster is being serviced, production work continues with the other cluster in the Grid configuration after virtual tape device addresses are online to the cluster.
- ▶ Unplanned outage of a cluster. Because of the copy policy to make a copy of a volume when it is being closed, all jobs that completed prior to the outage will have a copy of their data available on the other cluster. For jobs that were in progress on the cluster that failed, they can be re-issued after virtual tape device addresses are online on the other cluster (if they were not already online) and an ownership takeover mode has been established (either manually or through AOTM). For jobs that were writing data, the written data is not accessible and the job must start again.

Note: Fast Ready Categories and Data Classes work at the System level and are unique for all clusters in a Grid. This means that if you modify them in one cluster, it will apply to all clusters on that Grid.

3.5 Planning for software implementation

In this section, we provide information for planning tasks related to host configuration and software requirements for use with the TS7700 Virtualization Engine.

3.5.1 Host configuration definition

You must define the hardware to the host using the hardware configuration definition (HCD) facility. Specify the LIBRARY-ID and LIBPORT-ID in your HCD/IOCP definitions even in a single cluster configuration.

LIBRARY-ID

In a Grid configuration used with the TS7700 Virtualization Engine, each virtual device attached to a System z host reports back the same five character hexadecimal library sequence number, known as the *Composite Library ID*. With the Composite Library ID, the host considers the Grid configuration as a single library.

Each cluster in a Grid configuration possesses a unique five character hexadecimal library sequence number, known as the *Distributed library ID*, which identifies it among the clusters in its Grid. This distributed library ID is reported to the System z host upon request and is used to distinguish one cluster from another in a given Grid.

LIBPORT-ID

Each logical control unit, or 16-device group, must present a unique subsystem identification to the System z host. This ID is a 1-byte field that uniquely identifies each logical control unit within the cluster and is called the LIBPORT-ID. The value of this ID cannot be 0. Table 3-10 shows the definitions of the LIBPORT-IDs in a multi-cluster grid with 1024 virtual devices.

For more details, see 5.1.4, “Library names, Library IDs, and port IDs” on page 266.

Table 3-10 Subsystem identification definitions

Cluster	Logical CU (hex)	LIBPORT-ID (hex)
0	0-F	01-10
1	0-F	41-50
2	0-F	81-90
3	0-F	C1-D0

Virtual tape drives

The TS7700 Virtualization Engine presents a tape drive image of a 3490 C2A, identical to the current VTS and Peer-to-Peer (PTP) subsystems. Command sets, responses to inquiries, and accepted parameters match the defined functional specifications of a 3490E drive.

Each TS7700 Virtualization Engine Cluster provides 256 virtual devices. When two clusters are connected to a Grid configuration, the Grid provides 512 virtual devices. When three clusters are connected to a Grid configuration, the Grid provides up to 768 virtual devices. A four-cluster grid configuration can provide 1024 virtual devices.

3.5.2 Software requirements

The TS7700 Virtualization Engine is supported with the following operating system levels:

- ▶ z/OS V1R10 and later (lower release level support considered through RPQ process)
- ▶ With z/OS, the TS7700 Virtualization Engine is, for the most part, transparent to host software. However, to exploit new functionality in the TS7700 Virtualization Engine, corresponding APARs/PTFs might be needed. In general, refer to the existing VTS and PTP PSP buckets for the latest information about software maintenance for the TS7700 Virtualization Engine.

The software maintenance for the existing virtualization products also applies to the TS7700 Virtualization Engine. In addition to the PSP buckets, also refer to the new 3957 PSP bucket for APARs/PTFs that are specific to the TS7700 Virtualization Engine. Often a set of APARs/PTFs correspond to each release level of the TS7700 Virtualization Engine and in general, be sure to install any host support noted for the TS7700 Virtualization Engine.

- ▶ z/VM V5.4.0 and later

With z/VM, the TS7700 Virtualization Engine is transparent to host software. z/VM V5R1, or later, is required for both guest and native VM support. DFSMS/VM Function Level 221 is required for native VM tape library support and EREP V3.5 plus PTFs.

- ▶ z/VSE V4.1

With z/VSE, the TS7700 Virtualization Engine is transparent to host software. Version 4.1 or later is required.

The z/VSE does not support TS7700 multi-cluster grid and is unable to handle unsolicited messages and actions (for example, intervention required and hardware messages). Therefore, a Peer-to-Peer VTS requires an active z/OS to support and VSE. This requirement has been eliminated with the advent of the TS7700 Virtualization Engine and its management interface.

- ▶ TPF V4.1 and z/TPF V1.1 and later

The category *reserve and release* functions are not supported by the TS7740 in a grid environment or the TS7720 in a grid or stand-alone environment.

3.5.3 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 Integrated 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. The assigned Storage Group determines which logical partitions in the tape library are

to be used. Through the Storage Group construct, you can direct logical volumes to specific physical volume pools.

It is not enough to define the new SMS classes in z/OS; you must define the new SMS classes on the TS7700 Virtualization Engine through the management interface (MI). Otherwise, the constructs will be created on the TS7700 Virtualization Engine using the default construct parameters and displayed as a line, as shown in Figure 3-9.

Select	Name	Second	Description	Retain Copy M	"Archie[0]" (#BA96A)	"Veronica[1]" (#BA96B)
<input type="checkbox"/>	-----	0	The default Management Class	No	Deferred	Deferred
<input type="checkbox"/>		0	The default Management Class	No	Rewind Unload (RUN)	Rewind Unload (RUN)

Page 1 of 1 Rows: 10 Show Total: 2 Filtered: 2 Displayed: 2

Figure 3-9 Default construct

3.5.4 Sharing and partitioning considerations

If you plan to connect different sysplexes and even DR sites, it is better to decide in advance which scratch categories to assign to which sysplex. Table 3-11 shows an example of scratch categories assignment in a multi-sysplex center.

Table 3-11 Scratch categories assignment

PROD1	PROD2	PRE-PROD	TEST	DR site
0002	0012	0022	0032	0042

Partitioning between multiple hosts

The virtual drives and virtual volumes in a TS7700 Virtualization Engine can be partitioned just like physical drives and real volumes. Any virtual volume can go to any physical stacked volume. The TS7700 Virtualization Engine places no restrictions on the use and management of those resources. You have the ability to partition your stacked media in up to 32 separate pools. This is achieved by assigning a Storage Group to a defined range of stacked volumes prior to insertion time.

Sharing a TS7700 Virtualization Engine

A FICON-attached TS7700 Virtualization Engine supports two or four 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 TS7700 Virtualization Engine.

Use a FICON Director when connecting the TS7700 Virtualization Engine to more than one system.

The TS7700 Virtualization Engine places no limitations on the number of hosts that can use those channel paths, or the types of hosts, or their operating system environments. This is the same as for any tape technologies that are supported in IBM Tape Libraries.

An operating environment, however, through its implementation, does impose limits. z/OS DFSMS can support up to 32 systems or groups of systems.

Basically, anything that can be done with native drives in an IBM Tape Library can be done with the virtual drives in a TS7700 Virtualization Engine.

The TS7700 Virtualization Engine attaches to the host system or systems through two or four FICON channels. Each FICON channel provides 256 logical paths (starting with TS7700 Virtualization Engine Release 1.4). A four-FICON configuration results in a total of 1024 logical paths for each TS7700 Virtualization Engine.

For more information, see 5.1.3, “Logical path considerations” on page 265.

Note: Fast Ready Categories and Data Classes work at the system level and are unique for all clusters in a Grid. Therefore, if you modify them in one cluster, the modification will apply to all clusters on that Grid.

3.6 Planning for logical and physical volumes

Before you define logical volumes to the TS7700 Virtualization Engine, consider the total number of logical volumes required, the volume serial ranges to define, and the number of volumes within each range.

The VOLSERS must be unique throughout an SMSplex and throughout all storage hierarchies, such as DASD, tape, and optical storage media.

The VOLSERS for logical volumes and physical volumes must be unique.

3.6.1 Logical volumes

Determine the number of logical volumes that are required to handle the workload you are planning for the TS7700 Virtualization Engine. With Release 1.6 and later, the TS7700 Virtualization Engine provides support for logical WORM volumes.

Consider the size of your logical volumes, the number of scratch volumes you will need per day, the time that is required for return-to-scratch processing, how often scratch processing is performed, and whether you need to define logical WORM volumes.

Size of logical volumes

The TS7700 Virtualization Engine supports logical volumes with maximum sizes of 400, 800, 1000, 2000, 4000, and 6000 MB although effective sizes can be larger if data is compressed. For example, if your data compresses with a 3:1 ratio, the effective maximum logical volume size for a 4,000 MB logical volume is 12,000 MB.

Depending on the logical volume sizes that you choose, you might see the number of volumes required to store your data grow or shrink depending on the media size from which you are converting. If you have data sets that fill native 3590 volumes, even with 4,000 MB logical volumes, you will need more TS7700 Virtualization Engine logical volumes to store the data, which will be stored as multi-volume data sets.

The 400 MB CST-emulated cartridges or 800 MB with ECCST-emulated cartridges are the two types you can specify when adding volumes to the TS7700 Virtualization Engine. You can use these sizes directly or use policy management to override them to provide for the 1,000, 2,000, or 4,000 MB sizes.

A logical volume size may be set by VOLSER and can change dynamically using the DFSMS Data Class storage construct. The amount of data copied to the stacked cartridge is only the amount of data that has been written to a logical volume. The choice between all available

logical volume sizes does not affect the real space used in either the TS7700 Virtualization Engine cache or the stacked volume.

In general, unless you have a special need for CST emulation (400 MB), specify the ECCST media type when inserting volumes in the TS7700 Virtualization Engine.

In planning for the number of logical volumes needed, first determine the number of private volumes that make up the current workload you will be migrating. One way to do this is by looking at the amount of data on your current volumes and then matching that to the supported logical volume sizes. Match the volume sizes, taking into account the compressibility of your data. If you do not know the average ratio, use a conservative value of 2:1, for example.

If you choose to only use the 800 MB volume size, the total number needed increases because some of the current volumes that contain more than 800 MB compressed need to expand to a multi-volume set. Taking that into account for the example would increase the number of logical volumes required by 22,180 (for each 1,000 MB volume called out in the table, one additional 800 MB volume is needed for spanning, for each 2,000 MB volume, two additional volumes are needed, and for each 4,000 MB volume, four additional volumes are needed) for a total of 72,953.

Now that you know the number of volumes you need for your current data, you can estimate the number of empty scratch logical volumes you have to add. Based on your current operations, determine a nominal number of scratch volumes from your nightly use. If you have an existing VTS installed, you might have already determined this number and are able to set a scratch media threshold with that value through the ISMF Library Define window.

Next, multiply that number by the value that provides a sufficient buffer (typically 2×) and by the frequency with which you want to perform returns to scratch processing.

A suggested formula to calculate the number of logical volumes needed is as follows:

$$V_v = C_v + Tr + (Sc)(Si + 1)$$

The formula contains the following values:

Vv	Total number of logical volumes needed
Cv	Number of logical volumes needed for current data rounded up to the nearest 10,000
Tr	Threshold value from the ISMF Library Define window for the scratch threshold for the media type used (normally MEDIA2); usually set to equal the number of scratch volumes used per night
Sc	Number of empty scratch volumes used per night, rounded up to the nearest 500
Si	Number of days between scratch processing (return-to-scratch) by the tape management system

For example, assuming the current volume requirements (using all the available volume sizes), and you use of 2,500 scratch volumes per night and performance of return-to-scratch processing every other day, you need to plan on the following number of logical volumes in the TS7700 Virtualization Engine:

$$75,000 \text{ (current, rounded up)} + 2,500 + (2,500)(1+1) = 82,500 \text{ logical volumes}$$

If you define more volumes than you need, you can always delete the additional volumes. Unused logical volumes do not consume space.

The TS7700 Virtualization Engine supports up to 1,000,000 logical volumes. This is the maximum number either in a stand-alone or in a Grid configuration.

Consideration: Up to 10,000 logical volumes may be inserted at one time. This applies to both inserting a range of logical volumes and inserting a quantity of logical volumes. Attempting to insert amounts over 10,000 will return an error.

Number of scratch volumes needed per day

As you run your daily production workload, you will need enough logical volumes in scratch status to support the data that will be written to the TS7700 Virtualization Engine. This can be hundreds or thousands of volumes, depending on your workload. You will likely want more than one day's worth of scratch volumes available at any point in time.

Return-to-scratch processing

Return-to-scratch processing involves running a set of tape management tools that identify the logical volumes that no longer contain active data, and then communicating with the TS7700 Virtualization Engine to change the status of those volumes from private to scratch. The amount of time the process takes depends on the type of tape management system being employed, and how busy the TS7700 Virtualization Engine is when it is processing the volume status change requests and whether a Grid configuration is being used.

You can see elongated elapsed time in any Tape Management Systems return-to-scratch process when you migrate to or install a multi-cluster configuration solution.

If the number of logical volumes used on a daily basis is small (less than a few thousand), you might choose to only perform return-to-scratch processing every few days. A good rule is to plan for no more than a 4-hour time period to run return to scratch. By ensuring a nominal run time of four hours, enough time exists during first shift to run the process twice should problems be encountered during the first attempt. Unless there are specific reasons, execute return-to-scratch processing once per day.

With z/OS V1.9 or later, return to scratch in DFSMSrmm has been enhanced to speed up this process. To reduce time required for housekeeping, it is now possible to run several return-to-scratch processes in parallel. For additional information about the enhanced return-to-scratch process, see the *z/OS DFSMSrmm Implementation and Customization Guide*, SC26-7405.

Volume serial numbering

When you insert volumes, you do that by providing starting and ending volume serial number range values.

The TS7700 Virtualization Engine determines how to establish increments of VOLSER values based on whether the character in a particular position is a number or a letter. For example, inserts starting with ABC000 and ending with ABC999 will add logical volumes with VOLSERs of ABC000, ABC001, ABC002...ABC998, and ABC999 into the inventory of the TS7700 Virtualization Engine. You might find it helpful to plan for growth by reserving multiple ranges for each TS7700 Virtualization Engine you expect to install.

If you have multiple partitions, it is better to plan in advance which ranges will be used in which partitions, for example, A* for first sysplex, B* for second sysplex, and so on. If you need more than a range, you can select A* and B* for first sysplex, C* and D* for second sysplex, and so on.

3.6.2 Logical WORM (LWORM)

TS7700 Virtualization Engine, Release 1.6 and later, supports the logical Write Once, Read Many (WORM) function through TS7700 Virtualization Engine software emulation. The z/OS host software supports physical WORM media prior to z/OS Version 1. The logical WORM enhancement can duplicate most of the 3592 WORM behavior. The host views the TS7700 Virtualization Engine as an LWORM-compliant library that contains WORM-compliant 3490E logical drives and media.

3.6.3 Physical volumes for TS7740 Virtualization Engine

Determine the number of physical volumes that are required to accommodate the workload you are planning for the TS7700 Virtualization Engine.

Consider the following information:

- ▶ The number of logical volumes you define
- ▶ The average amount of data on a volume
- ▶ The average compression ratio achieved for the data
- ▶ If the Selective Dual Copy function is to be used
- ▶ Whether the Delete Expired Volume Data setting is to be used
- ▶ Whether the Copy Export function is to be used
- ▶ The reclaim threshold settings, scratch physical volumes, and the number of physical volume pools

Ordering new tape cartridges

Order new physical cartridges with consideration for the complete fulfilment cycle. New cartridges can be ordered prelabelled with a specified bar code range representing the volume ID (VOLID).

Reuse of existing cartridges

Reuse of existing 3592 cartridges is permitted. However, the cartridges must be re-initialized to ensure that the external bar code label matches the internal VOLID. If the external label does not match the internal label, the cartridge is rejected.

Number of logical volumes

Because the data on a logical volume is stored on a physical volume, the number of logical volumes has a direct effect on the number of physical volumes required.

Average amount of data on a volume

The TS7700 Virtualization Engine only stores the amount of data you write to a logical volume plus a small amount of metadata.

Average compression ratio achieved for the data

The data that a host writes to a logical volume might be compressible. The space required on a physical volume is determined after the results of compression. If you do not know the average compression ratio for your data, assume a conservative 2:1 ratio.

Selective Dual Copy function

If you use this function for some or all of your data, a second physical copy of the data is written to a physical volume.

For critical data that only resides on tape, you may currently make two copies of the data on physically separate tape volumes either manually through additional job steps or through their applications. Within a TS7700 Virtualization Engine, you might need to be able to control where a second copy of your data is placed so that it is not stacked on the same physical tape as the primary copy. Although this can be accomplished through the logical volume affinity functions, it simplifies the management of the tape data and make better use of the host CPU resources to have a single command to the TS7700 Virtualization Engine subsystem direct it to selectively make two copies of the data contained on a logical volume.

If you activate Dual Copy for a group of data or a specific pool, consider that all tasks and properties, which are connected to this pool, are duplicated:

- ▶ The number of reclamation tasks
- ▶ The number of physical drives used
- ▶ The number of cartridges used
- ▶ The number of writes to the cartridges: one from cache to the primary pool and another to the backup pool

You have to plan for additional throughput and capacity. You do not need more logical volumes because the second copy uses an internal volume ID.

Number of Copy Export volumes

If you are planning to use the Copy Export function, plan for a sufficient number of physical volumes for the Copy Export function and sufficient storage cells for the volumes in the library destined for Copy Export or in the Copy Export state. The Copy Export function defaults to a maximum of 2000 physical volumes to be in the Copy Export state. This amount includes off-site volumes, the volumes still in the physical library that are in the Copy Export state, and the empty, filling, and full physical volumes that will eventually be set to the Copy Export state. With R1.6 and later the default value can be adjusted through the management interface to a maximum value of 10000. After your Copy Export operations reach a steady state, approximately the same number of physical volumes are being returned to the library for reclamation as there are those being sent off-site as new members of the Copy Export set of volumes.

Delete Expired Volume Data setting

If the Delete Expired Volume Data setting on the management interface is not used, logical volumes occupy space on the physical volumes even after they have been returned to scratch. In that case, only when a logical volume is rewritten is the old data released to reduce the amount of active data on the physical volume.

With the Delete Expired Volume Data setting, the data associated with volumes that have been returned to scratch are deleted after a time period and their old data is released. For example, assume that you have 20,000 logical volumes in scratch status at any point in time and that the average amount of data on a logical volume is 400 MB and that the data compresses at a 2:1 ratio. The space occupied by the data on those scratch volumes is 4,000,000 MBs or the equivalent of 14 JA cartridges (when using J1A Emulation mode). By using the Delete Expired Volume Data setting, you can reduce the number of cartridges required in this example by 14.

Reclaim threshold settings

The number of physical volumes also depends on the Reclaim Threshold Percentage that you have specified for the minimum amount of active data on a volume before it becomes eligible for reclamation. The default is set to 10%. The reclaim threshold setting can have a large impact on the number of physical volumes required. Physical volumes will hold between the threshold value and 100% of data. On average, the percentage of active data on the physical volume is $(100\%+10\%)/2$ or 55% (assuming a reclamation setting of 10%).

Having too low a setting results in more physical volumes being needed; having too high a setting might impact the ability of the TS7740 Virtualization Engine to perform host workload, because it is using its resources to perform reclamation. You might need to experiment to find a threshold that matches your needs.

We suggest, as a good starting point, to use 35%. We think that this will accommodate most installations.

Number of physical volume pools

You should plan for at least 10 scratch physical volumes to be available in the common scratch pool.

For each physical volume pool you are planning on using, you should have at least three scratch physical volumes. These are in addition to the number of physical volumes calculated to hold the data on the logical volumes.

The suggested formula to calculate the number of physical volumes needed is as follows:

$$Pv = (((Lv + Lc) \times Ls/Cr) / (Pc \times (((100+Rp)/100)/2))$$

The formula contains the following values:

Pv	Total number of physical volumes needed
Lv	Number of logical volumes defined
Lc	The number of logical volumes that will be dual copied
Ls	The average logical volume size in MB
Cr	Compression ratio of your data
Rp	The reclamation percentage
Pc	Capacity of a physical volume in MB

To this number, you then add scratch physical volumes based on the common media pool and the number of physical volume pools you plan on using. For example, use the following assumptions:

Lv	82,500 (see "Number of logical volumes" on page 145)
Lc	10,000 (logical volumes)
Ls	400 MB
Cr	2
Rp	10
Pc	300,000 MB (capacity of a 3592 J1A written JA volume) 500,000 MB (capacity of a TS1120 written JA volume) 700,000 MB (capacity of a TS1120 written JB volume)
Common scratch pool	10
Volume Pools	5 (with three volumes per pool)

Note: The calculated number is the required minimum value. It does not include any spare volumes to allow data growth from the first installation phase.

Therefore, add enough extra scratch volumes for future data growth.

Using the suggested formula and the assumptions, plan to use the following number of physical volumes in your TS7740 Virtualization Engine:

$$(((82,500 + 10,000) \times 400 \text{ MB})/2)/(300,000 \times (((100 + 10)/100)/2)) + 10 + 5 \times 3 = 137 \text{ physical volumes}$$

If you insert more physical volumes in the TS7740 Virtualization Engine than you need, you can eject them at a later time.

Pooling considerations

Pooling might have an effect on the throughput of your TS7740 Virtualization Engine. If you are using physical volume pooling at your site, consider the following possibilities:

- ▶ The possible increase of concurrent tape drive usage within the TS7740 Virtualization Engine

Depending on the number of pools and the amount of logical volume data being created per pool, you need to ensure that sufficient drives are available to handle:

- Premigration of logical volumes from the Tape Volume Cache
- Recall of logical volumes from physical stacked volumes to the Tape Volume Cache
- The amount of logical volume data being dual copied

- ▶ The reclamation process

Reclamation is done at the pool level and each reclamation task will use two drives. To minimize the effects of the reclamation process, ensure you maintain a sufficient amount of physical TS7740 Virtualization Engine scratch cartridges so that the reclamation process is performed within the reclaim scheduled time.

- ▶ An increase in the amount of cartridges being used
- ▶ Library slot capacity
- ▶ TS7740 Virtualization Engine processing/cache capacity

Out of scratch for physical stacked volumes

Monitor the number of empty stacked volumes in a library. If the library is close to running out of a physical volume media type, actions should be taken to either expedite the reclamation of physical stacked volumes or add additional ones. The Host Console Request function can be used to obtain the physical volume counts within the TS7740 Virtualization Engine. You can also use the Bulk Volume Information Retrieval (BVIR) function to obtain the physical media counts for each library. The information obtained includes the empty physical volume counts by media type for the common scratch pool and each defined pool.

See “Out of Physical Volumes” on page 573 for more information about how to handle an out of scratch situation. For more information about BVIR, see 9.10.5, “Interpreting the BVIR response data” on page 665. For more information about the Host Console Request function, see 2.5.6, “Logical WORM (LWORM) support and characteristics” on page 75.

3.6.4 Data compression

When writing data to a virtual volume, the host compression definition is honored. Compression is turned on or off by the JCL parameter DCB=TRTCH=COMP (or NOCOMP), the Data Class parameter COMPACTION=YES/NO, or the COMPACT=YES/NO definition in the DEVSUPxx PARMLIB member. The TRTCH parameter overrides the Data Class definition and both override the PARMLIB definition.

Important: To achieve the optimum throughput, verify your definitions to make sure that you have specified compression for data written to the TS7700 Virtualization Engine.

3.6.5 Secure Data Erase

Expired data on a physical volume remains readable until the volume has been completely overwritten with new data. Some customers are concerned that a court order could expose them to liability and cost if they have to find an old version of a data volume. Another concern is security of old data on physical media.

TS7700 Virtualization Engine Release 1.3 added physical volume erasure on a physical volume pool basis controlled by an additional reclamation policy. It utilizes the Outboard Policy Management (OPM) feature, which is standard on the TS7700 Virtualization Engine. With the Secure Data Erase function, all reclaimed physical volumes in that pool are erased by writing a random pattern across the whole tape prior to being reused. In the case of a physical volume that has encrypted data, the erase might involve just “shredding” the encryption keys on the volume to accomplish the erasure. A physical cartridge is not available as a scratch cartridge as long as its data is not erased.

The Secure Data Erase function is performed in one of two ways:

- ▶ For an unencrypted physical volume, or a physical volume that is encrypted, but its previous use was unencrypted, the erasure is performed by writing a random data pattern on the physical volume being reclaimed. In the case of the encrypted volume, the encrypted data keys are also erased.
- ▶ For an encrypted physical volume whose previous use was also encrypted, the physical volume is erased by erasing the encrypted data keys. By erasing the data keys, the data on the volume cannot be decrypted, rendering it “erased”. This is a much faster method because only the encrypted data keys need to be erased. This process takes only a few minutes compared to a multi-hour process for writing the random pattern across the whole tape.

Consider that if you choose this “erase” functionality, TS7740 Virtualization Engine needs a lot of time to reclaim every physical tape. Therefore, the TS7740 Virtualization Engine will need more time every day to do reclamation.

As part of this data erase function, an additional reclamation policy is added. The policy specifies the number of days a physical volume may contain invalid logical volume data before the physical volume becomes eligible to be reclaimed. The data associated with a logical volume is considered invalidated as follows:

- ▶ A host has assigned the logical volume to a scratch category. The volume is subsequently selected for a scratch mount and data is written to the volume. The older version of the volume is now invalid.
- ▶ A host has assigned the logical volume to a scratch category that has the fast-ready attribute set, the category has a non-zero delete expired data parameter value, the parameter value has been exceeded, and the TS7740 Virtualization Engine has deleted the logical volume.
- ▶ A host has modified the contents of the volume. This can be a complete rewrite of the volume or appending to it. The new version of the logical volume will be migrated to a different physical location, and the older version is now invalid.

The TS7740 Virtualization Engine keeps track of the amount of active data on a physical volume. It starts at 100% when a volume becomes full. Although the granularity of the percent

of full TS7740 Virtualization Engine tracks is 0.01%, it rounds down, so even one byte of inactive data will drop the percent to 99.9%.

TS7740 Virtualization Engine keeps track of the time that the physical volume went from 100% full to less than 100% full by doing the following tasks:

- ▶ Checking, on an hourly basis, for volumes in a pool that have a non-zero setting
- ▶ Comparing this time against the current time to determine whether the volume is eligible for reclamation

This data erase function is enabled on a per pool basis. It is enabled when a non-zero value is specified for the data erase reclamation policy. When enabled, all physical volumes in the pool are erased as part of the reclamation process, independent of the reclamation policy under which the volume became eligible for reclamation.

Any physical volume that has a status of read-only is not subject to this function and is not designated for erasure as part of read-only recovery.

If you use the eject stacked volume function, no attempt is made to erase the data on the volume prior to ejecting the cartridge. The control of expired data on an ejected volume becomes a customer responsibility.

Volumes that are tagged for erasure cannot be moved to another pool until erased, but they can be ejected from the library, because such volumes are usually removed for recovery actions.

Using the Move function of the Integrated Library Manager will also cause a physical volume to be erased, even though the number of days specified has not yet elapsed. This includes returning borrowed volumes.

The TS7740 Virtualization Engine historical statistics are updated with the number of physical mounts for data erasure. The pool statistics are updated with the number of volumes waiting to be erased and the value for the days (number of days) until erasure reclamation policy.

As soon as you decide to implement Secure Data Erase for a limited group of data separated on a dedicated pool, the number of additional reclamation tasks plus data erase tasks will increase. Less physical drives might be available even during times when you have inhibited reclamation.

The Inhibit Reclaim Schedule specification only partially applies to Secure Data Erase:

- ▶ No new cartridges are reclaimed during this time.
- ▶ Cartridges already reclaimed can be erased during this time.

This means that, although you do not allow reclamation during your peak hours to have all your drives available for recall and premigration, Secure Data Erase will not honor your settings and thus will run up to two concurrent erasure operations per physical drive type as long as there are physical volumes to be erased.

Because the first logical volume that expires triggers the physical volume to be erased, an almost full physical cartridge will be first reclaimed and then erased.

Group logical volumes that require secure erase after they are expired so that no unnecessary reclamation and subsequent erasure operations take place. Pooling by expiration date might help reduce unnecessary reclamation. Although proper grouping reduces the amount of reclamation that needs to be done, it will not eliminate the erasure step.

3.6.6 Copy Policy Override Settings

With the TS7700 Virtualization Engine, you can define and set the optional override settings that influence the selection of the I/O Tape Volume Cache (TVC) and replication responses. The settings are specific to a cluster in a multi-cluster grid configuration, meaning that each cluster can have different settings if you want. The settings take effect for any mount requests that are received after the settings were saved. Mounts already in progress are not affected by a change in the settings. You can define and set the following settings:

- ▶ Prefer local cache for fast ready mount requests
- ▶ Prefer local cache for non-fast ready mount requests
- ▶ Force volumes mounted on this cluster to be copied to the local cache
- ▶ Allow fewer RUN consistent copies before reporting RUN command complete

Use the available settings to meet specific performance and availability objectives. You can view and modify these settings from the TS7700 Virtualization Engine management interface by clicking **Configuration** → **Copy Policy Override**, as shown in Figure 3-10.

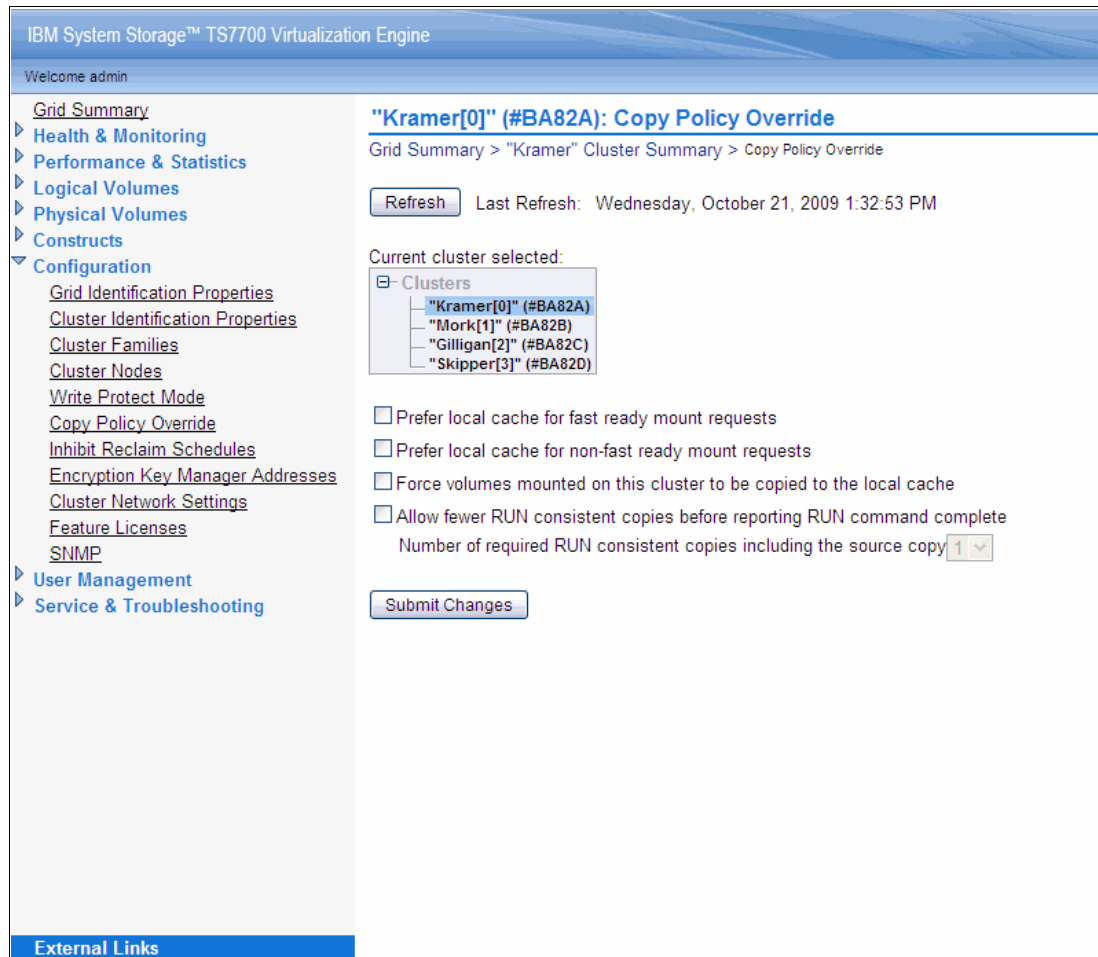


Figure 3-10 Copy Policy Override

The settings you can select in the MI window shown in Figure 3-10 on page 151 are:

- ▶ Prefer local cache for fast ready mount requests

A fast ready mount selects a local copy as long as a copy is available and a cluster copy consistency point is not specified as No Copy in the Management Class for the mount. The cluster is not required to have a valid copy of the data.
- ▶ Prefer local cache for non-fast ready mount requests

This override causes the local cluster to satisfy the mount request as long as the cluster is available and the cluster has a valid copy of the data, even if that data is only resident on physical tape. If the local cluster does not have a valid copy of the data, then default cluster selection criteria applies.
- ▶ Force volumes mounted on this cluster to be copied to the local cache

For a non-fast ready mount, this override causes a copy to be performed on the local cluster as part of mount processing. For a fast ready mount, this setting has the effect of overriding the specified Management Class with a copy consistency point of rewind/unload for the cluster. This setting does not change the definition of the Management Class, but serves to influence the replication policy.
- ▶ Allow fewer RUN consistent copies before reporting RUN command complete

If selected, the value entered for “Number of required RUN consistent copies including the source copy” is used to determine the number of copies to override before the Rewind/Unload (RUN) operation reports as complete. If this option is not selected, the Management Class definitions are to be used explicitly. Thus, the number of RUN copies can be from one to the number of clusters in the Grid.

Note: In a Geographically Dispersed Parallel Sysplex® (GDPS), all three Copy Policy Override settings must be selected on each cluster to ensure that wherever the GDPS primary site is, this TS7700 Virtualization Engine Cluster is preferred for all I/O operations.

If the TS7700 Virtualization Engine Cluster of the GDPS primary site fails, perform the following recovery actions:

1. Vary virtual devices from a remote TS7700 Virtualization Engine Cluster online from the primary site of the GDPS host.
2. Manually invoke, through the TS7700 Virtualization Engine management interface, a Read/Write Ownership Takeover, unless Automated Ownership Takeover Manager (AOTM) has already transferred ownership.

3.7 Planning for encryption in the TS7740 Virtualization Engine

The importance of data protection has become increasingly apparent with news reports of security breaches, loss and theft of personal and financial information, and government regulation. Encryption of backstore tapes helps control the risks of unauthorized data access without excessive security management burden or subsystem performance issues.

Encryption on the TS7740 Virtualization Engine is controlled on a storage pool basis. “Storage Group” and “Management Class” DFSMS constructs specified for logical tape volumes determine, through mapping in the Library Manager, which physical volume pools are used for the primary and backup (if used) copies of the logical volumes. The storage pools, originally created for management of physical media, have been enhanced to include encryption characteristics.

The encryption solution for tape virtualization consists of several components:

- ▶ The IBM tape encryption solutions all use an Encryption Key Manager (EKM) or the Tivoli Lifecycle Key Manager (TKLM) as a central point from which all encryption key information is managed and served to the various subsystems. The EKM or TKLM communicates with the TS7740 Virtualization Engine, and tape libraries, control units, and Open Systems device drivers.
- ▶ The TS1130 Tape Drive and the encryption-enabled TS1120 Tape Drive, as the other common component of the various data encryption solutions, provide hardware that performs the cryptography function without reducing the data-transfer rate.
- ▶ The TS7740 Virtualization Engine provides the means to manage the use of encryption and what keys are used on a storage-pool basis. It also acts as a proxy between the tape drives and the EKMs and TKLMs, using Ethernet to communicate with the EKMs or TKLMs and in-band Fibre Channel connections with the drives. Encryption support is enabled with Feature Code 9900.

As an extension of TS7740 Virtualization Engine encryption capabilities, TS7740 Virtualization Engine R1.7 supports the use of Default Keys. This support simplifies the management of the encryption infrastructure because no future changes are required at the TS7740 Virtualization Engine. After a pool is defined to use the Default Key, the management of encryption parameters is performed at the key manager.

TS7740 Virtualization Engine back-end drive encryption was introduced with Release 1.2 of the TS7740 Virtualization Engine. There are no host software updates required for this function, because the TS7740 Virtualization Engine controls all aspects of the encryption solution.

Tapes encrypted in the TS7740 Virtualization Engine backstore use a “wrapped key” model. The data on each cartridge is encrypted with a random 256-bit Advanced Encryption Standard (AES-256) Data Key (DK). The Data Key is stored on the cartridge in an encrypted, or “wrapped”, form. Four instances of these wrapped data keys are stored on each cartridge.

Although the feature for encryption support is customer-installable, some of the prerequisites might require additional hardware installation or configuration by an IBM System Service Representative.

3.7.1 Pool encryption settings for TS7740 Virtualization Engine

Encryption on the TS7740 Virtualization Engine is defined by storage pools. Storage Group DFSMS constructs that are specified for logical tape volumes determine which physical volume pools are used for the primary and optional backup copies of the logical volumes. The storage pools, originally created for management of physical media, have been enhanced to include encryption characteristics. Unique keys can be assigned to each physical volume pool in environments where the media is required to meet specific security policies. For centralized management of encryption, the TS7740 Virtualization Engine R1.7 and later support the use of default keys.

Use the window, shown in Figure 3-11, for viewing and modifying storage pool encryption settings on the TS7740 Virtualization Engine. The TS7740 Virtualization Engine allows encryption by storage pool. If you are planning to enable encryption for one or more pools, use this window to modify the encryption settings as follows:

1. Select the pool number you want to see or modify, then select **Encryption Settings** in the Select Action drop-down menu and click **Go**.

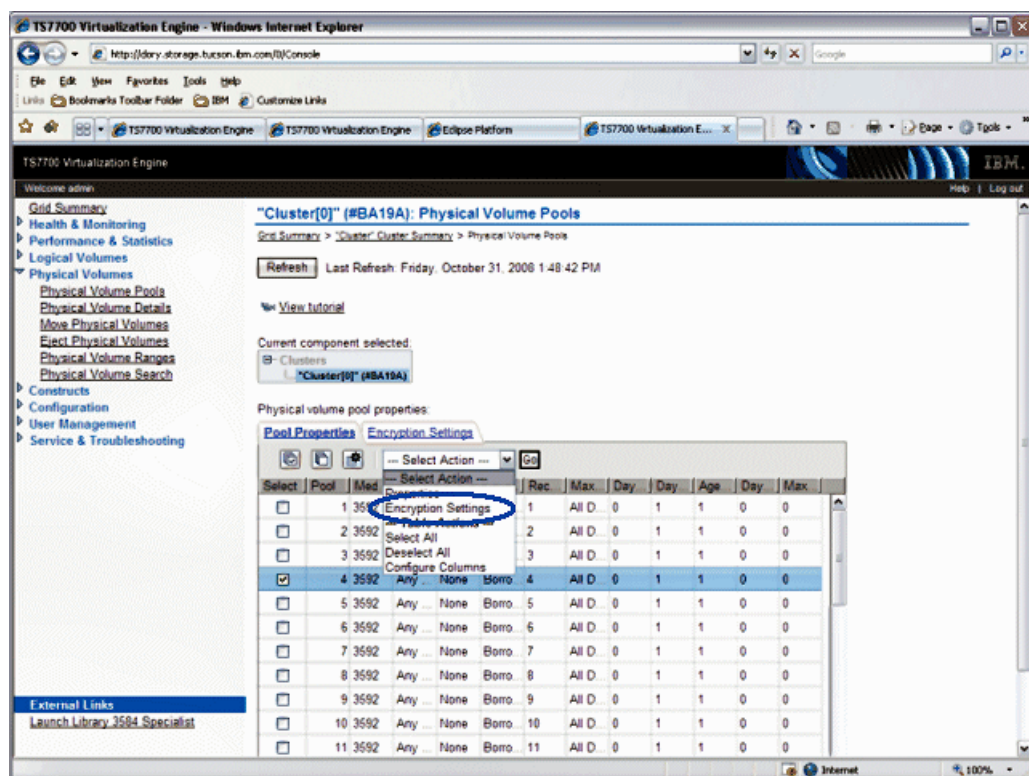


Figure 3-11 Modify encryption settings step 1

You receive the following message, as shown in Figure 3-12 on page 155:

HYDME0595W

Not all the drives are encryption-enabled. Any drives that are not encryption-enabled will become unavailable when encryption is enabled. These drives can be restored only by an IBM Service Representative. Do you wish to proceed?

2. Click **OK**.

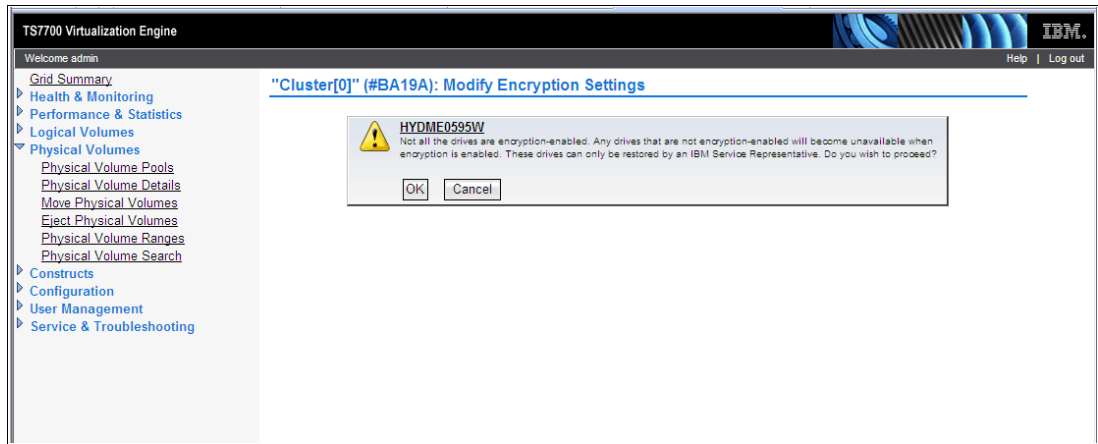


Figure 3-12 Modify encryption settings step 2

The window shown in Figure 3-13 opens.

3. Enable and modify your selected pool encryption settings.

Note: Pool encryption settings are *disabled* by default.

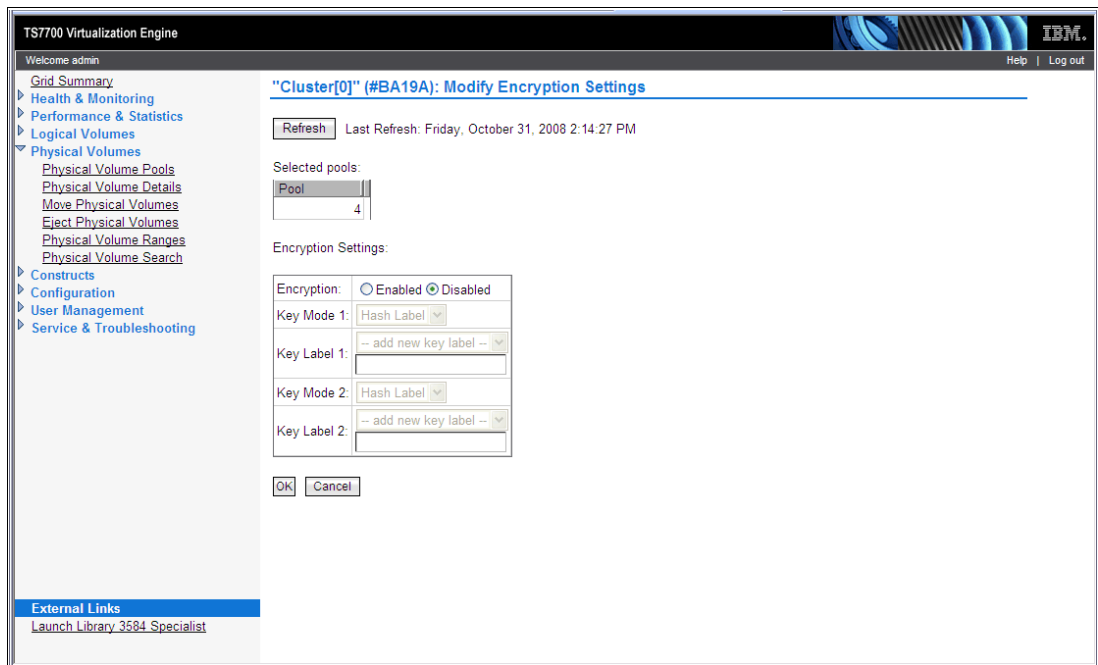


Figure 3-13 Modify encryption settings step 3

3.7.2 Encryption Key Manager (EKM)

Your EKMs should be installed, configured, and operational before installing the encryption feature on the TS7740 Virtualization Engine. You should also create the certificates and keys you plan to use for encrypting your backstore tape cartridges.

Although you can run with a single EKM, we strongly suggest that you have two EKMs for use by the TS7740 Virtualization Engine. Each EKM should have all the required keys in their respective keystores. The EKMs should have independent power and network connections to maximize the chances that at least one of them is reachable from the TS7740 Virtualization Engine when needed. If the TS7740 Virtualization Engine is unable to contact either EKM when required, you might temporarily lose access to migrated logical volumes and will not be able to move logical volumes in encryption-enabled storage pools out of cache.

See *IBM Encryption Key Manager Component for the Java Platform Introduction, Planning, and User's Guide*, GA76-0418 for details about installing and configuring your EKMs.

Because the TS7740 Virtualization Engine maintains TCP/IP connections with the EKMs at all times, the EKM configuration file should have the following setting to prevent the EKM from timing out on these always-on connections:

```
TransportListener.tcp.timeout = 0
```

3.7.3 Tivoli Key Lifecycle Manager

As an enhancement and follow-on to the Encryption Key Manager, the Tivoli Key Lifecycle Manager can be used to encrypt your data with the TS1120 and TS1130 Tape Drives. Like the Encryption Key Manager, the Tivoli Key Lifecycle Manager serves data keys to the tape drive. The first release of Tivoli Key Lifecycle Manager focuses on ease of use and provides a new graphical user interface (GUI) to help with the installation and configuration of the key manager. It also allows for the creation and management of the key encrypting keys (certificates). If you already use the existing Encryption Key Manager, you can migrate easily to the new Tivoli Key Lifecycle Manager.

For additional information about the Tivoli Key Lifecycle Manager, see the following address:

<http://publib.boulder.ibm.com/infocenter/tivihelp/v2r1/index.jsp?topic=/com.ibm.tk1m.doc/welcome.htm>

3.7.4 Tape drives

Because data encryption is performed on the tape drives themselves, TS1120 Model E05 encryption-capable tape drives must be attached to the TS7740 Virtualization Engine. They also must be running in native (E05) mode rather than J1A emulation mode.

TS1130 Model E06 Tape Drives can also be attached to the TS7740 Virtualization Engine. Intermixing with TS1120 Tape Drives is not allowed.

To enable encryption on a TS7740 Virtualization Engine with 3592 Model J1A drives attached, plan your back-end drive capacity with the expectation that these drives will not be used. The 3592 Model J1A tape drives are not encryption capable and should be detached. The TS7740 Virtualization Engine does not allow a mixture of drive types to be used. The J1A drives can be redeployed in other subsystems or used as direct-attached drives for open-systems hosts. If you have a mixture of J1A and E05 drives attached to your TS7740 Virtualization Engine and cannot detach the J1A drives right away, you may proceed as long as you have a minimum of four encryption-capable E05 drives attached. Be aware, however,

that the J1A drives will not be used by the TS7740 Virtualization Engine after the E05 drives are put into native mode.

All TS1120 Tape Drives with Feature Code 5592 or 9592 are encryption-capable.

3.7.5 TS7740 Virtualization Engine

The TS7740 Virtualization Engine must be running Release 1.2 or later. Feature Code 9900 must be installed to access the encryption settings.

The TS7740 Virtualization Engine must not be configured to force the TS1120 drives into “J1A” mode. This setting may only be changed by your IBM System Service Representative. If you need to update the microcode level, be sure the IBM System Service Representative checks and changes this setting if needed.

3.7.6 Encryption Key Manager IP addresses

The Encryption Key Manager assists encryption-enabled tape drives in generating, protecting, storing, and maintaining encryption keys that are used to encrypt information being written to and decrypt information being read from tape media.

The panel shown in Figure 3-14 on page 158 is used to set the encryption key manager addresses in the TS7700 Virtualization Engine. this panel is not available on the TS7720 Virtualization Engine management interface.

A video tutorial is available from this panel. The tutorial presents the properties of the Encryption Key Manager; click the **View tutorial** link.

If the cluster does not possess a physical library, the Encryption Key Manager Addresses window is not visible on the TS7700 Virtualization Engine management interface.

The Encryption Key Manager assists encryption-enabled tape drives in generating, protecting, storing, and maintaining encryption keys that are used to encrypt information being written to and decrypt information being read from tape media.

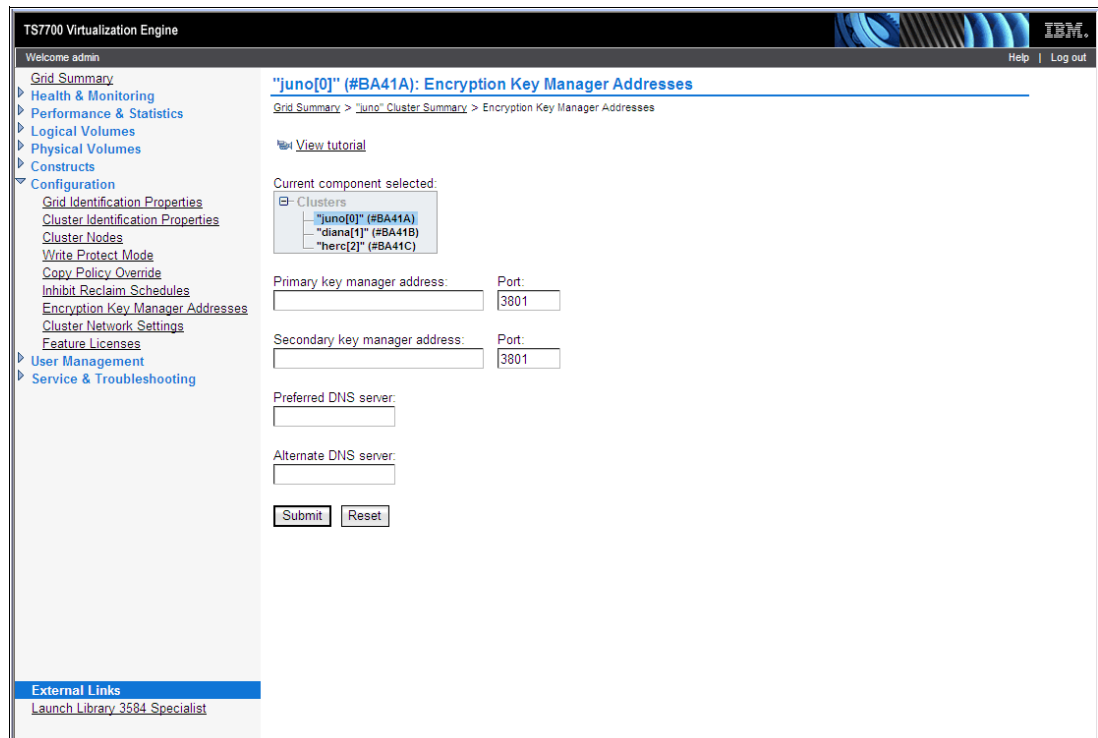


Figure 3-14 Encryption Key Manager Addresses

The following settings are used to configure the TS7700 Virtualization Engine connection to an Encryption Key Manager (see Figure 3-14 as needed):

- ▶ Primary key manager address: The key manager server name or IP address that is primarily used.
- ▶ Port (for the primary key manager address): The port number of the primary key manager. The default value is 3801. This field is only required if a primary key address is used.
- ▶ Secondary key manager address: The key manager server name or IP address that is used when the primary key manager is unavailable.
- ▶ Port (for the secondary key manager address): The port number of the secondary key manager. The default value is 3801. This field is only required if a secondary key address is used.
- ▶ Preferred DNS server: The Domain Name Server (DNS) that is primarily used. DNS addresses are needed only if you specify a symbolic domain name for one of the key manager addresses rather than a numeric IP address. If you need to specify a DNS, be sure to specify both a primary and an alternate so you do not lose access to your Encryption Key Manager because of one of the DNS servers being down or inaccessible. This address can be in IPv4 format.
- ▶ Alternate DNS server: The Domain Name Server that is used in case the preferred DNS server is unavailable. If a preferred DNS server is specified, specify an alternate DNS also. This address can be in IPv4 format.
- ▶ Using the Ping Test: Use the Ping Test buttons (in the Service & Troubleshooting selection) to check component network connection to a key manager after changing a component's address or port. If you change a key manager address or port and do not submit the change before using the Ping Test button, the following warning is issued:
 "In order to perform a ping test you must first submit your address and port changes."

After the ping test has been issued, one of the following messages is issued:

- “The ping test against the address <key manager address> on port <port> was successful.”
- “The ping test against the address <key manager address> on port <port> from “<cluster>” has failed. The error returned was: <error text>.”

Click **Submit** to save changes to any of the settings. To discard changes and return the field settings to their current values, click **Reset**.

Note: The two key manager servers must be set up on separate machines to provide redundancy. Connection to a key manager is required to read encrypted data.

3.7.7 Encryption key management

The importance of data protection has become increasingly apparent with news reports of security breaches, loss and theft of personal and financial information, and government regulation. Encryption of backstore tapes minimizes the risk of unauthorized data access without excessive security management burdens or subsystem performance issues. The TS7740 Virtualization Engine R1.7 and later allow the use of both specific encryption keys and a default key allowing for flexibility and ease of use.

The Virtualization Engine TS7740 Virtualization Engine provides the means to manage the use of encryption and what keys are used on a storage pool basis. The Storage Group DFSMS construct specified for a logical tape volume determines which storage pool is used for the primary and optional secondary copies in the TS7740 Virtualization Engine. Storage pool encryption parameters are configured through the TS7740 Virtualization Engine management interface under Physical Volume Pools.

With TS7700 Virtualization Engine R1.7, each pool can be defined to use Specific Encryption Keys or the Default Encryption Keys, which are defined at the key manager server:

► Specific Encryption Keys

Each pool in defined in the TS7740 Virtualization Engine can have its own unique encryption key. As part of enabling a pool for encryption, enter one or two key labels for the pool and a key mode. A key label can be up to 64 characters. Key labels do not have to be unique per pool and the management interface provides the capability to assign the same key label to multiple pools. For each key, a key mode can be specified. The supported key modes are Label and Hash. As part of the encryption configuration through the management interface, provide IP addresses for a primary and an optional secondary key manager.

► Default Encryption Keys

As an extension of TS7740 Virtualization Engine encryption capabilities, TS7740 Virtualization Engine R1.7 supports the use of a default key. This key simplifies the management of the encryption infrastructure because no future changes are required at the TS7740 Virtualization Engine. After a pool is defined to use the default key the management of encryption parameters is performed at the key manager.

For additional details about encryption and setting up you encryption key management solution, see *IBM System Storage Data Encryption*, SG24-7797.

3.7.8 Installation

FC9900 provides explicit instructions about setting up the drives and activating the feature on the TS7740 Virtualization Engine. Briefly, the installation steps are as follows:

1. Determine which TS1120 or TS1130 drives are attached to the TS7740 Virtualization Engine.
2. Verify that the tape drives are encryption-capable.
3. Configure the drives to be encryption-enabled by specifying the “System-Managed” encryption method.
4. Install the FC9900 License Key.
5. Set up the EKM or TKLM IP addresses and port information.
6. Verify connection to the EKMs or TKLMs.

An essential step is that you configure the tape drives for System-Managed encryption. The TS7740 Virtualization Engine uses the drives in this mode only and does not support Library-Managed or Application-Managed encryption.

Important: After the TS7740 Virtualization Engine is using drives for encrypted physical tape volumes, it will put drives that are not properly enabled for encryption offline to the subsystem. TS3500 Tape Library operators must be made aware that they should leave attached TS7740 Virtualization Engine drives in System-Managed encryption mode at all times so that drives are not taken offline.

For a comprehensive TS7740 Virtualization Engine encryption implementation plan, see “Implementing TS7700 Tape Encryption” in *IBM System Storage Tape Encryption Solutions*, SG24-7320.

3.8 Tape analysis and sizing the TS7700 Virtualization Engine

In this section, we document the process of using various tools to analyze your current tape environment and to size the TS7700 Virtualization Engine to meet your requirements. We show you how to get access to a tools library that offers many jobs to analyze your current environment and a procedure to unload specific SMF records for a comprehensive sizing with *BatchMagic*, which has to be done by your IBM System Service Representative or Business Partner.

3.8.1 IBM tape tools

Most of the IBM tape tools are available to customers, but some, such as, for example, BatchMagic, are only available to IBM personnel and Business Partners. You may download the tools that are generally available from the following address:

<ftp://ftp.software.ibm.com/storage/tapetool/>

A page opens to a list of TXT, PDF, and EXE files. To start, open the OVERVIEW.PDF file to see a brief description of all the various tool jobs. All jobs are found in the IBMT00LS.EXE file, which is a self-extracting ZIP file that will, after it has been downloaded to your PC, expand into four separate:

- ▶ IBMJCL.XMI: JCL for current tape analysis tools
- ▶ IBMCNTL.XMI: Parameters needed for job execution
- ▶ IBMLOAD.XMI: Load library for executable load modules
- ▶ IBMPAT.XMI: Data pattern library, which is only needed if you run the QSAMDRVR utility

Two areas of investigation can assist in tuning your current tape environment by identifying the impact of factors that influence the overall performance of the TS7700 Virtualization Engine. These weak points are bad block sizes, that is, smaller than 16 KB, and low compression ratios, both of which can affect performance in a negative way.

Tape compression analysis for TS7700 Virtualization Engine

By analyzing the Miscellaneous Data Records (MDR) from SYS1.LOGREC data set or the EREP history file, you can see how well current tape volumes are compressing.

The following job stream has been created to help analyze these records. See the installation procedure in member \$\$INDEX file:

- ▶ EREPMDR: JCL to extract MDR records from EREP history file
- ▶ TAPECOMP: A program that reads either SYS1.LOGREC or the EREP history file and produces reports on current compression ratios and MB transferred per hour.

Additionally, since z/OS R1.7, the SMF 21 records have been recording both channel-byte and device-byte information. The TapeWise tool calculates data compression ratios for each volume. The following reports show compression ratios:

- ▶ HRS
- ▶ DSN
- ▶ MBS
- ▶ VOL

MDR analysis for bad TS7700 Virtualization Engine block sizes

Again, by analyzing the MDR from SYS1.LOGREC or the EREP history file, you can identify tape volumes that are writing small blocks to the TS7700 Virtualization Engine and causing extended job run times.

The following job stream has been created to help analyze these records. See the installation procedure in the member \$\$INDEX file:

- ▶ EREPMDR: JCL to extract MDR records from EREP history file.
- ▶ BADBLKSZ: A program that reads either SYS1.LOGREC or the EREP history file, finds volumes writing small block sizes, then gets the job name and data set name from a tape management system (TMS) copy.

Data collection and extraction

To correctly size the TS7700 Virtualization Engine, you have to analyze your current workload. The SMF records that are required to perform the analysis are record types 14, 15, and 21. IOSTATS uses records 14, 15, 21, and 30.

Collect the stated SMF records for all z/OS systems that share the current tape configuration and might have data migrated to the TS7700 Virtualization Engine. The data collected should span one month (to cover any month-end processing peaks) or at least those days that

represent the peak load in your current tape environment. Check in SYS1.PARMLIB in member SMF to see whether the required records are being collected. If they are not being collected, arrange for collection.

The steps shown in Figure 3-15 are as follows:

1. The TMS data and SMF data collection using the FORMCATS and SORTSMF procedure selects only the required tape processing-related SMF records and the TMS catalog information.
2. The files created are compressed by BMPACKT and BMPACKS procedures.
3. Download the packed files (ZIP file format) to your PC and send them by email to your IBM System Service Representative (SSR); if the size is too large for email transmission, copy the files to a CD and send it to the SSR.

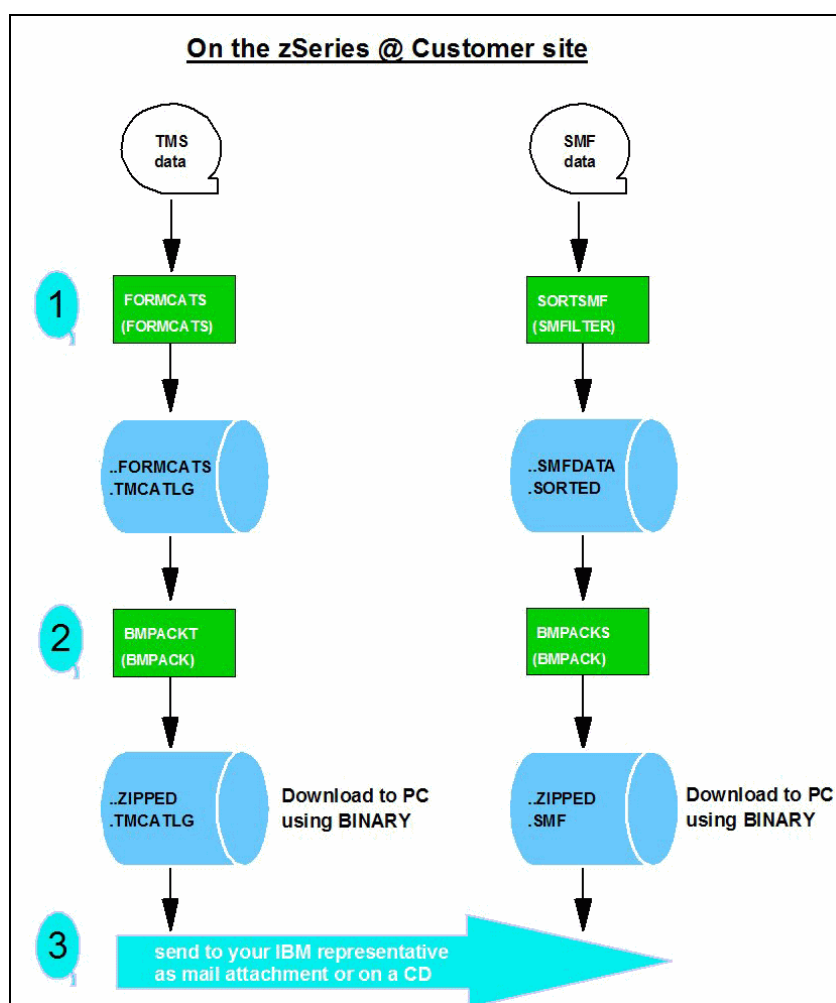


Figure 3-15 Unload process for TMS and SMF data

In addition to the extract file, the following information is useful for sizing the TS7700 Virtualization Engine:

- ▶ Number of volumes in your current tape library
This number includes all your tapes (those in automated libraries, on shelves, and off site). If you provide the unloaded Tape Management Catalog data, you do not have to collect the number of volumes.
- ▶ Criteria for identifying volumes
Because volumes that go off site, that are used as backup, or that should be ignored are of particular interest, their identification is important. Identifiers such as high-level qualifiers (HLQs), program names, or job names, must be documented.
- ▶ Number and type of tape control units installed
This information provides a good understanding of the current configuration and will help you identify the reasons for any apparent workload bottlenecks.
- ▶ Number and type of tape devices installed
This information will also help you identify the reasons for any apparent workload bottlenecks.
- ▶ Number and type of host channels attached to tape subsystems
This information will also help you identify the reasons for any apparent workload bottlenecks.

3.8.2 BatchMagic

This tool gives a comprehensive view of the current tape environment and predictive modeling of workloads and technologies. The general methodology behind this tool involves analyzing customer SMF type 14, 15, and 21 records, and data extracted from the customer tape management system. The TMS data is required only if you want to make a precise forecast of the cartridges to be ordered based on the current cartridge utilization that is stored in the TMS catalog.

A BatchMagic “run” involves data extraction, grouping data into workloads, and then targeting workloads to individual or multiple IBM tape technologies. BatchMagic examines Tape Management System catalogs and projects cartridges required with new technology and it models the operation of a TS7700 Virtualization Engine and 3592 drives (for TS7740 Virtualization Engine) and projects required resources. The reports from BatchMagic give a clear understanding of your current tape activities, and even more important, make projections for a TS7700 Virtualization Engine solution together with its major components, such as 3592 drives, which cover your overall sustained and peak throughput requirements.

BatchMagic is specifically for IBM internal and IBM Business Partner use.

3.8.3 Workload considerations

As with the prior generation of IBM Virtual Tape Server products, the TS7700 Virtualization Engine appears to the host systems as sixteen 3490E strings with a total of 256 devices attached per cluster. Any data that can reside on a 3480, 3490, 3590, or 3592, prior generations of VTS systems, or cartridges from other vendors, can reside on the TS7700 Virtualization Engine. However, processing characteristics of workloads differ, so some data is more suited for the TS7700 Virtualization Engine than other data. This topic highlights

several important considerations when you are deciding what workload to place in the TS7700 Virtualization Engine:

- **Throughput**

The TS7700 Virtualization Engine has a finite bandwidth capability, as does any other device attached to a host system, so plan to put workloads into the TS7700 Virtualization Engine that fit within its capabilities. With the 4 Gb FICON channels and large TS7700 Cache, there are few workloads that are not suitable for the TS7700 Virtualization Engine based on throughput.

- **Drive concurrency**

Each TS7700 Virtualization Engine appears to the host operating system as 256 3490E drives. If there are periods of time during the day when your tape processing jobs are limited by drive availability, the TS7700 Virtualization Engine might help your processing considerably. The design of the TS7700 Virtualization Engine allows transparent access to multiple logical volumes on the same stacked physical volume, because access to the logical volumes is solely through the TS7700 Virtualization Engine Cache. If you need access to more than one logical volume on a physical volume, it is provided without requiring any user involvement, unlike some alternatives, such as stacking by using JCL.

- **Allocation considerations**

For a discussion of scratch and specific allocation considerations in a TS7700 Virtualization Engine, see the “Load Balancing Considerations” section in *z/OS DFSMS OAM Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC35-0427.

- **Cartridge capacity utilization**

A key benefit of the TS7740 Virtualization Engine is its ability to fully use the capacity of the 3592 cartridges independent of the data set sizes written and to manage that capacity effectively without host or user involvement. A logical volume can contain up to 4 GB of data (12 GB, assuming data compressibility of 3:1) using the extended logical volume sizes. The actual size of a logical volume is only the amount of data written by the host. Therefore, even if an application only writes 20 MB to a 4 GB volume, only the 20 MB is kept in the TS7700 Virtualization Engine Cache or on a TS7740 Virtualization Engine managed physical volume. Very large data sets might gain little from the ability of the TS7740 Virtualization Engine to stack data, so you might want to leave them on native 3590 or 3592 cartridges.

- **Volume caching**

Often, one step of a job is writing a tape volume and a subsequent step (or job) is reading it. A major benefit can be gained using the TS7700 Virtualization Engine, because the data is cached in the TS7700 Virtualization Engine Cache, which effectively removes the rewind time, the robotics time, and load or thread times for the mount.

Figure 3-16 on page 165 shows the effect that a TS7700 Virtualization Engine can have on a job and drive assignment as compared to a native drive. The figure is a freehand drawing; it is not to scale. It shows typical estimated elapsed times for elements that make up the reading of data from a tape. When comparing the three timelines in Figure 3-16 on page 165, notice that the TS7700 Virtualization Engine Cache hit time line does not include robotics, load, or thread time at the beginning of the time line, and it does not include any rewind or unload time at the end.

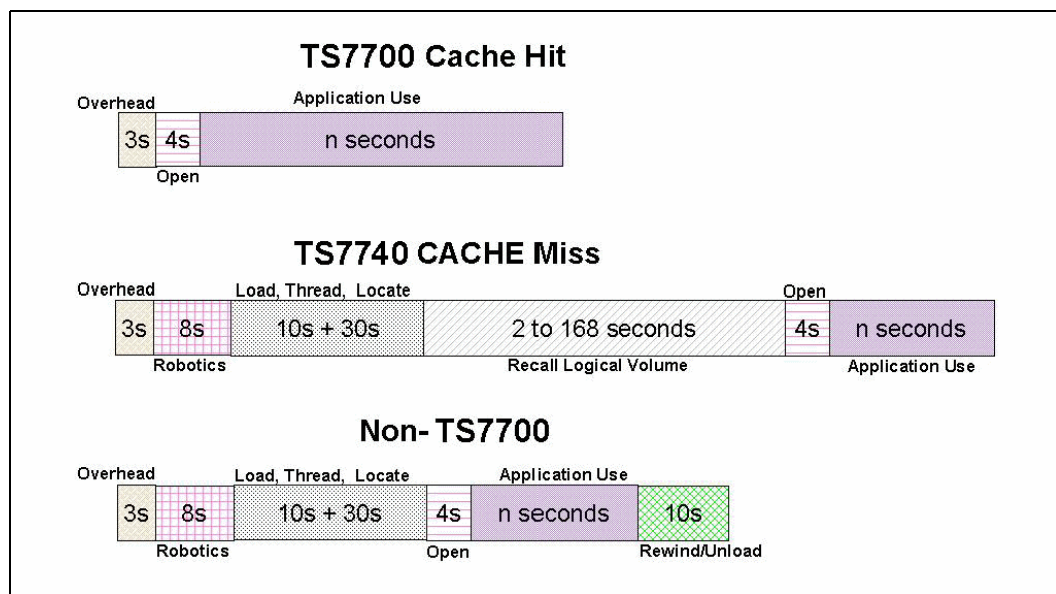


Figure 3-16 Tape processing time comparison (not to scale)

In this example, the TS7700 Virtualization Engine Cache hit results in a savings in tape processing elapsed time of 40 seconds.

The time reduction in the tape processing has two effects:

- It reduces the elapsed time of the job processing the tape.
- It frees up a drive earlier so the next job needing a tape drive can access it sooner, because there is no rewind or unload and robotics time after closing the data set.

When a job attempts to read a volume that is not in the TS7740 Virtualization Engine Cache, the logical volume is recalled from a stacked physical volume back into the cache. When a recall is necessary, the time to access the data is greater than if it were already in the cache. The size of the cache and the use of the cache management policies can reduce the number of recalls. Too much recall activity can negatively impact the overall throughput of the TS7740 Virtualization Engine.

During normal operation of a TS7700 Virtualization Engine Grid configuration, logical volume mount requests can be satisfied from the local TVC or a remote TVC. TS7700 Virtualization Engine algorithms can evaluate the mount request and determine the most effective way to satisfy the request from within the TS7700 Virtualization Engine Grid. If the *local* TVC does not have a current copy of the logical volume and a remote TVC does, the TS7700 Virtualization Engine can satisfy the mount request through the Grid by accessing the volume in the TVC of a *remote* TS7700 Virtualization Engine. The result is that in a Multi Cluster Configuration, the Grid combines the TS7700 Virtualization Engine TVCs to produce a larger effective cache size for logical mount request.

Notes:

- By *local*, we mean the TS7700 Virtualization Engine Cluster that is performing the logical mount to the host.
- By *remote*, we mean any other TS7700 Virtualization Engine that is participating in the same Grid as the local cluster.

► Scratch mount times

When a program issues a scratch mount to write data, the TS7700 Virtualization Engine completes the mount request without having to recall the logical volume into the cache. For workloads that create many tapes, this significantly reduces volume processing times and improves batch window efficiencies. The effect of using the Fast Ready attribute on TVC performance of scratch mounts is significant, because no physical mount is required. The performance for scratch mounts is the same as for TVC read hits. The comparison between the time taken to process a mount request on a subsystem with cache to a subsystem without cache is made in Figure 3-16 on page 165.

► Disaster recovery

The TS7700 Virtualization Engine Grid configuration is a perfect integrated solution for your disaster recovery data. The TS7700 Virtualization Engine Clusters in a multi-cluster grid may be separated over long distances and interconnected using a TCP/IP infrastructure to provide for automatic data replication. Data written to a local TS7700 Virtualization Engine is accessible at the remote TS7700 Virtualization Engine just as though it had been created there. Flexible replication policies make it easy to tailor the replication of the data to your business needs.

The Copy Export function, introduced with TS7700 Virtualization Engine Release 1.3 for a single-cluster grid and Release 1.4 for a multi-cluster grid, provides another disaster recovery method. The Copy Exported physical volumes can be used in an empty TS7700 Virtualization Engine to recover from a disaster. See 2.5.4, “Copy Export” on page 70 for more details.

► Multifile volumes

If you stack multiple files onto volumes today, by using JCL constructs or some other method, the reason you are stacking is most likely to better use cartridge capacity. Automatic utilization of physical cartridge capacity is one of the primary attributes of the TS7740 Virtualization Engine. Therefore, you should find that in many cases manual stacking of data sets onto volumes is no longer required. If you are planning a new application that uses JCL to stack data sets onto a volume, the TS7740 Virtualization Engine makes this JCL step unnecessary. Multifile volumes that are moved to the TS7740 Virtualization Engine can also work without changing the stacking. However, the TS7740 Virtualization Engine recalls the complete logical volume to the TS7740 Virtualization Engine Cache if the volume is not in cache, rather than moving each file as you access it. Therefore, in some cases, a possible advantage is to let the TS7740 Virtualization Engine do the stacking automatically for you. It can save you manual management overhead and, in some cases, host CPU cycles, host channel bandwidth, DASD space, or a combination of all of these items.

► Interchange or off-site storage

As currently delivered, the TS7740 Virtualization Engine does not support a capability to remove a stacked volume to be used for interchange. Native 3490, 3590, or 3592 tapes are better suited to your data for interchange. The Copy Export function can be used for off-site storage of data for the purposes of disaster recovery. See 2.5.4, “Copy Export” on page 70 for more details.

With the wide range of capabilities that the TS7700 Virtualization Engine provides, unless your data sets are very large or require interchange, the TS7700 Virtualization Engine is likely a suitable place to put your data.

3.9 Education and training

If you have never had an IBM library installed before, plan for education about the library for your storage administrators and for the operating personnel. Education topics should include handling of the TS3500 Tape Library Specialist, usage of the TSSC, and operation of the library itself. The amount of education and training your staff requires on the TS7700 Virtualization Engine depends on a number of factors:

- ▶ Are you installing the TS7740 Virtualization Engine in an existing TS3500 Tape Library environment?
- ▶ Are both the TS7740 Virtualization Engine and the library new to your site?
- ▶ Are you installing the TS7700 Virtualization Engine into an existing composite library?
- ▶ Is the library or the TS7700 Virtualization Engine shared among multiple host systems?
- ▶ Do you have existing tape drives at your site?
- ▶ Are you installing the TS7720 Virtualization Engine Disk Only solution?

3.9.1 Adding a TS7740 Virtualization Engine to an existing TS3500 Tape Library

If you are adding a TS7740 Virtualization Engine to an existing TS3500 Tape Library, the training needed for your operators, system programmers, and storage administrators is minimal. The operator intervention and help pages are on the management interface menu and contain the actions necessary to resolve the conditions.

For an existing TS3500 Tape Library, training for your operations staff should include training on the TS7740 Virtualization Engine management interface.

Your storage administrators and system programmers should receive the same training as the operations staff, plus:

- ▶ Software choices and how they affect the TS7700 Virtualization Engine
- ▶ Disaster recovery considerations

For additional information about this topic, see *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789.

3.9.2 Sharing the TS7740 Virtualization Engine with other systems within a TS3500 Tape Library

In addition to the material in 3.8, "Tape analysis and sizing the TS7700 Virtualization Engine" on page 160, see *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789 for information about this topic.

3.9.3 Implementation services

A range of services is available to assist you with your TS7700 Virtualization Engine. IBM can deliver end-to-end storage services to help you throughout all phases of the IT life cycle, including:

- ▶ **Assessment**

Provides an analysis of the tape environment and an evaluation of potential savings and benefits of installing new technology, such as tape automation, virtual tape, and tape mounting management.

- ▶ **Planning**

Assists in the collection of information required for tape analysis, analysis of the customer's current environment, and the design of the ATL environment, including coding and testing of customized Data Facility Storage Management Subsystem (DFSMS) Automatic Class Selection routines.

- ▶ **Implementation**

- TS7700 Virtualization Engine implementation provides technical consultation, software planning, and assistance and operator education to customers implementing an IBM TS7700 Virtualization Engine. Options include Data Analysis and SMS Tape Design for analysis of tape data in preparation and design of a DFSMS tape solution, New Allocations for assistance and monitoring of tape data migration through new tape volume allocations, and Static Data for migration of existing data to a TS7700 Virtualization Engine or traditional automated tape library.
- Automated Tape Library (ATL) implementation provides technical consultation, software planning assistance, and operational education to customers implementing an ATL.
- Tape Copy Service performs copying of data on existing media into an ATL. This service is generally performed subsequent to an Automated Library, TS7700 Virtualization Engine, or Grid implementation.

- ▶ **Support**

Support Line provides access to technical support professionals who are experts in all IBM tape products.

For more information about storage services and IBM Global Services, contact your IBM sales representative, or visit the following address:

<http://www.ibm.com/services>

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.

IBM Integrated Technology Services

IBM Integrated Technology Services include business consulting, outsourcing, hosting services, applications, and other technology management tasks.

These services help you learn about, plan, install, manage, or optimize your IT infrastructure to be an On Demand business. They can help you integrate your high-speed networks, storage systems, application servers, wireless protocols, and an array of platforms, middleware, and communications software for IBM and many non-IBM offerings.

For details about available services, contact your IBM System Service Representative or visit the following address:

<http://www.ibm.com/services/>

For details about available IBM Business Continuity and Recovery Services, contact your IBM representative or visit the following address:

<http://www.ibm.com/services/continuity>

For details about education offerings related to specific products, visit the following address:

<http://www.ibm.com/services/learning/index.html>

Select your country, and then select the product as the category.

3.10 Planning considerations

In this section, we consider the most important steps to be performed: from initial planning up to the complete installation or migration, going through hardware, software, educational, and performance monitoring activities.

Table 3-12 can help you with planning TS7700 Virtualization Engine preinstallation and sizing. Use the table as a checklist for the main tasks needed to complete the TS7700 Virtualization Engine installation.

Table 3-12 Main checklist

State	Task	Owner	Date	Reference
	Initial meeting			
	Physical planning			3.2, "Hardware installation and infrastructure planning" on page 125 and your IBM SSR
	Host connectivity			3.3, "Remote installations and FICON switch support" on page 132
	Hardware installation			Specific hardware manuals and your IBM SSR
	IP connectivity			"TS7740 Virtualization Engine configuration upgrades" on page 111

State	Task	Owner	Date	Reference
	HCD			5.2, "Hardware configuration definition" on page 269
	Maintenance check (PSP)			Preventive Service Planning buckets
	SMS			Appendix D, "JES3 examples and information" on page 821
	OAM			<i>z/OS DFSMS OAM Planning, Installation, and Storage Administration Guide for Tape Libraries</i> , SC35-0427
	RMM			<i>z/OS V1 DFSMSrmm Implementation and Customization Guide</i> , SC26-7405
	TS7700 Virtualization Engine customization			Chapter 8, "Operation" on page 409
	Set up BVIR			9.2, "TS7700 components and tasks distribution" on page 604
	Specialists training			
	DR implications			Chapter 10, "Failover and disaster recovery scenarios" on page 725
	Functional/performance test			Chapter 9, "Performance and monitoring" on page 589

State	Task	Owner	Date	Reference
	Cutover to production			
	Post-installation tasks (if any)			9.11, "IBM Tape Tools" on page 706
	Data migration (if required)			Chapter 7.1, "Migration to TS7700 Virtualization Engine" on page 350



Part 2

Implementation and migration

In this part, we provide you with all the information required to implement the IBM Virtualization Engine TS7740 R1.7 and the IBM Virtualization Engine TS7720 in your environment, or to migrate from another tape solution to the IBM Virtualization Engine TS7700 R1.7.

This part contains the following chapters:

- ▶ Chapter 4, “Hardware implementation” on page 175
- ▶ Chapter 5, “Software implementation” on page 263
- ▶ Chapter 6, “Upgrade considerations” on page 315
- ▶ Chapter 7, “Migration scenarios” on page 349



Hardware implementation

In this chapter, we describe the hardware-related implementation steps for the IBM Virtualization Engine TS7700.

We cover all implementation steps that relate to the setup of the following products:

- ▶ IBM System Storage TS3500 Tape Library
- ▶ TS7700 Virtualization Engine

Note: IBM 3494 Tape Library is only supported through RPQ at Release 1.7.

For details about host software implementation, see Chapter 5, “Software implementation” on page 263, where we describe the hardware configuration definition (HCD) steps on the host and operating system-related definitions.

For information about earlier release upgrades to Release 1.7, see Chapter 7, “Migration scenarios” on page 349.

4.1 TS7700 Virtualization Engine implementation and installation considerations

The following sections discuss the implementation and installation tasks to set up the TS7700 Virtualization Engine. TS7700 Virtualization Engine includes both the IBM Virtualization Engine TS7720 and the IBM Virtualization Engine TS7740, so we use the specific names if a certain task only applies to one of both in the remainder of this chapter, because there are slight differences between the TS7720 Virtualization Engine and the TS7740 Virtualization Engine.

The TS7720 Virtualization Engine does not have a tape library attached, so the implementation steps related to a physical tape library, TS3500 Tape Library, do not apply to the TS7720 Virtualization Engine.

You can install the TS7740 Virtualization Engine together with your existing tape library. As the Library Manager functions now reside inside of TS7700 Virtualization Engine microcode, the existing 3494 Library Manager it is not needed for physical volume content management, only for physical mount and dismount operations.

With the TS3500 Tape Library, the TS7740 itself manages all necessary operations, so the IBM 3953 Tape System is no longer required to attach a TS7740 Virtualization Engine in a TS3500 Tape Library to a IBM System z host.

Note: System z attachment of native 3592 tape drives through a tape controller still requires the IBM 3953 Tape System. See *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789 for more information.

You may also install a new TS3500 Tape Library and a new TS7740 Virtualization Engine at the same time.

4.1.1 Implementation tasks

The TS7700 Virtualization Engine implementation can be logically separated into three major sections:

- TS7700 Virtualization Engine and Tape Library setup

Using the TS7700 Virtualization Engine and the TS3500 Tape Library (for TS7740 Virtualization Engine) interfaces for these setup steps:

- Defining the physical properties of the TS7740 Virtualization Engine partition, such as physical tape drives and cartridges, using the TS3500 Tape Library Specialist, which is the web browser interface to the TS3500 Tape Library.
- Defining specific customer settings, such as encryption, and to insert logical volumes into the TS7740 Virtualization Engine. You may also use the management interface (MI) to define logical volumes, management policies, and volume categories.

In this chapter, we provide details of these implementation steps.

- ▶ Hardware I/O configuration definition

This section relates to the system generation. It consists of processes such as FICON channel attachment to the host, HCD/IOCP definitions, and Missing Interrupt Handler (MIH) settings. This activity can be done prior to the physical hardware installation and be part of the preinstallation planning and activity.

We describe these installation steps in Chapter 5, “Software implementation” on page 263.

- ▶ TS7700 Virtualization Engine software definition

Here is where you define the new tape library to the individual host operating system.

In a System z environment with DFSMS/MVS, this includes updating DFSMS Automatic Class Selection (ACS) routines, Object Access Method (OAM), and your tape management system during this phase. You also define Data Class (DC), Management Class (MC), Storage Class (SC), and Storage Group (SG) constructs and selection policies, which are passed to the TS7700 Virtualization Engine.

We describe these installation steps in Chapter 5, “Software implementation” on page 263.

These three groups of implementation tasks can be done in parallel or sequentially; the HCD and host definitions can be completed before or after the actual hardware installation.

4.1.2 Installation tasks

The tasks outlined in this section are specific to the simultaneous installation of a TS3500 Tape Library and a TS7740 Virtualization Engine.

Note: 3494 Tape Library is only supported through RPQ at Release 1.7.

If you are installing a TS7740 Virtualization Engine in an existing 3494 or TS3500 Tape Library environment, some of these tasks might not apply to you:

- ▶ Hardware related activities (completed by your IBM System Service Representative):
 - Install the IBM TS3500 Tape Library.
 - Install any native drives that will not be TS7740 Virtualization Engine controlled.
 - Install the TS7740 Virtualization Engine Frame and corresponding D2x Frame(s) in the TS3500 Tape Library.
- ▶ Define drives to hosts:
 - z/OS
 - z/VM
 - z/VSE
 - TPF and z/TPF
- ▶ Software-specific activities:
 - Apply maintenance for the TS3500 Tape Library.
 - Apply maintenance for the TS7700 Virtualization Engine.
 - Verify or update exits for the tape management system (if applicable) and define logical volumes to it.
 - See *z/VM V5R4.0 DFSMS/VM Removable Media Services*, SC24-6090, and 5.4, “Software implementation in z/VM and z/VSE” on page 303.

- ▶ Specific TS7700 Virtualization Engine activities:
 - Define policies and constructs using the TS7700 Virtualization Engine management interface (MI).
 - Define the logical VOLSER ranges of the logical volumes through the TS7700 Virtualization Engine MI.
- ▶ Specific TS7740 Virtualization Engine installation activities:
 - Define the TS7740 Virtualization Engine environment using the TS3500 Tape Library Specialist.
 - Define pool properties.
 - Define the physical VOLSER ranges for TS7740 Virtualization Engine-owned physical volumes to the Library Manager and TS3500 Tape Library (if applicable).
 - Insert TS7740 Virtualization Engine-owned physical volumes in the tape library.
- ▶ If you will be using encryption, you must specify the System Managed encryption method in the TS3500 Web management interface, and set all drives for the Logical Library as Encryption Capable.

Note: Encryption does not work with tape drives in emulation mode. Tape drives must be set to Native mode.

These tasks are further, including the suggested order of events, later in this chapter.

After your TS7740 Virtualization Engine is installed on the TS3500 Tape Library, perform the following post-installation tasks:

- ▶ Schedule and complete operator training.
- ▶ Schedule and complete storage administrator training.

4.2 TS3500 Tape Library definitions (TS7740 Virtualization Engine only)

We now describe how to implement the TS3500 Tape Library in a System z environment. Use this section if you are implementing the TS7740 Virtualization Engine in a TS3500 Tape Library.

If your TS7700 virtualization engine does not have an associated Tape Library (TS7720) see 4.3, “Setting up the TS7700 Virtualization Engine” on page 198.

Your IBM System Service Representative (SSR) performs the hardware installation of the library, its associated tape control units, and the frames. This installation does not require customer involvement other than the appropriate planning (see Chapter 3, “Preinstallation planning and sizing” on page 107 for details).

However, we do cover the following topics in this section:

- ▶ Defining a logical library
- ▶ Cartridge Assignment Policies
- ▶ Six-character-VOLSER support
- ▶ Assigning drives and creating control paths
- ▶ Defining an encryption method and encryption capable drives

After the SSR has physically installed the library hardware, you can use the TS3500 Tape Library Specialist to set up the logical library, which is attached to the System z host.

Note: The steps described in the following section relate to the installation of a new IBM TS3500 Tape Library with all the required features, such as Advanced Library Management System (ALMS), already installed. If you are attaching an existing IBM TS3500 Tape Library that is already attached to Open Systems hosts to System z hosts also, see *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789 for additional actions that might be required.

4.2.1 Defining a logical library

The TS3500 Tape Library Specialist is required to define a logical library and perform the following tasks. Therefore, you should make sure that it is set up properly and working. For access through a standard-based web browser, an IP address must be configured, which will be done initially by the SSR during hardware installation at the TS3500 Tape Library operator window.

Important:

- ▶ Each TS7740 Virtualization Engine (from Release 1.5 on) requires its own Logical Library in a TS3500 Tape Library.
- ▶ The Advanced Library Management System (ALMS) feature must be installed and enabled to define a logical library partition in the TS3500 Tape Library.

Ensure ALMS is enabled

Prior to enabling ALMS, the ALMS License Key has to be entered through the TS3500 Tape Library operator window, because ALMS is a chargeable feature.

You may check the status of ALMS with the TS3500 Tape Library Specialist by selecting **Library** → **ALMS**, as shown in Figure 4-1.

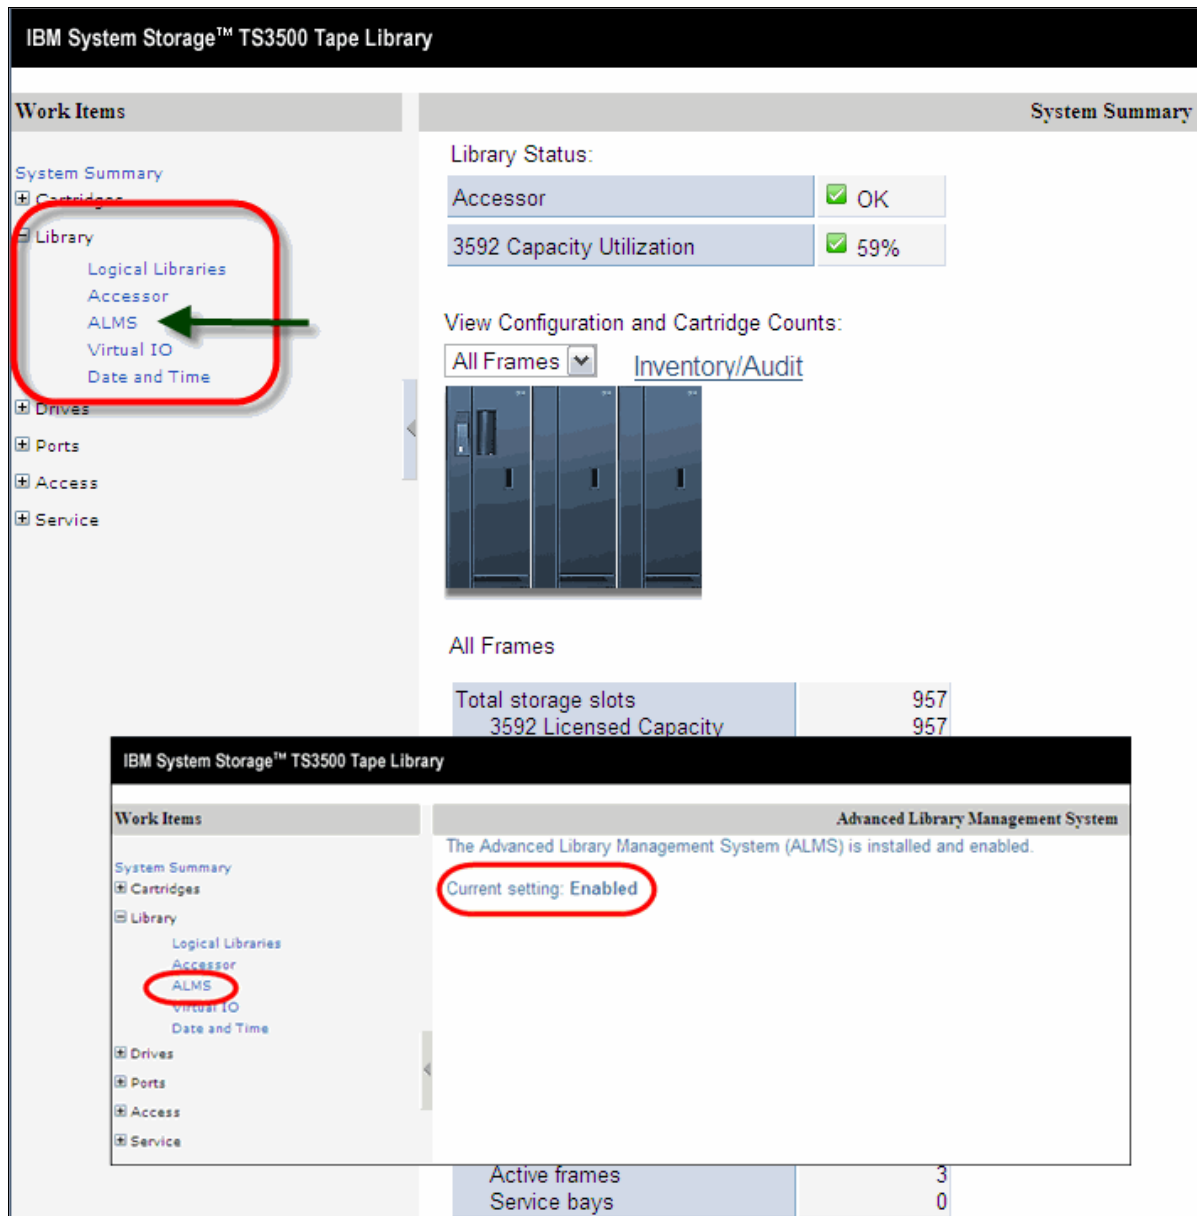


Figure 4-1 TS3500 Tape Library Specialist System Summary and Advanced Library Management System windows

As you can see in the ALMS window (at the bottom of the picture), ALMS is enabled for this TS3500 Tape Library.

When ALMS is enabled for the first time in a partitioned TS3500 Tape Library, the contents of each partition will be migrated to ALMS logical libraries. When enabling ALMS in a non-partitioned TS3500 Tape Library, cartridges that already reside in the library are migrated to the new ALMS single logical library.

Creating a new logical library with ALMS

This function is valid and available only if ALMS is enabled.

Note: You can create or remove a logical library from the TS3500 Tape Library by using the Tape Library Specialist web interface.

From the main section of TS3500 Tape Library Specialist Welcome window, go to the work items on the left side of the window and select **Library** → **Logical Libraries**, as shown in Figure 4-2. From the Select Action drop-down menu, select **Create** and click **Go**.

IBM System Storage™ TS3500 Tape Library

Work Items

- Welcome Page
- Cartridges
- Library
 - Frames
 - Logical Libraries
 - Accessor
 - ALMS
 - Virtual IO
 - Date and Time
- Drives
- Ports
- Access
- Service

Manage Logical Libraries

Refresh Last Refresh: 10/13/2008 17:02:24

Total Logical Libraries: 4

Select	Logical Library	Type	# Drives		# Cartridges		# VIO Slots		Exports	Encryption Method	Volser
			Dedicated	Shared	Assigned	Max.	Assigned	Max.			
<input type="checkbox"/>	ELC	3592	14	0	126	150	1	16	Show	N/A	8
<input type="checkbox"/>	heiko	LTO	0	0	0	6887	4	255	Hide	None	8
<input type="checkbox"/>	Lothar 75X0Kurs	3592	0	0	0	559	1	255	Show	None	8
<input type="checkbox"/>	CV7_lower_LL1	3592	3	0	7	559	3	16	Show	None	8

Figure 4-2 Create Logical Library starting window

An additional window, named Create Logical Library, opens; both windows are shown in Figure 4-3.

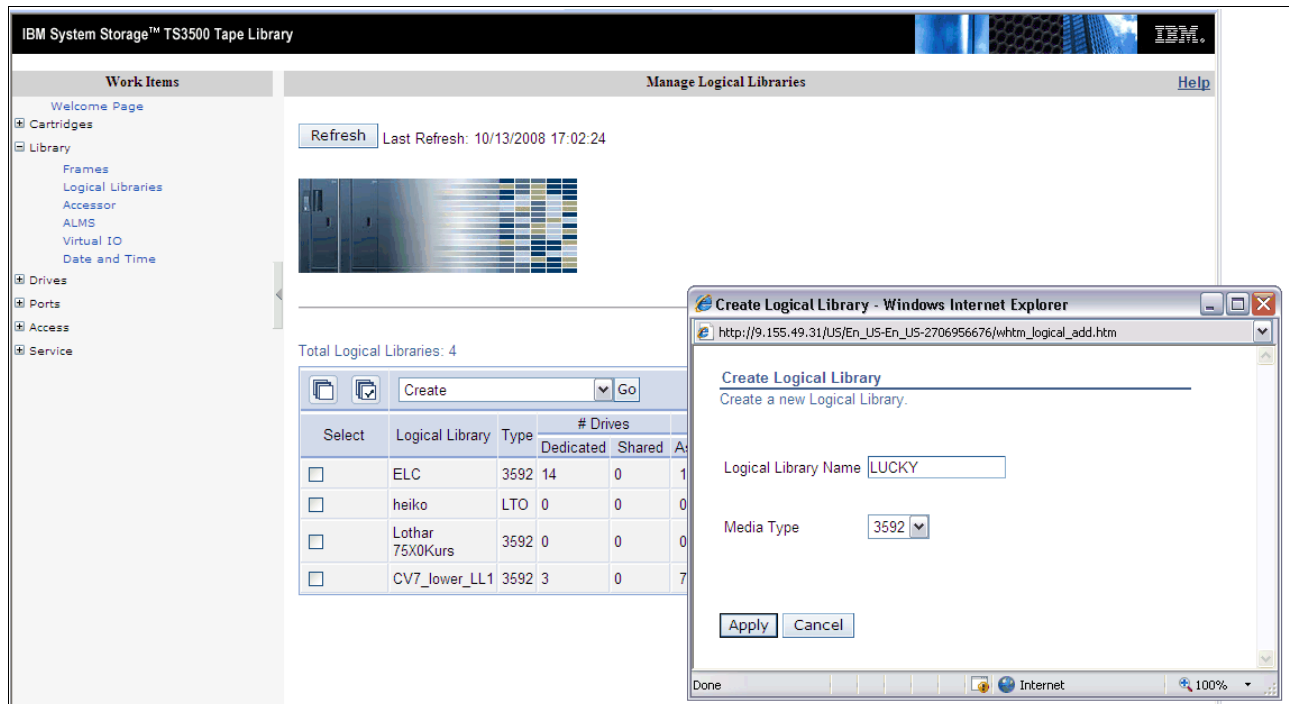


Figure 4-3 Create Logical Library windows

Type in the logical library name (up to 15 characters), select the media type (3592 for TS7740), and then click **Apply**. The new logical library is created and will appear in the logical library list when the window is refreshed.

After the logical library is created, you may display its characteristics by selecting **Library** → **Logical Libraries** under work items on the left side of the window, as shown in Figure 4-4 on page 183. From the Select Action drop-down menu, select **Detail** and then click **Go**.

In the Logical Library Details window, you see the element address range. The starting element address of each newly created logical library starts one element higher, such as:

- ▶ Logical Library 1: Starting SCSI element address is 1025.
- ▶ Logical Library 2: Starting SCSI element address is 1026.
- ▶ Logical Library 3: Starting SCSI element address is 1027.

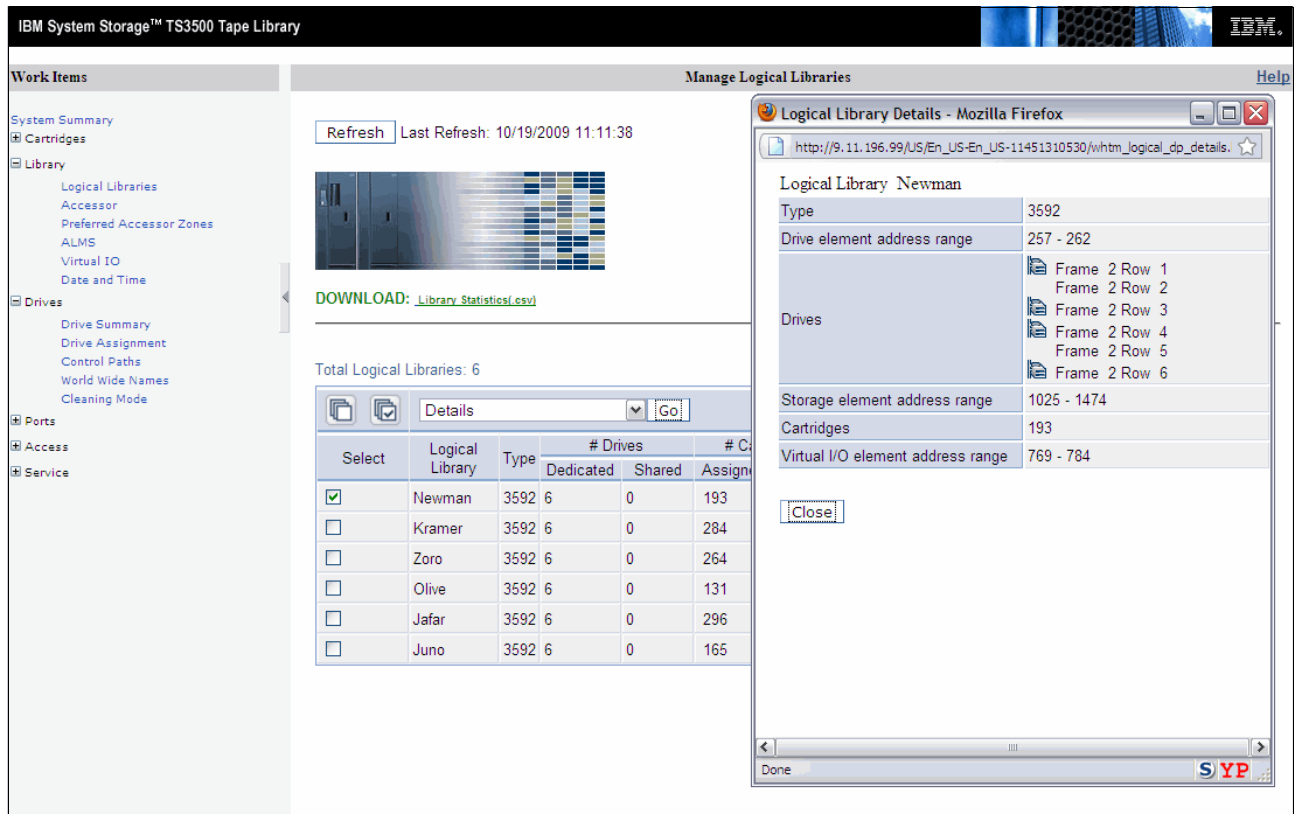


Figure 4-4 Recording of the starting SCSI element address of a logical library

Setting the maximum cartridges for the logical library

Define the maximum number of cartridge slots for the new logical library. If multiple logical libraries are defined, you can define the maximum number of cartridge slots for each logical library. This allows a logical library to grow without a reconfiguration each time you want to add empty slots. To define the quantity of cartridge slots, select a new logical library from the list, and from drop-down menu, select **Maximum Cartridges**, and then click **Go**.

Figure 4-5 shows the web interface windows.

IBM System Storage™ TS3500 Tape Library

Work Items

System Summary

Cartridges

Data Cartridges

Cleaning Cartridges

I/O Station

Cartridge Assignment Policy

Barcode Encryption Policy

Key Label Mapping

Insert Notification

Library

Logical Libraries

Accessor

Preferred Accessor Zones

ALMS

Virtual IO

Date and Time

Drives

Drive Summary

Drive Assignment

Control Paths

World Wide Names

Cleaning Mode

Ports

Fibre Channel Summary

Access

Service

Refresh Last Refresh: 10/19/2009 18:13:29

DOWNLOAD: [Library Statistics.csv](#)

Total Logical Libraries: 7

Maximum Cartridges Go

Select	Logical Library	Type	# Drives		# Cartridges		# VIO Slots		Exports	Encryption Method	Volser
			Dedicated	Shared	Assigned	Max.	Assigned	Max.			
<input type="checkbox"/>	Newman	3592	6	0	193	450	0	16	Show	System-Managed	8
<input type="checkbox"/>	Kramer	3592	6	0	284	452	0	16	Show	System-Managed	8
<input checked="" type="checkbox"/>	New_Library	3592	0	0	0	2037	0	255	Show	None	8
<input type="checkbox"/>	Zoro	3592	6	0	264	650	1	255	Show	System-Managed	8
<input type="checkbox"/>	Olive	3592	6	0	131	2037	0	255	Show	System-Managed	8
<input type="checkbox"/>	Jafar	3592	6	0	296	698	0	255	Show	System-Managed	8
<input type="checkbox"/>	Juno	3592	6	0	165	2037	0	255	Show	System-Managed	8

Change Maximum Number of Cartridges - ...

http://9.11.196.99/US/En_US-En_US-11451310530/whtm_logical_max_cartridges.htm

Change Maximum Number of Cartridges

Update the Maximum Number of Cartridges for the logical library.

The value must be between 1 and 20087

Logical Library New_Library

Current Maximum Number of Cartridges 2037

New Maximum Number of Cartridges 1

Apply Cancel

Physical slots of TS3500

Figure 4-5 Defining maximum number of cartridges

Figure 4-2 on page 181 shows an example of a maximum number of 150 cartridges specified for the logical library ELC.

Note: If multiple logical libraries are defined, you might consider capping the maximum size of some libraries, keeping growth under control. Resizing of logical libraries might require reconfiguration of the host applications for the selected logical library.

Setting eight-character Volser option in the new logical library

Check the new logical library for eight-character Volser reporting option. Go to Manage Logical Libraries screen under Library in the Tape Library Web Specialist, as shown in Figure 4-6.

IBM System Storage™ TS3500 Tape Library

Work Items

System Summary

Cartridges

Library

Logical Libraries

Accessories

Library Preferences

ALMS

Virtual IO

Date and Time

Drives

Ports

Access

Service

Manage Logical Libraries

Refresh Last Refresh: 9/15/2010 08:25:34

DOWNLOAD: [Library_Statistics.csv](#)

Total Logical Libraries: 6

Select	Logical Library	Type	# Drives		# Cartridges		# VIO Slots		Exports	Encryption Method	Volser
			Dedicated	Shared	Assigned	Max	Assigned	Max			
<input type="checkbox"/>	Newman	3592	4	0	140	450	0	16	Show	System-Managed	8
<input type="checkbox"/>	Kramer	3592	12	0	308	452	0	16	Show	System-Managed	8
<input type="checkbox"/>	Scotty	3592	12	0	260	650	0	255	Show	System-Managed	8
<input type="checkbox"/>	Olive	3592	12	0	140	2037	0	255	Show	System-Managed	8
<input type="checkbox"/>	Jafar	3592	4	0	296	698	7	255	Show	System-Managed	8
<input type="checkbox"/>	Juno	3592	12	0	164	2037	0	255	Show	System-Managed	8

Figure 4-6 Check Volser Reporting option

If the new logical library does not show 8 in the **Volser** column, correct the information:

1. Select a logical library.
2. Click the **Selection Action** drop-down menu.
3. Select **Modify Volser Report**.
4. Click **Go**.

Figure 4-7 illustrates the sequence.

IBM System Storage™ TS3500 Tape Library

Work Items

System Summary

Cartridges

Library

Logical Libraries

Accessories

Library Preferences

ALMS

Virtual IO

Date and Time

Drives

Ports

Access

Service

Manage Logical Libraries

Refresh Last Refresh: 9/15/2010 08:25:34

DOWNLOAD: [Library_Statistics.csv](#)

Total Logical Libraries: 6

1 Select Logical Library

2 Click here

3 Select Modify Volser Reporting

4 Click GO

Select	Logical Library	Type	# Drives		# Cartridges		# VIO Slots		Exports	Encryption Method	Volser
			Dedicated	Shared	Assigned	Max	Assigned	Max			
<input type="checkbox"/>	Newman	3592	4	0	140	450	0	16	Show	System-Managed	8
<input type="checkbox"/>	Kramer	3592	12	0	308	452	0	16	Show	System-Managed	8
<input type="checkbox"/>	Scotty	3592	12	0	260	650	0	255	Show	System-Managed	8
<input type="checkbox"/>	Olive	3592	12	0	140	2037	0	255	Show	System-Managed	8
<input type="checkbox"/>	Jafar	3592	4	0	296	698	7	255	Show	System-Managed	8
<input type="checkbox"/>	Juno	3592	12	0	164	2037	0	255	Show	System-Managed	8

Figure 4-7 Modifying Volser Reporting Length

In the pop-up window, select the correct VOLSER length (8-character) and apply it, as shown in Figure 4-8.

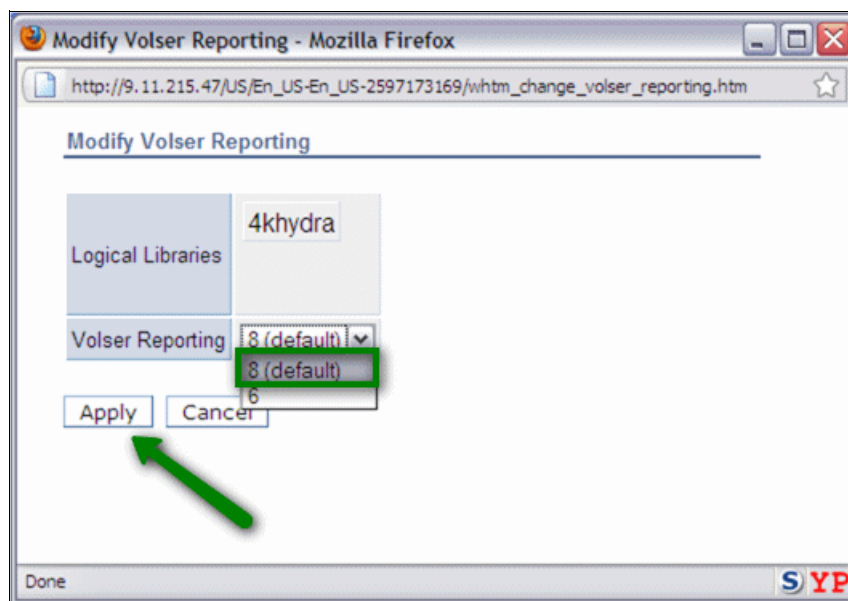


Figure 4-8 Setting 8-character Volser

4.2.2 Adding drives to the logical library

From the Logical Libraries window shown in Figure 4-4 on page 183, use the work items on the left side of the window to navigate to the requested web page by selecting **Drives** → **Drive Assignment**.

Note: You cannot intermix E06/EU6 with J1A or E05 drives.

This link takes you to a filtering window where you can select to have the drives displayed by drive element or by Logical Library. Upon your selection, a window opens so that you may add a drive to or remove it from a library configuration. It also enables you to share a drive between Logical Libraries and define its drive as a control path.

Note: Do not share drives belonging to a TS7400 (or any other Control Unit). They must be exclusive.

Figure 4-9 on page 187 shows the drive assignment window of a logical library that has all drives assigned.

Unassigned drives would appear in the Unassigned column with the box checked, so to assign them, check the appropriate drive box under the logical library name and click **Apply**.

Click the **Help** link at the top right corner of the window, shown in Figure 4-9, to see extended help information, such as detailed explanations of all the fields and functions of the window. The other TS3500 Tape Library Specialist windows provides similar help support.

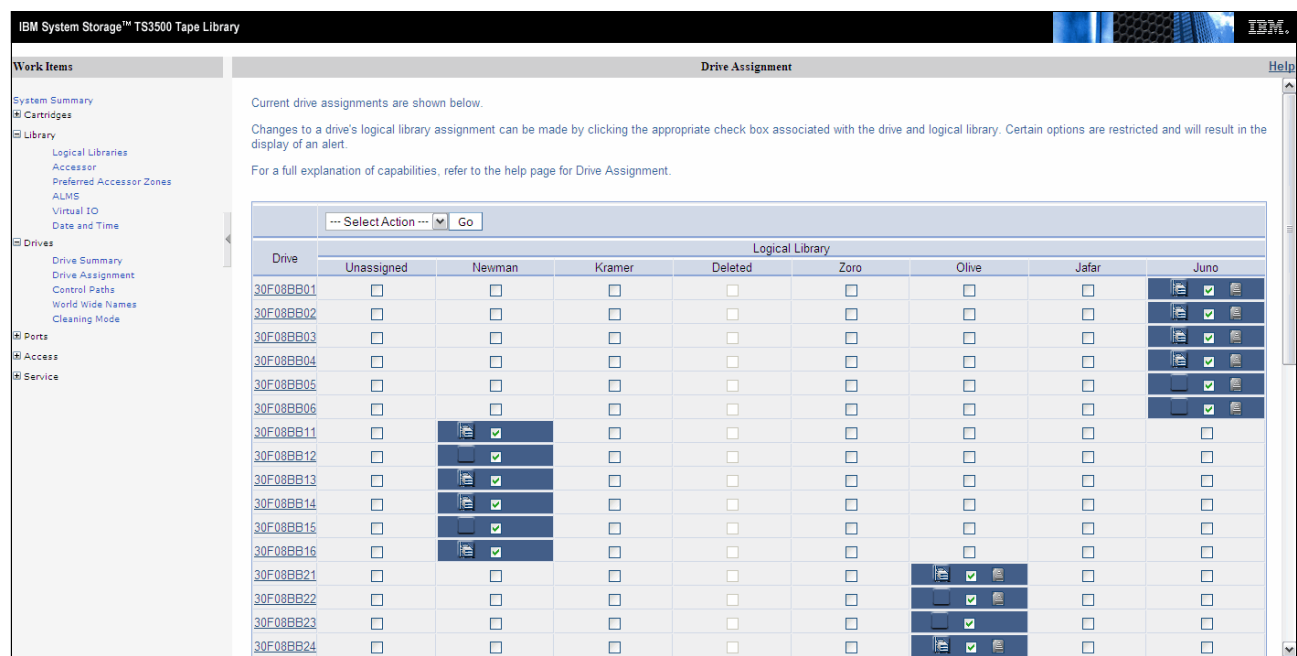


Figure 4-9 Drive Assignment window

In a multi-platform environment, you see logical libraries as shown in Figure 4-9, and you can reassign physical tape drives from one logical library to another. You can easily do this for the Open Systems environment, where the tape drives attach directly to the host systems without a tape controller or VTS/TS7700 Virtualization Engine.

Notes:

- ▶ Drive connections to a new host must be provided when reassigning a drive in a open system.
- ▶ Do not change drive assignments if they belong to an operating TS7740 or tape controller. Work with your IBM System Service Representative if necessary.

In a System z environment, a tape drive always attaches to one tape controller or VTS/TS7740 Virtualization Engine only. If you reassign a tape drive from an IBM 3953 Library Manager partition to an Open Systems partition temporarily, you must also physically detach the tape drive from the TS7740 Virtualization Engine or tape controller first and then attach the tape drive to the Open Systems host. Only IBM System Service Representatives (SSR) should perform these tasks to protect your tape operation from unplanned outages. The reconfiguration might require host software updates or TS7700 Virtualization Engine reconfiguration so that nothing can seriously interfere with your System z production work.

Important: In a System z environment, use the Drive Assignment window only to:

- ▶ *Initially* assign the tape drives from TS3500 Tape Library Web Specialist to a logical partition.
- ▶ Assign additional tape drives after they have been attached to the TS7740 Virtualization Engine or a tape controller.
- ▶ Remove physical tape drives from the configuration *before* they are physically detached from the TS7740 Virtualization Engine or tape controller.

In addition, never disable ALMS at the TS3500 Tape Library after it has been enabled for System z host support and System z tape drive attachment.

4.2.3 Defining the Encryption Method for the new Logical Library

After adding tape drives to the new Logical Library, you must specify the Encryption Method for the new Logical Library (if applicable).

Reminders:

- ▶ When using encryption, tape drives must be set to Native mode.
- ▶ To activate encryption, FC9900 must have been ordered for TS7400 and license key factory installed. Also the associated tape drives must be Encryption Capable 3592-E05 or 3592-E06 (although supported, 3592-J1A is not able to encrypt data).

Perform the following steps:

1. Check the drive mode by opening the Drive Summary window in the TS3500 management interface, as in Figure 4-10, and look in Mode column (this column is displayed only if drives in the tape library are emulation-capable).

IBM System Storage™ TS3500 Tape Library

Work Items

System Summary

Cartridges

Library

Drives

Drive Summary

Drive Assignment

Control Paths

World Wide Names

Cleaning Mode

Ports

Access

Service

Drives

Refresh Last Refresh: 9/17/2010 10:23:13

Select a Frame Select a Logical Library

All Frames OR All Libraries

DOWNLOAD: Drive Statistics(.csv)

Drive Details Go

Select	Drive	Location	Frame	Row	Logical Library	Element Address	Type	Contents	SCSI/Loop ID	Mode	Drive Interface	Status	Drive Display
<input type="checkbox"/>	30F092511	2	1	Archie	269	3592-E05	Empty	33	E05	Fibre	Online	DRIVE 23	
<input type="checkbox"/>	30F092512	2	2	Archie	270	3592-E05	Empty	34	E05	Fibre	Online	DRIVE 22	
<input type="checkbox"/>	30F092513	2	3	Archie	271	3592-E05	Empty	35	E05	Fibre	Online	DRIVE 31	
<input type="checkbox"/>	30F092514	2	4	Archie	272	3592-E05	Empty	36	E05	Fibre	Online	DRIVE 30	
<input type="checkbox"/>	30F092515	2	5	Archie	273	3592-E05	Empty	37	E05	Fibre	Online	DRIVE 29	
<input type="checkbox"/>	30F092516	2	6	Archie	274	3592-E05	Empty	38	E05	Fibre	Online	DRIVE 28	
<input type="checkbox"/>	30F092517	2	7	Archie	275	3592-E05	Empty	39	E05	Fibre	Online	DRIVE 27	
<input type="checkbox"/>	30F092518	2	8	Archie	276	3592-E05	Empty	40	E05	Fibre	Online	DRIVE 26	
<input type="checkbox"/>	30F092519	2	9	Archie	277	3592-E05	Empty	41	E05	Fibre	Online	DRIVE 25	
<input type="checkbox"/>	30F09251A	2	10	Archie	278	3592-E05	Empty	42	E05	Fibre	Online	DRIVE 24	
<input type="checkbox"/>	30F09251B	2	11	Archie	279	3592-E05	Empty	43	E05	Fibre	Online	DRIVE 21	
<input type="checkbox"/>	30F09251C	2	12	Archie	280	3592-E05	Empty	44	E05	Fibre	Online	FID2 C1 A511D7	

Checking drive mode

Figure 4-10 Drive Mode

2. If necessary change the drive mode to Native mode: In the Drive Summary window, select a drive and select **Change Emulation Mode**, as shown in Figure 4-11

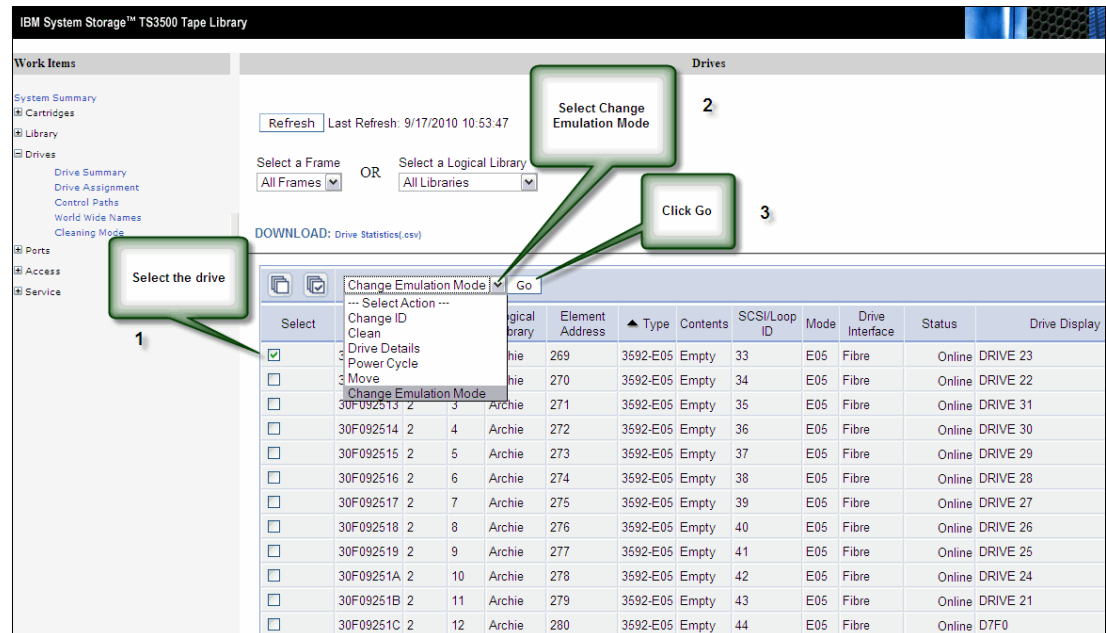


Figure 4-11 Changing drive emulation

3. In the next window that opens, select the native mode for the drive.
After drives are at desired mode, proceed with the Encryption Method definition.
4. In the TS3500 management interface select **Library** → **Logical Libraries**, select the Logical Library you are working with, select **Modify Encryption Method**, and then click **Go**. See Figure 4-12

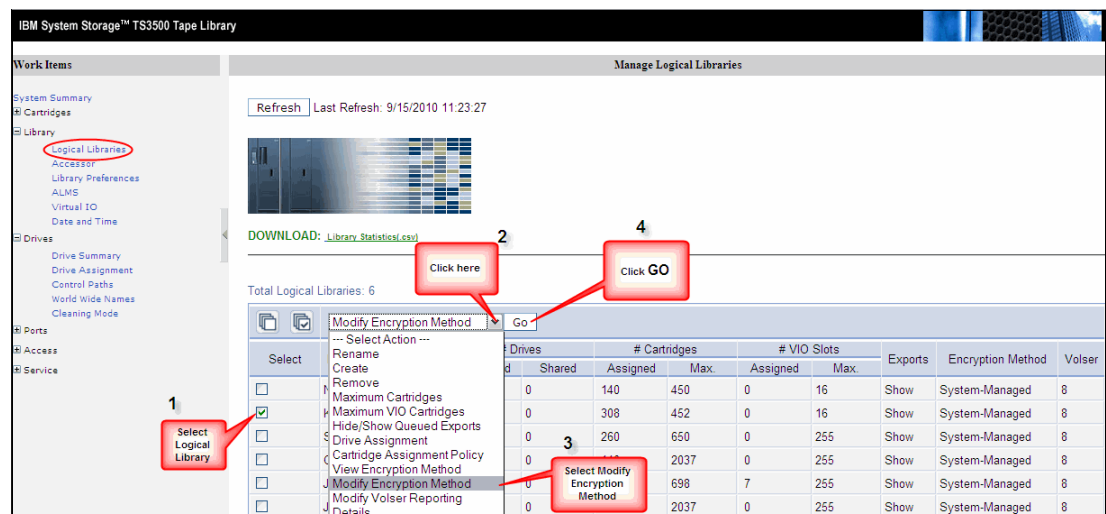


Figure 4-12 Selecting the Encryption Method

5. In the window that opens, select **System Managed** for the chosen method, and select all drives for this partition. See Figure 4-13.

Encryption Method

Encryption Method
System-Managed

Select drives to encrypt

Select	Drive	Encryption Capable
<input checked="" type="checkbox"/>	30F08BB31	Yes
<input checked="" type="checkbox"/>	30F08BB32	Yes
<input checked="" type="checkbox"/>	30F08BB33	Yes
<input checked="" type="checkbox"/>	30F08BB34	Yes
<input checked="" type="checkbox"/>	30F08BB35	Yes
<input checked="" type="checkbox"/>	30F08BB36	Yes
<input checked="" type="checkbox"/>	30F08BB37	Yes
<input checked="" type="checkbox"/>	30F08BB38	Yes
<input checked="" type="checkbox"/>	30F08BB39	Yes
<input checked="" type="checkbox"/>	30F08BB3A	Yes
<input checked="" type="checkbox"/>	30F08BB3B	Yes
<input checked="" type="checkbox"/>	30F08BB3C	Yes

Advanced Encryption Settings (for Engineering Support use only)

Advanced Method	No Advanced Setting
Advanced Policy	No Advanced Setting
Density Code	No Advanced Setting
Key Path	No Advanced Setting

Apply
Cancel

Figure 4-13 Setting the Encryption Method

To make encryption fully operational in the TS7740 configuration, additional steps are necessary. Work with your IBM System Service Representative to configure the Encryption parameters into TS7740 during installation process.

Attention: Keep the Advanced Encryption Settings as NO ADVANCED SETTING, unless specifically set otherwise by IBM Engineering.

4.2.4 Defining control path drives

Each S7740 must have at least one control path drive defined, and should have the optimum number of four control path drives. If possible, distribute the control path drives over more than one TS3500 Tape Library frame to avoid single points of failure.

Important: Define four control path drives for the TS7740 Virtualization Engine.

In a logical library, you may designate any empty, dedicated drive to become a control path drive. A drive that is loaded with a cartridge cannot become a control path until you remove the cartridge. Similarly, any drive that is a control path cannot be disabled until you remove the cartridge that it contains.

The definition of the control path drive is specified on the Drive Assignment window shown in Figure 4-14. Notice that drives, defined as control paths, are identified by the symbol on the left side of the drive box. You may change the control path drive definition by selecting or deselecting this symbol.




Figure 4-14 Control Path symbol

4.2.5 Defining Cartridge Assignment Policies (CAP)

The Cartridge Assignment Policy (CAP) of the TS3500 Tape Library is where you can assign ranges of physical cartridge volume serial numbers to specific logical libraries. If you have previously established a CAP and place a cartridge with a VOLSER that matches that range into the I/O station, the library automatically assigns that cartridge to the appropriate logical library.

Select **Cartridge Assignment Policy** from the Cartridges work items to add, change, and remove policies. The maximum quantity of Cartridge Assignment Policies for the entire TS3500 Tape Library must not exceed 300.

Figure 4-15 shows the VOLSER ranges defined for logical libraries.

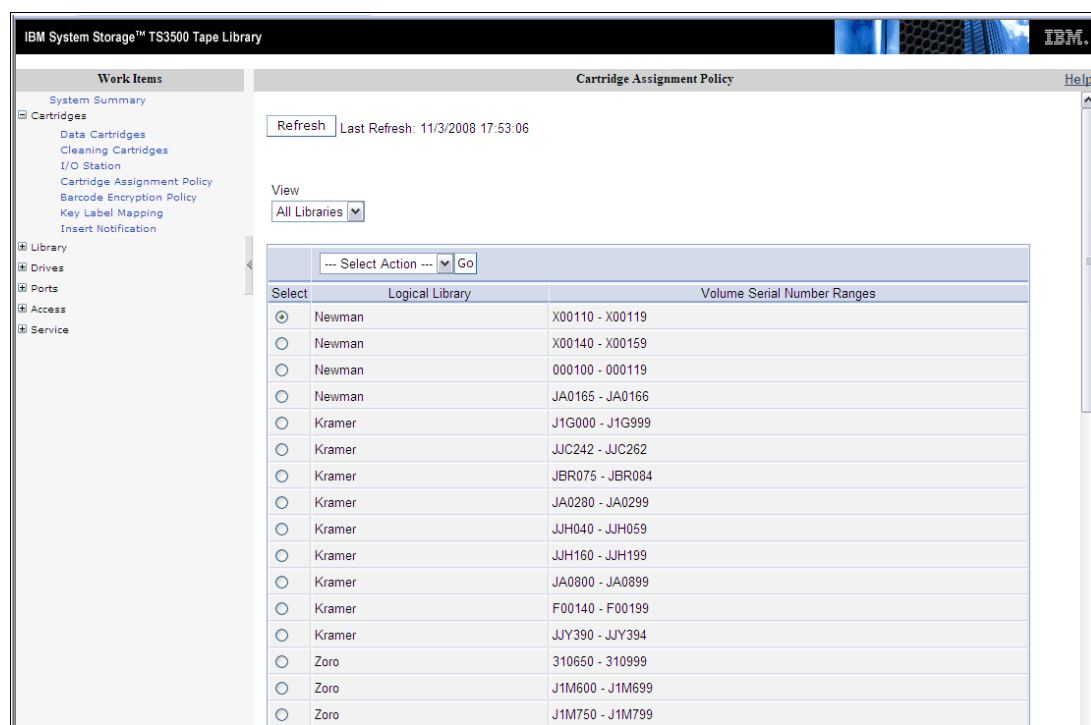


Figure 4-15 TS3500 Tape Library Cartridge Assignment Policy

The TS3500 Tape Library allows duplicate VOLSER ranges for different media types only. For example, Logical Library 1 and Logical Library 2 contain LTO media, and Logical Library 3 contains IBM 3592 media. Logical Library 1 has a Cartridge Assignment Policy of ABC100-ABC200. The library will reject an attempt to add a Cartridge Assignment Policy of ABC000-ABC300 to Logical Library 2 because the media type is the same (both LTO). However, the library does allow an attempt to add a Cartridge Assignment Policy of ABC000-ABC300 to Logical Library 3 because the media (3592) is different.

In an SMS-managed z/OS environment, all VOLSER identifiers across all storage hierarchies are required to be unique. Follow the same rules across host platforms also, whether or not you are sharing a TS3500 Tape Library between System z and Open Systems hosts.

Note: The Cartridge Assignment Policy does not reassign an already assigned tape cartridge. If needed, you must first make it unassigned, and then manually reassign it.

4.2.6 Inserting TS7740 Virtualization Engine physical volumes

The TS7740 Virtualization Engine subsystem manages both logical and physical volumes. The Cartridge Assignment Policy (CAP) of TS3500 Tape Library only affects the physical volumes associated with this TS7740 Virtualization Engine Logical Library. Logical Volumes are managed from TS7700 Virtualization Engine management interface (MI) only.

Perform the following steps to add physical cartridges:

1. Define Cartridge Assignment Policies at the *TS3500 Tape Library level* by using ALMS through the Web Specialist. This way ensures that all TS7740 Virtualization Engine ranges are recognized and assigned to the correct TS3500 Tape Library logical library partition (the logical library created for this specific TS7740 Virtualization Engine), before you begin any TS7700 Virtualization Engine MI definitions.
2. *Physically insert* volumes into the library by using the I/O station or by opening the library and placing cartridges in empty storage cells. Cartridges are assigned to the Library Manager partitions, according to the definitions.

These procedures ensure that TS7700 Virtualization Engine back-end cartridges will never be assigned to a host by accident. Figure 4-16 shows the flow of physical cartridge insertion and assignment to logical libraries for TS7740 Virtualization Engine R1.6.

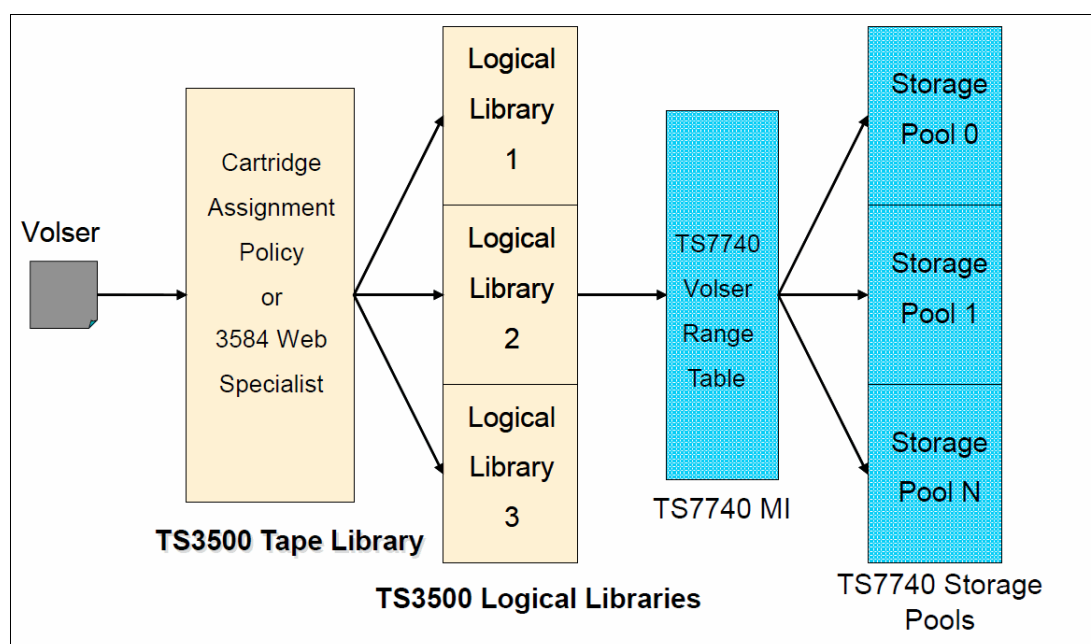


Figure 4-16 Volume assignment

Inserting physical volumes into the TS3500 Tape Library

Two methods are available for inserting physical volumes into the TS3500 Tape Library:

- Opening the library doors and inserting the volumes directly into the Tape Library storage empty cells (bulk loading)
- Using the TS3500 Tape Library I/O station

Insertion directly into storage cells

Note: With the Advanced Library Manager System (ALMS) enabled, cartridges that are not in a cartridge assignment policy will not be added to any logical library.

Use the operator panel of the TS3500 to pause it. Open the door and insert the cartridges into any empty slot, except those reserved for diagnostic cartridges, that is Frame 1, Column 1 in first Row (F01,C01,R01) in a single media-type library.

After completing the new media insertion, close the doors. After a little while (approximately 15 seconds) the TS3500 automatically inventories the frame or frames of the door you opened. During the inventory, the message INITIALIZING is displayed on the Activity screen on operator panel. When the inventory completes, TS3500 operator panel displays a Ready state. The TS7740 Virtualization Engine uploads its Logical Library inventory and updates its Integrated Library Manager inventory accordingly. After completing this operation, TS7740 Virtualization Engine Library reaches the Auto state.

Note: Place cartridges only in a frame that has an open front door. Do not add or remove cartridges from an adjacent frame.

Insertion by using the I/O station

With the ALMS, your TS3500 can be operating with or without virtual I/O being enabled. The procedure varies slightly depending on which mode is active in the library.

Basically, with virtual I/O enabled, TS3500 will move the cartridges from the physical I/O station into the physical library by itself. In the first moment, the cartridge leaves the physical I/O station and goes into a slot mapped as virtual I/O - SCSI elements between 769 (X'301') and 1023 (X'3FF') for that logical library pointed to by Cartridge Assignment Policy (CAP).

Each logical library has its own set of VIO slots, up to 256 (this is defined in the logical library creation, and can be altered later if needed).

With virtual I/O disabled, TS3584 does not move cartridges from the physical I/O station unless commanded to (by the TS7400 Virtualization Engine or any other host in control).

In both cases, TS3500 detects the volumes inserted when the I/O station door is closed and scans all I/O cells by using the barcode reader. CAP decides to which Logical Library those cartridges belong and then performs *either* of the following tasks:

- ▶ Moves them to that logical library's virtual I/O slots, if virtual I/O is enabled.
- ▶ Waits for a host command in this logical partition. Cartridges stay in the I/O station after barcode scan.

Given that inserted cartridges belong to a defined range in CAP of this Logical Library, and those ranges were defined into TS7740 Virtualization Engine Physical Volume Range, as explained in 4.3.3, "Defining VOLSER ranges for physical (stacked) volumes" on page 203, those cartridges will be assigned to this logical library. If any VOLSER is not in the range defined by CAP, the operator must identify the correct logical library as the destination using the Insert Notification window at the operator panel. If Insert Notification is not answered, the volume remains unassigned.

Note: Insert Notification is not supported on a high-density library. If a cartridge outside CAP defined ranges is inserted, it will remain unassigned without any notification.

Verify that the cartridges were correctly assigned by using the TS3500 management interface: Click **Cartridges** → **Data Cartridges** and select the appropriated Logical Library. If everything is correct, the inserted cartridges are listed. Alternatively, displaying the Unassigned/Shared volumes show none. See Figure 4-17 on page 195.

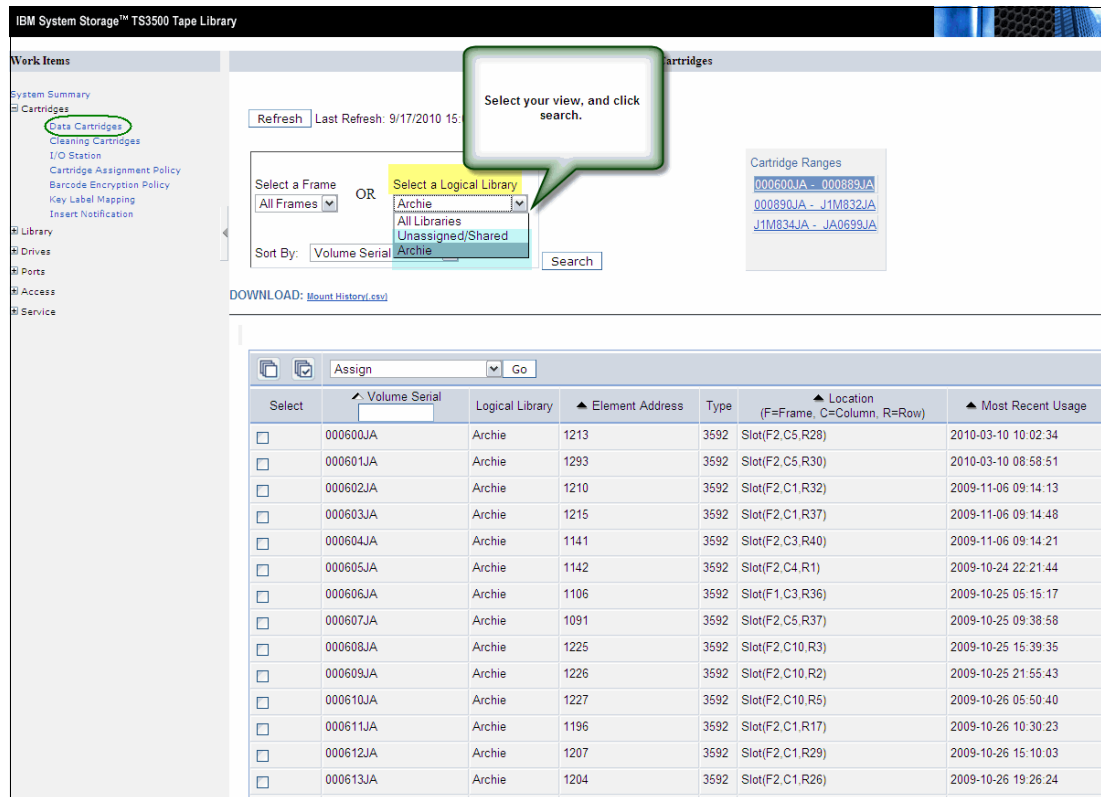


Figure 4-17 Checking volume assignment

Unassigned volumes in the TS3500 Tape Library

If a volume does not match the definitions in the CAP, and during the Insert Notification process no owner was specified, the cartridge remains unassigned in the TS3500 Tape Library. You can check for unassigned cartridges by using the TS3500 management interface and selecting **Cartridges** → **Data Cartridges**. In the drop-down menu, select **Unassigned/Shared** (shown in Figure 4-17).

Note: A cleaning cartridge is normally unassigned.

You can then assign the cartridges to the TS7740 Virtualization Engine logical library partitions using the TS3500 Tape Library Specialist web interface:

- ▶ For a Native or VTS partition, you must access the ETL and request an upload of the inventory to make sure that the now-assigned volumes are transferred to the 3593 Library Manager. This is because the manual assignment of the cartridges will not automatically process an upload of the inventory to the 3593 Library Manager.
- ▶ For a TS7740 Virtualization Engine cluster, select **Management Interface** → **Physical Volumes** → **Physical Volume Ranges** and click **Inventory Upload**, as shown in Figure 4-18 on page 196.

Important: Unassigned cartridges can exist in the TS3500 Tape Library, and in the TS7700 Virtualization Engine MI. But “unassigned” has separate meanings and requires separate actions from the operator in each system.

Unassigned volumes in TS7740 Virtualization Engine

A physical volume will be put in the Unassigned category by the TS7740 Virtualization Engine if it does not fit in any defined range of Physical Volume for this TS7740 Virtualization Engine. Defined Ranges and Unassigned volumes can be checked in Physical Volume Ranges window, shown in Figure 4-18. If an unassigned volume should belong to this TS7740 Virtualization Engine, a new range that includes this volume must be created, as discussed in 4.3.3, “Defining VOLSER ranges for physical (stacked) volumes” on page 203. If this volume was wrongly assigned to the TS7740 Virtualization Engine, you must eject and reassign it to the proper logical library in the TS3500 Tape Library. Also, double-check the CAP definitions in the TS7500 Tape Library.

The screenshot shows the IBM System Storage™ TS7700 Virtualization Engine interface. The main window is titled "dohnut[2]" (#BA96C): Physical Volume Ranges. It displays a list of clusters under "Current cluster selected:" including "archie[0]" (#BA96A), "veronica[1]" (#BA96B), "dohnut[2]" (#BA96C), and "veronica[3]" (#BA96D). Below this, there is a table for "Physical Volume Ranges" with columns for Start Volser, End Volser, Media Type, and Home Pool. To the right, there is a table for "Unassigned Volumes" with columns for Select, Volser, Media Type, and Pending Eject. The "Unassigned Volumes" table lists several volumes (J1M782, J1M784, J1M786, J1M798, J1M797, J1M820, J1M821, J1M822) all with Media Type JA-ETC and Pending Eject status No. The interface also includes a sidebar with navigation options like Grid Summary, Health & Monitoring, Performance & Statistics, Logical Volumes, Physical Volumes, and Constructs.

Figure 4-18 TS7740 Unassigned volumes

4.2.7 Assigning cartridges in TS3500 Tape Library to logical library partition

This procedure is necessary only if a cartridge was inserted, but a CAP was not provided in advance. To use this procedure you must assign the cartridge manually to a logical library in the TS3500 Tape Library.

Notes:

- ▶ Insert Notification is not supported in a high-density library. CAP must be correctly configured to provide automated assignment of all inserted cartridges.
- ▶ A cartridge that has being manually assigned to the TS7740 Logical Library will not show up automatically in the TS7740 inventory. An Inventory Upload is needed to refresh it.
- ▶ Cartridge assignment to a logical library is available only through the TS3500 Tape Library Specialist web interface. The operator panel does not provide this function.

Assigning a data cartridge

To assign a data cartridge to a logical library in the TS3500 Tape Library, perform the following steps:

1. Open the Tape Library Specialist web interface (type the library Ethernet IP address or the library URL on the URL line of a standard browser and press Enter.) The Welcome window opens.
2. Select **Cartridges** → **Data Cartridges**. The Data Cartridges window opens.
3. Select the logical library to which the cartridge is currently assigned and select how you want the cartridge range to be sorted. (The library can sort the cartridge by volume serial number, SCSI element address, or frame, column, and row location.) Select **Search**. The Cartridges window opens and shows all the ranges for the logical library that you specified.
4. Select the range that contains the data cartridge that you want to assign.
5. Select the data cartridge and then select **Assign**.
6. Select the logical library partition to which you want to assign the data cartridge.
7. Click **Next** to complete the function.

Note: After this procedure, an inventory upload from TS7740 Management interface (Physical Volume Ranges) is required so that the TS7740's inventory is updated.

Inserting a cleaning cartridge

Each drive in the TS3500 Tape Library will eventually need cleaning. Tape drives used by TS7740 subsystem can request a cleaning action when necessary. This cleaning is carried out by the TS3500 Tape Library automatically. However, you must provide the necessary cleaning cartridges.

Notes:

- ▶ The Advanced Library Management System (ALMS) must be enabled in a Library that is connected to a TS7740 Virtualization Engine. As a result, the cleaning mode is set to automatic and the library will manage drive cleaning.
- ▶ A cleaning cartridge is good for 50 cleaning actions.

The process to insert cleaning cartridges varies depending on the setup of the TS3500 Tape Library. A cleaning cartridge can be inserted by using the web interface or from the operator panel. As many as 100 cleaning cartridges can be inserted in a TS3500 Tape Library.

To insert a cleaning cartridge using the TS3500 Tape Library Specialist, perform the following steps:

1. Open the door of the I/O station and insert the cartridge so that the bar code label faces the interior of the library and the write-protect switch is on the right.
2. Close the door of the I/O station.
3. Type the Ethernet IP address on the URL line of the browser and press Enter. The Welcome Page opens.
4. Select **Cartridges** → **I/O Station**. The I/O Station window opens.
5. Follow the instructions in the window.

To insert a cleaning cartridge by using the operator panel, perform the following steps:

1. From the Library's Activity touchscreen, press **MENU** → **Manual Operations** → **Insert Cleaning Cartridges** → **Enter**. The library displays the message "Insert Cleaning Cartridge into I/O station before you continue. Do you want to continue?"
2. Open the I/O station and insert the cleaning cartridge correctly (if you do it wrong, the I/O station will not close properly. Do not force it).
3. Close the I/O station and press **YES**. The Tape Library will scan the I/O station for the cartridges and move them to an appropriate slot. The Tape Library displays the message "Insertion of Cleaning Cartridges has completed".
4. Press **Enter** to return to Manual Operations menu, and **Back** until you return to Activity Screen.

Note: A cleaning cartridge are normally unassigned.

4.3 Setting up the TS7700 Virtualization Engine

This section uses the TS7700 management interface (MI) to continue with the TS7700 Virtualization subsystem setup. If you have a TS 7720, you might skip the entire session in 4.2, "TS3500 Tape Library definitions (TS7740 Virtualization Engine only)" on page 178 and start here.

In this section, we describe the definitions and settings to apply to the TS7700 Virtualization Engine. The major tasks are as follows:

- ▶ Definitions that are made by the IBM System Service Representative (SSR) during the installation of the TS7700 Virtualization Engine, at your request
- ▶ Definitions that are made through the TS7740 Virtualization Engine MI
- ▶ Insertion of logical volumes through the TS7740 Virtualization Engine MI

4.3.1 TS7700 Virtualization Engine definition with the management interface

The TS7700 Virtualization Engine management interface (MI) is a web-based user interface to the TS7700 for customer information and management. It is accessed through any standard web browser, using the TS7700 Virtualization Engine IP address. During the installation process, the service representative set up TCP/IP on the TS7700 to use the customer-assigned TCP/IP host name and TCP/IP address.

Using the MI you can perform the following functions:

- ▶ Monitor the status of the TS7700 functions and hardware components.
- ▶ Monitor the performance of the TS7700 and Grid.
- ▶ Manage the TS7700 logical volumes.
- ▶ Configure the TS7700 and Grid.
- ▶ Manage the operation of the TS7700 and Grid.
- ▶ Manage user access to the TS7700.
- ▶ Access the 3953 or 3494 Library Manager web interface.
- ▶ Access the TS7700 Information Center.

To log on to the TS7700 Virtualization Engine management interface, enter a user ID (the default is admin) and a password (the default is admin). You may change these values

through the MI at any time. Using the admin/admin login, you can add and administer other users.

Figure 4-19 shows the Grid Summary window after logging into a four-cluster grid of a TS7700 Virtualization Engine.

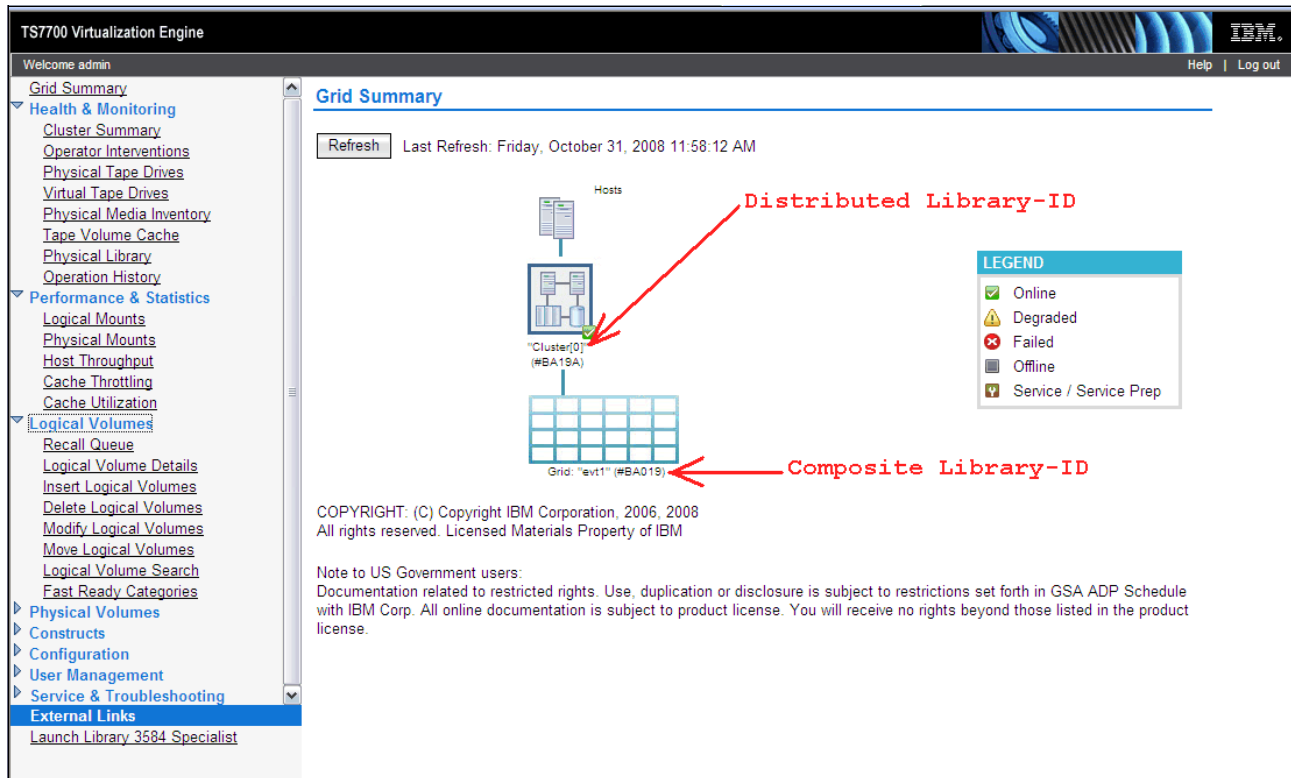


Figure 4-19 Welcome window for a single cluster

As Figure 4-19 shows, even if you are accessing a single-cluster grid TS7700 Virtualization Engine, it always displays as a Grid, which in this case is a single-cluster grid.

Compare the Grid in Figure 4-19 on page 199 with the one shown in Figure 4-20, which shows a four-cluster grid.

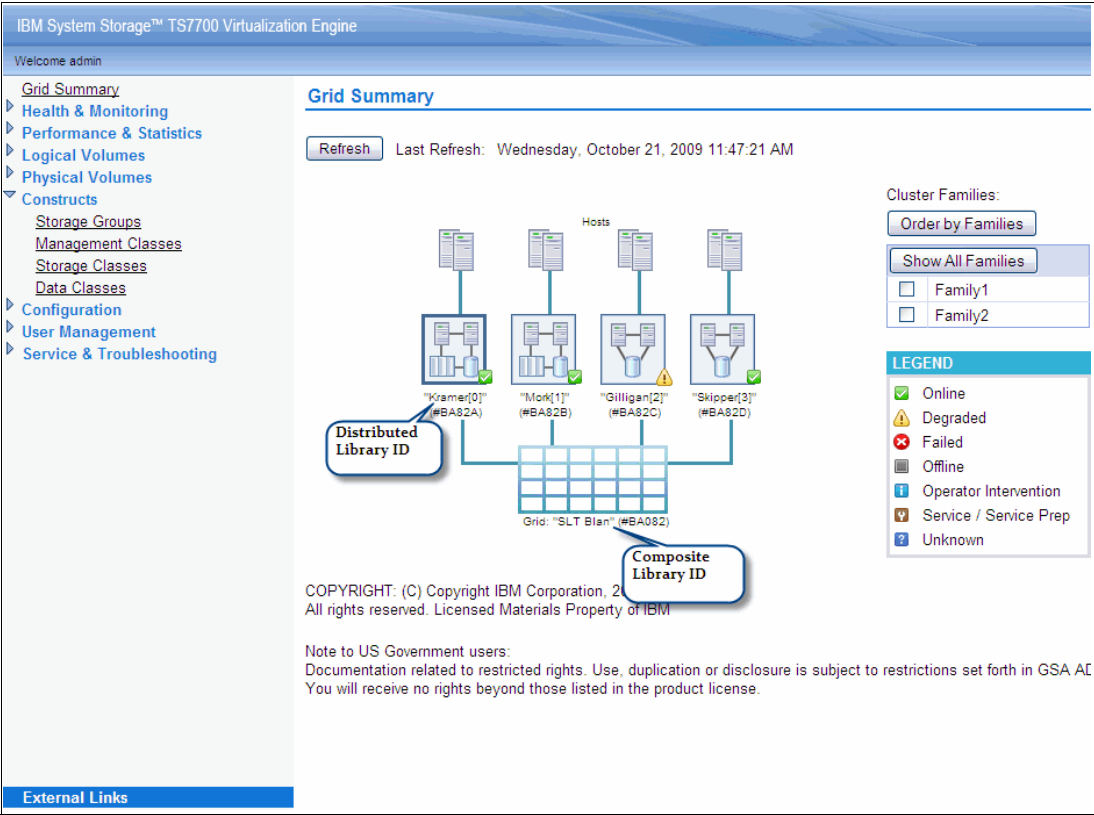


Figure 4-20 Four-cluster grid summary

Observe that Composite Library Sequence Number, Distributed Library sequence number, and Cluster Number are shown in this window, no matter whether it is a single cluster or multi-cluster configuration.

Composite and Distributed Library sequence number, and Cluster number values are defined to each TS7700 cluster by the IBM System Service Representative (SSR) during hardware installation.

Note: The Composite and Distributed Library sequence numbers must use hexadecimal characters.

Name your grid and cluster (or clusters) by using the management interface as described in this section. Be sure to use meaningful names to make resource identification as easy as possible to anyone who might be managing or monitoring this grid through the MI.

Note: A best practice is to make the grid name the same as the composite library name that was defined through DFSMS.

To set the grid and cluster names, access the TS7700 management interface. Go to **Configuration** → **Grid Identification Properties**. See Figure 4-21 on page 201 for the Grid Identification Properties window. Enter the appropriate information as required.

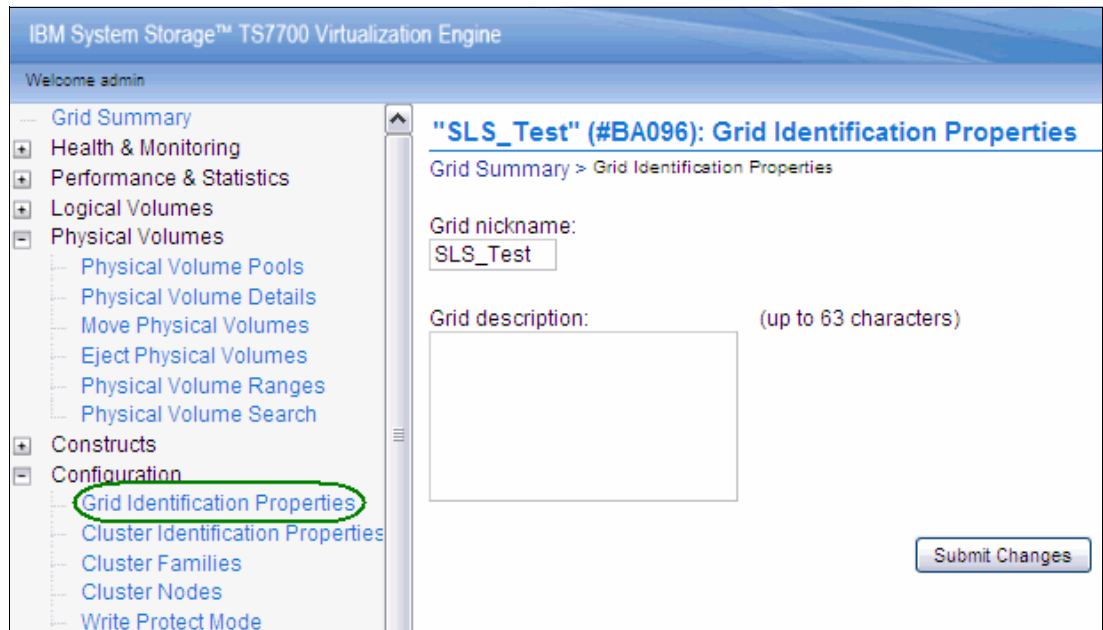


Figure 4-21 Grid identification Properties

To set the Cluster or Cluster Names, go to **Configuration** → **Cluster Identification Properties**. See Figure 4-22 on page 202 for the Cluster Identification Properties window.

Note: We suggest using the same denomination in cluster nickname as in DFSMS distributed library for the same cluster.

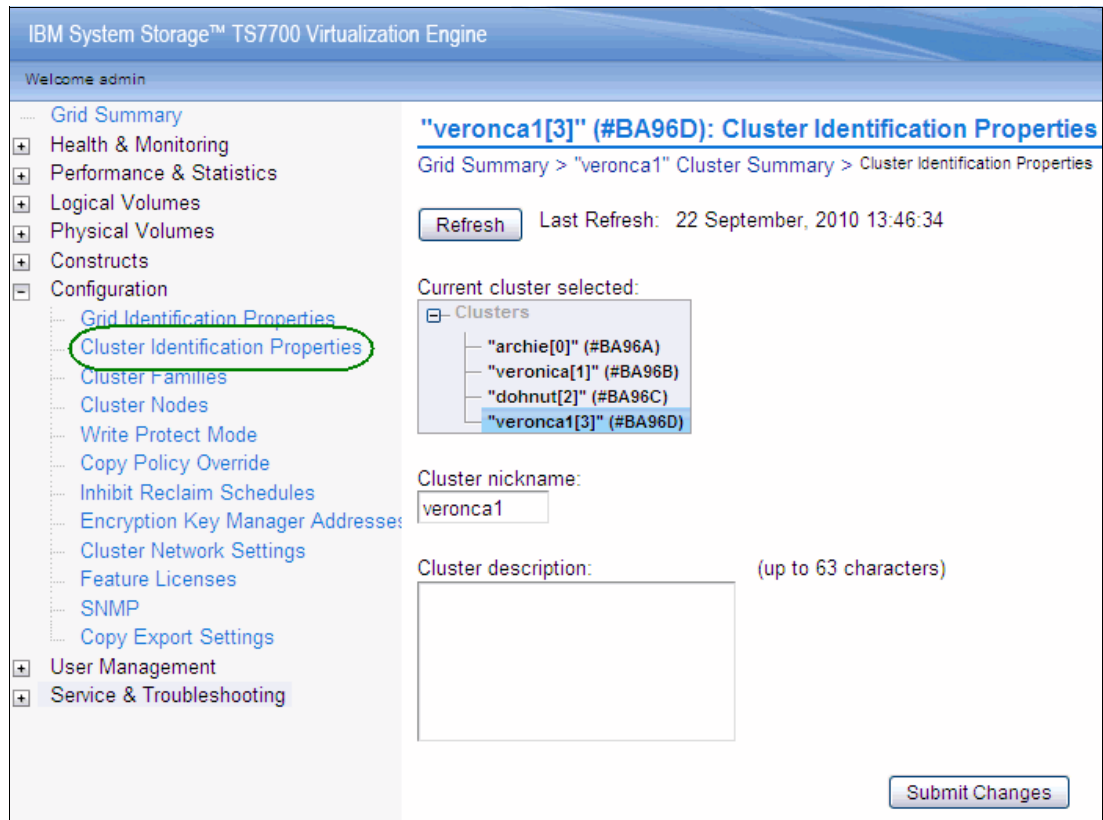


Figure 4-22 Cluster Identification Properties

Both cluster or grid nickname must be one to eight characters in length composed of alphanumeric characters. Blank spaces and the following special characters are also allowed:

@ . - +

Blank spaces cannot be used in first or last character positions.

In the **Grid description** or **Cluster description** field, a brief description (up to 63 characters) is enough.

4.3.2 Verifying the composite and distributed library sequence numbers

You must ensure that distributed and composite libraries sequence numbers are unique within the Grid.

Both composite and distributed library sequence numbers are set into each TS7700 Virtualization Engine Cluster by the IBM System Service Representative during installation. Composite library sequence number defined into TS7700 must match HCD LIBRARY-ID definitions and TS7700's Distributed Library sequence number must match to LIBRARY-ID listed in the ISMF Tape Library definition panels.

As explained before, even a single-cluster grid (stand-alone) TS7700 Virtualization Engine requires the definition of a Composite and a Distributed Library ID (visualize a Grid compounded by a lone cluster). The Composite Library ID is used for the host definition of the TS7700 Virtualization Engine logical library; the Distributed Library ID is used to link to the hardware aspects of the TS7700 Virtualization Engine, such as displaying scratch stacked volumes.

Check the Distributed and Composite library sequence numbers in the Grid Summary screen of the management interface, as shown in Figure 4-23.

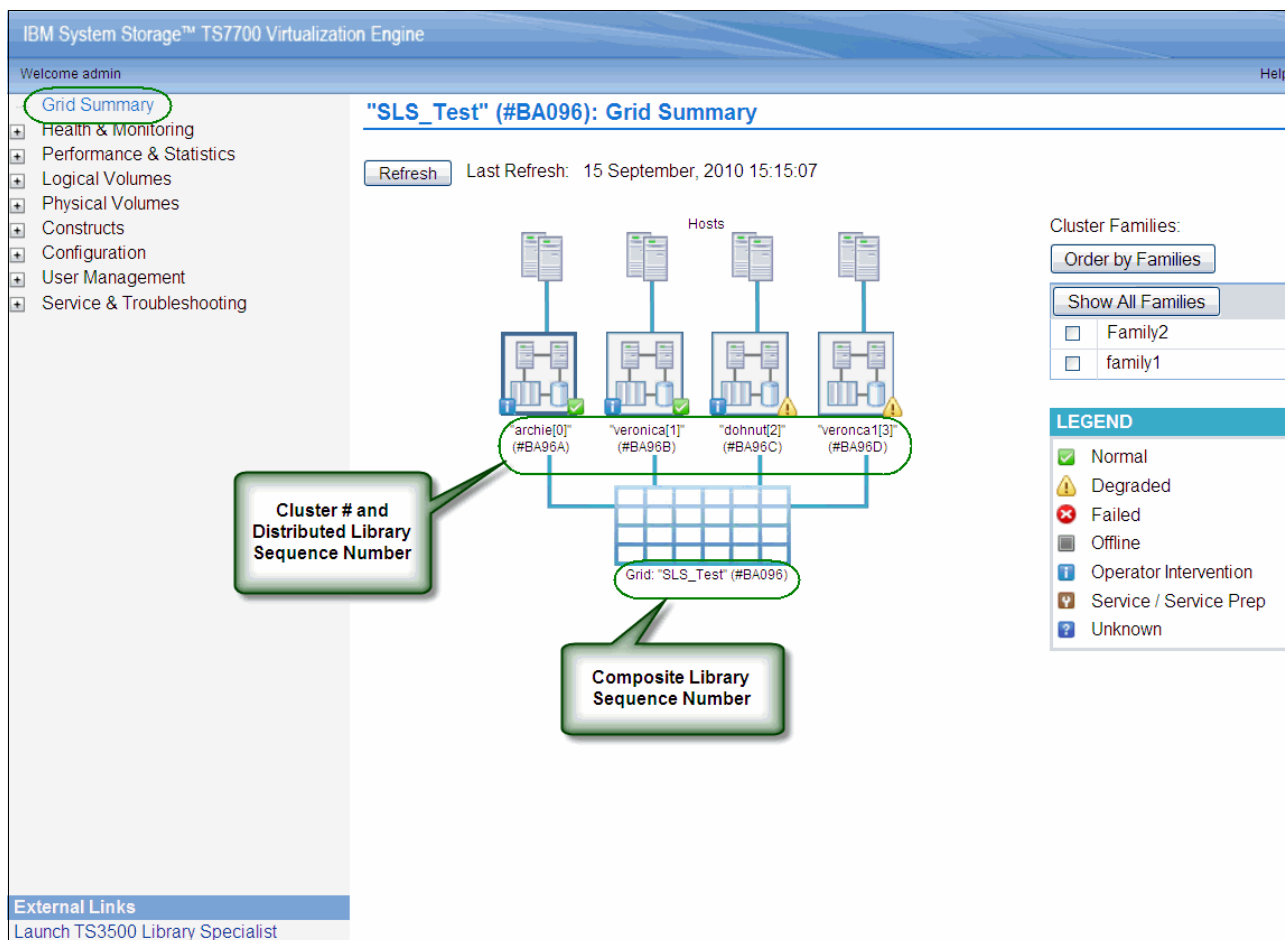


Figure 4-23 Composite and Distributed Library sequence numbers

Note: Do not use distributed library names starting with the letter V, because on the z/OS host, a library name cannot start with the letter V.

4.3.3 Defining VOLSER ranges for physical (stacked) volumes

After a cartridge is assigned to a logical library that is associated to a TS7740 Virtualization Engine by Cartridge Assignment Policies (CAP), it will be presented to the TS7740 Virtualization Engine Integrated Library Manager. Integrated Library Manager uses the VOLSER ranges that are defined in its VOLSER Ranges Table to set it to a proper Library Manager category. Define the proper policies in the VOLSER Ranges Table *before* inserting the cartridges into the Tape Library.

Important:

- ▶ When using a TS3500 Tape Library, you must assign CAP at the library hardware level before using the library with System z hosts.
- ▶ When using a TS3500 Tape Library, the TS7740 Virtualization Engine, physical volumes must fall within ranges that are assigned by CAP to this TS7740 Virtualization Engine Logical Library in the TS3500 Tape Library.

Use the window shown in Figure 4-24 to add, modify, or delete physical volume ranges. Also, unassigned physical volumes can be observed in this window. When you observe an unassigned volume that belongs to this TS7740 Virtualization Engine, add a range that includes that volume to fix it. If an unassigned volume does not belong to this TS7740 Virtualization Engine, you must eject it and reassign it to the proper logical library in TS3500 Tape Library. In this case, re-check CAP for proper range definition in TS7700 Virtualization Engine.

The screenshot displays the 'Physical Volume Ranges' window in the IBM System Storage TS7700 Virtualization Engine. The left sidebar contains a tree view with 'Physical Volume Ranges' highlighted. The main area shows the 'dohnut[2]' (#BA96C) cluster summary. Below this, there is a 'Refresh' button and a 'Last Refresh' timestamp. A section titled 'Current cluster selected:' lists four clusters: 'archie[0]' (#BA96A), 'veronica[1]' (#BA96B), 'dohnut[2]' (#BA96C), and 'veronica[3]' (#BA96D). Below this is an 'Inventory Upload' button. The 'Physical Volume Ranges' section includes a 'Select Action' dropdown and a 'Go' button. Below this is a table with columns 'Start Volser', 'End Volser', 'Media Type', and 'Home Pool'. The 'Unassigned Volumes' section also includes a 'Select Action' dropdown and a 'Go' button. Below this is a table with columns 'Select', 'Volser', 'Media Type', and 'Pending Eject'. The table shows 8 unassigned volumes, with the first 7 having 'JA-ETC' media type and 'No' pending eject status. The 8th volume has 'JA-ETC' media type and 'No' pending eject status. The bottom of the window shows 'Page 1 of 14', '1' rows, 'Go', 'Rows: 8', 'Show', and 'Total: 111 Displayed: 8'.

Figure 4-24 Physical Volume Ranges

Click the **Inventory Upload** button to upload the inventory from TS3500 and update any range or ranges of physical volumes that were recently assigned to that logical library.

The VOLSER Ranges Table displays the list of defined VOLSER ranges for a given component.

You can use the VOLSER Ranges Table to create a new VOLSER range or modify or delete a predefined VOLSER range.

Note: Operator intervention will be required to resolve unassigned volumes.

Figure 4-24 on page 204 shows status information that is displayed in the VOLSER Ranges Table as follows:

- ▶ From: The first VOLSER in a defined range
- ▶ To: The last VOLSER in a defined range
- ▶ Media Type: The media type for all volumes in a given VOLSER range. Possible values are as follows:
 - JA-ETC: Enterprise Tape Cartridge
 - JB-ETCL: Enterprise Extended-Length Tape Cartridge
 - JJ-EETC: Enterprise Economy Tape Cartridge
- ▶ Home Pool: The home pool to which the VOLSER range is assigned

Use the drop-down menu in the VOLSER Ranges Table to add a new VOLSER range or modify or delete a predefined range:

- ▶ To add a new VOLSER range, select **Add** from the drop-down menu. Complete the fields for information that will be displayed in the VOLSER Ranges Table, as defined previously.
- ▶ To modify a predefined VOLSER range, click the radio button from the **Select** column that appears in the same row as the name of the VOLSER range you want to modify. Select **Modify** from the drop-down menu and make your changes to the information that will be displayed in the VOLSER Ranges Table.

The VOLSER entry fields must contain six characters. The characters can be letters, numerals, or a space. The two VOLSERs must be entered in the same format. *Corresponding characters* in each VOLSER must both be either alphabetic or numeric. For example, AAA998 and AAB004 are of the same form, but AA9998 and AAB004 are not. The VOLSERs that fall within a range are determined as follows: The VOLSER range is increased such that alphabetic characters are increased alphabetically, and numeric characters are increased numerically. For example, VOLSER range ABC000 - ABD999 result in a range of 2000 VOLSERs (ABC000 - ABC999 and ABD000 - ABD999).

Note: The VOLSER ranges you define on the IBM TS3500 Tape Library apply to physical cartridges only. You define logical volumes only through the TS7700 Virtualization Engine management interface. See 4.3.12, “Inserting logical volumes” on page 240 for more information.

For the TS7700 Virtualization Engine, no additional definitions are required at the hardware level other than setting up the proper VOLSER ranges at the TS3500 library.

Note: Although you could now enter cartridges into the TS3500 library, complete the required definitions at the host before you insert any physical cartridges into the Tape Library.

In previous generations of the TS7700 Virtualization Engine, as with the IBM TotalStorage Virtual Tape Server (VTS), the logical volumes were defined and inserted through the Library Manager console or the IBM ETL Specialist.

This process has changed for the TS7700 Virtualization Engine; this task is now performed through the TS7700 Virtualization Engine MI, which is described in detail in 4.3.12, “Inserting logical volumes” on page 240.

4.3.4 Defining physical volume pools

Physical Volumes pooling was first introduced as part of advanced policy management advanced functions in the IBM TotalStorage Virtual Tape Server (VTS).

Pooling physical volume allows you to separate your data into separate sets of physical media, treating each media group in a specific way. For instance, you might want to segregate production data from test data, or encrypt part of your data, or yet have data in different media type, and more. All this can be accomplished defining physical volume pools appropriately. Also, you can define the reclaim parameters for each specific pool to best suit specific needs. The TS7700 Virtualization Engine management interface is used for pool property definitions.

Note: Items under Physical Volumes in the Management interface only apply to clusters with an associated Tape Library (TS7740 Virtualization Engine). Trying to access those screens from a TS7720 will result in the following HYDME0995E message:

This cluster is not attached to a physical tape library.

Use the window shown in Figure 4-25 to view or modify settings for physical volume pools, which manage the physical volumes used by the TS7700 Virtualization Engine.

The screenshot displays the 'Physical Volume Pools' configuration page for cluster 'Juno[0]' (#BA41A). The interface includes a sidebar with navigation links such as 'Grid Summary', 'Health & Monitoring', 'Performance & Statistics', 'Logical Volumes', 'Physical Volumes', 'Constructs', 'Configuration', 'User Management', and 'Service & Troubleshooting'. The main content area shows the cluster name, a 'Refresh' button, and a 'View tutorial' link. Below this, the 'Current component selected' section lists three clusters: 'Juno[0]' (#BA41A), 'diana[1]' (#BA41B), and 'herc[2]' (#BA41C). The 'Physical volume pool properties' section features two tabs: 'Pool Properties' and 'Encryption Settings'. The 'Pool Properties' tab is active, displaying a table with columns for 'Select', 'Pool', 'Me', 'Fir', 'Se', 'Bor', 'Re', 'Ma', 'Da', 'Da', 'Ag', 'Da', and 'Ma'. The table contains 9 rows of data, each representing a physical volume pool.

Select	Pool	Me	Fir	Se	Bor	Re	Ma	Da	Da	Ag	Da	Ma
<input type="checkbox"/>	1	3592	Any...	None	Borr...	11	All ...	0	0	1	0	0
<input type="checkbox"/>	2	3592	JA...	JJ-E...	Borr...	12	All ...	0	1	0	0	0
<input type="checkbox"/>	3	3592	Any...	None	Borr...	13	All ...	0	0	0	1	95
<input type="checkbox"/>	4	3592	Any...	None	Borr...	4	All ...	0	0	0	0	0
<input type="checkbox"/>	5	3592	Any...	None	Borr...	5	All ...	0	0	0	0	0
<input type="checkbox"/>	6	3592	Any...	None	Borr...	6	All ...	0	0	0	0	0
<input type="checkbox"/>	7	3592	Any...	None	Borr...	7	All ...	0	0	0	0	0
<input type="checkbox"/>	8	3592	Any...	None	Borr...	8	All ...	0	0	0	0	0
<input type="checkbox"/>	9	3592	Any...	None	Borr...	9	All ...	0	0	0	0	0

Figure 4-25 Physical Volume Pools

The Physical Volume Pool Properties table displays the encryption setting and media properties for every physical volume pool defined for a given cluster in the Grid.

You may use the Physical Volume Pool Properties table to view encryption and media settings for all installed physical volume pools. To view and modify additional details of pool properties, select a pool or pools from this table and then click either **Pool Properties** or **Encryption Settings** from the drop-down menu.

Note: Pools 1 - 32 are preinstalled. Pool 1 functions as the default pool and is used if no other pool is selected. All other pools must be defined before they can be selected.

The Physical Volume Pool Properties table displays the media properties and encryption settings, and for every physical volume pool defined for a given cluster in the Grid. This table contains two tabs: Pool Properties and Encryption Settings.

► Under Pool Properties tab:

- Pool: Lists the pool number, which is a whole number in the range of 1 - 32, inclusive.
- Media Class: Lists the supported media class of the storage pool is 3592.
- First Media (Primary): The primary media type that the pool can borrow or return to the common scratch pool (Pool 0). Possible values are as follows:

Any 3592	Any 3592
JA	Enterprise Tape Cartridge (ETC)
JB	Enterprise Extended-Length Tape Cartridge (ETCL)
JJ	Enterprise Economy Tape Cartridge (EETC)

To modify pool properties, select the check box next to one or more pools listed in the Physical Volume Pool Properties table and select **Properties** from the drop-down menu. The Pool Properties Table will be displayed. You will be able to modify the fields Media Class and First Media, defined previously, and the following fields:

- Second Media (Secondary): Lists the second choice of media type that the pool can borrow from. Options listed exclude the media type selected for First Media. Possible values are as follows:

Any 3592	Any 3592
JA	Enterprise Tape Cartridge (ETC)
JB	Enterprise Extended-Length Tape Cartridge (ETCL)
JJ	Enterprise Economy Tape Cartridge (EETC)
None	The only option available if the Primary Media type is Any 3592.

- Borrow Indicator: Defines how the pool is populated with scratch cartridges. Possible values are as follows:

Borrow, Return	A cartridge is borrowed from the Common Scratch Pool (CSP) and returned when emptied.
Borrow, Keep	A cartridge is borrowed from the CSP and retained, even after being emptied.
No Borrow, Return	A cartridge may not be borrowed from CSP, but an emptied cartridge is placed in CSP. This setting is used for an empty pool.
No Borrow, Keep	A cartridge may not be borrowed from CSP, and an emptied cartridge will be retained.

- Reclaim Pool: Lists the pool to which logical volumes are assigned when reclamation occurs for the stacked volume on the selected pool.
- Maximum Devices: Lists the maximum number of physical tape drives that the pool can use for premigration.
- Export Pool: Lists the type of export that is supported if the pool is defined as an Export Pool, which is the pool from which physical volumes are exported. (From the Physical

Volume Pools page, click the Pool Properties tab; select the check box next to each pool to be modified; select **Modify Pool Properties** from the Physical volume pools drop-down menu; and click **Go** to open the Modify Pool Properties page.) Possible values are as follows:

Not Defined The pool is not defined as an Export pool

Copy Export The pool is defined as a Copy Export pool.

- Days Before Secure Data Erase: Lists the number of days a physical volume that is a candidate for Secure Data Erase can remain in the pool without access to a physical stacked volume. Each stacked physical volume possesses a timer for this purpose, which is reset when a logical volume on the stacked physical volume is accessed. Secure Data Erase occurs at a later time, based on an internal schedule. Secure Data Erase renders all data on a physical stacked volume inaccessible. The valid range of possible values is 1 - 365. Clearing the check box deactivates this function.
- Days Without Access: Lists the number of days the pool can persist without access to set a physical stacked volume. Each physical stacked volume has a timer for this purpose, which is reset when a logical volume is accessed. The reclamation occurs at a later time, based on an internal schedule. The valid range of possible values is 1 - 365. Clearing the check box deactivates this function.
- Age of Last Data Written: Lists the number of days the pool has persisted without write access to set a physical stacked volume as a candidate for reclamation. Each physical stacked volume has a timer for this purpose, which is reset when a logical volume is accessed. The reclamation occurs at a later time, based on an internal schedule. The valid range of possible values is 1 - 365. Clearing the check box deactivates this function.
- Days Without Data Inactivation: Lists the number of sequential days the pool's data ratio has been higher than the Maximum Active Data to set a physical stacked volume as a candidate for reclamation. Each physical stacked volume has a timer for this purpose, which is reset when data inactivation occurs. The reclamation occurs at a later time, based on an internal schedule. The valid range of possible values is 1-365. Clearing the check box will deactivate this function. If deactivated, this field is not used as a criteria for reclamation.
- Maximum Active Data: Lists the ratio of the amount of active data in the entire physical stacked volume capacity. This field is used in conjunction with Days Without Data Inactivation. The valid range of possible values is 5 - 95(%). This function is disabled if Days Without Data Inactivation is not selected.
- Reclaim Threshold: Lists the percentage that is used to determine when to perform reclamation of free storage on a stacked volume. When the amount of active data on a physical stacked volume drops below this percentage, a reclaim operation will be performed on the stacked volume. The valid range of possible values is 5 - 95(%). The default value is 10%. Clearing the check box deactivates this function.
- Under Encryption Settings tab:
 - Pool: Lists the pool number. This number is a whole number 1 - 32, inclusive.
 - Encryption: Lists the encryption state of the pool. Possible values are Enabled or Disabled.
 - Key Mode 1: Lists Encryption Mode used with Key Label 1. Possible values for this field are as follows:
 - Clear Label: The data key is specified by the key label in clear text.
 - Hash Label: The data key is referenced by a computed value corresponding to its associated public key.

- None: Key Label 1 is disabled.
- Dash (-): The default key is in use.
- Key Label 1: Lists the current encryption key Label 1 for the pool. The label must consist of ASCII characters and cannot exceed 64 characters. Leading and trailing blanks are removed, but an internal space is allowed. Lowercase characters are internally converted to uppercase upon storage, thus, key labels are reported using uppercase characters.

If the encryption state indicates Disabled, this field is blank. If the default key is used, the value in this field is “default key.”

- Key mode 2: Lists the Encryption Mode used with Key Label 2. Possible values for this field are as follows:
 - Clear Label: The data key is specified by the key label in clear text.
 - Hash Label: The data key is referenced by a computed value corresponding to its associated public key.
 - None: Key Label 2 is disabled.
 - Dash (-): The default key is in use.
- Key Label 2: The current encryption key Label 2 for the pool. The label must consist of ASCII characters and cannot exceed 64 characters. Leading and trailing blanks are removed, but an internal space is allowed. Lowercase characters are internally converted to uppercase upon storage, thus, key labels are reported using uppercase characters.

If the encryption state is “Disabled”, this field is blank. If the default key is used, the value in this field is “default key”.

To modify encryption settings for one or more physical volume pools use the following steps (and see Figure 4-26 and Figure 4-27 on page 211 for reference):

Note: You can click on the provided link to watch a tutorial showing how to modify encryption properties.

1. Open the Physical Volume Pools page (Figure 4-26)




Figure 4-26 Modifying encryption parameters for a pool

2. Perform the following steps:
 - a. Click the **Encryption Settings** tab.
 - b. Select the check box next to each pool to be modified.
 - c. Click **Select Action** → **Modify Encryption Settings**.
 - d. Click **Go**.

The Modify Encryption Settings page opens (Figure 4-27 on page 211).




Figure 4-27 Modify encryption settings parameters

In this window, you may modify values for any of the following controls:

► Encryption:

This field is the encryption state of the pool and can have the following values:

- Enabled: Encryption is enabled on the pool.
- Disabled: Encryption is not enabled on the pool.

When this value is selected, key modes, key labels, and check boxes are disabled.

► Use Encryption Key Manager default key

Select this check box to populate the Key Label field by using a default key provided by the encryption key manager.

Note: Your encryption key manager software must support default keys to use this option.

This check box occurs prior to both Key Label 1 and Key Label 2 fields; you must select this check box for each label to be defined using the default key.

If this check box is selected, the following fields are disabled:

- Key Mode 1
- Key Label 1
- Key Mode 2
- Key Label 2

► Key Mode 1

This field is the Encryption Mode that is used with Key Label 1. It can have the following possible values:

- Clear Label: The data key is specified by the key label in clear text.
- Hash Label: The data key is referenced by a computed value corresponding to its associated public key.
- None: Key Label 1 is disabled. The default key is in use.

► Key Label 1

This field is the current encryption key Label 1 for the pool. The label must consist of ASCII characters and cannot exceed 64 characters. Leading and trailing blanks are removed, but an internal space is allowed. Lowercase characters are internally converted to uppercase upon storage, thus, key labels are reported using uppercase characters.

► Key Mode 2

This field is the Encryption Mode used with Key Label 2. Possible values for this field are as follows:

- Clear Label: The data key is specified by the key label in clear text.
- Hash Label: The data key is referenced by a computed value that corresponds to its associated public key.

► None

Indicates that the Key Label 2 is disabled. The default key is in use.

► Key Label 2

This field is the current encryption key Label 2 for the pool. The label must consist of ASCII characters and cannot exceed 64 characters. Leading and trailing blanks are removed, but an internal space is allowed. Lowercase characters are internally converted to uppercase upon storage, thus, key labels are reported using uppercase characters.

To complete the operation click **OK**. To abandon the operation and return to the Physical Volume Pools page, click **Cancel**.

Reclaim thresholds

To optimize utilization of the subsystem resources, such as CPU cycles and tape drive usage, you can inhibit space reclamation during predictable busy periods of time and adjust reclamation thresholds to the optimum point in your TS7740 through the management interface. *Reclaim threshold* is the percentage used to determine when to perform reclamation of free space in a stacked volume. When the amount of active data on a physical stacked volume drops below this percentage, this volume becomes eligible to reclamation. Reclamation values can be in the range of 5 - 95%; default value is 10%. Clearing the check box deactivates this function.

Throughout data life cycle, new logical volumes are created and old logical volumes become obsolete. Logical volumes are migrated to physical volumes, occupying real space there. When a logical volumes becomes obsolete, that space becomes a waste of capacity in that physical tape. In other words, the active data level in that volume goes lower and lower.

TS7740 actively monitor the active data in its physical volumes. Whenever this active data level crosses the reclaim threshold that is defined in the Physical Volume Pool in which that volume belongs, the TS7400 sets that volume in a candidate list for reclamation.

Reclamation will copy active data from that volume to another stacked volume in the same pool. When the copy finishes and volume became empty, it will be returned to available scratch status. This cartridge is now available for use and will be returned to CSP or directed to the specified reclaim pool, according to the Physical Volume Pool definition.

Note: Each reclamation task uses two tape drives (source and target) in a tape-to-tape copy function. TS7740 cache (TVC) is not used for reclamation.

Multiple reclamation processes can run in parallel. The maximum number of reclaim tasks is limited by the TS7740, based on the number of available back-end drives, as shown in the Table 4-1

Table 4-1 Installed drives versus maximum reclaim tasks

Number of available drives	Maximum number of reclaims
3	1
4	1
5	1
6	2
7	2
8	3
9	3
10	4
11	4
12	5
13	5
14	6
15	6
16	7

The reclamation level for the physical volumes must be set using the Physical Volume Pools window in the TS7740 management interface. See Figure 4-28.

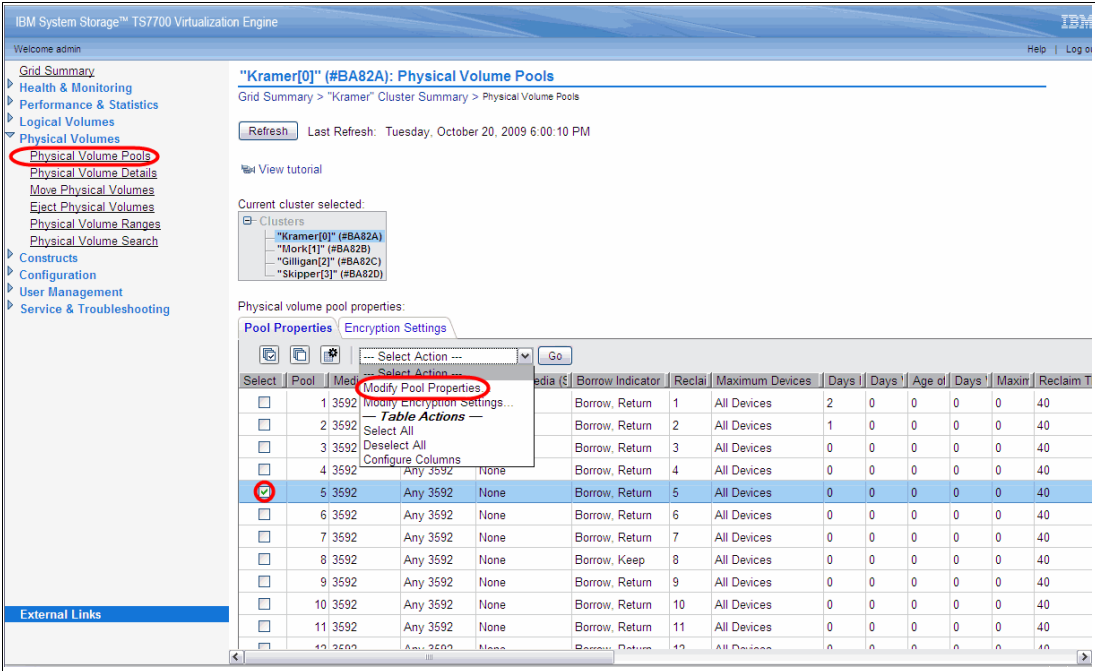


Figure 4-28 Physical Volume Pools

Select a pool and click **Modify Pool Properties** in the drop-down menu to set reclamation level and other policies for that pool. See Figure 4-29 on page 215.

In this example, reclamation level is set to 40%, borrow-return policy is in effect for this pool, and reclaimed physical cartridges stay in the same Pool 5, except borrowed volumes, which will be returned to the original pool.

IBM System Storage™ TS7700 Virtualization Engine

Welcome admin

[Grid Summary](#)

▶ [Health & Monitoring](#)

▶ [Performance & Statistics](#)

▶ [Logical Volumes](#)

▼ [Physical Volumes](#)

[Physical Volume Pools](#)

[Physical Volume Details](#)

[Move Physical Volumes](#)

[Eject Physical Volumes](#)

[Physical Volume Ranges](#)

[Physical Volume Search](#)

▶ [Constructs](#)

▶ [Configuration](#)

▶ [User Management](#)

▶ [Service & Troubleshooting](#)

[External Links](#)

"Kramer[0]" (#BA82A): Modify Pool Properties

Last Refresh: Tuesday, October 20, 2009 6:05:17 PM

Selected pools:

Pool
5

Pool Properties:

Media Class:	3592 ▼
First Media (Primary):	Any 3592 ▼
Second Media (Secondary):	None ▼
Borrow Indicator:	Borrow, Return ▼
Reclaim Pool:	5 ▼
Maximum Devices:	All Devices ▼
Export Pool:	Not Defined ▼
<input type="checkbox"/> Days Before Secure Data Erase:	0
<input type="checkbox"/> Days Without Access:	0
<input type="checkbox"/> Age of Last Data Written:	0
<input type="checkbox"/> Days Without Data Inactivation:	0
Maximum Active Data:	5% ▼
Reclaim Threshold Percentage:	40% ▼

Figure 4-29 Pool properties

Reclamation enablement

To minimize any impact on TS7700 Virtualization Engine activity, the storage management software monitors resource utilization in the TS7700 Virtualization Engine, and enables or disables reclamation as appropriate. You can optionally prevent reclamation activity at specific times of day by specifying an Inhibit Reclaim Schedule in the TS7740 Virtualization Engine management interface (Figure 4-30 on page 218 shows an example). However, the TS7740 Virtualization Engine determines whether reclamation is to be enabled or disabled once an hour depending on the number of available scratch cartridges and will ignore the Inhibit Reclaim Schedule, if the TS7740 Virtualization Engine is close to running out of scratch category.

Note: The maximum number of inhibit reclaim schedules is 14.

Using the Bulk Volume Information Retrieval (BVIR) process, you can run the query for PHYSICAL MEDIA POOLS to monitor the amount of active data on stacked volumes to help you plan for a reasonable and effective reclaim threshold percentage. You can also use the Host Console Request function to obtain the physical volume counts.

Although reclamation is enabled, stacked volumes going might not always be going through the process all the time. Other conditions must be met, such as stacked volumes that meet one of the reclaim policies and drives available to mount the stacked volumes.

Reclamation for a volume is stopped by the TS7700 Virtualization Engine internal management functions if a tape drive is needed for a recall or copy (because these are of a higher priority) or a logical volume is needed for recall off a source or target tape being used in the reclaim process. If this happens, reclamation is stopped for this physical tape after the current logical volume move is complete.

Reclamation is enabled or disabled according to a set of rules. To enable the TS7700 Virtualization Engine to compensate for the pooling structure and scratch thresholds these rules are required. Pooling is enabled as a standard feature of the TS7700 Virtualization Engine, even if you are only using one pool. Reclamation can occur on multiple volume pools at the same time, and processing multiple tasks for the same pool. One of the reclamation methods selects the volumes for processing based on the percentage of active data. For example, if the reclaim threshold was set to 30% generically across all volume pools, the TS7700 Virtualization Engine would select all the stacked volumes from 0% to 29% of remaining active data. The reclaim tasks would then process the volumes from least full (0%) to most full (29%) up to the defined reclaim threshold of 30%.

Individual pools can have different reclaim policies set. The number of pools can also influence the reclamation process, because the Library Manager always evaluates the stacked media starting with Pool 1.

The scratch count for physical cartridges also affects reclamation. The *scratch state* of pools is assessed as follows:

1. A pool enters a *Low scratch state* when it has access to less than 50 but two or more stacked volumes.
2. A pool enters a *Panic scratch state* when it has access to less than two empty stacked volumes.

Access to includes any borrowing capability, which means that if the pool is configured for borrowing, and if there are more than 50 cartridges in the Common Scratch Pool (CSP), the pool will not enter the *Low scratch state*.

Whether borrowing is configured or not, as long each pool has two scratch cartridges, the Panic Reclamation mode is not entered. Panic Reclamation mode is entered when a pool has less than two scratch cartridges and no more scratch cartridges can be borrowed from any other pool defined for borrowing. Borrowing is described in “Physical volume pooling” on page 63.

Note: Be aware that a Physical Volume Pool running out of scratch may stop mounts in the TS7400, impacting customer operation. Mistakes in Pool Configuration (media type, borrow and return, home pool, and so on) or some unusual operation with an empty Common Scratch Pool (CSP) may lead to this situation.

Consider that one reclaim task can consume two drives for the data move, and CPU cycles. When a reclamation starts, these drives are busy until the volume being reclaimed is empty. If you raise the reclamation threshold level too high, the result is larger amounts of data to be moved, with resultant penalty in resources that are needed for recalls and premigration. Default setting for the reclamation threshold level is 10%, and we suggest that you operate with reclamation threshold level in the range of 10 - 30%. Also see Chapter 8, “Operation” on page 409 to fine-tune this function considering your peak load using new host functions.

Pools in either scratch state (Low or Panic state) get priority for reclamation. Table 4-2 summarizes the thresholds.

Table 4-2 Reclamation priority table

Priority	Condition	Reclaim schedule honored	Active data threshold% honored	Number of concurrent reclaims	Comments
1	Pool in Panic scratch state	No	No	1, regardless of idle drives	
2	Priority move	Yes or No	No	1, regardless of idle drives	<p>If a volume is within 10 days of a Secure Data Erasure (SDE) and still has active data on it, it will be reclaimed at this priority. An SDE priority move will honor the inhibit reclaim schedule.</p> <p>For a LLbrary Manager operator initiated priority move, the option to honor the inhibit reclaim schedule is given to the operator.</p>
3	Pool in Low scratch state	Yes	Yes	1, regardless of idle drives	Volumes that are subject to reclaim because of Maximum Active Data, Days Without Access, Age of Last Data Written, and Days Without Data Inactivation will use priority 3 or 4 reclamation.
4	Normal reclaim	Yes	Yes, pick from all eligible pools	(Number of idle drives divided by 2) minus 1 8 drv: 3 max 16 drv: 7 max	Volumes that are subject to reclaim because of Maximum Active Data, Days Without Access, Age of Last Data Written, and Days Without Data Inactivation will use priority 3 or 4 reclamation.

Notes:

- A physical drive is considered *idle* when no activity has occurred for the previous ten minutes.
- The Inhibit Reclaim Schedule is not honored by the Secure Data Erase function for a volume that has no active data.

Inhibit Reclaim Schedule

The Inhibit Reclaim Schedule defines when the TS7700 Virtualization Engine should refrain from reclaim operations. During times of heavy mount activity, it might be desirable to make all of the physical drives available for recall and premigration operations. If these periods of

heavy mount activity are predictable, you can use the Inhibit Reclaim Schedule to inhibit reclaim operations for the heavy mount activity periods.

To define the Inhibit Reclaim Schedule, select **Management Interface** → **Configuration**, which opens the window shown in Figure 4-30.

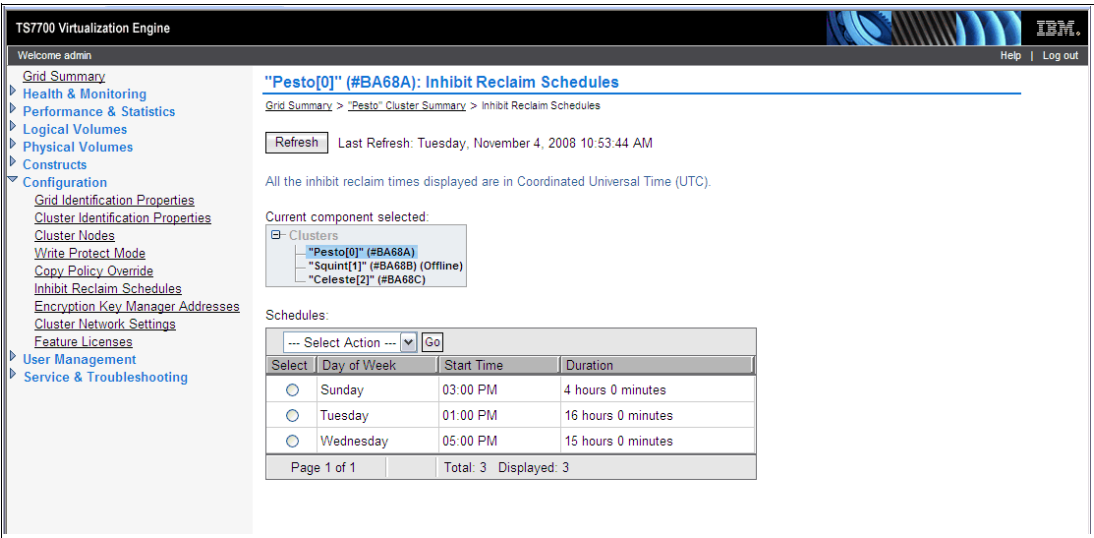


Figure 4-30 Inhibit Reclaim Schedules

The schedules table (Figure 4-31) displays the day, time, and duration of any scheduled reclamation interruption. All inhibit reclaim dates and times are first displayed in Coordinated Universal Time (UTC) and then in local time. Use the drop-down menu on the schedules table to add a new Reclaim Inhibit schedule or modify or delete an existing schedule, as shown in Figure 4-30.

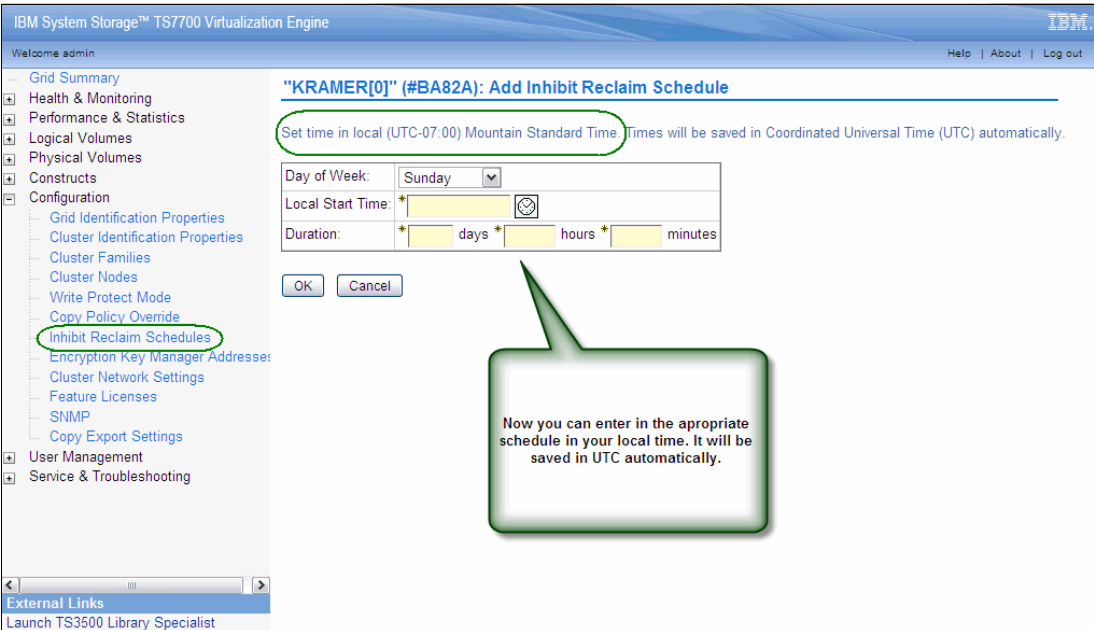


Figure 4-31 Add Inhibit Reclaim Schedule

Reclaim Threshold Percentage

Note: Inhibit Reclaim Schedules are *not* permitted to overlap.

The reclaim threshold is the percentage value used as a reference for the active data in the stacked volumes in one pool. Whenever the amount of active data on a physical stacked volume drops below this percentage, this cartridge becomes eligible to reclamation. See Figure 4-32 for reference. Each pool has its own reclaim threshold level setting.

IBM System Storage™ TS7700 Virtualization Engine

Welcome admin

Grid Summary

- Health & Monitoring
- Performance & Statistics
- Logical Volumes
- Physical Volumes
 - Physical Volume Pools**
 - Physical Volume Details
 - Move Physical Volumes
 - Eject Physical Volumes
 - Physical Volume Ranges
 - Physical Volume Search
- Constructs
- Configuration
- User Management
- Service & Troubleshooting

"dohnut[2]" (#BA96C): Modify Pool Properties

Refresh Last Refresh: 23 September, 2010 11:20:51

Selected pools:

Pool
1

Pool Properties:

Media Class:	3592
First Media (Primary):	Any 3592
Second Media (Secondary):	None
Borrow Indicator:	Borrow, Return
Reclaim Pool:	1
Maximum Devices:	All Devices
Export Pool:	Not Defined
<input type="checkbox"/> Days Before Secure Data Erase:	0
<input type="checkbox"/> Days Without Access:	0
<input type="checkbox"/> Age of Last Data Written:	0
<input type="checkbox"/> Days Without Data Inactivation:	0
Maximum Active Data:	5%
Reclaim Threshold Percentage:	35%

External Links

Launch TS3500 Library Specialist

OK Cancel

Figure 4-32 Reclaim Threshold Percentage

During the reclamation process, all active data from the original stacked volume is moved to another cartridge. After all active data has been moved, the original cartridge is made available for reuse, complying with the defined rules for that pool: whether returning to common scratch pool or going to the designated reclaim pool.

The Reclaim Threshold Percentage is initially set at 10%. This percentage directly affects the amount of data to be moved by the reclaim operation. As a general rule, try not to go above 30%.

You can set a Reclaim Threshold Percentage of 0% to prevent reclamation in this pool (assuming that all of the other reclamation criteria are also set to 0).

Be aware that a multiple of two drives are involved in the reclamation process and because of this resource usage, you should not specify high percentages. BVIR reports can help you adjust the percentages. See 9.10, “Bulk Volume Information Retrieval” on page 657 for details.

4.3.5 Defining Fast Ready categories

To take advantage of the scratch mount performance of the TS7700 Virtualization Engine and to prevent recalls for scratch mounts, you must assign the Fast Ready attribute to the categories used by the host for scratch volumes.

The MOUNT FROM CATEGORY command is not exclusively used for scratch mounts, therefore, the TS7700 Virtualization Engine cannot assume that any MOUNT FROM CATEGORY is for a scratch volume.

The Fast Ready attribute provides a definition of a category to supply scratch mounts and for z/OS it depends on the definitions. The Fast Ready definition is done through the management interface. Figure 4-33 on page 221 shows the Fast Ready Categories window.

The actual category hexadecimal number depends on the software environment and on the definitions in the SYS1.PARMLIB member DEVSUPxx for library partitioning. Also, the DEVSUPxx member must be referenced in IEASYSxx member to be activated.

Note: Any category up to and including the X'FEFF' may be overridden in DEVSUPxx, but be aware that categories may only be overridden if the TS7700 Virtualization Engine will not be accessed by the owning operating systems, which are described in Appendix G, “Library Manager volume categories” on page 913.

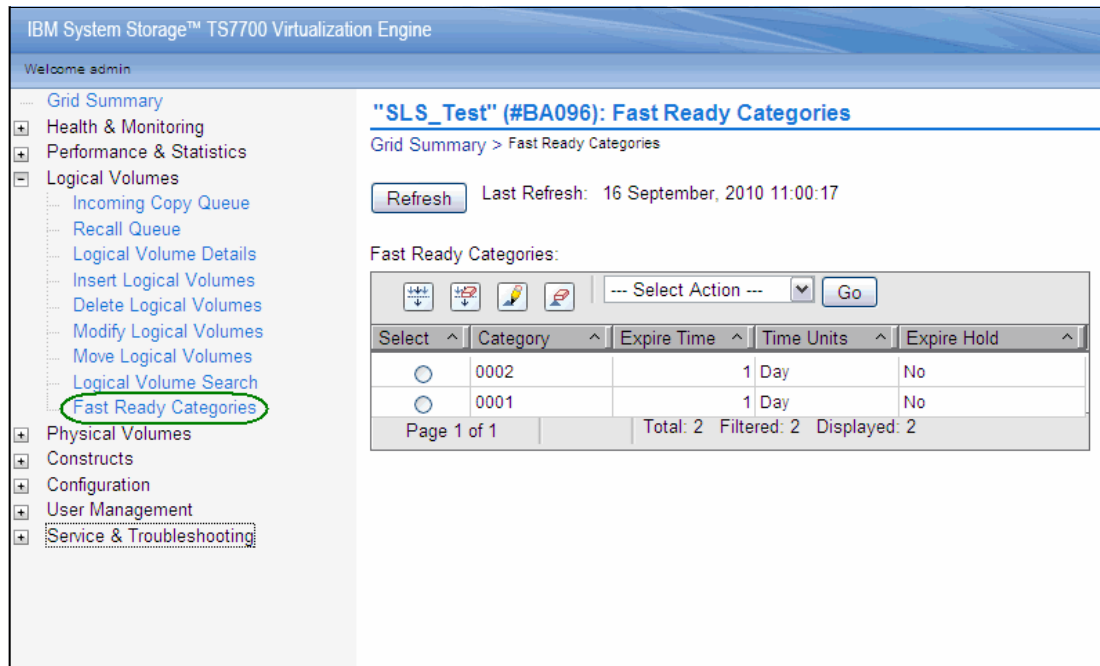


Figure 4-33 Fast Ready Categories

In Figure 4-33, the Fast Ready Categories Table lists all Fast Ready categories and their attributes. If no Fast Ready Categories are defined, this table is not visible.

Note: You may not use category name 0000 or FFxx (where xx equals 0 - 9 or A - F) because 0000 represents a null value, and “FFxx” is reserved for hardware.

Figure 4-34 shows the Add Category window (which opens by selecting **Add** in the Action drop-down menu shown in Figure 4-33).

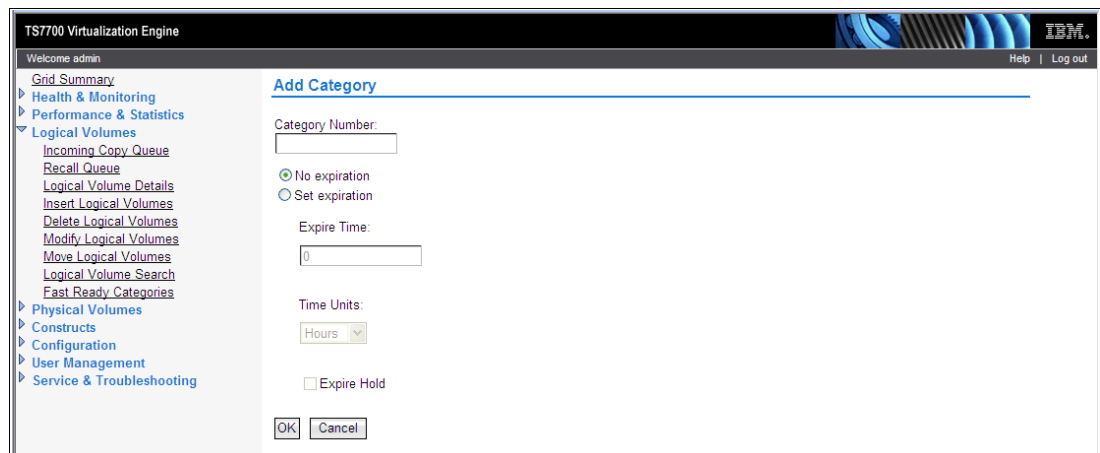


Figure 4-34 Fast Ready Categories: Add Category

When defining a Fast Ready category, also define whether volumes in this category will expire or not. If **Set expiration** is selected, define the expiration time for this category (00 will be rejected). Selecting the **Expire Hold** check box puts the non expired volumes in this category in “hold” state (meaning that they cannot be mounted or assigned to other

categories until expiration time is due). See 4.3.6, “Defining the logical volume expiration time” on page 222 for details concerning this aspect of Fast Ready categories.

Note: Add a comment to DEVSUPnn to make sure that the Library Manager Fast Ready categories are updated when the category values in DEVSUPnn are changed. They need to be in sync at all times.

See Appendix G, “Library Manager volume categories” on page 913 for the scratch mount category for each software platform. In addition to the z/OS DFSMS default value for the scratch mount category, you can define your own scratch category to the Library Manager. In this case, you should also add your own scratch mount category to the Fast Ready category list.

4.3.6 Defining the logical volume expiration time

You define the expiration time from the management interface window shown in Figure 4-34 on page 221. If the Delete Expired Volume Data setting is not used, logical volumes which have been returned to scratch will still be considered active data, allocating physical space in tape cartridges. In that case, only rewriting this logical volume will expire the old data, thus allowing that physical space occupied by old data to be reclaimed at a later time. With the Delete Expired Volume Data setting, the data that is associated with volumes that have been returned to scratch are expired after specified time period and their physical space in tape can be reclaimed.

For example, assume that you have 20,000 logical volumes in scratch status at any point in time, that the average amount of data on a logical volume is 400 MB, and that the data compresses at a 2:1 ratio. The space occupied by the data on those scratch volumes is 4,000,000 MB or the equivalent of 14 3592 JA cartridges. By using the Delete Expired Volume Data setting, you could reduce the number of cartridges required in this example by 14. The Parameter Expire Time specifies the amount of time in hours, days, or weeks. The data continues to be managed by the TS7700 Virtualization Engine after a logical volume is returned to scratch before the data associated with the logical volume is deleted. A minimum of 24 and a maximum of 32,767 hours (approximately 194 weeks) can be specified.

Specifying a value of zero used to work as the No Expiration option in older levels. Zero in this field causes an error message as shown in Figure 4-35 on page 223. You see the message because the data associated with the volume is managed as it was prior to the addition of this option, meaning that it is never deleted. In essence, specifying a value (other than zero) provides a “grace period” from when the logical volume is returned to scratch until its associated data is eligible for deletion. A separate Expire Time can be set for each category defined as fast-ready.

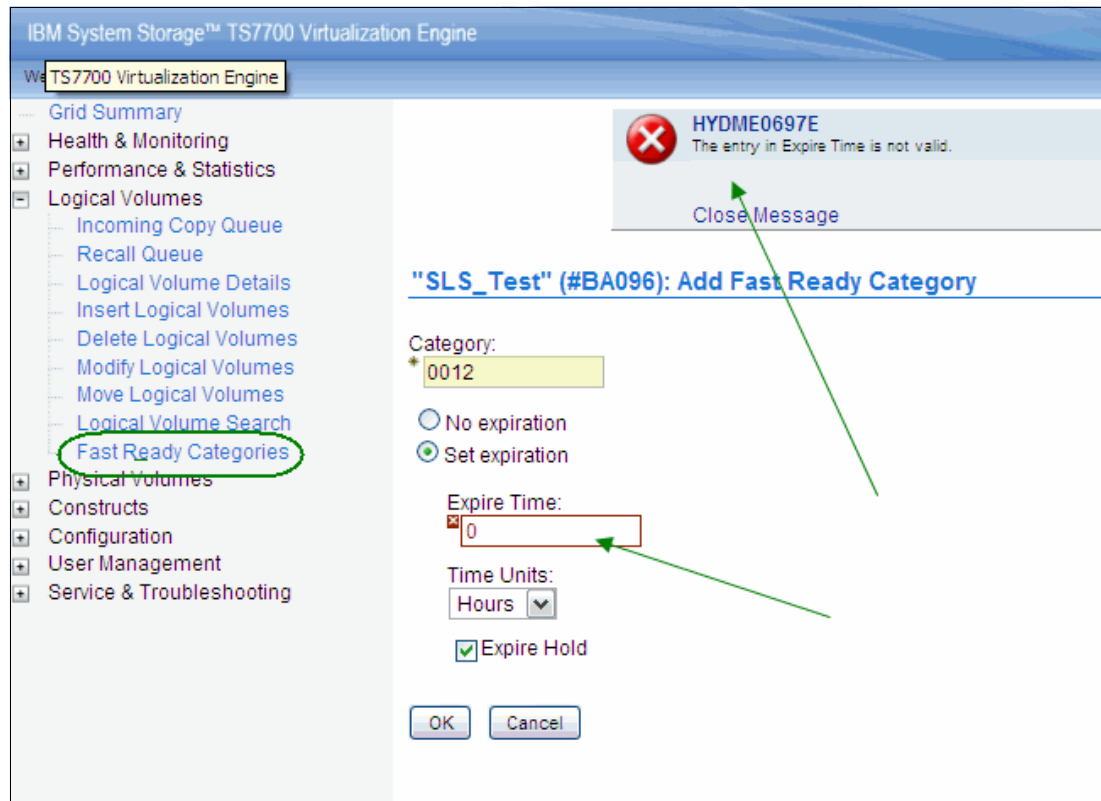


Figure 4-35 Invalid expire time value

Expire Time

Figure 4-34 on page 221 shows the number of hours or days in which logical volume data categorized as Fast Ready will expire. If the field is set to 0, the categorized data will never expire. The minimum Expire Time is 1 hour and the maximum Expire Time is 195 weeks, 1365 days, or 32,767 hours. The Expire Time default value is 24 hours.

Establishing the Expire Time for a volume occurs as a result of specific events or actions. The possible events or actions and their effect on the Expire Time of a volume are as follows:

- A volume is mounted.

The data that is associated with a logical volume will not be deleted, even if it is eligible, if the volume is mounted. Its Expire Time is set to zero, meaning it will not be deleted. It will be re-evaluated for deletion when its category is subsequently assigned.

- A volume's category is changed.

Whenever a volume is assigned to a category, including assignment to the same category that it currently is in, it is re-evaluated for deletion.

- Expiration.

If the category has a non-zero Expire Time, the volume's data is eligible for deletion after the specified time period, even if its previous category had a different non-zero Expire Time.

- No action.

If the volume's previous category had a non-zero Expire Time or even if the volume was already eligible for deletion (but has not yet been selected to be deleted) and the category it is assigned to has an Expire Time of zero, the volume's data is no longer eligible for deletion. Its Expire Time is set to zero.

- A category's Expire Time is changed.

If a user changes the Expire Time value through the Fast-ready Definition window on the Library Manager, the volumes assigned to the category are re-evaluated for deletion.

- Expire Time is changed from non-zero to zero.

If the Expire Time is changed from a non-zero value to zero, volumes assigned to the category that currently have a non-zero Expire Time are reset to an Expire Time of zero. If a volume was already eligible for deletion, but had not been selected for deletion, the volume's data is no longer eligible for deletion.

- Expire Time is changed from zero to non-zero.

Volumes currently assigned to the category continue to have an Expire Time of zero. Volumes subsequently assigned to the category will have the specified non-zero Expire Time. Non-zero to non-zero.

Volumes maintain their current Expire Time. Volumes subsequently assigned to the category will have the updated non-zero Expire Time.

After a volume's Expire Time has been reached, it is eligible for deletion. Not all data that is eligible for deletion will be deleted in the hour it is first eligible. Once an hour, on the half hour, the Library Manager selects up to 500 eligible volumes for data deletion. The volumes are selected based on the time that they became eligible, with the oldest ones being selected first. Up to 500 eligible volumes for the TS7700 Virtualization Engine in the library are selected first.

The Hold function, which you can enable by setting the check mark, essentially “locks” the volume from being used during the grace period. This means that it will not be selected to satisfy either a scratch or specific mount. In addition, it cannot be moved out of the category without IBM service intervention.

Expire Hold

Expire hold is on for a volume and expire time has not passed. An unexpired volume in a category with the hold attribute set will not be selected for a mount.

Checking Expire Hold (see Figure 4-34 on page 221) means that volumes in this fast-ready category may not be moved to another category or mounted before its data expires. Not checking it means no expire hold, letting volumes within this fast-ready category be freely mounted or moved to other categories before its data expires.

Note: Trying to mount a non-expired volume that belongs to a fast-ready category with Expire Hold on will result in error.

4.3.7 Operator Interventions

Verify that all raised Operator Interventions are being reported to the attached System z hosts, by selecting **Management Interface** → **Health & Monitoring** → **Operator Interventions** and checking the Send Future Interventions to Host table, which indicates that the “Current setting” is Enabled, as shown in Figure 4-36, so that the information can be sent to the hosts.

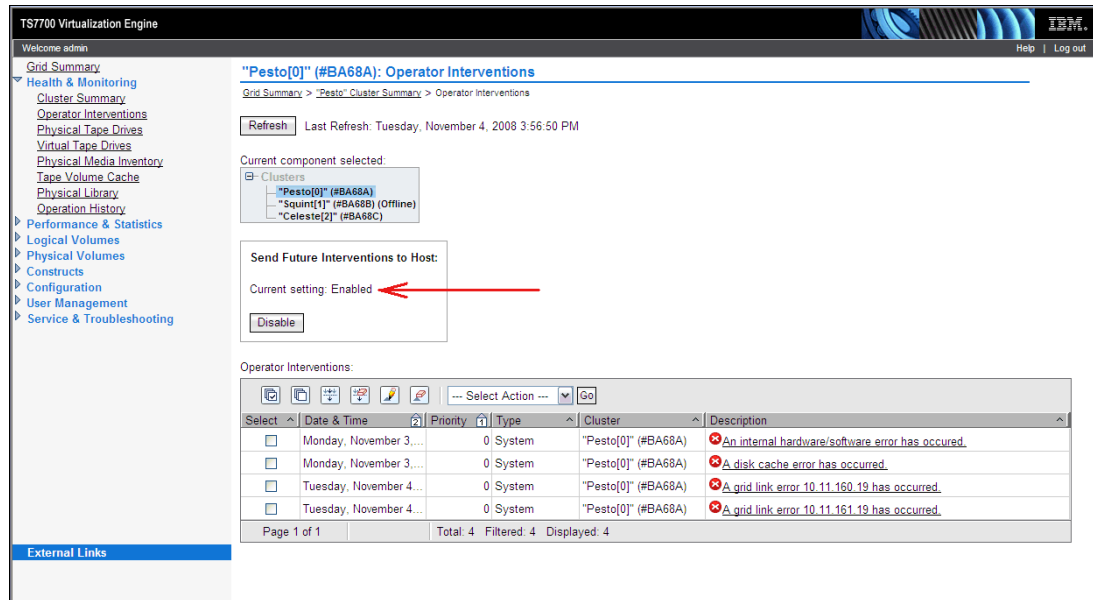


Figure 4-36 Operator Interventions

4.3.8 Defining TS7700 Constructs

To exploit the Outboard Policy Management functions, you must define four constructs:

- ▶ Storage Group (SG)
- ▶ Management Class (MC)
- ▶ Storage Class (SC)
- ▶ Data Class (DC)

These construct names are passed down from the z/OS host and stored with the logical volume. The actions defined for each construct are performed by the TS7700 Virtualization Engine.

Storage Groups

On the z/OS host, the Storage Group construct determines into which tape library a logical volume is written. Within the TS7740 Virtualization Engine, the Storage Group construct allows you to define the Storage Pool to which you want the logical volume.

Even before you define the first storage group, there is always at least one storage group present: the default storage group, which is identified by eight dashes (-----). This storage group cannot be deleted, but you can modify it to point to a different Storage Pool. You can define up to 256 storage groups, including the default.

Use the window shown in Figure 4-37 to add, modify, or delete a storage group used to define a primary pool for logical volume premigration.

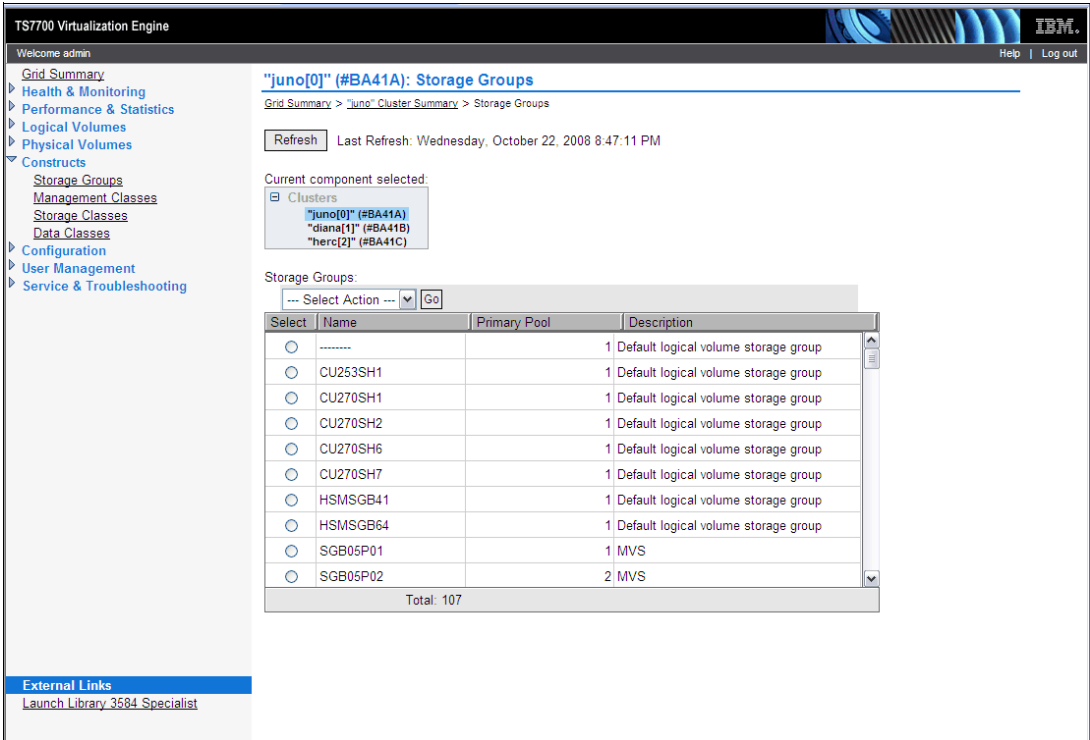


Figure 4-37 Storage Groups

The Storage Groups table displays all existing storage groups available for a given cluster.

You can use the Storage Groups table to create a new storage group, modify an existing storage group, or delete a storage group. The following status information is listed in the Storage Groups table:

- ▶ **Name:** The name of the storage group

Each storage group within a cluster must have a unique name. Valid characters for this field are as follows:

A - Z	Alphabetic characters
0 - 9	Numerals
\$	Dollar sign
@	At sign
*	Asterisk
#	Number sign
%	Percent
- ▶ **Primary Pool:** The primary pool for premigration

Only validated physical primary pools may be selected. If the cluster does not possess a physical library, this column will not be visible and the management interface will categorize newly created storage groups using pool 1.
- ▶ **Description:** A description of the storage group

Use the drop-down menu in the Storage Groups table to add a new storage group or modify or delete an existing storage group.

To add a new storage group, select **Add** from the drop-down menu. Complete the fields for the information that is displayed in the Storage Groups table.

Note: If the cluster is not attached to a physical library, the Primary Pool field will not be available in the Add or Modify menu options.

To modify an existing storage group, click the radio button from the Select column that appears adjacent to the name of the storage group you want to modify. Select **Modify** from the drop-down menu. Complete the fields for information that will be displayed in the Storage Groups table.

To delete an existing storage group, select the button in the Select column next to the name of the storage group you want to delete. Select **Delete** from the drop-down menu. You are prompted to confirm your decision to delete a storage group. If you select **Yes**, the storage group will be deleted. If you select **No**, your request to delete is cancelled.

Creating Management Classes

You can define, through the Management Class, whether you want to have a dual copy of a logical volume within the same TS7700 Virtualization Engine. In a Grid configuration, you will most likely choose to copy logical volumes over to the other TS7700s instead of creating a second copy in the same TS7700 Virtualization Engine. In a Single Cluster configuration, however, you might want to protect against media failures by using the dual copy capability. The second copy of a volume can be in a pool that is designated as an Export Copy pool. See 2.5.4, “Copy Export” on page 70 for more information.

If you want to have dual copies of selected logical volumes, you must use at least two Storage Pools, because the copies cannot be written to the same Storage Pool as the original logical volumes.

A default Management Class is always available. It is identified by eight dashes (-----) and cannot be deleted. You can define up to 256 Management Classes, including the default. Use the window shown in Figure 4-38 on page 228 to define, modify, or delete the Management Class that defines the TS7700 Virtualization Engine copy policy for volume redundancy.

The Current Copy Policy Table displays the copy policy in force for each component of the Grid. If no Management Class is selected, this table will not be visible. You must select a Management Class from the Management Classes Table to view copy policy details.

IBM System Storage™ TS7700 Virtualization Engine

Welcome admin

Grid Summary
 Health & Monitoring
 Performance & Statistics
 Logical Volumes
 Physical Volumes
 Constructs
 Storage Groups
 Management Classes
 Storage Classes
 Data Classes
 Configuration
 User Management
 Service & Troubleshooting

"Kramer[0]" (#BA82A): Management Classes
 Grid Summary > "Kramer" Cluster Summary > Management Classes
 Refresh Last Refresh: Thursday, October 22, 2009 8:26:23 AM

Current cluster selected:
 Clusters
 "Kramer[0]" (#BA82A)
 "Mork[1]" (#BA82B)
 "Gilligan[2]" (#BA82C)
 "Skipper[3]" (#BA82D)

Management Classes:

Select	Name	Secord	Description	Retain Copy Mo	"Kramer[0]" (#BA82A)	"Mork[1]" (#BA82B)	"Gilligan[2]" (#BA82C)	"Skipper[3]" (#BA82D)
<input type="checkbox"/>	-----	0	Default management class	No	Rewind Unload (RUN)	No Copy	No Copy	No Copy
<input type="checkbox"/>	DDDDNNNN	0	D D D D	No	Deferred	Deferred	Deferred	Deferred
<input type="checkbox"/>	DDDDNNNN	0	D D D I	No	Deferred	Deferred	Deferred	Rewind Unload (RUN)
<input type="checkbox"/>	DDDDNNNN	0	D D D N	No	Deferred	Deferred	Deferred	No Copy
<input type="checkbox"/>	DDIDNNNN	0	D D I D	No	Deferred	Deferred	Rewind Unload (RUN)	Deferred
<input type="checkbox"/>	DDIDNNNN	0	D D I I	No	Deferred	Deferred	Rewind Unload (RUN)	Rewind Unload (RUN)
<input type="checkbox"/>	DDIDNNNN	0	D D I N	No	Deferred	Deferred	Rewind Unload (RUN)	No Copy
<input type="checkbox"/>	DDNDNNNN	0	D D N D	No	Deferred	Deferred	No Copy	Deferred
<input type="checkbox"/>	DDNDNNNN	0	D D N I	No	Deferred	Deferred	No Copy	Rewind Unload (RUN)
<input type="checkbox"/>	DDNNNNNN	0	D D N N	No	Deferred	Deferred	No Copy	No Copy

Page 1 of 9 1 Go Rows: 10 Show Total: 83 Filtered: 83 Displayed: 10

Figure 4-38 Management Classes

The Management Classes Table (Figure 4-38) displays defined Management Class copy policies that may be applied to a cluster.

You can use the Management Classes Table to create a new Management Class, modify an existing Management Class, or delete one or more existing Management Classes. The default Management Class may be modified, but cannot be deleted. The default name of Management Class uses eight dashes (-----).

Status information displayed in the Management Classes Table is as follows:

- ▶ **Name:** The name of the Management Class
Valid characters for this field are A - Z, 0 - 9, \$, @, *, #, and %. The first character of this field may not be a number. This is the only field that may not be modified after it is added.
- ▶ **Secondary Pool:** The target pool in the volume duplication
If the cluster does not possess a physical library, this column will not be visible and the management interface will categorize newly created storage groups using pool 0.
- ▶ **Description:** A description of the Management Class definition
The value in this field must be between 1 and 70 characters in length.
- ▶ **Retain Copy Mode (Yes or No)**
Retain Copy Mode honors the original Copy Consistency Policy that is in place in the cluster where the volume was created. With this, undesired copies are not created throughout the grid. See Chapter 2, "Architecture, components, and functional characteristics" on page 15 for more details.
- ▶ **Cluster Copy Policy** provides the ability to define where and when copies are made

Use the drop-down menu in the Management Classes table to add a new Management Class, modify an existing Management Class, or delete one or more existing Management Classes.

To add a new Management Class, select **Add** from the drop-down menu and click **Go**. Complete the fields for information that will be displayed in the Management Classes Table. You can create up to 256 Management Classes per TS7700 Virtualization Engine Grid.

Note: If cluster is not attached to a physical library, the Secondary Pool field will not be available in the **Add** option.

The Copy Action drop-down menu is adjacent to each cluster in the TS7700 Virtualization Engine Grid. Use the Copy Action menu to select, for each component, the copy mode to be used in volume duplication. Actions available from this menu are as follows:

- ▶ No Copy: No volume duplication will occur if this action is selected.
- ▶ Rewind Unload (RUN): Volume duplication occurs when the Rewind Unload command is received. The command return only after the volume duplication completes successfully.
- ▶ Deferred: Volume duplication occurs at a later time based on the internal schedule of the copy engine.

Figure 4-39 shows the Add Management Class window.

Clusters	Copy Mode
"archie[0]" (#BA96A)	Deferred
"veronica[1]" (#BA96B)	Deferred
"dohnut[2]" (#BA96C)	Deferred
"veronca1[3]" (#BA96D)	Deferred

Figure 4-39 Add Management Class

To modify an existing Management Class, select the check box (in the Select column) that is in the same row as the name of the Management Class you want to modify. You can modify only one Management Class at a time. Select **Modify** from the drop-down menu and click **Go**. Of the fields listed in the Management Classes Table, or available from the Copy Action drop-down menu, you are able to change all of them except the Management Class name.

Figure 4-40 shows the Modify Management Class window.

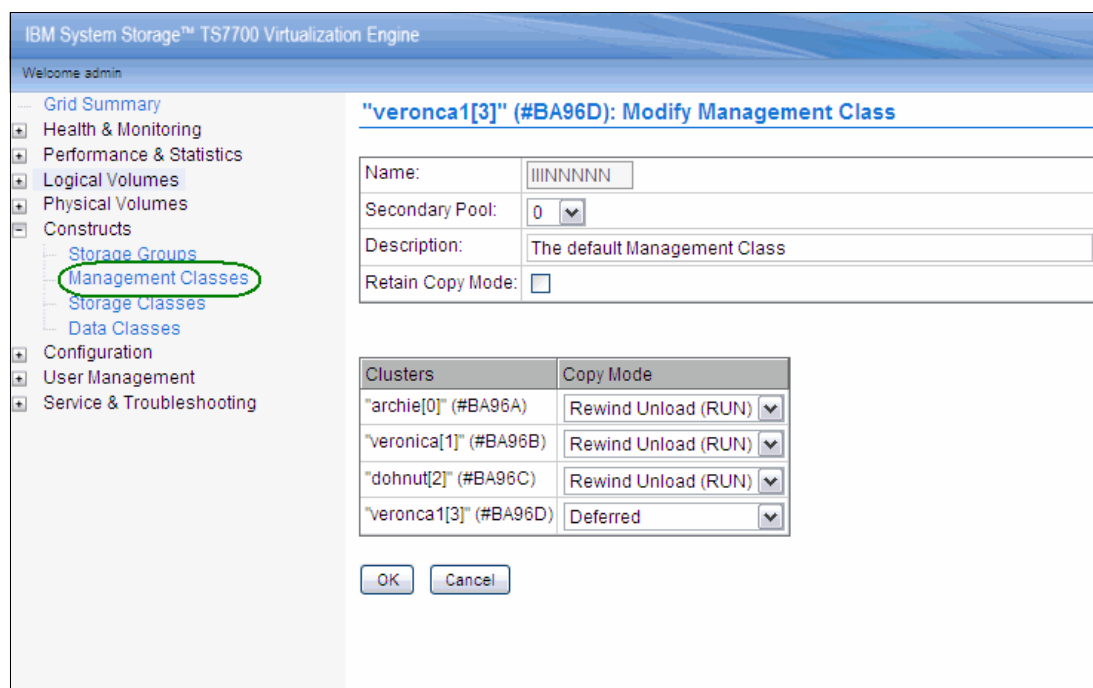


Figure 4-40 modify management class

To delete one or more existing Management Classes, select the check box (in the Select column) that is in the same row as the name of the Management Class you want to delete. Select multiple check boxes to delete multiple Management Classes. Select **Delete** from the drop-down menu and click **Go**.

Note: You may not delete the default Management Class.

Creating Storage Classes

By using the Storage Class construct, you can influence when a logical volume is removed from cache.

A default Storage Class is always available. It is identified by eight dashes (-----) and cannot be deleted. Use the window shown in Figure 4-41 on page 231 to define, modify, or delete a storage class used by the TS7700 Virtualization Engine to automate storage management through classification of data sets and objects.

The Storage Classes table lists storage classes that are defined for each component of the Grid.

Note: Storage Class constructs do not apply to TS7700 Virtualization Engine Cluster that is not attached to a physical Library.

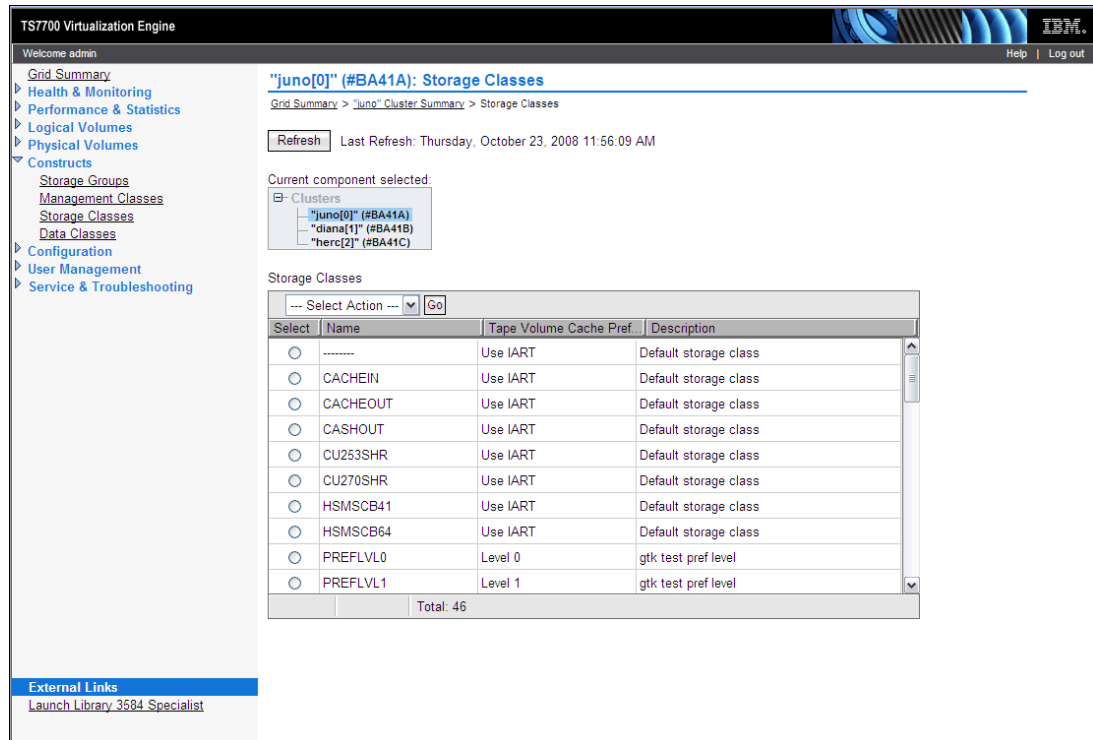


Figure 4-41 Storage Classes

The Storage Classes table (Figure 4-41) lists defined storage classes that are available to control data sets and objects within a cluster.

Use the Storage Classes table to create a new storage class or modify or delete an existing storage class. The default storage class may be modified, but cannot be deleted. The default storage class uses eight dashes as the name (-----).

The following status information is displayed in the Storage Classes table:

- ▶ **Name:** The name of the storage class
Valid characters for this field are A - Z, 0 - 9, \$, @, *, #, and%. The first character of this field may not be a number. The value in this field must be 1 - 8 characters in length.
- ▶ **Tape Volume Cache Preference:** The preference level for the storage class
It determines how soon volumes are removed from cache following their copy to tape. Possible values are as follows:
 - Use IART** Volumes are removed according to the Initial Access Response Time (IART) of the TS7700 Virtualization Engine.
 - Level 0** Volumes are removed from the tape volume cache as soon as they are copied to tape.
 - Level 1** Copied volumes remain in the tape volume cache until additional space is required, then are the first volumes removed to free space in the cache.
- ▶ **Description:** A description of the storage class definition
The value in this field must be 1 - 70 characters in length.

Use the drop-down menu in the Storage Classes Table to add a new storage class or modify or delete an existing storage class.

To add a new storage class, select **Add** from the drop-down menu. Complete the fields for the information that will be displayed in the Storage Classes Table.

Note: You can create up to 256 storage classes per TS7700 Virtualization Engine Grid.

To modify an existing storage class, click the radio button from the Select column that appears in the same row as the name of the storage class you want to modify. Select **Modify** from the drop-down menu. Of the fields listed in the Storage Classes Table, you will be able to change all of them except for the storage class name.

To delete an existing storage class, click the radio button from the Select column that appears in the same row as the name of the storage class you want to delete. Select **Delete** from the drop-down menu. A dialog box opens where you confirm the storage class deletion. Select **Yes** to delete the storage class, or select **No** to cancel the delete request.

Figure 4-42 shows the Add Storage Class panel.

The screenshot shows the 'Add Storage Class' dialog box in the IBM System Storage TS7700 Virtualization Engine. The left-hand navigation pane lists various system components, with 'Storage Classes' under the 'Constructs' section highlighted with a green circle. The main area of the dialog is titled '"archie[0]" (#BA96A): Add Storage Class'. It contains three input fields: 'Name:' with a text box containing a single character, 'Tape Volume Cache Preference:' with a dropdown menu currently showing 'Use IART', and 'Description:' with an empty text box. At the bottom of the dialog are 'OK' and 'Cancel' buttons.

Figure 4-42 Add Storage Class

Creating Data Classes

From a z/OS perspective (SMS managed tape) the DFSMS Data Class defines the following information:

- ▶ Media type parameters
- ▶ Recording technology parameters
- ▶ Compaction parameters
- ▶ Performance Scaling
- ▶ Segmentation

For the TS7700 Virtualization Engine, only the Media type, Recording technology, and Compaction parameters are used. The use of larger logical volume sizes is controlled through Data Class.

A default Data Class is always available. It is identified by eight dashes (-----) and cannot be deleted.

Use the window, shown in Figure 4-43 on page 233, to define, modify, or delete a TS7700 Virtualization Engine Data Class used to automate storage management through classification of data sets.

Select	Name	Logical Volume Size (MiB)	Logical WORM	Description
<input type="radio"/>	-----	Insert Media Class	No	The default Data Class
<input type="radio"/>	HSMDCLS1	Insert Media Class	No	The default Data Class
<input type="radio"/>	HSMDCLS2	Insert Media Class	No	The default Data Class
<input type="radio"/>	HSMDCLS6	Insert Media Class	No	The default Data Class
<input type="radio"/>	HSMDCLS7	Insert Media Class	No	The default Data Class
<input type="radio"/>	M1C1GIG	1,000	Yes	Media 1 Compact 1GB
<input type="radio"/>	M1C2GIG	2,000	No	Media 1 Compact 2GB
<input type="radio"/>	M1C4GIG	4,000	No	Media 1 Compact 4GB
<input type="radio"/>	M1N1GIG	1,000	No	Media 1 Non compact 1GB
<input type="radio"/>	M1N2GIG	2,000	No	Media 1 Non compact 2GB

Figure 4-43 Data Classes

The Data Classes table (Figure 4-43) displays the list of Data Classes defined for each cluster of the Grid.

You can use the Data Classes table to create a new data class or modify or delete an existing Data Class. The default Data Class may be modified, but cannot be deleted. The default Data Class shows the name as eight dashes (-----).

Status information displayed in the Data Classes table is as follows:

- **Name:** The name of the Data Class
Valid characters for this field are A - Z, 0 - 9, \$, @, *, #, and %. The first character of this field may not be a number. The value in this field must be between 1 and 8 characters in length.
- **Logical Volume Size (mebibyte, MiB):** The logical volume size of the Data Class
It determines the maximum number of MiB for each logical volume in a defined class. Possible values are as follows:
 - Insert Media Class** Logical volume size is not defined; the Data Class will not be defined by a maximum logical volume size. You can use 1000 MiB, 2000 MiB, or 4000 MiB
 - Logical WORM (Yes or No)** Whether Logical WORM (write once, read many) is set for the Data Class. Logical WORM is the virtual equivalent of WORM tape media, achieved through Licensed internal code emulation. This setting is available only when all clusters in a Grid operate a code level of 8.6.0.xx or later. This is a new function introduced in R1.6.
- **Description:** A description of the Data Class definition
The value in this field must be at least 0 and no greater than 70 characters in length.

Use the drop-down menu on the Data Classes table to add a new Data Class or modify or delete an existing Data Class. See Figure 4-44.

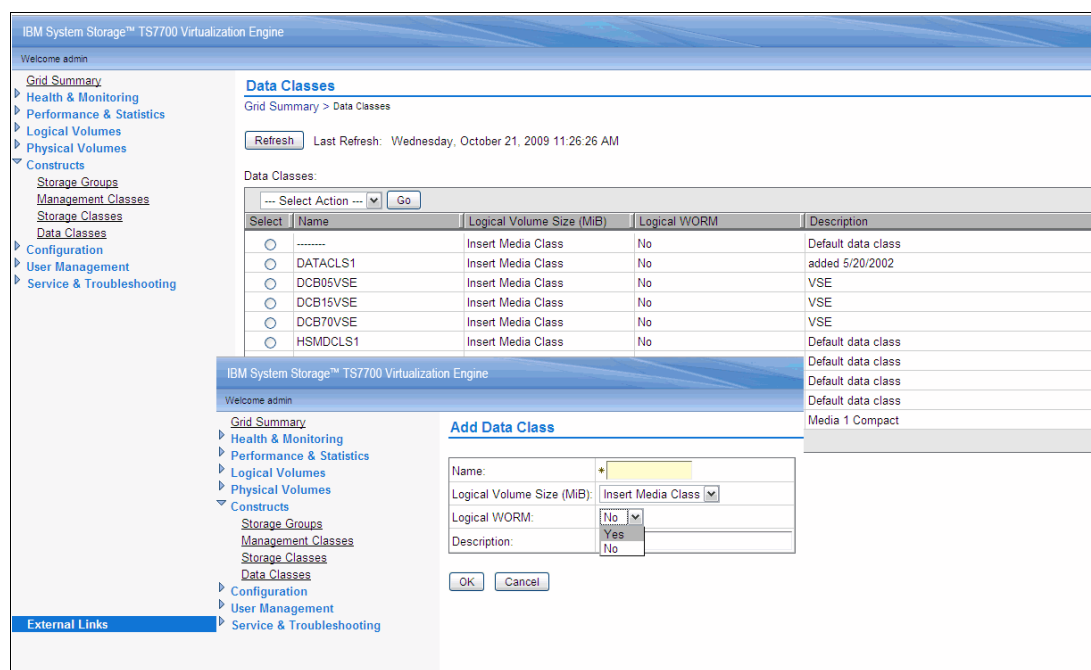


Figure 4-44 Data Classes and Add Data Class windows

To add a new Data Class, select **Add** from the drop-down menu and click **Go**. Complete the fields for information that will be displayed in the Data Classes Table.

Note: You can create up to 256 Data Classes per TS7700 Virtualization Engine Grid.

To modify an existing Data Class, select the check box (in the Select column) that appears in the same row as the name of the Data Class you want to modify. Select **Modify** from the drop-down menu and click **Go**. Of the fields listed in the Data Classes Table, you will be able to change all of them except the default Data Class name.

To delete an existing Data Class, click the radio button from the Select column that appears in the same row as the name of the Data Class you want to delete. Select **Delete** from the drop-down menu and click **Go**. A dialog box opens where you can confirm the Data Class deletion. Select **Yes** to delete the Data Class, or select **No** to cancel the delete request.

4.3.9 TS7700 licensing

This section describes several licensing options with the TS7700 and cache enablement performance increment license key entry.

Note: Cache enablement performance increment license key entry applies only on a TS7740 Virtualization Engine configuration, because a TS7720 Virtualization Engine does not have a 1-TB cache enablement feature (FC5267).

The amount of disk cache capacity and performance capability are enabled using feature license keys. You will receive feature license keys for the features that you have ordered.

Each feature increment allows you to tailor the subsystem to meet your disk cache and performance needs.

You use the Feature Licenses window (Figure 4-45) to activate feature licenses in the TS7700 Virtualization Engine. To open the window, select **Activate New Feature License** from the list and click **Go**. Enter the license key into the fields provided and select **Activate**.

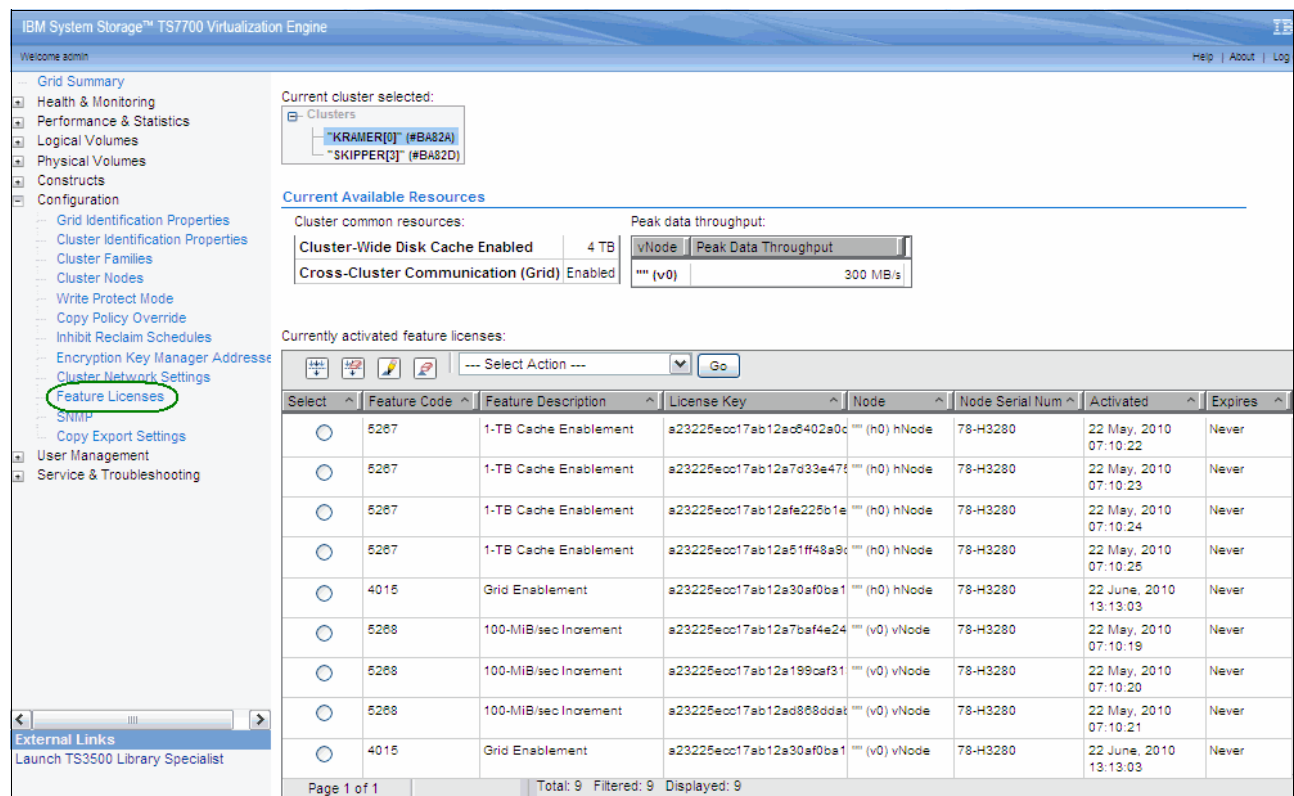


Figure 4-45 Feature Licenses window

To remove a license key, select the feature license to be removed, select **Remove Selected feature License** from the list, and click **Go**.

Attention: do not remove any installed peak data throughput features because removal can affect host jobs.

FC5627 1 TB Cache Enablement is not removable after activation.

When you select **Activate New Feature License**, the Feature License entry window opens, as shown in Figure 4-46. When you enter a valid feature license key and click **Activate**, the feature will be activated.

Note: Performance Increments become active immediately. Cache Increments become active within 30 minutes.

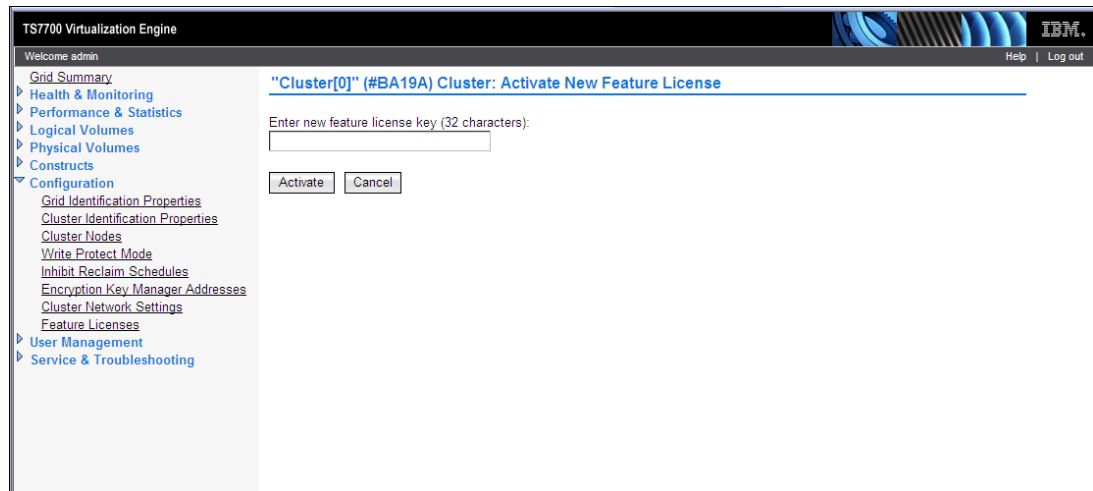


Figure 4-46 Activate New Feature Licenses

4.3.10 Defining Encryption Key Manager addresses

Use the window (shown in Figure 4-48 on page 237) to set the encryption key manager addresses in the TS7700 Virtualization Engine, if required.

To watch a tutorial that shows the properties of the Encryption Key Manager, click the **View tutorial** link. If the cluster is not attached to a physical library, the screen displays an error message, as in Figure 4-47 on page 237.

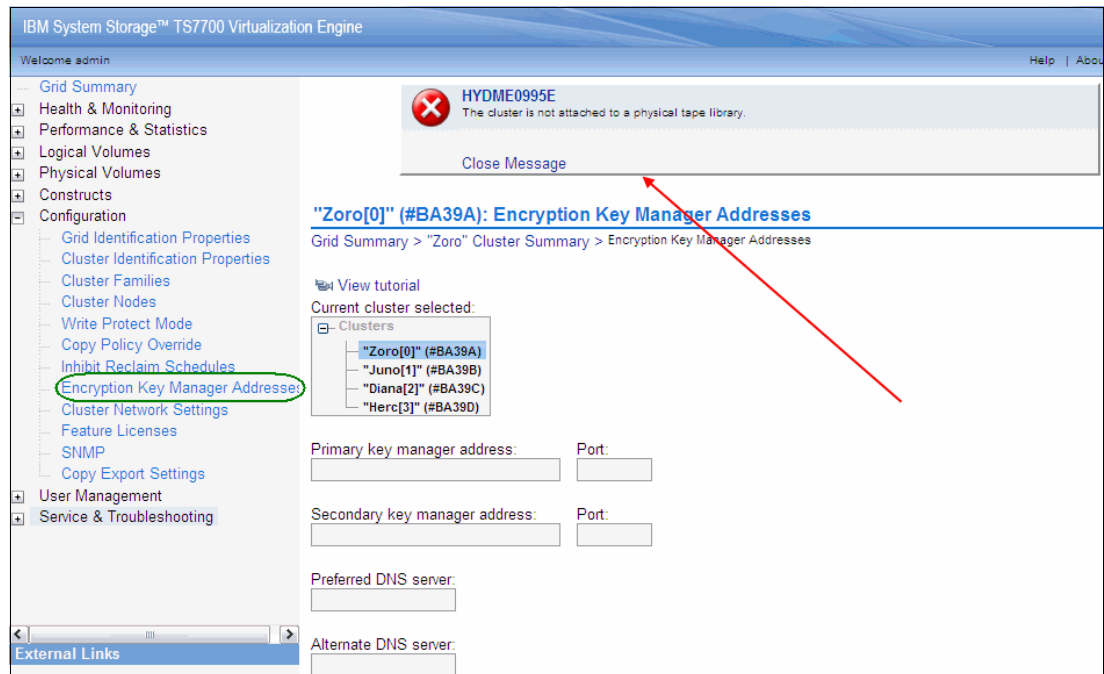


Figure 4-47 Encryption Key Manager Addresses error message

The encryption key manager assists encryption-enabled tape drives in generating, protecting, storing, and maintaining encryption keys that are used to encrypt information being written to and decrypt information being read from tape media (tape and cartridge formats).

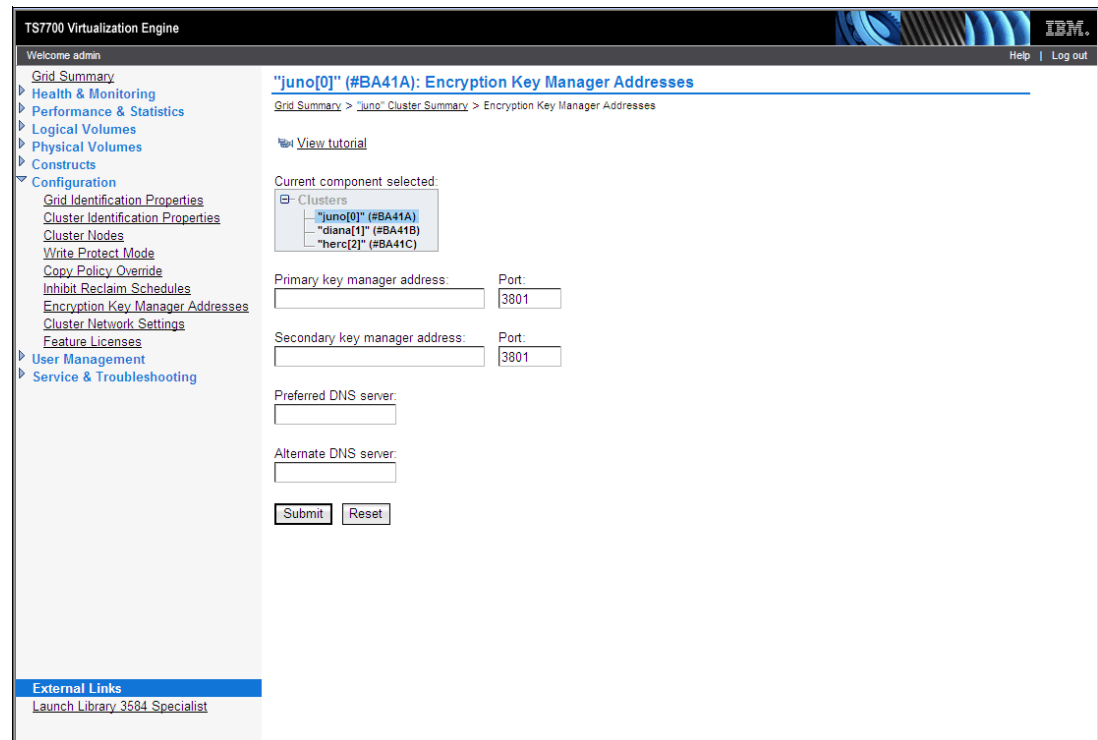


Figure 4-48 Encryption Key Manager Addresses

The following settings are used to configure the TS7700 Virtualization Engine connection to an Encryption Key Manager (see Figure 4-48 on page 237 as needed):

- ▶ Primary key manager address: The key manager server name or IP address that is primarily used.
- ▶ Primary key manager port: The port number of the primary key manager. The default value is 3801. This field is only required if a primary key address is used.
- ▶ Secondary key manager address: The key manager server name or IP address that is used when the primary key manager is unavailable.
- ▶ Secondary key manager port: The port number of the secondary key manager. The default value is 3801. This field is required only if a secondary key address is used.
- ▶ Preferred DNS server: The Domain Name Server (DNS) that is primarily used. DNS addresses are only needed if you specify a symbolic domain name for one of the key manager addresses rather than a numeric IP address. If you need to specify a DNS, be sure to specify both a primary and an alternate so you do not lose access to your Encryption Key Manager because of one of the DNS servers being down or inaccessible. This address can be in IPv4 format.
- ▶ Alternate DNS server: The Domain Name Server that is used in case the preferred DNS server is unavailable. If a preferred DNS server is specified, specify an alternate DNS also. This address can be in IPv4 format.
- ▶ Using the Ping Test: Use the Ping Test button to check component network connection to a key manager after changing a component's address or port. If you change a key manager address or port and do not submit the change before using the Ping Test button, you receive the following warning:

In order to perform a ping test you must first submit your address and port changes.

After the ping test has been issued, you receive one of the following messages:

- “The ping test against the address <key manager address> on port <port> was successful.”
- “The ping test against the address <key manager address> on port <port> from “<cluster>” has failed. The error returned was: <error text>.”

Click the **Submit** button to save changes to any of the settings. To discard changes and return the field settings to their current values, click the **Reset** button.

Note: The two EKMs must be set up on separate systems to provide redundancy. Connection to an EKM is required to read encrypted data.

4.3.11 Defining SNMP

Use the window shown in Figure 4-49 to view or modify the simple network management protocols (SNMP) configured on a IBM Virtualization Engine TS7700 Cluster.

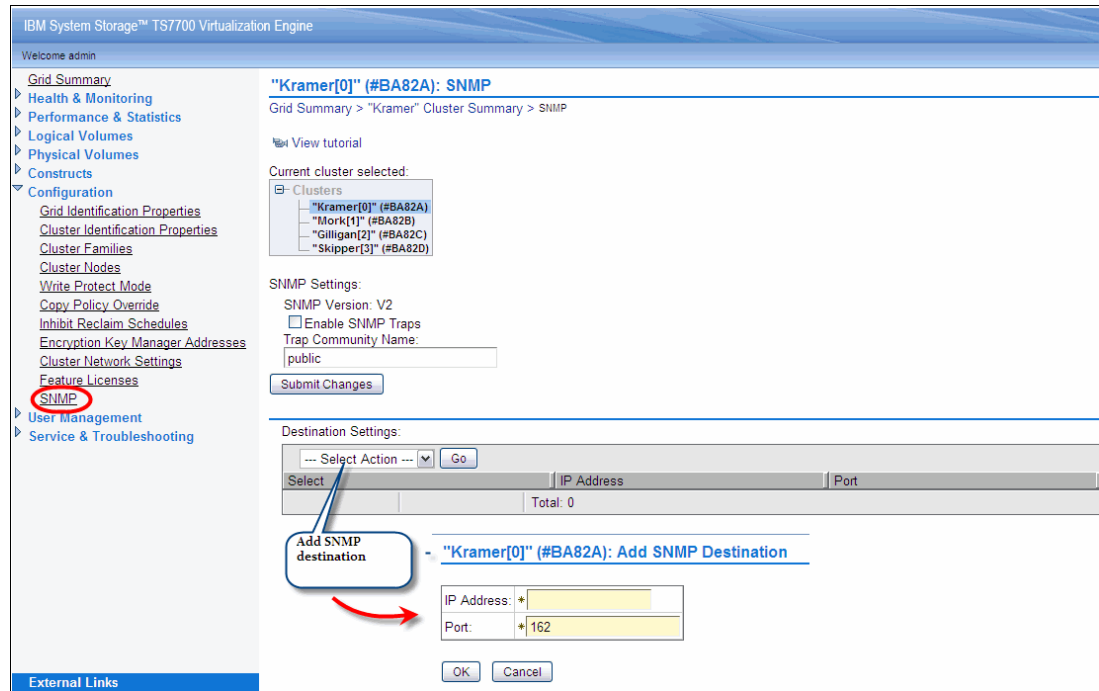


Figure 4-49 SNMP settings

Use the window to configure SNMP traps that will log operation history events such as login occurrences, configuration changes, status changes (vary on or off and service prep), shut down, and code updates. SNMP is a networking protocol that allows a IBM Virtualization Engine TS7700 to automatically gather and transmit information about alerts and status to other entities in the network.

SNMP Settings

Use this section to configure global settings that apply to SNMP traps on an entire cluster. Settings that may be configured are as follows:

- ▶ **SNMP Version:** The SNMP version defines the protocol used in sending SNMP requests and is determined by the tool you are using to monitor SNMP traps. Various versions of SNMP traps work with different management applications. Selecting a value of V1, the recommended trap version is compatible with the greatest number of management applications.
- ▶ **Enable SMP Traps:** This check box enables or disables SNMP traps on a cluster. If the check box is selected, SNMP traps on the cluster are enabled, if the check box is *not* selected (the default) SNMP traps on the cluster are disabled.
- ▶ **Trap Community Name:** This name identifies the trap community and is sent along with the trap to the management application. This value behaves as a password; the management application will not process an SNMP trap unless it is associated with the correct community. This value must be 1 - 15 characters in length and composed of Unicode characters.

Destination Settings

Use the Destination Settings table to add, modify, or delete a destination for SNMP trap logs. You can add, modify, or delete a maximum of 16 destination settings at one time. Settings that may be configured are as follows:

- ▶ **IP Address:** This setting is the IP address of the SNMP server in IPv4 format. A value in this field is required.
- ▶ **Port:** This port is where the SNMP trap logs are sent. This value must be a number between 0 and 65535. A value in this field is required.

Use the **Select Action** drop-down menu in the Destination Settings table to add, modify, or delete an SNMP trap destination:

- ▶ **Add SNMP destination:** Select this menu item to add an SNMP trap destination for use in the IBM Virtualization Engine TS7700 Grid.
- ▶ **Modify SNMP destination:** Select this menu item to modify an SNMP trap destination that is used in the IBM Virtualization Engine TS7700 Grid.

4.3.12 Inserting logical volumes

Use the Insert Logical Volumes window (Figure 4-50 on page 241) to insert a range of logical volumes in the TS7700 Virtualization Engine Grid. Logical volumes inserted into an individual cluster will be available to all clusters within a Grid configuration.

IBM System Storage™ TS7700 Virtualization Engine

Welcome admin

Grid Summary

- Health & Monitoring
- Performance & Statistics
- Logical Volumes
 - Incoming Copy Queue
 - Recall Queue
 - Logical Volume Details
 - Insert Logical Volumes**
 - Delete Logical Volumes
 - Modify Logical Volumes
 - Move Logical Volumes
 - Logical Volume Search
 - Fast Ready Categories
- Physical Volumes
- Constructs
- Configuration
- User Management
- Service & Troubleshooting

External Links

Insert Logical Volumes

Grid Summary > Insert Logical Volumes

Last Refresh: Wednesday, October 21, 2009 3:05:36 PM

Current availability across entire grid:

Currently Inserted:	203,999
Maximum Allowed:	1,000,000
Available Slots:	796,001

Show inserted volume ranges between and

Insert a new logical volume range:

*Starting volser: *Quantity

☐ Ending volser

Initially owned by:

Media type:
☐ Cartridge System Tape (400 MiB)
☒ Enhanced Capacity Cartridge System Tape (800 MiB)

☒ Set Constructs

Storage Groups: Management Classes: Storage Classes: Data Classes:

Figure 4-50 TS7700 Virtualization Engine MI Insert Logical Volumes

During logical volumes entry processing on z/OS, even if the library is online and operational for a given host, at least one device needs to be online (or have been online) for that host for the library to be able to send the volume entry attention interrupt to that host. If the library is online and operational, yet there are no online devices to a given host, that host will not receive the attention interrupt from the library unless a device had previously been varied online. To work around this limitation, ensure that at least one device is online (or had been online) to each host or use the LIBRARY RESET,CBRUXENT command to initiate cartridge entry processing from the host. This task is especially important if you only have one host attached to the library that owns the volumes being entered. In general, after you have entered volumes into the library, if you do not see the expected CBR36xxI cartridge entry messages being issued, you can use the LIBRARY RESET,CBRUXENT command from MVS to initiate cartridge entry processing. The LIBRARY RESET,CBRUXENT command causes the host to ask for any volumes in the insert category.

Up to now, as soon as OAM started, and if volumes were in the Insert category, entry processing started without giving the customer the chance to stop entry processing the first time OAM is started.

Now, the LI DISABLE, CBRUXENT command can be used without starting the OAM address space. This approach gives you the chance to stop entry processing before the OAM address space initially starts.

The table at the top of Figure 4-50 on page 241 shows the current information about the number of logical volumes in the TS7700 Virtualization Engine:

- ▶ **Currently Inserted:** The total number of logical volumes inserted into the TS7700 Virtualization Engine
- ▶ **Maximum Allowed:** The total maximum number of logical volumes that can be inserted
- ▶ **Available Slots:** The available slots remaining for logical volumes to be inserted, which is obtained by subtracting the Currently Inserted logical volumes from the Maximum Allowed

To view the current list of logical volume ranges in the TS7700 Virtualization Engine Grid, enter a logical volume range and click **Show**.

Use the following fields if you want to insert a new logical volume range action:

- ▶ **Starting VOLSER:** This is the first logical volume to be inserted. The range for inserting logical volumes begins with this VOLSER number.
- ▶ **Quantity:** Select this option to insert a set amount of logical volumes beginning with Starting VOLSER. Enter the quantity of logical volumes to be inserted in the adjacent text field. You can insert up to 10,000 logical volumes at one time.
- ▶ **Ending VOLSER:** Select this option to insert a range of logical volumes. Enter the ending VOLSER number in the adjacent text field.
- ▶ **Initially owned by:** Indicates the name of the cluster that will own the new logical volumes. Select a cluster from the drop-down menu.
- ▶ **Media type:** Indicates the media type of the logical volume (volumes). Possible values are as follows:
 - Cartridge System Tape (400 MiB)
 - Enhanced Capacity Cartridge System Tape (800 MiB)
- ▶ **Set Constructs:** Select this check box to specify constructs for the new logical volume (or volumes), then use the drop-down menu below each construct to select a predefined construct name. You may specify the use of any or all of the following constructs:
 - Storage Group
 - Storage Class
 - Data Class
 - Management Class

Note: When using z/OS, do not specify constructs when the volumes are added, instead they are assigned during job processing when a volume is mounted.

To insert a range of logical volumes, complete the fields listed and click **Insert**. You are prompted to confirm your decision to insert logical volumes. To continue with the insert operation, click **Yes**. To abandon the insert operation without inserting any new logical volumes, click **No**.

Note: You can insert up to ten thousand (10,000) logical volumes at one time. This applies to both inserting a range of logical volumes and inserting a quantity of logical volumes.

4.4 Virtualization Engine multi-cluster grid definitions

This section covers several available configuration definitions with the TS7700.

4.4.1 TTS7700 Grid configuration considerations: two-cluster grid

The TS7700 Grid configuration provides for the automatic replication of data and may be used in a variety of high availability and disaster recovery situations. When connected together at the same location, a Grid configuration may also help facilitate higher availability to data. The topics in this section provide information for you to consider in planning for a disaster recovery or high availability solution using the Virtualization Engine TS7700 Grid configuration.

Note: All clusters in a multi-cluster grid must have FC4015 (Grid Enablement) installed. This includes existing stand-alone clusters that are being used to create a multi-cluster grid.

Figure 4-51 shows a basic two-way TS7740 Grid configuration.




Figure 4-51 TS7740 Virtualization Engine in a two-cluster grid configuration

Note: In the case of a TS7720 Virtualization Engine Grid configuration, disregard the physical library attachment.

Figure 4-52 shows a two-cluster hybrid example:




Figure 4-52 TS7700 Virtualization Engine in a hybrid two-cluster grid Configuration

Each TS7700 Cluster provides two or four FICON host attachments and 256 virtual tape device addresses. The clusters in a Grid configuration are connected together through dual, 1 Gb Ethernet links. The 1 Gb Ethernet links are used for the replication of data between clusters, passing control information and access between a local cluster's virtual tape device and a logical volume's data in a remote cluster's TS7700 Cache.

4.4.2 Data access and availability characteristics

A TS7700 Grid configuration provides the following data access and availability characteristics:

- ▶ Direct host attachment to a Virtualization Engine TS7700
- ▶ Accessibility of logical volumes through virtual device addresses on the TS7700s in the Grid configuration
- ▶ Selection of location of logical volume for access
- ▶ Whether a copy is available at another TS7700 cluster
- ▶ Volume Ownership and Ownership Takeover
- ▶ Service Prep/Service Mode
- ▶ Consequences of a link failure in your Grid

These grid operational characteristics must have been carefully considered and thoroughly planned by you and your IBM representative. The infrastructure needs must be addressed in advance.

4.4.3 Defining Grid copy mode control

When upgrading a stand-alone cluster to a grid, FC4015 must be installed on all clusters in the grid. Also, you must set up the copy consistency points in the Management Class definitions on all clusters in the new grid.

The data consistency point is defined in the Management Classes construct definition through the management interface. You can perform this task only for an existing Grid system; in a single cluster configuration, you will see only your single cluster in the Modify Management class definition. Figure 4-53 shows the Modify Management Class panel.

Note: With a single cluster configuration, copy mode must be set to Rewind Unload.

See “Creating Management Classes” on page 227 for more information.

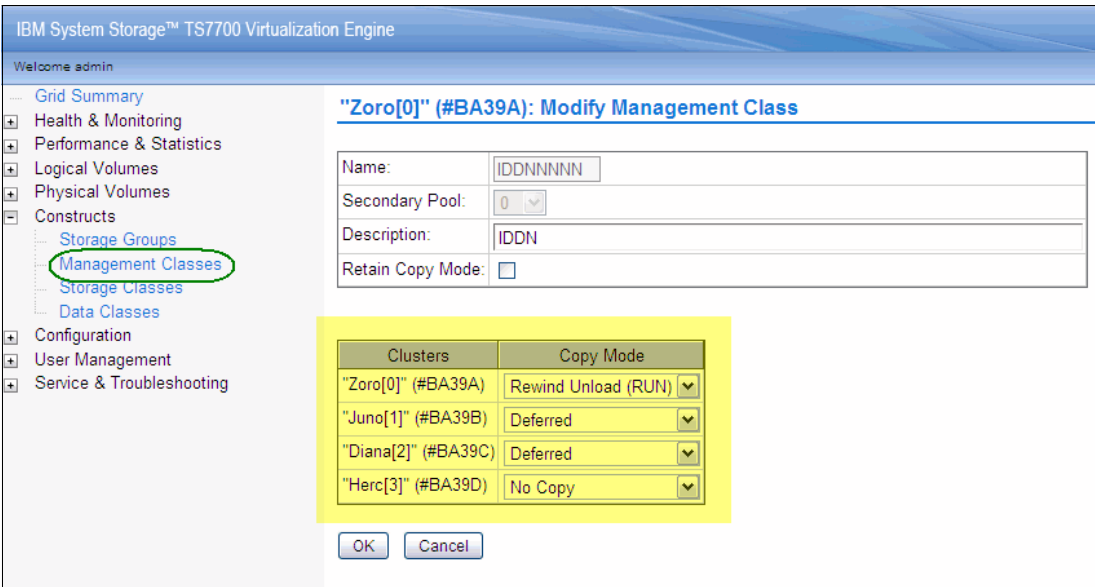


Figure 4-53 Example of Modify Management Class page

To open the Management Classes window (Figure 4-54 on page 246), click **Constructs** → **Management Classes** under the Welcome Admin menu. Mark the Management Class name and select **Modify** from the Select Action drop-down menu.

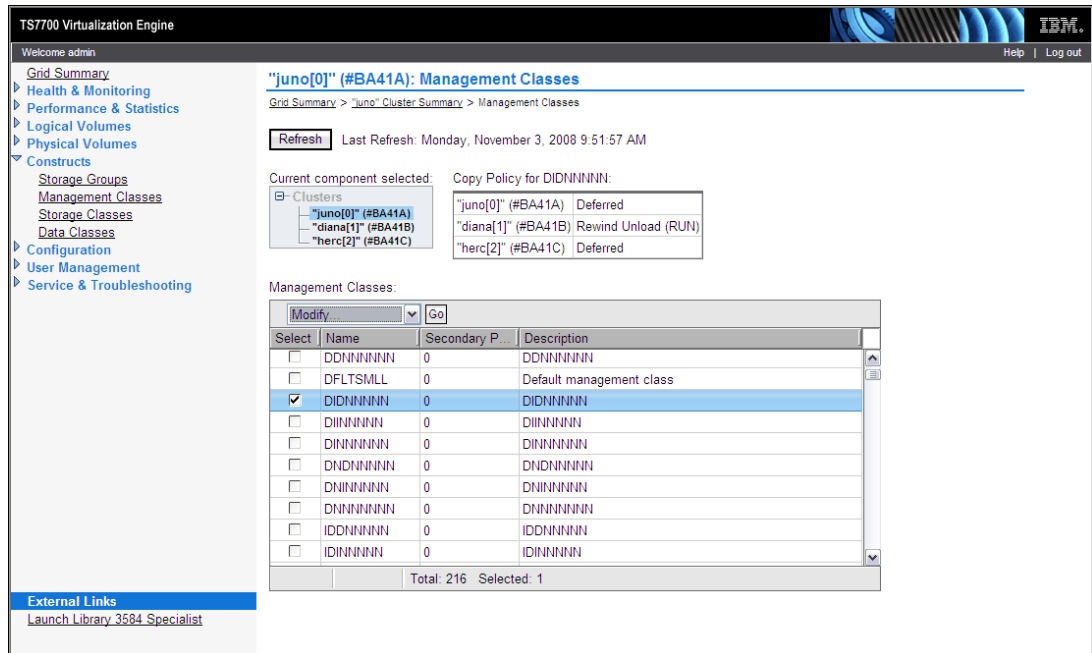


Figure 4-54 Management Classes

Click **Go**. The next window (Figure 4-55) is where you may modify the copy consistency by using the Copy Action table, and then clicking **OK**. In this example, the TS7700 Virtualization Engine (named VE - Partition 1) is part of a multi-cluster grid configuration. This additional drop-down menu is displayed only if a TS7700 Virtualization Engine is part of a multi-cluster grid environment.

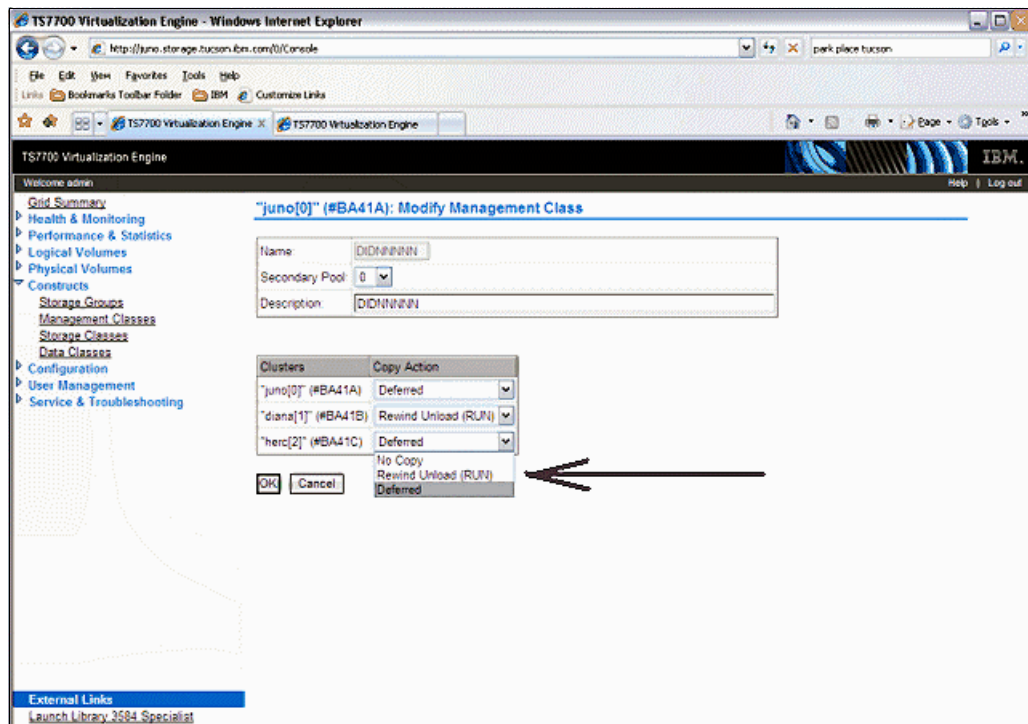


Figure 4-55 Modify Management Classes

As shown in Figure 4-55 on page 246, you can choose between three consistency points per cluster:

- ▶ No Copy: No copy (NC) is made to this cluster.
- ▶ Rewind Unload (RUN): A valid version of the logical volume has been copied to this cluster as part of the volume unload processing.
- ▶ Deferred: A replication of the modified logical volume is made to this cluster after the volume had been unloaded (DEF).

Table 4-3 provides an overview of the possible Cluster Copy Data Consistency Points for the Management Class in a two-cluster grid, and their consequences.

Table 4-3 Possible settings of the Data Consistency Points combinations: two-cluster grid

Cluster 0	Cluster 1	Consequence
RUN	RUN	Use this setting to have both clusters maintain a consistent copy when the Rewind/Unload is acknowledged back to the host. The Grid will manage the utilization of each cluster.
RUN	DEF	This ensures that Cluster 0's site will have a valid copy of the logical volume at job completion. Cluster 1 might have a valid copy at job completion, or sometime after job completion, depending on the virtual drive address selected and the override settings.
DEF	RUN	This ensures that Cluster 1's site will have a valid copy of the logical volume at job completion. Cluster 0 might have a valid copy at job completion, or sometime after job completion, depending on the virtual drive address selected and the override settings.
DEF	DEF	Use this setting if you do not care to which cluster the initial creation of the virtual volume is directed. A copy will be made as soon as possible after Rewind/Unload is acknowledged back to the host.
RUN	NC	Use this setting to specify Cluster 0 to have the initial valid copy when the Rewind/Unload is acknowledged back to the host. No copy will be made to Cluster 1. Note that the "force" override and virtual device address might have the system create a copy at Cluster 1.
NC	RUN	Use this setting to specify Cluster 1 to have the initial valid copy when the Rewind/Unload is acknowledged back to the host. No copy will be made to Cluster 0.
DEF	NC	Use this setting to specify Cluster 0 to have the initial valid copy when the Rewind/Unload is acknowledged back to the host. No copy will be made to Cluster 1.
NC	DEF	Use this setting to specify Cluster 1 to have the initial valid copy when the Rewind/Unload is acknowledged back to the host. No copy will be made to Cluster 0.
NC	NC	This setting is not supported, because it implies that the volumes should not be consistent anywhere. You will receive an error message when trying to specify this combination on the Library Manager.
RUN: Copy is created at rewind-unload time. DEF: Copy is created deferred. NC: No copy is created.		

Note: A single-cluster grid (stand-alone TS7700 Virtualization Engine) always uses Rewind/Unload (RUN) as the Data Consistency Point.

All Data Consistency rules shown in the Table 4-3 on page 247 apply to three-cluster grid and four-cluster grid configuration.

The Management Class window shows all clusters in Grid and the Data Consistency settings for each defined management class. See Figure 4-56 for an example of a Management Class definition in a four-cluster grid. The following terminology applies:

- ▶ **RUN:** The designated cluster has a consistent copy when rewind-unload operation is acknowledged back to the host.
- ▶ **DEF:** The designated cluster has a consistent copy sometime after job completion.
- ▶ **NC:** No copy is made to the designated cluster.

IBM System Storage™ TS7700 Virtualization Engine

Welcome admin

Grid Summary > "Kramer[0]" (#BA82A): Management Classes

Grid Summary > "Kramer[0]" Cluster Summary > Management Classes

Refresh Last Refresh: Thursday, October 22, 2009 8:26:23 AM

Current cluster selected:

- Clusters
 - "Kramer[0]" (#BA82A)
 - "Morik[1]" (#BA82B)
 - "Gilligan[2]" (#BA82C)
 - "Skipper[3]" (#BA82D)

Management Classes:

Select	Name	Secord	Description	Retain Copy Mcl	"Kramer[0]" (#BA82A)	"Morik[1]" (#BA82B)	"Gilligan[2]" (#BA82C)	"Skipper[3]" (#BA82D)
<input type="checkbox"/>	-----	0	Default management class	No	Rewind Unload (RUN)	No Copy	No Copy	No Copy
<input type="checkbox"/>	DDDDNNNN	0	D D D D	No	Deferred	Deferred	Deferred	Deferred
<input type="checkbox"/>	DDDDNNNN	0	D D D I	No	Deferred	Deferred	Deferred	Rewind Unload (RUN)
<input type="checkbox"/>	DDDDNNNN	0	D D D N	No	Deferred	Deferred	Deferred	No Copy
<input type="checkbox"/>	DDIDNNNN	0	D D I D	No	Deferred	Deferred	Rewind Unload (RUN)	Deferred
<input type="checkbox"/>	DDIINNIN	0	D D I I	No	Deferred	Deferred	Rewind Unload (RUN)	Rewind Unload (RUN)
<input type="checkbox"/>	DDIINNIN	0	D D I N	No	Deferred	Deferred	Rewind Unload (RUN)	No Copy
<input type="checkbox"/>	DDNDNNNN	0	D D N D	No	Deferred	Deferred	No Copy	Deferred
<input type="checkbox"/>	DDNDNNNN	0	D D N I	No	Deferred	Deferred	No Copy	Rewind Unload (RUN)
<input type="checkbox"/>	DDNNNNNN	0	D D N N	No	Deferred	Deferred	No Copy	No Copy

Page 1 of 9 1 Go Rows: 10 Show Total: 83 Filtered: 83 Displayed: 10

Figure 4-56 Defined Management classes in a four-cluster grid

Resources for information

See the following resources for details, configuration examples, and best practices for a multi-cluster grid:

- ▶ *IBM Virtualization Engine TS7700 Series Best Practices - TS7700 Hybrid Grid Usage:*
<http://w3.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101656>
- ▶ *IBM Virtualization Engine TS7700 Series Best Practices - Copy Consistency Points:*
<http://w3.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101230>

Define Copy Policy Override settings

With the TS7700 Virtualization Engine, you can define and set the optional override settings that influence the selection of the I/O Tape Volume Cache (TVC) and replication responses. The settings are specific to a cluster in a multi-cluster grid configuration, meaning that each cluster can have different settings if desired. The settings take effect for any mount requests received after the settings were saved. Mounts already in progress are not affected by a change in the settings. You can define and set the following settings:

- ▶ Prefer local cache for fast ready mount requests

- ▶ Prefer local cache for non-fast ready mount requests
- ▶ Force volumes mounted on this cluster to be copied to the local cache
- ▶ Allow fewer RUN consistent copies before reporting RUN command complete. You can view and modify these settings from the TS7700 Virtualization Engine management interface when you select **Copy Policy Override** from the Configuration drop-down menu, as shown in Figure 4-57.



Figure 4-57 Copy Policy Override

The settings you can select in the MI window are as follows:

- ▶ Prefer local cache for fast ready mount requests
A fast ready mount selects a local copy as long as a copy is available and a cluster copy consistency point is not specified as No Copy in the Management Class for the mount. The cluster is not required to have a valid copy of the data.
- ▶ Prefer local cache for non-fast ready mount requests
This override causes the local cluster to satisfy the mount request as long as the cluster is available and the cluster has a valid copy of the data, even if that data is only resident on physical tape. If the local cluster does not have a valid copy of the data, then default cluster selection criteria applies.
- ▶ Force volumes mounted on this cluster to be copied to the local cache
For a non-fast ready mount, this override causes a copy to be performed on the local cluster as part of mount processing. For a fast-ready mount, this setting has the effect of overriding the specified Management Class with a copy consistency point of rewind-unload for the cluster. This does not change the definition of the Management Class, but serves to influence the replication policy.
- ▶ Allow fewer RUN consistent copies before reporting RUN command complete
If selected, the value entered for "Number of required RUN consistent copies including the source copy" will be used to determine the number of copies to override before the Rewind/Unload operation reports as complete. If this option is not selected, the

Management Class definitions are to be used explicitly. Thus, the number of RUN copies can be from one to the number of clusters in the Grid.

Note: In a Geographically Dispersed Parallel Sysplex (GDPS), all three Copy Policy Override settings (cluster overrides for certain I/O and copy operations) must be selected on each cluster to ensure that wherever the GDPS primary site is, this TS7700 Virtualization Engine Cluster is preferred for all I/O operations.

If the TS7700 Virtualization Engine Cluster of the GDPS primary site fails, you must perform the following recovery actions:

1. Vary virtual devices from a remote TS7700 Virtualization Engine Cluster online from the primary site of the GDPS host.
2. Manually invoke, through the TS7700 Virtualization Engine management interface, a Read/Write Ownership Takeover, unless Automated Ownership Takeover Manager (AOTM) has already transferred ownership.

4.4.4 Defining cluster families

If you have a Grid with three or more clusters, you may define cluster families.

This function introduces a concept of grouping clusters together into families. Using cluster families, you will be able to define a common purpose or role to a subset of clusters within a Grid configuration. The role assigned, for example production or archive, will be used by the TS7700 microcode to make improved decisions for tasks such as replication and tape volume cache selection. For example, clusters in a common family are favored for tape volume cache selection, or replication can source volumes from other clusters within its family prior to using clusters outside of its family.

To view or modify cluster family settings, first verify that these permissions are granted to your assigned user role. If your user role includes cluster family permissions, click the **Modify** button to perform the following actions:

- ▶ Add a family
- ▶ Move a cluster
- ▶ Delete a family
- ▶ Save changes

Note: You cannot view or modify cluster family settings if permission to do so is not granted to your assigned user role.

Add a family

Click **Add** to create a new cluster family. A new cluster family placeholder is created to the right of any existing cluster families. Enter the name of the new cluster family in the active Name text box. Cluster family names must be 1 - 8 characters in length composed of Unicode characters. Each family name must be unique. To add a cluster to the new cluster family, move a cluster from the Unassigned Clusters area by following instructions in "Move a cluster" on page 251.

Note: You can create a maximum of eight cluster families.

Move a cluster

You can move one or more clusters between existing cluster families, to a new cluster family from the Unassigned Clusters area, or to the Unassigned Clusters area from an existing cluster family:

- ▶ Select a cluster: A selected cluster is identified by its highlighted border. Select a cluster from its resident cluster family or the Unassigned Clusters area by using one of the following ways:
 - Clicking on the cluster.
 - Pressing the Spacebar.
 - Pressing Shift while selecting clusters to select multiple clusters at one time.
 - Pressing Tab to switch between clusters before selecting a cluster.
- ▶ Move the selected cluster or clusters by one of the following ways:
 - Clicking on a cluster and dragging it to the destination cluster family or the Unassigned Clusters area.
 - Using the keyboard arrow keys to move the selected cluster or clusters right or left.

Note: An existing cluster family cannot be moved within the Cluster families page.

Delete a family

You can delete an existing cluster family. Click the **X** in the top, right corner of the cluster family you want to delete. If the cluster family you attempt to delete contains any clusters, a warning message is displayed. Click **OK** to delete the cluster family and return its clusters to the Unassigned Clusters area. Click **Cancel** to abandon the delete action and retain the selected cluster family.

Save changes

Click **Save** to save any changes made to the Cluster families page and return it to read-only mode.

Note: Each cluster family must contain at least one cluster. If you attempt to save changes and a cluster family does not contain any clusters, an error message is displayed and the Cluster families page remains in edit mode.

Cluster family configuration

Figure 4-58 illustrates the actions to create a family.

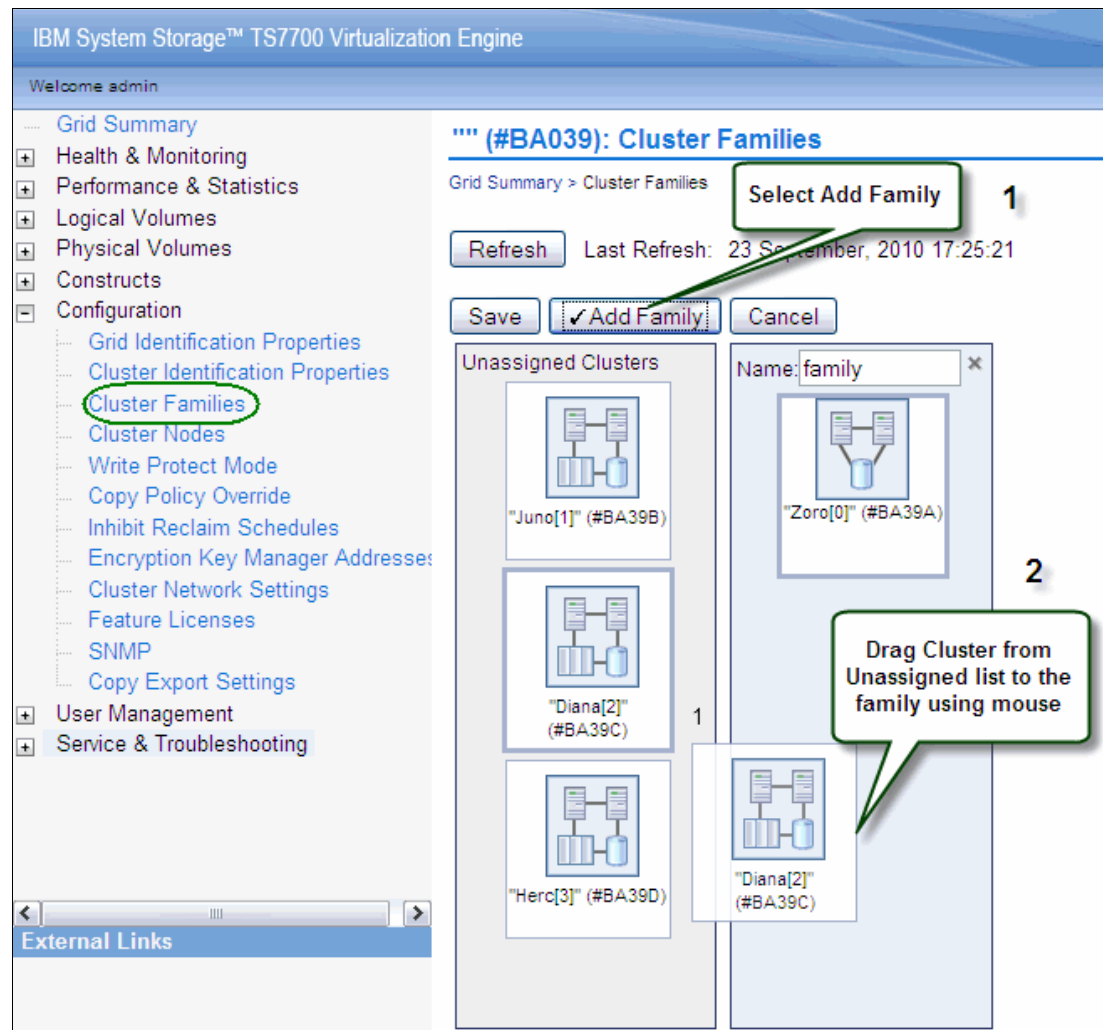


Figure 4-58 Creating a cluster family

Figure 4-59 shows an example of a cluster family configuration.

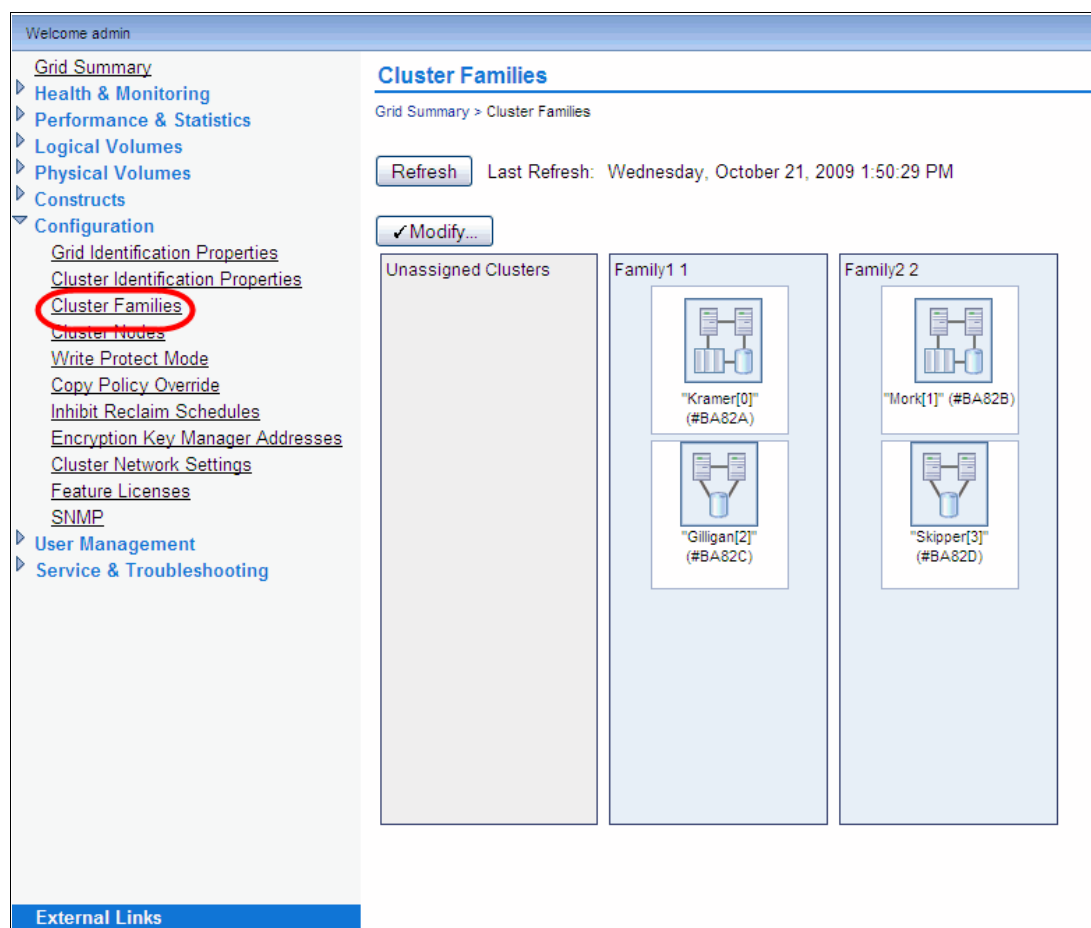


Figure 4-59 Cluster families

Note: Each cluster family should contain at least one cluster.

4.4.5 Retain Copy Mode

Retain Copy Mode is an optional setting where a volume's existing copy consistency points are honored instead of applying the copy consistency points defined at the mounting cluster. This applies to private volume mounts for reads or write appends. It is used to prevent more copies of a volume being created in the grid than desired. This is important in a grid with three or more clusters that has two or more clusters online to a host.

This parameter is set into Management Class panel for each Management Class. Figure 4-60 shows the Add Management Class panel and the Retain Copy Mode check box.

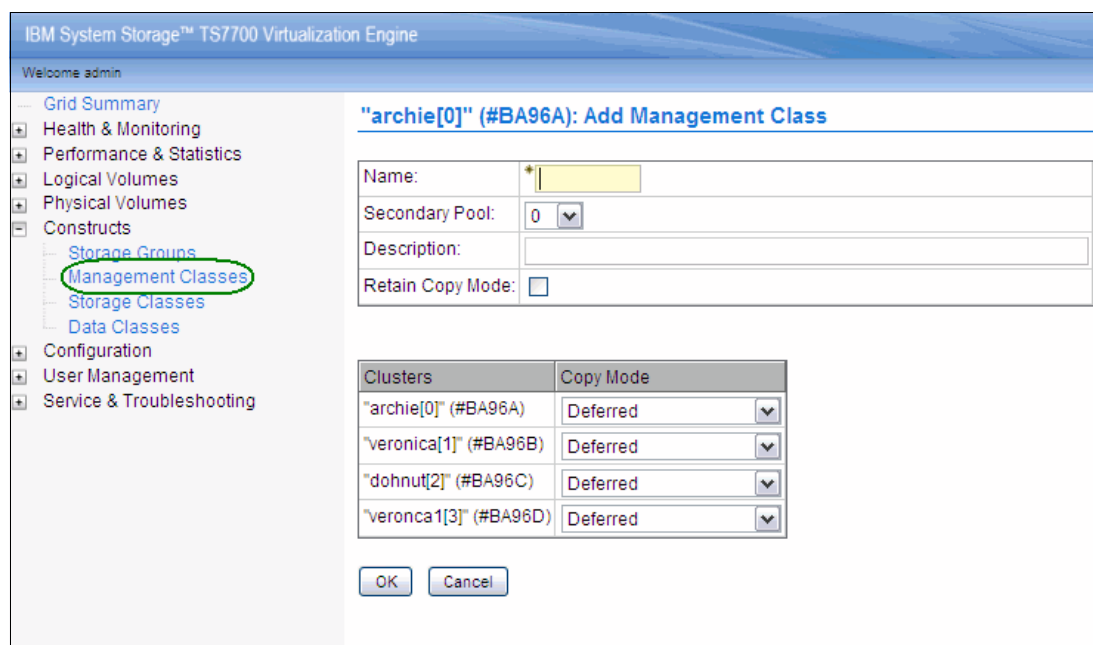


Figure 4-60 Retain Copy Mode in the Add Management Class screen

4.4.6 Automatic and Enhanced Removal Policies for grids

This TS7720 policy is to support grid configurations where there is a mixture of TS7720 and TS7740 clusters. Because the TS7720 has a maximum storage space that is the size of its tape volume cache, after that cache fills, this removal policy allows logical volumes to be automatically removed from cache as long as there is another consistent copy in the grid, such as on physical tape associated with a peer TS7740 or in another TS7720 tape volume cache. In essence, when coupled with the copy policies, it provides an automatic data removal function for the TS7720s.

Release 1.7 introduces enhancements that build upon the TS7700 Hybrid removal policy introduced in release 1.6. These enhancements allow you more control over the removal of content from a TS7720 as the active data reaches full capacity. You can now specify a minimum retention time for which a volume will remain in a TS7720 disk cache and not be removed. You also can tag each volume within a TS7720 with a cache preference:

- ▶ Pinned: Never remove from cache.
- ▶ Prefer Remove: When space is needed, you may now specify a minimum retention time which a volume will remain in a TS7720 disk cache and not be removed.
- ▶ Prefer Keep: When space is needed and all “Prefer Remove” candidates have been removed, these volumes will be the next candidates for removal in least recently used (LRU) order. The minimum retention time must have elapsed before either a “Prefer Remove” or “Prefer Keep” volume becomes a candidate for removal. Also, the TS7720 removes all volumes, which have been returned to scratch first prior to removing any “Prefer Remove” or “Prefer Keep” volumes.

See the following resources:

- ▶ *IBM Virtualization Engine TS7700 Series Best Practices Cache Management in the TS7720 white paper:*

<http://w3.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101382>

- ▶ Chapter 3, “Preinstallation planning and sizing” on page 107 for more details

The removal policy is set using the Storage Class panel on the TS7720 management interface. Figure 4-61 shows several definitions in place.

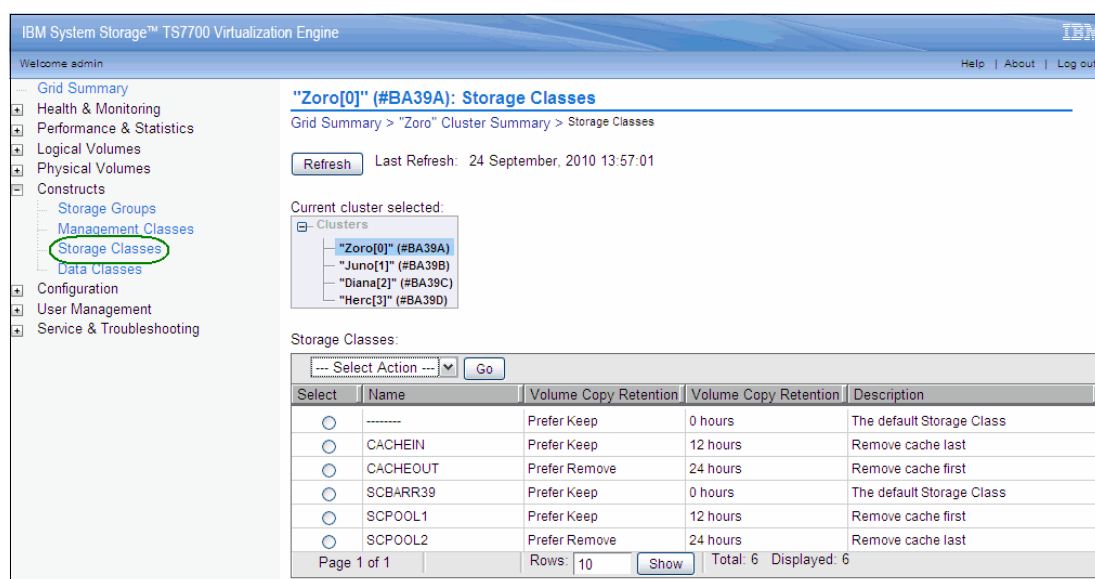


Figure 4-61 Storage Classes in TS7720 with Removal Policies.

To add or change an existent Storage Class, select appropriate action in the drop-down menu and click **Go**. See Figure 4-62.

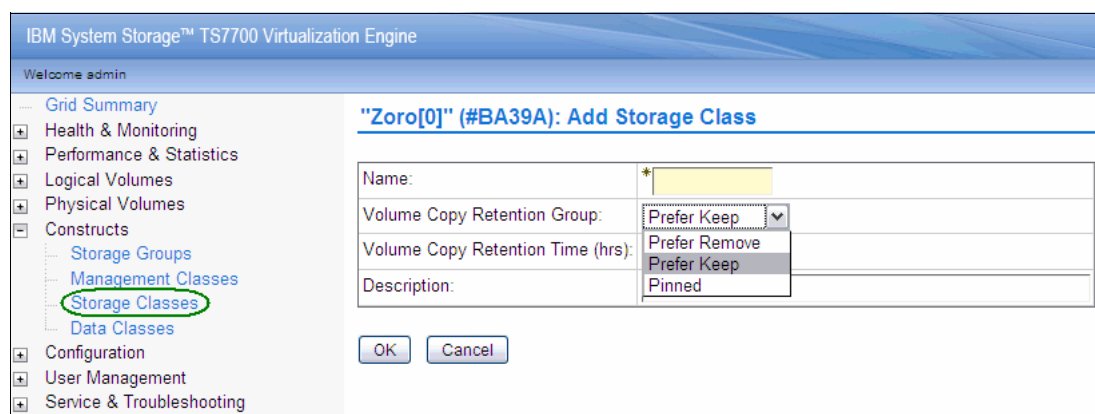


Figure 4-62 Defining a new Storage Class with TS7720

4.4.7 Temporary Removal Threshold

This setting is used in a hybrid grid to allow the TS7740 to be taken into service mode for a period without having the TS7720 cache fill up. This temporary threshold value sets a lower threshold for when the Automatic Removal Policy removes volumes from the TS7720 cache.

This setting is typically used when the last or only TS7740 in the grid is to be taken down. With R 1.7 it can also be used in a 7720 only grid.

See the following resources:

- ▶ *IBM Virtualization Engine TS7700 Series Best Practices Cache Management in the TS7720 white paper:*
<http://w3.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101382>
- ▶ Chapter 3, “Preinstallation planning and sizing” on page 107

The temporary removal threshold is set independently for each TS7720 in the grid by using the management interface. Go to the Tape Volume Cache in the TS7720 cluster, set the appropriate Temporary Removal Threshold, and click **Submit Changes**. See Figure 4-63.

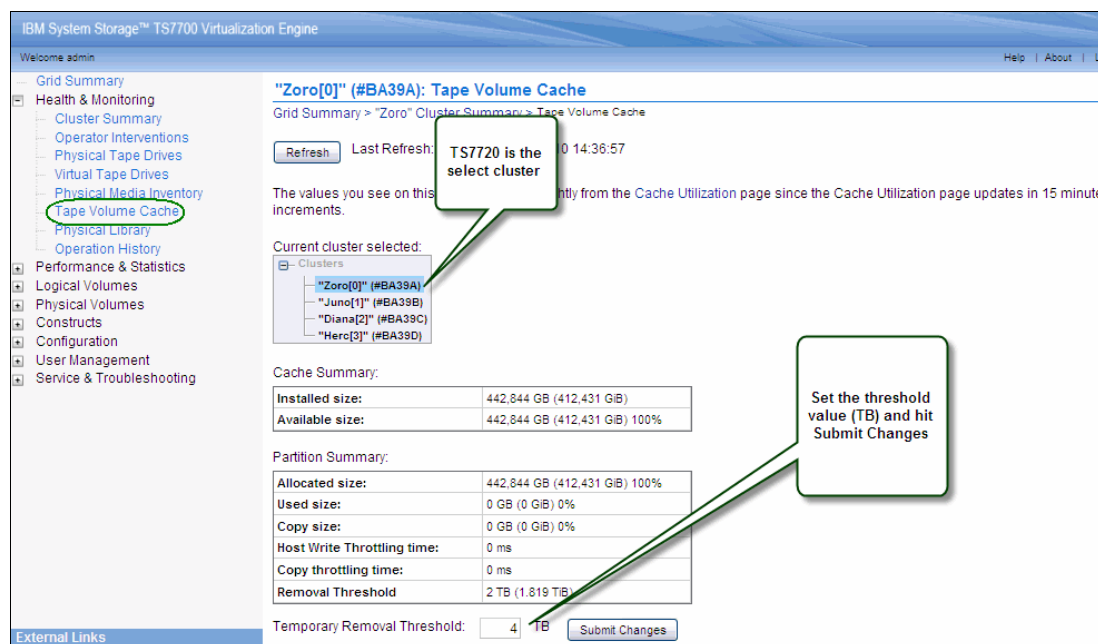


Figure 4-63 Setting Temporary Remove Threshold in a TS7720

After Temporary Removal Threshold is already defined into TS7720, you can use the function in preparation for TS7740 planned outage.

Using management interface, go to the Service Mode page. Be sure to select the TS7740 cluster.

See Figure 4-64 on page 257 for reference. The Service Mode panel contains a Lower Threshold button. If you click this button, the Confirm Lower Threshold window opens. If confirmed, the volume removal on TS7720 begins.

Note: The Lower Threshold button is active only on TS7740 clusters.

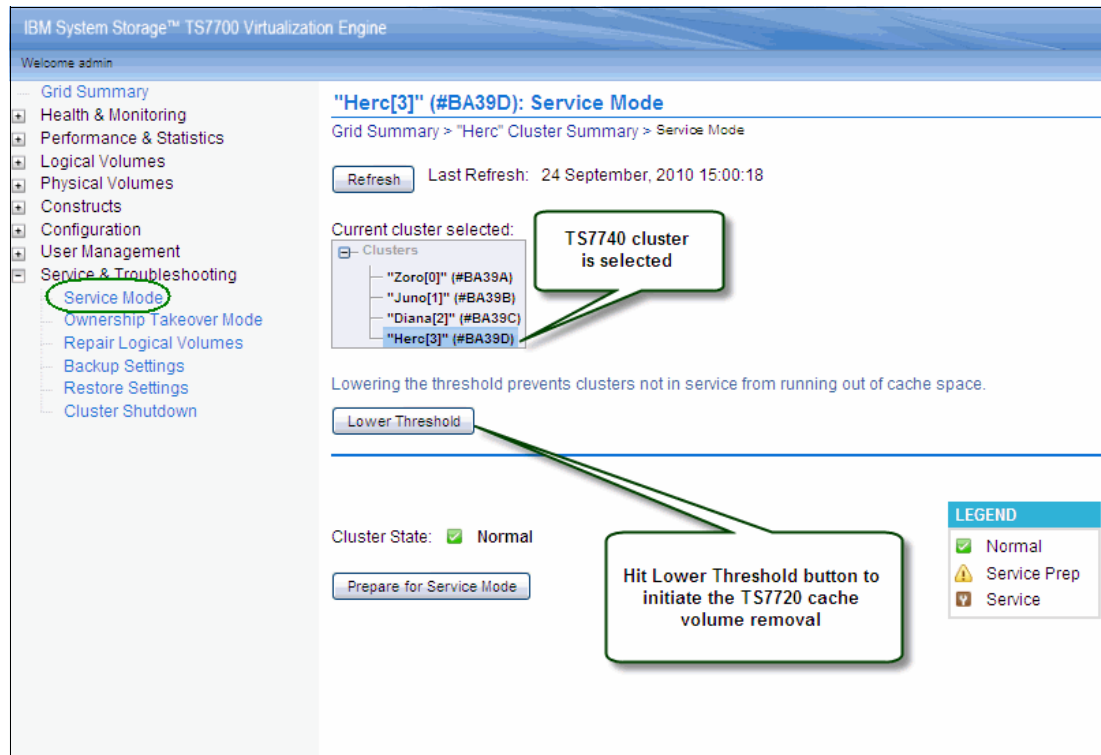


Figure 4-64 Lower Threshold button

Volume Removal Operation progress can be viewed using the Operation History page in the cluster that operation was initiated, as in Figure 4-65.

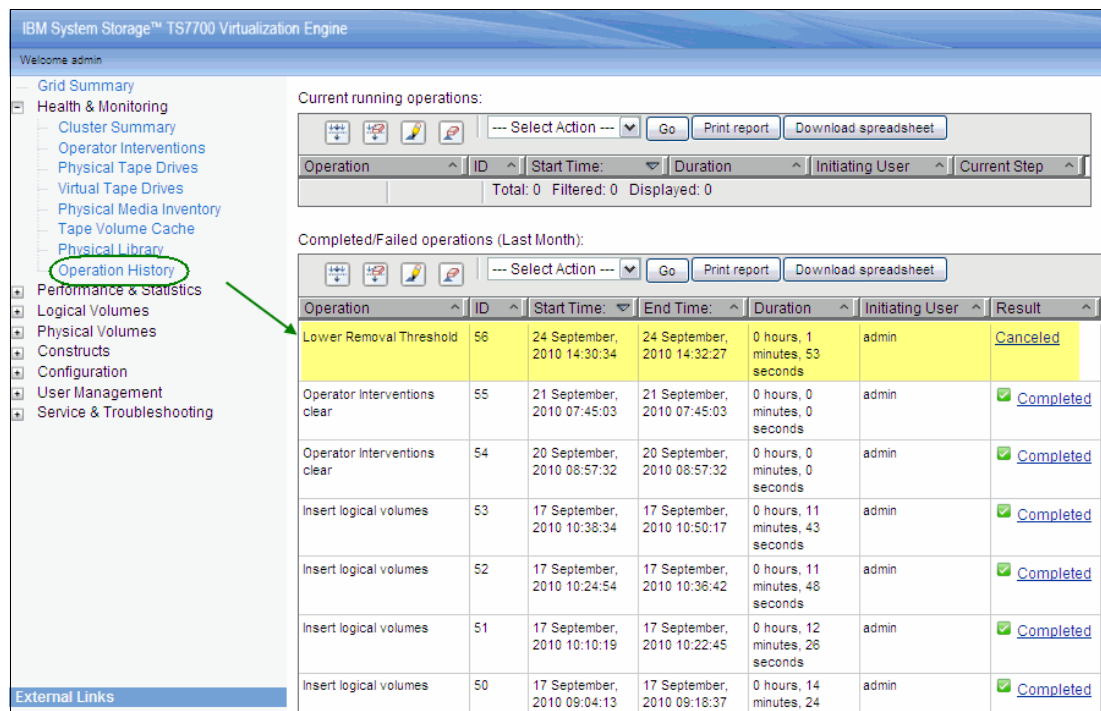


Figure 4-65 Progress panel

4.4.8 Backup and restore of construct definitions

Using the management interface, you can back up and restore fast-ready categories. Physical Volume Pools are also All constructs definitions, as shown in Figure 4-66.

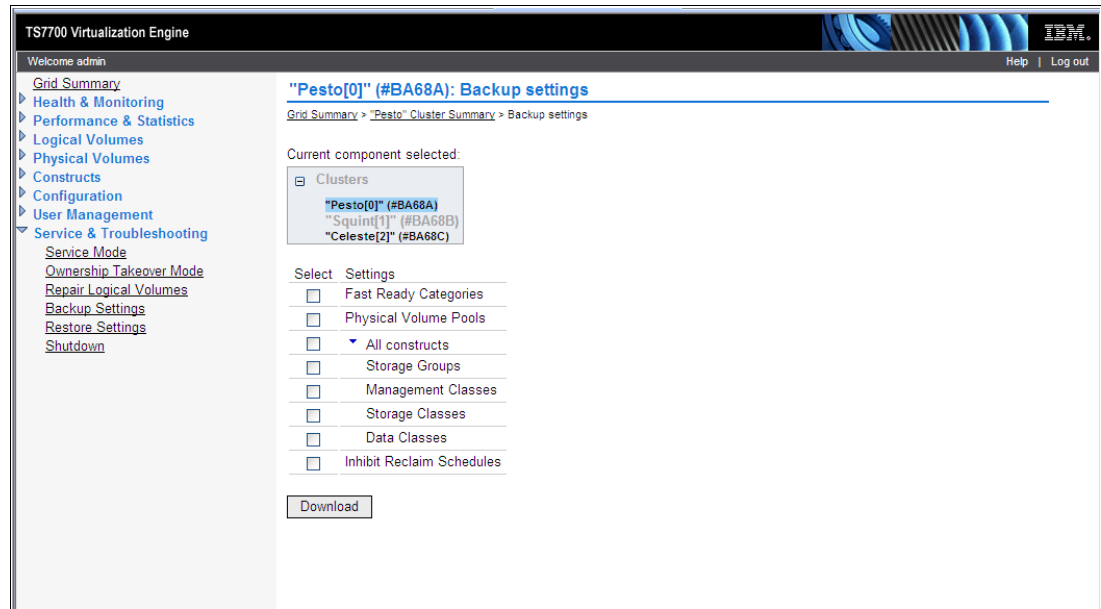


Figure 4-66 Backup settings

Note: See Chapter 8, "Operation" on page 409 for complete details of this window.

4.4.9 Data management settings (TS7740 Virtualization Engine only)

The following settings for the Virtualization Engine TS7740 Virtualization Engine are optional. Your IBM System Service Representative configures these settings during installation of the TS7740 Virtualization Engine or at a later time through the TS7740 Virtualization Engine SMIT menu. These data management settings are as follows:

- ▶ Copy Files Preferred to Reside in Cache
- ▶ Recalls Preferred for Cache Removal

Copy Files Preferred to Reside in Cache

Normally, the Tape Volume Cache (TVC) in both Virtualization Engine TS7740 Virtualization Engines in a multi-cluster grid is managed as one to increase the likelihood that a needed volume will be in cache. By default, the volume on the TS7740 Virtualization Engine selected for I/O operations is preferred to stay in cache on that TS7740 Virtualization Engine, whereas the copy made on the other TS7740 Virtualization Engine is preferred to be removed from cache:

- ▶ Preferred to stay in cache means that when space is needed for new volumes, the oldest volumes are removed first. This algorithm is called the least recently used (LRU) algorithm. This is also referred to as Preference Group 1 (PG1).
- ▶ Preferred to be removed from cache means that when space is needed for new volumes, the largest volumes are removed first, regardless of when they were written to the cache. This is also referred to as Preference Group 0 (PG0).

For a TS7740 Virtualization Engine running in a dual production multi-cluster grid configuration, both TS7740 Virtualization Engines are selected as the I/O TVCs and will have the original volumes (newly created or modified) preferred in cache while the copies to the other TS7740 Virtualization Engine will be preferred to be removed from cache. The result is that each TS7740 Virtualization Engine TVC is filled with unique, newly created or modified volumes, thereby roughly doubling the amount of cache seen by the host.

For a TS7740 Virtualization Engine running in a multi-cluster grid configuration used for business continuance, particularly when all I/O is preferenced to the local Tape Volume Cache (TVC), this default management method might not be desired. In the case where the remote site of the multi-cluster grid is used for recovery, the recovery time is minimized by having most of the needed volumes already in cache. What is really needed is to have the most recent copy volumes remain in the cache, not being preferred out of cache.

Based on customer requirements, an IBM System Service Representative (SSR) can set or modify this control through the TS7740 Virtualization Engine SMIT menu for the remote TS7740 Virtualization Engine, where:

- ▶ The default is off.
- ▶ When set to off, copy files are managed as Preference Group 0 volumes (prefer out of cache first by largest size).
- ▶ When set to on, copy files are managed based on the Storage Class construct definition.

Recalls Preferred for Cache Removal

Normally, a volume recalled into cache is managed as though it were newly created or modified, because it resides in the TS7740 Virtualization Engine selected for I/O operations on the volume. A recalled volume will displace other volumes in cache.

In the case where the remote TS7740 Virtualization Engine is used for recovery, the recovery time is minimized by having most of the needed volumes in cache. However, it is not likely that all of the volumes to restore will be resident in the cache and that some amount of recalls will be required. Unless you can explicitly control the sequence of volumes to be restored, it is likely that recalled volumes will displace cached volumes that have not yet been restored from, resulting in further recalls at a later time in the recovery process.

After a restore has been completed from a recalled volume, that volume is no longer needed, and such volumes should be removed from the cache after they have been accessed so that they minimally displace other volumes in the cache.

Based on customer requirements, the IBM System Service Representative (SSR) can set or modify this control through the TS7700 Virtualization Engine SMIT menu of the remote TS7740 Virtualization Engine, where:

- ▶ When off the default, recalls are managed as Preference Group 1 volumes (LRU).
- ▶ When on, recalls are managed as Preference Group 0 volumes (prefer out of cache first by largest size).

This control is independent of and not affected by cache management controlled through the Storage Class SMS construct also. Storage Class cache management affects only how the volume is managed in the I/O Tape Volume Cache (TVC).

4.5 Implementing Outboard Policy Management for non-z/OS hosts

Outboard Policy Management and its constructs are exploited only in DFSMS host environments where OAM has knowledge of the construct names and dynamically assigns and resets them. z/VM, z/VSE, TPF, z/TPF, and other hosts do not have knowledge of the construct names and therefore cannot change them. In addition, non-z/OS hosts use multiple Library Manager (LM) categories for scratch volumes and therefore can use multiple logical scratch pools on the Library Manager, as shown in Table 4-4.

Table 4-4 Scratch pools and Library Manager volume categories

Host software	Library Manager scratch categories	Number of scratch pools	Library Manager private categories
VM (+ VM/VSE)	X'0080' - X'008F'	16	X'FFFF'
BTLS	X'0FF2' - X'0FF8', X'0FFF'	8	X'FFFF
Native VSE	X'00A0' - X'00BF'	32	X'FFFF
Unassigned (can be used by Open Systems)	X'012F' - X'0FF1' X'0FF9' - X'0FFE' X'F00F' - X'FEFF'		

Note: In a TPF environment, manipulation of construct names for volumes can occur when they are moved from scratch through a user exit. The user exit allows the construct names and clone VOLSER to be altered. If the exit is not implemented, TPF does not alter the construct names.

TPF use of categories is flexible and is difficult to fully explain. TPF allows each drive to be assigned a scratch category. Concerning private categories, each TPF has their own category that volumes are assigned to when they are mounted.

For more information about this topic, see the zTPF Information Center:

<http://publib.boulder.ibm.com/infocenter/tpfhelp/current/index.jsp>

Because the hosts do not know about constructs, they ignore static construct assignment, and the assignment is kept even when the logical volume is returned to scratch. Static assignment means that at insert time of logical volumes, they are assigned construct names also. Construct names can also be assigned later at any time.

Note: In a z/OS environment, OAM controls the construct assignment and will reset any static assignment made before using the TS7700 Virtualization Engine management interface. Construct assignments are also reset to blank, when a logical volume is returned to scratch.

To implement Outboard Policy Management for non-z/OS hosts attached to a TS7700 Virtualization Engine, perform the following steps:

1. Define your pools and constructs.
2. Insert your logical volumes into groups through the TS7700 Virtualization Engine management interface (MI), as described in “Inserting logical volumes” on page 240. You can assign the required static construct names during the insertion, as shown at the bottom part of the window in Figure 4-67.

On the left side of the MI, click **Logical Volumes** → **Insert Logical Volumes**. The panel in Figure 4-67 opens. Use the panel to insert logical volumes. Select the **Set Constructs** check box and fill in the construct names.

TS7700 Virtualization Engine

Welcome admin Help | Log out

[Grid Summary](#)
[Health & Monitoring](#)
[Performance & Statistics](#)
[Logical Volumes](#)
[Incoming Copy Queue](#)
[Recall Queue](#)
[Logical Volume Details](#)
[Insert Logical Volumes](#)
[Delete Logical Volumes](#)
[Modify Logical Volumes](#)
[Move Logical Volumes](#)
[Logical Volume Search](#)
[Fast Ready Categories](#)
[Physical Volumes](#)
[Constructs](#)
[Configuration](#)
[User Management](#)
[Service & Troubleshooting](#)

Insert Logical Volumes
[Grid Summary](#) > [Insert Logical Volumes](#)

[Refresh](#) Last Refresh: Freitag, 16. Januar 2009 13:06:56

Current availability across entire grid:

Currently Inserted:	140.772
Maximum Allowed:	1.000.000
Available Slots:	859.228

Show inserted volume ranges between and [Show](#)

Insert a new logical volume range:

*Starting volser: *Quantity
 ☐ Ending volser

Initially owned by:
 [v](#)

Media type:
☐ Cartridge System Tape (400 MiB)
☒ Enhanced Capacity Cartridge System Tape (800 MiB)

☒ Set Constructs

Storage Groups: [v](#) Storage Classes: [v](#) Data Classes: [v](#) Management Classes: [v](#)

[Insert](#)

External Links
[Launch Library 3584 Specialist](#)

Figure 4-67 Insert logical volumes by assigning static construct names

- If you want to modify existing VOLSER ranges and assign the required static construct names to the logical volume ranges through the “change existing logical volume” function, select **Logical Volumes** → **Modify Logical Volumes** to open the window shown in Figure 4-68, which allows you to change existing logical volumes.

TS7700 Virtualization Engine

Welcome admin Help | Log out

Grid Summary

- Health & Monitoring
- Performance & Statistics
- Logical Volumes**
 - Incoming Copy Queue
 - Recall Queue
 - Logical Volume Details
 - Insert Logical Volumes
 - Delete Logical Volumes
 - Modify Logical Volumes
 - Move Logical Volumes
 - Logical Volume Search
 - Fast Ready Categories
- Physical Volumes
- Constructs
- Configuration
- User Management
- Service & Troubleshooting

Modify Logical Volumes

Grid Summary > Modify Logical Volumes

Last Refresh: Freitag, 16. Januar 2009 13:09:24

The modify logical volumes function is only valid for logical volumes used by non-MVS hosts. MVS hosts will automatically assign constructs for logical volumes.

Show inserted volume ranges between and

From	To	Quantity	Media Type	Capacity
K00000	K29999	30.000	CST	400 MiB
K30000	K39999	10.000	ECCST	800 MiB
Z00000	Z09999	10.000	ECCST	800 MiB
Total: 17				

Volume Range:

*From: *To:

Storage Groups: Storage Classes: Data Classes: Management Classes:

External Links

[Launch Library 3584 Specialist](#)

Figure 4-68 Change existing logical volumes from the ETL Specialist

Define groups of logical volumes with the same construct names assigned and, during insert processing, direct them to different volume categories so that all volumes in one LM volume category have identical constructs assigned.

Host control is given through usage of the appropriate scratch pool. By requesting a scratch mount from a specific scratch category, the actions defined for the constructs assigned to the logical volumes in this category are executed at Rewind/Unload of the logical volume.



Software implementation

In this chapter, we describe how to implement the IBM Virtualization Engine TS7700 on the IBM System z host systems. The tasks you have to complete include defining the hardware to the hosts by using the hardware configuration definition (HCD) menus, and other, operating system-specific tasks, depending on the System z software you are using.

From a software perspective, differences exist between the IBM Virtualization Engine TS7740 and the IBM Virtualization Engine TS7720. If we do not indicate any specific differences, the implementation steps apply to both; otherwise, we explain what the differences are in each relevant step.

In a z/OS environment, DFSMS provides automated system management of tape libraries. In addition to the implementation steps described in this chapter, you can find a step-by-step implementation process in Appendix B, “IBM Virtualization Engine TS7700 implementation steps” on page 797.

5.1 Host implementation considerations

From a host perspective, in a Single Cluster configuration the TS7700 Virtualization Engine, as a subsystem, supports 16 tape control units, each with 16 IBM 3490E tape drives, for a total of 256 tape drives. In a two-cluster grid, 32 tape control units provide 512 virtual tape drives. A three-cluster grid provides the capability for 48 tape control units, with a total of 768 tape drives. With the release of TS7700 Virtualization Engine R1.6 and later, four-cluster grid configurations provide the capability for 64 tape control units, with a total of 1024 tape drives. Each TS7700 Virtualization Engine can be host attached through two or four FICON channels.

The host does not know whether it is dealing with “physical” 3490E tape drives or with the virtual 3490E tape drives of the TS7700 Virtualization Engine. Therefore, the TS7700 Virtualization Engine with virtual 3490E tape drives is defined just like multiple physical IBM 3490-C2A controllers with 16 addresses through the hardware configuration definition (HCD) interface.

Before you can use the TS7700 Virtualization Engine, you need to define it to the System z host through HCD. Because the virtual tape drives of the TS7700 Virtualization Engine are library resident, you must specify `LIBRARY=YES` in the define device parameters. If FICON directors are being installed for the first time, the directors themselves also have to be defined in the IOCP and HCD input/output definition file (IODF).

In a z/OS environment, you then must define the TS7700 Virtualization Engine logical library to SMS. Update the SMS constructs and ACS routines to direct mounts to the TS7700 Virtualization Engine. See 5.3, “TS7700 Virtualization Engine software definitions for z/OS” on page 285 for more details about the implementation in a System z environment. You might need to update Missing Interrupt Handler (MIH) values also, as described in 5.2.6, “Set values for the Missing Interrupt Handler” on page 284.

The software implementation steps for z/VM and z/VSE are described in 5.4, “Software implementation in z/VM and z/VSE” on page 303. For TPF-related implementation details, see 5.5, “Software implementation in Transaction Processing Facility” on page 310.

After defining the TS7700 Virtualization Engine to a system by whichever method, verify that the devices can be brought online. Also plan to update the expected IPL configuration so that this is automatic and coincides with the implementation of the TS7700 Virtualization Engine and host software.

5.1.1 Sharing the TS7700 Virtualization Engine by multiple hosts

Each logical library has its own library sequence number, which is used to define the logical library to the host. Each logical library, identified as a Composite Library when dealing with a TS7700 Virtualization Engine, looks like a separate library to the host. A TS7700 Virtualization Engine can be shared by multiple System z systems, VM VSE, and TPF systems.

Sharing can be achieved in two ways:

- ▶ By logically dividing it into different partitions (partitioning)
- ▶ By allowing all attached systems to sequentially access all physical and logical volumes (sharing)

Sharing of an IBM Automated Tape Library means that all attached hosts have the same access to all volumes in the tape library. To achieve this true sharing, you need to share the host control data sets, that is, the tape management system inventory, the catalog information, and the tape configuration database (TCDB), among the attached hosts. In a non-SMS environment, all systems must share the ICF catalog that contains the BTLS inventory.

In general, these requirements can be met only in a single-platform environment. In this configuration, only one global tape volume scratch pool is available.

5.1.2 Partitioning the TS7700 Virtualization Engine between multiple hosts

Partitioning is the solution if you need to dedicate the use of volume ranges to certain systems or complexes or different host platforms. Dividing one or more libraries into logical libraries is the easiest way to allow different hosts to access them. Each host or complex owns its own set of drives and volumes, which another system or complex cannot access without manual intervention. Each system knows only about its part of the library.

This partition is sustained in DEVSUPxx category definitions. Up to now, to modify a category value, you needed to change DEVSUPxx member and IPL system. A new DS QLIB,CATS, allow us to modify these category values without IPL. See “DEVSERV QUERY LIBRARY command” on page 282 for further information.

Partitioning is also appropriate for the attachment to a z/OS logical partition (LPAR) for testing. If there is a need for running a test environment with a date different from the actual date, as it was the case during Y2K tests, you should have a separate TCDB and tape management system inventory for the test complex.

5.1.3 Logical path considerations

The TS7700 Virtualization Engine attaches to the host system or systems through two or four FICON channels. Each FICON channel provides 256 logical paths. With a four FICON configuration, that will result in a total of 1024 logical paths per TS7700 Virtualization Engine.

To calculate the number of logical paths required in an installation use the following formula:

Number of logical paths per FICON channel = number of LPARs x number of CU

This formula assumes all LPARs access all control units in the TS7700 Virtualization Engine with all channel paths.

For example, if one logical partition (LPAR) has 16 CU defined, this means that we are using 16 logical paths of the 256 logical paths available on each FICON adapter. See Table 5-1 to see the how many logical paths are used in various scenarios.

Table 5-1 Logical paths per FICON channel

Number of CUs defined	Number of LPARs	Logical paths used	Maximum paths
16	8	128	256
16	16	256	256
8	32	256	256

If you feel that you could have more than 16 LPARs, or even if you have more, do not define more than eight CUs per LPAR.

The *FICON Native Implementation and Reference Guide*, SG24-6266, covers the planning and implementation of FICON channels and operating in FICON native (FC) mode. It also discusses the FICON and Fibre Channel architectures, terminology, and supported topologies.

Define one tape control unit (CU) in the HCD dialog for every 16 virtual devices. Up to eight channel paths may be defined to each control unit. A logical path might be thought of as a three element entity: a host port, a TS7700 Virtualization Engine port, and a logical control unit in the TS7700 Virtualization Engine.

Note: A reduction in the number of physical paths will reduce the throughput capability of the TS7700 Virtualization Engine and the number of available logical paths. A reduction in control units will reduce the number of virtual devices available for any individual host.

5.1.4 Library names, Library IDs, and port IDs

Library names, Library IDs, and port IDs are used to define the TS7700 Virtualization Engine to the host at the hardware, operating system, and storage management subsystem level. Some of these identifiers are also used by the IBM System Services Representative (SSR) in the hardware configuration phase of installation.

On the host side, this means that several definitions have to be made in HCD and others in SMS. See Table 5-2 for an example, and create a similar one during your planning phase. It will be used in later steps.

Table 5-2 lists examples of the Library names and IDs needed in a z/OS implementation.

Table 5-2 Sample of library names and IDs: Four Cluster TS7700 Virtualization Engine implementation

TS7700 Virtualization Engine virtual library names	SMS name ^a	LIBRARY-ID	HCD Define	SMS Define
IBMC1 (Composite)	IBMC1	C7401	Yes	Yes
IBMD1TU (Distributed Tucson)	IBMD1TU	D1312	No	Yes
IBMD1PH (Distributed Phoenix)	IBMD1PH	D1307	No	Yes
IBMD1SJ (Distributed San Jose)	IBMD1SJ	D1300	No	Yes
IBMD1AT (Distributed Atlanta)	IBMD1AT	D1963	No	Yes

a. The SMS name cannot start with a “V”.

Distributed Library name and Composite Library name

The Distributed Library name and the Composite Library name are defined to z/OS and DFSMS/MVS. The Composite Library name is linked to the Composite Library ID when defining the tape library to DFSMS, as shown in Figure 5-10 on page 287. In the same manner, the Distributed Library name is linked to the Distributed Library ID, as shown in Figure 5-11 on page 287. Use names similar to those listed in Table 5-2. Use the letter “C” to indicate the Composite Library Names and the letter “D” to indicate the Distributed Library Names. The Composite Library name and the Distributed Library name cannot start with the letter “V”.

The Distributed Library name and the Composite Library name are not directly tied to configuration parameters used by the IBM System Services Representative (SSR) during installation of the TS7700 Virtualization Engine. These names are not defined to the TS7700 Virtualization Engine hardware. However, to make administration easier, it is useful to associate the LIBRARY-IDs with the SMS Library names through the nickname setting in the TS7700 Virtualization Engine management interface (MI).

Note: As a best practice, match the Distributed and Composite Library names entered at the host with the aliases defined at the TS7700 Virtualization Engine MI. Although they do not have to be the same, it will simplify management of the subsystem.

LIBRARY-ID and LIBPORT-ID

LIBRARY-ID and LIBPORT-ID are z/OS HCD parameters that allow HCD to provide the composite library configuration information that is normally obtained by the operating system at IPL time. If the devices are unavailable during IPL, the HCD information allows the logical tape devices to be varied online (when they subsequently become available to the system) without reactivating the IODF.

Note: As a best practice, specify the LIBRARY-ID and LIBPORT-ID in your HCD/IOCP definitions, even in a Single-Cluster configuration. This reduces the likelihood of having to reactivate the IODF when the library is not available at IPL, and providing enhanced error recovery in some cases. It might also eliminate the need to IPL when you make changes to your I/O configuration. In a Multi Cluster configuration, the LIBRARY-ID and LIBPORT-IDs must be specified in HCD, as shown Table 5-6 on page 277.

Distributed Library ID

During installation planning, each cluster is assigned a unique, five-digit hexadecimal number, (that is, the sequence number). This is used during subsystem installation procedures by the IBM SSR. This is the *Distributed Library ID*. This sequence number is arbitrary, and can be selected by you. It should start with the letter D.

In addition to the letter D, you may use the last four digits of the hardware serial number if it only consists of hexadecimal characters. For each Distributed Library ID, it would be the last four digits of the TS7700 Virtualization Engine serial number.

If you are installing a new multi-cluster grid configuration, you might consider choosing LIBRARY-IDs that clearly identify the cluster and the Grid. The Distributed Library IDs of a four-cluster grid configuration could be:

Cluster 0	DA01A
Cluster 1	DA01B
Cluster 2	DA01C
Cluster 3	DA01D

The Composite Library ID for this four-cluster grid could then be CA010.

Note: Whether you are using your own or our nomenclature, the important point is that the subsystem identification should be clear. We use the “D” in the Distributed Library ID. Because the identifier that appears in all system messages is the SMS library name, we believe that it important to distinguish the source of the message through the SMS Library name.

The Distributed Library ID is not used in defining the configuration in HCD.

Composite Library ID

The Composite Library ID is defined during installation planning and is arbitrary. The LIBRARY-ID is entered by the IBM SSR into the TS7700 Virtualization Engine configuration during hardware installation. All TS7700 Virtualization Engines participating in a Grid will have the same Composite Library ID. In our example in “Distributed Library ID”, the Composite Library ID starts with a “C” for this five hex-character sequence number. The last four characters can be used to uniquely identify each Composite Library in a meaningful way. The sequence number must match the LIBRARY-ID used in the HCD library definitions and the LIBRARY-ID listed in the ISMF Tape Library definition panels.

Note: In all configurations, each LIBRARY-ID, whether Distributed or Composite, must be unique.

LIBPORT-ID

The LIBPORT-ID reflects the order in which the tape control units are configured to the TS7700 Virtualization Engine across all Distributed Libraries participating in a Composite Library. It also provides the tape drive pool ID, which is transparent and only used by allocation and JES3.

DEVICE SERVICES QTAPE command

In an existing installation, you can use the DEVSERV QTAPE system command to discover what to specify. All tape drives (logical or physical) connected to a given logical control unit (LCU, CU) have the same LIBPORT-ID. Therefore, you only have to issue the DS QT command once per control unit (for any logical device number in that string of 16).

The command syntax is shown in Example 5-1.

Example 5-1 QTAPE command

```
DS QT,devnum,1,RDC
```

Values in the command are as follows:

DS	Device service
QT	Query tape
devnum	Device address
1	Number of devices to be displayed
RDC	Read device characteristics

Figure 5-1 shows the output of a DS QT system command.

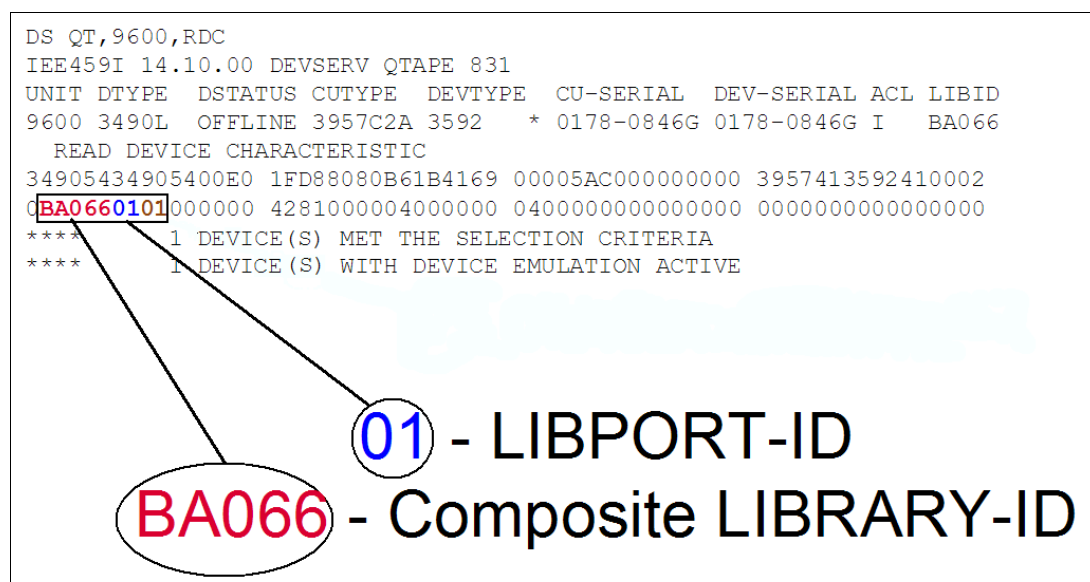


Figure 5-1 Sample DEVSERV QT command output

Note: The Distributed Library number or cluster index number for a given logical drive can be determined with the DS QT command. As identified in Figure 5-1, the response shows LIBPORT-ID 01 for logical drive 9600. LIBPORT-ID 01 is associated with Cluster 0. We will discuss the association between Distributed Libraries and LIBPORT-IDs in 5.2.1, “Defining devices through HCD for Cluster 0” on page 270.

From the DS QT command in Figure 5-1, you can derive the LIBRARY-ID for the Composite Library and the LIBPORT-ID of the logical control unit presenting the logical device. The real device type of the physical devices is unknown to the host and DEVSERV always shows 3592 as DEVTYPE. The LIBID field identifies the Composite Library ID associated with the device.

Note: You can get the real device type from the Host Console Request function LIBRARY REQUEST,<Distributed Library Name>,PDRIVE located in the Distributed Library.

The short form of LIBRARY REQUEST is LI REQ.

5.2 Hardware configuration definition

In this section, we describe the process of defining the TS7700 Virtualization Engine through the hardware configuration definition (HCD) interface for Cluster 0. The most important points to observe are as follows:

- ▶ HCD definitions are required.
- ▶ Up to sixteen 3490 tape control units can be defined per TS7700 Virtualization Engine, with 16 x 3490E drives each.
- ▶ Keep the link address blank when no FICON director is used.
- ▶ Specify LIBRARY=YES when using system-managed tape.

Usually HCD definitions are made by z/OS system administrators. A helpful approach is to fill in a table with all the definitions that the administrators need to be able to create HCD, and then give the table to the administrators. Table 5-3 is an example for Cluster 0. In general, all the blank cells should be filled in by system administrators, because they know what channels are free, what CU numbers are free, and so on.

Table 5-3 HCD definitions table for Cluster 0

CHPID	CU	CUADD	Link	Devices	ADD	LIB-ID	Libport
		0			00-0F		01
		1			00-0F		02
		2			00-0F		03
		3			00-0F		04
		4			00-0F		05
		5			00-0F		06
		6			00-0F		07
		7			00-0F		08
		8			00-0F		09
		9			00-0F		0A
		A			00-0F		0B
		B			00-0F		0C
		C			00-0F		0D
		D			00-0F		0E
		E			00-0F		0F
		F			00-0F		10

5.2.1 Defining devices through HCD for Cluster 0

You can define up to 16 control units with 16 devices each per cluster in the Grid configuration. Use CUADD=0 through CUADD=7 and LIBPORT-IDs of 01 through 08 for the first eight control units, as shown in Table 5-4.

Table 5-4 CUADD and LIBPORT-ID for the first set of 256 virtual devices

CU	1	2	3	4	5	6	7	8
CUADD	0	1	2	3	4	5	6	7
LIBPORT-ID	01	02	03	04	05	06	07	08

For the ninth to sixteenth control units, use CUADD=8 through CUADD=F and LIBPORT-IDs of 09 through 10, as shown in Table 5-5.

Table 5-5 CUADD and LIBPORT-ID for the second set of virtual devices

CU	9	10	11	12	13	14	15	16
CUADD	8	9	A	B	C	D	E	F
LIBPORT-ID	09	0A	0B	0C	0D	0E	0F	10

Figure 5-2 and Figure 5-3 on page 272 show the two important screens for specifying a tape control unit.

----- Add Control Unit -----

CBDPCU10

Specify or revise the following values.

Control unit number 0440 +

Control unit type 3490 +

Serial number _____

Description _____

Connected to switches . . . 01 01 01 01 _ _ _ _ +

Ports D6 D7 D8 D9 _ _ _ _ +

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

F1=Help F2=Split F3=Exit F4=Prompt F5=Reset F9=Swap

F12=Cancel

Figure 5-2 Adding the first TS7700 Virtualization Engine control unit through HCD: Part 1

Specify the control unit number and the type here (3490), as shown in Figure 5-2 on page 271, then press Enter. The screen shown in Figure 5-3 will be displayed; choose the processor to which the control unit is to be connected.

```

----- Add Control Unit -----
CBDPCU12

Specify or revise the following values.

Control unit number . : 0440          Type . . . . . : 3490
Processor ID . . . . . : PROC1        This is the main processor
Channel Subsystem ID . : 0

Channel path IDs . . . . 40   50   60   70   _   _   _   _   +
Link address . . . . . D6   D7   D8   D9   _   _   _   _   +

Unit address . . . . . 00   _   _   _   _   _   _   _   +
Number of units . . . . 16   _   _   _   _   _   _   _

Logical address . . . . 0   + (same as CUADD)

Protocol . . . . . _   + (D,S or S4)
I/O concurrency level . 2   + (1, 2 or 3)

F1=Help      F2=Split    F4=Prompt   F5=Reset    F9=Swap     F12=Cancel

```

Figure 5-3 Adding the first TS7700 Virtualization Engine control unit through HCD: Part 2

Note: When the TS7700 Virtualization Engine is not attached through FICON directors, the link address fields would be blank.

Repeating the previous process, you would define the second through sixteenth TS7700 Virtualization Engine virtual tape control unit, specifying the logical unit address (CUADD)=1 to F, in the Add Control Unit screens. The Add Control Unit summary screen is shown in Figure 5-3.

To define the TS7700 Virtualization Engine virtual drives, you need to use the Add Device screen shown in Figure 5-4.

```

----- Add Device -----
CBDPDV10

Specify or revise the following values.
Device number . . . . . 0A40 (0000 - FFFF)
Number of devices . . . . . 16
Device type . . . . . 3490 +

Serial number . . . . . 
Description . . . . . 

Connected to CUs . . 0440 +

F1=Help    F2=Split    F3=Exit    F4=Prompt    F5=Reset    F9=Swap F12=Cancel

```

Figure 5-4 Adding the first 16 drives through HCD

After entering the required information, you can specify to which processors and operating systems the devices are connected. Figure 5-5 shows the screen used to update the processor's view of the device.

```

----- Define Device / Processor-----
CBDPDV12

Specify or revise the following values.

Device number . : 0A40          Number of devices . . . . : 16
Device type . . : 3490
Processor ID. . : PROC1        This is the main processor

Unit address . . . . . 00 +(only necessary whrn different from
                           the last 2 digits of device number)
Time-Out . . . . . No (Yes or No)
STADET . . . . . No (Yes or No)

Preferred CHPID . . . . . +
Explicit device candidate list . No (Yes or No)

F1=Help    F2=Split    F4=Prompt    F5=Reset    F9=Swap    F12=Cancel

```

Figure 5-5 HCD Define Device / Processor screen

After entering the required information and specifying to which operating systems the devices are connected, the screen in Figure 5-6 is displayed, where you can update the device parameters.

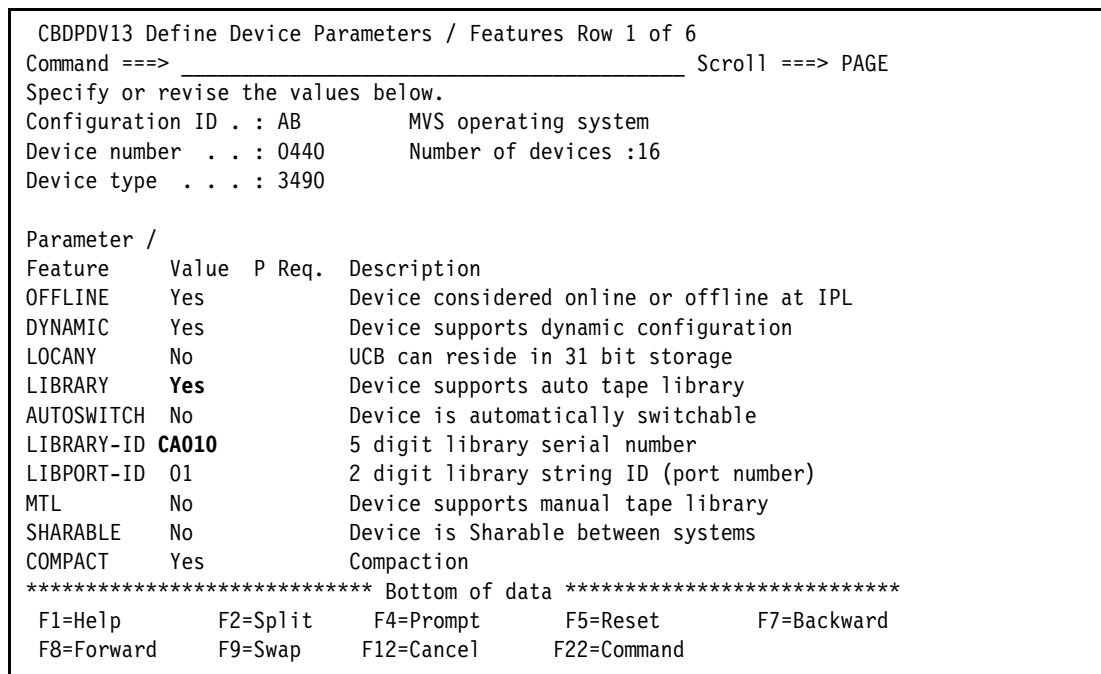


Figure 5-6 Define Device Parameters HCD screen

Notes:

- ▶ If you are defining drives that are installed in a system-managed IBM Tape Library, such as the TS7700 Virtualization Engine, you must specify LIBRARY=YES.
- ▶ If more than one System z host will be sharing the virtual drives in the TS7700 Virtualization Engine, specify SHARABLE=YES. This will force OFFLINE to YES. It is up to the installation to ensure proper serialization from all attached hosts.
- ▶ For a Single Cluster configuration, specify LIBRARY-ID and LIBPORT-ID. This will position the site for possible future Multi Cluster implementations.
- ▶ You must use the Composite Library ID of the TS7700 Virtualization Engine in your HCD definitions.
- ▶ The Distributed Library IDs are not defined in HCD.

To define the remaining TS7700 Virtualization Engine 3490E virtual drives, you need to repeat this process for each control unit in your implementation plan.

5.2.2 Activate the I/O configuration

There are differences in the concurrent IODF activation process between a new tape library implementation and a configuration change made to an existing library. Changes to the virtual devices address range of an existing library is an example of where concurrent IODF activation would be useful.

As an alternative to the procedures described below, you can always IPL the system.

Installing a new tape library

If you are installing a TS7700 Virtualization Engine for the first time, from a host software definition point of view, this is an installation of a new library. When you are activating the IODF for a new tape library, the following steps must be performed to get the tape library or TS7700 Virtualization Engine online without IPLing your systems:

1. Activate the IODF.
2. Run MVS VARY ONLINE to vary online the devices in the library. This will create some of the control blocks. You will see the following message:

```
IEA437I TAPE LIBRARY DEVICE(ddd), ACTIVATE IODF=xx, IS REQUIRED
```

3. Do the final ACTIVATE. This is required to build the Eligible Device Table (EDT) for MVS Allocation.

You can check the details using the DEVSERV QTape command, which provides information about Unit Control Block (UCB), UCB prefix, UCB common extension, Device Class Extension (DCE), and Read Device Characteristics (RDC) data and Read Configuration Data (RCD) data, which are data buffers acquired directly from the device.

Note: If you are just adding additional device address ranges to an existing TS7700 Virtualization Engine, you can use the same process as for a new tape library.

Modifications to an existing tape library

When you are modifying an existing tape library so that existing device addresses are to be changed, perform the following steps:

1. Activate an IODF that deletes all devices from the library.
2. Activate an IODF that defines all of the devices of the modified library.
3. Run MVS VARY ONLINE to vary online the devices in the library. This will create some of the control blocks. You will see the following message:

```
IEA437I TAPE LIBRARY DEVICE(ddd), ACTIVATE IODF=xx, IS REQUIRED
```

4. Do the final ACTIVATE.

Alternatively, you can use DS QL,nnnnn,DELETE (where nnnnn is the LIBID) command to delete the library's dynamic control blocks. If you have IODEF with LIBID and LIBPORT coded already, perform the following steps:

1. Use QLIB LIST to see if the INACTIVE control blocks have been deleted.
2. Use ACTIVATE IODF to redefine the devices.
3. Use QLIB LIST to verify that the ACTIVE control blocks are properly defined.

If LIBRARY-ID (LIBID) and LIBPORT-ID are not coded, perform the following steps:

1. MVS VARY ONLINE the devices in the library. This will create some control blocks, and you will see the following message:

```
IEA437I TAPE LIBRARY DEVICE(ddd), ACTIVATE IODF=xx, IS REQUIRED
```

2. Use QLIB LIST to verify that the ACTIVE control blocks are properly defined.

5.2.3 HCD considerations for multi-cluster grid operation

Each TS7700 Virtualization Engine presents 256 virtual device images for a total of 1024 virtual devices if used in a four-cluster grid configuration. Each TS7700 Virtualization Engine has 256 virtual devices with 16 logical control units (LCU 0-F). The host generates the physical paths to each site separately, so the host sees one Composite Library image, and four Distributed libraries. An example of a TS7700 Virtualization Engine three-cluster grid configuration is shown in Figure 5-7.

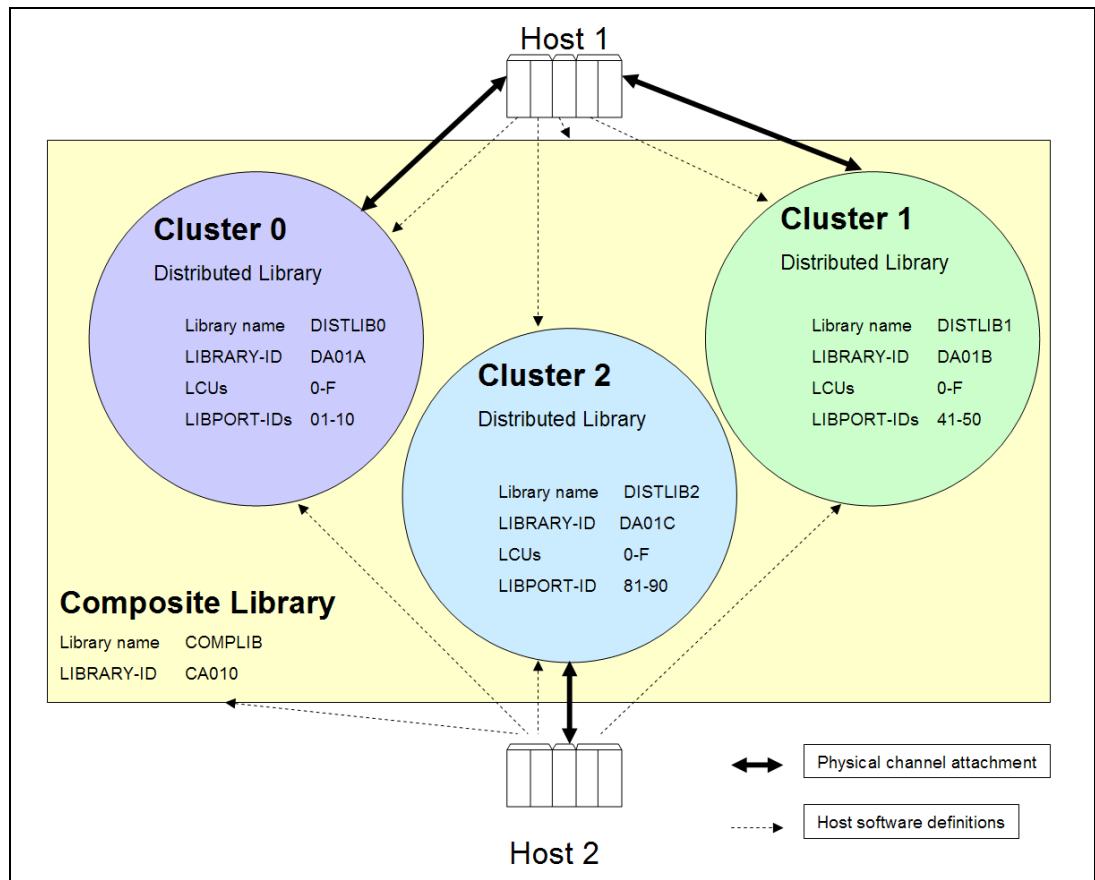


Figure 5-7 Example of a TS7700 Virtualization Engine three-cluster grid configuration

Using Figure 5-7 as an example, define Host 1 as having physical connections to Cluster 0 and Cluster 1. Cluster 2 and Host 2 are probably far away in a disaster recovery site and Host 2 has only physical connections to Cluster 2. You then configure Host 1 with 512 (2×256) 3490E drives and Host 2 with 256 3490E drives. In HCD, you use the LIBRARY-ID from the Composite Library (CA010). Host 1 and Host 2 have three Distributed Libraries and one Composite Library defined to SMS. Internally, the three clusters are connected with a customer supplied IP-network.

Logical control units and physical paths are defined on a vNode/gNode boundary, similar to the Virtual Tape Controllers (VTCs) in the previous generation of the Peer-to-Peer VTS. All of them are part of the same Composite Library image presented to the host. Table 5-6 shows the subsystem IDs (LIBPORT-IDs) that you must use for each cluster in a four-cluster grid configuration.

Table 5-6 LIBPORT-ID for each cluster in a four-cluster grid configuration

Cluster	Logical control units (LCUs)	LIBPORT-IDs (hex)
0	0-7	01-08
	8-F	09-10
1	0-7	41-48
	8-F	49-50
2	0-7	81-88
	8-F	89-90
3	0-7	C1-C8
	8-F	C9-D0

The definition steps are essentially the same as for a single-cluster grid configuration. The important difference is that you need to specify the range 41-50 for the LIBPORT-IDs of Cluster 1, LIBPORT-IDs 81-90 for Cluster 2, and LIBPORT-IDs C1-D0 for Cluster 3.

Note: The virtual device allocation for each cluster in a Dual Cluster Grid configuration is managed by the host. That means the host randomly picks a device from each cluster for an I/O operation based upon a host device allocation algorithm. Referencing Figure 5-7 on page 276, if the remote Cluster 2 is now attached through a limited bandwidth FICON connection to Host 1 or Host 2, it might negatively affect I/O performance. The possibility would exist that the remote Cluster 2 might be selected as the I/O cluster, even if the data is residing in the TVC of Cluster 0 or Cluster 1.

To avoid those situations, we recommend varying offline the remote virtual devices from each host's point for normal operation. Only in the case of disaster recovery should those remote virtual drives be varied online from the host.

In your installation, you might have to review the FICON switch redundancy and performance objectives. Policies are based on requirements for data availability and accessibility. Your local policies might require that all FICON equipment must be attached through two FICON switches, with half of the connections on each.

If you have two data centers and FICON switch equipment at both sites connected to the hosts, use a cascading FICON switch configuration to attach to your tape subsystems.

Note: For the latest information about supported FICON directors and TS7700 Virtualization Engine Licensed Internal Code levels, go to the following address:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/FQ116133>

An example of how to do cascading FICON attachment with two sites can be found at the following address:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS2844>

5.2.4 More HCD and IOCP examples: two-cluster grid

This section shows IOCP examples to help you get started.

Example 1: Cluster 0, Composite LIBID = C0001

Figure 5-8 shows the configuration for IOCP definition.

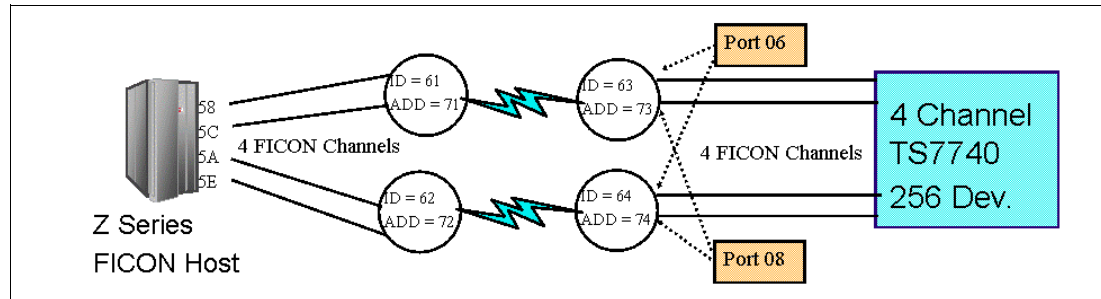


Figure 5-8 Example configuration used for IOCP definition

Example 5-2 shows Cluster 0 IOCP.

Example 5-2 Cluster 0 IOCP

```
RESOURCE PARTITION=(LPAR0,1)
CHPID PATH=(58),SHARED,SWITCH=61,TYPE=FC
CHPID PATH=(5A),SHARED,SWITCH=62,TYPE=FC
CHPID PATH=(5C),SHARED,SWITCH=61,TYPE=FC
CHPID PATH=(5E),SHARED,SWITCH=62,TYPE=FC
CNTLUNIT CUNUMBER=(440),PATH=(58,5A,5C,5E),UNIT=TAPE,UNITADD=((00,16)),LINK=(7306,7406,7308,7408),CUADD=0
IODEVICE ADDRESS=(2A00,16),UNIT=TAPE,CUNUMBER=(440),UNITADD=00,LIBID=C0001,LIBPORT=01
CNTLUNIT CUNUMBER=(441),PATH=(58,5A,5C,5E),UNIT=TAPE,UNITADD=((00,16)),LINK=(7306,7406,7308,7408),CUADD=1
IODEVICE ADDRESS=(2A10,16),UNIT=TAPE,CUNUMBER=(441),UNITADD=00,LIBID=C0001,LIBPORT=02
CNTLUNIT CUNUMBER=(442),PATH=(58,5A,5C,5E),UNIT=TAPE,UNITADD=((00,16)),LINK=(7306,7406,7308,7408),CUADD=2
IODEVICE ADDRESS=(2A20,16),UNIT=TAPE,CUNUMBER=(442),UNITADD=00,LIBID=C0001,LIBPORT=03
CNTLUNIT CUNUMBER=(443),PATH=(58,5A,5C,5E),UNIT=TAPE,UNITADD=((00,16)),LINK=(7306,7406,7308,7408),CUADD=3
IODEVICE ADDRESS=(2A30,16),UNIT=TAPE,CUNUMBER=(443),UNITADD=00,LIBID=C0001,LIBPORT=04
CNTLUNIT CUNUMBER=(444),PATH=(58,5A,5C,5E),UNIT=TAPE,UNITADD=((00,16)),LINK=(7306,7406,7308,7408),CUADD=4
IODEVICE ADDRESS=(2A40,16),UNIT=TAPE,CUNUMBER=(444),UNITADD=00,LIBID=C0001,LIBPORT=05
CNTLUNIT CUNUMBER=(445),PATH=(58,5A,5C,5E),UNIT=TAPE,UNITADD=((00,16)),LINK=(7306,7406,7308,7408),CUADD=5
IODEVICE ADDRESS=(2A50,16),UNIT=TAPE,CUNUMBER=(445),UNITADD=00,LIBID=C0001,LIBPORT=06
CNTLUNIT CUNUMBER=(446),PATH=(58,5A,5C,5E),UNIT=TAPE,UNITADD=((00,16)),LINK=(7306,7406,7308,7408),CUADD=6
IODEVICE ADDRESS=(2A60,16),UNIT=TAPE,CUNUMBER=(446),UNITADD=00,LIBID=C0001,LIBPORT=07
CNTLUNIT CUNUMBER=(447),PATH=(58,5A,5C,5E),UNIT=TAPE,UNITADD=((00,16)),LINK=(7306,7406,7308,7408),CUADD=7
IODEVICE ADDRESS=(2A70,16),UNIT=TAPE,CUNUMBER=(447),UNITADD=00,LIBID=C0001,LIBPORT=08
CNTLUNIT CUNUMBER=(448),PATH=(58,5A,5C,5E),UNIT=TAPE,UNITADD=((00,16)),LINK=(7306,7406,7308,7408),CUADD=8
IODEVICE ADDRESS=(2A80,16),UNIT=TAPE,CUNUMBER=(448),UNITADD=00,LIBID=C0001,LIBPORT=09
CNTLUNIT CUNUMBER=(449),PATH=(58,5A,5C,5E),UNIT=TAPE,UNITADD=((00,16)),LINK=(7306,7406,7308,7408),CUADD=9
IODEVICE ADDRESS=(2A90,16),UNIT=TAPE,CUNUMBER=(449),UNITADD=00,LIBID=C0001,LIBPORT=0A
CNTLUNIT CUNUMBER=(44A),PATH=(58,5A,5C,5E),UNIT=TAPE,UNITADD=((00,16)),LINK=(7306,7406,7308,7408),CUADD=A
IODEVICE ADDRESS=(2AA0,16),UNIT=TAPE,CUNUMBER=(44A),UNITADD=00,LIBID=C0001,LIBPORT=0B
CNTLUNIT CUNUMBER=(44B),PATH=(58,5A,5C,5E),UNIT=TAPE,UNITADD=((00,16)),LINK=(7306,7406,7308,7408),CUADD=B
IODEVICE ADDRESS=(2AB0,16),UNIT=TAPE,CUNUMBER=(44B),UNITADD=00,LIBID=C0001,LIBPORT=0C
CNTLUNIT CUNUMBER=(44C),PATH=(58,5A,5C,5E),UNIT=TAPE,UNITADD=((00,16)),LINK=(7306,7406,7308,7408),CUADD=C
IODEVICE ADDRESS=(2AC0,16),UNIT=TAPE,CUNUMBER=(44C),UNITADD=00,LIBID=C0001,LIBPORT=0D
CNTLUNIT CUNUMBER=(44D),PATH=(58,5A,5C,5E),UNIT=TAPE,UNITADD=((00,16)),LINK=(7306,7406,7308,7408),CUADD=D
IODEVICE ADDRESS=(2AD0,16),UNIT=TAPE,CUNUMBER=(44D),UNITADD=00,LIBID=C0001,LIBPORT=0E
CNTLUNIT CUNUMBER=(44E),PATH=(58,5A,5C,5E),UNIT=TAPE,UNITADD=((00,16)),LINK=(7306,7406,7308,7408),CUADD=E
IODEVICE ADDRESS=(2AE0,16),UNIT=TAPE,CUNUMBER=(44E),UNITADD=00,LIBID=C0001,LIBPORT=0F
CNTLUNIT CUNUMBER=(44F),PATH=(58,5A,5C,5E),UNIT=TAPE,UNITADD=((00,16)),LINK=(7306,7406,7308,7408),CUADD=F
IODEVICE ADDRESS=(2AF0,16),UNIT=TAPE,CUNUMBER=(44F),UNITADD=00,LIBID=C0001,LIBPORT=10
```


Example 2: Cluster 1, Composite LIBID = C0001

Figure 5-9 shows cluster configuration for IOCP.

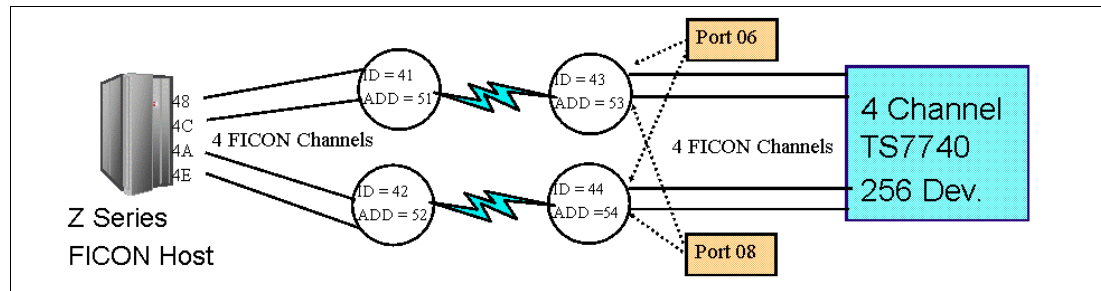


Figure 5-9 Cluster configuration for IOCP

Example 5-3 shows Cluster 1 IOCP.

Example 5-3 Cluster 1 IOCP

```
RESOURCE PARTITION=(LPAR0,1)
CHPID PATH=(48),SHARED,SWITCH=41,TYPE=FC
CHPID PATH=(4A),SHARED,SWITCH=42,TYPE=FC
CHPID PATH=(4C),SHARED,SWITCH=41,TYPE=FC
CHPID PATH=(4E),SHARED,SWITCH=42,TYPE=FC
CNTLUNIT CUNUMBER=(480),PATH=(48,4A,4C,4E),UNIT=TAPE,UNITADD=((00,16)),LINK=(5306,5406,5308,5408),CUADD=0
IODEVICE ADDRESS=(2B00,16),UNIT=TAPE,CUNUMBER=(480),UNITADD=00,LIBID=C0001,LIBPORT=41
CNTLUNIT CUNUMBER=(481),PATH=(48,4A,4C,4E),UNIT=TAPE,UNITADD=((00,16)),LINK=(5306,5406,5308,5408),CUADD=1
IODEVICE ADDRESS=(2B10,16),UNIT=TAPE,CUNUMBER=(481),UNITADD=00,LIBID=C0001,LIBPORT=42
CNTLUNIT CUNUMBER=(482),PATH=(48,4A,4C,4E),UNIT=TAPE,UNITADD=((00,16)),LINK=(5306,5406,5308,5408),CUADD=2
IODEVICE ADDRESS=(2B20,16),UNIT=TAPE,CUNUMBER=(482),UNITADD=00,LIBID=C0001,LIBPORT=43
CNTLUNIT CUNUMBER=(483),PATH=(48,4A,4C,4E),UNIT=TAPE,UNITADD=((00,16)),LINK=(5306,5406,5308,5408),CUADD=3
IODEVICE ADDRESS=(2B30,16),UNIT=TAPE,CUNUMBER=(483),UNITADD=00,LIBID=C0001,LIBPORT=44
CNTLUNIT CUNUMBER=(484),PATH=(48,4A,4C,4E),UNIT=TAPE,UNITADD=((00,16)),LINK=(5306,5406,5308,5408),CUADD=4
IODEVICE ADDRESS=(2B40,16),UNIT=TAPE,CUNUMBER=(484),UNITADD=00,LIBID=C0001,LIBPORT=45
CNTLUNIT CUNUMBER=(485),PATH=(48,4A,4C,4E),UNIT=TAPE,UNITADD=((00,16)),LINK=(5306,5406,5308,5408),CUADD=5
IODEVICE ADDRESS=(2B50,16),UNIT=TAPE,CUNUMBER=(485),UNITADD=00,LIBID=C0001,LIBPORT=46
CNTLUNIT CUNUMBER=(486),PATH=(48,4A,4C,4E),UNIT=TAPE,UNITADD=((00,16)),LINK=(5306,5406,5308,5408),CUADD=6
IODEVICE ADDRESS=(2B60,16),UNIT=TAPE,CUNUMBER=(486),UNITADD=00,LIBID=C0001,LIBPORT=47
CNTLUNIT CUNUMBER=(487),PATH=(48,4A,4C,4E),UNIT=TAPE,UNITADD=((00,16)),LINK=(5306,5406,5308,5408),CUADD=7
IODEVICE ADDRESS=(2B70,16),UNIT=TAPE,CUNUMBER=(487),UNITADD=00,LIBID=C0001,LIBPORT=48
CNTLUNIT CUNUMBER=(488),PATH=(48,4A,4C,4E),UNIT=TAPE,UNITADD=((00,16)),LINK=(5306,5406,5308,5408),CUADD=8
IODEVICE ADDRESS=(2B80,16),UNIT=TAPE,CUNUMBER=(488),UNITADD=00,LIBID=C0001,LIBPORT=49
CNTLUNIT CUNUMBER=(489),PATH=(48,4A,4C,4E),UNIT=TAPE,UNITADD=((00,16)),LINK=(5306,5406,5308,5408),CUADD=9
IODEVICE ADDRESS=(2B90,16),UNIT=TAPE,CUNUMBER=(489),UNITADD=00,LIBID=C0001,LIBPORT=4A
CNTLUNIT CUNUMBER=(48A),PATH=(48,4A,4C,4E),UNIT=TAPE,UNITADD=((00,16)),LINK=(5306,5406,5308,5408),CUADD=A
IODEVICE ADDRESS=(2BA0,16),UNIT=TAPE,CUNUMBER=(48A),UNITADD=00,LIBID=C0001,LIBPORT=4B
CNTLUNIT CUNUMBER=(48B),PATH=(48,4A,4C,4E),UNIT=TAPE,UNITADD=((00,16)),LINK=(5306,5406,5308,5408),CUADD=B
IODEVICE ADDRESS=(2BB0,16),UNIT=TAPE,CUNUMBER=(48B),UNITADD=00,LIBID=C0001,LIBPORT=4C
CNTLUNIT CUNUMBER=(48C),PATH=(48,4A,4C,4E),UNIT=TAPE,UNITADD=((00,16)),LINK=(5306,5406,5308,5408),CUADD=C
IODEVICE ADDRESS=(2BC0,16),UNIT=TAPE,CUNUMBER=(48C),UNITADD=00,LIBID=C0001,LIBPORT=4D
CNTLUNIT CUNUMBER=(48D),PATH=(48,4A,4C,4E),UNIT=TAPE,UNITADD=((00,16)),LINK=(5306,5406,5308,5408),CUADD=D
IODEVICE ADDRESS=(2BD0,16),UNIT=TAPE,CUNUMBER=(48D),UNITADD=00,LIBID=C0001,LIBPORT=4E
CNTLUNIT CUNUMBER=(48E),PATH=(48,4A,4C,4E),UNIT=TAPE,UNITADD=((00,16)),LINK=(5306,5406,5308,5408),CUADD=E
IODEVICE ADDRESS=(2BE0,16),UNIT=TAPE,CUNUMBER=(48E),UNITADD=00,LIBID=C0001,LIBPORT=4F
CNTLUNIT CUNUMBER=(48F),PATH=(48,4A,4C,4E),UNIT=TAPE,UNITADD=((00,16)),LINK=(5306,5406,5308,5408),CUADD=F
IODEVICE ADDRESS=(2BF0,16),UNIT=TAPE,CUNUMBER=(48F),UNITADD=00,LIBID=C0001,LIBPORT=50
```

5.2.5 Display and control your settings

In a z/OS environment, you can use the DISPLAY SMS and DEVSERV QUERY LIBRARY commands to check some portions of your definitions.

DISPLAY SMS command in TS7740 Virtualization Engine

You can use the DISPLAY SMS command to display and check the settings in the DEVSUPxx member for the scratch categories, as shown in the following example:

```
DISPLAY SMS,LIBRARY(xxxxxxxx),DETAIL
```

Note: The scratch count of MEDIA2 does not necessarily match the number of scratch volumes of your tape management system when you use the Expire Hold function in the TS7700 Virtualization Engine. OAM displays the scratch count it receives from the TS7700 Virtualization Engine.

Example 5-4 shows the sample output of a DISPLAY SMS,LIBRARY command against the Composite Library; in this case, for a TS7700 Virtualization Engine, the OPM function is always supported by this type of library. The display shows that MEDIA2 uses category 0002 and MEDIA1 uses category 0001.

Example 5-4 Display SMS,LIB from a Composite Library

```

D SMS,LIB(COMPLIB),DETAIL
F OAM,D,LIB,COMPLIB,L=ST6T10-Z
CBR1110I OAM LIBRARY STATUS: 141
TAPE      LIB  DEVICE    TOT  ONL  AVL  TOTAL  EMPTY  SCRTCH  ON OP
LIBRARY   TYP  TYPE      DRV  DRV  DRV  SLOTS  SLOTS  VOLS
COMPLIB   VCL  3957-V06   768  768  287    0      0 368298  Y  Y
-----
MEDIA      SCRATCH      SCRATCH      SCRATCH
TYPE       COUNT      THRESHOLD      CATEGORY
MEDIA1     170345             0          0001
MEDIA2     197953             0          0002
-----
DISTRIBUTED LIBRARIES:  DISTLIB0      DISTLIB1      DISTLIB2
-----
LIBRARY ID: 10001
OPERATIONAL STATE:  AUTOMATED
ERROR CATEGORY SCRATCH COUNT:             33
CORRUPTED TOKEN VOLUME COUNT:             0
-----
LIBRARY SUPPORTS IMPORT/EXPORT.
LIBRARY SUPPORTS OUTBOARD POLICY MANAGEMENT.

```

Example 5-5 shows the sample output of a DISPLAY SMS,LIBRARY command against the Distributed Library.

Example 5-5 Display SMS,LIB output from a Distributed Library

```

D SMS,LIB(DISTLIB1),DETAIL
F OAM,D,LIB,DISTLIB1,L=ST6T10-Z
CBR1110I OAM LIBRARY STATUS: 062
TAPE      LIB  DEVICE    TOT  ONL  AVL  TOTAL  EMPTY  SCRTCH  ON OP
LIBRARY   TYP  TYPE      DRV  DRV  DRV  SLOTS  SLOTS  VOLS
DISTLIB1  VDL  3957-V06    0    0    0   1348    819    0  Y  Y
-----
COMPOSITE LIBRARY:      COMPLIB
-----
LIBRARY ID: 10011
OPERATIONAL STATE:  AUTOMATED
SCRATCH STACKED VOLUME COUNT:             222
PRIVATE STACKED VOLUME COUNT:             108
-----
LIBRARY SUPPORTS IMPORT/EXPORT.
LIBRARY SUPPORTS OUTBOARD POLICY MANAGEMENT.
CONVENIENCE I/O STATION INSTALLED.
CONVENIENCE I/O STATION IN OUTPUT MODE.
BULK INPUT/OUTPUT NOT CONFIGURED.

```

Four states can be reported (if the library is in the associated state) with the DISPLAY SMS command. These are:

- ▶ Limited Cache Free Space - Warning State (TS7720 Virtualization Engine only)
- ▶ Out of Cache Resources - Critical State (TS7720 Virtualization Engine only)
- ▶ Forced Pause Occurred
- ▶ Grid Links Degraded

DISPLAY SMS command for the TS7720 Virtualization Engine

For the TS7720 Virtualization Engine, the output differs slightly. Example 5-6 shows the output of D SMS,LIB(BARR60),DETAIL for Composite Library BARR60, which consists of two TS7720 Virtualization Engines.

Example 5-6 Display SMS,LIB from a TS7720 Virtualization Engine composite library

```

D SMS,LIB(BARR60),DETAIL
CBR1110I OAM LIBRARY STATUS: 672
TAPE      LIB  DEVICE    TOT  ONL  AVL  TOTAL  EMPTY  SCRTCH  ON OP
LIBRARY   TYP  TYPE      DRV  DRV  DRV  SLOTS  SLOTS  VOLS
BARR60    VCL  3957-VEA  512   0   0    0      0    495  Y  Y
-----
MEDIA      SCRATCH      SCRATCH      SCRATCH
TYPE       COUNT      THRESHOLD    CATEGORY
MEDIA2      495          0          3002
-----
DISTRIBUTED LIBRARIES:  BARR60A  BARR60B
-----
LIBRARY ID: BA060
OPERATIONAL STATE:  AUTOMATED
ERROR CATEGORY SCRATCH COUNT:          0
CORRUPTED TOKEN VOLUME COUNT:          0
-----
LIBRARY SUPPORTS OUTBOARD POLICY MANAGEMENT.

```

The only difference is in the Device Type field. In a Grid consisting of TS7720 Clusters only, it is 3957-VEA, where in a pure TS7740 Grid, a device type of 3957-V06 is displayed. In a multi-cluster hybrid grid, the device type is obtained from the library of the last drive that was initialized in the Grid and might be shown as 3957-VEA or as 3957-V06.

In a TS7720 Virtualization Engine Distributed Library, we see fields that are important. Example 5-7 shows the output of D SMS,LIB(BARR66A),DETAIL for a distributed library.

Example 5-7 The Display SMS,LIB from a TS7720 Virtualization Engine distributed library

```

D SMS,LIB(BARR60A),DETAIL
CBR1110I OAM LIBRARY STATUS: 088
TAPE      LIB  DEVICE    TOT  ONL  AVL  TOTAL  EMPTY  SCRTCH  ON OP
LIBRARY   TYP  TYPE      DRV  DRV  DRV  SLOTS  SLOTS  VOLS
BARR60A    VDL  3957-VEA    0   0   0    0      0    0  Y  Y
-----
COMPOSITE LIBRARY:      BARR60
-----
LIBRARY ID: BA60A
CACHE PERCENTAGE USED:  77
OPERATIONAL STATE:  AUTOMATED
-----
LIBRARY SUPPORTS OUTBOARD POLICY MANAGEMENT.
BULK INPUT/OUTPUT NOT CONFIGURED.

```

The example shows the following information:

- ▶ The device type is 3957-VEA.
- ▶ The amount of cache used is listed.
- ▶ The number of TOTAL SLOTS and EMPTY SLOTS is always 0.
- ▶ As there are no cartridges to manage, there is no I/O station.
- ▶ The SCRATCH and PRIVATE STACKED lines are suppressed.

DEVSERV QUERY LIBRARY command

When making changes to the libraries, the DEVSERV QUERY LIBRARY or DS QL command should always be used to query your library configuration before and after an activation of your IODF. You can query both a Composite and a Distributed library.

The new command DS QL,CATS allows us to change logical VOLSER categories without need to IPL the system.

Example 5-8 shows how you would list all categories used in a system.

Example 5-8 Sample output of DEVSERV QL,CATS

```
DS QL,CATS
IEE459I 09.56.42 DEVSERV QL,LIB 600
0001 0002 0003 0004 0005 0006 0007 0008 0009 000A 000E 000F
```

After you have the actual categories, you can change them. To perform that task, change the first three digits of the category, however, the last digit must remain unchanged, because it represents the media type.

Example 5-9 shows the command that changes all categories to 111 for the first three digits.

Example 5-9 Sample output of DEVSERV QL,CATS(111)*

```
DS QL,CATS(111*)
IEE459I 09.57.01 DEVSERV QL,LIB 603
1111 1112 1113 1114 1115 1116 1117 1118 1119 111A 111E 111F
```

Ensure that this change will be done in the DEVSUPxx PARMLIB member; otherwise, the next initial program load (IPL) will revert categories to what they were in DEVSUPxx.

Note: Be aware of the potential risks of changing scratch categories in a running system. If you alter the media2 scratch category and there are no logical volumes in this altered scratch category in the TS7700 Virtualization Engine, all scratch mounts will fail.

Example 5-10 shows how you would list all libraries using the DS QL command. Those LIBRARY-IDs (LIBIDs) marked with asterisks (*) are actually attached to the host.

Example 5-10 DEVSERV QL,LIST

```
DS QL,LIST
IEE459I 14.57.36 DEVSERV QL,LIB 708
The following libids are defined in the ACTIVE configuration:
*C0323 *BA094 *BA055 *CA002 *BA012 *BA091 BA049 *BA022 *BA044 *059C8
*BA095 *BA048 BA092 *BA010 *BA008 *BA060 *BA036 *11975 *B0009 BA069
CA022 C0159 11974 *C0076 BA009 *CA003 BA056 12087 BA066 BA035
BA071 21252 BA072 BA042 BA046 BA063 BA041 BA040 BA061 BA070
BA047 BA034 BA033 BA013 BA096 BA067
#NOTE: Asterisks indicate libraries that are actually attached to the host.
```

Example 5-11 shows a DEVSERV QLIB command output for Composite Library BA060. It shows the virtual devices and the Distributed Libraries belonging to this Composite Library.

Example 5-11 Sample output of DEVSERV QLIB for a Composite library

```

DS QL,BA060
IEE459I 15.04.45 DEVSERV QLIB 306
The following are defined in the    ACTIVE configuration:
LIBID  PORTID      DEVICES
BA060  02          901A* 901B* 901C* 901D* 901E* 901F* 9010* 9011*
          9012* 9013* 9014* 9015* 9016* 9017* 9018* 9019*
          04          903A* 903B* 903C* 903D* 903E* 903F* 9030* 9031*
          9032* 9033* 9034* 9035* 9036* 9037* 9038* 9039*
          44          913A* 913B* 913C* 913D* 913E* 913F* 9130* 9131*
          9132* 9133* 9134* 9135* 9136* 9137* 9138* 9139*
          06          905A* 905B* 905C* 905D* 905E* 905F* 9050* 9051*
          9052* 9053* 9054* 9055* 9056* 9057* 9058* 9059*
          08          907A* 907B* 907C* 907D* 907E* 907F* 9070* 9071*
          9072* 9073* 9074* 9075* 9076* 9077* 9078* 9079*
          42          911A* 911B* 911C* 911D* 911E* 911F* 9110* 9111*
          9112* 9113* 9114* 9115* 9116* 9117* 9118* 9119*
DISTRIBUTED LIBID(S)
  BA60A* BA60B*
#NOTE: Asterisks indicate devices that are currently attached.
```

Example 5-12 shows a detailed list of one single Library using the DS QL,<library-ID>, DETAIL command. Check that no duplicate port IDs are listed and that each port has 16 devices. This is the correct output for a TS7700 Virtualization Engine.

Example 5-12 Sample output of the DEVSERV QLIB,<library-ID>,ACTIVE command

```

DS QL,BA094,ACTIVE
IEE459I 15.02.46 DEVSERV QLIB 774
The following are defined in the    ACTIVE configuration:
LIBID  PORTID      DEVICES
BA094  04          293B* 293C* 293D* 293E* 293F* 2930* 2931* 2932*
          2933* 2934* 2935* 2936* 2937* 2938* 2939* 293A*
          02          291B* 291C* 291D* 291E* 291F* 2910* 2911* 2912*
          2913* 2914* 2915* 2916* 2917* 2918* 2919* 291A*
          06          295B* 295C* 295D* 295E* 295F* 2950* 2951* 2952*
          2953* 2954* 2955* 2956* 2957* 2958* 2959* 295A*
          08          297B* 297C* 297D* 297E* 297F* 2970* 2971* 2972*
          2973* 2974* 2975* 2976* 2977* 2978* 2979* 297A*
          03          2920* 2921* 2922* 2923* 2924* 2925* 2926* 2927*
          2928* 2929* 292A* 292B* 292C* 292D* 292E* 292F*
          01          2900* 2901* 2902* 2903* 2904* 2905* 2906* 2907*
          2908* 2909* 290A* 290B* 290C* 290D* 290E* 290F*
          05          2940* 2941* 2942* 2943* 2944* 2945* 2946* 2947*
          2948* 2949* 294A* 294B* 294C* 294D* 294E* 294F*
          07          2960* 2961* 2962* 2963* 2964* 2965* 2966* 2967*
          2968* 2969* 296A* 296B* 296C* 296D* 296E* 296F*
#NOTE: Asterisks indicate devices that are currently attached.
```

You can display the command syntax as follows:

```
DS QL,?
```

For a complete description of the QLIB command, see the following resources:

- ▶ Appendix H, “DEVSERV QLIB command” on page 923
- ▶ *z/OS MVS System Commands*, SA22-7627

5.2.6 Set values for the Missing Interrupt Handler

The TS7700 Virtualization Engine emulates 3490E devices and does not automatically communicate the Missing Interrupt Handler (MIH) timeout values to the host operating system in the Read Configuration Data Channel Control Word (CCW).

Important: An MIH value of 45 minutes is recommended for the virtual devices in a multi-cluster grid when a copy consistency for the remote clusters is set to RUN.

You must specify the MIH timeout value for IBM 3490E devices. The value applies only to the virtual 3490E drives and not to the real IBM TS1130/TS1120/3592 drives that the TS7740 Virtualization Engine manages in the back end. Remember that the host only knows about logical 3490E devices.

Table 5-7 summarizes the minimum values, which might need to be increased, depending on specific operational factors.

Table 5-7 Tape device MIH values

Tape device	MIH
3480, 3490 with less than eight devices per CU or low usage	3 minutes
3480, 3490 with eight devices per CU or heavy usage	5 minutes
3490E or 3480, 3490 with ESCON	10 minutes
3490E with ECST	20 minutes
TS7700 Virtualization Engine single-cluster grid with 3490E emulation drives	20 minutes
TS7700 Virtualization Engine multi-cluster grid with 3490E emulation drives	45 minutes

Specify the MIH values in PARMLIB member IECIOSxx. Alternatively, you can also set the MIH values through the System z operator command SETIOS. This setting is available until it is manually changed or until the system is initialized.

Use the following statements in PARMLIB, or manual commands to display and set your MIH values:

- ▶ You can specify the MIH value in the IECIOSxx PARMLIB member:
MIH DEV=(0A40-0A7F),TIME=45:00
- ▶ To manually specify MIH values for emulated 3490E tape drives, use:
SETIOS MIH,DEV=(0A40-0A7F),TIME=45:00
 - To display the new settings, use:
D IOS,MIH,DEV=0A40
 - To check the current MIH time, use:
D IOS,MIH,TIME=TAPE

The settings of the SETIOS and the MIH values in the IECIOSxx member change the value for the primary timeouts, but you cannot change the secondary timeouts. Those are delivered by the self-describing values from the device itself.

More information about MIH settings is in *MVS Initialization and Tuning Reference*, SA22-7592.

When specifying time intervals, consider the following possibilities:

- ▶ The MIH detects a missing interrupt condition within 1 second of the time interval that you specify.
- ▶ If the time interval is too short, a false missing interrupt can occur and cause early termination of the channel program. For example, if a 30-second interval is specified for a tape drive, a rewind might not complete before the MIH detects a missing interrupt.
- ▶ If the time interval is too long, a job or system could hang because the MIH has not yet detected a missing interrupt. For example, if a 15-minute time interval is specified for a tape drive used as an IMS™ log tape, the MIH could delay IMS for 15 minutes because of MIH detection.

During IPL (if the device is defined to be ONLINE) or during the VARY ONLINE process, some devices might present their own MIH timeout values through the *primary/secondary* MIH timing enhancement contained in the self-describing data for the device. The *primary* MIH timeout value is used for most I/O commands, but the *secondary* MIH timeout value may be used for special operations, such as long-busy conditions or long running I/O operations.

Any time a user specifically sets a device or device class to have an MIH timeout value that is different from the IBM-supplied default for the device class, that value will override the device-established primary MIH timeout value. This implies that if an MIH timeout value that is equal to the MIH default for the device class is explicitly requested, IOS will not override the device-established primary MIH timeout value. To override the device-established primary MIH timeout value, you must explicitly set a timeout value that is not equal to the MIH default for the device class.

Note: Overriding the device-supplied primary MIH timeout value might adversely affect MIH recovery processing for the device or device class.

See the specific device's reference manuals to determine if the device supports self-describing MIH timeout values.

5.3 TS7700 Virtualization Engine software definitions for z/OS

In this section, we describe the software definition considerations for implementing the TS7700 Virtualization Engine in z/OS, VM/ESA, and z/VSE environments. From a software point of view, the TS7700 Virtualization Engine is the same as an IBM 3494 Enterprise Tape Library with IBM 3490E Tape Drives.

The TS7700 Virtualization Engine must be defined as a new tape library with emulated 3490E Tape Drives from the host system. See *IBM TotalStorage 3494 Tape Library: A Practical Guide to Tape Drives and Tape Automation*, SG24-46322 for more information about defining this configuration.

The software levels required to support a TS7700 Virtualization Engine are explained in 3.5.2, “Software requirements” on page 140.

Tape management systems

From the host perspective, a TS7700 Virtualization Engine is a single automated tape subsystem whether it is a Single Cluster configuration or a multi-cluster grid configuration. The tape management system sees only the Composite Library and logical drives. There is no difference from the tape management system's point of view between a multi-cluster grid TS7700 Virtualization Engine installation (Peer-to-Peer) and a Single Cluster installation (stand-alone).

See "RMM and Copy Export" on page 299 for Copy Export functions.

See 7.7, "DFSMSrmm and other tape management systems" on page 395 for other new improvements in DFSMSrmm functions.

5.3.1 z/OS and DFSMS/MVS System-managed tape

To define the TS7700 Virtualization Engine to DFSMS, use the ISMF screens to create a new definition of the TS7700 Virtualization Engine logical tape library to be recognized from the host. This definition is done in the same way for a new installation of a single-cluster grid as it is done for a new installation of a multi-cluster grid TS7700 Virtualization Engine, with the exception of Distributed Library definitions:

- ▶ In a Single Cluster configuration, you define one Composite and one Distributed Library.
- ▶ In a Multi Cluster configuration, you define one Composite Library and two, three, or four Distributed Libraries.

To use the TS7700 Virtualization Engine, at least one Storage Group must be created to allow the TS7700 Virtualization Engine logical tape library virtual drives to be allocated by the ACS routines. Because all of the logical drives and volumes are associated with the Composite Library, only the Composite Library can be defined in the Storage Group. The distributed libraries must not be defined in the Storage Group.

See the following resources for information about host software implementation tasks for IBM tape libraries:

- ▶ *z/OS DFSMS OAM Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC35-0427
- ▶ *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789

Perform the following steps to define the TS7700 Virtualization Engine Tape Library in a z/OS environment:

1. Modify the SYS1.PARMLIB member (such as IEFSSNxx, IGDSMSxx, LOADxx, DEVSUPxx, or COMMANDxx).
2. Create the Tape Configuration Database (TCDB), which is an ICF catalog type VOLCAT.
3. Running the IDCAMS IMPORT CONNECT command to connect the TCDB to the other system is required when the tape library sharing capability is used.
4. Add the procedure to start the OAM address space.
5. Define the tape library as a DFSMS resource. Define the Composite Library and one or more distributed libraries. Remember that library names may not start with a "V".

Figure 5-10 shows the definition of a Composite Library.

Panel	Utilities	Scroll	Help

TAPE LIBRARY DEFINE		Page 1 of 2	
Command ==>_			
SCDS Name . : SCDS.TEMP.PRIMARY			
Library Name : IBMC1			
To Define Library, Specify:			
Description ==> TS7700 Grid Composite library			
==>			
Library ID CA010 (00001 to FFFFF)			
Console Name LIB1CON			
Entry Default Data Class DCATLDS			
Entry Default Use Attribute . . S (P=PRIVATE or S=SCRATCH)			
Eject Default K (P=PURGE or K=KEEP)			
Media Type: Scratch Threshold			
Media1 100 Media3 0 (0 to 999999)			
Media2 400 Media4 0 (0 to 999999)			
Use ENTER to Perform Verification; Use DOWN Command to View Next Panel;			
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.			

Figure 5-10 Composite Library definition

Figure 5-11 shows a sample screen that defines one of the Distributed Libraries.

Panel	Utilities	Scroll	Help

TAPE LIBRARY DEFINE		Page 1 of 2	
Command ==>_			
SCDS Name . : SCDS.TEMP.PRIMARY			
Library Name : IBMD1TU			
To Define Library, Specify:			
Description ==> TS7700 Distributed library A			
==>			
Library ID D1312 (00001 to FFFFF)			
Console Name			
Entry Default Data Class			
Entry Default Use Attribute . . (P=PRIVATE or S=SCRATCH)			
Eject Default (P=PURGE or K=KEEP)			
Media Type: Scratch Threshold			
Media1 0 Media3 0 (0 to 999999)			
Media2 0 Media4 0 (0 to 999999)			
Use ENTER to Perform Verification; Use DOWN Command to View Next Panel;			
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.			

Figure 5-11 Distributed Library definition

Note: Library ID is the only field that applies for the Distributed Libraries. All other fields can be blank or left as the default.

6. Create or update the Data Classes (DCs), Storage Classes (SCs), and Management Classes (MCs) for the TS7700 Virtualization Engine. Make sure that these defined construct names are the same as those you have defined at the TS7700 Virtualization Engine management interface (MI), especially in a Grid configuration, because outboard policy management is being used for multi-cluster grid copy control.
7. Create the Storage Groups (SGs) for the TS7700 Virtualization Engine. Make sure that these defined construct names are the same as those you have defined at the TS7700 Virtualization Engine MI.

The Composite Library must be defined in the Storage Group. Do *not* define the distributed libraries in the Storage Group.

Note: At OAM address space initialization, if a Distributed Library is defined to a Storage Group, the warning message CBR3017I is issued indicating that the Distributed Library is incorrectly defined to the Storage Group.

8. Create ACS routines to assign the constructs. Translate, test, and validate ACS routines.
9. Activate the new Source Control Data Set (SCDS).
10. SCDS activation will initiate an OAM restart if parameter RESTART=YES is specified in the OAM startup procedure in PROCLIB. If RESTART=NO is used, you must issue an OAM restart command manually using the F OAM,RESTART command. Example 5-13 shows a sample OAM procedure where RESTART=NO is used.

Example 5-13 OAM start-up procedure from PROCLIB

```
//OAM PROC
//IEFPROC EXEC PGM=CBROAM,REGION=0M,
// PARM=('OSMC=NO,APLAN=CBROAM,OAM=60,MAXS=2,RESTART=NO')
//SYSABEND DD SYSOUT=A
```

11. Vary the Composite Library and Distributed Libraries online.
12. Vary the TS7700 Virtualization Engine virtual drives online.

For more detailed information about defining a tape subsystem in a DFSMS environment, see *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789.

Note: The CBRXLCS FUNC=PTPDATA and FUNC=PTPMC programming interfaces are not supported in the TS7700 Virtualization Engine. The command is accepted but is treated as a no operation (NOP).

5.3.2 Implementing Copy Export

Copy Export function allows a copy of selected logical volumes written to the TS7700 to be removed and taken off site for disaster recovery purposes. The benefits of volume stacking, which places many logical volumes on a physical volume, are retained with this function. In addition, because the data being exported is a copy of the logical volume, the logical volume data remains accessible by the production host systems.

In this section, we describe how to implement and execute Copy Export. For more details and error messages that are related to the Copy Export function, see *IBM Virtualization Engine TS7700 Series Copy Export Function User's Guide*, which is available on the Techdocs website at the following address:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101092>

Copy Export overview

The basic steps for using Copy Export are as follows:

1. Setting up data management definitions.

Determine the data that is needed for off-site disaster recovery purposes. Through the automatic class selection routines, assign a management class that specifies a secondary volume pool to any data that is to be copy exported. As logical volumes are written to the TS7700 and then closed, they will be scheduled to be copied to a physical tape volume. With a management class that specifies a secondary copy pool, the data is copied both to a primary pool physical volume and a secondary pool physical volume.

2. Creating the Export List File Volume

When it is time to remove the physical volumes that have a copy of the data on them:

- a. First, write an Export List File volume to the TS7700 that is to perform the Copy Export operation. The export list file volume contains instructions for the Copy Export operation, one of the key instructions is the number of the physical volume pool to export. The export list file volume also holds a status file that contains records generated by the TS7700 as part of the Copy Export operation.
- b. Second, after the export list file volume has been written and closed, issue the Library Export z/OS console command.

Typically, you execute the Copy Export operation on a periodic basis. Because the purpose is to get a copy of the data off site for disaster recovery purposes, performing it soon after the data is created minimizes the time for the recovery point objective.

3. Executing the Copy Export operation

In processing a Copy Export operation, the TS7700 first copies any logical volumes that are assigned to the pool being exported that only have a copy of the logical volume in the cache. Next, each of the physical volumes in the specified pool that contains active data is mounted and a database backup of the TS7700 is written to them. Status records are written to the export list file volume and then the physical volume is ejected from the library. After all the physical volumes in the pool have been processed, the Copy Export operation is complete.

On the completion of the Copy Export operation, a summary status console message is generated. The status file on the export list file volume contains detailed information about the execution of the Copy Export operation and should be retained for future reference.

4. Removing the physical volumes.

An operator removes the physical volumes from the I/O station of the library and prepares them for shipment to an off-site location.

See the following sections for further information:

- If a Copy Export operation must be stopped, it can be canceled either through a host command or through the TS7700. See “Cancelling a Copy Export operation” on page 297.
- For information about how to use the copy exported volumes for Disaster Recovery testing of for recovery in case of a real disaster, see Chapter 10, “Failover and disaster recovery scenarios” on page 725.

- ▶ For requirements to set up and execute the Copy Export functions that are summarized in steps 1 - 3, see the following topics:
 - “Setting up data management definitions” on page 292
 - “Validating before activating the Copy Export function” on page 293
 - “Executing the Copy Export operation” on page 297,

However, we first explain general consideration that apply when you use the Copy Export function.

General considerations for Copy Export

Consider the following information when you use Copy Export:

- ▶ The library associated with the TS7700 executing the Copy Export operation must have an I/O station feature for the operation to be accepted. Empty the I/O station prior to executing Copy Export and prevent it from going to the full state.
- ▶ Only one Copy Export operation can be performed at a time.
- ▶ During the Copy Export operation, if the TS7700 cannot access the primary version of a logical volume and the secondary version is in a pool defined for Copy Export, regardless of whether that pool is involved in the current Copy Export operation, that secondary version is not accessible and the mount will fail. When a Copy Export operation is not being performed, if the primary version of a logical volume cannot be accessed and a secondary version exists, the secondary becomes the primary.
- ▶ Copy Export and insertion of logical volumes are mutually exclusive functions in a TS7700 or Grid.
- ▶ The export list file volume cannot be assigned to the secondary copy pool that is specified for the operation. If it is, the Copy Export operation will fail.
- ▶ A minimum of four physical tape drives must be available to the TS7700 for the Copy Export operation to be performed. The operation will be terminated by the TS7700 when fewer than four physical tape drives are available. Processing for the physical stacked volume in progress when the condition occurred will be completed and the status file records reflect what was completed before the operation was terminated.
- ▶ If a scratch physical volume is needed during a Copy Export operation, the secondary physical volume pool must have an available scratch volume or access to borrow one for the operation to continue. If a scratch volume is not available, the TS7700 indicates this through a console message and wait for up to 60 minutes. If a scratch volume is not made available to the secondary physical volume pool within 60 minutes, the Copy Export operation is terminated.
- ▶ Only one secondary physical volume pool can be specified per export operation and it must have been previously defined as a Copy Export pool.
- ▶ Specific logical volumes are not specified as part of a Copy Export operation, rather all valid logical volumes on the physical volumes in the specified secondary pool are considered for export. After the first time Copy Export is performed for a pool, the logical volumes that will be exported are only the ones for that pool that have been newly written or modified since the last export began. For recovery, all exported physical volumes that still contain active data from a source TS7700 need to be included as not all of the logical volumes created are going to be on the last set exported.
- ▶ During execution, if the TS7700 determines that a physical volume assigned to the specified secondary pool contains one or more primary logical volumes, that physical volume and any secondary logical volume on it are excluded from the Copy Export operation.

- ▶ Logical volumes are stacked on physical volumes managed by a TS7700. They are not directly readable by a physical tape drive. They can only be accessed through a TS7700.
- ▶ The primary copy of the logical volumes exported remains in the inventory of the TS7700 and Grid.
- ▶ A physical volume written in a pre-TS7700 format cannot be exported. The Physical Volume Status request of the BVIR function can be used to obtain information about the physical volumes in a pool. Included in that information is the format the volume was written in. If any of the volumes are still in the pre-TS7700 format, use the physical volume move function of the library manager to force a conversion of the data on the volume to the TS7700 format. Also see the description in section Copy Export Operation, pre-execution checking for more details.
- ▶ When a Copy Export operation is initiated, only those logical volumes assigned to the secondary pool specified in the Export List File Volume that are resident on a physical volume of the pool or in the cache of the TS7700 performing the export operation will be considered for export. For a Grid configuration, if a logical volume is to be copied to the TS7700 that will be performing the Copy Export operation, but that copy had not yet completed when the export is initiated, it will not be included in the current export operation.
- ▶ Logical volumes to be exported that are resident only in the cache and not mounted when the Copy Export operation is initiated will be copied to stacked volume in the secondary pool as part of the Copy Export operation.
- ▶ Any logical volume assigned to the specified secondary pool in the TS7700 after the Copy Export operation is initiated is not part of the export and will be written to a physical volume in the pool but will not be exported. This includes host and copy sourced data.
- ▶ To minimize the number of physical tapes used for Copy Export, use the highest capacity media and physical drive format that is compatible with the recovery TS7700. You may also want to reduce the number of concurrent tape devices that the TS7700 will use when copying data from cache to the secondary copy pool used for Copy Export.
- ▶ All copy-exported volumes that are exported from a source TS7700 must be placed in a library for recovery. The source TS7700 limits the number of physical volumes that can be copy-exported to approximately 2000 to ensure that they will all fit into the receiving library. At code release level R1.6 and later, the value can be changed to a maximum of 10000.
- ▶ The recovery TS7700 must have physical tape drives that are capable of reading the physical volumes from a source TS7700. If a source TS7700 writes the volumes using the native E05 format, then the recovery TS7700 must also have 3592 E05 drives running in native format mode. If the exporting pool on the source TS7700 is set up to encrypt the data, then the recovery TS7700 must also be set up to handle encrypted volumes and have access to the encryption key manager with replicated keys from the production site. If the source TS7700 writes the volumes in J1A or emulated J1A mode, then any 3592 model drive in the recovery TS7700 can read the data.
- ▶ The recovery TS7700 cannot contain any previous data and the recovery process cannot merge data from more than one source TS7700 together. As a part of Copy Export recovery, an option is provided to erase any previous data on the TS7700. This allows a TS7700 that is used for disaster recovery testing to be re-used for testing of different source TS7700's data.
- ▶ The physical and logical volumes exported during a Copy Export operation are not known by the tape management system. You can use the export status file to keep track of the volumes or use the information to manually update tape management system records.

Setting up data management definitions

To set up the data management definitions, perform the following steps:

1. Decide on Management Class construct name (or names).

As part of the plan for using the copy export function, you must decide on at least one Management Class construct name. A good practice is to make the name be meaningful and perhaps relate to the type of data to reside on the pool or the location where the data will be sent. For example, if the pool is to be used to send data to the primary disaster recovery site in Atlanta, a name like "MCPRIATL" could be used. "MC" for management class, "PRI" indicates its for the primary recovery site and "ATL" indicates Atlanta. Up to an eight-character name can be defined.

2. Define the Management Class names to DFSMS.

After the Management Class names are selected, the names must be defined to DFSMS and to the TS7700. For details about defining the Management Class in DFSMS, see *z/OS DFSMSdfp Storage Administration Reference*, SC26-7402.

None of the settings are actually used for system managed tape. All settings associated with a Management Class name are defined through the TS7700, not the DFSMS panels.

3. Define the Management Class names to the TS7700.

You should also define the Management Class names on the TS7700, because you are not using the Default Management Class settings for Copy Export volumes. Define a Secondary Pool for the copies to be exported.

See, "Creating Management Classes" on page 227 for details of how to add a Management Class.

4. Define the VOLSER ranges for the 3592 media.

You must define the VOLSER range (or ranges) for the physical volumes that are to be used for Copy Export, if you plan to use a specific VOLSER range. Make sure that you define the same pool you used in the Management Class definition as the Home Pool for this VOLSER range.

For the physical volumes that you use for Copy Export, defining a specific VOLSER range to be associated with a secondary pool on a source TS7700 may simplify the task of knowing the volumes to use in recovery and in returning a volume that no longer has active data on it to the TS7700 that manages it.

See 4.3.3, "Defining VOLSER ranges for physical (stacked) volumes" on page 203 for details about how to define the VOLSER ranges.

5. Define the characteristics of the physical volume pools used for Copy Export.

For the pool or pools, that you plan to use for Copy Export and that you have specified previously in the Management Class definition and, optionally, in the VOLSER range definition, you must select **Copy Export** in the Export Pool field.

See 4.3.4, "Defining physical volume pools" on page 206 for more details on how to change physical volume pool properties.

6. Code or modify the Management Class Automatic Class Selection (ACS) routine.

Add selection logic to the management class Automatic Class Selection (ACS) routine to assign the new Management Class name, or names, as appropriate.

7. Activate the new construct names and ACS routines.

Before new allocations will be assigned the new management class, the source control data set (SCDS) with the new management class definitions and ACS routines must be activated.

Validating before activating the Copy Export function

When the logical volumes are to be exported, there are some general validations that should be done. Before you initiate the operation, you should check that the TS7700 Virtualization Engine has the required physical drives and scratch physical volume resources. Verify that the TS7700 Virtualization Engine is not near the limit of the number of physical volumes that can have a status of Copy Exported (2000 physical volumes), and modify the value if required. Depending on your production environment, you might want to automate these validation steps.

The validation steps are as follows:

1. Determine whether data is in an older format. If you had migrated from a B10 or B20 VTS to the TS7700 Virtualization Engine using the outboard migration method, you might have data that is still in the older VTS format. The TS7700 Virtualization Engine cannot export data in the old format, so you must determine whether any of the data to export was written with the old format.
2. Validate that the TS7700 Virtualization Engine has at least four available physical tape drives. You can use the Library Request host console command that specifies the PDRIVE request. This returns the status of all physical drives attached to the TS7700 Virtualization Engine. If fewer than the required numbers of physical drives are available, you must call for service to repair drives prior to performing the Copy Export operation.

See Example 5-14 for the output of the PDRIVE request.

Example 5-14 Data returned by the PDRIVE request

```
LI REQ,BARR68A,PDRIVE
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR68A,PDRIVE.
CBR1280I LIBRARY BARR68A REQUEST. 896
KEYWORDS: PDRIVE
-----
PHYSICAL DRIVES V1
  SERIAL NUM   TYPE  MODE AVAIL  ROLE  POOL   PVOL   LVOL
000007878161 3592E05E E05E Y    RCLS  01   S70470 Z09866
000007878175 3592E05E E05E Y    MIGR  01   JA8149 Z04381
000001365176 3592E05E E05E Y    RCLT  01   S70421 Z09866
000001365177 3592E05E E05E Y    MIGR  01   JA8145 Z08629
000001365137 3592E05E E05E Y    RCLS  03   310112 XC4487
000001365181 3592E05E E05E Y    IDLE  01   310451
000007878194 3592E05E      Y    IDLE  00
000007878312 3592E05E E05E Y    RCLT  03   S70479 XC4487
```

In the response shown in Example 5-14, you can see the following items:

- Eight drives are defined. Their serial numbers (SN) are shown in the left column.
- All TS1120 drives are encryption-capable, as indicated by TYPE 3592E05E. TS1120 Tape Drives that are not encryption-capable will show as 3592E05. TS1130 Tape Drives would be listed as 3592E06E, because all TS1130 Tape Drives are encryption-capable.
- All eight drives are available (AVAIL=Y).
- MODE indicates the format in which the drive is currently operating. The information is only available when a tape is mounted on the drive.

- ROLE describes what the drive is doing at the moment, for example, RCLS is reclaim source and RCLT is reclaim target. In our example, we can see that two reclamation and two premigration operations are running:
 - Logical volume Z09866 is being reclaimed from physical volume S70470 mounted on a drive with SN-7878161 to physical volume S70421 mounted on a drive with SN-1365176. Both stacked volumes reside in Pool 01.
 - Logical volume XC4487 is being reclaimed from physical volume 310112 mounted on drive SN-1365137 to physical volume S70479 mounted on drive SN-7878312. Both stacked volumes reside in Pool 03.
 - Logical volume JA8149 is being written to physical volume Z04381 mounted on drive SN-7878175. Logical volume JA8145 is being written to physical volume Z08629 mounted on drive SN-1365177.
- Six drives are in use and two are IDLE, meaning ready for use. Serial number 1365181 is IDLE, but a physical volume is still mounted.

Note: The Copy Export operation requires a single drive to write the TS7700 Virtualization Engine database to the volumes being exported. Be sure to consider this situation when analyzing workload and drive utilization. See Chapter 9, “Performance and monitoring” on page 589 for more information about workload and drive utilization.

3. Check that the pool to be exported has sufficient scratch physical volumes and that the TS7700 Virtualization Engine is under the 2000 volume limit for copy exported volumes in all pools. You can use the Library Request host console command that specifies the POOLCNT request. See Example 5-15 for the response to the LI REQ, <library-ID>, POOLCNT command.

Example 5-15 Data returned from POOLCNT command

```
LI REQ,BARR68A,POOLCNT
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR68A,POOLCNT.
CBR1280I LIBRARY BARR68A REQUEST. 919
KEYWORDS: POOLCNT
```

PHYSICAL MEDIA COUNTS V1

POOL	MEDIA	EMPTY	FILLING	FULL	ERASE	ROR	UNAVAIL	CXPT
0	JA	164						
0	JJ	38						
1	JA	2	6	12	0	0	1	0
9	JJ	0	4	22	0	0	0	45

Pool 00 is the Common Scratch Pool. Pool 9 is the one used for Copy Export.

Example 5-15 shows the command POOLCNT; the response that is listed per pool is as follows:

- The media type used for each pool
- The number of empty physical volumes
- The number of physical volumes in the filling state
- The number of full volumes
- The number of physical volumes that have been reclaimed, but need to be erased
- The number of physical volumes in read-only recovery state
- The number of volumes unavailable or in a destroyed state (1 in Pool 1)
- The number of physical volumes in the copy exported state (45 in Pool 9)

You should determine when you usually want to start the Copy Export operation.

Thresholds could be the number of physical scratch volumes or other values that you

define. These thresholds could even be automated by creating a program that interprets the output from the Library Request commands PDRIVE and POOLCNT, and acts based on the required numbers.

For more information about the Library Request Command, see 8.3.3, “Host Console Request function” on page 544.

4. Create an Export List volume that provides the TS7700 Virtualization Engine with information about which data to export, and options to use during the operation.

If you use a multi-cluster grid, be sure to create the Export List volume only on the same TS7700 Virtualization Engine that is used for Copy Export, but not on the same physical volume pool as used for Copy Export. If more than one TS7700 Virtualization Engine in a multi-cluster grid configuration contains the Export List volume, the Copy Export operation will fail. Figure 5-12 shows the setting of a Management Class on the MI for the Export List volume in a two-cluster grid configuration. RN means one copy locally at RUN (R) and no copy (NN) on the other cluster.

Clusters	Copy Action
"Pesto[0]" (#BA68A)	Rewind Unload (RUN)
"Squint[1]" (#BA68B)	No Copy

Figure 5-12 Management Class settings for the Export List volume

5. Create an Export List volume, as shown in Example 5-16.

Example 5-16 Sample JCL to create an Export List volume

```
//*****
//* FILE 1: EXPORT LIST
//*****
//STEP1 EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSIN DD DUMMY
//SYSUT2 DD DSN=HILEVELQ.EXPLIST,
// UNIT=VTS1,DISP=(NEW,KEEP),LABEL=(1,SL),
// VOL=(,RETAIN),
// DCB=(RECFM=FB,BLKSIZE=80,LRECL=80,TRTCH=NOCOMP)
//SYSUT1 DD *
EXPORT LIST 03
EXPORT PARAMETERS PHYSICAL POOL TO EXPORT:09
OPTIONS1,COPY,EJECT
/*
//*****
//* FILE 2: RESERVED FILE
//*****
//STEP2 EXEC PGM=IEBGENER,COND=(4,LT)
```

```

//SYSPRINT DD SYSOUT=*
//SYSIN DD DUMMY
//SYSUT2 DD DSN=HILEVELQ.RESERVED,MGMTCLAS=MCNOCOPY,
// UNIT=VTS1,DISP=(NEW,KEEP),LABEL=(2,SL),
// VOL=(,RETAIN,REF=*.STEP1.SYSUT2),
// DCB=*.STEP1.SYSUT2
//SYSUT1 DD *
RESERVED FILE
/*
/*****
/* FILE 3: EXPORT STATUS FILE
/*****
//STEP3 EXEC PGM=IEBGENER,COND=(4,LT)
//SYSPRINT DD SYSOUT=*
//SYSIN DD DUMMY
//SYSUT2 DD DSN=HILEVELQ.EXPSTATS,
// UNIT=VTS1,DISP=(NEW,CATLG),LABEL=(3,SL),
// VOL=(,REF=*.STEP1.SYSUT2),
// DCB=*.STEP1.SYSUT2
//SYSUT1 DD *
EXPORT STATUS 01
/*

```

The information required in the Export List file is, as for BVIR, provided by writing a logical volume that fulfills the following requirements:

- ▶ That logical volume must have a standard label and contain three files:
 - An Export List file, as created in STEP1 in Example 5-16 on page 295. Notice that we want to export Pool 09. Option EJECT in record 2 tells the TS7700 Virtualization Engine to eject the stacked volumes upon completion. With OPTIONS1,COPY, the physical volumes will be placed in the export-hold category for later handling by an operator.
 - A Reserved file, as created in STEP2 in Example 5-16 on page 295. This file is reserved for future use.
 - An Export Status file, as created in STEP3 in Example 5-16 on page 295. In this file, the information is stored from the Copy Export operation. It is essential that you keep this file because it contains information related to the result of the Export process.
- ▶ All records must be 80 bytes in length.
- ▶ The Export List file must be written without compression. Therefore, you must assign a Data Class that specifies COMPRESSION=NO or you can overwrite the Data Class specification by coding TRTCH=NOCOMP in the JCL.
- ▶ Make sure that the files are assigned a Management Class that specifies that only the local TS7700 Virtualization Engine has a copy of the logical volume. You can either have the ACS routines assign this Management Class, or you can specify it in the JCL. We recommend that these files have the same expiration dates as the longest of the logical volumes you export because they must be kept for reference.

Executing the Copy Export operation

We can now initiate the Copy Export operation based on the logical volume, created by using the JCL in Example 5-16 on page 295. The process is as follows:

1. The Copy Export operation is initiated by issuing the LIBRARY EXPORT command. In this command, *logical VOLSER* is a variable and is the logical volume used in creating the Export List file volume. The command syntax is shown in Example 5-17.

Example 5-17 library export command

```
LIBRARY EXPORT,logical VOLSER
```

2. The host sends a command to the Composite Library and from there it is routed to the TS7700 Virtualization Engine where the Export List VOLSER resides.
3. The executing TS7700 Virtualization Engine validates the request, checking for required resources, and if all is acceptable, the Copy Export continues.
4. Logical volumes related to Pool 09 that still reside only in cache can delay the process. They will be copied to physical volumes in pool 9 as part of Copy Export execution.
5. Messages about the progress are sent to the system console. All messages are in the format shown in Example 5-18.

Example 5-18 Library message format

```
CBR3750I Message from library library-name: message text.
```

See the white paper *IBM Virtualization Engine TS7700 Series Copy Export Function User's Guide*, which is available at the Techdocs Library website:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101092>

After a successful completion, all physical tapes related to Pool 09 (in our example) are ejected. The operator can empty the I/O station and transport the tapes to another location.

Note: We describe how to perform Copy Export Recovery for disaster recovery or disaster recovery testing in Chapter 10, “Failover and disaster recovery scenarios” on page 725.

Cancelling a Copy Export operation

Examine the export status file records to see what has been processed prior to the cancellation request. Any physical volumes that completed the export process must be processed as though the export operation had completed.

Many reasons exist for cancelling a Copy Export operation:

- ▶ After initiating a Copy Export operation, you might realize that the pool being processed for export is incorrect.
- ▶ Other, more critical workloads must be run on the TS7700 Virtualization Engine and the extra impact of running the export operation is not desirable.
- ▶ A problem is encountered with the export that cannot be quickly resolved, for example, there are no physical scratch volumes available to add to the library.
- ▶ A problem is encountered with the library that requires it to be taken offline for service.

An export operation can be canceled from a host or through the TS7700 Virtualization Engine management interface.

A request to cancel an export operation can be initiated from any host attached to the TS7700 Virtualization Engine subsystem by using one of the following methods:

- ▶ Use the host console command `LIBRARY EXPORT,XXXXXX,CANCEL`, where `XXXXXX` is the volume serial number of the Export List File Volume.
- ▶ Use the Program Interface of the Library Control System (LCS) external services `CBRXLCS`.

If an export operation must be canceled and there is no host attached to the TS7700 Virtualization Engine that can issue the `CANCEL` command, you can cancel the operation through the TS7700 Virtualization Engine management interface. After confirming the selection, a cancel request is sent to the TS7700 Virtualization Engine processing the Copy Export operation.

Regardless of whether the cancellation originates from a host or the management interface, the TS7700 Virtualization Engine can process it as follows:

- ▶ If the processing of a physical volume has reached the point where it has been mounted to receive a database backup, the backup completes and the volume placed in the export-hold or eject category before the cancel processing can continue. Status file records are written for all logical and physical volumes that completed export processing.
- ▶ All physical resources (drives, stacked volumes, and exported stacked volumes) are made available for normal TS7700 Virtualization Engine subsystem processing.
- ▶ A completion message is sent to all hosts attached to the TS7700 Virtualization Engine indicating that the export was canceled by a host request. The message contains information about how much export processing completed prior to the execution of the cancellation request.

Host completion message

At the completion of the Copy Export operation, a completion message is broadcast to all hosts attached to the TS7700 Virtualization Engine. For z/OS, console messages are generated that provide information about the overall execution status of the operation.

Messages differ, depending on what the TS7700 Virtualization Engine encountered during the execution of the operation:

- ▶ If no errors or exceptions were encountered during the operation, message `CBR3855I` is generated. The message has the format shown in Example 5-19.

Example 5-19 CBR3855I message format

```
CBR3855I Export operation for logical list volume 'volser' in library 'library-name'
completed successfully. Requested: 'requested-number' Exportable: 'exportable-number'
Exported: 'exported-number' Stacked volumes: 'stacked-number' MBytes Exported:
'MBytes-exported' MBytes Moved: 'MBytes-moved'
```

- ▶ If error or exceptions were encountered during the operation, message `CBR3856I` is generated. The message has the format shown in Example 5-20.

Example 5-20 CBR3856I message format

```
CBR3856I Export operation for logical list volume 'volser' in library 'library-name'
completed with exceptions or errors. Requested: 'requested-number' Exportable:
'exportable-number' Exported: 'exported-number' Stacked volumes: 'stacked-number' MBytes
Exported: 'MBytes-exported' MBytes Moved: 'MBytes-moved'
```

If message `CBR3856I` is generated, you must examine the Export Status file to determine what errors or exceptions were encountered.

Either of the completion messages provides statistics on what was processed during the operation. The statistics reported are as follows:

- ▶ Requested-number: This is the number of logical volumes associated with the secondary volume pool specified in the export list file. Logical volumes associated with the specified secondary volume pool that were previously exported are not considered part of this count.
- ▶ Exportable-number: This is the number of logical volumes that are considered exportable. A logical volume is exportable if it is associated with the secondary volume pool specified in the export list file and it has a valid copy resident on the TS7700 Virtualization Engine performing the export. Logical volumes associated with the specified secondary volume pool that were previously exported are not considered to be resident in the TS7700 Virtualization Engine.
- ▶ Exported-number: This is the number of logical volumes that were successfully exported.
- ▶ Stacked-number: This is the number of physical volumes that were successfully exported.
- ▶ MBytes Exported: This is the number of MB contained in the logical volumes that were successfully exported. If the data on the logical volumes is compressed, the number includes the effect of compression.

Note: The number of megabytes (MB) exported is the sum of the MB integer values of the data stored on each Exported Stacked Volume. The MB integer value for each Exported Stacked Volume is the full count by bytes divided by 1,048,576 bytes. If the result is less than 1, the MB integer becomes 1, and if greater than 1 MB, the result is truncated to the integer value (rounded down).

- ▶ MBytes Moved: For Copy Export at code release level R1.3, this field reports the same number as the MBytes Exported field. For Copy Export at code release level R1.4 and later, this value is 0.

For more details and error messages related to the Copy Export function, see the white paper *IBM Virtualization Engine TS7700 Series Copy Export Function User's Guide*, which is available at the Techdocs Library website. Search for TS7700 at the following address:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs>

RMM and Copy Export

When a stacked volume associated with a Copy Export operation is ejected from a library (placed in export-hold or is physically ejected from the library), status message E0006 is issued by the library. RMM intercepts this message and:

- ▶ If the stacked volume is predefined to RMM, RMM will mark the volume as “ejected” or “in-transit” and set the movement/store date associated with the stacked volume.
- ▶ If the stacked volume is not predefined to RMM and the STACKEDVOLUME(YES) option in RMM is specified, RMM will automatically add the stacked volume to its Control Data Set (CDS).

To have DFSMSrmm policy management manage the retention and movement for volumes created by Copy Export processing, you must define one or more volume Vital Record Specifications (VRS). For example, assuming all copy exports are targeted to a range of volumes STE000-STE999, you could define a VRS as shown in Example 5-21.

Example 5-21 VRS definition

```
RMM AS VOLUME(STE*) COUNT(99999) LOCATION(location)
```

As a result of this, all matching stacked volumes that are set in AUTOMOVE will have their destination set to the required location and your existing movement procedures can be used to move and track them.

In addition to the support we listed, a copy exported stacked volume can become eligible for reclamation based on the reclaim policies defined for its secondary physical volume pool or through the Host Console Request function (LIBRARY REQUEST). When it becomes eligible for reclamation, the exported stacked volume no longer contains active data and can be returned from its off-site location for reuse.

For users that use DFSMSrmm, when you have stacked volume support enabled, DFSMSrmm automatically handles and tracks the stacked volumes created by Copy Export; however, there is no way to track which logical volume copies are on the stacked volume. You should retain the updated export list file created by you and updated by the library so that you have a record of what logical volumes were exported and on what exported stacked volume they reside.

For more details and error messages related to the Copy Export function in RMM, see the *z/OS V1 DFSMSrmm Implementation and Customization Guide*, SC26-7405.

Manage copy exported physical volumes

The default limit is 2000 copy exported physical volumes per TS7700. Prior to performing the operation, a good plan is to determine how many volumes have already been exported. You can use the Library Request function to view the number of copy exported volumes by pool in a specific library. With TS7700 R.16 and later, the limit can be changed through the MI. The value can be adjusted to a maximum of 10000.

You must set up procedures for when exported physical volumes can return to the library. Normally, the physical volumes must return when all logical volumes are expired. To determine when exported volumes have expired, the simplest way is to use the Bulk Volume Information Retrieval (BVIR) facility. Use BVIR with *Physical Volume Status Pool xx* as input. With BVIR, you can retrieve information for all physical volumes in a pool. See Chapter 9, “Performance and monitoring” on page 589 for more information about how to use BVIR. The copy exported physical volumes continue to be managed by the source TS7700 in terms of space reclamation. You can allow the TS7700 to determine when to reclaim a Copy Export physical volume or you can manually cause a volume to be reclaimed.

Another reason for bringing back a physical volume applies to a single-cluster grid, where the original physical volume is destroyed, meaning the exported copy would be the only available copy available.

When you insert the cartridges through the I/O station, these volumes are directly added to the insert category or the private category (if they contain active logical volumes) in the TS7740 Virtualization Engine. The TS7740 Virtualization Engine recognizes the volumes it owns and return them to the pool in which they were previously. The state of the volumes is changed to read-write and, if they are empty and without active logical volumes, they will have an empty status also. When insertion of the volumes is finished, the entire Copy Export operation for a single physical volume life cycle is complete. For more information about managing unassigned volumes, see “Unassigned volumes in the TS3500 Tape Library” on page 195.

Tape Encryption is supported for stacked volumes that are used for Copy Export, as described in 4.2.3, “Defining the Encryption Method for the new Logical Library” on page 188. The TS1130 and TS1120 drives must be encryption-enabled, the tape library must be at a Licensed Internal Code level that supports tape encryption, and the Encryption Key Manager

must be available at the disaster recovery site to be able to process encrypted stacked volumes that have been exported.

Grid considerations for Copy Export

Copy export is supported in a Grid configuration, but several considerations exist regarding that support.

Consideration 1

The first consideration has to do with how Copy Export is performed. In a Grid configuration, a Copy Export operation is performed against an individual TS7700, not across all TS7700 Virtualization Engines. In setting up to use Copy Export in a Grid plan for the following information:

- ▶ Decide on which TS7700(s) in a Grid configuration is going to be used to export a specific set of data. Although you can set up more than one TS7700 to export data, only the data from a single source TS7700 can be used in the recovery process. You cannot merge copy exported volumes from more than one source TS7700 in the recovery TS7700.
- ▶ For each specific set of data to export, define a management class name. On the TS7700 that will be used to export that data, define a secondary physical volume pool for that management class name and also make sure that you indicate that it is an export pool. Although you will need to define the management class name on all TS7700s in the Grid configuration, specify only the secondary physical volume pool on the TS7700 that is to perform the export operation. Specifying it on the other TS7700s in the Grid configuration does not interfere with the Copy Export operation, but it is a waste of physical volumes. The exception to this approach would be if you want one of the TS7700s in the Grid configuration to have a second physical copy of the data if the primary copies on other TS7700s are not accessible.
- ▶ As you are defining the management class name for the data, also make sure that the TS7700 to perform the export operation has a copy policy specifying it is to have a copy.
- ▶ When the Copy Export operation is executed, the export list file volume must only be valid on the TS7700 performing the operation. You will need to define a unique management class name to be used for the export list file volume. For that management class name, you will need to define its copy policy such that a copy is only on the TS7700 that is to perform the export operation. If the VOLSER specified for the export list file volume when the export operation is initiated is resident on more than one TS7700, the Copy Export operation will fail.
- ▶ On the TS7700 on which a Copy Export operation is initiated, do not set items of Copy Policy Override checked through the management interface panel. Activating Copy Policy Override, allows a copy of export list volume created on the TS7700 regardless of the management class setting and causes Copy Export failure. For instance, assume three Clusters (C0, C1, C2) exist in a grid configuration:
 - a. A Copy Export with the export list volume EXP000 is initiated from a host connected to the C0, and the Copy Export runs on the C2.
 - b. The copy mode of EXP000 must be [N,N,D] or [N,N,R] indicating that the only copy of EXP000 exists on C2.
 - c. If Copy Policy Override is activated on the C0 and the Copy Export is initiated from the host attached to C0, a copy of EXP000 is created both on the C0 and C1.
 - d. Grid detects that a copy of EXP000 exists on two clusters (C0 and C2) and does not start the Copy Export.
 - e. Copy export fails.

For example, assume that the TS7700 that is to perform the Copy Export operation is Cluster 1. The pool on that cluster to export is pool 8. You would need to set up a management class for the data that is to be exported such that it will have a copy on Cluster 1 and a secondary copy in pool 8. To ensure the data is on that cluster and is consistent with the close of the logical volume, you would want to have a copy policy of Rewind/Unload (RUN). You would define the following information:

- Define a management class, for example, of MCCEDATA, on Cluster 1 as follows:

Secondary Pool	8
Cluster 0 Copy Policy	RUN
Cluster 1 Copy Policy	RUN

- Define this same management class on Cluster 0 without specifying a secondary pool.
- To ensure that the export list file volume gets written to Cluster 1 and only exists there, define a management class, for example, of MCELFVOL, on Cluster 1 as follows:

Cluster 0 Copy Policy	No Copy
Cluster 1 Copy Policy	RUN

- Define this management class on Cluster 0 as follows:

Cluster 0 Copy Policy	No Copy
Cluster 1 Copy Policy	RUN

A Copy Export operation can be initiated through any virtual tape drive in the TS7700 Grid configuration; it does not have to be initiated on a virtual drive address in the TS7700 that is to perform the Copy Export operation. The operation will be internally routed to the TS7700 that has the valid copy of the specified export list file volume. Operational and completion status will be broadcast to all hosts attached to all of the TS7700s in the Grid configuration.

Only the logical volumes resident on the TS7700 performing the operation, at the time it is initiated, are exported. If a logical volume has not been copied yet or completed its copy operation when the export operation is initiated, it is not considered for export during the operation. It is assumed that Copy Export is performed on a regular basis and logical volumes, whose copies were not complete when a Copy Export was initiated, will be exported the next time Copy Export is initiated. You can check the copy status of the logical volumes on the TS7700 that is to perform the Copy Export operation prior to initiating the operation by using the Volume Status function of the BVIR facility. You can then be sure that all critical volumes will be exported during the operation.

Consideration 2

The next consideration has to do with how Copy Export recovery is performed. Copy export recovery is always to a stand-alone TS7700. As part of the recovery process, the recovery TS7700 processes all grid-related information in the database, converting it to look like a single TS7700. This conversion means that the recovery TS7700 will have volume ownership of all volumes. It is possible that one or more logical volumes become inaccessible because they were modified on a TS7700 other than the one that performed the Copy Export operation, and the copy did not complete prior to the start of the operation. The last has to do with returning copy exported physical volumes when they are empty. Remember that each copy exported physical volume remains under the management of the TS7700 from which it was exported.

Normally, you would return the empty physical volumes to the library I/O station that associated with the source TS7700 and re-insert them. They would then be reused by that

TS7700. If you want to move them to another TS7700, whether in the same Grid configuration or another, consider two important points:

- ▶ Ensure that the VOLSER ranges you had defined for that TS7700 matches the VOLSERS of the physical volume(s) you want to move.
- ▶ Use the host console request function, Library Request, to have the original TS7700 stop managing them:

```
LIBRARY REQUEST,libname,COPYEXP,volsr,DELETE
```

5.4 Software implementation in z/VM and z/VSE

This section explains how to implement and run the TS7700 Virtualization Engine under z/VM and z/VSE. It covers the basics for software requirements, implementation, customization, and platform-specific considerations about operations and monitoring. For more detailed information, see *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789.

5.4.1 General support information

Not all IBM Tape Libraries and TS7700 Virtualization Engine solutions are supported in all operating systems. Table 5-8 shows a summary of several supported tape solutions for non-z/OS environments. For information about support for TPF, see 5.5, “Software implementation in Transaction Processing Facility” on page 310.

Table 5-8 Supported tape solutions for non-z/OS platforms in System z environments

Platform / Tape System	IBM System Storage TS3500 Tape Library	TS7700 Virtualization Engine	3592 drives
z/VM 5.4 and 6.1 native	Yes	Yes ^a	Yes
z/VSE V4.1 native z/VSE V4.2 native z/VSE V4.3 native ^b	Yes	Yes	Yes
z/VSE V4.1 under z/VM z/VSE V4.2 under z/VM z/VSE V4.3 under z/VM ^b	Yes	Yes ^a	Yes
zTPF	Yes	Yes	Yes

a. With restrictions; see “Restrictions in all TS7700 Virtualization Engine environments” on page 303.

b. Includes support for Logical WORM

Even if z/VM and z/VSE can use the TS7700 Virtualization Engine, there are some restrictions you should consider.

Restrictions in all TS7700 Virtualization Engine environments

z/VSE is not able to provide SMS constructs to TS7700. Therefore, the functions provided by outboard policy management for the TS7700 Virtualization Engine cannot be used on this platform.

However, one possibility is to use dedicated physical pools in a TS7700 Virtualization Engine environment. After insert processing of virtual volumes completes, you can define a default

construct to the volume range, as described in 4.5, “Implementing Outboard Policy Management for non-z/OS hosts” on page 260.

TS7700 Virtualization Engine multi-cluster grid environments

z/VSE does not support TS7700 multi-cluster grid. z/VSE is unable to handle unsolicited messages and actions (for example, intervention required and hardware messages); therefore, a Peer-to-Peer VTS requires an active z/OS to support z/VSE. This requirement has been eliminated with the advent of the TS7700 Virtualization Engine and its management interface.

5.4.2 z/VM native support using DFSMS/VM

DFSMS/VM Function Level 221 (FL221) is the only way for a z/VM system to communicate with a TS7700 Virtualization Engine. DFSMS/VM FL221 is part of z/VM. The removable media services (RMS) function of DFSMS/VM FL221 provides TS7700 Virtualization Engine support in z/VM, as described in *DFSMS/VM Function Level 221 Removable Media Services*, SC24-6090.

Tape management

Although the RMS functions themselves do not include tape management system services, such as inventory management and label verification, RMS functions are designed to interface with a tape management system that can perform these functions. Additional information about third-party tape management systems that support the TS7700 Virtualization Engine in the and z/VM environment can be found in *IBM TotalStorage 3494 Tape Library: A Practical Guide to Tape Drives and Tape Automation*, SG24-4632.

Figure 5-13 shows the z/VM native support for the TS7700 Virtualization Engine.

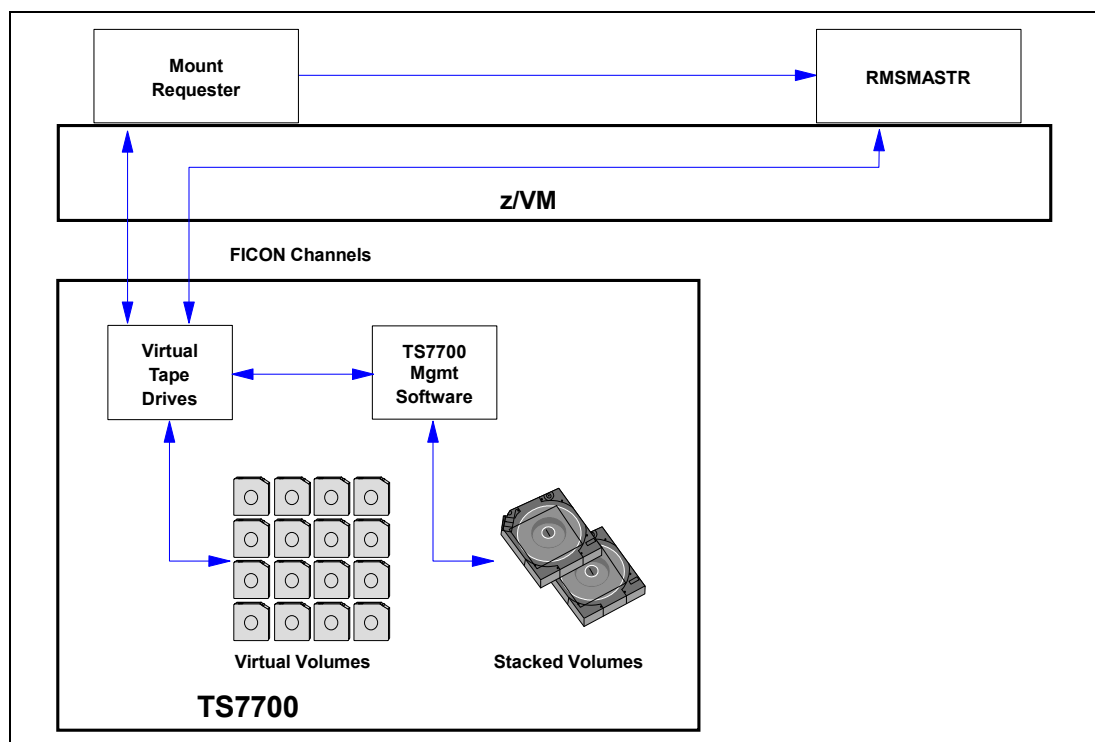


Figure 5-13 TS7700 Virtualization Engine in a native z/VM environment using DFSMS/VM

When you use the TS7700 Virtualization Engine in a VM environment, consider that many VM applications or system utilities use specific mounts for scratch volumes, so every time a mount request is issued from the host, the TS7700 Virtualization Engine has to recall the requested logical volume from the stacked cartridge if it is not already in the Tape Volume Cache (TVC). This can lead to performance degradation when writing data in a VM environment. In addition, VM backups usually require off-site movement, so the TS7700 Virtualization Engine is not the best candidate for this data.

DFSMS/VM

After you have defined the new TS7700 Virtualization Engine Tape Library through HCD, you must define the TS7700 Virtualization Engine to DFSMS/VM if the VM system is to use the TS7700 Virtualization Engine directly. You define the TS7700 Virtualization Engine Tape Library through the DFSMS/VM DGTVCNTL DATA control file. Also, you define the available tape drives through the RMCONFIG DATA configuration file. See *z/VM V6R1 DFSMS/VM Removable Media Services*, SC24-6185 for more information.

You have the RMS as a component of DFSMS/VM. To allow RMS to perform automatic insert bulk processing, you must create the RMBnnnnn DATA file in the VMSYS:DFSMS CONTROL directory, where nnnnn is the five-character tape library sequence number that is assigned to the TS7700 Virtualization Engine during hardware installation.

Note: The outboard policy management functions are currently not supported with z/VM.

For details about implementing DFSMS/VM and RMS, see *DFSMS/VM Function Level 221 Removable Media Services User's Guide and Reference*, SC35-0141. If the TS7700 Virtualization Engine is shared by your VM system and other systems, additional considerations apply. See *Guide to Sharing and Partitioning IBM Tape Library Data*, SG24-4409 for further information.

5.4.3 Native z/VSE

Native support is provided for the single-cluster grid TS7700 Virtualization Engine configuration in z/VSE Version 3.1.2, Version 4.1, and Version 4.2, and Version 4,3 which supports all IBM TS1130/TS1120/3592-J1A configurations (J70, C06, and TS7700 Virtualization Engine) without APARs in all automation offerings, including 3494 Enterprise Tape Library and TS3500 Tape Library configurations.

Note: The outboard policy management functions are currently not supported with z/VSE.

z/VSE supports the TS3500 Tape Library/3953 natively through its Tape Library Support (TLS). In addition to the old Tape Library support, a function has been added to allow the Tape Library to be supported through the S/390 channel command interface commands, thus eliminating any XPCC/APPC communication protocol required by the old interface. The external interface (LIBSERV JCL and LIBSERV macro) remains unchanged.

Define library support

First, define the type of support you are using by specifying the SYS ATL statement. You may define the following types:

- TLS** TLS Tape Library Support, which provides full VSE LPAR support
- VSE** LCDD, which does not support TS1130/TS1120/3592 (only IBM 3490E and 3590) and therefore not the TS3500 Tape Library
- VM** VM Guest Support, which when running z/VSE under z/VM and a TS7700 is used by both operating systems, where VGS (VSE Guest server) and DFSMS is needed (see 5.4.5, “z/VSE as a z/VM guest using a VSE Guest Server (VGS)” on page 308)

For native support under VSE, where TS7700 is used only by z/VSE select TLS, where at last one tape drive must be permanently assigned to VSE.

Define tape libraries

Next, define your tape library or libraries. This is done through a batch job, as shown in Example 5-22. Use skeleton member TLSDEF from ICCF Lib 59.

Example 5-22 Define tape libraries

```
* $$ JOB JNM=TLSDEF,CLASS=0,DISP=D
* $$ LST CLASS=A
// JOB TLSDEF
// EXEC LIBR,PARM='MSHP'
ACCESS S=IJSYSRS.SYSLIB
CATALOG TLSDEF.PROC REPLACE=YES
  LIBRARY_ID TAPELIB1 SCRDEF=SCRATCH00 INSERT=SCRATCH00          --- default library
  LIBRARY_ID TAPELIB2                                * SECOND LIB DEF
  DEVICE_LIST TAPELIB1 460:463                        * DRIVES 460 TO 463
  DEVICE_LIST TAPELIB2 580:582                        * DRIVES 580 TO 582
  QUERY_INV_LISTS LIB=TLSINV                          * MASTER INVENTORY FILES
  MANAGE_INV_LISTS LIB=TLSMAN                        * MANAGE FROM MASTER
/+
```

LIBSERV

The communication from the host to the TS7700 Virtualization Engine goes through the LIBSERV JCL or macro interface. Example 5-23 shows a sample job using LIBSERV to mount volume 123456 for write on device address 480 and, in a second step, to release the drive again.

Example 5-23 Sample LIBSERV JCL

```
$$ JOB JNM=BACKUP,CLASS=0,DISP=D
$$ JOB BACKUP
// ASSGN SYS005,480
// LIBSERV MOUNT,UNIT=480,VOL=123456/W
// EXEC LIBR
BACKUP S=IJSYSRS.SYSLIB TAPE=480
/*
// LIBSERV RELEASE,UNIT=480
/&
$$ EOJ
```

LIBSERV provides the following functions:

Query all libraries for a volume	LIBSERV AQUERY,VOL=123456
Mount from category	LIBSERV CMOUNT,UNIT=480,SRCCAT=SCRATCH01
Mount a specific volume	LIBSERV MOUNT,UNIT=480,VOL=123456
Demount a volume	LIBSERV RELEASE,UNIT=480
Query count of volumes	LIBSERV CQUERY,LIB=TAPELIB1,SRCCAT= SCRATCH01
Query device	LIBSERV DQUERY,UNIT=480
Query inventory of library	LIBSERV IQUERY,LIB=TAPELIB1,SRCCAT=SCRATCH01
Query library	LIBSERV LQUERY,LIB=TAPELIB1
Manage inventory	LIBSERV MINVENT,MEMNAME=ALL,TGTCAT=SCRATCH01
Change category	LIBSERV SETVCAT,VOL=123456,TGTCAT=SCRATCH01
Query library for a volume	LIBSERV SQUERY,VOL=123456,LIB=TAPELIB1

For additional information, see *z/VSE System Administration Guide*, SC33-8224 and *z/VSE System Macros Reference*, SC33-8230.

5.4.4 VM/ESA and z/VM guest support

In this section, we describe two host environments that allow you to use an IBM TS7700 Virtualization Engine while running it as a guest host system under z/VM.

Note: When z/OS is installed as or z/VM guest on a virtual machine, you must specify the following statement in the virtual machine directory entry for the VM user ID under which the z/OS guest operating system is IPLed:

```
STDEVOPT LIBRARY CTL
```

z/OS guests

The environments described in 5.3.1, “z/OS and DFSMS/MVS System-managed tape” on page 286 can operate when z/OS is running as a guest of z/VM Release 5.4 or later.

In this environment, additional software products are not required.

The STDEVOPT statement specifies the optional storage device management functions available to a virtual machine. The LIBRARY operand with CTL tells the control program that the virtual machine is authorized to issue Tape Library commands to an IBM Automated Tape Library Dataserver. If the CTL parameter is not explicitly coded, the default of NOCTL is used. NOCTL specifies that the virtual machine is not authorized to issue commands to a Tape Library, and this results in an I/O error (command reject) when MVS tries to issue a command to the library. For further information about the STDEVOPT statement, go to the following address:

<http://www.vm.ibm.com/zvm540/>

z/VSE guests

Some VSE tape management systems require VSE Guest Server (VGS) support and also DFSMS/VM RMS for communication with the TS7700 Virtualization Engine.

If the VGS is required, define the LIBCONFIG file and FSMRMVGC EXEC configuration file on the VGS service machine's A disk. This file simply cross-references the z/VSE guest's tape library names with the names that DFSMS/VM uses. To enable z/VSE guest exploitation of inventory support functions through the LIBSERV-VGS interface, the LIBRCMS part must be installed on the VM system.

If VGS is to service inventory requests for multiple z/VSE guests, you must edit the LIBRCMS SRV NAMES cross-reference file. This file enables the inventory support server to access Librarian files on the correct VSE guest machine. For further information, see the following resources:

- ▶ 5.4.5, “z/VSE as a z/VM guest using a VSE Guest Server (VGS)” on page 308
- ▶ 7.6, “VSE Guest Server Considerations” in *Guide to Sharing and Partitioning IBM Tape Library Data*, SG24-4409.

CA DYNAM/TM-VSE does not use the VGS machine.

5.4.5 z/VSE as a z/VM guest using a VSE Guest Server (VGS)

When a z/VSE guest machine uses a tape drive in the TS7700 Virtualization Engine, the virtual tape drive must be attached to that machine and the virtual tape volume must be mounted on the drive. Because, as a virtual machine, z/VSE cannot communicate with the TS7700 Virtualization Engine to request a tape mount; RMSMASTR (a z/VM machine) must attach the tape drive and mount the volume. However, z/VSE cannot use RMSMASTR directly, because RMS functions run only in CMS mode.

Therefore, some z/VSE guest scenarios use the CMS service machine, called the VGS, to communicate with RMSMASTR. VGS uses the standard facilities of RMS to interact with TS7700 Virtualization Engine and the virtual drives.

Figure 5-14 shows the flow and connections of a TS7700 Virtualization Engine in a z/VSE environment under VM.

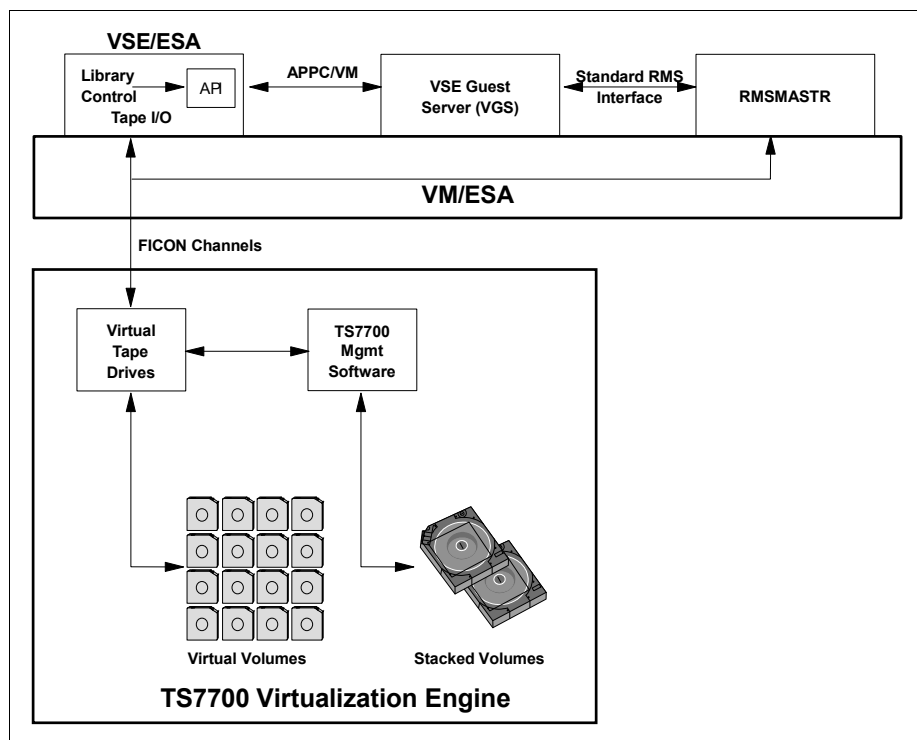


Figure 5-14 TS7700 Virtualization Engine in a z/VSE environment as a VM guest

Tape management systems

As with the VM/ESA® native environment (see 5.4.2, “z/VM native support using DFSMS/VM” on page 304), the tape management system is responsible for keeping an inventory of volumes in the TS7700 Virtualization Engine that belong to z/VSE. Some vendor tape management support scenarios do not use VGS. Instead, they communicate directly with RMSMASTR through CSL calls. Figure 5-15 on page 310 shows the case of CA-DYNAM/T VSE.

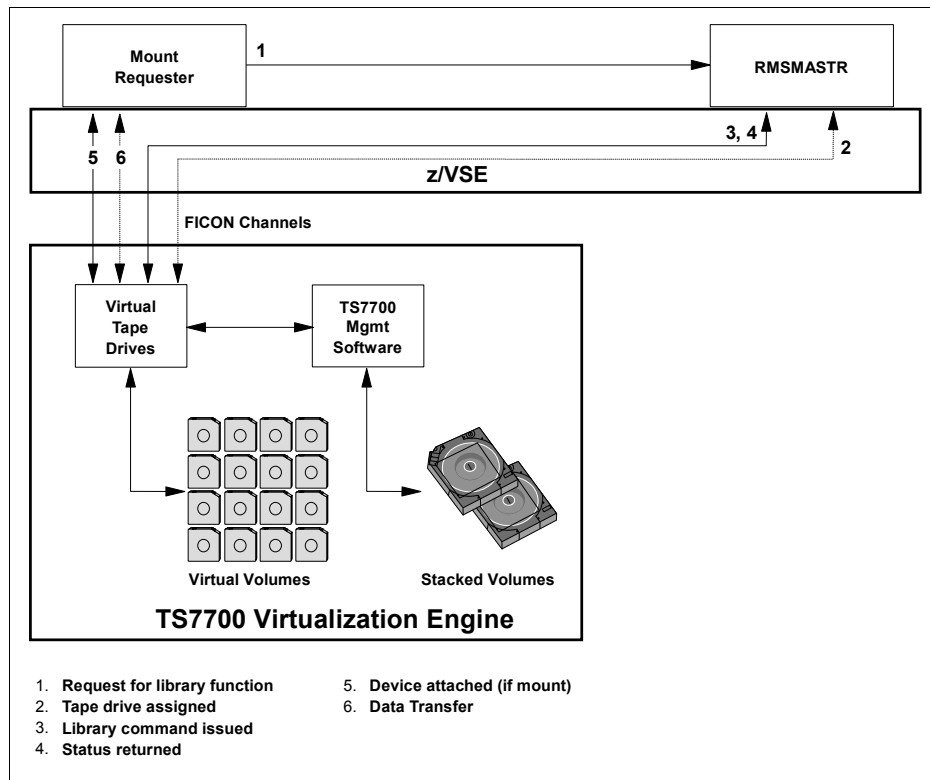


Figure 5-15 TS7700 Virtualization Engine in a z/VSE environment as a VM guest (no VGS)

VSE uses OEM tape management products that support scratch mounts, so if you are using VSE under VM, you have the benefit of using the fast-ready attribute for the VSE library's scratch category.

For more information about z/VSE, see *z/VSE V4R1.0 Administration*, SC33-8304.

5.5 Software implementation in Transaction Processing Facility

This section describes the support for a TS7700 Virtualization Engine in a native Transaction Processing Facility (TPF) environment with TPF V4.1 or z/TPF. The TPF control program and several new and modified TPF E-type programs support the TS7740 Virtualization Engine and TS7720 Virtualization Engine. The support is limited to a command-based interface.

Because TPF does not have a tape management system or a tape catalog system, customers usually rely on z/OS to manage this function. In a TPF environment, most tape data is passed between the systems. In general, 90 percent of the tapes are created on TPF and read on z/OS; the remaining 10 percent are usually created on z/OS and read on TPF.

Be sure to use the normal z/OS and TS7700 Virtualization Engine installation process.

5.5.1 Usage considerations for TS7700 with TPF

TPF uses virtual volumes from the z/OS scratch pools and shares the TS7700 Virtualization Engine scratch categories with z/OS. The z/OS host performs the insert processing for these virtual volumes and continues to manage them based on the input obtained from TPF. TPF has a set of commands (**ztp1f**) to make loading the volumes in TPF-allocated virtual drives possible.

After a volume is loaded into a TPF drive, you have an automated solution in place that passes the volume serial number (VOLSER), the tape data set name, and the expiration date over to z/OS to have it processed automatically.

On z/OS, you must update the tape management system's catalog and the TCDB so that z/OS can process virtual volumes that have been created by TPF. After the TPF-written volumes have been added to the z/OS tape management system catalog and the TCDB, normal expiration processing applies. When the data on a virtual volume has expired, and the volume is returned to scratch, the TS7700 Virtualization Engine internal database is updated to reflect the volume information maintained by z/OS.

Specifics for TPF and z/OS with a shared TS7700 Virtualization Engine

From the virtual drive side, TPF must be allocated certain drive addresses. This information depends on what tape functions are needed on TPF and can vary with the customer set. Therefore, the TS7700 Virtualization Engine will have tape addresses allocated to multiple TPF and z/OS systems and can be shared by dedicating device addresses to certain systems.

Tapes that are created on z/OS and read into TPF

Tapes that are created on z/OS and read into TPF use the same z/OS process for creating tapes. Now when TPF wants to read this z/OS-created tape, it does a specific mount of the tape VSN into a TPF-allocated drive using the TPF (**ztp1f**) commands.

TS7700 Virtualization Engine performance for TPF

You may use the normal TPF Data Collection and Reduction reports that summarize read and write activity to the TPF-allocated drive. For TS7700 Virtualization Engine-specific performance, use the normal TS7700 Virtualization Engine statistics that are off-loaded to z/OS through the TS7700 Virtualization Engine BVIR function.

Support of large virtual volumes for TPF (2 GB and 4 GB)

TPF itself does not use functions such as Data Class to control the logical volume size for specific mounts. User exits provide the ability to set construct names for a volume. If you are not using the user exits, you can set the default size through the TS7700 Virtualization Engine management interface during logical volume insertion, as described in 4.5, "Implementing Outboard Policy Management for non-z/OS hosts" on page 260.

Consider the following information when you implement a TS7700 Virtualization Engine in a TPF environment:

- ▶ Reserving a tape category does not prevent another host from using that category. The user is responsible to monitor the use of reserved categories.
- ▶ Automatic insert processing is not provided in TPF.
- ▶ Currently, no IBM tape management system is available for TPF.

Advanced Policy Management is supported in TPF through a user exit. The exit is called any time a volume is loaded into a drive. At that time, the user can specify, through the TPF user

exit, whether the volume should inherit the attributes of an existing volume using the clone VOLSER attribute or the code can elect to specifically set any or all of the storage group, Management Class, storage class, or data class construct names. If the exit is not coded, the volume attributes remain unchanged, as the volume is used by TPF.

For the two levels of TPF, two separate APARs are required for this support:

- ▶ For TPF V4.1 the APAR number is PJ31643.
- ▶ For zTPFV 1.1, the APAR number is PJ31394.

Library interface

TPF's only operator interface to the TS7700 Virtualization Engine is a TPF functional message, ZTPLF. The various ZTPLF functions allow the operator to manipulate the tapes in the library as operational procedures require. These functions include Reserve, Release, Move, Query, Load, Unload, and Fill. For more information, see *IBM TotalStorage 3494 Tape Library: A Practical Guide to Tape Drives and Tape Automation*, SG24-4632.

Control data sets

The TPF host does not keep a record of the volumes in the TS7700 Virtualization Engine Tape Library or manage the tape volumes in it. You can use the QUERY command to obtain information about the tape volumes held in the TS3500/3952 Tape Library.

SIM and MIM presentation

SIMs and MIMs report hardware-related problems to the operating system. See the *Statistical Analysis and Reporting System User Guide*, which you can access on the web at the following address:

<http://www-1.ibm.com/support/docview.wss?uid=ssg1S7000247>

SIM and MIM are represented in TPF by EREOP reports and the following messages:

- ▶ CEFR0354
- ▶ CEFR0355W
- ▶ CEFR0356W
- ▶ CEFR0357E
- ▶ CEFR0347W
- ▶ CDFR0348W
- ▶ CDFR0349E

5.5.2 Performance considerations for TS7700 Multi-Cluster Grids with TPF

TPF customers who use a TS7700 Grid need an understanding of how the clusters within the grid communicate with one another and how this communication can influence the way that these devices should be used within TPF.

Overview

When clusters are operating within a TS7700 Grid, they share information about the status of volumes and devices. Certain operations initiated by TPF require all the clusters in the Grid to communicate with one another. Under normal conditions, this communication occurs without delay and no impact to TPF. In addition, if one cluster has failed and the other clusters in the Grid have recognized that condition, the communication with that cluster is no longer needed.

The issue with TPF arises because the period that clusters wait before recognizing that another cluster in the Grid has failed exceeds the timeout values on TPF. This issue also means that during this recovery period, TPF is unable to perform any ZTPLF commands that

change the status of a volume, including the loading of tapes or changing the category of a volume through a **ZTPLF** command or through the tape category user exit in segment CORU. The recovery period when a response is still required from a failing cluster can be as long as six minutes. Attempting to issue a tape library command to any device in the grid during this period can render that device inoperable until the recovery period has elapsed even if the device is on a cluster that is not failing.

To protect against timeouts during a cluster failure, TPF systems must be configured to avoid issuing tape library commands to devices in a TS7700 grid along critical code paths within TPF. This can be accomplished through the tape category change user exit in segment CORU. To isolate TPF from timing issues, the category for a volume should not be changed if the exit has been called for a tape switch. Be sure that the exit changes the category when a volume is first loaded by TPF and then not changed again.

To further protect TPF against periods in which a cluster is failing, TPF must keep enough volumes loaded on drives that have been varied on to TPF so that the TPF system can operate without the need to load an additional volume on any drive in the Grid until the cluster failure has been recognized. TPF must have enough volumes loaded so that it can survive the six-minute period where a failing cluster prevents other devices in that grid from loading any new volumes.

Important: Read and write operations to devices in a grid do not require communication between all clusters in the grid. By eliminating the tape library commands from the critical paths in TPF, helps TPF tolerate the recovery times of the TS7700 and read or write data without problems if a failure of one cluster occurs within the grid.

Another configuration consideration pertains to volume ownership. Each volume in a TS7700 Grid is owned by one of the clusters in the Grid. When a scratch volume is requested from a category for a specific device, then a volume that is owned by the cluster that device belongs to is selected if possible. TPF systems should always be configured so that any scratch category is populated with volumes that are owned by each cluster in the Grid. In this manner TPF will have access to a scratch tape that is owned by the cluster that was given the request for a scratch volume. If all the volumes in a grid are owned by one cluster, then a failure on that cluster will require a cluster takeover (which can take tens of minutes) before volume ownership can be transferred to a surviving cluster.

Guidelines

When TPF applications use a TS7700 multi-cluster grid represented by the Composite Library, the following usage and configuration guidelines can help you meet the TPF response time expectations on the storage subsystems:

- ▶ The best configuration is to have the active and standby TPF devices and volumes on separate Composite Libraries (either single- or multi-cluster grid). This configuration prevents a single event on a Composite Library from affecting both the primary and secondary devices.
- ▶ If the active and standby TPF devices/volumes will be configured on the same composite library in a grid configuration, be sure to use the following guidelines:
 - Change the category on a mounted volume only when it is first mounted through ZTPLF LOAD command or as the result of a previous ZTPLF FILL command.

This change can be accomplished through the tape category change user exit in segment CORU. To isolate TPF from timing issues, the category for a volume must never be changed if the exit has been called for a tape switch. Be sure that the exit changes the category when a volume is first loaded by TPF, and then not changed again.
 - TPF must keep enough volumes loaded on drives that have been varied on to TPF so that the TPF system can operate without the need to load additional volumes on any drive in the grid until a cluster failure has been recognized and the cluster isolated. TPF must have enough volumes loaded so that it can survive the six-minute period when a failing cluster prevents other devices in that grid from loading any new volumes.
 - TPF systems must always be configured so that any scratch category is made up of volumes that are owned throughout all the various clusters in the grid. This way assures that during cluster failures, volumes on other clusters are available for use without having ownership transfers.
- ▶ The RUN Copy Consistency point should be used only for the cluster that will be used as the TVC. All other clusters must be configured with the Deferred consistency point to avoid timeouts on the close of the volume.



Upgrade considerations

In this chapter, we explain the component upgrades for IBM Virtualization Engine TS7700:

- ▶ TS7700 Virtualization Engine components upgrades, that is, adding hardware or functionality to existing TS7700 Virtualization Engines
- ▶ TS7700 Virtualization Engine Release 1.7 of Licensed Internal Code upgrade for existing IBM Virtualization Engine TS7740 environments
- ▶ Adding more Clusters to a grid

6.1 TS7700 Virtualization Engine component upgrades

Several field-installable upgrades give an existing TS7700 Virtualization Engine additional functions or capacities. In this section, we review the TS7700 Virtualization Engine component upgrades.

6.1.1 TS7700 Cache Upgrade options

TS7700 R1.7 introduces cache upgrades for both TS7720 and TS7740 configurations, which we discuss in this section.

Note: All cache upgrades, other than for CC6/CX6 cache configurations, require the TS7700 to be at R1.7.

Figure 6-1 shows possible cache configurations that are available for already-installed TS7740 configurations by adding cache modules, which include the larger disk drives.

1.5 TB	CX6 – 150GB		
1.5 TB	CX6 – 150GB		
1.5 TB	CX6 – 150GB		
1.5 TB	CX6 – 150GB		
1.5 TB	CC6 – 150GB		
3.4 TB	CX7 – 300GB		
3.4 TB	CX7 – 300GB		
3.4 TB	CX7 – 300GB		
3.4 TB	CC7 – 300GB		
7.04 TB	CX7 – 600GB		
7.04 TB	CX7 – 600GB		
7.04 TB	CX7 – 600GB		
7.04 TB	CX7 – 600GB		
3.4 TB	CC7 – 300GB		
3.4 TB	CC7 – 300GB		

Figure 6-1 TS7740 Cache configuration upgrade options

The upgrade options depicted in Figure 6-1 are as follows:

- ▶ Existing configurations with a Model CC6 controller installed (top left corner) are limited to a maximum of 9 TB. Up to 6 cache models CX6 can be installed.
- ▶ Fully configured configurations with a Model CC7 Cache Controller and three Model CX7 cache modules (top right corner) cannot be upgraded further and are limited to a maximum of 1.2 TB.
- ▶ Existing configurations with a Model CC7 cache controller and one Model CX7 cache module can be upgraded with two additional Model CX7 Cache Modules containing the new, larger disk drives for a maximum of 20.16 TB.
- ▶ Existing configurations with one Model CC7 cache controller can be upgraded with one or three additional Model CX7 Cache Modules containing the new, larger disk drives for a maximum of 24.16 TB.

TS7720 Virtualization Engine configurations at R1.7 can be upgraded as shown in Figure 6-2. On the left side, you see the upgrade options of an TS7720 currently configured with 40 TB of cache. On the right side, you see the upgrade options for a TS7720 configured with 70 TB of cache.

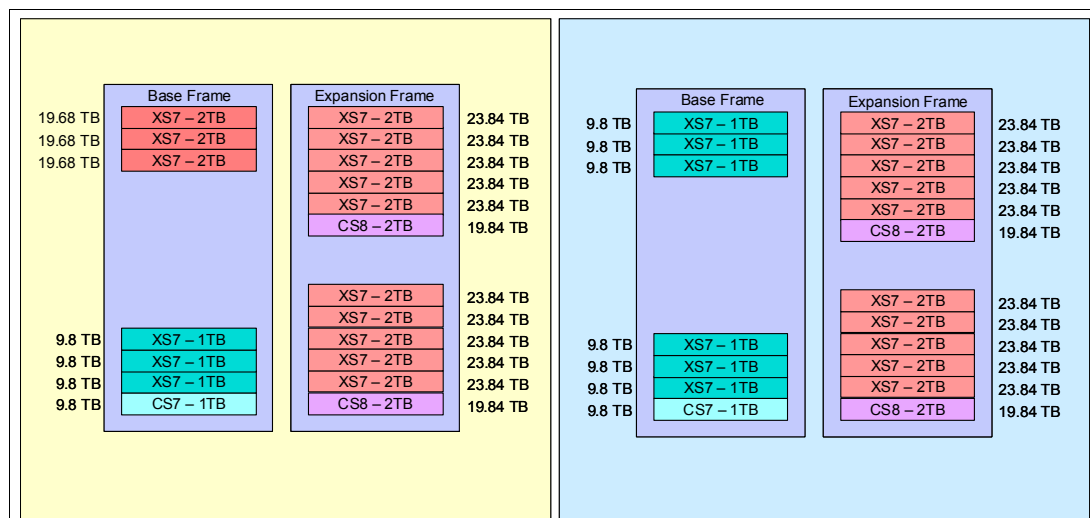


Figure 6-2 TS7740 Cache configuration upgrade options

The left part of Figure 6-2 shows that you can upgrade an existing TS7720 Base Frame from about 40 TB to up to 98.24 TB in increments of one Model XS7 cache module, which represents 19.98 TB. When you add an Expansion Frame, a minimum of two Model CS8 Cache Controllers must be installed in this Expansion Frame. Additional XS7 Cache Modules can be added in increments of one up to a total capacity of the TS7720 of 376.32 TB.

If you already have a fully configured TS7720 frame installed, you can add an Expansion Frame equipped with a minimum of two Model CS8 Cache Controllers. Additional XS7 Cache Modules can be added in increments of one up to a total capacity of the TS7720 of 346.32 TB.

As already mentioned, these cache upgrades require that the TS7700 Virtualization Engines are at the R.1.7 Licensed Internal Code level.

6.1.2 TS7700 Virtualization Engine concurrent system component upgrades

Concurrent system upgrades can be installed while the TS7700 Virtualization Engine is online and operating. The following component upgrades can be made concurrently to an existing, on-site TS7700 Virtualization Engine:

- FICON adapters

You can install FICON adapters to convert from a two to a four-FICON configuration, or replace one pair of FICON adapters of one type with a pair of another type.

- TS7740 Virtualization Engine incremental disk cache capacity

You can add 1 TB increments of disk cache to store logical volumes, with up to 9 TB for 3956-CC6/CX6, up to 14 TB for 3956-CC7/CX7 configurations, and up to 28 TB for 3956-CC8/CX7 configurations.

- ▶ Incremental data throughput

You can add a 100 Mbps increment of peak data throughput for TS7700 Virtualization Engine as follows:

- From one to six to the TS7740 Virtualization Engine
- From one to five to the IBM Virtualization Engine TS7720, where the first increment is installed from the plant

- ▶ TS7740 Encryption

With TS1130 tape drives installed and with TS7740 at R1.5 or later, implementing encryption is nondisruptive

- ▶ TS7720 Expansion frame

When adding an expansion frame, the cache upgrade is performed concurrently. Adding cache controllers or cache modules to the TS7720 Base or Expansion Frame still is disruptive.

6.1.3 TS7700 non-concurrent component upgrades

Non-concurrent upgrades require the TS7700 Virtualization Engine to be brought offline prior to installation. In some instances, the targeted component must be reconfigured before the upgrade will take effect. The component upgrades listed in the following sections must be made non-concurrently to an existing, on-site TS7700 Virtualization Engine.

Memory size

With TS7700 R1.7, the main processor memory of the TS7720 and the TS7740 can be upgraded from 8 GB to 16 GB. All systems available after June 2010 already include the larger memory (FC9461); existing installations can be upgraded by installing FC3461 (8 GB Memory Upgrade - Field) on the 3957 Model V06 or VEA. TS7700 R1.7 is required for the memory upgrade, but the upgrade to the R1.7 Licensed Internal Code level does not require the memory upgrade.

Install the memory upgrade in TS7700 Grids that have more than 500,000 logical volumes defined.

Cache capacity

The following upgrade options are available:

- ▶ TS7720 Virtualization Engine upgrades from four and more cache drawers
 - Cache upgrade from 40 TB to 59.68 TB, 79.36 TB, or 98.24 TB
 - Cache upgrades by adding additional cache modules to the TS7720 Expansion Frame

Note: Adding a TS7720 Expansion Frame including two CS8 Cache Controllers and from one to ten XS7 Cache Modules can be done concurrently.

- ▶ TS7740 Virtualization Engine one to two cache modules
 - Cache upgrade from 1 to 3 TB for 3956-CC6 and 3956-CX6 configurations
 - Cache upgrade from 3.4 to 10.8 TB for 3956-CC7 and 3956-CX7 configurations

- ▶ TS7740 Virtualization Engine two to four Cache Drawers
 - Cache upgrade from 3 to 6 TB for 3956-CC6 and 3956-CX6 configurations
 - Cache upgrade from 6.8 to 20.88 TB for 3956-CC7 and 3956-CX7 configurations
- ▶ TS7740 Virtualization Engine four to six cache modules
 - Cache upgrade from 6 to 9 TB for 3956-CC6 and 3956-CX6 configurations

Note: Cache modules include Cache Controllers and Cache Drawers.

Grid connections

The following upgrade options are available:

- ▶ Shortwave Fibre Ethernet
 - You can add a 1 Gb shortwave Fibre Ethernet adapter for Grid communication between TS7700 Virtualization Engines.
 - You can upgrade to a dual-ported Ethernet adapter for both TS7740 Virtualization Engines and TS7720 Virtualization Engines (FC1032 and FC1033).
- ▶ Copper Ethernet
 - You can add a 1 Gb copper Ethernet adapter for Grid communication between TS7700 Virtualization Engines.

See 3.1.6, “TS7700 Virtualization Engine upgrades” on page 120 for more information about planning for upgrades.

6.2 TS7700 Virtualization Engine upgrade to Release 1.7

Existing TS7740 Virtualization Engine machines can be upgraded to Release 1.7. To migrate to Release 1.7, the existing TS7740 Virtualization Engine must be at least at the 8.4.0.39 (R1.4 PGA4) level or later. Existing TS7720 Virtualization Engines can be upgraded to R1.7 without restrictions.

6.2.1 Planning for the upgrade

A Release 1.7 Licensed Internal Code upgrade is a disruptive activity in a stand-alone cluster. A Licensed Internal Code update is done by an IBM System Service Representative (SSR). Some preinstallation planning is needed, and also an outage schedule. If you are planning this Licensed Internal Code upgrade, contact your SSR to plan the activities.

When updating code in a cluster in a Grid configuration, planning an upgrade to minimize the time that a Grid will operate at different code levels is important.

Prior to starting a code upgrade, all devices in this cluster must be quiesced and varied offline. If the devices are in a Grid configuration, the cluster must be put into Service Prep mode.

Note: Within the Grid, new functions or features are *not* usable until all clusters within the Grid are updated to the same Licensed Internal Code level.

The management interface in the cluster being updated is not accessible during installation. You can use a web browser to access the remaining clusters if necessary.

Apply the required software support prior to performing the Licensed Internal Code upgrade. Although no new software support for TS7700 R1.7 exists, be sure to check whether all software support for previous releases are installed.

Software support for R1.5 included support, for example, for TS7720, workload balancing enhancements, and Device Allocation Assist. Support was provided in the following z/OS APARs:

- ▶ OA24964
- ▶ OA24966
- ▶ OA24967
- ▶ OA24968
- ▶ OA24969
- ▶ OA24970
- ▶ OA26414
- ▶ OA30718

TS7700 R1.6 included support, for example, for logical WORM and four-cluster grid management, and was provided in the following APARs:

- ▶ OA28632
- ▶ OA28633
- ▶ OA28634
- ▶ OA28635
- ▶ OA28636
- ▶ OA28637

Note: Although no new APARs for TS7700 R1.7 are available, check the D/T3957 PSP bucket for any recommended maintenance.

See 3.5.2, “Software requirements” on page 140 for more details.

We describe the implementation of new functionality, introduced with previous releases, in Chapter 5, “Software implementation” on page 263. TS7700 R1.7 includes enhanced support for TS7720 cache management in a hybrid multi-cluster grid configuration, which is defined in the Management Class on the TS7700. See Chapter 4, “Hardware implementation” on page 175 for more details.

6.2.2 Upgrading from TS7700 R1.5 or later

If you are upgrading from R1.5 or from R1.6, the code upgrade activity will be performed by the SSR. In a multi-cluster grid configuration, the code upgrade can be performed concurrently by taking one cluster after the other into service and upgrading to R1.7. After resynchronization of all clusters in the Grid, the update can be done on the next cluster in the Grid.

6.2.3 Upgrading from TS7700 R1.4

If you are upgrading from R1.4, you must be aware that, as part of the upgrade, the external Library Manager in the 3953-F05 frame (if an IBM System Storage TS3500 Tape Library is attached) and the Library Manager Web Specialist will no longer be used for the TS7740 Virtualization Engine, and all construct information and configuration will be done through the TS7740 Virtualization Engine management interface (MI).

If the TS7740 Virtualization Engine is attached to an IBM System Storage 3494 Tape Library, the 3494 Library Manager is still used to mount physical volumes and continues to manage for TS7700 Virtualization Engine physical drives. The MI is used for management of constructs and logical volumes. For TS7740 releases 1.5 to 1.7, the IBM 3494 is supported only through RPQ.

The TS7700 Virtualization Engine Release 1.7 upgrade supports both stand-alone TS7700 Virtualization Engines and TS7700 Virtualization Engines within a Grid configuration, hybrid or homogeneous. In this section, we describe the upgrade process in detail.

Note the following information about attachment to a TS3500 Tape Library:

- ▶ A 3953-L05 Integration is required if the existing TS7740 Virtualization Engine was running R1.4 level.
- ▶ The upgrade migrates all of the targeted TS7740 Virtualization Engines 3953-L05 data to the TS7740 Virtualization Engine.
- ▶ At the end of the upgrade, the TS7740 Virtualization Engine will no longer be attached to or use the external 3953-L05 Library Manager in the 3953-F05 frame.

See “Replacing the 3953-L05 Library Manager” on page 321 for more information.

Replacing the 3953-L05 Library Manager

Upgrading to TS7700 Virtualization Engine Release 1.7 for TS7740 Virtualization Engine machines that are attached to a TS3500 Tape Library through the 3953-L05 Library Manager requires additional planning. As we mentioned, with Release 1.5, the 3953-F05 Library Manager is not needed anymore for the TS7740 Virtualization Engine.

To better understand the hardware changes that are associated with the upgrade to R1.5 and later, see Figure 6-3 on page 322. This figure illustrates a TS7740 Virtualization Engine attached to a TS3500 Tape Library for a pre-Release 1.5 configuration.

The Fibre Channel switches used by the TS7740 Virtualization Engine still reside in the 3953-F05 frame. The second and third clusters are optional. Also, the tape drives and tape library are not shown for the second and third clusters.

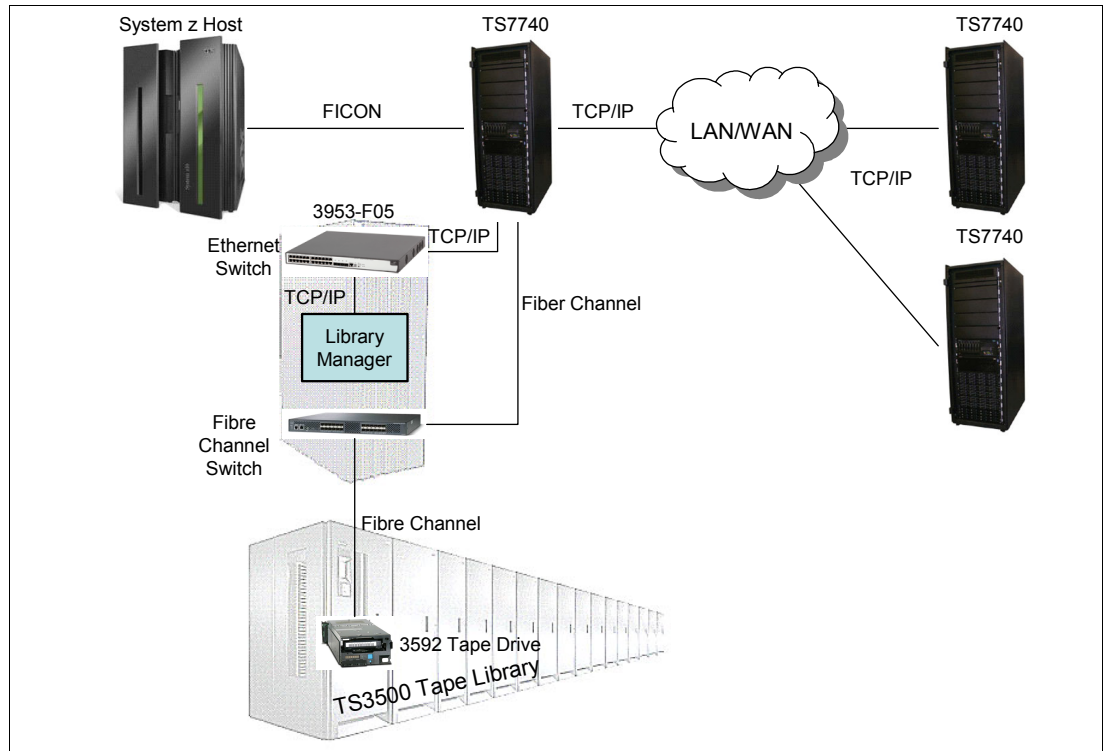


Figure 6-3 Pre-Release R1.5: TS7740 Virtualization Engine, TS3500 Tape Library, and 3953

The 3953-L05 Library Manager is used to manage the logical and physical volumes of the TS7740 Virtualization Engine up to R1.4a. The following diagrams illustrate the TS7740 Virtualization Engine attached to a TS3500 Tape Library for a Release 1.5 or later configuration. The 3953-F05 Library Manager is not part of this configuration, because it is no longer used.

The Fibre Channel switches that are used by the TS7740 Virtualization Engine can either remain in the 3953-F05 (Figure 6-4 on page 323) frame, or reside in the TS3500 Tape Library (Figure 6-5 on page 323). The second to fourth clusters are optional. Also, the tape drives and tape library are only shown for the first cluster.

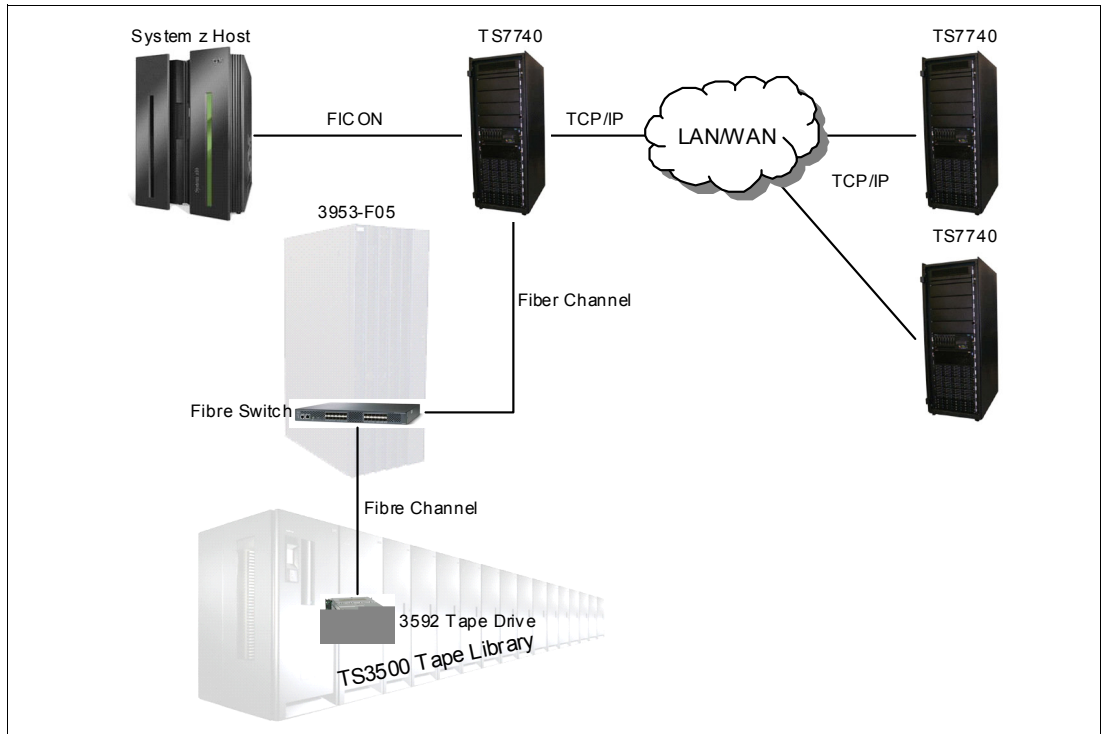


Figure 6-4 R1.5 or later: TS7740, TS3500, and Fibre Channel Switches in 3953

If the TS3500 Tape Library also includes TS110x tape drives attached the mainframe host through an IBM 3592 Tape Controller, the 3953 Library Manager remains required for these native, mainframe-attached drives and the cartridges not managed by the TS7740. In this case, the configuration shown in Figure 6-5 is not valid.




Figure 6-5 R1.5 or later: TS7740 and Fibre Channel Switches in TS3500 Tape Library

If the TS7740 Virtualization Engine is the only system attached to the 3953-L05 and the Fibre Channel Switches are moved to a frame in the TS3500 Tape Library, the 3953-F05 frame is no longer needed and can be uninstalled, which will result in the configuration shown in Figure 6-5 on page 323.

To remove the 3953-F05 frame from the configuration, consider the following aspects:

► Fibre Switches

Depending on the initial and the target configuration, a combination of the following feature codes must be ordered to move the Fibre Channel Switches from the 3953-F05 frame to one TS3500 Tape Library drive frame:

- FC4748, Remove 4 Gb Switch: Removes the 4 Gb Fibre Channel Switches from the 3953-F05. If FC4748 is selected, the switches removed from the 3953 frame can be inserted into the TS3500 Tape Library frame (FC4873) along with the switch mounting hardware (FC4871).
- FC4871, TS7700 Virtualization Engine BE Switch Mounting HW: Switch Mounting hardware for the TS3500 Tape Library frame. FC4871 is always needed whether switches are being reused from the 3953 (removed with FC4748) or new switches are ordered (using FC4872).
- FC4872, TS7700 Virtualization Engine BE 4 Gb Switch: 4 Gb 20 Port Fibre Switch. If the switches removed from the 3953 (by FC4748) will not be re-inserted into the 3584 frame, then new switches must be ordered using FC4872.
- FC4873, Reinstall TS7700 Virtualization Engine BE Switch: Reinstalls “removed” switches in a TS3500 Tape Library frame.
- FC4874, Adjacent Frame Support for BE Switch: Provides cabling to support an adjacent drive frame.

► TS3000 System Console (TSSC)

The TS7740 Virtualization Engine requires the TSSC; therefore if the TSSC connected is installed in the 3953-F05 that is being discontinued/removed, then one of the following actions is required:

- Reconfigure the TS7740 Virtualization Engine to another TSSC at the customer site.
- Order a new TSSC to be installed in the TS7740, consisting of:
 - FC2730, TS3000 System Console: Provides an 1-U server, keyboard, display, mouse, and Ethernet switch for installation in a TS7740 Virtualization Engine frame.
 - FC2719 Console Upgrade: Provides memory and network adapter upgrades for FC2730.
 - FC1904 Redundant AC power: Provides additional power connectors for FC2730.

Note: The TSSC installed in the 3953-F05 frame is a part of that system and must not be removed. Installation of a TSSC into the TS7740 Virtualization Engine frame requires additional hardware (KVM assembly and a network switch). This hardware must *not* be removed from the 3953-F05.

If other systems (a TS1120 control unit or VTSs) are attached to the 3953-L05 Library Manager, the 3953-L05 is still needed for those systems. In this case, the 3953-F05 frame cannot be uninstalled and TS7740 Virtualization Engine Fibre Channel Switches can remain installed inside the frame.

In either case, the relocation of TS7740 Virtualization Engine back-end Fibre Channel Switches inside one of the TS3500 Tape Library drive frame is optional. Relocate the

switches to the TS3500 Tape Library if the TS7740 Virtualization Engine is the only system attached to the 3953-L05 Library Manager; this action simplifies the environment,

The TS7740 Virtualization Engine back-end drives can reside in the same frame as the Fibre Channel Switches or within three frames from the frame containing the switches.

Upgrade options and tasks

Depending on the existing physical configuration, two possible options exist:

- ▶ If the TS7740 Virtualization Engine is the *only* subsystem on the 3953-L05, the following information applies:
 - The upgrade uses the existing 3584 logical library that is used by the 3953-L05.
 - The SSR is instructed to disconnect the network cables between the TS7740 Virtualization Engine and the 3953-L05.
 - If the back-end Fibre Channel Switches are to be converted or TSSC changes are needed, those tasks are performed at this time.
 - After the tasks are complete, the SSR is directed back to the TS7740 Virtualization Engine Licensed Internal Code upgrade process to await for the TS7740 Virtualization Engine Licensed Internal Code activation to complete.
- ▶ If the TS7740 Virtualization Engine shares the 3953-L05 with other control units, another VTS, or another TS7740 Virtualization Engine, the following information applies:
 - The upgrade requires a new 3584 logical library specifically for the TS7740 Virtualization Engine.
 - The SSR is instructed to disconnect the network cables between the TS7740 Virtualization Engine and the 3953-L05.
 - If the back-end Fibre Channel Switches are to be converted or TSSC changes are needed, those tasks are performed at this time.
 - The SSR is instructed to use the TS3500 Tape Library Web Specialist to do the following tasks:
 - Create a new logical library.
 - Set up or change the cartridge assignment policy.
 - Reassign the physical tape devices/control paths from the current 3953-L05 logical library to the TS7740 Virtualization Engine logical library.
 - Reassign all physical tape cartridges used by the TS7740 Virtualization Engine from the 3953-L05 logical library to the TS7740 Virtualization Engine logical library.
 - After the tasks are complete, the SSR is directed back to the TS7740 Virtualization Engine Licensed Internal Code upgrade process to await for the TS7740 Virtualization Engine Licensed Internal Code activation to complete.

After the outage, reconfigure the 3953-L05 to remove the TS7740 Virtualization Engine configuration data.

Discontinue the 3953-L05 if it is no longer being used.

6.3 Adding clusters to a grid

The TS7700 Virtualization Engine Clusters can be installed in stand-alone or in multi-cluster grid configurations. In this section, we describe the available options and the required steps to merge clusters or to add clusters to a grid.

6.3.1 TS7700 Virtualization Engine multi-cluster grid considerations

A TS7700 Virtualization Engine Grid refers to one, two, three, or four physically separate TS7700 Virtualization Engine Clusters, connected by means of a *customer-supplied* TCP/IP network.

TCP/IP

The TCP/IP infrastructure connecting a TS7700 Virtualization Engine multi-cluster grid with two or three clusters is known as the *Grid Network*. The term *Grid* refers to the code and functionality that provides replication and management of logical volumes and their attributes in cluster configurations. A multi-cluster grid provides remote logical volume replication and can be used to provide disaster recovery and high availability solutions. A disaster recovery solution is achieved when multiple clusters are geographically distant from one another.

Migrations to a TS7700 Virtualization Engine multi-cluster grid configuration require using the customer-supplied TCP/IP network. Be sure you have the network prepared at the time the migration starts. The TS7700 Virtualization Engine provides two independent 1 Gb copper (RJ-45) shortwave fiber Ethernet links (single- or dual-ported) for Grid network connectivity. Be sure to connect each one through an independent WAN interconnection to be protected from a single point of failure that would disrupt service to both WAN paths from a node. See the following sections:

- ▶ 3.2.3, “TCP/IP configuration considerations” on page 128
- ▶ 3.2.4, “LAN/WAN requirements” on page 131

Multi-cluster grid configuration

ts7720 and ts7740 virtualization engines can be combined in the same multi-cluster grid. The following system configuration upgrades are available for an existing TS7700 single or multi-cluster grid:

- ▶ Two-cluster TS7700 Virtualization Engine Grid
 - You can add a new, empty TS7700 Virtualization Engine Cluster to an existing TS7700 Virtualization Engine Cluster to create a TS7700 Virtualization Engine two-cluster grid.
 - You can combine two existing TS7700 Virtualization Engine Clusters to create a TS7700 Virtualization Engine two-cluster grid.
- ▶ Three-cluster TS7700 Virtualization Engine Grid
 - You can add a new, empty TS7700 Virtualization Engine Cluster to an existing TS7700 Virtualization Engine two-cluster grid to create a TS7700 Virtualization Engine three-cluster grid.
 - You can add an existing TS7700 Virtualization Engine Cluster to an existing TS7700 Virtualization Engine two-cluster grid to create a TS7700 Virtualization Engine three-cluster grid.

- ▶ Four-cluster TS7700 Virtualization Engine Grid
 - You can add a new, empty TS7700 Virtualization Engine Cluster to an existing TS7700 Virtualization Engine three-cluster grid to create a TS7700 Virtualization Engine four-cluster grid.
 - You can add an existing TS7700 Virtualization Engine Cluster to an existing TS7700 Virtualization Engine three-cluster grid to create a TS7700 Virtualization Engine four-cluster grid.
- ▶ Additional configuration options
 - You can add a new TS7700 Cluster at R1.7 to an existing R1.6 Grid consisting of up to three clusters. This option allows you to quickly integrate a new cluster in an existing Grid without having to update all other clusters first. However, this Grid will be limited to the functionality of R1.6 until all its clusters have been upgraded to R1.7.
 - You cannot directly merge two TS7700 two-cluster grids to create a four-cluster grid. If you want to join two existing two-cluster grids, additional steps are required, as described in 6.3.4, “Four-cluster grid configurations” on page 340.
 - In certain configurations and in specific environments, an upgrade to a five- or six-cluster grid may be available through an RPQ. If you need a configuration with more than four clusters, contact your IBM Sales representative for submitting the request.

Notes: For any of these upgrades, FC4015, Grid Enablement, must be installed on all TS7700 Virtualization Engines that are part of a Grid.

Tape Library attachment

Although you can attach as many TS7740 Virtualization Engines to a TS3500 Tape Library as you can create logical libraries containing the minimum number of physical tape drives required for the TS7740, consider attaching clusters that are members of the same Grid, to separate physical TS3500 tape libraries. If all TS7740 Virtualization Engines are attached to the same physical tape library, service on this tape library can affect your entire tape operation. Attachment to separate TS3500 tape libraries can help ensure a high availability environment where the Grid remains to work although one physical library is not available.

An important point is that two TS7740 Virtualization Engines connected to a single 3494 Tape Library cannot be combined to form a two-cluster grid. Each TS7740 Virtualization Engine Cluster in a TS7740 Virtualization Engine two-cluster grid configuration must possess its own tape library.

6.3.2 Merge two TS7700 Single Clusters into a TS7700 two-cluster grid

This section describes how to add a TS7720 Virtualization Engine or TS7740 Virtualization Engine to another existing or new cluster forming a two-cluster grid for the first time. This process is also called a merge when clusters already containing data are joined into a Grid.

Considering the host configuration changes that are needed before you attempt to use the newly joined cluster is important. Before doing host configuration changes, you might need to access the cluster, for example, to remove duplicate volumes that would prevent a successful Grid join. For this reason, the host configuration changes must not be performed until you are absolutely sure that you do not need to access both clusters as stand-alone libraries any more.

This migration scenario applies to TS7720 Virtualization Engine and TS7740 Virtualization Engine with Release 1.7 and is disruptive to both TS7700 Virtualization Engines that are being merged. TS7720 and TS7740 can now be joined in a Hybrid Cluster.

With this scenario, an existing stand-alone single cluster is merged with another existing stand-alone single cluster to create a cluster 0 / cluster 1 Grid configuration. Figure 6-6 shows this scenario.

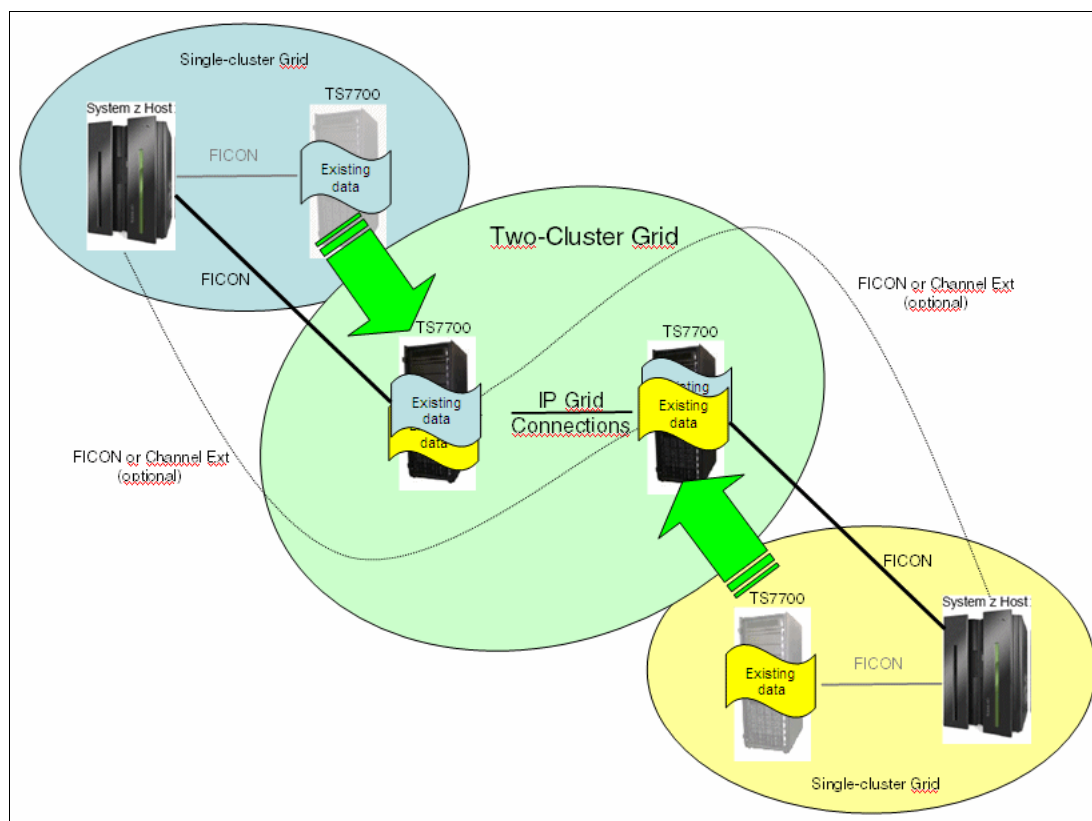


Figure 6-6 Merge of two existing single cluster Grids into one two-cluster grid

Preparation

Before starting the merge process, make sure that all preparatory activities have been completed. Review with your IBM System Service Representative the prerequisites that must be covered before attempting to merge two existing TS7700 Virtualization Engine Single Cluster Grids. The prerequisites might include Licensed Internal Code levels, planning activities, and gathering information from the current environment. A TS7700 Virtualization Engine supports up to one million logical volumes, with a maximum not to exceed two million logical volumes per library manager where attachment to a library manager is required. The one million logical volume limit can be raised to two million with a RPQ. The Grid itself is limited to the maximum number of one million logical volumes, no matter how many clusters compose the Grid. If the current combined amount of logical volumes in the clusters to be joined exceed one million, some logical volumes must be moved to another library or deleted to reach the allowed Grid capacity.

Determine if any conflicting information exists between the TCDB, RMM, and the TS7700. An example of the job is shown in Example F-14 on page 910. You need to check for duplicate logical and physical volumes. Logical volume ranges in a TS7700 Virtualization Engine must be unique. New logical volumes cannot be inserted from the time you start checking for duplicate logical volumes until the join is complete. If duplicate volumes are found during the join process, it has to be stopped to remove the duplicate volumes.

Migration steps

Perform the following steps to accomplish this migration scenario:

Note: The IBM SSR performs steps 5 - 8 as part of the merge. We list them for informational purposes only.

1. Stop host activity on both clusters to be merged:

- a. Complete or cancel all host jobs.
- b. Vary the virtual drives offline.
- c. Vary the existing channels offline.

2. Make changes to HCD channel definitions.

Define the new channels and the device units address in HCD. Define channels and units with OFFLINE=YES, and then vary the channels and units online manually, when and if they are to be used. See Chapter 5, "Software implementation" on page 263 for more information about HCD.

3. Make changes to SMS and TCDB.

For the two-cluster grid, you need one Composite Library ID and two separate Distributed Libraries. Each stand-alone TS7700 Virtualization Engine must have defined one Composite Library and one Distributed Library. When merging two stand-alone Clusters, they will each keep their Distributed Library ID, but they will share one Composite Library Name as depicted in Figure 6-7.

Keep the Composite Library ID of the Cluster that had the most logical volumes defined. For all logical volumes in the other cluster, the volumes entries in the TCDB must be updated to contain the new Library-ID.

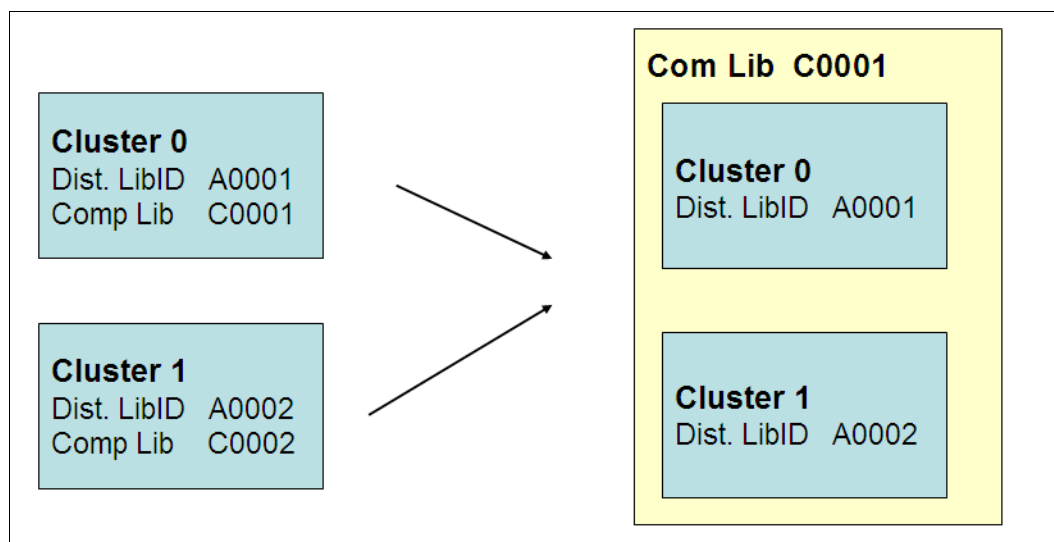


Figure 6-7 Library IDs when merging two clusters

Update RMM and TCDB to match the respective storage groups to be able to have continued access to the original volumes created when the systems were configured as stand-alone clusters. You can get a list of all volumes in the TCDB by running a job, as shown in Example F-16 on page 910.

Updating all the volume entries can be a time-consuming process. Because of the time used, you could change all existing scratch tapes first to prepare for non-specific mounts (write) from production as fast as possible and, after that, take all the rest. See

Example F-17 on page 911 for an example of the JCL for changing one scratch volume and one private volume.

If you are using RMM, an RMM command for each volume is also needed. The command is shown in Example F-18 on page 911. The FORCE parameter is used only when needed and requires access to a specific Facility class in Resource Access Control Facility (RACF®) named STGADMIN.EDG.FORCE. Verify that you have the required authorization.

Delete the SMS library definition of the second Composite Library ID using ISMF.

4. Activate the IODF and the SMS definitions and issue an OAM restart (if it was not done after the SMS activation). For more information, see Chapter 5, “Software implementation” on page 263. Cluster 2 is not ready yet, and will go online in a later step.

Note: If the new SCDS is activated before the new library is ready, the host cannot communicate with the new library yet. Expect message CBR3006I to be issued:

CBR3006I Library library-name with Library ID library-ID unknown in I/O configuration.

5. Configure the TS7700 Virtualization Engine for Grid.
6. Configure and test the TS7700 Virtualization Engine Grid communications.
7. Join the TS7700 Virtualization Engine Clusters:
 - a. Change the distributed library sequence number.
 - b. Change the TS7700 Virtualization Engine Vital Product Data (VPD).
 - c. Merge the TS7700 Virtualization Engine databases' data.
8. Configure and test Autonomic Ownership Takeover Manager (AOTM).
9. Vary devices online to all connected hosts. Now you are ready to validate the two-cluster grid.
10. Run test jobs to read and write to volumes from both original Grids.
11. Modify Copy Policies and Retain Copy Mode in the Management Class definitions according to your needs.

6.3.3 Three-cluster grid

In this section, we cover the migration aspects to join an existing two-cluster grid TS7700 Virtualization Engine with a Single Cluster Grid TS7700 Virtualization Engine to form a three-cluster grid.

Planning for a three-cluster grid

The TS7700 Virtualization Engine Grid configuration may be used to facilitate disaster recovery and high data availability. The TS7700 Virtualization Engine Grid configuration provides support for automatic replication of data and may be used in a variety of disaster recovery situations. If connected at the same location, a Grid configuration might also help facilitate high availability of data.

To configure a three-cluster grid for disaster recovery, you must plan for the following items:

- ▶ Access from your local site's hosts to the FICON channels on the TS7700 Virtualization Engine Cluster located at the disaster recovery site (or sites). This might involve connections using Dense Wavelength-Division Multiplexing (DWDM) or channel extension equipment, depending on the distance separating the sites. If the local TS7700 Virtualization Engine Cluster becomes unavailable, you would use this remote access to continue your operations using a remote TS7700 Virtualization Engine Cluster.
- ▶ Because the virtual devices on a remote TS7700 Virtualization Engine Cluster are connected to the host through channel extensions, there might be a difference in read or write performance compared to the virtual devices on the local TS7700 Virtualization Engine Cluster. If performance differences are a concern, consider using only the virtual device addresses in a remote TS7700 Virtualization Engine Cluster when the local TS7700 Virtualization Engine is unavailable. If these differences are an important consideration, in addition to the ownership takeover procedure, you would need to provide operator procedures to vary the virtual devices in a remote TS7700 Virtualization Engine from online to offline.
- ▶ You might want to maintain separate copy consistency policies for disaster recovery data and data that requires high availability.

The three-cluster grid is built on an existing Grid. It extends the configuration to three clusters combined into a single Composite Library, as shown in Figure 6-8.

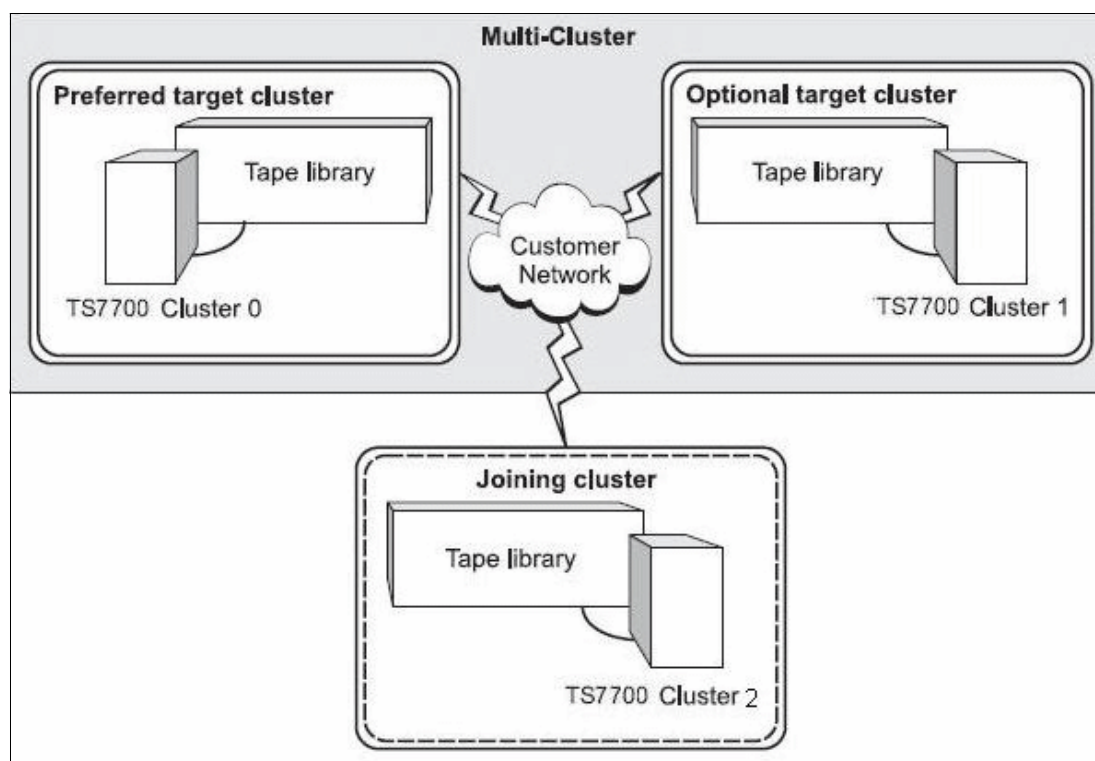


Figure 6-8 TS7700 Virtualization Engine three-cluster grid configuration

The figure shows the required steps for joining a TS7700 Virtualization Engine and a two-cluster grid configuration to form a three-cluster grid configuration.

You must first verify the items discussed in the following sections.

Three-cluster grid configurations

It is possible to configure a TS7700 Virtualization Engine Grid to provide for both disaster recovery and high availability solutions.

The assumption is that two or three TS7700 Virtualization Engine Clusters will reside in different locations, separated by a distance dictated by your company's requirements for disaster recovery. In a three-cluster grid configuration, disaster recovery and high availability can also be achieved simultaneously by ensuring that two local, high availability clusters possess RUN volume copies and have shared access to the host, while the third and remote cluster possesses deferred volume copies for disaster recovery. During a Single Cluster outage, the three-cluster grid solution maintains no single points of failure, which would prevent customers from accessing their data, assuming that copies exist on other clusters as defined in the copy consistency point. See 4.4.3, "Defining Grid copy mode control" on page 245 for more details.

We consider the following two configuration scenarios:

- ▶ Configuration 1:
 - Two of the three clusters will be contained within the same building or campus, as a high availability pair, and both clusters will be connected to one or more production hosts. We confirm that the rewind-unload (RUN) consistency point has been configured between the two-cluster grid.
 - The third cluster is remote and outside the immediate region, and it will have backup host connectivity and devices varied offline. We confirm that the third site that will join the two-cluster grid to form the three-cluster grid has the copy consistency policy configured as deferred.
- ▶ Configuration 2:
 - Two independent production sites exist, where both clusters are connected to independent production host configurations. No copy consistency point has been configured between the production sites, and no copies have been done between the two clusters.
 - The third cluster is remote and outside both regions of the production clusters, and it will have backup host connectivity with the devices varied offline. Copy consistency has been configured as deferred to the third site from both production clusters.

Implementing Configuration 1

In this section, we describe the scenario of going from a two-cluster grid TS7700 Virtualization Engine to a three-cluster grid, as shown in Figure 6-9 on page 333.

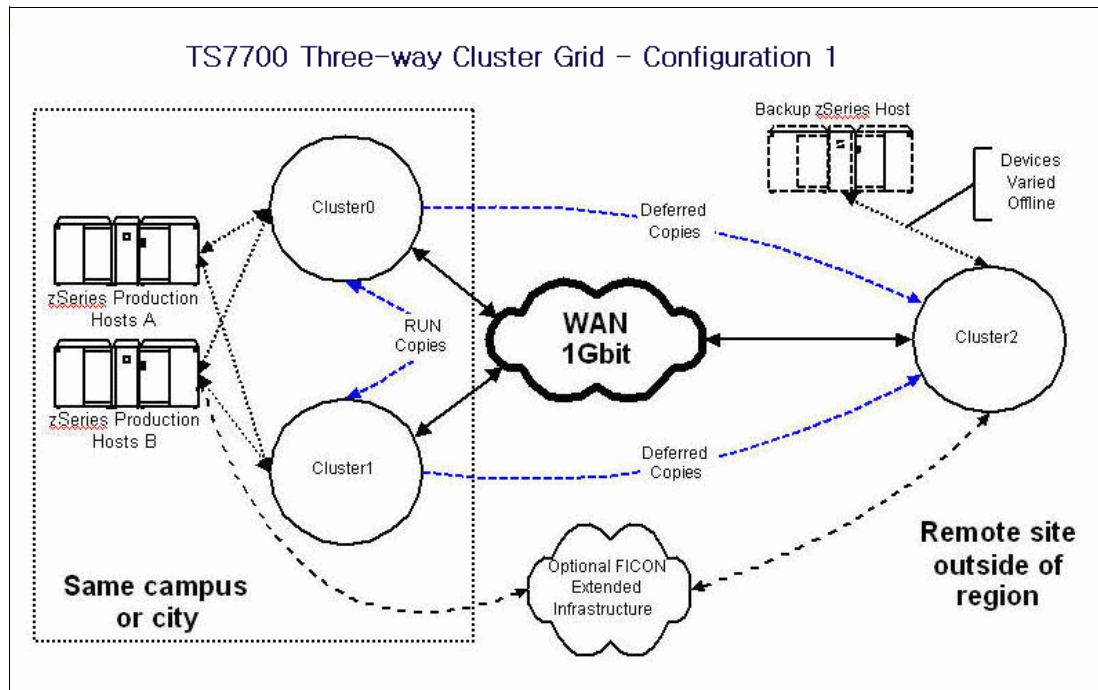


Figure 6-9 TS7700 Virtualization Engine three-cluster grid Configuration 1

In the configuration shown in Figure 6-9, two clusters are in the same campus location or in the same city. The clusters can have one of these Copy Consistency points specified:

- **Rewind/Unload (RUN) copy consistency point**

If a data consistency point of RUN is specified, the data created on one TS7700 Virtualization Engine Cluster is copied to the other TS7700 Virtualization Engine Cluster as part of successful Rewind/Unload command processing, meaning that for completed jobs, a copy of the volume will exist on both TS7700 Virtualization Engine Clusters. Access to data written by completed jobs (successful Rewind/Unload) prior to the failure is maintained by other TS7700 Virtualization Engine Cluster. Access to data of incomplete jobs that were in process at the time of the failure is not provided. We assume that the two existing clusters are using RUN Copy Consistency points specified in the Management Class storage construct, when volumes are written.

- **Deferred copy consistency point**

If a data consistency point of Deferred is specified, the data created on one TS7700 Virtualization Engine Cluster is copied to the other TS7700 Virtualization Engine Cluster after successful Rewind/Unload command processing. Access to the data through the other TS7700 Virtualization Engine Cluster is dependent on when the copy completes. Because some amount of delay occurs in performing the copy, access might not be available when a failure occurs. These are the copy consistency policies we assume are assigned to the third cluster that joins the two-cluster grid with a third cluster to form the three-cluster grid configuration.

- **No Copy consistency point**

If a data consistency point of No Copy is specified, the data created on one TS7700 Virtualization Engine Cluster is not copied to the other TS7700 Virtualization Engine Cluster. If the TS7700 Virtualization Engine Cluster that data was written to fails, the data for that logical volume is inaccessible until that TS7700 Virtualization Engine Cluster's operation is restored. In this configuration, none of the clusters are set up for No Copy.

The new cluster that will join the two-cluster grid to form the three-cluster grid must already be installed. Every cluster in the system requires two network connections to a customer WAN for site to site operations, and the WAN connections between the three clusters in the three-cluster grid must be completed. The Grid network on the new cluster must be configured, containing the IP addresses of the three clusters.

The following tasks are done by the customer and also the SSR, where indicated:

1. Stop host activity on Cluster 1, which must go into Service Prep mode:
 - a. Vary the unit addresses offline.
 - b. Complete or cancel all jobs for Cluster 1.

2. IBM SSR: Set Cluster 1 to Service Prep mode.

3. Make changes to SMS.

With a three-cluster grid, you need one Composite Library and three Distributed Libraries. You must now define the third Distributed Library in SMS. Make sure to enter the correct Library-ID delivered by the SSR.

4. Make changes to HCD.

Define the new channels and the 256 units in HCD. Define channels and units with OFFLINE=YES, and then vary the channels and units online manually, when and if they are to be used. See Chapter 5, "Software implementation" on page 263 for more information about HCD.

5. Set missing Interrupt Handler (MIH) values.

If you are defining specific address ranges in the IECIOSxx member in SYS1.PARMLIB, make sure that MIH values for the new devices are set. Proper values are described in 5.2.6, "Set values for the Missing Interrupt Handler" on page 284.

6. Activate the IODF and the SMS definitions and issue an OAM restart (if it was not done after the SMS activation). For more information, see Chapter 5, "Software implementation" on page 263. Cluster 2 is not ready yet, and will go online in a later step.

Note: If the new SCDS is activated before the new library is ready, the host cannot communicate with the new library yet. Expect message CBR3006I to be issued:

CBR3006I Library library-name with Library ID library-ID unknown in I/O configuration.

7. IBM SSR: Configure the local Grid network from the TS7700 Virtualization Engine System Management Interface Tool (SMIT) window.
8. IBM SSR: Join Cluster 2 to Cluster 1 using SMIT.
 - a. Merge Vital Product Data (VPD).
 - b. Merge the TS7700 Virtualization Engine databases' data.
9. IBM SSR: Exit Service Prep mode for Cluster 1.
10. Vary devices from Cluster 1 online for all connected hosts.
11. After a new cluster is joined to a cluster in an existing Grid, all clusters in the existing Grid are automatically joined. In the previous example, we joined new Cluster 2 to Cluster 1 of an existing Grid. This is enough to make the existing two-cluster grid into a three-cluster grid. Cluster 0 in the example could be operational and available for customer operation during that time.

12. Now you are ready to validate the new three-cluster grid:

- a. Vary Cluster 2 defined FICON channels online, making sure that Cluster 2 can be accessed through the channel paths defined in HCD.
- b. Vary logical devices for Cluster 2 online.
- c. Vary Cluster 2 online to the hosts.
- d. Using the **D SMS,LIB(libraryname),DETAIL** commands, validate that the relationship between the Composite and Distributed Libraries is correct, as shown in Example 6-1 and Example 6-2.

Example 6-1 Display Composite Library

```

D SMS,LIB(COMPLIB),DETAIL
F OAM,D,LIB,COMPLIB,L=ST6T10-Z
CBR1110I OAM LIBRARY STATUS: 141

```

TAPE	LIB	DEVICE	TOT	ONL	AVL	TOTAL	EMPTY	SCRATCH	ON	OP
LIBRARY	TYP	TYPE	DRV	DRV	DRV	SLOTS	SLOTS	VOLS		
COMPLIB	VCL	3957-V06	768	768	287	0	0	368298	Y	Y

```

MEDIA          SCRATCH          SCRATCH          SCRATCH
TYPE           COUNT            THRESHOLD        CATEGORY
MEDIA1         170345            0              0001
MEDIA2         197953            0              0002

```

```

DISTRIBUTED LIBRARIES:  DISTLIB0    DISTLIB1    DISTLIB2

```

```

LIBRARY ID: C0001
OPERATIONAL STATE:  AUTOMATED
ERROR CATEGORY SCRATCH COUNT:          33
CORRUPTED TOKEN VOLUME COUNT:          0

```

```

LIBRARY SUPPORTS IMPORT/EXPORT.
LIBRARY SUPPORTS OUTBOARD POLICY MANAGEMENT.

```

Example 6-2 Display Distributed Library

```

D SMS,LIB(DISTLIB1),DETAIL
F OAM,D,LIB,DISTLIB1,L=ST6T10-Z
CBR1110I OAM LIBRARY STATUS: 062

```

TAPE	LIB	DEVICE	TOT	ONL	AVL	TOTAL	EMPTY	SCRATCH	ON	OP
LIBRARY	TYP	TYPE	DRV	DRV	DRV	SLOTS	SLOTS	VOLS		
DISTLIB1	VDL	3957-V06	0	0	0	1348	819	0	Y	Y

```

COMPOSITE LIBRARY:      COMPLIB

```

```

LIBRARY ID: 10001
OPERATIONAL STATE:  AUTOMATED
SCRATCH STACKED VOLUME COUNT:          222
PRIVATE STACKED VOLUME COUNT:          108

```

```

LIBRARY SUPPORTS IMPORT/EXPORT.
LIBRARY SUPPORTS OUTBOARD POLICY MANAGEMENT.
CONVENIENCE I/O STATION INSTALLED.
CONVENIENCE I/O STATION IN OUTPUT MODE.
BULK INPUT/OUTPUT NOT CONFIGURED.

```

- e. Vary the logical devices for Cluster 2 offline again so that they will be ready to test if the original two-cluster grid still works.

13. Modify Copy Policies defined in the LM Management Class.

The copy consistency points on all three clusters (Table 6-1) must be modified to support a RUN copy between Cluster 0 and Cluster 1 and also a deferred copy from Cluster 0 and Cluster 1 to Cluster 2. The values must be updated in the Management Classes using the TS7700 Virtualization Engine management interface. Make sure that the definitions will work when logical units are allocated from Cluster 2. See 4.4.3, “Defining Grid copy mode control” on page 245 for more information. Table 6-2 on page 339 describes the settings needed for the scenario shown in Figure 6-9 on page 333.

Table 6-1 Copy Consistency point on Management Class: three-cluster grid Configuration 1

Management Class	Cluster 0	Cluster 1	Cluster 2
	From → To	From → To	New cluster
MC Hosts A	RR → RRD	RR → RRD	DDD
MC Hosts B	RR → RRD	RR → RRD	DDD

14. Run test jobs writing to and reading from the two original clusters part of the two-cluster grid.
15. Vary logical devices for Cluster 0 and Cluster 1 offline to be ready to validate the use of Cluster 2 as though there were a disaster, and set up copy consistency points that support the customer requirements, for example, there might be a deferred copy mode.
16. Test write and read with Cluster 2 and validate the result.
17. IBM SSR: The migration is done. Return to normal production mode.
18. If you want part or all of the existing logical volumes to have a logical copy on Cluster 2, this can be done in different ways. IBM has tools, such as PRESTAGE, to support these actions. The logical volumes must be read or referred to retrieve the new management policies that you define. The tools are available at the following address:

<ftp://ftp.software.ibm.com/storage/tapetool/>

Implementing Configuration 2

In this configuration, two clusters are in independent campuses, locations, or cities. We assume that the two existing clusters are running with Deferred Copy Consistency points. Figure 6-10 shows an overview of the configuration.

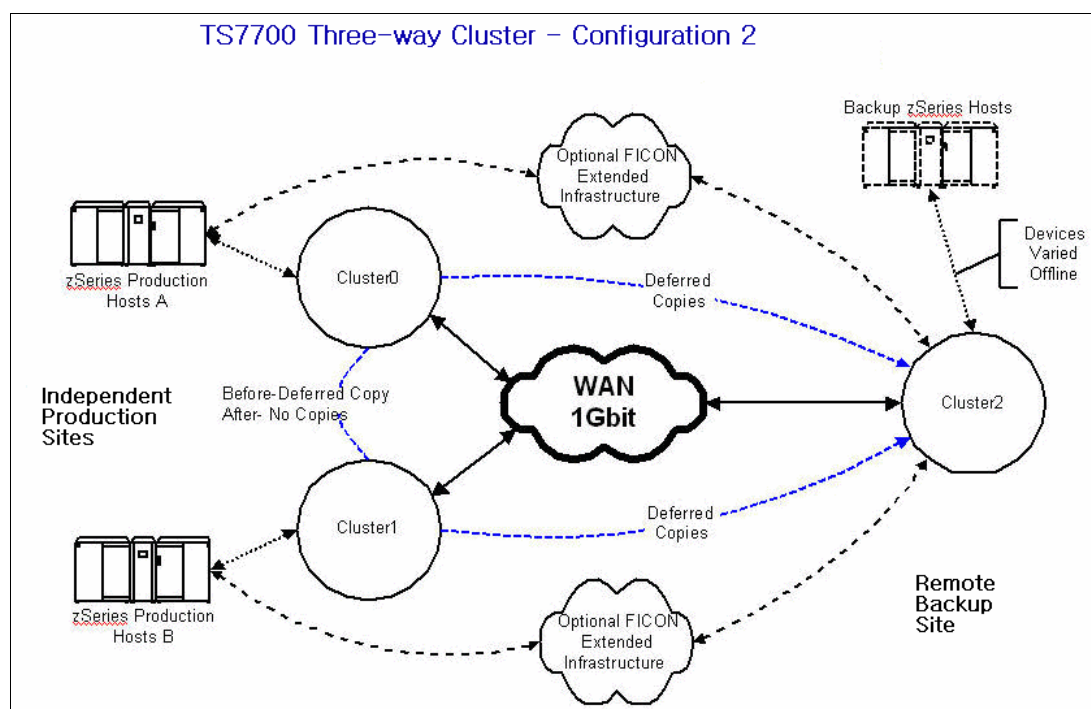


Figure 6-10 TS7700 Virtualization Engine three-cluster grid Configuration 2

The new cluster that will join the two-cluster grid to form the three-cluster grid must be installed in advance. Every cluster in the system requires two network connections to a customer WAN for site to site operations, and the WAN connections between the three clusters in the three-cluster grid must be completed. The Grid network on the new cluster must be configured and contain the IP addresses of the three clusters.

The following are tasks done by the customer and also the SSR, where indicated:

Note: The SSR performs steps 2, 3, 8 - 10, and 14 - 18 as part of the installation. We list them for informational purposes only.

1. Stop host activity on Cluster 1, which must go in Service Prep mode. Cluster 0 stays operational.
2. Vary the unit addresses for Cluster 1 offline. Complete or cancel all jobs for Cluster 1.
3. IBM SSR: Set Cluster 1 to Service Prep mode.
4. Make changes to SMS.

With a three-cluster grid, you need one Composite Library and three Distributed Libraries. You must now define the third Distributed Library in SMS. Make sure to enter the correct Library-ID delivered by the SSR. If the connected hosts are in separate SMS configurations, perform this step twice.

5. Make changes to HCD.

Define the new channels and the 256 units in HCD. Define channels and units with OFFLINE=YES, and then vary the channels and units online manually, when and if they

are to be used. See Chapter 5, “Software implementation” on page 263 for more information about HCD.

6. Set missing Interrupt Handler (MIH) values.

If you are defining specific address ranges in the IECIOSxx member in SYS1.PARMLIB, make sure that MIH values for the new devices are set. Proper values are described in 5.2.6, “Set values for the Missing Interrupt Handler” on page 284. If the connected hosts used different SYS1.PARMLIBs, this step must be repeated.

7. Activate the IODF and the SMS definitions and issue an OAM restart if it was not done after the SMS activation. For more information, see Chapter 5, “Software implementation” on page 263. Cluster 2 is not ready yet, and will go online in a later step. Again, this step must be repeated to reach all hosts.

Note If the new SCDS is activated before the new library is ready, the host cannot communicate with the new library yet. Expect message CBR3006I to be issued:

CBR3006I Library library-name with Library ID library-ID unknown in I/O configuration.

8. IBM SSR: Configure the local Grid network using the TS7700 Virtualization Engine System Management Interface Tool (SMIT) window.

9. IBM SSR: Join Cluster 2 to Cluster 1 using SMIT:

- a. Merge Vital Product Data (VPD).
- b. Merge the TS7700 Virtualization Engine databases' data.

10. IBM SSR: Take Cluster 1 out of Service Prep mode.

11. Vary devices from Cluster 1 online to all connected hosts.

12. Run test jobs to read and write from the original Cluster 0

13. Run test jobs to read and write from the original Cluster 1.

Note: The following steps apply to all connected hosts.

14. Now you are ready to validate the new three-cluster grid:

- a. Vary Cluster 2 defined FICON channels online, making sure that Cluster 2 can be accessed through the channel paths defined in HCD.
- b. Vary logical devices for Cluster 2 online.
- c. Vary Cluster 2 online to the hosts.
- d. Evaluate, using **D SMS, LIB(libraryname), DETAIL** commands, that the relationship between Composite and Distributed Libraries is correct, as shown in Example 6-1 on page 335 and Example 6-2 on page 335.
- e. Vary logical devices for Cluster 2 offline again and be ready to test and see that the original two-cluster grid still works.

15. Modify Copy Policies defined in the TS7740 Virtualization Engine management interface.

The Copy Consistency points on all three clusters must be modified to support deferred copy from Clusters 0 and 1 to Cluster 2. The values must be updated in the Management Classes (from the Library Manager Interface). Make sure that the definitions work when logical units are allocated from Cluster 2. See 4.4.3, “Defining Grid copy mode control” on page 245 for more information.

Table 6-2 describes the setting needed for the scenario shown in Figure 6-10 on page 337.

Table 6-2 Copy Consistency point on Management Class: three-cluster grid Configuration 2

Management Class	Cluster 0	Cluster 1	Cluster 2
	From → To	From → To	New cluster
MC Hosts A	RD → RND	RD → RND	DND
MC Hosts B	RD → NRD	RD → NRD	NDD

16. Vary logical devices for Cluster 0 and Cluster 1 offline to be ready to validate the use of Cluster 2 as though there were a disaster, and set up Copy Consistency points that support the customer requirements. There might be a deferred copy mode.
17. Test write and read with Cluster 2 and validate the result.
18. IBM SSR: The migration is done; return to normal production mode.
19. If you want part or all of the existing logical volumes to have a logical copy on Cluster 2, this can be done in different ways. IBM has tools, such as PRESTAGE, that will help you with this task. The logical volumes must be read or referred to retrieve the new management policies that you define. The tools are available at the following address:

<ftp://ftp.software.ibm.com/storage/tapetool/>

Implementing a three-cluster hybrid grid configuration

The TS7700 three-cluster grid can also be a hybrid configuration, for example, with two TS7720 Virtualization Engines in two separated production centers and one TS7740 Virtualization Engine in a backup center. This configuration can deliver high availability at the production site. Old data will be accessed remotely, unless a host connection is provided to the remote site.

From a host perspective, the implementation steps are the same for a hybrid grid and for a homogeneous grid configuration consisting only of TS7740 or TS7720 Virtualization Engines. The same restrictions and recommendations apply. However, for a hybrid grid, be sure to understand that, after a logical volume has been migrated out of the cache of a TS7720, it can be recalled only through the TS7740. Therefore, be especially careful when you plan for Copy Consistency Points and Cache Management Policies. To avoid unnecessary copies to be created, be sure that you define the Retain Copy Mode setting in the Management Class definition.

Figure 6-11 on page 340 shows an example of a three-cluster hybrid grid configuration.




Figure 6-11 Three-cluster hybrid grid

In this case, the customer has two separate production sites, A and B, and one remote Disaster Recover site. You can implement this scenario using the process discussed in “Implementing Configuration 2” on page 337.

6.3.4 Four-cluster grid configurations

Up to four TS7700 Virtualization Engine Clusters can be configured in one Grid, and TS7740 and TS7720 Virtualization Engines can be intermixed in the same Grid, which is called a hybrid grid.

Configuration and upgrade guidelines

Before you start planning for the migration, consider the following general guidelines for a multi-cluster grid:

- ▶ When merging multiple clusters, or when adding more clusters to an existing Grid, only one can be done at a time. For example, to migrate from a two-cluster grid to a four-cluster grid requires two steps that can be performed within one maintenance window:
 - a. You start with adding one cluster to the existing two-cluster grid. This cluster can be an empty Cluster that is being added, or a Cluster containing data that is being merged.
 - b. After you have successfully tested this three-cluster grid configuration, and after the three clusters have been synchronized, the next cluster is added to upgrade from a three-cluster to a four-cluster grid.

Adding more clusters requires the same steps as adding a second, third or fourth cluster to a grid. These steps are described in detail in “Implementing a four-cluster hybrid grid” on page 342.

- ▶ The merging of two two-cluster grids with data requires additional steps to be taken, because you cannot directly merge them. Consider using the following outline of migration steps:
 - a. Decide which Grid you want to use as the base (Grid A) and which one you want to join (Grid B). Make sure, no duplicate volume serial numbers exist in both Grids. Then, redirect all host activity from Grid B to Grid A and stop all host activity to Grid B.
 - b. Make sure, all logical volumes in Grid B are available in one of the Clusters (Cluster 2).
 - c. Remove the other cluster (Cluster 3) from Grid B and perform a Factory Cleanup on this Cluster.
 - d. Merge the remaining Cluster 2 containing all the logical volumes to Grid A resulting in a three-cluster grid.
 - e. After resynchronization, add the remaining Cluster (Cluster 3) to the Grid resulting in a four-cluster grid.

The IBM SSR performs the actual merge. See 6.3.2, “Merge two TS7700 Single Clusters into a TS7700 two-cluster grid” on page 327 for more details about merging clusters.

- ▶ Whether or not you are implementing a hybrid grid configuration is not really important from an implementation perspective, because you perform the steps related to whether it is a TS7740 or a TS7720 Cluster. However, it is important how you define those parameters that influence cache management and data movement in the Grid:
 - Copy consistency points define where and when a copy of a logical volumes is created.
 - Retain Copy Policy influences whether unnecessary copies are made.
 - Cluster Families can help to avoid unnecessary traffic on the Grid links.
 - Using the software enhancements for workload balancing and Device Allocation Assist influence overall Grid performance.
 - Using the enhanced cache removal policies for TS7720 can improve the cache hit ratio and the overall logical volume mount times.

Four-cluster hybrid configuration

A possible configuration of a four-cluster hybrid grid is shown in Figure 6-12. Two TS7720 Virtualization Engine clusters are installed at the production site, plus two remote TS7740 Virtualization Engines at a backup site.




Figure 6-12 Four-cluster hybrid grid: Configuration 1

A TS7720 Virtualization Engine is used at the production sites. For the remote disaster recovery center, two TS7740 Virtualization Engines are used. The four data centers are interconnected through both FICON directors/DWDM and high-speed WAN.

With this configuration, you can achieve high availability and high performance, because the TS7720 Virtualization Engines provide for high cache hit rates.

Long-term retention data that has been created in the production clusters will be migrated to TS7740 Virtualization Engine clusters in disaster recovery sites, and ultimately written to physical tapes into a physical library at both DR sites. This migration is automatically performed as defined in the enhanced removal policies introduced with TS7720 R1.7.

Implementing a four-cluster hybrid grid

The steps necessary for implementing a four-cluster hybrid grid configuration, as shown on Figure 6-12, are essentially the same steps already described in “Implementing Configuration 1” on page 332 and “Implementing Configuration 2” on page 337. All recommendations regarding HCD, SMS, and so on, apply here.

In the following section, we describe the steps that are required to install the four-cluster hybrid grid shown in Figure 6-12. We start with an existing single-cluster grid, assuming that the existing Cluster 0 is the TS7740 Virtualization Engine in Backup Site X.

Preparation steps

Use the same approach as described in “Implementing Configuration 1” on page 332 for the preparation. In our scenario, we have only one operational cluster, so the first step, going from a single-cluster to a two-cluster grid, is disruptive.

You can perform the following steps before starting the actual migration:

- ▶ The hardware of the new clusters that will join the original cluster to form the two-cluster, three-cluster, and finally the four-cluster grid must already be installed. A new TS7740 Virtualization Engine must be already connected to the new TS3500 Tape Library, with all proper definitions already in place (such as CAP in the TS3500, VOLSER ranges, and Logical Libraries).

See Chapter 4, “Hardware implementation” on page 175 for details about the installation of a TS7740 Virtualization Engine and a TS3500 Tape Library.

- ▶ Every cluster in the system requires two network connections to a customer WAN for site to site operations, and the WAN connections between the four clusters in the four-cluster grid must be completed. The Grid network on the new clusters must be configured and tested, and contain the IP addresses of the four clusters.
- ▶ Upgrade all clusters to R1.7 before starting the actual merge. Although the hybrid four-cluster grid is already supported with TS7700 R1.6, Enhanced TS7720 Cache Removal Policies are only available with TS7700 R1.7.
- ▶ HCD and system definitions must be planned in advance to show the final (four-cluster grid) configuration. Define the new channels and the 256 devices in HCD for each of the three new clusters. Define channels and units as OFFLINE=YES. Channels and devices must be varied online manually only when they are needed in the migration process.
- ▶ Missing Interrupt Handler (MIH) values might need to be updated. If you are defining specific address ranges in the IECIOSxx member in SYS1.PARMLIB, make sure that MIH values for the new devices are set. The proper values are described in 5.2.6, “Set values for the Missing Interrupt Handler” on page 284.
- ▶ Define the new Distributed Libraries using ISMF. With a four-cluster grid, you need one Composite Library, which should remain the same one you already have in place for the first cluster already installed, and four Distributed Libraries. When defining the three new Distributed Libraries, make sure to match your definitions in SMS with the Library-IDs set by the SSR at the new clusters.
- ▶ Plan for the constructs you will be using on the three new clusters. Use the MI to obtain backup files containing the settings to be restored on the TS7720 clusters and the information to be restored on the TS7740 cluster. You can use these files as a basis for the definitions to be made on the new clusters. Figure 6-13 on page 344 shows the settings that are used on the TS7720 (left) and for the TS7740 (right).




Figure 6-13 Backup construct settings through the MI

- Plan for copy consistency points and how to set them during the migration steps. Remember that you can test only the Grid functionality, if you change the default definition of No Copy (N) to either Deferred Copy (D) or RUN Copy (R). Table 6-3 shows example settings for the final configuration.

Table 6-3 Sample Copy Consistency Point definitions 1

Management Class	Cluster 0 (Site X)	Cluster 1 (Site A)	Cluster 2 (Site Y)	Cluster 3 (Site B)
Host in Site A	D	R	D	D
Host in site B	D	D	D	R
Host in Site X	R	D	D	D
Host in Site Y	D	D	R	D

These settings result in each host writing an immediate copy (R) to its local cluster and a deferred copy to each of the remote clusters. You must make these updates of the Management Class on the TS7700 for each Cluster after it has been added to the Grid.

- Plan for the enhanced cache management settings on the TS7720. In the Storage Class construct, define the following items:
 - Volume Copy Retention Group
 - Volume Copy Retention Time

If you merge a TS7720 with data into a Grid, these volumes are automatically assigned to Volume Copy Retention Group *Prefer Keep*. To change this setting, use the Host Console Request function. See 8.3.3, “Host Console Request function” on page 544.

- During the process of adding clusters to the Grid, only one cluster is required to be offline. After you have added Cluster 1, you may use Cluster 0 for normal operation. However, no copies are written to any of the other clusters. In our sample configuration, only one copy to Cluster 1 is initiated and made subsequently when Cluster 1 comes back online, after Cluster 2 has been added. Because no CCPs have been defined for the remaining clusters, no copies are made to Cluster 2 and 3 and will have to be initiated manually, after these clusters have been joined and the four-cluster grid is complete.

The same is true after adding Cluster 2: You may then use Clusters 0 and 1 for normal operation, but no copy will be made to Cluster 3. Therefore, plan carefully regarding at which point in time during the upgrade process to restart your normal tape operation to the Grid.

Implementation steps for Cluster 1

To implement the four-cluster hybrid grid shown in Figure 6-12 on page 342, perform the following steps; steps that are performed by the IBM SSR are marked accordingly:

1. Back up settings of Cluster 0 from the management interface as described in Figure 6-13 on page 344.
2. Stop all host activity on the TS7740 Virtualization Engine (Cluster 0, the only one). Vary all logical devices offline and wait for the processing jobs to finish or cancel them.
3. IBM SSR: Vary Cluster 0 offline. Because it is a Single Cluster, Cluster 0 cannot be set to Service mode, but must be shut down.
4. Activate the IODF and the SMS definitions and issue an OAM restart (if it was not done after the SMS activation). For more information, see Chapter 5, “Software implementation” on page 263. Cluster 1 is not ready yet, and will go online in a later step.

Note: If a new IOCDS is activated before the new library is ready, the host cannot communicate with the new library yet. Expect message CBR3006I to be issued:

CBR3006I Library library-name with Library ID library-ID unknown in I/O configuration.

5. IBM SSR: Configure the local Grid network from TS7700 Virtualization Engine System Management Interface Tool (SMIT) window.
6. IBM SSR: Join Cluster 1 (the TS7720 at Site A) to Cluster 0 using SMIT panels.
Vital Product Data (VPD) will be merged, and databases will be exchanged between those two clusters in the process.
7. IBM SSR: After the merge, we have a two-cluster grid. Bring both clusters online using the SMIT panel.
8. Proceed to validate the new two-cluster grid:
 - a. On Cluster 1, restore Constructs and Fast Ready Categories obtained from Cluster 0 using the management interface and make the adjustments necessary. Or define the constructs you want to use on this Cluster.
 - b. On Cluster 0 and 1, define Copy Consistency Points in the Management Classes intended for testing. Define TS7720 Cache removal Policies in the Storage Classes intended for testing.
 - c. Vary Cluster 0 attached FICON channels online to all connected hosts.
 - d. Vary logical devices online to the hosts. Vary Cluster 0 online to the hosts.
Consider only varying online the number of drives you need for testing.
 - e. Write and read from Cluster 0, making sure all hosts and partitions have proper access to it. Use D SMS,LIB(libraryname),DETAIL commands and check for the proper relation between Composite and Distributed Libraries, as described in Example 6-1 on page 335.
 - f. Cluster 0 was validated, so it can be returned to cluster operations.
 - g. Vary new Cluster 1 defined channels online for all hosts, making sure that all paths are validated from all hosts and partition (new Hardware and HCD definition now).

- h. Vary Cluster 1 logical devices online to the hosts.
- i. Vary Cluster 1 online to the hosts. Write and read from Cluster 1, making sure all hosts and partitions have proper access to it. Use `D SMS,LIB(libraryname),DETAIL` commands and check for the proper relation between Composite and Distributed Libraries, as described in Example 6-1 on page 335.

Implementation steps for Cluster 2

Use the new Cluster 1 as a foundation to add the next cluster to this (at this point) two-cluster grid. Perform the following steps; steps that are performed by the IBM SSR are marked accordingly:

1. Verify that Cluster 2, the second TS7740 Virtualization Engine in this Grid, is ready to work and that all required definitions on the TS3500 Tape Library and the TS7740 have been made as described in Chapter 4, "Hardware implementation" on page 175. This step includes restoring the definitions and settings from the backup obtained from Cluster 0 and also any necessary modifications.
2. Vary Cluster 1 logical devices offline.
3. IBM SSR: Set Cluster 1 to Service mode (through the management interface or through SMIT window).
4. IBM SSR: Vary Cluster 1 offline (through the SMIT window).
5. IBM SSR: Make sure that the grid network is configured and operational between Cluster 2 and the rest of the Grid by using SMIT.
6. IBM SSR: Join Cluster 2 to Cluster 1 (through the Cluster 2 SMIT windows). Vital Product Data (VPD) will be merged, and databases will be exchanged between those two clusters in the process.
7. IBM SSR: If the merge succeeds, you will have a three-cluster grid. Bring both Clusters 1 and 2 online, and check the management interface for a three-cluster grid in the Summary window.
8. Proceed to validate the new three-cluster grid:
 - a. On all three clusters, update the Copy Consistency Points in the Management Classes intended for testing.
 - b. Vary Cluster 1 attached FICON channels online to all connected hosts.
 - c. Vary logical devices online to the hosts. Vary Cluster 1 online to the hosts.
 - d. Write and read from Cluster 1, making sure all hosts and partitions have proper access to it. Use `D SMS,LIB(libraryname),DETAIL` commands and check for the proper relation between Composite and Distributed Libraries, as shown in Example 6-1 on page 335
 - e. Because Cluster 1 is validated, leave it as operational in the Grid with Cluster 0.
 - f. Vary new Cluster 2 defined channels online for all hosts, making sure that all paths are validated from all hosts and partitions (new Hardware and HCD definition now).
 - g. Vary Cluster 2 logical devices online to the hosts.
Consider only varying online the number of drives you need for testing.
 - h. Vary Cluster 2 online to the hosts. Write and read from Cluster 2, making sure all hosts and partitions have proper access to it. Use `D SMS,LIB(libraryname),DETAIL` commands and check for the proper relationship between Composite and Distributed Libraries, as shown in Example 6-1 on page 335.

Implementation steps for Cluster 3

Now we can use the new Cluster 2 as the base to add Cluster 3 to the configuration. Perform these steps; steps that are performed by the IBM SSR are marked accordingly:

1. Vary Cluster 2 logical devices offline.
2. IBM SSR: Set Cluster 2 to Service mode (through the management interface or SMIT window).
3. IBM SSR: Vary Cluster 2 offline.
4. IBM SSR: Make sure that the Grid network is configured and operational between Cluster 3 and the rest of the Grid.
5. IBM SSR: Join Cluster 3 to Cluster 2. Vital Product Data (VPD) are merged, and databases are exchanged between those two clusters in the process.
6. IBM SSR: If the merge succeeds, we have the four-cluster grid. Bring both Clusters 2 and 3 online, and check the management interface to view the four-cluster grid in the Summary window.
7. Proceed with the validation of the four-cluster grid:
 - a. On Cluster 3, restore the constructs and Fast-Ready categories from the backups obtained earlier from Cluster 0.
 - b. Update CCPs in the Management Class definitions of all clusters. Define the TS7720 cache management policies on Cluster 3.
 - c. Vary Cluster 2 attached FICON channels online to all connected hosts.
 - d. Vary logical devices online to the hosts. Vary Cluster 2 online to the hosts. Write and read from Cluster 2, making sure all hosts and partitions have proper access to it. Use D SMS,LIB(libraryname),DETAIL commands and check for the proper relationship between Composite and Distributed Libraries, as shown in Example 6-1 on page 335
 - e. Cluster 2 is validated, so it can be left operational in the Grid with Clusters 0 and 1.
 - f. Vary new Cluster 3 defined channels online for all hosts, making sure that all the paths are validated from all the hosts and partitions (new Hardware and HCD definition now).
 - g. Vary Cluster 3 logical devices online to the hosts.
8. Vary Cluster 3 online to the hosts. Write and read from Cluster 3, making sure all hosts and partitions have proper access to it. Use D SMS,LIB(libraryname),DETAIL commands and check for the proper relationship between Composite and Distributed Libraries, as shown in Example 6-1 on page 335.

Completing the implementation

Additionally, you can group clusters in families. In our hypothetical Configuration 1, we can make TS7720 Virtualization Engine Clusters 1 and 3 (Production Sites A and B) one production family.

Similarly, we can group TS7740 Virtualization Engines Clusters 0 and 2 (Backup Sites X and Y) as another family, for example, a DR family.

See 4.3.12, “Inserting logical volumes” on page 240 for more details about defining cluster families.



Migration scenarios

In this chapter, we explain scenarios for migrating to a TS7700 Virtualization Engine environment from VTS or another tape drive system. We consider different options, depending on where you start.

We provide guidance to help you achieve the migration scenario that best fits your needs. For this reason, we cover additional topics that are related to methods, tools, and software products that can help you with the migration in an easier way.

7.1 Migration to TS7700 Virtualization Engine

Various scenarios exist for migrating to a TS7740 Virtualization Engine. Depending on the source configuration, which can be for an IBM VTS, native tape drives (not IBM), or some other source, one of these scenarios applies. In this chapter, we describe the possible scenarios to migrate from a non-TS7700 technology to a TS7700 Virtualization Engine configuration.

Migrations to the TS7720 Virtualization Engine can be done only by using the host. The TS7720 Virtualization Engine does not have any attached back-end tape drives, so data must be copied into the TS7720 Virtualization Engine by using host programs.

Migration of VTS Model B10 or B20 hardware to a TS7740 Virtualization Engine, also called outboard VTS migration, is possible depending on the target configuration. It provides an upgrade path for existing B10 or B20 VTS models to a TS7740 Virtualization Engine as long as the VTS system contains only 3592-formatted data. IBM provides service offerings to support you in the migration to TS7700 Virtualization Engine.

When migrating data from a VTS to a TS7740 Virtualization Engine, and if the TS7740 Virtualization Engine is connected to the existing library and the physical cartridges are not moved, this process is called *merge*. If the data migration is performed from VTS to the TS7740 Virtualization Engine and the TS7740 Virtualization Engine is connected to another library, and the physical cartridges are moved from the existing library to the other, this process is called *move*.

Migration from VTS with 3590 Tape Drives, or native tape drives to the TS7740 Virtualization Engine, always requires host involvement to copy the data into the TS7700 Virtualization Engine. See 7.5, “Moving data in and out of the TS7700 Virtualization Engine” on page 376 for more information about the methods you can use.

Similar to previous VTS model migrations, the hardware migration scenarios have the following aspects:

- ▶ Software changes in SMS, HCD, TCDB and RMM
- ▶ Connectivity for the new configuration
- ▶ Migration of the Library Manager database from an existing tape library to the TS7700 Virtualization Engine
- ▶ Physical swap of the B10/B20 VTS to the TS7700 Virtualization Engine hardware
- ▶ Migration of the database from B10/B20 VTS to TS7700 Virtualization Engine

7.2 Migration planning from VTS

In this section, we describe the major considerations for migration from VTS Model B10 or B20 to a TS7700 Virtualization Engine and provide you with important information for the migration.

Replacing the hardware is only one of the items we consider here. The main aspects to consider are as follows:

- ▶ Planning costs and benefits

Determine the best migration method for you. The various methods have temporary infrastructure implications. You can maintain the existing channel path configuration or duplicate the channel paths during the migration to help minimize system outage during the migration. Additional channel paths means the need for additional FICON director ports, additional channel ports on the host system, the need for more addresses in the unit control block (UCB), and more FICON cables.

- ▶ Planning software

You might have to verify any software prerequisites necessary for the migration, prepare a new dynamic HCD generation and definitions, new DFSMS definitions, SMS, TCDB, RMM, and discover what is needed, depending on the scenario we are considering.

- ▶ Planning environment

One consideration is whether to use additional channel paths for disaster recovery or to minimize the system outages. Then, new cables can be laid out in the environment. Consider new matrix configurations for the FICON Directors you might plan to update or reconfigure.

- ▶ Planning activities

Plan to perform the hardware migration activities at a time when you minimize the impact on the production activities. Consider that the hardware migrations from B10/B20 VTS to TS7700 Virtualization Engine are disruptive, and therefore, with the assistance of the IBM System Service Representative, plan all the aspects of the activities for minimizing a system outage.

Note: After you start writing data into the TS7700 Virtualization Engine after the migration, the TS7700 Virtualization Engine starts normal background operations, such as reclamation and Secure Data Erase. As soon as changes are made to the logical or physical volumes, there is no way to go back to using the B10/B20 VTS.

7.2.1 VTS migration procedures and times

The VTS migration procedures involve the replacement of an existing B10 or B20 VTS with a TS7700 Virtualization Engine. This procedure is performed by IBM as a service offering. Contact your SSR for details about this offering and how to request the service.

The database on the existing VTS is moved to a new TS7700 Virtualization Engine. The replacement can be accomplished in two ways:

- ▶ The TS7700 Virtualization Engine is attached to the same tape library as the VTS that is being replaced.
- ▶ The TS7700 Virtualization Engine is attached to a different library than the VTS that is being replaced. All physical tapes that are associated with the VTS partition are also moved to the new library.

All the migration procedures are disruptive, including those that involve existing PTP VTSs. Plan for a downtime of 8 - 12.5 hours, depending upon the migration type. During this time, you do not have access to the data in the VTS that is being migrated. Plan a systems outage to perform the necessary steps, because the outboard migration is not concurrent with the system production activities.

If the VTS shares an IBM 3953 or 3494 Library Manager with another subsystem, or with native drives, those subsystems are also affected because the Library Manager must be “re-taught.” Those other subsystems are not affected for the entire outage time, but only for the amount of time needed to vary the Library Manager offline, perform the re-teaching, and return it again to the online status.

If you have more than one VTS installed, reduce the workload and, if possible, shift new workloads away from the VTS that will be migrated next, a day prior to the migration. This approach gives VTS time to copy all of the logical volumes in its Tape Volume Cache (TVC) to the physical tape. This approach might reduce the outage time needed for the migration.

To minimize the impact of the migration, consider temporarily redirecting your tape workload to a different library and VTS, TS7700 Virtualization Engine, or native drives. This can be easily done by changing the ACS routines and enables you to write tape data. However, you are not able to read the logical volumes in the VTS being migrated.

7.2.2 IBM 3494 Tape Library attachment

Note: Support for TS7740 attachment to the IBM 3494 Tape Library is provided only through RPQ. Although the TS7700 Releases 1.5, 1.6, and 1.7 still support the 3494, be sure that you consider migration to a TS3500 Tape Library when you migrate from a VTS to a TS7740.

The attachment of the TS7700 Virtualization Engine to the 3494 Tape Library involves the following connections:

- Power control

The 3494 Tape Library contains power sequencing and control equipment that remotely turns the power on or off to the TS7700 Virtualization Engine. This mode of power control for the TS7700 Virtualization Engine is normal.

- Library Manager

The 3494 Library Manager is used in a TS7740 Virtualization Engine configuration only for management of the physical volumes used by the TS7740 Virtualization Engine.

In a VTS environment, the 3494 Library Manager is used for configuration and management of virtual and stacked volumes, and management policies for the VTS.

- Tape Drives

A VTS supports 3590, 3592-J1A, TS1120 (3592-E05) and 3592-E06 Tape Drives.

The TS7700 Virtualization Engine can only attach to 3592 Tape Drives (3592-J1A, TS1120, or TS1130). Up to 12 IBM 3592 Tape Drives in a single model D22 frame in the 3494 Tape Library can be attached to the TS7700 Virtualization Engine. A maximum of two TS7700 Virtualization Engines can be connected to a single 3494 Tape Library.

Note: If you are planning to migrate a B20 or B10 VTS with 3590 Tape Drives to a TS7700 Virtualization Engine, consider copying the VTS data under host control into a TS7700 configuration as described in 7.5, “Moving data in and out of the TS7700 Virtualization Engine” on page 376. Alternatively, you must first upgrade the 3590 drives to TS1120 drives (which have been withdrawn from marketing and cannot be ordered anymore) and migrate all data to the new media. For more information about upgrading tape drives, see the following resources:

- ▶ *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring*, SG24-2229
- ▶ *IBM TotalStorage Peer-to-Peer Virtual Tape Server Planning and Implementation Guide*, SG24-6115

The following feature codes are required on the 3494 Tape Library for connection to a TS7700 Virtualization Engine:

- ▶ FC5124 Second Disk Drive for the Library Manager
- ▶ FC5246 Dual Control Path Facility
- ▶ FC5047 or FC9047 The LAN PCI Library Manager
- ▶ FC9013 on LM for TS7700 Virtualization Engine Attached
- ▶ FC9218 on TS7700 Virtualization Engine to indicate 3494 Tape Library attachment

7.2.3 Licensed Internal Code levels

To support the upgrade, the B10/B20 VTS must be at least at Licensed Internal Code level 2.32.747.xx (R7.4+7). That code level supports a method to make a copy of the VTS database that can be converted by the TS7700 Virtualization Engine into its DB2 database. Also, because the TS7700 Virtualization Engine does not support 2 Gb Fibre Channel Switches to the 3592 or to TS1120 drives, configurations that have 2 Gb switches must have the switches replaced with 4 Gb switches as part of the hardware change. We describe this in detail in 7.4.1, “Migrating a stand-alone VTS to a TS7740 single-cluster grid” on page 357 and in 7.4.3, “PTP VTS to TS7700 two-cluster grid” on page 370.

Starting with TS7700 R1.5, the 3953 Library Manager is no longer required, and the 3494 Library Manager controls only physical movement of the cartridges. Management of the logical volumes is done by the TS7700 Controller, and the TS7700 management interface serves as the primary user interface.

7.2.4 DFSMS definitions

Because a TS7700 Virtualization Engine, even as a single-cluster grid system, is defined with a Distributed and a Composite Library ID, the DFSMS library definitions must be updated as part of the migration. Retain the current library name (8-character name defined at DFSMS) as the Composite Library name and define a new name for the underlying Distributed Library.

Keeping the same library name for the Composite Library avoids having to change the library name that is associated with all of the logical volumes in the TCDB and the tape management system catalog. It might also eliminate the need to change the storage group definitions. However, you can only keep the same Composite Library name, if you are not using the old and the new configuration in parallel.

7.2.5 HCD definitions

Prepare the IODF in advance with the LIBRARY-ID and LIBPORT-IDs of the Composite Library. We provide an example of Distributed Library ID, the Logical Control Unit, and the LIBPORT-ID relationship, in Table 7-1. For more details, see 5.1.3, “Logical path considerations” on page 265.

Table 7-1 LIBPORT-ID

Distributed Library	Logical Control Unit	Libport/Subsystem IDs
Cluster 0	0-F	01-10
Cluster 1	0-F	41-50
Cluster 2	0-F	81-90
Cluster 3	0-F	C1-D0

You might need additional FICON director ports unless the B10/B20 ports can be reused. Because the TS7700 Virtualization Engine supports only FICON attachment, reuse of existing paths is not possible for ESCON-attached B10/B20 VTSs.

7.3 Hardware processes

This section covers general considerations regarding the hardware processes involved in the hardware migration from the B10/B20 VTS to the TS7740 Virtualization Engine.

7.3.1 Data format conversion

The physical data formats that are used by the B10/B20 VTS and the TS7740 Virtualization Engine differ. The TS7740 Virtualization Engine format was designed to allow the physical tapes to be directly exportable in the future. The TS7740 Virtualization Engine can read the old VTS format, but writes data only in the new format.

The conversion of the physical data format is handled completely transparently to the customer's applications. Data written in the old data format is automatically converted to the new data format when the data is accessed. The access can occur because of a host recall of the data or as part of reclamation processing for a physical volume. Other than to be ready for the export functions, there is no need to convert the data, because the data is completely accessible in the old format. Figure 7-1 on page 355 illustrates this concept.

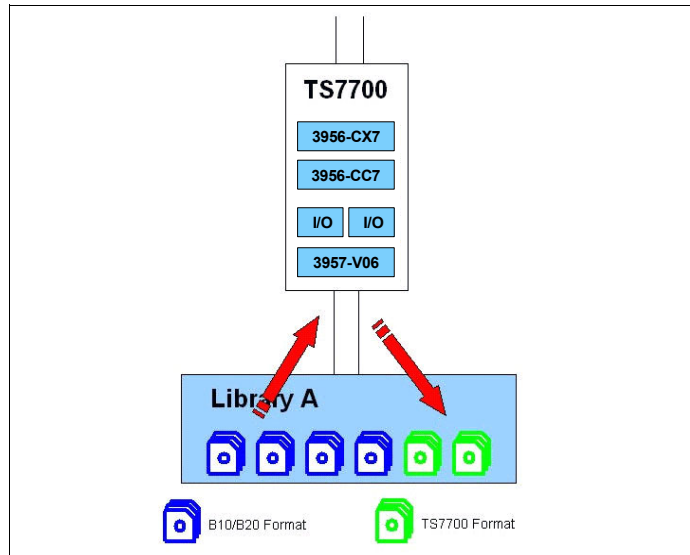


Figure 7-1 Conversion of the physical data format

If you want, you can expedite the conversion by changing the reclamation criteria for the pool that the old format volumes are in and defining a new pool for the new format data (including new writes from the host). The target pool for the old data would be the new pool. Consider adding extra physical scratch volumes prior to the migration to be sure that there will be enough scratch media on the new TS7700 Virtualization Engine configuration.

7.3.2 Back up VTS data

This procedure is used to extract data from the existing B10/B20 VTS to be exported on the TS7700 Virtualization Engine, using outboard migration. This procedure is a task that is initiated by the SSR. The process we show creates a number of files that can be used by the TS7700 Virtualization Engine to build its DB2 database and to determine the configuration of the new machine.

The backup process does the following tasks (see Figure 7-2 on page 356):

- ▶ Migrates all data from B10/B20 VTS cache to tape.
- ▶ Reconciles the B10/B20 VTS database.
- ▶ Extracts data from the B10/B20 VTS database to comma-separated files.
- ▶ Gathers logical volume data from the previously active Library Manager.




Figure 7-2 B10/B20 VTS data backup

7.3.3 Restore VTS data

This procedure exports data into the TS7740 Virtualization Engine, using outboard migration. This task is initiated by an SSR. The process we show restores the data from the previously created backup tapes to the TS7740 Virtualization Engine. The VTS files are post-processed, and a series of DB2 import files are created.

The restore process does the following tasks (see Figure 7-3):

- ▶ Converts the B10/B20 VTS data into importable DB2 files.
- ▶ Creates an empty database if one does not exist.
- ▶ Cleans up existing databases if they are not needed.
- ▶ Merges the data with the existing data if needed.




Figure 7-3 TS7740 Virtualization Engine data restore

7.4 TS7700 Virtualization Engine hardware migration scenarios

In this section, we describe in detail the hardware migration from a B10/B20 VTS or existing TS7700 Virtualization Engine to a new TS7700 Virtualization Engine. We provide various migration scenarios, depending on the existing VTS configuration and the final TS7700 Virtualization Engine configuration, which might be a single cluster or a multi-cluster grid, but we cannot cover every single configuration. If your migration scenario is not described, combining information from the listed scenarios and derive information from them for your own scenario is possible.

Table 7-2 shows the matrix of migration scenarios.

Table 7-2 Migration scenarios to TS7700 Virtualization Engine

Source configuration	Target configuration
One or more Stand-alone VTS B10 or B20	One TS7700 Virtualization Engine single-cluster grid
One Peer-to-Peer VTS B10 or B20	One TS7700 Virtualization Engine two-cluster grid

Note: Migrations from a VTS B10 or B20 to a TS7700 Virtualization Engine and migrations from a TS7700 Virtualization Engine to another TS7700 Virtualization Engine through an outboard migration process are available for TS7740 Virtualization Engine configurations only. You must request this service from IBM if necessary. Starting with R1.6, you can intermix TS7720 Virtualization Engine clusters and TS7740 Virtualization Engine clusters within the same Grid configuration.

7.4.1 Migrating a stand-alone VTS to a TS7740 single-cluster grid

In this section, we describe the data and hardware migration from a stand-alone VTS B10 or B20 to a TS7740 Virtualization Engine connected to a 3494 Tape Library or to a TS3500 Tape Library.

Note: Support for TS7740 attachment to the IBM 3494 Tape Library is provided only through RPQ. Although the TS7700 Releases 1.5, 1.6 and 1.7 still support the 3494, be sure you consider migration to a TS3500 Tape Library when you migrate from a VTS to a TS7740.

We describe the following two options:

- Where the TS7740 Virtualization Engine connects to the same tape library as the VTS (see “Option 1: Existing tape library” on page 357).
- Where the TS7700 Virtualization Engine connects to a different library (see “Option 2: New target tape library” on page 361).

Option 1: Existing tape library

If you are migrating the data from a stand-alone B10/B20 VTS to a TS7740 Virtualization Engine single-cluster grid, where the TS7740 Virtualization Engine will be connected to the same tape library as the VTS, use this procedure. Figure 7-4 on page 358 shows an overview of the migration process.




Figure 7-4 Stand-alone VTS to a TS7740 single-cluster grid in the same tape library

Note: This migration does not change the tape drive configuration. Therefore, in a 3494 Tape Library, the VTS to be replaced must already have 3592 Tape Drives installed, and all data must have already been migrated to 3592 cartridges.

Support for TS7740 attachment to the IBM 3494 Tape Library is provided only through RPQ.

The migration steps are as follows:

Note: Your SSR will perform steps 1, and 6 - 10 of the migration steps listed. We show them for informational purposes only.

1. IBM SSR: Start the installation of the TS7740 Virtualization Engine hardware a few days prior to the outage window.
2. Validate whether any conflicting information exists in the TCDB, the tape management system catalog, and the Library Manager database. This step is needed only if you plan to change the library name, as described in step 4. Appendix F, "Sample JCLs" on page 889 offers a sample job for using the RMM utility EDGUTIL for verification.
3. Stop all host activity:
 - a. If another VTS or TS7740 Virtualization Engine system is available to the host during the migration, you can change the ACS routines to direct the allocations to that other system.
 - b. Complete or cancel all host jobs for the VTS.
 - c. Vary offline all device addresses associated with the VTS for all attached hosts.
 - d. Vary the VTS to be migrated offline to all hosts.
 - e. Vary the channel paths to the VTS offline.
4. Prepare the software changes. (All of the following steps can be done concurrently with the hardware, starting with step 6 on page 360.)
 - a. Define the library names to SMS.

When you define the single-cluster grid TS7740 Virtualization Engine, you must define a Composite Library and one Distributed Library in SMS. Reuse the existing library name of the VTS as the Composite Library name, as summarized in the top half of Figure 7-5 on page 359. Reusing the library name has the advantage that you do not

have to update the library name in the volume records in the TCDB and the tape management systems catalog. In this case, you simply update the existing library definition `MyLib` and define the new distributed library `DistLib`.

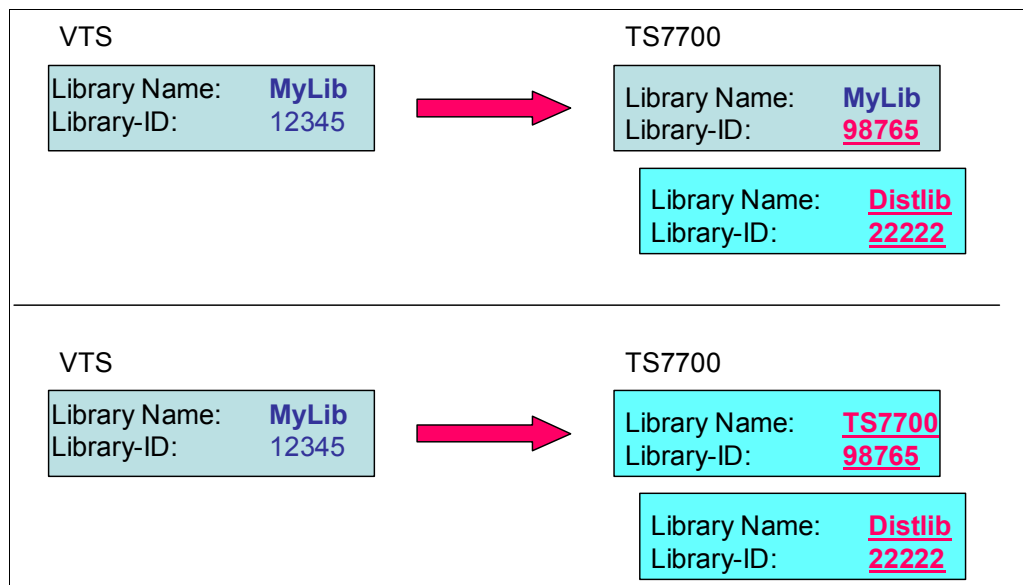


Figure 7-5 Defining library names

If you do not keep existing library names, as shown in the bottom half of Figure 7-5, you have to change the library names in all volume entries in the TCDB and your tape management system catalog to the new names before you can use the volumes. Delete the old library definition `MyLib` and define the two new libraries `TS7700` and `DistLib` in SMS through ISMF.

Remember to write the new Library-IDs in the definitions. If a new Composite Library name is used, then update existing Storage Group definitions to relate to that new library name. When you delete a library in the SCDS through ISMF, the change must be reflected in the TCDB.

b. Make changes to HCD channel definitions.

You must plan appropriately if you want to reuse existing host channels and FICON adaptor connections, or plan to define new channels, or maybe even switch from ESCON to FICON with this process. If you define the devices as offline in HCD and use any product for device sharing, you need to define the new addresses to that product also.

c. Make changes to LCU definitions.

Changes are needed because you change the Library-ID. If the VTS is a model with 64 or 128 logical units, you must also define more LCUs to enable the use of 256 logical units supported on a single-cluster grid. See Chapter 5, "Software implementation" on page 263 for more information.

d. Set missing Interrupt Handler (MIH) values.

If you are defining specific address ranges in the `IECIOSxx` member in `SYS1.PARMLIB`, make sure that MIH values for the new devices are set. Proper values are described in 5.2.6, "Set values for the Missing Interrupt Handler" on page 284.

e. Complete the software definitions.

If you have redirected workload to other systems during the migration, you might need to change the ACS routines back. If you are planning to use new SMS constructs and

implement functions provided through outboard policy management, define them on the host now.

Activate the IODF and SMS definitions and issue an OAM restart (if it was not done after the SMS activation). For more information, see Chapter 5, “Software implementation” on page 263.

Note: If the new SCDS is activated ahead of the new library being ready, the host cannot communicate with the new library yet. Expect message CBR3006I to be issued:

```
CBR3006I Library library-name with Library ID library-ID unknown in I/O configuration.
```

5. Change volume entries in the TCDB.

If the Composite Library name remains the same as the previous name, no changes are required and you can continue with Step 6.

If you change the Library name, you must change all volume entries in the TCDB to relate to the new name for every single volume.

Updating all of the volume entries can be a time-consuming process. To reduce the update time, you can change all existing scratch tapes first to prepare for non-specific mounts (write) from production as fast as possible and, after that, take all the rest. However, do not start using the TS7740 Virtualization Engine before all the TCDB volume records have been updated with the new library name.

If you are using RMM as your tape management system, a command to RMM for each volume is also needed.

6. IBM SSR: Drain the VTS Tape Volume Cache.

7. IBM SSR: Extract the VTS database, as explained in “Back up VTS data” on page 355.

8. IBM SSR: Disconnect the VTS from its Tape Library and from its TS1120 or 3592 Tape Drives.

9. IBM SSR: If the VTS connects to the existing 3953 or 3494 Library Manager using 2 Gb Fibre Channel Switches, they must be replaced with the 4 Gb switches at this time. If the TS7740 Virtualization Engine has R1.5 and the VTS is attached to a 3953, the switches can be moved to the 3584. If the switches were 2 Gb, new 4 Gb switches can be installed in the 3584 frame.

10. IBM SSR: Complete the installation of the TS7740 Virtualization Engine:

- a. Connect the TS7700 Virtualization Engine to the TS1120 or 3592 drives.
- b. If the VTS was the only subsystem configured in this 3953 Library Manager, we can use the same Logical Library already defined in TS3500 Tape Library for the new TS7740 Virtualization Engine. In this case, the 3953-F05 frame can be discontinued, if Fibre Channel Switches have been moved to the drive frame in TS3500 Tape Library and a new TSSC has been ordered for TS7740 Virtualization Engine.

If other subsystems were configured in this 3953 Library Manager, they continue to use the same Logical Library Partition as before. Therefore, a new Logical Library must be created to accommodate the TS7740 Virtualization Engine. Also, the physical volumes belonging to the old VTS must be reassigned to the new Logical Library, and tape drives formerly used by VTS. The ELC Library Manager configuration must be updated. Also, all physical and logical volumes belonging to that VTS must be removed from the ELC Library Manager database.

- c. Restore the VTS database to the TS7740 Virtualization Engine, as explained in “Restore VTS data” on page 356.
11. Vary the TS7740 Virtualization Engine online from the TS7700 Virtualization Engine System Management Interface Tool (SMIT) screen by selecting **3957-V06/VEA Online/Offline Menus → Vary 3957-V06/VEA Online**.
 12. Now you are ready to test the new single-cluster grid:
 - a. Vary the defined channels online to the host.
 - b. Vary the logical devices online to the host.
 - c. Vary the Composite and Distributed Libraries online to the host.
 - d. Run test jobs to read and write from the TS7700 Virtualization Engine single-cluster grid.

Option 2: New target tape library

Migration from a stand-alone B10/B20 VTS and tape library to a new TS7700 Virtualization Engine installed in a separate IBM tape library includes moving the cartridges from the source library to the target library, as illustrated in Figure 7-6.




Figure 7-6 Stand-alone VTS to TS7700 Virtualization Engine in a different tape library with cartridge move

In one option, the TS7740 Virtualization Engine will be connected to the same TS3500 Tape Library where the VTS is attached but in a separate logical library; for this reason, it is also considered in this section.

If this migration scenario occurs within the same physical TS3500 Tape Library, where one B10/B20 VTS is attached to one logical library with its own Library Manager, and the TS7700 Virtualization Engine is attached to another logical library, then no physical stacked cartridges that belong to the B10/B20 VTS will be moved, but the logical associations of the physical volumes from the VTS to the TS7700 Virtualization Engine will be transferred. You have to reassign the physical volume serial numbers that are dedicated to the B10/B20 VTS logical library to the TS7700 Virtualization Engine logical library by using the TS3500 Tape Library Specialist. For details, see *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789.

The time that is required for the physical move of the cartridges from one library to another might vary depending on the number of cartridges and the distance they are transported. Be sure to plan extra outage time for this activity.

If you are installing new drives together with the TS7740 Virtualization Engine in the new library, the TS7740 Virtualization Engine can be connected and tested in advance to minimize

the system outage window needed for the migration. If you are reusing the existing VTS-attached drives in the new library, plan for an extra outage to remove the drives from the existing library and install them in the new library. For this scenario, we assume that you are not moving the existing VTS drives.

Perform the following steps to accomplish this migration scenario:

Note: The IBM SSR will perform steps 1, 6, 7, 8, 10, 11, 13, and 14 as part of the installation. We list them for informational purposes only.

1. IBM SSR: Back up the Library Manager Administrative Data of the library constructs from the source library. You may perform this activity prior to the outage window.
2. Determine whether there is any conflicting information in the TCDB, the tape management system catalog, and the Library Manager database. This step is needed only if you plan to change the library name, as described in step 4a on page 358.
3. Stop all host activity.
 - a. If another VTS or TS7700 Virtualization Engine system is available to the host during the migration, you can change the ACS constructs to direct allocation to that system.
 - b. Complete or cancel all host jobs for the VTS.
 - c. Vary off all device addresses associated with the library for all attached hosts.
 - d. Vary the existing VTS offline to all hosts.
 - e. Vary the existing channels offline.
4. Prepare the software changes. All of the following steps can be done concurrent with the hardware changes that follow:
 - a. Defining the library names to SMS.

When you define the single-cluster grid TS7740 Virtualization Engine, you must define a Composite Library and one Distributed Library in SMS. We recommend that you reuse the existing library name of the VTS as the Composite Library name, as shown in Figure 7-5 on page 359.
 - b. Make changes to HCD channel definitions.

If you plan to reuse existing host channels and FICON adaptor connections, or plan to define new channels, or maybe even switch from ESCON to FICON with this process, complete these planning considerations at this time. You could define new FICON channels to ease the process, for example.
 - c. Make changes to HCD LCU definitions.

Changes are needed because you change the Library-ID. If the VTS is a model with 64 or 128 logical units, you also need to define more LCUs, to enable the use of 256 logical units supported on a single-cluster grid. See Chapter 5, “Software implementation” on page 263 for more information.

If you define the devices as offline in HCD and use any product for device sharing, you need to reflect the new addresses to that product.
 - d. Set the missing interrupt values.

If you are defining specific address ranges in the IECIOSxx member in SYS1.PARMLIB, make sure that MIH values for the new devices are set. Proper values are described in 5.2.6, “Set values for the Missing Interrupt Handler” on page 284.

- e. Complete the software definitions.

If you have redirected workload to other systems during the migration, you might need to change the ACS routines back. If you are planning to use new SMS constructs and implement functions provided through outboard policy management, define them on the host now.

Activate the IODF and the SMS definitions and issue an OAM restart (if it was not done after the SMS activation). For more information, see Chapter 5, “Software implementation” on page 263.

Note: If the new SCDS is activated before the new library is ready, the host cannot communicate with the new library yet. Expect message CBR3006I to be issued:

CBR3006I Library library-name with Library ID library-ID unknown in I/O configuration.

5. Change volume entries in the TCDB.

If the Composite Library name is the same as the previous name, no changes are required, and you can continue with the next step. If you had to change the library name, see Step 5 on page 360 for more information about how to update the TCDB volume entries.

6. IBM SSR: Drain the VTS Tape Volume Cache.
7. IBM SSR: Extract the VTS database (see “Back up VTS data” on page 355 for details).
8. IBM SSR: Back up 3494 or 3953 Library Manager data, including the Volume Serial range for the media types.
9. Remove the physical cartridges that belong to the VTS being migrated from the source library.
10. IBM SSR: Restore the VTS database in the TS7700 Virtualization Engine, as explained in “Restore VTS data” on page 356.
11. IBM SSR: Restore 3494 or 3953 Library Manager data. This includes the Volume Serial range for the media types.
12. Insert the physical cartridges previously removed from the source library into the target library.
13. IBM SSR: Perform an inventory of the library to allow the TS3500 and the TS7740 to recognize the physically inserted cartridge.
14. IBM SSR: Vary TS7740 Virtualization Engine online from the TS7700 Virtualization Engine System Management Interface Tool (SMIT) screen by selecting **3957-V06/VEA Online/Offline Menus → Vary 3957-V06/VEA Online**.
15. Now you are ready to test the new single-cluster grid:
 - a. Vary the defined channels online.
 - b. Vary the logical devices online.
 - c. Vary the Composite and Distributed Libraries online to the host.
 - d. Run test jobs to read and write from the TS7740 Virtualization Engine single-cluster grid.

7.4.2 Merge multiple VTSs to a TS7700 single-cluster grid

In this section, we describe the migration of two stand-alone VTSs to a TS7700 Virtualization Engine single-cluster grid.

You can migrate stand-alone VTSs systems to a single-cluster grid TS7700 Virtualization Engine without the need to use host utilities to copy the data.

As with the previous migration scenario, two options are available:

- The TS7740 Virtualization Engine can be attached to the same 3494 or TS3500 Tape Library as the VTS systems (see “Option 1: Existing tape library” on page 364).

Note: Support for TS7740 attachment to the IBM 3494 Tape Library is only provided through RPQ. Although the TS7700 Releases 1.5, 1.6 and 1.7 still support the 3494, we recommend that you consider migration to a TS3500 Tape Library when you migrate from a VTS to a TS77406.

- The TS7740 Virtualization Engine can be attached to a different TS3500 Tape Library (see “Option 2: New target tape library” on page 367).

If two stand-alone VTS systems are migrated to a single-cluster grid TS7740 Virtualization Engine, the assumption is that the physical volumes for one VTS system will remain in the current tape library to be used by the TS7700 Virtualization Engine, and that physical volumes for the other VTS system will reside in a separate tape library. The physical volumes belonging to the other VTS will be moved to the library where the TS7700 Virtualization Engine is attached. If both VTS systems share the same Library Manager, physical transfer of physical volumes does not occur. Instead, all the associations of the partitions to which the physical volumes belong are transferred.

Option 1: Existing tape library

Figure 7-7 shows an overview of the process.

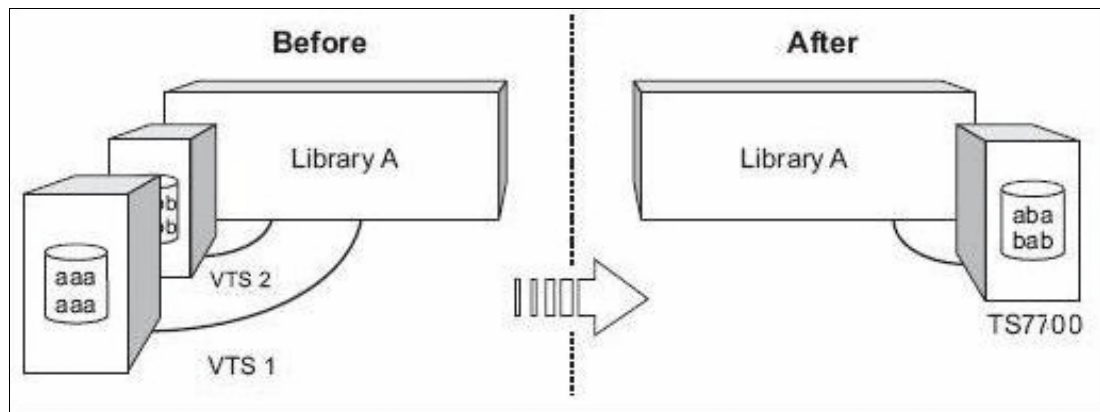


Figure 7-7 Two stand-alone VTS to a single-cluster grid TS7700 Virtualization Engine in the same tape library

Perform the following steps:

Note: Your SSR performs steps 1, and 7 - 13 as part of the installation. We list them for informational purposes only.

1. IBM SSR: Start the installation of the TS7740 Virtualization Engine hardware a few days prior to the outage window. Installation cannot be completed if TS7740 Virtualization Engine does not have at least four physical tape drives. In advance, consider whether or not TS7740 Virtualization Engine will use the same Logical Library definition that is in use by the 3953 Library Manager (and VTSSs).
2. Logical volume ranges in a TS7740 Virtualization Engine must be unique. Verify that the logical volume ranges for the two VTSSs are unique before you start the merge procedure.
3. Validate that no conflicting information exists between the TCDB, RMM, and the Library Manager. This step is important, because in this migration scenario, the library names of volumes belonging to one of the two existing library names need to be changed. Only one of the existing library names can continue to exist after you merge two VTSSs to one TS7740 Virtualization Engine.
4. Stop all host activity:
 - a. If there is another VTS or TS7700 Virtualization Engine system available to the host during the migration, you can change the ACS constructs to direct allocation to that system.
 - b. Complete or cancel all host jobs for the two VTSSs.
 - c. Vary off all device addresses associated with the library for all attached hosts.
 - d. Vary the existing VTSSs offline to all hosts.
 - e. Vary the existing channels offline.
5. Prepare the software changes. Make sure that you apply the changes to all connected hosts. All of the following steps can be done concurrently with the hardware changes that follow:
 - a. Changes to SMS
 - When you define a single-cluster grid, you need to define a Composite Library and one Distributed Library. Reuse the existing library name of one of the two existing VTSSs as Composite Library name for the new TS7700 Virtualization Engine. For the VTS that is removed, you need to change all volume entries in the TCDB and in RMM to the remaining Composite Library name before you can use the volumes. If the Composite Library name is not reused, you must change the information for all volumes from both VTSSs.
 - Delete both of the old libraries and define the two new libraries in SMS through ISMF. Remember to write the new Library-IDs in the definitions. You must also remember to then relate all existing Storage Group definitions from the library name that is not reused to that remaining library name.

Note: If you are reusing an existing library name, you only need to update its Library-ID in the library definition and do not need to delete it.

- b. Changes to HCD channel definitions

If you plan to reuse existing host channels and FICON adaptor connections, or plan to define new channels, or maybe even switch from ESCON to FICON with this process,

those are planning considerations that should be completed at this time. You could define new FICON channels to ease the process, for example.

c. Changes to LCU definitions

Changes are needed because you change the Library-ID. If one or both VTS models have 64 logical units, you must also define more LCUs to enable the use of 256 logical units supported on a single-cluster grid. See Chapter 5, “Software implementation” on page 263 for more information.

d. Device sharing

If you define the devices as offline in HCD and use any product for device sharing, you must reflect the new addresses to that product.

e. Missing interrupt values

If you are defining specific address ranges in the IECIOSxx member in SYS1.PARMLIB, make sure that MIH values for the new devices are set. Proper values are described in 5.2.6, “Set values for the Missing Interrupt Handler” on page 284.

Activate the IODF and SMS definitions and issue an OAM restart (if it was not done after the SMS activation). For more information, see Chapter 5, “Software implementation” on page 263.

Note: If the new SCDS is activated before the new library is ready, the host cannot communicate with the new library yet. Expect message CBR3006I to be issued:

CBR3006I Library library-name with Library ID library-ID unknown in I/O configuration.

6. Make changes to volume entries in the TCDB.

The Composite Library name cannot remain the same for one of the old Library names. You must change all volume entries within the TCDB to relate to the new name by issuing a command to every single volume entry.

An update of all volume entries can be a time-consuming process. Because of the time used, you could change all existing scratch tapes first to prepare for non-specific mounts (write) from production as fast as possible and then take all the rest.

7. IBM SSR: Drain the Tape Volume Cache of the first VTS and the second VTS partitions.

8. IBM SSR: Extract the first VTS and the second VTS partitions database; see “Back up VTS data” on page 355 for details.

9. IBM SSR: Disconnect the first VTS and the second VTS hardware from the Tape Library and the TS1120 or 3592 Tape Drives.

10. IBM SSR: If the first VTS and the second VTS partitions connect to the existing 3953 or 3494 Tape Libraries using 2 GB Fibre Channel Switches, they must be replaced with the 4 GB switches at this time.

11. IBM SSR: Complete the installation of the TS7740 Virtualization Engine single-cluster grid:

a. Connect the TS7700 Virtualization Engine to the TS1130 or 3592 drives.

b. If J70 or C06 were part of the configuration, you must reconfigure 3953 Library Manager by taking both VTSs out of the configuration and deleting all logical and physical volumes belonging to them from the 3953 Library Manager. Also, the Library Manager needs a Logical Library within TS3500 Tape Library. You can consider moving this native partition to a new Logical Library (leaving the old Logical Library to be used by the new TS7740 Virtualization Engine) or keeping the 3953 Library

Manager in the same Logical Library, depending on how difficult it would be to reassign the physical cartridges.

- c. Restore the first VTS and the second VTS partitions database into the TS7740 Virtualization Engine; see “Restore VTS data” on page 356 for details.

12. IBM SSR: Vary TS7740 Virtualization Engine online from the TS7700 Virtualization Engine System Management Interface Tool (SMIT) screen by selecting **3957-V06/VEA Online/Offline Menus** → **Vary 3957-V06/VEA Online**.

13. IBM SSR: Now the SSR is ready to test the new single-cluster grid:

- a. Vary the defined channels online.
- b. Vary the logical devices online.
- c. Vary the Composite and Distributed Libraries online to the host.
- d. Run test jobs to read and write from the TS7740 Virtualization Engine single-cluster grid.

Note: This procedure applies also for migration from more than two VTS into a TS7740 Virtualization Engine single-cluster grid; the limiting factor is the amount of virtual devices and volumes that can be managed for the TS7740 Virtualization Engine. The TS7740 Virtualization Engine can also reuse the Library name of only one of the VTSs, so you need to change the volume entries in TCDB for the other VTSs.

Even if this procedure is valid, if possible, migrate one VTS at a time.

Option 2: New target tape library

If you are merging the data existing in two stand-alone B10/B20 VTSs into a single-cluster grid TS7700 Virtualization Engine, and the merged TS7700 Virtualization Engine will be attached to a new target Tape Library, and physical cartridges will be moved to the target library (*move*), use the procedure we describe here. Figure 7-8 shows an overview for this migration.




Figure 7-8 Two stand-alone VTSs to a single-cluster grid with different tape library cartridges move

When physically moving cartridges from one library to another, the time required for this task might vary, depending on the number of cartridges and the distance they are transported.

The TS7700 Virtualization Engine can be connected to the new Tape Library and drives. This step minimizes the system outage window needed for the migration.

Perform the following steps to accomplish the migration scenario described previously:

Note: Your SSR performs steps 1, 6, 7, 8, 10, and 12 as part of the installation. We list them for informational purposes only.

1. IBM SSR: Back up the Library Manager administrative data, including the library constructs from the source library. You may perform this activity prior to the outage window.
2. Determine whether any conflicting information exists between the TCDB, RMM, and the Library Manager. This step is important, because in this migration scenario, volumes belonging to one of the two existing library names must be handled. Only one of the existing library names can continue to exist after you merge two VTSSs into one TS7740 Virtualization Engine.
3. Stop host activity:
 - a. If another VTS or TS7700 Virtualization Engine system is available to the host during the migration, you can change the ACS constructs to direct allocation to that system.
 - b. Complete or cancel all host jobs for the two VTSSs.
 - c. Vary off all device addresses associated with the library for all attached hosts.
 - d. Vary the existing VTSSs offline to all hosts.
 - e. Vary the existing channels offline.
4. Prepare the software changes. All the following steps could be done concurrently with all the hardware changes that follow:
 - a. Changes to SMS

When you define a single-cluster grid, you must define a Composite Library and one Distributed Library. Reuse the existing library name of one of the two existing VTSSs as the Composite Library name for the new TS7740 Virtualization Engine. For the VTS that is removed, you must change all volume entries in the TCDB and in RMM to the remaining Composite Library name before you can use the volumes. If the Composite Library name is not reused, you must change information for all volumes from both VTSSs.

Delete both old libraries and define the two new libraries in SMS through ISMF. Remember to write the new Library-IDs in the definitions. You must also remember to then relate all existing Storage Group definitions from the library name that is not reused to that remaining library name. When you delete a library in the SCDS through ISMF, the change needs to be reflected in the TCDB.
 - b. Changes to HCD channel definitions

If you plan to reuse existing host channels and FICON adaptor connections, or plan to define new channels, or maybe even switch from ESCON to FICON with this process, these are planning considerations that should be completed at this time. You may define new FICON channels to ease the process, for example.
 - c. Changes to HCD LCU definitions

Changes are needed because you change the Library-ID. If one or both VTS models have 64 logical units, you also must define more LCUs to enable the use of

Note: If you are reusing one of the existing library names, you only need to update its Library-ID in the library definition and do not need to delete it.

256 logical units supported on a single-cluster grid. See Chapter 5, “Software implementation” on page 263 for more information.

d. New addresses

If you define the devices as offline in HCD and use any product for device sharing, you must reflect the new addresses to that product.

e. Missing interrupt values

If you are defining specific address ranges in the IECIOSxx member in SYS1.PARMLIB, be sure that MIH values for the new devices are set. Proper values are described in 5.2.6, “Set values for the Missing Interrupt Handler” on page 284.

f. IODF and SMS definitions

Activate the IODF and SMS definitions and issue an OAM restart (if it was not done after the SMS activation). For more information, see Chapter 5, “Software implementation” on page 263.

Note: If the new SCDS is activated before the new library is ready, the host cannot communicate with the new library yet. Expect message CBR3006I to be issued:

CBR3006I Library library-name with Library ID library-ID unknown in I/O configuration.

5. Change volume entries in TCDB.

The Composite Library name cannot remain the same for one of the old library names. You must change all volume entries within the TCDB to relate to the new name by issuing a command to every single volume entry.

Updating all the volume entries can be a time-consuming process. To reduce the update time, you could change all existing scratch tapes first to prepare for non-specific mounts (write) from production as fast as possible and then take all the rest.

If you are using RMM, an RMM command for each volume is also needed. The FORCE parameter is only used when needed and requires access to a specific Facility class in Resource Access Control Facility (RACF) named STGADMIN.EDG.FORCE. Verify that you have the required authorization.

6. IBM SSR: Drain the Tape Volume Cache of the first VTS and the second VTS partitions.

7. IBM SSR: Extract the first VTS and the second VTS partitions database; see “Back up VTS data” on page 355 for details.

8. IBM SSR: Back up 3494 or 3953 Library Manager data; this includes the volume serial range for the media types.

9. Remove the physical cartridges that belong to the first VTS and the second VTS partitions you are migrating to TS7700 Virtualization Engine from the source library.

10. IBM SSR: Restore the first VTS and the second VTS partitions database into the TS7740 Virtualization Engine; see “Restore VTS data” on page 356 for details.

11. Insert the physical cartridges into the target library that belong to the first VTS and the second VTS partitions you are migrating to TS7740 Virtualization Engine that were previously removed from the source library.

12. IBM SSR: Vary TS7700 Virtualization Engine online from the TS7700 Virtualization Engine system Management Interface Tool (SMIT) screen by selecting **3957-V06/VEA Online/Offline Menus** → **Vary 3957-V06/VEA Online**.
13. Now you are ready to test the new single-cluster grid:
 - a. Vary the defined channels online.
 - b. Vary the logical devices online.
 - c. Vary the Distributed and the Composite Library online to the hosts.
 - d. Run test jobs to read and write from the TS7700 Virtualization Engine single-cluster grid.

7.4.3 PTP VTS to TS7700 two-cluster grid

In this section, we cover the migration scenario from two VTSs that are in a Peer-to-Peer configuration to two new TS7740 Virtualization Engines in a two-cluster grid configuration.

Two options are available:

- ▶ The new TS7740 Virtualization Engine Grid will be installed in the existing tape libraries, in which case no physical cartridge movement will occur (see “Options 1: Existing source tape libraries” on page 371).
- ▶ The TS7700 Virtualization Engine Grid will be installed in new target tape libraries, in which case physical cartridges will be moved from the source library to the target library (see “Option 2: New target tape libraries” on page 373).

Figure 7-9 illustrates the data migration process from one VTS to the corresponding, empty TS7740 Virtualization Engine.




Figure 7-9 Data migration PTP VTS to a TS7700 Virtualization Engine two-cluster grid

Note: Both installation teams must be on site at the same time and be synchronized with regard to the installation instruction steps.

Options 1: Existing source tape libraries

In this section, we show the steps to perform the hardware migration from the Peer-to-Peer VTS B10 or B20 to the TS7700 Virtualization Engines in an IBM 3494 or TS3500 Tape Library. In this section, we assume that the tape libraries remain the same and you reconnect the TS7740 Virtualization Engines in the existing tape libraries where the Peer-to-Peer VTSs were previously attached (see Figure 7-10).




Figure 7-10 Migration VTS PTP to TS7740 Virtualization Engine two-cluster grid with the same tape libraries

Perform the following steps to accomplish the migration scenario:

Note: The IBM SSR performs steps 1, and 6 - 13 as part of the installation. We list them for informational purposes only.

1. IBM SSR: Prepare the TS7740 Virtualization Engines a few days prior to the outage window.
2. Determine whether conflicting information exists between the TCDB, RMM, and the Library Manager. This step is required only if you plan to change the library name, as described in Step 4 on page 358.
3. Stop host activity:
 - a. If another VTS or TS7740 Virtualization Engine system is available to the host during the migration, you can change the ACS constructs to direct allocation to that system.
 - b. Complete or cancel all host jobs for the PTP VTS.
 - c. Vary off all device addresses associated with the library for all attached hosts.
 - d. Vary the existing VTSs offline to all hosts.
 - e. Vary the existing channels offline.
4. Prepare the software changes. The following steps could be done concurrently with all the hardware changes that follow:
 - a. Changes to SMS

When you define a two-cluster grid, you need to define a Composite Library and two Distributed Libraries, which is the same as for PTP VTS. Reuse the existing library name as the Composite Library name so that you only need to write the new Library-IDs in the existing library definitions. If you do not keep the existing library name, you must delete the old names and add the new names and then relate all existing Storage Group definitions to that new name. After that task is done, you must

change all volume entries in the TCDB and in RMM to the new name before you can use the volumes.

b. Changes to HCD channel definitions

If you plan to reuse existing host channels and FICON adaptor connections, or plan to define new channels, or maybe even switch from ESCON to FICON with this process, these are planning considerations that should be completed at this time. You could define new FICON channels to ease the process, for example.

c. Changes to LCU definitions

Changes are needed because you change the Library-ID of the Composite Library. If the PTP VTS model has 64 or 128 logical units, you also need to define more LCUs to enable the use of 256 logical units supported on a single-cluster grid. See Chapter 5, “Software implementation” on page 263 for more information.

d. New addresses

If you define the devices as offline in HCD and use any product for device sharing, you need to propagate the new addresses to that product.

e. Missing interrupt values

If you are defining specific address ranges in the IECIOSxx member in SYS1.PARMLIB, make sure that MIH values for the new devices are set. Proper values are described in 5.2.6, “Set values for the Missing Interrupt Handler” on page 284.

f. IODF and SMS definitions

Activate the IODF and SMS definitions and issue an OAM restart (if it was not done after the SMS activation). For more information, see Chapter 5, “Software implementation” on page 263.

Note: If the new SCDS is activated before the new library is ready, the host cannot communicate with the new library yet. Expect message CBR3006I to be issued:
CBR3006I Library library-name with Library ID library-ID unknown in I/O configuration.

5. Change volume entries in the TCDB.

If the Composite Library is the same as the previous name, there are no changes and you can continue with Step 6 on page 360. If you change the name, be sure to change all volumes within the TCDB to propagate the new name by issuing a command to every single volume.

An update of all volume entries can be a time-consuming process. Because of the time used, you can change all existing scratch tapes, first to prepare for non-specific mounts (write) from production as fast as possible and, after that, take all the rest.

If you are using RMM, an RMM command for each volume is also needed. The FORCE parameter is only used when needed and requires access to a specific Facility class in RACF named STGADMIN.EDG.FORCE. Verify that you have the required authorization.

6. IBM SSR: Place the non-master PTP VTS in service preparation. The team on the other site's master VTS does not need to perform this function; it must wait for the non-master VTS to be in service. When the non-master VTS is in service, both teams start Step 7 in synchronization with each other.

7. IBM SSR: Drain the PTP VTSs Tape Volume Cache; this is done by both teams on the PTP sites.

8. IBM SSR: Extract the PTP VTSs database; see “Back up VTS data” on page 355 for details. This step is done by both teams on the PTP sites.
9. IBM SSR: Disconnect the VTSs from their Tape Library, either 3953/TS3500 Tape Library or 3494 Tape Library, and TS1120 or 3592 Tape Drives; both teams perform this step.
10. IBM SSR: If the PTP VTSs connect to the existing 3953 or 3494 Tape Libraries using 2 GB Fibre Channel Switches, they must be replaced with the 4 GB switches at this time; both teams must verify this action. If the TS7740 Virtualization Engine has R1.5 and the VTS is attached to a 3953, the switches can be moved to the 3584. If the switches were 2 Gb, new 4 Gb switches can be installed in the 3584 frame.
11. IBM SSR: Complete the installation of the TS7740 Virtualization Engines on both sites; this is done by both installation teams.
 - a. Connect the TS7740 Virtualization Engines to the tape libraries.
 - b. Connect the TS7740 Virtualization Engines to the TS1120s or 3592 drives.
 - c. If J70 or C06 were part of the original configuration, you must reconfigure the 3953 Library Manager by taking both VTSs out of the configuration and deleting all logical and physical volumes belonging to them from the 3953 Library Manager. Also, the Library Manager will still need a Logical Library within TS3500 Tape Library. You can consider moving this native partition to a new Logical Library (leaving the old Logical Library to be used by the new TS7740 Virtualization Engine) or keeping the 3953 Library Manager in the same Logical Library, depending on how difficult would be reassigning the physical cartridges one way or other. This has to be performed at both sites.
 - d. The TS7740 Virtualization Engine has to go online after installation completes. This action is performed at both sites.
 - e. Both installation teams establish the WAN connection between the TS7700 Virtualization Engines to form the two-cluster grid configuration and join the clusters into a Grid.
 - f. Vary both TS7740 Virtualization Engines offline.
12. IBM SSR: Restore the PTP VTSs database into the TS7740 Virtualization Engines; see “Restore VTS data” on page 356 for details. Both installation teams perform this procedure on the PTP sites.
13. IBM SSR: Vary TS7740 Virtualization Engines online from the TS7700 Virtualization Engine system Management Interface Tool (SMIT) screen by selecting **3957-V06/VEA Online/Offline Menus** → **Vary 3957-V06/VEA Online**.
14. Now you are ready to test the new two-cluster grid:
 - a. Vary the defined channels online.
 - b. Vary the logical devices online.
 - c. Vary the Distributed and Composite Libraries online to the hosts.
 - d. Run test jobs to read and write from the two-cluster grid TS7700 Virtualization Engine.

Option 2: New target tape libraries

In this section, we describe the hardware migration from the Peer-to-Peer VTS to TS7700 Virtualization Engine Grids and the movement of the physical cartridges from the existing library to the new source TS3500 Tape Libraries. The old libraries can be 3494 Tape Libraries or 3953 Library Controllers connected to TS3500 Tape Libraries. This process is described in Figure 7-11 on page 374 and in the following steps.

Note: Support for TS7740 attachment to the IBM 3494 Tape Library is provided only through RPQ. Although the TS7700 Releases 1.5, 1.6 and 1.7 still support the 3494, be sure you consider migration to a TS3500 Tape Library when you migrate from a VTS to a TS7740.

The time required for physically moving cartridges from one library to another might vary, depending on the number of cartridges and the distance that they are transported.

New TS7740 Virtualization Engines and new TS3500 Tape Libraries must be completely installed, interconnected, and tested. Also, Gigabit Ethernet must be ready and new TS7740 Virtualization Engine clusters must be joined into a two-cluster grid in preparation for the outage. You must define the Composite and Distributed Libraries and management interface in advance; they must be in place before starting the migration.

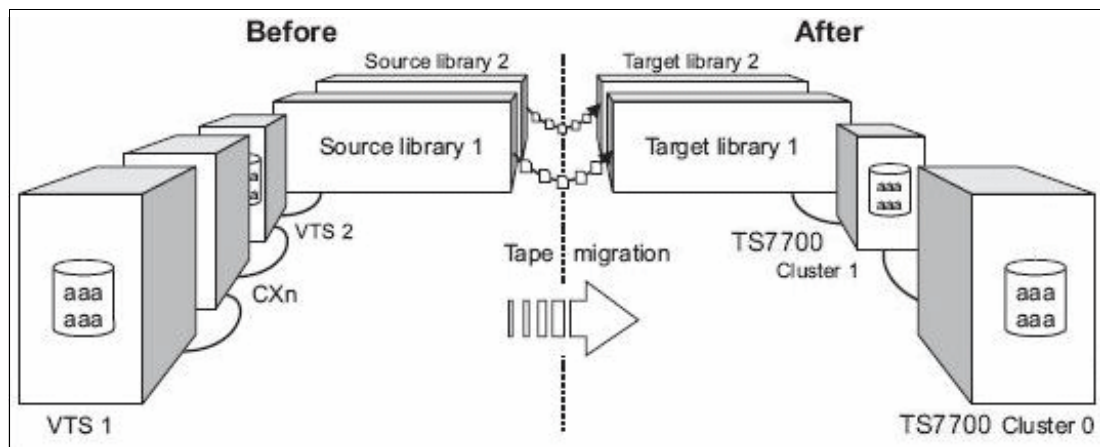


Figure 7-11 Migration of VTS PTP to TS7740 Virtualization Engine two-cluster grid with separate tape libraries

Perform the following steps to accomplish the migration scenario:

Note: The IBM SSR performs steps 2, 6 - 9, 1, 13, and 14 as part of the installation. We list them for informational purposes only.

1. Back up the PTP VTS Library Managers Administrative Data of the library constructs from the PTP VTS source libraries. You may perform this activity prior to the outage window
2. Determine if any conflicting information exists between the TCDB, RMM, and Library Manager. This step is needed only if you plan to change the library name, as described in Step 4 on page 358.
3. Stop host activity:
 - a. If another VTS or TS7740 Virtualization Engine system is available to the host during the migration, you can change the ACS constructs to direct allocation to that system.
 - b. Complete or cancel all host jobs for the PTP VTS.
 - c. Vary off all device addresses associated with the library for all attached hosts.
 - d. Vary the existing VTSs offline to all hosts.
 - e. Vary the existing channels offline.

4. Prepare the software changes. The following steps must be performed simultaneously with the hardware changes listed in step 5 on page 376:

- a. Changes to SMS

When you define a two-cluster grid, you must define a Composite Library and two Distributed Libraries, which is the same as for PTP VTS. Reuse the existing library name as the Composite Library name; then you only have to write the new Library-IDs in the existing library definitions.

If you do not keep the existing library names, you must delete the old names, add the new names, and then relate all existing Storage Group definitions to those new names. After that, you must change all volume entries in the TCDB and in RMM to the new name before you can use the volumes.

- b. Changes to HCD channel definitions

If you plan to reuse existing host channels and FICON adaptor connections, or plan to define new channels, or maybe even switch from ESCON to FICON with this process, complete these planning considerations at this time. You could define new FICON channels to ease the process, for example.

- c. Changes to LCU definitions

Changes are needed because you change the Library-ID of the Composite Library. If the PTP VTS is a model with 64 or 128 logical units, you also need to define more LCUs to enable the use of 256 logical units supported on a single-cluster grid. See Chapter 5, "Software implementation" on page 263 for more information.

- d. New addresses

If you define the devices as offline in HCD and use any product for device sharing, you need to propagate the new addresses to that product.

- e. Missing interrupt values

If you are defining specific address ranges in the IECIOSxx member in SYS1.PARMLIB, be sure that MIH values for the new devices are set. Proper values are described in 5.2.6, "Set values for the Missing Interrupt Handler" on page 284.

- f. IODF and SMS definitions and OAM restart

Activate the IODF and the SMS definitions and issue an OAM restart (if it was not done after the SMS activation). For more information, see Chapter 5, "Software implementation" on page 263.

Note: If the new SCDS is activated before the new library is ready, the host cannot communicate with the new library yet. Expect message CBR3006I to be issued:

```
CBR3006I Library library-name with Library ID library-ID unknown in I/O
configuration.
```

- g. Changes to volume entries in TCDB

If the Composite Library has the same name, no changes are necessary and you can continue with Step 6 on page 360. If you change the name, you must change all volumes within the TCDB to relate to the new name by issuing a command to every single volume.

An update of all volume entries can be a time-consuming process. Because of the time spent, you could change all existing scratch tapes first to prepare for non-specific mounts (write) from production as fast as possible, and then after that, take all the rest.

If you are using RMM, an RMM command for each volume is also needed. The FORCE parameter is only used when needed and requires access to a specific Facility class in RACF named STGADMIN.EDG.FORCE. Verify that you have the required authorization.

5. Place the non-master PTP VTS in service preparation. The team on the other site at the master VTS does not need to perform this function, but must wait for the non-master VTS to be in service. When the non-master VTS is in service, both teams start Step 7 in synchronization with each other.
6. Drain the PTP VTSs Tape Volume Cache; this step is done by both teams on the PTP sites.
7. Extract the PTP VTSs database; see “Back up VTS data” on page 355 for details. This step is done by both teams on the PTP sites.
8. Leave both VTSs offline and shut down. They will be discontinued at a later time.
9. Remove physical cartridges for those VTSs from the source libraries and move them to the target new libraries.

Note: New TS7740 Virtualization Engine clusters must be configured as a two-cluster grid prior to restoring PTP VTS data.

10. Restore the PTP VTSs database into the TS7740 Virtualization Engines; see “Restore VTS data” on page 356 for details. Both installation teams perform this procedure on the PTP sites.
11. Insert the physical cartridges from the source tape libraries into the target new libraries.
12. Inventory new libraries and determine whether volumes are correctly assigned to the proper TS7740 Virtualization Engine logical library. Cartridge Assignment Policies (CAP) should have been defined in advance. If definitions are correct, you see volumes that are assigned to the correct TS7740 Virtualization Engine Logical Partition. TS7740 Virtualization Engine will request an inventory upload startup and sort it out into its own database. This step must be performed at both sites.
13. Vary the TS7740 Virtualization Engines online from the TS7700 Virtualization Engine System Management Interface Tool (SMIT) screen by selecting **3957-V06/VEA Online/Offline Menus → Vary 3957-V06/VEA Online**.
14. Now you are ready to test the new two-cluster grid:
 - a. Vary the defined channels online.
 - b. Vary the logical devices online.
 - c. Vary the Distributed and Composite Libraries online to the hosts.
 - d. Run test jobs to read and write from the two-cluster grid TS7740 Virtualization Engine.

7.5 Moving data in and out of the TS7700 Virtualization Engine

When moving data into the TS7700 Virtualization Engine, it is not possible to simply move the cartridges out of an 3494-Bxx VTS and insert them into a Single Cluster TS7700 Virtualization Engine and copy the control data sets. This migration approach is only supported for the scenarios described in 7.4, “TS7700 Virtualization Engine hardware migration scenarios” on page 357 (for merge or move). In all other scenarios, migrating data into the TS7700 Virtualization Engine requires that the TS7700 Virtualization Engine and the

existing environment remain installed in parallel until the data has been migrated through the host attached to them.

Examples of this type of configuration are native tape drives, VTS with 3590 Tape Drives, or other vendor tape solutions. Although they can all be migrated to TS7700 Virtualization Engine, the process requires host involvement to copy the data into the TS7700 Virtualization Engine.

In this section, we discuss techniques for moving data in and out of the TS7700 Virtualization Engine. You can start using the TS7700 Virtualization Engine by moving data into it. The best method depends on the application you want to manage with the TS7700 Virtualization Engine. We describe two methods:

- ▶ **Phased method**

This method consists of using the TS7700 Virtualization Engine with new allocations. The migration of data takes longer, but it can be more controlled and flexible.

- ▶ **Quick method**

Use this method when you want to move existing data into the TS7700 Virtualization Engine. It is called quick because it swiftly puts all data you want to move under TS7700 Virtualization Engine control.

Some hints about how to move data out of the TS7700 Virtualization Engine are provided in 7.5.6, “Moving data out of the TS7700 Virtualization Engine” on page 383. However, the TS7700 Virtualization Engine is a closed-storage method, so you must be careful about selecting data to move into it. You do not want to store a large amount of data in the TS7700 Virtualization Engine that will need to be moved back out.

7.5.1 Moving data into the TS7700 Virtualization Engine: Phased method

The data movement techniques outlined here depend more on changes in parameters, routines, or procedures than on overt data movement.

Select the data

If you select DFSMSHsm-owned data, you can group your data as listed according to any or all of the items in the following list:

- ▶ Migration data (DFSMSHsm level 2)
- ▶ Backup copies (user data, CDS data, or both)
- ▶ Dump copies

You may select data based on data set name, by application, or by any other variable that you can use in the ACS routines. You can also select data based on type, such as SMF data, or DASD DUMP data.

Update the applicable parameters

If you select DFSMSHsm-owned data, review the ARCCMDxx member according to the guidelines in 7.6, “Migration of DFSMSHsm-managed data” on page 386 and update the following definitions:

- ▶ The Data Class ACS routines (if used)
- ▶ The Management Class ACS routines (if used)
- ▶ The Storage Class ACS routines (required)
- ▶ The Storage Group ACS routines (required)
- ▶ For BTLS, the unit parameter in the JCL

For DFSMSDss, update the following definitions:

- ▶ The Data Class ACS routines (if used)
- ▶ The Management Class ACS routines (if used)
- ▶ The Storage Class ACS routines (required)
- ▶ The Storage Group ACS routines (required)
- ▶ For BTLS, the unit parameter in the JCL

If you use database data, such as logs or image copy, direct new allocations into the TS7700 Virtualization Engine by updating the following definitions:

- ▶ The Data Class ACS routines (if used)
- ▶ The Management Class ACS routines (if used)
- ▶ The Storage Class ACS routines (required)
- ▶ The Storage Group ACS routines (required)

For data other than DFSMSHsm and DFSMSDss, if you are using SMS tape, update the ACS routines to include the data you want to move. You decide what you filter for and how you write the ACS routines. You can also migrate based on the UNIT parameter in the JCL to reflect the applicable unit for the TS7700 Virtualization Engine.

Update the tape management system

Although you are not overtly copying data in this option, be sure to update the tape management system catalog to reflect the changes that you expect. Check the retention rules and limits and update accordingly. If you change data set names when moving to TS7700 Virtualization Engine, you must validate changes against retention rules in your tape management system. See 7.7, “DFSMSrmm and other tape management systems” on page 395 for more information.

Watch the data move to the TS7700 Virtualization Engine

Data movement using this option does not involve overt actions, such as COPY, RECYCLE, or DUMP. When you activate the ACS routines containing the code for the TS7700 Virtualization Engine, all new data allocations for the data you have selected are written to the TS7700 Virtualization Engine. You simply verify that data is going where you expect it to go and add code to the ACS routines to manage more data as you see fit.

You can select data types that create large quantities of data, like SMF records or DASD DUMPS, and you can also select data types that create many small data sets. By observing how the TS7700 Virtualization Engine handles each type of data, you become familiar with the TS7700 Virtualization Engine, its functions, and capabilities.

7.5.2 Moving data into the TS7700 Virtualization Engine: Quick method

The steps outlined in this section involve overt actions on your part to move data into the TS7700 Virtualization Engine. As with the techniques outlined in 7.5.1, “Moving data into the TS7700 Virtualization Engine: Phased method” on page 377, you choose the data you want to move to the TS7700 Virtualization Engine.

Select the data to copy

The data you select influences all the subsequent steps in this process. If you select DFSMSHsm-owned data, the process for moving the data to the TS7700 Virtualization Engine differs from the process that you use for moving DFSMSDss data. You can select data based on the data's attributes, such as the expiration date. For example, you could select data that you keep for seven years. Probably the best method for selecting data to copy in the TS7700 Virtualization Engine is based on the data set name, application by application.

Be aware that some applications have knowledge of the VOLSER where the data is stored. There are special considerations for these applications. If you change the VOLSER that the data is on, the application has no way of knowing where the data resides. For more information about this topic, see 7.5.4, “Considerations for static VOLSERs” on page 382.

An easy method is to obtain information from the tape management system database. Some reports can give you details about the data you actually have in the tape shop, which helps you select the input volumes.

If you are using DFSMSrmm, you can easily acquire data from an RMM EXTRACT file, which is normally created as part of the regular housekeeping. Then, using a REXX EXEC or ICETOOL JCL program, you extract the information needed, such as data set name, VOLSER, and file sequence of the input volumes.

Move data to new TS7700 Virtualization Engine

Although you should use the Tape Copy Utility Tool to move data to TS7700 Virtualization Engine, not all z/OS data can be moved using these tools. In general, this tool is able to move any data on tape, except those products that own the data, for example, DFSMSshm, DFSMSoam, or IBM Tivoli Storage Manager.

When using SMS tape, the first step is to update the ACS routines to create all new data in the TS7700 Virtualization Engine. With this minimal change, tapes are created in a new TS7700 Virtualization Engine, so moving it again is not necessary.

If you move DFSMSshm-owned data, use the DFSMSshm recycle process to move the data to the TS7700 Virtualization Engine. Use a COPYDUMP job to move DFSMSdss data to the TS7700 Virtualization Engine.

The utility to use depends on the data selected. In most cases, it is sequential data that can be copied using the IEBGENER utility, DITTO/ESA. If you have DFSORT or a similar utility, ICEGENER and ICETOOL can probably give better performance.

You must use a specific utility when the input data is in a special format, for example, DFSMSdss dump data. DFSMSdss uses a 64 KB blocksize and only the proper DSS utility, such as COPYDUMP, can copy with that blocksize. Also, be careful when copying multifile and multivolume chains. You might want to separate these files, because there is no penalty when they are in a TS7700 Virtualization Engine.

In general, all other data (except data owned by an application, such as DFSMSshm) belongs to batch and backup workloads. Use EXPDT and RETPD from the DFSMSrmm Extract file to find what tapes have a distant expiration date, begin moving these tapes, and leave the short time retention tapes to the last phase of data movement (they likely will be moved by the everyday process).

Update the tape management system with the correct retention information

When the manual copy operation has been successful, update the tape management system catalog. The following data must be updated on the output volume:

- ▶ File sequence number
- ▶ Creation date and time
- ▶ Last read and last write date
- ▶ Jobname

Optionally, you can also update the following items:

- ▶ Stepname
- ▶ DDname
- ▶ Account number
- ▶ Device number

In RMM, this step can be done with a CHANGEDATASET command that has special authority to update O/C/EOV recorded fields.

To avoid this time-consuming process, use a tape copy tool, because they are ready to make all the necessary changes in a tape management system.

Update the ICF catalog with the correct output volume

The next step is to uncatalog the input data sets (if they were cataloged) and recatalog the output data sets with the new volume information.

Tape copy tools recatalog tapes during movement.

Release the input volume for SCRATCH processing

This final step must be done after you are sure the data has been correctly copied. You must also verify that the retention and catalog information is correct.

Using this quick-method sequence, you can copy every kind of tape data, including GDGs, without modifying the generation number.

In an RMM environment, you can use REXX CLIST and RMM commands, listing data from the input volumes and then using the RMM REXX variables with the CD command to update the output. After, call IDCAMS to update the ICF catalog. When the operation completes and all errors have been corrected, use the RMM DELETEVOLUME command to release the input volumes.

See *z/OS DFSMSrmm Guide and Reference*, SC26-7404 for more information about RMM commands and REXX variables. If you are using a tape management system other than RMM, refer to the appropriate product functions to obtain the same results.

Migrating data inside the TS7700 Virtualization Engine can be made easier by using products such as DFSMSHsm or IBM Tivoli Storage Manager. If you are planning to put DFSMSHsm or IBM Tivoli Storage Manager data in the TS7700 Virtualization Engine, see the following sections:

- ▶ 7.6, “Migration of DFSMSHsm-managed data” on page 386
- ▶ 7.8, “IBM Tivoli Storage Manager” on page 396

With DFSMSHsm, you can change the ARCCMDxx tape device definitions to an esoteric name with TS7700 Virtualization Engine virtual drives (in a BTLIS environment) or change SMS ACS routines to direct DFSMSHsm data in the TS7700 Virtualization Engine. The DFSMSHsm RECYCLE command can help speed the movement of the data.

A similar process can be used with IBM Tivoli Storage Manager, changing the device class definitions for the selected data to put in the TS7700 Virtualization Engine and then invoking the space reclamation process.

If you are moving DB2 data into the TS7700 Virtualization Engine, be sure that, when copying the data, the DB2 catalog is also updated with the new volume information. You may use the DB2 MERGECOPY utility to speed up processing, using TS7700 Virtualization Engine virtual volumes as output.

In general, DB2 Imagecopies and Archlog are not retained for a long time; after all new write activity goes to the TS7740 Virtualization Engine, we can expect that this data is moved by the everyday process.

7.5.3 Products to simplify the task

You might want to consider using a product designed to copy data from one medium to another. The first choice is the IBM enhancement for DFSMSrmm called Tape Copy Tool (see Table 7-3 on page 382). The Tape Copy Tool function of the internal IBM ADDONS package is designed to copy all types of MVS tape data sets from one or more volumes or volume sets to a new tape volume or tape volume set. Any tape media that is supported by DFSMSrmm is supported by this tool. The input tape media can be different than the output tape media. Do not use tool to copy tape data sets owned by Hierarchical Storage Manager (DFSMSHsm) or IBM Tivoli Storage Manager or similar products, where information of old VOLSERs is kept within the product itself and not reflected after a copy is made. This challenge typically applies to products where tapes are not cataloged in an ICF catalog.

The DFSMSrmm Tape Copy Tool cannot be used when you have a Tape Management System other than DFSMSrmm. You must choose another Tape Copy Tool from Table 7-3 on page 382.

Consider the following factors when you evaluate a tape copy product:

- ▶ Interaction with your tape management system
- ▶ Degree of automation of the process
- ▶ Speed and efficiency of the copy operation
- ▶ Flexibility in using the product for other functions, such as duplicate tape creation
- ▶ Ease of use
- ▶ Ability to create a pull list for any manual tape mounts
- ▶ Ability to handle multivolume data sets
- ▶ Ability to handle volume size changes, whether from small to large or large to small
- ▶ Ability to review the list of data sets before submission
- ▶ Audit trail of data sets already copied
- ▶ Ability to handle failures during the copy operation, such as input volume media failures
- ▶ Flexibility in filtering the data sets by wild cards or other criteria, such as expiration or creation date

Table 7-3 on page 382 lists several common tape copy products. You may choose one of these products or perhaps use your own utility for tape copy. Certainly, you do not need any of these products, but a tape copy product can make your job easier if you have many tapes to move into the TS7700 Virtualization Engine.

Table 7-3 Selection of tape copy tools

Product name	Vendor name	For more information
Tape Copy Tool/DFSMSrmm	IBM	Contact your IBM Representative for more information about this service offering.
Tape Optimizer	IBM	Go to the following address: http://www.ibm.com/software/tivoli/products/tape-optimizer-zos/
Beta55	Beta Systems Software AG	Go to the following address: http://www.betasystems.com
CA-1/TLMS Copycat	Computer Associates International, Inc.	Go to the following address: http://www.cai.com
Tape/Copy	OpenTech Systems, Inc.	Go to the following address: http://www.opentechsystems.com/tape-copy.php
TelTape	Cartagena Software Ltd.	Go to the following address: http://www.cartagena.com
Zela	Software Engineering of America	Go to the following address: http://www.seasoft.com/zela.asp

In addition to using one of these products, consider using IBM Global Services Global Technology Services (GTS) to assist you in planning and moving the data into the TS7700 Virtualization Engine. For more information about these services, see 3.9.3, “Implementation services” on page 168.

7.5.4 Considerations for static VOLSERs

Some applications have knowledge of the VOLSER of the volume where the data is stored. DFSMSHsm is one such application. When moving data for these applications, you have two choices:

- ▶ You can utilize instructions from the application author to copy the data.
- ▶ You can copy the data to a volume with the same VOLSER.

For assistance with DFSMSHsm tapes, see 7.5.2, “Moving data into the TS7700 Virtualization Engine: Quick method” on page 378.

The preferred method for moving this type of data is to use instructions from the application author. If however, you must copy the data to a volume with the same VOLSER, consider the following information:

- ▶ The source and target media might not be the exact same size.
- ▶ You cannot mount two volumes with the same VOLSER at the same time.
- ▶ If the source tape is a system-managed tape, you cannot have two volumes with the same VOLSER.

As much as possible, avoid having data-set-to-VOLSER relationships. In general, these relationships are only necessary in old backup tapes or tapes that are retained for legal reasons. Using only one scratch pool greatly simplifies tape management from all perspectives.

This method is not preferred for moving data to the TS7700 Virtualization Engine. This method applies only if you have to maintain the data-set-to-VOLSER relationship; it has limitations and weaknesses.

To move data to the TS7700 Virtualization Engine, perform the following steps:

1. Copy the non-TS7700 Virtualization Engine tape volumes to DASD or other tape volumes.
2. If the source volume is resident on a 3494 Tape Library, eject the cartridge from the 3494 Tape Library using the LIBRARY EJECT command or the ISMF EJECT line operator command from the ISMF screen.
3. Delete the ejected volume from the tape management system.
4. Define the VOLSER range, including the once-duplicated number, to the TS7700 Virtualization Engine.
5. Update ACS routines so that the data is directed to the TS7700 Virtualization Engine. For BTLS, update the UNIT parameter in the JCL.
6. Create a job to copy the data currently on DASD to the TS7700 Virtualization Engine.
7. Run the copy job.
8. Update the tape management system records and any other catalog structures.

7.5.5 Combining methods to move data into the TS7700 Virtualization Engine

You will most likely want to use a combination of the phased and quick methods for moving data into the TS7700 Virtualization Engine. One approach is to classify your data as static or dynamic.

Static data is information that will be around for a long time. This data can be moved into the TS7700 Virtualization Engine only with the quick method. You must decide how much of this data will be moved into the TS7700 Virtualization Engine. One way to make this decision is to examine expiration dates. You can then set a future time when all volumes, or a subset, are copied into the TS7700 Virtualization Engine. There might be no reason to copy volumes that are going to expire in two months. By letting these volumes go to scratch status, you can save yourself some work.

Dynamic data is of a temporary nature. Full volume backups and log tapes are one example. These volumes typically have a short expiration period. You can move this type of data with the phased method. There is no reason to copy these volumes if they are going to expire soon.

7.5.6 Moving data out of the TS7700 Virtualization Engine

There are many reasons why you would want to move data out of the TS7700 Virtualization Engine. The most common is for disaster recovery or data interchange. You can move data out of the TS7700 Virtualization Engine in three ways:

- ▶ Host-based copy tools
- ▶ Copy Export
- ▶ DFSMSHsm ABARS

Host-based copy tools

You can use a host-based tool to copy the data from the TS7700 Virtualization Engine to the target.

With this method, the data is reprocessed by the host and copied to another medium. This method is described in 7.5.1, “Moving data into the TS7700 Virtualization Engine: Phased method” on page 377. The only difference is that you need to address the TS7700 Virtualization Engine as input and the non-TS7700 Virtualization Engine drives as output.

Copy Export

You can use the Copy Export function to copy the data.

With this function, a copy of selected logical volumes that is written to the TS7700 Virtualization Engine can be removed and taken off site for disaster recovery purposes. The benefits of volume stacking, which places many logical volumes on a physical volume, are retained with this function. In addition, because the data being exported is a copy of the logical volumes, the logical volumes data remains accessible by the production host systems. We describe this method in detail in the following sections:

- ▶ 5.3.2, “Implementing Copy Export” on page 288
- ▶ 10.4, “Disaster recovery using Copy Export” on page 742

The Copy Export function for stand-alone configurations was introduced with TS7700 Virtualization Engine code level 8.3.x.x and Library Manager code level 534.x. For Grid configurations, the Copy Export function was introduced with TS7700 Virtualization Engine code level 8.4.x.x. Although no host software updates required to support the Copy Export function are required, other functions are supported in the TS7700 Virtualization Engine that do require a later level of host software for support. One of those, host console request, does require z/OS support, which is provided with z/OS V1R6 and later. Refer to OAM APAR OA20065 and device services APARs OA20066, OA20067, and OA20313.

Note: The expectation is that customers will execute the Copy Export operation on a periodic basis, possibly even more than once a day. Because the purpose is to get a copy of the data off site for disaster recovery purposes, performing it soon after the data is created minimizes the time for the recovery point objective.

DFSMShsm ABARS

The third way is to copy the data with the DFSMShsm ABARS function.

ABARS is the command-driven DFSMShsm function that backs up a user-defined group (called an aggregate group) of data sets (usually for recovery purposes) at another computer site or at the same site. ABARS can be used to back up and recover both SMS-managed and non-SMS-managed data, on DASD and on tape. Using the DFSMShsm ABARS function, group the data you want to move outside the TS7700 Virtualization Engine, then start addressing other tape drives outside the TS7700 Virtualization Engine, or use the Copy Export function. In this way, you obtain an exportable copy of the data that can be put in an off-site location.

The process is as follows:

1. Create a selection data set
2. Define an aggregate group
3. Execute the ABACKUP VERIFY command
4. Execute the ABACKUP EXECUTE command

Create a selection data set

Before you can run an aggregate backup, create one or more selection data sets. The selection data set lists the names of the data sets to be processed during aggregate backup.

You can identify the data set names in a single selection data set, or you can divide the names among as many as five selection data sets. You can specify six types of data set lists in a selection data set. The type you specify determines which data sets are backed up and how they are recovered.

An INCLUDE data set list is a list of data sets to be copied by aggregate backup to a tape data file where they can be transported to the recovery site and recovered by aggregate recovery. The list can contain fully qualified data set names or partially qualified names with place holders. DFSMSHsm expands the list to fully qualified data set names.

Using a selection data set with the names of the data sets you want to export from the TS7700 Virtualization Engine, obtain a list of files on logical volumes that the ABARS function copies to non-TS7700 Virtualization Engine drives.

You can also use the Copy Export function to move the ABAR tapes to a Data Recovery site, out of the library.

Define an aggregate group

Define an aggregate group and related Management Class to specify exactly which data sets are to be backed up.

Define the aggregate group and Management Class used for aggregate backup to DFSMS through ISMF screens.

The aggregate group lists the selection data set names, instruction data set name, and additional control information used by aggregate backup in determining which data sets are to be backed up.

Execute the ABACKUP VERIFY command

You may use the ABACKUP command to verify the contents of the aggregate backup without actually backing up any data sets, which is the same as performing a test run of aggregate backup. An example of the ABACKUP command is as follows:

```
HSEND ABACKUP agname VERIFY UNIT(non_TS7700_unit) PROCESSONLY(USERTAPE)
```

With the PROCESSONLY(USERTAPE) keyword, only tape data sets are processed. In this way, therefore, you can be sure that only the input data from TS7700 Virtualization Engine logical volumes is used.

Execute the ABACKUP EXECUTE command

When you are ready, perform the actual backup by using the following command:

```
HSEND ABACKUP agname EXECUTE UNIT(non_TS7700_unit) PROCESSONLY(USERTAPE)
```

When you issue the ABACKUP command with the EXECUTE option, the following tape files are created for later use as input for aggregate recovery:

- ▶ Data file: Contains copies of the data sets that have been backed up.
- ▶ Control file: Contains control information needed by aggregate recovery to verify or recover the application's data sets.
- ▶ Instruction/activity log file: Contains the instruction data set, which is optional.

Summary

At the end of this process, you obtain an exportable copy of the TS7700 Virtualization Engine data, which can be used for disaster recovery and stored off site using other physical tapes.

Consider using the Copy Export function, which allows you to move a copy of the original logical volumes to an off-site location without reading the tape data twice. The Copy Export function operates on another Physical Volume Pool in the library and creates the copy in the background without any process required on the host. However, Copy Export requires an empty TS7700 Virtualization Engine at your disaster site.

For more information, see the following resources:

- ▶ For Copy Export, see 5.3.2, “Implementing Copy Export” on page 288.
- ▶ For Copy Export Recovery, see 10.4, “Disaster recovery using Copy Export” on page 742.
- ▶ For using the DFSMSHsm ABARS function, see *DFSMSHsm Storage Administration Guide*, SC35-0421.

7.6 Migration of DFSMSHsm-managed data

The z/OS DFSMS Hierarchical Storage Manager (DFSMSHsm) has its own functions that allow full utilization of the capacity of tape cartridges. Therefore, at initial consideration, using the TS7700 Virtualization Engine in a DFSMSHsm environment does not bring any special advantage from a capacity point of view, because you create virtual volumes, which, after copied, fill a stacked volume completely. You can achieve the same results by writing directly to a physical tape drive, leaving the TS7700 Virtualization Engine storage for other applications that cannot use the full cartridge capacity and are therefore better candidates for TS7700 Virtualization Engine management.

Although DFSMSHsm is an application that is capable of using the full cartridge capacity, for various reasons you might want to consider using the TS7700 Virtualization Engine instead of native physical drives for DFSMSHsm data. For example, when writing ML2 data onto a cartridge with an uncompressed capacity of 300 GB, chances are higher that a recall request needs exactly this cartridge that is currently being written to by a space management task. This incident is known as *recall takeaway*.

The effects of recall takeaway can be a real disadvantage when writing ML2 data onto native, high capacity cartridges, because the space management task must set aside its output tape to make it available to the recall task. Although the partially-filled output tape remains eligible for subsequent selection the next time that space management runs, it is possible to accumulate a number of partial tapes beyond DFSMSHsm's needs if recall takeaway activity occurs frequently. Excess partial tapes created by recall takeaway activity result in poor utilization of native cartridges. In addition, because recall takeaway activity does not cause the set-aside tape to be marked full, it is not automatically eligible for recycling, despite its poor utilization.

High capacity cartridges are more likely to experience both frequent recall takeaway activity, and also frequent *piggy-back recall* activity, in which recalls for multiple data sets on a single tape are received while the tape is mounted. Although piggy-back recalls have a positive effect by reducing the number of mounts required to perform a given number of recalls, you must also consider that multiple recalls from the same tape must be performed serially by the same recall task. Were those same data sets to reside on different tapes, the recalls could potentially be performed in parallel, given a sufficient number of recall tasks. In addition, the persistence of the virtual tape in the Tape Volume Cache after it has been demounted allows

DFSMSHsm to perform ML2 recalls from the disk cache for a period of time without requiring that a physical tape be mounted.

Other reasons also exist for directing DFSMSHsm data into a TS7700 Virtualization Engine. The amount of native drives limits the number of DFSMSHsm tasks that can run concurrently. With the large number of up to 256 virtual drives in a single-cluster grid configuration or 512 virtual drives in a two-cluster grid configuration, you can dedicate a larger number of virtual drives to each DFSMSHsm function and allow for higher throughput during your limited backup and space management window.

When increasing the number of DFSMSHsm tasks to take advantage of the large number of virtual drives in a TS7700 Virtualization Engine, consider adding more DFSMSHsm auxiliary tasks (MASH), rather than simply increasing the number of functional tasks within the existing started tasks. Each DFSMSHsm started task can support up to 15 AUTOBACKUP tasks.

Other reasons for using the TS7700 Virtualization Engine with DFSMSHsm are the greatly reduced run times of DFSMSHsm operations that process the entire volume, such as AUDIT MEDIACONTROLS and TAPECOPY.

DFSMSHsm can benefit from TS7700 Virtualization Engine's high throughput and from its large tape volume cache size, allowing for long periods of peak throughput.

DFSMSHsm data is well suited for the TS7700 Virtualization Engine, given the appropriate tailoring of those parameters that can affect DFSMSHsm performance. The subsequent sections describe this tailoring in more detail.

For more details, see *DFSMSHsm Storage Administration Guide*, SC26-7402.

7.6.1 Volume and data set sizes

The size of user data sets is important when you choose between a TS7700 Virtualization Engine or native drives such as 3592. DFSMSHsm Migration, Backup, and Recycle use only Single File Format to write to tape cartridges.

z/OS supported data set sizes

Different data set sizes are supported for disk and tape data sets, based on the data set organization and the number of volumes a single data set can span:

- ▶ DASD data sets are limited to 59 volumes, except for PDS and PDSE data sets, which are limited to one volume.
- ▶ A data set on a VIO-simulated device is limited to 65,535 tracks and to one volume.
- ▶ Tape data sets are limited to 255 volumes.

Table 7-4 lists the maximum data set sizes supported in z/OS environments.

Table 7-4 Maximum supported data set sizes

Storage medium	Maximum volume size	Maximum number of volumes	Maximum data set size
DASD: IBM System Storage DS8000®	65,520 CYL = 54 GB	59	3.18 TB
Tape: TS1120	700 GB x 2.5 Compression	40	70 TB
Tape: TS7700 Virtualization Engine	4,000 MB x 2.5 Compression	40	400GB

DFSMSHsm-supported data set sizes

Single file format, as used by DFSMSHsm, reduces I/O and system serialization because only one label is required for each connected set (as opposed to multiple file format tapes that require a label for each data set). The standard-label tape data set that is associated with the connected set can span up to the allocation limit of 255 tapes. This standard-label tape data set is called the DFSMSHsm tape data set. Each user data set is written in 16 K logical blocks to the DFSMSHsm tape data set.

Important: A single DFSMSHsm user data set can span up to 40 tapes. This limit is for migration, backup, and recycle.

After DFSMSHsm writes a user data set to tape, it checks the volume count for the DFSMSHsm tape data set. If the volume count is greater than 215, the DFSMSHsm tape data set is closed, and the currently mounted tape is marked full and is de-allocated. The number 215 is used so that a data set spanning 40 tapes fits within the 255-volume allocation limit. DFSMSHsm selects another tape, and then starts a different DFSMSHsm tape data set.

Data set spanning can be reduced by using the SETSYS TAPESPANSIZE command.

DFSMSHsm and large logical volumes

The TS7700 Virtualization Engine supports logical volume sizes of 400, 800, 1000, 2000, and 4000 MB. With a maximum of 40 volumes supported and assuming a compression ratio of 2.5:1, the maximum user data set size for 800 MB volumes is as follows:

$$800 \text{ MB} \times 2.5 \times 40 = 80 \text{ GB}$$

Let us assume that you have a very large data set of 300 GB. Such a data set does not fit on 40 volumes of 800 MB each, but it would fit on 4000 MB large virtual volumes, as shown in the following example:

$$4000 \text{ MB} \times 2.5 \times 40 = 400 \text{ GB}$$

Any single user data set larger than 400 GB is a candidate for native TS1120 Tape Drives; assuming a compression rate of 2.5:1, they might not fit onto the supported number of 40 volumes. In this case, consider using native TS1130 Tape Drives rather than TS7700 Virtualization Engine.

IDCAMS DCOLLECT BACKUPDATA and MIGRATEDATA can be used to determine the maximum size of backed-up and migrated data sets, respectively, in DFSMSHsm's inventory.

Important: DFSMSHsm can have more than one address space on one LPAR, Multi Address Space Support (MASH), or have a separate setup on separate LPARs in your PARMLIB member ARCCMDxx and separated by ONLYIF statements. One DFSMSHsm address space can have a MIGUNIT(3590-1) and the other address space a MIGUNIT(TS7700 Virtualization Engine). The same is true for BUUNIT. The DFSMSHsm instance that has the 3592 as migration or backup unit can run space management or auto backup only for that Storage Group (SG), where all your large data sets, such as z/FS, reside. The other DFSMSHsm instance would migrate and back up all the smaller data sets onto TS7700 Virtualization Engine. You can use a command such as F DFSMS2,BACKDS or F DFHSM2,BACKVOL(SG2) to issue the command to the second address space of DFSMSHsm:

- ▶ SETSYS TAPEMIGRATION(ML2TAPE(TAPE(unittyp)))
- ▶ SETSYS RECYCLEOUTPUT(MIGRATION(unittyp))
- ▶ SETSYS BACKUP(TAPE(unittyp))
- ▶ SETSYS RECYCLEOUTPUT(BACKUP(unittyp))

Migration to a different logical volume size

To make sure that DFSMSHsm starts using larger data sets, you must mark as full any empty or partially filled tapes that are written using the previous logical volume size. To identify these tapes, issue the following DFSMSHsm command:

```
LIST TTOC SELECT(NOTFULL)
```

Each tape identified as being empty or partially filled must be marked full by using one of the following DFSMSHsm commands:

```
DELVOL volser MIGRATION(MARKFULL)
DELVOL volser BACKUP(MARKFULL)
```

As DFSMSHsm migrates data and creates backup copies, it prefers to add to an existing migration/backup volume. As the volume nears full, it handles spanning of data sets, as described in “Tape spanning” on page 391. If a data set spans across DFSMSHsm volumes, it becomes a “connected set” in DFSMSHsm terms.

A key point, however, is that whether or not the data set spans, DFSMSHsm uses FEOV processing to get the next volume mounted. So, the system believes that the volume is part of a multi-volume set regardless of whether DFSMSHsm identifies it as a connected set. Because of the EOV processing, the newly mounted DFSMSHsm volume will use the same Data Class and other SMS constructs as the previous volume.

With the DFSMSHsm SETSYS PARTIALTAPE MARKFULL option, DFSMSHsm marks the last output tape full, even though it has not reached its physical capacity. By marking the last volume full, the next time processing starts, DFSMSHsm will use a new volume, starting a new multi-volume set and allowing for the use of a new Data Class and other SMS constructs. If the volume is not marked full, the existing multi-volume set will continue to grow and to use the old constructs.

Use the SETSYS PARTIALTAPE MARKFULL option, because it reduces the number of occasions in which DFSMSHsm will append to a partial tape, which results not only in the need to mount a physical tape, but also in the invalidation of the existing virtual tape, which will eventually need to be reclaimed by the TS7700 Virtualization Engine.

This discussion is relevant to outboard policy management and the implementation of different logical volume sizes. If all volumes have been marked full, you can simply update

your ACS routines to assign a new Data Class and other SMS constructs and, from then on, each new migration or backup volume will use the new size.

7.6.2 TS7700 Virtualization Engine implementation considerations

This section summarizes some DFSMSHsm implementation considerations with regard to the TS7700 Virtualization Engine.

Mount wait time

If you direct DFSMSHsm data into the TS7700 Virtualization Engine, modify your DFSMSHsm mount wait timer to be 12 minutes. This modification allows for possibly needed extra time on specific mounts for the TS7700 Virtualization Engine to stage the data back into cache.

Logical volume size

Consider using large logical volumes, such as 4000 MB, for backup and smaller logical volumes for migration. If you have a high recall rate from ML2, you might not even want to use the entire capacity of a MEDIA1 or MEDIA2 virtual volume. Installations in which recalls from ML2 are rare and installations in which extremely large data sets are migrated that could result in the 40-volume limit being reached should use the maximum capacity of the virtual volume. You must select another SMS DATACLAS for backup than for migration through your ACS routine.

See Table 7-5 on page 393 when tailoring the ARCCMDxx SETSYS parameters. From TS7700 Virtualization Engine V1.4 onwards, with the host support added with APAR OA24969, HSM is aware of the large virtual volume capacity. Therefore, it is not necessary to use high PERCENTFULL values anymore, to tune capacity of tapes, from a DFSMSHsm point of view. The maximum PERCENTFULL value that can be defined is 110%.

In TS7700 Virtualization Engine V1.4, this support was added to report larger capacity through the read buffered log. However, the host support to use this new function was not added until TS7700 Virtualization Engine V1.5. Logical volume sizes have not changed since then.

Other applications might have a similar existing TAPECAPACITY-type specification or a PERCENTFULL-type specification enabling applications to write beyond the default volume sizes for MEDIA1 (cartridge system tape) and MEDIA2 (enhanced capacity cartridge system tape).

In the case of OAM's Object Tape Support, the TAPECAPACITY parameter in the SETOAM statement of the CBROAMxx PARMLIB member is used to specify the larger logical volumes sizes. Because of this new functionality, defining TAPECAPACITY in the CBROAMxx PARMLIB member is not necessary. For additional information about outboard policy management, see *z/OS DFSMS Object Access Method Planning, Installation and Storage Administration Guide for Tape Libraries*, SC35-0427.

Multi-system considerations

If multiple TS7700 Virtualization Engines are eligible for a request, also consider that the same logical volume size is used for the request across all libraries. When displaying the volumes through your tape management system, the tape management system might continue to display the volume capacity based on the default volume size for the media type with the volume usage (or a similar parameter) showing how much data has actually been written to the volume reflecting its larger capacity.

Scratch volumes

The default volume size is overridden at the library through the Data Class policy specification and is assigned or reassigned when the volume is mounted for a scratch mount. Using a global scratch pool, you benefit from a fast mount time by using the fast-ready attribute for the scratch category, as explained in 4.3.5, “Defining Fast Ready categories” on page 220. Consider using the following definitions to benefit from the fast scratch mount times:

```
SETSYS SELECTVOLUME(SCRATCH)
SETSYS TAPEDELETION(SCRATCHTAPE)
SETSYS PARTIALTAPE(MARKFULL)
```

The first of these commands requests DFSMSHsm to use volumes from the common scratch pool, and the second one defines that DFSMSHsm returns tapes to the common scratch pool. The last one defines that an DFSMSHsm task will mark the last tape it used in a cycle to be full, thus avoiding a specific mount during the next cycle. The MARKFULL parameter does not mean a waste of space using TS7700 Virtualization Engine, because the stacked volume contains only the written data of each logical volume copied and the same applies to the TVC.

Tape spanning

You can use the optional TAPESPANSIZE parameter of the SETSYS command to reduce the spanning of data sets across migration or backup tape volumes, for example:

```
SETSYS TAPESPANSIZE(4000)
```

The value in parentheses represents the maximum number of megabytes of tape (ML2 or backup) that DFSMSHsm may leave unused while it tries to eliminate spanning of data sets. To state this in another way, this value is the minimum size of a data set that is allowed to span tape volumes. Data sets whose size is less than the value do not normally span volumes. Only those data sets whose size is greater than or equal to the specified value are allowed to span volumes.

This parameter offers a trade-off: It reduces the occurrences of a user data set spanning tapes in exchange for writing less data to a given tape volume than its capacity would otherwise allow. The amount of unused media can vary from 0 to nnnn physical megabytes, but roughly averages 50% of the customer's median data set size. For example, if you specify 4000 MB and your median data set size is 2 MB, then on average only 1 MB of media is unused per cartridge.

Installations that currently experience an excessive number of spanning data sets might want to consider specifying a larger value in the SETSYS TAPESPANSIZE command. Using a high value reduces tape spanning. In a TS7700 Virtualization Engine, this value reduces the number of virtual volumes that need to be recalled to satisfy DFSMSHsm recall/recover requests. You can be generous with the value, because no space is wasted. For example, a TAPESPANSIZE of 4000 means that any data set with less than 4000 MB that does not fit on the remaining space of a virtual volume will be started on a fresh new virtual volume.

7.6.3 DFSMSHsm task-related considerations

To better understand the use of DFSMSHsm with TS7700 Virtualization Engine, we summarize the DFSMSHsm functions that use tapes and analyze the benefit of tape virtualization for these functions.

Backups of DFSMSHsm control data sets

The backup of DFSMSHsm control data sets (CDSs) can easily be done in a TS7700 Virtualization Engine, exploiting the benefit of using virtual volumes instead of physical volumes, which might otherwise be underutilized.

Volume dumps

When using TS7700 Virtualization Engine as output for the DFSMSHsm AUTODUMP function, *do not* specify the following parameters:

```
DEFINE DUMPCLASS(dclass STACK(nn))  
BACKVOL SG(sgname)|VOLUMES(volser) DUMP(dclass STACK(10))
```

These parameters were introduced to force DFSMSHsm to use the capacity of native physical cartridges. If used with TS7700 Virtualization Engine, they cause unnecessary multivolume files and reduce the level of parallelism possible when the dump copies are restored. Use the default value, which is NOSTACK.

Migrate or recall (DFSMSHsm Migration Level 2)

When using a TS7700 Virtualization Engine as DFSMSHsm Migration Level 2, consider the number of simultaneous recall processes. Consider how many recall tasks are started at the same time and compare that number with the number of physical drives that are assigned to your TS7700 Virtualization Engine.

For example, if your installation often has more than 10 tape recall tasks at one time, you probably need twelve TS1120 back-end drives to satisfy this throughput request, because all migrated data sets might already have been removed from the TVC and need to be recalled from tape.

Backup and recovery

Unlike the DFSMSHsm RECALL operation, RECOVERY usually has a lower frequency in an DFSMSHsm environment. Therefore, using TS7700 Virtualization Engine for DFSMSHsm backup and recovery functions benefits you without affecting DFSMSHsm performance. However, carefully review your DFSMSHsm performance requirements before moving DFSMSHsm BACKUP to the TS7700 Virtualization Engine.

TAPECOPY

The DFSMSHsm TAPECOPY function requires that original and target tape volumes are of the same media type and use the same recording technology. Using a TS7700 Virtualization Engine as the target for the TAPECOPY operation from a original volume that is not a TS7700 volume can cause problems in DFSMSHsm, because TS7700 Virtualization Engine virtual volumes have different volume sizes, even though they are defined as CST or ECCST.

Use the information in Table 7-5 to tailor your TAPECOPY environment.

Table 7-5 TAPECOPY utilization

ORIGINAL volume unit name	ALTERNATE volume unit name	Percent full to be defined (assuming 2:1 compression)
TS7700 (CST) - 400 MB	3490E (CST)	106%
TS7700 (ECCST) - 800 MB	3490E (ECCST)	107%
3490E (CST) - 400 MB	TS7700 CST - 400 MB	45%
3490E (ECCST) - 800 MB	TS7700 (ECCST)- 800 MB	45%
TS7700 (CST) - 400 MB	TS7700 (CST) - 400 MB	106%
TS7700 (CST) - 1 GB	TS7700 (CST) - 1 GB	100%
TS7700 (CST) - 2 GB	TS7700 (CST) - 2 GB	100%
TS7700 (CST) - 4 GB	TS7700 (CST) - 4 GB	100%
TS7700 (ECCST) - 800 MB	TS7700 (ECCST) - 800 MB	107%
TS7700 (ECCST) - 1 GB	TS7700 (ECCST) - 1 GB	100%
TS7700 (ECCST) - 2 GB	TS7700 (ECCST) - 2 GB	100%
TS7700 (ECCST) - 4 GB	TS7700 (ECCST) - 4 GB	100%

For example, if you are planning to put DFSMSHsm alternate copies into a TS7700 Virtualization Engine, a tape capacity of 45% might not be enough for the input non-TS7700 Virtualization Engine ECCST cartridges. TAPECOPY fails if the (virtual) output cartridge encounters EOV before the input volume has been copied completely.

However, using TS7700 Virtualization Engine logical volumes as the original and 3490E native as the TAPECOPY target might cause EOV at the alternate volume because of the higher IBMLZ1 compression seen on the virtual drive compared to the IDRC compression on the native drive.

For special situations where copying from standard to enhanced capacity media is needed, the following patch command can be used:

```
PATCH .MCVT.+4F3 BITS(.....1..)
```

Note: In TS7700 Virtualization Engine V1.4 and later with host support applied, 1 GB and larger sized, you do not need to use PERCENTFULL greater than 100%.

DUPLEX TAPE

For duplexed migration, both output tapes have to be of the exact same size and unit type. A good practice is to use a multi-cluster grid and let the hardware do the duplex rather than the DFSMSHsm software function. This way also helps to more easily manage the disaster side. You can use GDPS and switch to the remote DASD side and the tape VOLSER itself does not need to be changed. No TAPEREP or SETSYS DISASTERMODE commands are needed.

When using the TS1120 or TS1130 Tape Drive for duplexed migration output, performance is degraded because of the back-to-back SYNCDEV operations done for the original and the alternate tapes. APAR OA09928 provides a patch allowing syncs on the alternate tape to be disabled. The performance improvement varies with data set size, with the greatest

improvements seen for the smaller data sets. Performance improvements can be quite substantial.

When HSM writes ML2 data to tape, it deletes the source data as it goes along, but before the RUN is issued to the TS7700 Virtualization Engine. This means that, for a period of time until the copy is made, only one copy of the ML2 data may exist. The reason is because the TS7700 Virtualization Engine Grid, even with a Copy Consistency Point (CCP) of RR, makes a second copy at RUN time.

By using the appropriate Management Class settings in SMS, you can make sure that a data set is not migrated to ML2 before a valid backup copy of this data set exists. This way, there are always two valid instances from which the data set can be retrieved: one backup and one ML2 version. After the second copy is written at rewind-unload time, two copies of the ML2 data will exist in the Grid.

Another way to ensure that two copies of the ML2 data exist is to use HSM duplexing. This way creates two separate copies of the ML2 data before HSM deletes it. Ideally, with a multi-cluster grid, you want one copy of the data in one cluster and the second copy in another to avoid loss of data if one of the clusters experiences a disaster. You can use the CCPs to ensure that each copy of the duplexed data is sent to separate clusters.

RECYCLE

The DFSMSHsm RECYCLE function reduces the number of logical volumes inside the TS7700 Virtualization Engine, but when started, it can cause bottlenecks in the TS7700 Virtualization Engine recall process. If you have a TS7700 Virtualization Engine with four physical drives, use a maximum of two concurrent DFSMSHsm RECYCLE tasks. If you have a TS770 Virtualization Engine with six physical drives, use no more than five concurrent DFSMSHsm RECYCLE tasks.

Select the RECYCLEPERCENT carefully and consider the following information:

- ▶ You will free up logical volumes residing on a stacked volume with hundreds of other logical volumes.
- ▶ The space occupied by the logical volume will be freed up only if and when the logical volume is used (overwritten) again, unless you are using Expired Volume Management.
- ▶ To RECYCLE, the TS7700 Virtualization Engine has to load the input volumes into the TVC.

Use a RECYCLEPERCENT value that depends on the logical volume size, for example:

- ▶ 5 for 1000, 2000, or 4000 MB volumes
- ▶ 10 for 400 or 800 MB volumes

Using the following commands for RECYCLE input can be helpful while selecting and migrating data to and from a TS7700 Virtualization Engine:

- ▶ RECYCLE SELECT (INCLUDE(RANGE(nnnnn:mmmm)))
- ▶ RECYCLE SELECT (EXCLUDE(RANGE(nnnnn:mmmm)))

The immediate purpose is to enable you to set up volume ranges for various media types and emulation types, such as TS7700 Virtualization Engine logical volumes and 3490-emulated cartridges. There are no special data set names for RECYCLEOUTPUT, although you must code your ACS routines to route RECYCLEOUTPUT to the library using the &UNIT variable.

See *DFSMSHsm Primer*, SG24-5272 for more information about implementing DFSMSHsm.

7.7 DFSMSrmm and other tape management systems

No changes are required to any tape management system to support basic TS7700 Virtualization Engine. You need only review the retention and movement criteria for the data in the TS7700 Virtualization Engine. Besides, you must check your daily tape management process, to delete any step that will be related to EJECT activities.

DFSMSrmm accepts logical volumes capacity from an open close end-of volume (OCE) module. DFSMSrmm can now always list the actual reported capacity from TS7740 Virtualization Engine.

To start the low-on-scratch procedure, DFSMSrmm uses these messages:

- ▶ CBR3660A
- ▶ CBR3792E
- ▶ CBR3794A

When you direct allocations inside the TS7700 Virtualization Engine, the Vital Record Specifications (VRSs) or vault rules indicates to the tape management system that the data set will never be moved outside the library. During VRSEL processing, each data set and volume is matched to one or more VRSs and the required location for the volume is determined based on priority. The volume's required location is set. The volume is not moved unless DSTORE is run for the location pair that includes the current volume location and its required location. For logical volumes, this required location can be used to determine which volume should be exported. For Copy Export, the required location is only used for stacked volumes that have been copy exported.

Other tape management systems must modify their definitions in a similar way. For example, CA-1 Tape Management must modify their "RDS" and "VPD" definitions in CA/1 PARMLIB. Control-M/Tape (Control-T) must modify its "rules" definitions in the Control-T PARMLIB.

The DFSMSrmm return-to-scratch process has been enhanced to allow more parallelism in the return-to-scratch process. EDGSPLCS is a new option for the EDGHSKP SYSIN file EXPROC command that can be used to return to scratch tapes in an asynchronous way. With the most recent software support changes, EDGSPLCS can be used to run scratch processing in parallel across multiple libraries or in parallel within a library. The only necessary step is to run different instances of CBRSP LCS. For additional information about the enhanced return-to-scratch process, see *z/OS DFSMSrmm Implantation and Customization Guide*, SC26-7405.

Stacked volumes cannot be used by the host; they are managed exclusively by the TS7700 Virtualization Engine. Do not allow any host to either implicitly or explicitly address these stacked volumes. To indicate that the stacked VOLSER range is reserved and cannot be used by any host system, define the VOLSERS of the stacked volumes to RMM.

Use the following PARMLIB parameter, assuming that VT is the prefix of your stacked TS7700 Virtualization Engine cartridges:

```
REJECT ANYUSE(VT*)
```

This parameter causes RMM to deny any attempt to read or write those volumes on native drives. There are no similar REJECT parameters in other tape management systems.

You do not need to explicitly define the virtual volumes to RMM. During entry processing, the active RMM automatically records information about each volume in its control data set. RMM uses the defaults that you specified in ISMF for the library entry values if there is no existing RMM entry for an inserted volume. Set the default entry status to SCRATCH.

When adding 1,000,000 virtual volumes, the size of the RMM CDS and the amount of secondary space available must be checked. RMM uses 1 MB for every 1000 volumes defined in its CDS; an additional 1,000,000 volumes would need 1000 MB of space. However, do not add all the volumes initially. See 4.3.12, “Inserting logical volumes” on page 240 for more information.

To increase the size of the RMM CDS, you have to quiesce RMM activities, back up the CDS, then reallocate a new CDS with a larger size and restore the CDS from the backup copy. To calculate the correct size of the RMM CDS, see *z/OS DFSMSrmm Guide and Reference*, SC26-7404.

Other tape management systems, such as BrightStor, CA-1 Tape Management Copycat Utility (BrightStor CA-1 Copycat), and BrightStor CA-Dynam/TLMS Tape Management Copycat Utility (BrightStor CA-Dynam/TLMS Copycat) must reformat their database to add more volumes. This means that they must stop to define more cartridges.

Additionally, some tape management systems do not allow the specification of tape volumes with alphanumeric characters or require user modifications to do so. See the proper product documentation for this operation.

In both RMM and the other tape management systems, the virtual volumes do not have to be initialized. The first time a VOLSER is used, TS7700 Virtualization Engine marks the virtual volume with VOL1, HDR1, and a tape mark, as though it had been done by EDGINERS or IEHINITT.

7.8 IBM Tivoli Storage Manager

Customers who use IBM Tivoli Storage Manager with TS1120 or TS1130 tape drives might experience reduced performance of IBM Tivoli Storage Manager and TS7700 Virtualization Engine compared to the connection to native drives. In general, IBM Tivoli Storage Manager operations perform better with native drives than with a TS7700 Virtualization Engine. You should understand the trade-off between performance, economy, and the TS7700 Virtualization Engine functionality.

When using a multi-cluster grid, you can get two or more copies instantly without any use of the ITSM COPYSTORAGEPOOL function and CPU cycles needed to do the function. You can even move a copy of the logical volumes off site with the Copy Export function.

7.8.1 Guidance for TS7700 Virtualization Engine usage

IBM Tivoli Storage Manager, like DFSMSHsm, can automatically fill a native 3592 cartridge. It can use the tape up to EOV, independent of the media type.

Note: Starting with Tivoli Storage Manager 6.1, which was released in 2009, there is no Tivoli Storage Manager Server support for z/OS anymore. Tivoli Storage Manager Client support is still available. For the most current information about IBM Tivoli Storage Manager support, go to the following address:

<http://www.ibm.com/support/docview.wss?uid=swg21243309>

If you plan to store IBM Tivoli Storage Manager data into the TS7700 Virtualization Engine, consider the following suggestions for placing data on your TS7700 Virtualization Engine:

- ▶ Use TS7700 Virtualization Engine for IBM Tivoli Storage Manager archiving for archiving and backup of large files or databases for which you do not have a high performance requirement during backup and restore. TS7700 Virtualization Engine is ideal for IBM Tivoli Storage Manager archive or long term storage because archive data is not frequently retrieved. Archives and restores for large files can see less impact from the staging. Small files, such as individual files on file servers, can see performance impacts from the TS7700 Virtualization Engine staging. (If a volume is not in cache, the entire volume must be staged before any restore can be done.)
- ▶ Set IBM Tivoli Storage Manager reclamation off by setting the reclamation threshold to 100%. IBM Tivoli Storage Manager, like DFSMSHsm, has a reclamation function to consolidate valid data from tapes with a low valid data percentage onto scratch tapes so that tapes can be freed up for reuse. IBM Tivoli Storage Manager reclamation with TS7700 Virtualization Engine can be slower because all volumes have to be staged to the cache. Periodically set ITSM reclamation “on” by setting the threshold to a lower value to regain the use of TS7700 Virtualization Engine volumes with a small amount of valid data that will not expire for a longer period of time. IBM Tivoli Storage Manager reclamation must be scheduled for off-peak hours.
- ▶ Use collocation to reduce the number of TS7700 Virtualization Engine volumes required for a full restore. IBM Tivoli Storage Manager has a collocation function to group IBM Tivoli Storage Manager client data onto a minimum set of tapes to provide a faster restore and to provide separation of client data onto different physical tapes. Collocation with TS7700 Virtualization Engine does not minimize the physical tapes used, but minimizes the number of logical volumes used. Collocation with TS7700 Virtualization Engine can improve restore time for large amounts of data. TS7700 Virtualization Engine does not ensure physical tape separation when collocation is used because separate logical volumes can reside on the same physical tape.
- ▶ Use TS7700 Virtualization Engine for IBM Tivoli Storage Manager database backups that are to be used for recovery from local media and use TS7700 Virtualization Engine at a recovery site or native drives for backups that are to be used for recovery from off-site media. IBM Tivoli Storage Manager requires a separate tape for every backup of the IBM Tivoli Storage Manager database, so a large number of logical volumes with less data is created. When using the TS7700 Virtualization Engine, you do not have to worry about the unused capacity of logical volumes.
- ▶ Use TS7700 Virtualization Engine for backups of primary pools, noting that similar considerations apply to copy storage pools. If only one copy pool is used for local backups, then that storage pool should not be in the TS7700 Virtualization Engine, because there is no guarantee that data in the copy storage pools are on different physical volumes. If storage pools for local and off-site backups are used, the copy storage pools for local backups can be in the TS7700 Virtualization Engine. The copy storage pools for off-site backups should use native drives or a TS7700 Virtualization Engine at the recovery site.
- ▶ Use TS7700 Virtualization Engine in server-to-server configurations for multiple IBM Tivoli Storage Manager server implementations. If you are using a ITSM server-to-server configuration, the data from your remote ITSM servers are stored as virtual volumes, which appear as sequential media volumes on the source server and are actually stored as archive files on a target server. These are ideal candidates for a TS7700 Virtualization Engine.

7.8.2 Native drives

Use native drives for data that will be used for frequent individual file restores or have a requirement for high performance for backup and restore without any delays because of staging activity. IBM Tivoli Storage Manager uses EXPORT function to move data from one IBM Tivoli Storage Manager server to another. This way requires that both servers have compatible devices for the EXPORT media. Use native drives for IBM Tivoli Storage Manager EXPORT unless you have the advanced function IMPORT/EXPORT at both TS7700 Virtualization Engines.

7.8.3 IBM Tivoli Storage Manager parameter settings

The settings for the following parameters can affect the performance of IBM Tivoli Storage Manager with TS7700 Virtualization Engine:

- ▶ MAXSCRATCH (storage pool definition)

As for DFSMSHsm, IBM Tivoli Storage Manager should use a scratch pool for tapes, because you do not have to predefine tapes to Tivoli Storage Manager, and you can benefit from the faster TS7700 Virtualization Engine scratch mounts.

- ▶ MOUNTLimit (device class definition)

With one TS7700 Virtualization Engine, you have up to 256 virtual drives available. The number of drives available for IBM Tivoli Storage Manager use can probably be increased, taking into account TS7700 Virtualization Engine performance. Set MOUNTLimit high enough so that the number of available drives does not limit the performance of IBM Tivoli Storage Manager tape operations.

- ▶ MOUNTRetention (device class definition)

When storing data in the TS7700 Virtualization Engine, you can set this parameter to zero, because you have a greater chance of finding the virtual volume still in the TVC when IBM Tivoli Storage Manager needs it. This way avoids the need to keep the virtual volume mounted and frees a virtual drive for other users.

- ▶ MAXCAPacity (device class definition)

With this parameter, you can tailor the size of the data written in a virtual volume. Having smaller virtual volumes can speed up recall processing. Using the full capacity of the virtual volume can limit the number of volumes used by Tivoli Storage Manager.

- ▶ Backup DB (database backup)

Use SCRATCH=YES to use tapes from the Tivoli Storage Manager scratch pool and benefit from the faster TS7700 Virtualization Engine scratch mounts.

For details about setting up Tivoli Storage Manager, see *IBM Tivoli Storage Manager Administrators Guide* at the following address:

<http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp>

7.9 DFSMSdss

This section describes the uses of DFSMSdss with the TS7700 Virtualization Engine.

7.9.1 Full volume dumps

DFSMSdss full volume dumps can use the TS7700 Virtualization Engine well. You have to plan carefully, however, and make sure you can achieve the required throughput: A DFSMSdss full volume physical dump can easily provide a data transfer rate of 10 MBps and higher for a single job. However, with today's TS7700 Virtualization Engine throughput capabilities, the TS7700 Virtualization Engine throughput capabilities most likely will not be a limiting factor. In the past, the data rate was often limited by the bandwidth of the DASD subsystem as the weakest part in the chain.

With TS7700 Virtualization Engine, you fill the stacked cartridge completely without changing JCL, using multiple virtual volumes. TS7700 Virtualization Engine then moves the virtual volumes created onto a stacked volume.

The only problem you might experience when using TS7700 Virtualization Engine for the DSS volume dumps is related to the size of the virtual volumes. If a single dump does not fit onto five logical volumes, you can either use an SMS DATACLAS specification, Volume Count nn, to enable more than five volumes. A better way is if you have TS7700 Virtualization Engine Release 7.4 installed and can choose a 4000 MB logical volume through your SMS DATACLAS. This way prevents unneeded multi-volume files.

Using the COMPRESS keyword of the DUMP command, you obtain a software compression of the data at the host level. As data is compressed at the TS7700 Virtualization Engine before being written into the tape volume cache, host compression is not required, unless channel utilization is high already.

7.9.2 Stand-Alone Services

Stand-Alone Services of DFSMSdss provides a stand-alone restore function that enables you to restore vital system packs without needing to rely on a System z environment.

Stand-Alone Services supports the 3494 Tape Library and the Virtual Tape Server. You may use it to restore from native and virtual tape volumes in a TS7700 Virtualization Engine. With Stand-Alone Services, you specify the input volumes on the RESTORE command and send the necessary mount requests to the TS3500 Tape Library.

You can use an initial program load (IPL) of the Stand-Alone Services core image from a virtual tape device and use it to restore dump data sets from virtual tape volumes.

Stand-Alone Services is provided as a replacement to the previous DFDSS V2.5 and DFSMS/MVS V1 stand-alone functions. The installation procedure for Stand-Alone Services retains, rather than replaces, the existing stand-alone restore program so you do not have to immediately change your recovery procedures. Implement the procedures as soon as you can and start using the enhanced Stand-Alone Services.

To use Stand-Alone Services, create a stand-alone core image suitable for IPL by using the new BUILDSEA command of DFSMSdss. Create a new virtual tape as non-labeled and then put the stand-alone program on it.

Note: The BUILDSEA command does not write over the label. A tape that the TS7700 Virtualization Engine marked as labeled initially *cannot* be changed to unlabeled. It is not possible to alter an LVOL from labeled to unlabeled. This criteria is valid for all other stand-alone programs also.

Use the following steps to use an IPL of the Stand-Alone Services program from a virtual device and to restore a dump data set from virtual volumes (see “Modify Logical Volumes window” on page 449 for information about how to use the TS7700 management interface to set a device in stand-alone mode):

1. Ensure that the virtual devices you will be using are offline to other host systems. Tape drives to be used for stand-alone operations must remain offline to other systems.
2. Set in stand-alone mode the virtual device from which you will load the Stand-Alone Services program by selecting Virtual Drives on the TS7700 management interface of the cluster where you want to mount the logical volume as shown in Figure 7-12.

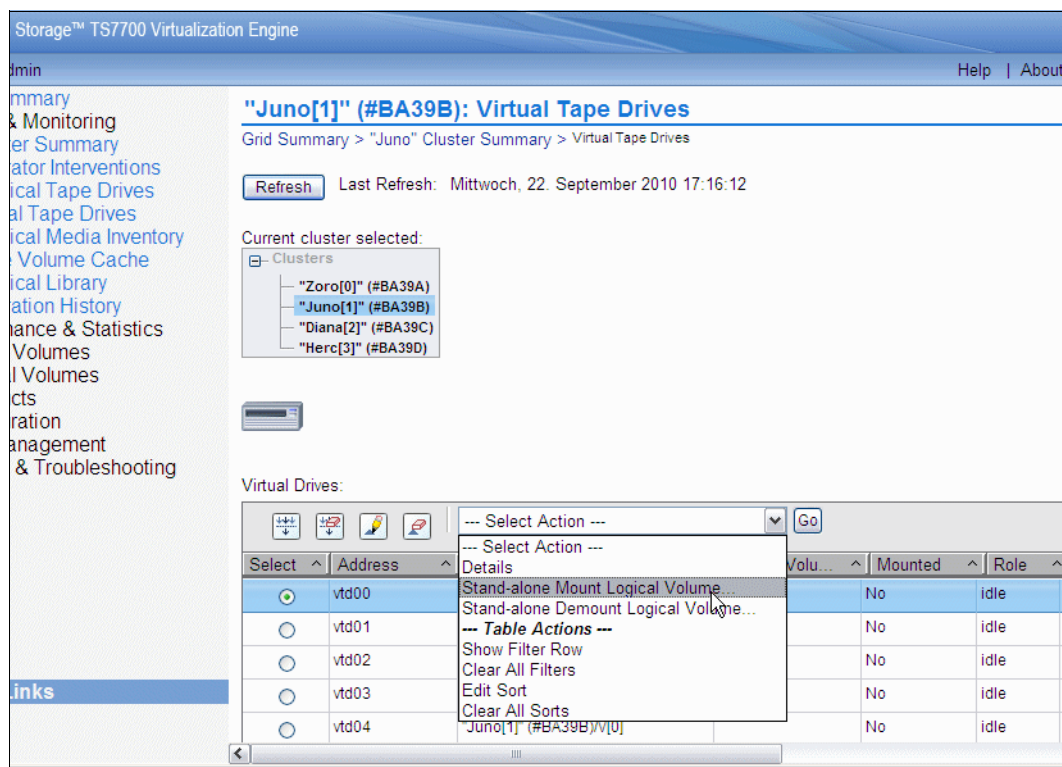


Figure 7-12 Virtual Tape Drive Panel

Select **Stand-alone Mount Logical Volume**, select a virtual device, and click **Go**. Figure 7-13 shows the panel that opens, where you specify the virtual volume that contains the core image for IPL.



Figure 7-13 Stand-alone Mount Logical Volume Panel

3. Load the Stand-Alone Services program from the device you just set in stand-alone mode. As part of this process, select the operator console and specify the input device for entering Stand-Alone Services commands.
4. When the IPL is complete, enter the Stand-Alone Services RESTORE command from the specified input device. Example 7-1 shows a group of statements for using this command.

Example 7-1 RESTORE command

```
RESTORE FROMDEV(TAPE) FROMADDR(0A40) TOADDR(0900) -
      NOVERIFY TAPEVOL((L00001),(L00002))
```

L00001 and L00002 are virtual volumes that contain the dump data set to be restored, 0A40 is the virtual device used for reading source volumes L00001 and L00002, and 0900 is the device address of the DASD target volume to be restored.

Stand-Alone Services requests the TS7700 to mount the source volumes in the order in which they are specified on the TAPEVOL parameter. It automatically unloads each volume, then requests the TS7700 Virtualization Engine to demount it and to mount the next volume.

5. When the restore is complete, unload and demount the IPL volume from the virtual device by using the TS7700 MI's Setup Stand-alone Device window.
6. Take the IPL device out of stand-alone mode by selecting **Stand-alone Demount Logical Volume** on the Virtual Drives panel (see Figure 7-12 on page 400).

Stand-Alone Services issues the necessary mount and demount orders to the library. If you are using another stand-alone restore program that does not support the mounting of library resident volumes, you would have to set the source device in stand-alone mode and manually instruct the TS7700 Virtualization Engine to mount the volumes using the Setup Stand-alone Device window.

For details about how to use Stand-Alone Services, see *z/OS DFSMSdss Storage Administration Reference*, SC35-0424.

7.10 Object Access Method

Tape cartridges provide a low-cost storage medium for storing primary and backup copies of Object Access Method (OAM) objects.

Allowing objects to be stored on tape volumes in conjunction with DASD and optical media provides flexibility and more efficiency within the storage management facility.

OAM stores objects on a TS7700 Virtualization Engine as in a normal TS3500 Tape Library, with up to 256 virtual drives and many virtual volumes available.

Consider using the TAPEPERCENTFULL parameter with object tape data, because the retrieval time of an OAM object is important. The recall time for smaller logical volumes can be reduced considerably.

TAPECAPACITY is not needed anymore; from TS7700 Virtualization Engine V1.4 onwards, with APAR OA24966 installed, OAM should know the real capacity of each virtual tape.

Virtual volumes in a TS7700 Virtualization Engine can contain primary or backup copies of OAM objects, addressing either OBJECT or OBJECT BACKUP Storage Groups. We advise addressing TS7700 Virtualization Engine with the OBJECT Storage Group and other non-TS7700 Virtualization Engine tape devices with the OBJECT BACKUP Storage Group.

A virtual volume can contain multiple OAM objects, separated by a buffer space. To optimize the use of TS7700 Virtualization Engine storing OAM object data, consider the following suggestions:

- ▶ Review the MOUNTWAITTIME parameter when using TS7700 Virtualization Engine to store OAM object tape data. The default (five minutes) should probably be increased. Twelve minutes is a better number in case you have to recall a logical volume to read object data and there are other recall requests queued at the time. The TS7700 Virtualization Engine might need to stage the data back into cache and this accounts for the extra mount time.
- ▶ Review the MAXTAPERETRIEVETASKS and MAXTAPESTORETASKS parameters when using TS7700 Virtualization Engine, because you have more virtual tape drives available.
- ▶ Other parameters, such as DEMOUNTWAITTIME, TAPEPERCENTFULL, and TAPEFULLTHRESHOLD, also might need to be reviewed when using TS7700 Virtualization Engine to store OAM data.

7.11 Database backups

Using a TS7700 Virtualization Engine as output confers several advantages to database backups. In this section, we provide a detailed description of these benefits for database products, such as DB2.

7.11.1 DB2 data

DB2 uses tapes for two purposes: for storing archive logs and for storing image copies. Either one may be created in multiple copies, to be stored both on site for local recovery and off site for disaster recovery purposes. To use DB2 tape data with the TS7700 Virtualization Engine, use the approaches described in this section.

Archive logs

DB2 keeps track of database changes in its active log. The active log uses up to 31 DASD data sets (up to 62 with dual logging) in this way: When a data set becomes full, DB2 switches to the next one and automatically offloads the full active log to an archive log.

Archive logs are sequential data sets that are allocated on either DASD or tape. When archiving to tape, a scratch tape volume is requested each time.

Archive logs contain unique information necessary for DB2 data recovery. Therefore, to ensure DB2 recovery, customers usually make backups of archive logs. You can use general backup facilities or DB2's dual archive logging function.

When creating a Dual Copy of the archive log, usually one is local and the other is for disaster recovery. The local copy can be written to DASD, then moved to tape, using Tape Mount Management (TMM). The other copy can be written directly to tape and then moved to an off-site location.

With TS7700 Virtualization Engine, you can write the local archive log directly inside the TS7700 Virtualization Engine. Avoiding the use of TMM saves DASD space, saves DFSMSshm CPU cycles, and simplifies the process. The disaster recovery copy must be created on non-TS7700 Virtualization Engine tape drives, so that it can be moved off site.

The size of an archive log data set varies from 150 MB to 1 GB. The size of a virtual volume on a TS7700 Virtualization Engine can be up to 4000 MB, so be sure that your archive log can

fit in only one virtual volume. We suggest using a single volume when unloading an archive log to tape. The size of a virtual volume on a TS7700 Virtualization Engine can be up to 12,000 MB, assuming a 3:1 compression ratio.

Tailoring the size and number of active log DASD data sets allows you to obtain an archive log on tape whose size does not exceed the virtual volume size.

Limiting data set size might increase the frequency of offload operations and reduce the amount of active log data on DASD. However, this should not be a problem, because TS7700 Virtualization Engine does not require manual operation, and archive logs will stay in the TVC for some time and be available for fast recovery.

One form of DB2 recovery is backward recovery, typically done after a processing failure, where DB2 backs out uncommitted changes to resources. When doing so, DB2 processes log records in reverse order, from the latest back toward the oldest.

If the application being recovered has a large data set and makes only a few commit operations, you probably need to read the old archive logs that are on tape. When archive logs are on tape, DB2 uses read-backward channel commands to read the log records. Read-backward is a slow operation on tape cartridges processed on real IBM 3480 (if IDRC enabled) and IBM 3490 tape drives. On a TS7700 Virtualization Engine, it is only about 20% slower than a normal I/O because data is retrieved from the TVC, so the tape drive characteristics are replaced by the random access disk characteristics. Another benefit TS7700 Virtualization Engine can provide for DB2 operations is the availability of up to 256 (single cluster) or 512 virtual drives (multi-cluster grid configuration), because DB2 often needs a large number of drives concurrently to perform recovery or backup functions.

Image copies

Image copies are backup copies of table spaces in a DB2 database. DB2 can create both full and incremental image copies. A full image copy contains an image of the whole table space at the time the copy was taken. An incremental image copy contains only those pages of a table space that have changed since the last full image copy was taken. Incremental image copies are typically taken daily, whereas full image copies are typically taken weekly.

DB2 provides the option for multiple image copies. You can create up to four identical image copies of a table space, one pair for local recovery use and one pair for off-site storage.

The size of the table spaces to be copied varies from a few megabytes to several gigabytes. The TS7700 Virtualization Engine solution is best for small and medium sized table spaces, because you need a higher bandwidth for large table spaces.

When a database is recovered from image copies, a full image copy and the subsequent incremental image copies need to be allocated at the same time. This can potentially tie up many tape drives and, in smaller installations, can prevent other work from being run. With one TS7700 Virtualization Engine, with its 256 virtual drives, this is not an issue.

The large number of tape drives is important also for creating DB2 image copies. Having more drives available allows you to run multiple copies concurrently and use the MERGECOPY DB2 utility without impact. An advisable solution is to run a full image copy of the DB2 databases once a week outside the TS7700 Virtualization Engine and run the incremental image copies daily using TS7700 Virtualization Engine. (The smaller incremental copy fits better with the TS7700 Virtualization Engine volume sizes.)

7.11.2 CICS and IMS

As with DB2, both CICS® and IMS use tapes to store logs and image copies of databases.

CICS is only a data communication product, whereas IMS has both the data communication and the database function (IMS-DL/1). CICS uses the same DL/1 database function to store its data.

CICS journals and IMS logs

CICS keeps track of database changes in its journal data sets. IMS keeps track of database changes in its online log data sets. After these data sets become full, both CICS and IMS offload the logs to tape.

CICS and IMS logs are sequential data sets. When offloading these logs to tape, you must request a scratch volume every time.

The logs contain the information necessary to recover databases and usually those logs are offloaded, as with DB2, in two copies, one local and one remote. You can write one local copy and then create the second for disaster recovery purposes later, or you can create the two copies in the same job stream.

With TS7700 Virtualization Engine, you can create the local copy directly on TS7700 Virtualization Engine virtual volumes, and then copy those volumes to non-TS7700 Virtualization Engine tape drives, or to a remote TS7700 Virtualization Engine.

Having a local copy of the logs written inside the TS7700 Virtualization Engine allows you faster recovery, because the data stays in the tape volume cache for some time.

When recovering a database, you can complete backout operations in significantly less time with the TS7700 Virtualization Engine, because when reading logs from tape, IMS uses the slow read backward operation (100 KBps) on real tape drives. With the TS7700 Virtualization Engine, the same operation is much faster, because the data is read from TVC. Lab measurements do not see much difference between read forward and read backward in a TS7700 Virtualization Engine; both perform much better than on physical drives. The reason is not just that the data is in the tape volume cache, but the TS7700 Virtualization Engine code also fully buffers the records in the reverse order that they are on the volume when in read backwards mode.

Another benefit TS7700 Virtualization Engine gives to recovery operations is the availability of up to 256 virtual drives per cluster, allowing you to mount several logs concurrently and therefore to back out the database to be recovered faster.

The IMS change accumulation utility is used to accumulate changes to a group of databases from several IMS logs. This implies the use of many input logs that will be merged into an output accumulation log. With the TS7700 Virtualization Engine, you can use more tape drives for this function.

Image copies

Image copies are backup copies of the IMS databases. IMS can create only full image copies. To create an image copy of a database, use a batch utility, copying one or more databases to tape.

With the TS7700 Virtualization Engine, you do not have to stack multiple small image copies to fill a tape cartridge. Using one virtual volume per database does not waste space, because the TS7700 Virtualization Engine then groups these copies into a stacked volume.

IMS, unlike DB2, has a batch function that works with databases through an IMS batch region. If you do not use logs when running an IMS batch region, then to recover the database, you must use an image copy taken before running the batch job. Otherwise, you can use logs and checkpoints, which allows you to restart from a consistent database image taken during the batch execution processing. Using TS7700 Virtualization Engine, you can access these image copies and logs at a higher speed.

The TS7700 Virtualization Engine volume stacking function is the best solution for every database backup, because it is transparent to the application and does not require any JCL procedure change.

7.11.3 Batch data

Other applications that write to tape and benefit from using the TS7700 Virtualization Engine are as follows:

- ▶ VSAM REPRO
- ▶ IEBGENER / IEBCOPY / ICETOOL
- ▶ DSS data set COPY or DUMP
- ▶ DFSMSrmm Tape Copy Tool (an IBM service offering)
- ▶ IBM Tivoli Tape Optimizer
- ▶ Any other tape copy utility

The amount of data from these applications can be huge if your environment does not use TMM or if you do not have DFSMSHsm installed. All such data benefits from using the TS7700 Virtualization Engine for output.

With TS7700 Virtualization Engine, the application can write one file per volume, using only part of the volume capacity and TS7700 Virtualization Engine takes care of completely filling the stacked cartridge for you, without JCL changes.

The only step you must remember is that if you need to move the data off site, you must address a device outside the local TS7700 Virtualization Engine, or use other techniques to copy TS7700 Virtualization Engine data onto other movable tapes, as described in 7.5.6, “Moving data out of the TS7700 Virtualization Engine” on page 383.



Part 3

Operation

In this part, we describe daily operations and the monitoring tasks that are related to the IBM Virtualization Engine TS7700 with R1.7. We also provide you with planning considerations and scenarios for disaster recovery testing.

This part contains the following chapters:

- ▶ Chapter 8, “Operation” on page 409
- ▶ Chapter 9, “Performance and monitoring” on page 589
- ▶ Chapter 10, “Failover and disaster recovery scenarios” on page 725



Operation

In this chapter, we describe operational considerations and usage guidelines unique to the IBM Virtualization Engine TS7700. For general guidance about how to operate the IBM System Storage TS3500 Tape Library, see the following publications:

- ▶ *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789
- ▶ *z/OS Object Access Method Planning, Installation and Storage Administration Guide for Tape Libraries*, SC35-0427

We provide information about how to operate the TS7700 Virtualization Engine by covering the following main topics:

- ▶ User interfaces
- ▶ IBM Virtualization Engine TS7700 management interface
- ▶ System-managed tape
- ▶ Basic operations
- ▶ Tape cartridge management
- ▶ Managing logical volumes
- ▶ Messages and displays
- ▶ Recovery scenarios
- ▶ IBM Virtualization Engine TS7720 considerations

8.1 User interfaces

To successfully operate your TS7700 Virtualization Engine, you must understand its concepts and components. In this chapter, we combine the components and functions of the TS7700 Virtualization Engine into two groups:

- ▶ The logical view
- ▶ The physical view

Each component and each function belongs to only one view.

The logical view is also the host view. From the host allocation point of view, there is only one library to deal with: the Composite Library. A Composite Library can have up to 1024 virtual addresses for tape mounts. The logical view includes virtual volumes and tape drives.

The host is only aware of the existence of the physical libraries because they are defined through ISMF in a z/OS environment. The term *distributed library* is used to denote the physical libraries and TS7700 Virtualization Engine components that are part of one cluster of the multi-cluster grid configuration. The physical view is the hardware view that deals with the hardware components of a Single Cluster or a multi-cluster grid configuration, with TS3500 Tape Libraries and 3592 J1A, TS1120, or TS1130 Tape Drives.

The operator interfaces for providing information about the TS7700 Virtualization Engine are:

- ▶ OAM commands are available at the host operator console, and these commands provide information regarding the TS7700 Virtualization Engine in stand-alone and Grid environments. This information represents the host view of the components within the TS7700 Virtualization Engine. Other z/OS commands can be used against the virtual addresses. These commands are not aware that the 3490E addresses are part of a TS7700 Virtualization Engine configuration.
- ▶ Web Specialist functions are available through web-based user interfaces:
 - You can access the different web interfaces with Microsoft Internet Explorer Version 6.0 or a fully compatible alternative browser with JavaScript and Java™ enabled.
 - There are two specialists available for tape library management:
 - The TS3500 Tape Library Specialist allows for management (configuration and status) of the TS3500 Library.
 - The TS7700 Virtualization Engine *management interface* (MI) is used to perform all TS7700 Virtualization Engine related configuration, setup, and monitoring actions.

For further information about using the operator windows, see *IBM TotalStorage Library Operator Manual*, GA32-0280 or *IBM System Storage TS3500 Tape Library Operator Guide*, GA32-0560.

In this chapter, we focus on the interfaces related to the operation of the TS7700 Virtualization Engine. For detailed information about the related TS3500 Tape Library operational tasks, see the respective operator's guide or to *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789.

8.1.1 TS3500 Tape Library Specialist (TS7740 Virtualization Engine only)

The IBM System Storage TS3500 Tape Library Specialist (TS3500 Tape Library Web specialist) interface lets you perform many library functions from the web.

Figure 8-1 shows the TS3500 Tape Library Specialist Welcome window with System Summary, where you can choose to view the status of complete library.



Figure 8-1 TS3500 Tape Library Specialist Welcome window

Figure 8-2 shows a flowchart of the functions that are available depending on the configuration of your TS3500 Tape Library.

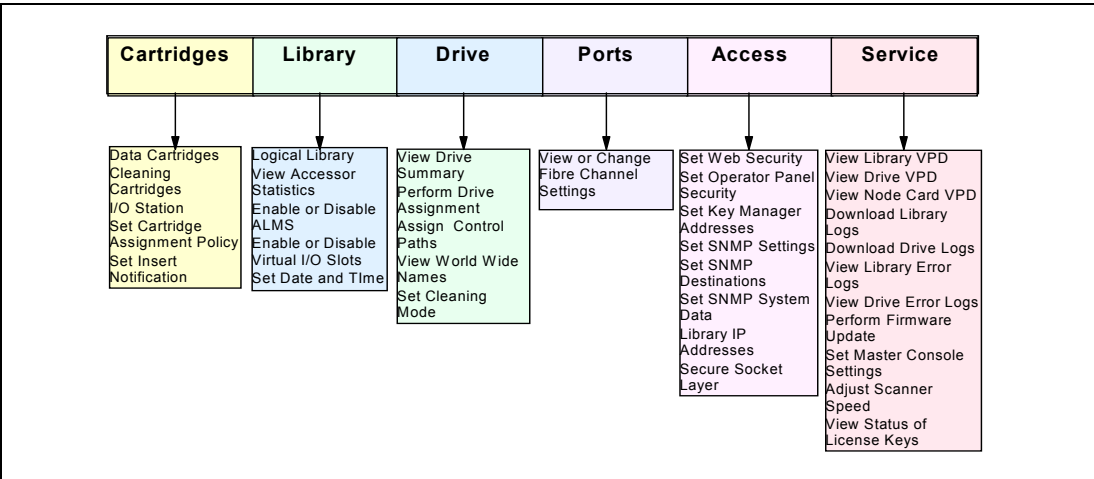


Figure 8-2 TS3500 Tape Library Specialist functions

8.1.2 Call Home (TS7740 Virtualization Engine with library attachment only)

The TS3500 Tape Library include several external interfaces not directly associated with data paths; rather, these interfaces are associated with library control, service, and status information. They support customer interaction and feedback, and attachment to IBM remote support infrastructure for product service and support.

The Call Home function generates a service alert automatically when a problem occurs with one of the following components:

- ▶ TS3500 Tape Library
- ▶ 3592 Model J70 and TS1120 Model C06 controller
- ▶ TS7700 Virtualization Engine

Error information is transmitted to the IBM System Storage TS3000 System Console for service, and then to the IBM Support Center for problem evaluation; the IBM Support Center can dispatch an IBM System Services Representative (SSR) to the client installation. Call Home can send the service alert to a pager service to notify multiple people, including the operator. The SSR can deactivate the function through service menus, if required.

8.1.3 Electronic Customer Care

Here we provide an overview of the communication path between a customer's purchased subsystems and IBM support. This method has been used for Call Home since the introduction of the TotalStorage System Controller (TSSC). It was previously known as the TotalStorage Master Console (TSMC).

The TSSC console uses analog phone lines and a modem to connect to the IBM Remote Technical Assistance Information Network, better known as RETAIN®. The code running in RETAIN then decides what to do with the information. A Problem Management Record (PMR) will be opened in the case of a problem. After the PMR is created, RETAIN automatically consults with the IBM Knowledge Based Systems (RKBS) to add information pertinent to the reported problem. Finally, in the case where RETAIN detects that the call home is a data package, the data is forwarded to catchers that move the data to an IBM Internal Server called DFS Cell. From there, it is pulled into the IBM Tape Call Home Database.

Figure 8-3 on page 413 describes the Electronic Customer Care Call Home function. We define three security zones:

- ▶ The *Red Zone* is defined as the customer data center. This zone is where the IBM Storage subsystem and the TSSC reside.
- ▶ The *Yellow Zone* is defined as the open Internet. This is open to all outside communication.
- ▶ The *Blue Zone* is defined as IBM Support. This sits inside the IBM intranet and is only accessible to IBM-authenticated users.

In Electronic Customer Care, the TSSC will use either a modem connection or a broadband connection to connect to an Electronic Customer Care gateway located in the Yellow Zone. This is an IBM-managed server used as a relay for IBM Support information. This gateway then forwards the information to the Inter Enterprise Process Directory (IEPD). IEPD is located within IBM Support in the Blue Zone. IEPD then checks the problem report that has been submitted and monitors from which system the report has come. The system's credentials are then checked to make sure a valid support contract exists. If the credentials pass for this system, it will consult Technical Services Knowledge Base System (TSKBS) for known information to add to the problem report. The problem report will then be used to open a Problem Management Record (PMR) in IBM RETAIN.

In the event of a data call home, the data is sent from the same TSSC connection to an IBM-managed server located in the Yellow Zone known as Testcase. Dumpers monitor this server for new information. When they see this new information, they move the data package to DFS space where it gets pulled into the RMSS Call Home Database.

Initial ECC handshaking communication utilizes HTTP communication. After initial handshaking, all information is sent using secure HTTPS communication. Because all traffic is outbound in nature and only uses the HTTP and HTTPS ports, connectivity should work through most customers' firewall without additional firewall rules being added by the customer.

Problem Reporting communication is then sent to IEPD, which consults Technical Services Knowledge Base System for the systems listed, and opens a PMR in RETAIN. Of course, all of these details are for informational purposes only. After initial setup, these details are utilized in the background, without knowledge of the user.

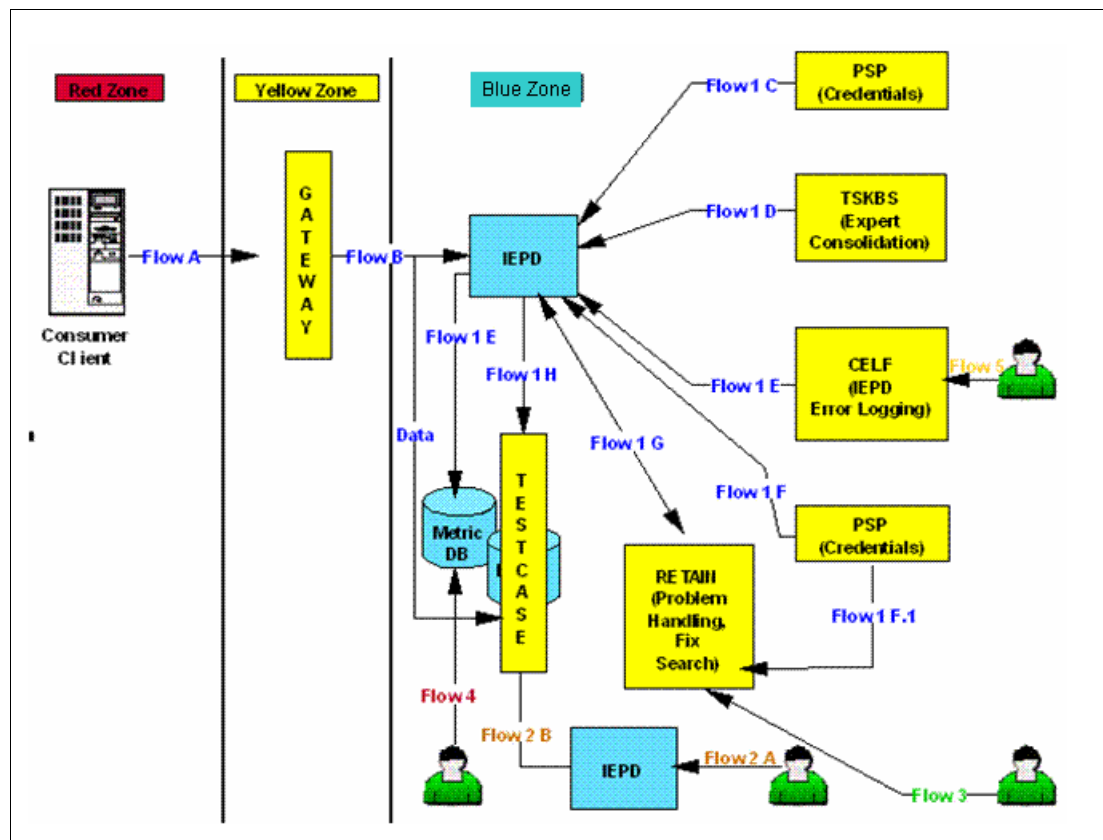


Figure 8-3 Electronic Customer Care

By default, all inbound connections by IBM service personnel are still through a dial-in modem connection.

The outbound communication associated with ECC call home can be through a customer-supplied Ethernet connection, a modem, or both in the form of a failover setup. The local subnet LAN connection between the TSSC and the attached subsystems remains the same; it is still isolated without any outside access.

ECC adds another Ethernet connection to the TSSC, bringing the total number to three. These connections are labeled:

- ▶ The External Ethernet Connection, which is the ECC Interface.
- ▶ The Grid Ethernet Connection, which is used for the TS7740 Virtualization Engine Autonomic Ownership Takeover Manager (AOTM).
- ▶ The Internal Ethernet Connection, used for the local attached subsystems subnet.

All of these connections are set up using the Console Configuration Utility User Interface located on the TSSC.

Tivoli Assist On-site

Enhanced support capabilities include the introduction of Tivoli Assist On-site (AOS) to augment maintenance capabilities. This is a service function managed by an IBM System Services Representative (SSR). AOS is a tool that allows an authenticated session for remote TSSC desktop connections over an external broadband ethernet adapter. An AOS session allows IBM remote support center representatives to troubleshoot issues with the machine.

AOS utilizes the same network as broadband call home and will work on either HTTP or HTTPS. With TS7700 Virtualization Engine R1.7 and later the AOS function is disabled by default. When enabled, the AOS can be configured to run in either attended or unattended modes.

- ▶ Attended mode requires that the AOS session be initiated at the TSSC console associated with the target TS7700 Virtualization Engine. This will require physical access by the IBM SSR to the TSSC.
- ▶ Unattended mode, also called Lights Out mode, allows a remote support session to be establish without manual intervention at the TSSC console associated with the target TS7700 Virtualization Engine.

All AOS connections are outbound. In unattended mode the session is establish by periodically connecting to regional AOS relay servers to determine if remote access is needed. If access has been requested, AOS will authenticate and establish the connection allowing a remote desktop access to the TSSC.

8.2 TS7740 Virtualization Engine management interface

The TS7740 Virtualization Engine management interface (MI) is the primary interface to monitor and administer the TS7700 Virtualization Engine.

8.2.1 Connecting to the MI

To connect to the TS7700 Virtualization Engine management interface, perform the following steps:

1. The TS7700 Virtualization Engine must first be installed and configured.
2. In the address field of a supported web browser, enter `http://x.x.x.x/Console` (where `x.x.x.x` is the virtual IP address that assigned during installation). Press the Enter key on your keyboard or click the **Go** button in your web browser. The virtual IP is one of three customer IP addresses given during installation; see 3.2.3, “TCP/IP configuration considerations” on page 128.

If you are using your own name server, where you can associate a name with the virtual IP address, you can use the name instead of the hardcoded address for reaching the MI.

3. The login page for the management interface loads, as shown in Figure 8-4. Enter the default login name as `admin` and the default password as `admin`.

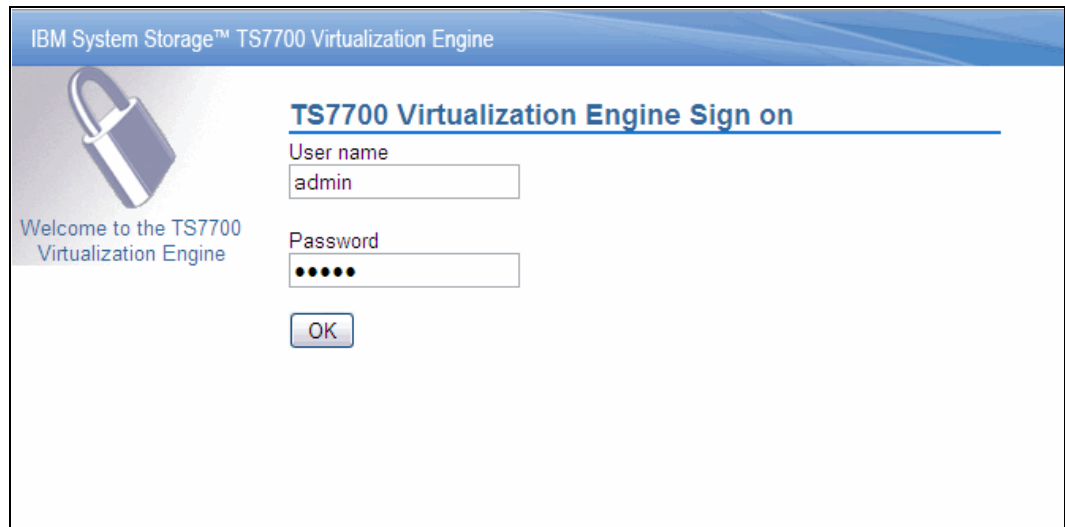


Figure 8-4 TS7700 Virtualization Engine management interface login

After entering your password, you see the first web page presented by the MI, the Virtualization Engine Grid Summary, as shown in Figure 8-5 on page 416.

After security policies have been implemented locally at the TS7700 Virtualization Engine cluster or through the use of centralized Role-Based Access Control (RBAC) a unique user identifier and password can be assigned by the administrator. The user profile can be modified to only provide functions applicable to the role of the user. All users may not have access to the same functions or views through the management interface.

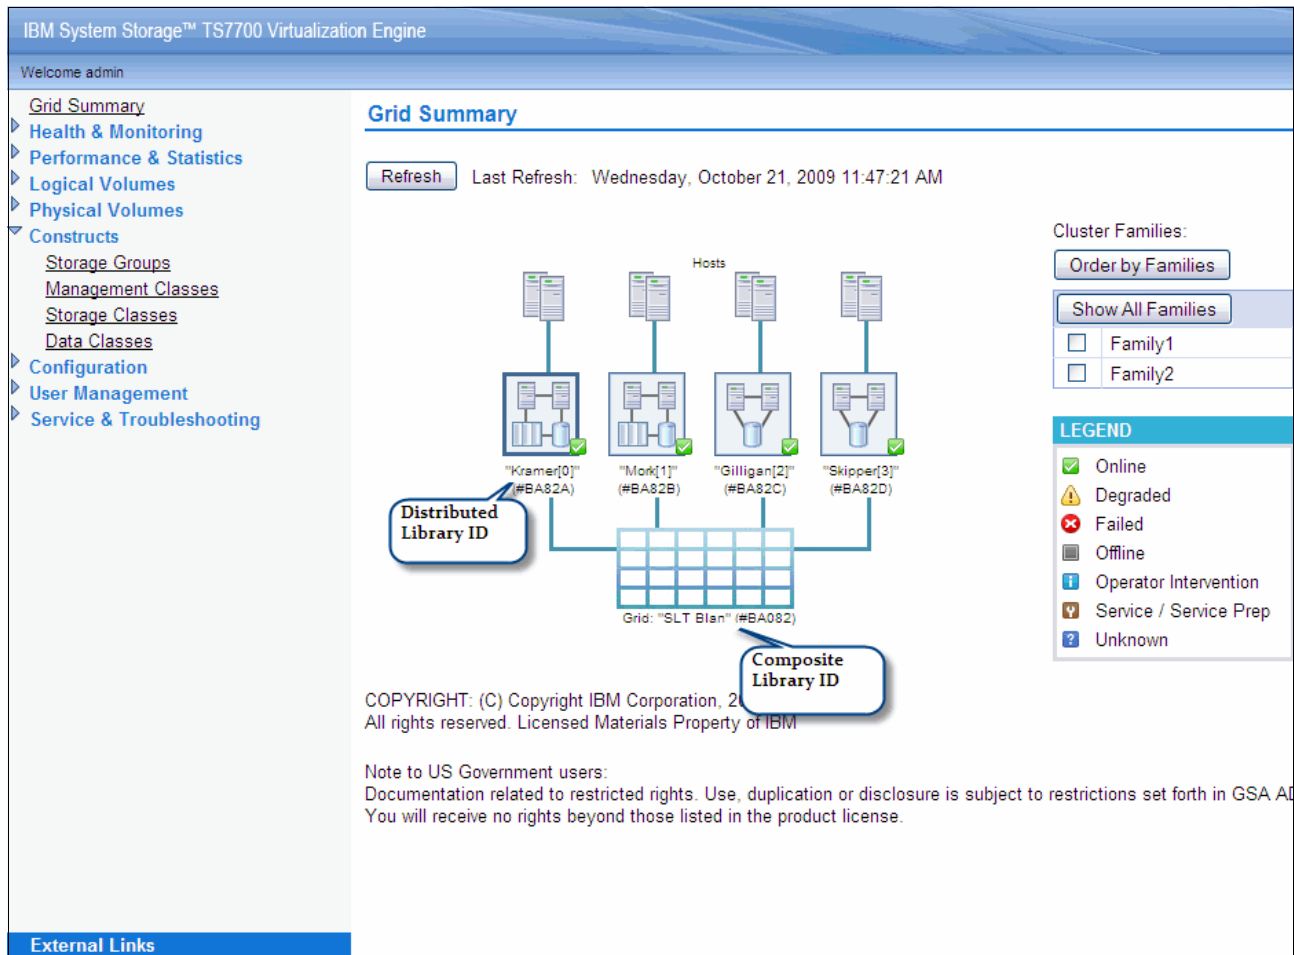


Figure 8-5 MI Virtualization Engine Grid Summary

Figure 8-5 shows a visual summary of the TS7700 Virtualization Engine Grid, its components, and health status. The TS7700 Virtualization Engine Grid is depicted as a blue grid on a white background. Attached to the Grid are individual TS7700 Virtualization Engine Clusters, each connected to a host.

Each cluster is represented by a single blue square icon and is named using the cluster's nickname, cluster ID, and distributed library sequence number. The cluster that you are currently connected to is identified by a dark blue border.

The health of the system is checked and updated automatically at times determined by the TS7700 Virtualization Engine. Data loaded on the page is not real time; the Last Refresh field reports the date and time the displayed data was retrieved from the TS7700 Virtualization Engine. To populate the summary with an updated health status, click **Refresh**. This operation can take some time to complete.

The health status of each cluster is indicated by a status sign affixed to its icon; the Legend explains the meaning of each status sign. To obtain additional information about a specific cluster, click that component's icon.

Component health status

The following terms describe the health status of clusters in a TS7700 Virtualization Engine Grid:

- ▶ Normal: The connection between the Grid and cluster is normal.
- ▶ Degraded: The connection between the Grid and cluster is working, but one of the component's redundant parts has stopped functioning.
- ▶ Failed: The connection between the Grid and cluster has failed.
- ▶ Offline; A cluster can be considered offline because of any of the following conditions: being in Service or Service Prep mode, all vNodes are offline, an active hNode is offline, or a cluster is in a pending online state.
- ▶ Service/Service Prep; A cluster is in the Service or Service Prep mode.

Customer network interfaces

The TS7700 Virtualization Engine provides two Ethernet connections to the customer's network for access to the TS7700 Virtualization Engine management interface (MI). Access to the functions provided by the TS7700 Virtualization Engine MI are secured with a user ID and password. There are two security policy methodologies available. The current method uses an internal, locally administered security policy management function. A centrally managed, role-based access control method using LDAP is implemented in TS7700 Virtualization Engine R.16.

Locally administered security policy management

Locally administered policy management is unique to each TS7700 Virtualization Engine. This method of managing user access has been available with the TS7700 Virtualization Engine MI from the initial product release.

The following user security policy settings are available with locally administered security policy management:

- ▶ Disable account expiration: User accounts will not expire.
- ▶ Enable account expiration: Allow accounts to expire after a set number of days from when their passwords are set. The amount of days can be set in the "Password expires after" text box.
- ▶ Disable account lockout: User accounts will not be locked out when a wrong password is entered.
- ▶ Enable account lockout: The account will be allowed a certain number of failed logon attempts before locking the account out of logging into the management interface. The number of failed attempts is entered in the "Number of successive incorrect password attempts allowed" text box.

Note: To unlock an account, an administrator can modify the user account and change the password from the Manage Users page.

Centralized Role-Based Access Control

Centralized Role-Based Access Control (RBAC) was introduced with TS7700 Virtualization Engine Release 1.6. RBAC provides the capability to manage user access from a central control point. This implementation also includes the capability to log user actions for audit purposes. RBAC is an effective method for managing a large number of subsystems, because the user profile is set up once for all subsystems. RBAC uses an IBM System Storage Productivity Center (SSPC). The SSPC forwards the request to a customer Lightweight Directory Access Protocol (LDAP) server for ID authentication and role validation.

An authenticated user ID will be associated with a defined role for accessing functions provided by the TS7700 Virtualization Engine subsystem MI.

Library Control with TS7700 Virtualization Engine management interface

The TS770 management interface also control the library operations.

In environments where the tape configuration is separated from the LAN-attached hosts or web clients by a firewall, these are the only ports that must be opened on the firewall; all others might be closed. See Table 8-1 for more information.

Table 8-1 Customer network interface firewall

Function	Port	Direction (from library)	Protocol
TS3500 Tape Library Specialist	80	Inbound	TCP/IP
SNMP Traps	161/162	Bi-directional	UDP/IP
Encryption Key Manager	1443	Outbound	SSL
Encryption Key Manager	3801	Outbound	TCP/IP

For more information about the customer network interface, see the following publications:

- ▶ *IBM System Storage TS3500 Tape Library Operator Guide*, GA32-0560
- ▶ *IBM System Storage TS3500 Tape Library Introduction and Planning Guide*, GA32-0559
- ▶ *IBM System Storage TS1120 Tape Drive and Controller Operator Guide*, GA32-0556
- ▶ *IBM System Storage TS1120 Tape Drive and Controller Introduction and Planning Guide*, GA32-0555
- ▶ *IBM Encryption Key Manager component for the Java platform Introduction, Planning, and User's Guide*, GA76-0418
- ▶ *IBM Virtualization Engine TS7700 Series Introduction and Planning Guide*, GA32-0567

8.2.2 Using the management interface

Before we explain in detail the tasks that you can perform from the TS7700 Virtualization Engine management interface, we describe some common page and table elements.

Standard page elements

The following elements might be present on a page and allow you to use the management interface more effectively:

Refresh button

Clicking **Refresh** at the top of the page refreshes the page's data. The date and time of the page's last refresh are displayed next to the Last refresh field.

Current cluster being accessed

For cluster-specific pages, the current cluster being accessed is displayed at the top of the page. Clicking the graphic or hyperlink for the cluster opens the Cluster Summary page. If you have a multi-cluster Grid, there will be an option to change the active cluster by selecting a new cluster in the drop-down menu of the current cluster being accessed.

Standard table elements

Tables in the management interface contain a number of features that allow for a variety of filtering and sorting options. These options are selectable through the Select Action drop-down menu and through buttons above the table header.

Edit Sort	Selecting this option allows you to select primary, secondary, and tertiary sorting categories, and whether the sort will be ascending or descending.
Clear All Sorts	Removes any current table sorting.
Show/Hide Filter Row	Shows or hides the Filter option available under a table header. Selecting Filter displays the option to filter that table column by entering a Condition and Text.
Clear All Filters	Clears any current filtering in place.

Note: To support Japanese input, a Japanese front-end processor needs to be installed on the computer where a web browser is accessing the management interface.

Standards compliance

The storage management interface (MI) for the TS7700 Virtualization Engine complies with industry standards that make it easier to manage devices from different manufacturers.

The implementation of storage management standards allows users to manage storage devices from different manufacturers. The TS7700 Virtualization Engine complies with the Storage Management Initiative - Specification (SMI-S) standard. The Storage Networking Industry Association (SNIA) promotes SMI-S as a means to standardize interfaces between management software and storage devices obtained from different manufacturers.

SMI-S incorporates two components:

- ▶ **Web-Based Enterprise Management:** Web-Based Enterprise Management (WBEM) is a set of management and Internet standard technologies developed to unify the management of distributed computing environments.
- ▶ **Common Information Model:** The Common Information Model (CIM) is the data model for WBEM. It provides a common definition of management information for systems, networks, applications, and services, and allows for vendor extensions.

In 8.2.3, “Health & Monitoring (TS7740 Virtualization Engine with Library attachment)” on page 419 through 8.2.10, “Service and troubleshooting” on page 520, we explain the MI windows that enable you to control and monitor the TS7700 Virtualization Engine.

For details about the Performance and Statistics selections, see Chapter 9, “Performance and monitoring” on page 589.

8.2.3 Health & Monitoring (TS7740 Virtualization Engine with Library attachment)

This section of the TS7740 Virtualization Engine management interface gives you the choice to query the current cluster status, and acquire logical and physical volume details. You can access the following selections if you navigate within the TS7700 Virtualization Engine management interface to the Health & Monitoring section:

- ▶ Cluster Summary
- ▶ Operator Interventions
- ▶ Physical Tape Drives

- ▶ Virtual Tape Drives
- ▶ Physical Media Inventory
- ▶ Tape Volume Cache
- ▶ Physical Library
- ▶ Operation History

Cluster Summary

Use this window (Figure 8-6) for a visual summary of the IBM Virtualization Engine TS7740 Virtualization Engine Cluster.

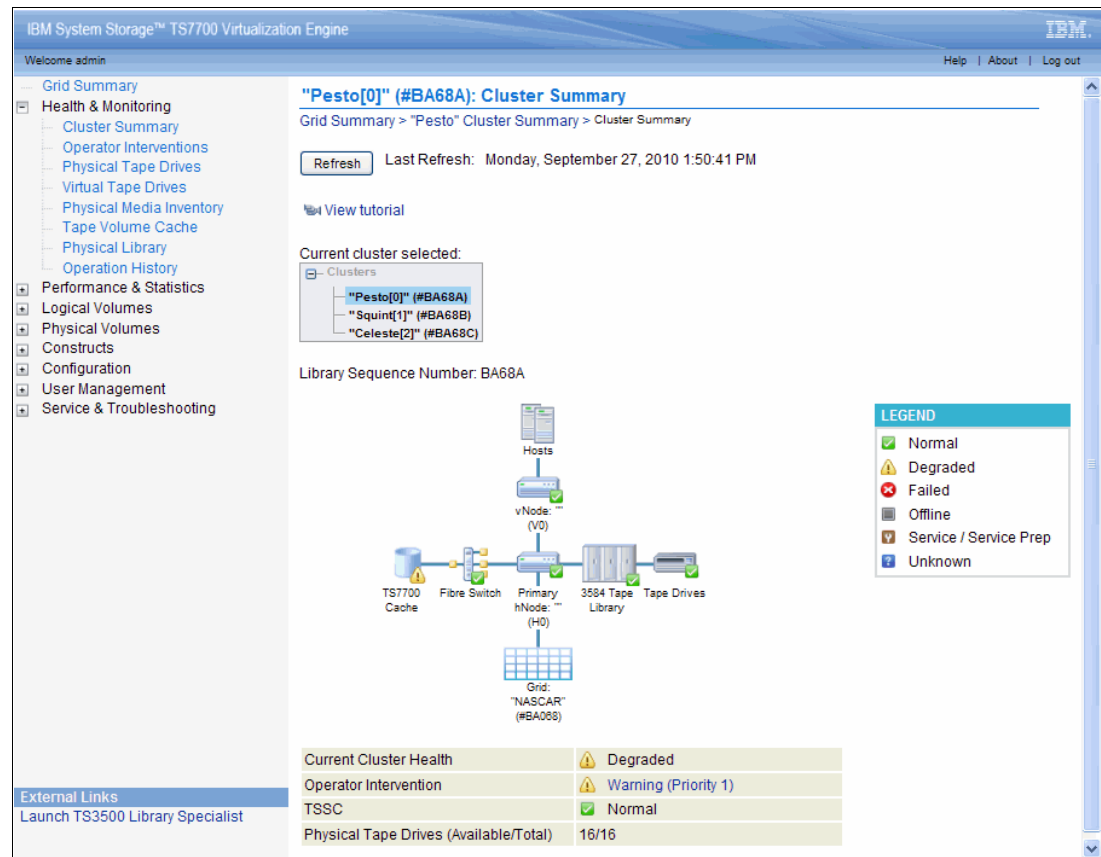


Figure 8-6 TS7700 Virtualization Engine management interface Cluster Summary

To watch a tutorial showing how to monitor system status, click the **View tutorial** link.

Following the Current Component Selected graphic, the cluster summary includes these cluster identifiers:

- ▶ Library Sequence Number: The library sequence number for this cluster.
- ▶ Management Interface IP Address: The IP address for the management interface, as defined in the Cluster Network Settings window. Click the IP address to display the Cluster Network Settings window where network settings for the cluster can be viewed or altered.

This window also contains a visual summary of the TS7700 Virtualization Engine Cluster selected from the Virtualization Engine Grid Summary window. The summary will show the cluster, its key components, such as the vNode, hNode, tape library and drives, and their connectivity. If the cluster does not possess a physical library, this graphic will not display a physical library, physical drives, or Fibre Channel Switch. The health of the system is checked and updated automatically at times determined by the TS7700 Virtualization Engine.

The status of each component is indicated by a status sign affixed to its icon; the Legend explains the meaning of each status sign. To obtain information about a failed or degraded component and how to correct its state, click the component's icon. To obtain more information concerning details of either the vNode or hNode, click the picture of the node you are interested in; the Cluster Nodes window will open. Cluster components are connected by a solid line, even if the component is offline. Click the picture of the Grid to return to the Virtualization Engine Grid Summary window.

Cluster components

The following components are displayed in the cluster summary (Figure 8-6 on page 420).

Components can be in a normal, failed, degraded, or offline state, as indicated by the status sign attached to each component's icon. Additionally, in preparation for maintenance activities, the vNode and hNode can also go into Service Prep or Service modes. A degraded state indicates that a component is working but one of its redundant parts has stopped functioning.

- ▶ vNode: vNode on the cluster.
- ▶ hNode: hNode located on the cluster.
- ▶ Disk cache: Disk cache located on the cluster.
- ▶ Fibre switch: Fibre Switch connecting the tape library and the cluster.
- ▶ Tape Library: Tape library connected to the cluster.
- ▶ Tape Drives: Physical drives located in the tape library.

Note: The Fibre Channel Switch, Tape Library, and Tape Drives image will not be visible if the cluster does not possess a physical library.

Table fields

The following information can be found at the bottom of the window (Figure 8-6 on page 420) below the cluster diagram.

- ▶ Cluster Health State: The health state of the cluster as a whole. Possible values are:
 - Normal
 - Degraded
 - Failed
 - Offline
 - Service or Service Prep
 - Unknown
- ▶ Virtual Tape Drives (Licensed/Total): Amount of virtual drives licensed for the cluster and total amount able to be licensed for the cluster. The amount of virtual drives can be increased through activating feature licenses.
- ▶ Physical Tape Drives (Available/Total): Available and total physical drives for this cluster.

Note: This image is not visible if the cluster does not possess a physical library.

- ▶ Operator Interventions: Indicates whether Operator Interventions is still active and actions have to be taken. Clicking this item will open the window shown in Figure 8-7 on page 422.

Operator Interventions

Use this window (Figure 8-7) to view all operator interventions performed on a given TS7700 Virtualization Engine Cluster.

The Send Future Interventions to Host function allows any interventions that occur to be sent to the host during the course of operations. Click the **Disable** or **Enable** toggle buttons to change the current setting for this function.

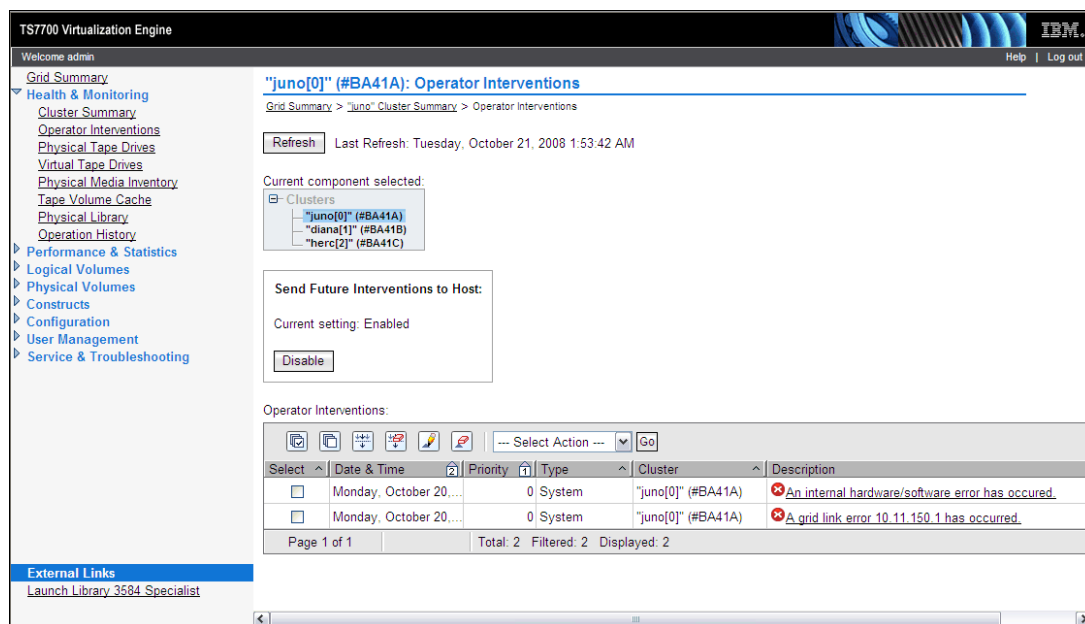


Figure 8-7 MI Operator Interventions

The Operator Interventions table

The Operator Interventions table displays the details of each operator intervention on the cluster.

You can use the Operator Interventions table to view, sort, and filter the details of operator interventions. You can also clear one or more intervention entries to remove them from this table.

Status information displayed in the Operator Interventions table includes:

- ▶ **Message ID:** The operator intervention code number that identifies the reason for the intervention and actions to take. Click this number to display additional information about the intervention.
- ▶ **Date & Time:** The date and time the intervention occurred.
- ▶ **Priority:** The priority of the intervention, on a scale of 0-2, where 0 is the highest priority. Priority is also indicated by an icon appearing at the beginning of the Description field.
 - **Priority 0:** The highest priority operator intervention. The representative icon is a red circle containing a white X.
 - **Priority 1:** The second-highest priority operator intervention. The representative icon is a yellow triangle containing a black exclamation point (!).
 - **Priority 2:** The lowest priority operator intervention. The representative icon is a blue square containing a white, lowercase letter i.

- ▶ **Type:** The type of intervention occurring. Possible values include:
 - **System:** An intervention that can either be cleared by you, or will be cleared automatically by the system when the condition causing the intervention has been resolved.
 - **User:** An intervention that you must clear after resolving the condition causing the intervention.
- ▶ **Cluster:** The cluster receiving the operator intervention.
- ▶ **Description:** A detailed description of the operator intervention action. To view the entire list of instructions that comprise this intervention, click the text of the description in this column. An icon at the beginning of this field denotes the intervention's priority. See the Priority status description for definitions of these icons.

Use the drop-down menu on the Operator Interventions table to control the appearance of data in the table or clear one or more operator intervention entries.

To clear an operator intervention, check the check box in the row of the intervention you want to clear, then select **Clear** from the drop-down menu and click **Go**. If you clear a system intervention, you will be presented with a confirmation window asking you to confirm the intervention to be cleared. Click **Yes** to clear the entry or **No** to retain the entry on the table.

Use the Select Table Action section of the drop-down menu to select or deselect multiple rows or control the appearance of the table through sort and filter actions.

Clicking **Operator Intervention Description** will give you the correct actions for the selected intervention, as shown in Figure 8-8.

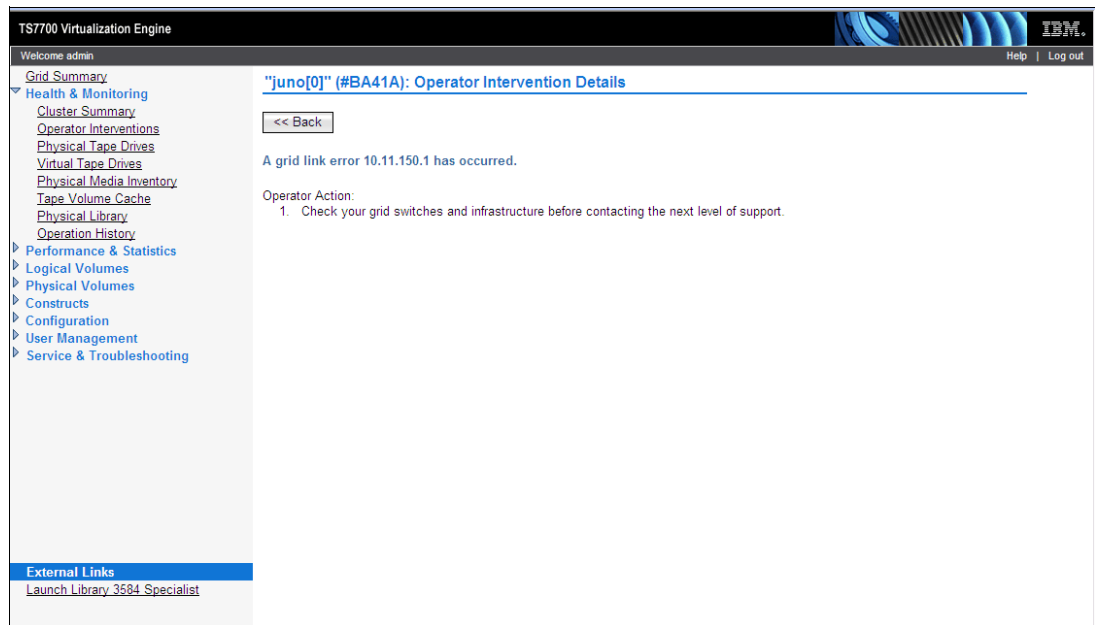


Figure 8-8 MI Operator Intervention Details

Physical Tape Drives

Use this window (Figure 8-9) to view the status of physical tape drives accessible to the TS7700 Virtualization Engine Cluster.

Note: If the cluster does not possess a physical library, the Physical Tape Drives window will not be visible on the TS7700 Virtualization Engine management interface.

Whenever you need to know specific details about a physical drive, such as its serial number, the drive type, the drive format, the drive encryption, or the drive location by Frame and Row inside library, this window allows you to see the desired information, as shown in Figure 8-9.

The screenshot shows the TS7700 Virtualization Engine management interface. The sidebar on the left contains navigation links under 'Health & Monitoring', including 'Physical Tape Drives'. The main content area is titled '"juno[0]" (#BA41A): Physical Tape Drives'. It includes a 'Refresh' button and a 'Last Refresh' timestamp. Below this is a 'Current component selected:' section showing a tree view with 'Clusters' expanded, listing '"juno[0]" (#BA41A)', '"diana[1]" (#BA41B)', and '"herc[2]" (#BA41C)'. A message states: 'The health status of the physical tape drives can only be retrieved every 20 minutes. The health status of the physical tape drives cannot be retrieved when the library is paused.' Below this is a 'Physical Tape Drives:' section with a table. The table has columns: Select, Serial N..., Drive Ty..., Drive Fo..., Encrypti..., Online, Frame, Row, and Health. The table contains six rows of data, all with a 'Health' status of 'OK'.

Select	Serial N...	Drive Ty...	Drive Fo...	Encrypti...	Online	Frame	Row	Health
	0000078988...	3592E05E	E05	Enabled	Yes		12	OK
	0000078984...	3592E05E	E05	Enabled	Yes		11	OK
	0000078968...	3592E05E	Not Available	Enabled	Yes		2	OK
	0000078969...	3592E05E	Not Available	Enabled	Yes		1	OK
	0000078985...	3592E05E	Not Available	Enabled	Yes		10	OK
	0000078987...	3592E05E	Not Available	Enabled	Yes		9	OK

Figure 8-9 TS7700 Virtualization Engine management interface Physical Tape Drives

Note: As stated in the middle of the window, "The health status of the physical tape drives can only be retrieved every 20 minutes. The health status of the physical tape drives cannot be retrieved when the library is paused".

You can sort the information by a single column heading by clicking the column heading you want to sort by; you may also sort using multiple column headings by using the Edit Sort table action. To access the Edit Sort action, use the pencil icon at the top of the Physical Drives table, or select **Edit Sort** from the drop-down menu and click **Go**. The Edit Sort action allows you to choose up to three columns of information by which to sort the Physical Drives table; you can also select whether your sort criteria are displayed in ascending or descending order. To clear an earlier sort, click the eraser icon at the top of the Physical Drives table or select **Clear All Sorts** from the drop-down menu and click **Go**.

Status information displayed in the Physical Drives table includes:

- ▶ Drive Serial Number: The serial number of the physical drive.
- ▶ Drive Type: The machine type and model number of the drive. Possible values are:
 - 3592J1A
 - 3592E05
 - 3592E05E (A 3592 E05 drive that is encryption capable.)
 - 3592E06
- ▶ Drive Format: The format the drive is operating in. Possible values are:
 - J1A
 - E05
 - E05E (The drive is operating with E05E encrypted data.)
 - E06
 - E06E (The drive is operating with E06E encrypted data.)
- ▶ Not Available: The format is unable to be determined because there is no physical media in the drive or the media is being erased.
- ▶ Unavailable: The format is unable to be determined because the health and monitoring checks have not yet completed. Refresh the current page to determine whether the format state has changed. If the Unknown state persists for one hour or longer, contact your IBM System Service Representative.
- ▶ Encryption: Whether encryption is enabled or disabled on the drive.

Note: If you are monitoring this field while changing the encryption status of a drive, the new status will not display until you bring the TS7700 Virtualization Engine Cluster offline and then return it to an online state.

- ▶ Online: Whether the drive is currently online to the TS7740 Virtualization Engine.
- ▶ Drive Health: The health of the physical drive. This value is obtained automatically at times determined by the TS7700 Virtualization Engine. Possible values are:
 - OK: The drive is fully functional.
 - WARNING: The drive is functioning but reporting errors. Action should be taken to correct the errors.
 - DEGRADED: The drive is functioning but at lesser redundancy and performance.
 - FAILURE: The drive is not functioning and immediate action should be taken to correct the fault.
 - OFFLINE/TIMEOUT: The drive is out of service or could not be reached within a certain time frame.

To view additional information pertaining to a specific drive in a cluster, in the Physical Drives table, click the radio button from the Select column that appears in the same row as the Serial Number of the drive. From the drop-down menu, select **Details** and click **Go** to display the Physical Drives Details table, as shown in Figure 8-10.

Physical Drive Details:	
Serial Number:	000007898825
Drive Type:	3592E05E
Drive Format:	E05
Encryption Enabled:	Yes
Encryption Capable:	Yes
Online:	Yes
Role:	Reclaim Source
Frame:	1
Row:	12
WWNN:	500507630f08bb0c
Pool:	2
Physical Volume:	320071
Logical Volume:	174125
Health:	OK

Figure 8-10 MI Physical Tape Drive details

In addition to the Serial Number, Drive Type, Drive Format, Encryption, Online, and Health information shown in the Physical Drives table, the Physical Drives Details table also displays the following status and identity information, as shown in Figure 8-9 on page 424:

- Encryption Enabled: Whether encryption is enabled on the drive.

Note: If you are monitoring this field while changing the encryption status of a drive, the new status will not display until you bring the TS7700 Virtualization Engine Cluster offline and then return it to an online state.

- Encryption Capable: Indicates whether the drive is capable of encryption.
- Role: The current role the drive is performing. Possible values are as follows:
 - IDLE: The drive is currently not in use.
 - MIGRATION: The drive is being used to copy a logical volume from the tape volume cache to a physical volume.
 - RECALL: The drive is being used to recall a logical volume from a physical volume to the tape volume cache.
 - RECLAIM SOURCE: The drive is being used as the source of a reclamation operation.
 - RECLAIM TARGET: The drive is being used as the target of a reclamation operation.
 - SECURE ERASE: The drive is being used to erase expired volumes from the physical volume securely and permanently.
- Frame: The frame in which the drive resides.
- Row: The row of the frame in which the drive resides.
- WWNN: The World Wide Node Name that is used to locate the drive.

- **Physical Pool:** The pool name of the physical volume mounted by the drive. This field will be blank if the drive is IDLE.
- **Physical Volume:** VOLSER of the physical volume mounted by the drive. This field will be blank if the drive is IDLE.
- **Logical Volume:** VOLSER of the logical volume being processed by the drive. This field will be blank if the drive is IDLE or if media is being erased.

Virtual Tape Drives window

Use this window to view a summary of the state of all virtual tape drives in a TS7700 Virtualization Engine Cluster, component, or Grid, as shown in Figure 8-11. This window also provides access to the Stand-Alone mount and demount functions for logical volumes.

The Virtual Tape Drives table displays status information for all drives accessible by the cluster, while the Virtual Drives Details table displays additional information for a specific, selected drive.

You can sort the information by a single column heading by clicking the column heading you want to sort by; you may also sort using multiple column headings by using the Edit Sort table action. To access the Edit Sort action, use the pencil icon at the top of the Virtual Drives table, or select **Edit Sort** from the drop-down menu and click **Go**. The Edit Sort action allows you to choose up to three columns of information by which to sort the Virtual Drives table; you can also select whether your sort criteria are displayed in ascending or descending order. To clear an earlier sort, click the eraser icon at the top of the Virtual Drives table or select **Clear All Sorts** from the drop-down menu and click **Go**.

The screenshot shows the 'Virtual Tape Drives' window in the TS7700 Virtualization Engine. The window title is '"Juno[0]" (#BA41A): Virtual Tape Drives'. Below the title, there is a 'Refresh' button and a 'Last Refresh' timestamp: 'Tuesday, October 21, 2008 5:03:21 AM'. The 'Current component selected' section shows a tree view with 'Clusters' expanded, listing 'Juno[0]' (#BA41A), 'diana[1]' (#BA41B), and 'herc[2]' (#BA41C). Below this, there is a 'Virtual Drives' section with a table of drives. The table has columns: Select, Address, Path, Logical, Online, Mounted, Role, and Standalone. The table contains 8 rows of data for drives vtd00 through vtd07. The 'Select' column has radio buttons for each row. The 'Path' column shows the path for each drive, all starting with '"Juno[0]" (#BA41A)/V[0]'. The 'Logical' column shows the logical volume number (e.g., 203586, 113162, etc.). The 'Online' column shows 'Yes' for all drives. The 'Mounted' column shows 'No' for all drives. The 'Role' column shows 'idle' for all drives. The 'Standalone' column shows 'No' for all drives.

Select	Address	Path	Logical	Online	Mounted	Role	Standalone
<input type="radio"/>	vtd00	"Juno[0]" (#BA41A)/V[0]	203586	Yes	No	idle	No
<input type="radio"/>	vtd01	"Juno[0]" (#BA41A)/V[0]	113162	Yes	No	idle	No
<input type="radio"/>	vtd02	"Juno[0]" (#BA41A)/V[0]	161990	Yes	No	idle	No
<input type="radio"/>	vtd03	"Juno[0]" (#BA41A)/V[0]	229163	Yes	No	idle	No
<input type="radio"/>	vtd04	"Juno[0]" (#BA41A)/V[0]	146083	Yes	No	idle	No
<input type="radio"/>	vtd05	"Juno[0]" (#BA41A)/V[0]	113204	Yes	No	idle	No
<input type="radio"/>	vtd06	"Juno[0]" (#BA41A)/V[0]	202976	Yes	No	idle	No
<input type="radio"/>	vtd07	"Juno[0]" (#BA41A)/V[0]	W68747	Yes	No	idle	No

Figure 8-11 MI Virtual Tape Drives

Status information displayed in the Virtual Tape Drives table includes:

- **Address:** The virtual drive address. This address takes the format vtdXX, where X is a hexadecimal number.
- **Path:** The path to the virtual drive. This path is used to show the location of the drive when viewing drive status at the component or Grid level. Path information is displayed as Component/Cluster/vNode.

- ▶ Logical Volume: VOLSER of the mounted logical volume. If the value of Mounted is No, then this value is the VOLSER of the most recently mounted logical volume.
- ▶ Online: The online state from the perspective of the cluster. Possible values are Yes and No.
- ▶ Role: The current role the drive is performing. Possible values are:
 - IDLE: The drive is currently not in use.
 - READ: The drive is being used to read a logical volume.
 - WRITE: The drive is being used to write a logical volume.
- ▶ Mounted: Whether the drive is mounted. Possible values are Yes and No.
- ▶ Stand-alone Mount: Whether the virtual drive is a stand-alone mount. Possible values are Yes and No.

To view additional information pertaining to a specific drive in a cluster, in the Virtual Drives table, select the radio button from the Select column that appears in the same row as the drive's path. From the drop-down menu, select **Details** and click **Go** to display the Virtual Drives Details table, as shown in Figure 8-12.

The screenshot shows the TS7700 Virtualization Engine web interface. On the left is a navigation menu with categories like Health & Monitoring, Performance & Statistics, Logical Volumes, Physical Volumes, Constructs, Configuration, User Management, and Service & Troubleshooting. The main content area is titled '"Juno[0]" (#BA41A): Virtual Tape Drives'. It includes a 'Refresh' button and a timestamp 'Last Refresh: Tuesday, October 21, 2008 8:04:25 AM'. Below this is a 'Virtual Drive Details' table.

Virtual Drive Details:	
Virtual Drive Address:	vt000
Path:	"juno[0]" (#BA41A)/V[0]
Logical Volume:	203586
Online:	Yes
Mounted:	No
Role:	idle
Primary Cache Mount Cluster:	"juno[0]" (#BA41A)
Time On The Drive:	0 hours, 0 minutes, 0 seconds
Bytes Read (Compressed/Raw):	0 KiB/0 KiB
Bytes Written (Compressed/Raw):	0 KiB/0 KiB
Logical Block Position:	0x0
Standalone Mount:	No

At the bottom of the interface, there is an 'External Links' section with a link to 'Launch Library 3584 Specialist'.

Figure 8-12 MI Virtual Tape Drive details

In addition to the Virtual Drive Address, Path, Logical Volume, Online, Mounted, Stand-alone Mount, and Role information shown in the Virtual Drives table, the Virtual Tape Drive Details table also displays the following status and identity information:

- ▶ Time on the Drive: The elapsed time that the logical volume has been mounted on the virtual tape drive. This value takes the format hhhh hours, mm minutes, ss seconds.
- ▶ Bytes Read (Compressed/Raw): The amount of data read from the logical volume, shown first as number of compressed bytes, then number of raw bytes. These values are displayed in binary bytes, in KiB units.

- **Bytes Written (Compressed/Raw):** The amount of data written to the mounted logical volume, shown first as number of compressed bytes, then number of raw bytes. These values are displayed in binary bytes, in KiB units.
- **Logical Block Position:** The position of the virtual drive on the tape surface, as represented by block number starting from the beginning of the tape. This value is displayed in hexadecimal form, such as 0x34. If the volume is not mounted, the value in this field is 0.

Use the drop-down menu on the Virtual Drives table to control the appearance of data in the table, mount a logical volume, or demount a logical volume. Use the Select Table Action section of the drop-down menu to select or deselect multiple rows or control the appearance of the table through sort and filter actions.

To mount a logical volume, select **Stand-alone Mount Logical Volume** from the drop-down menu and click **Go**. To demount a logical volume, select the radio button in the row of the volume you want to demount, select **Stand-alone Demount Logical Volume** from the drop-down menu, and click **Go**.

Physical Media Inventory window

Use this window (Figure 8-13) to obtain detailed information about a physical media inventory in the TS7740 Virtualization Engine.

Note: If the cluster does not possess a physical library, the Physical Media Inventory page will not be visible on the TS7700 Virtualization Engine management interface.

Whenever you need to know specific details about a physical media, such as its pool number, media type pool related, or the quantity of physical media in full status, this window allows you to see all that desired information, as shown in Figure 8-13.

The screenshot shows the TS7700 Virtualization Engine management interface. The left sidebar contains a navigation menu with categories like Health & Monitoring, Performance & Statistics, Logical Volumes, Physical Volumes, Constructs, Configuration, User Management, and Service & Troubleshooting. The main content area is titled '"Celeste[2]" (#BA68C): Physical Media Inventory'. It includes a 'Refresh' button and a 'Last Refresh' timestamp. Below this, a tree view shows the cluster structure with 'Celeste[2]' selected. A table titled 'Inventory of Physical Media Pools' displays the status of various media pools.

Pool	Media Type	Empty	Filling	Full	Queued for E...	ROR	Unavailable
0 JB		607	-	-	-	0	1
1 JB		2	18	37	0	0	0
11 JB		2	3	22	0	0	0
12 JB		2	0	0	0	0	0
13 JB		2	0	0	0	0	0
24 JB		0	2	0	0	0	0
30 JB		2	5	95	0	0	0

Figure 8-13 TS7700 Virtualization Engine management interface Physical Media Inventory

The following physical media counts are displayed for each media type in each storage pool:

- ▶ **Pool:** This is the number that identifies the pool for the specific media type.
- ▶ **Media Type:** The media type defined for the pool. A storage pool can have multiple media types and each media type will be displayed separately. Possible values are:
 - JA (ETC): Enterprise Tape Cartridge (ETC)
 - JB (EETC): Enterprise Extended-Length Tape Cartridge (EETC)
 - JJ (EETC): Enterprise Economy Tape Cartridge (EETC)
- ▶ **Empty:** The count of physical volumes that are empty for the pool.
- ▶ **Filling:** The count of physical volumes filling for the pool. This field is blank for Pool 0.
- ▶ **Full:** The count of physical volumes that are full for the pool. This field is blank for Pool 0.
- ▶ **Queued for Erase:** The count of physical volumes that are reclaimed but need to be erased before they can become empty. This field is blank for Pool 0.
- ▶ **ROR (Read Only Recovery):** The count of physical volumes that are in the Read Only Recovery state. That does not imply that there is an error on the physical stacked volume. The ROR is activated for the reclamation process.
- ▶ **Unavailable:** The count of physical volumes that are in the unavailable or destroyed state.

Note: Pool 0 is the Common Scratch Pool (CSP) and all cartridges are empty. These can be used as required by the pools 1-32

Tape Volume Cache window

Use this window (Figure 8-14) for obtaining information about the tape volume cache in the TS7700 Virtualization Engine.

The screenshot displays the IBM System Storage™ TS7700 Virtualization Engine management interface. The left sidebar contains a navigation menu with categories like Health & Monitoring, Performance & Statistics, Logical Volumes, Physical Volumes, Constructs, Configuration, User Management, and Service & Troubleshooting. The main content area is titled '"Archie[0]" (#BA96A): Tape Volume Cache'. It shows a breadcrumb trail: Grid Summary > "Dohnut" Cluster Summary > Tape Volume Cache. A 'Refresh' button and the last refresh time (Wednesday, November 04, 2009 10:42:03 AM) are present. A note states that values may differ from the Cache Utilization page due to 15-minute update increments. A 'Current cluster selected:' section shows a tree view with 'Archie[0]' selected. Below are two tables: 'Cache Summary' and 'Partition Summary'.

Cache Summary:

Installed size:	9,000 GB (8,382 GiB)
Available size:	5,000 GB (4,657 GiB) 56%

Partition Summary:

Allocated size:	5,000 GB (4,657 GiB) 100%
Used size:	0 GB (0 GiB) 0%
Preference Group 0 size:	0 GB (0 GiB) 0%
Preference Group 1 size:	0 GB (0 GiB) 0%
Pre-migration size:	0 GB (0 GiB) 0%
Copy size:	0 GB (0 GiB) 0%
Host Write Throttling time:	0 ms
Copy throttling time:	0 ms

External Links: [Launch Library 3584 Specialist](#)

Figure 8-14 TS7700 Virtualization Engine management interface Tape Volume Cache

The following information is displayed for the tape volume cache:

- ▶ **Cache summary:** The table contains status information about the summary of the cache of the TS7740 Virtualization Engine. The following information can be obtained:
 - Installed size: The cache size that is physically installed in gigabytes (GB).
 - Available size: The cache size that is logically enabled in gigabytes (GB).
- ▶ **Partition Summary:** The table contains status information related to the cache.
 - Allocated size: The amount of cache that has been allocated for this partition.
 - Used size: The amount of cache that is currently in use by the partition.
 - Preference Group 0 size: The size of preference group zero. Volumes in this group are preferred to be removed from cache over other volumes.

Note: The Preference Group size shown on the Tape Volume Cache page is the latest cache size; it is the same as the values returned through Bulk Volume Information Retrieval. By contrast, the Preference Group size displayed on the Cache Utilization page is the average of cache sizes sampled every 15 seconds, over a 15 minute period.

- Preference Group 1 size: The size of preference group one. Volumes in this group are preferred to be retained in cache over other volumes.
- Pre migration size: The amount of data that needs to be copied to a physical volume.
- Copy size: The amount of data that needs to be copied to another cluster.
- Migration throttling time: Indicates whether migration throttling is active for the partition. If throttling is active, a value in milliseconds is displayed. If throttling is not active, zero is displayed.
- Copy throttling time: Indicates whether copy throttling is active for the partition. If throttling is active, a value in milliseconds is displayed. If it is not active, zero is displayed.

Physical Library window with a TS7740 Virtualization Engine

Use this window (Figure 8-15) to view the status of a TS7740 Virtualization Engine physical library and adjust its settings. The Physical Library Properties table displays health and monitoring status and settings for a given physical library in a component, as shown in Figure 8-15.

Note: If the cluster does not possess a physical library, as in a TS7720 Virtualization Engine, the Physical Library window will not be visible on the TS7700 Virtualization Engine management interface.

The screenshot displays the IBM System Storage™ TS7700 Virtualization Engine management interface. The top navigation bar includes 'Welcome admin' and 'Help | Log out'. The left sidebar contains a tree view with categories like 'Health & Monitoring', 'Performance & Statistics', and 'External Links'. The main content area is titled '"Veronca1[3]" (#BA96D): Physical Library' and shows the 'Physical Library Properties' table.

Physical Library Properties:	
Operational Mode:	Auto
Online State:	Online
Health:	Normal
I/O Slots:	Empty
Frame Doors:	Closed
Physical Volumes:	59

Below the table, it indicates 'Library 3584 Specialist'.

Figure 8-15 MI Physical Library with a TS7740 Virtualization Engine

Status information displayed in the Physical Library Properties table includes:

- ▶ **Operational Mode:** The operational status of the physical library. Possible values include:
 - Auto: The library is operating in automatic mode.
 - Paused: The library's operations are paused.
 - Manual: The library is operating in manual mode.
- ▶ **Online State:** The online status of the physical library. Possible values include:
 - Online: Host applications can interact with the library.
 - Offline: Host applications cannot interact with the library.
- ▶ **Health:** The health status of the physical library. Possible values include:
 - Normal: The library is online and ready to interact with the host.
 - Degraded: There is an actual or potential problem with the library.
 - Failed: The library is offline or there is no response to communication from the host.
- ▶ **I/O Slots:** The status of the input/output station used to move cartridges into and out of the library. Possible values include:
 - Occupied: The I/O station is occupied.
 - Empty: The I/O station is not occupied.
 - Full: The I/O station is occupied and at full capacity.
- ▶ **Frame Doors:** The position of the doors on the physical library. Possible values include:
 - Open: One or more enclosure doors are open.
 - Closed: All enclosure doors are closed.
- ▶ **Physical Volumes:** The number of volumes allocated to the TS7740 Virtualization Engine by the physical library. This value is the remainder of the total number of empty slots subtracted from the total cell capacity.
- ▶ **Library Specialist Link:** This link appears under the heading External Links below the work area on the left. Click this link to access the TS3500 Library Specialist Web interface belonging to the physical library attached to this selected cluster. The link is not displayed until it is configured in the TS7740 Virtualization Engine by the SSR.

Operation History window

Use this window for viewing information about currently running and completed tasks on a TS7700 Virtualization Engine Cluster.

The following information is available for currently running and completed tasks and is presented in two tables, as shown in Figure 8-16 on page 434.

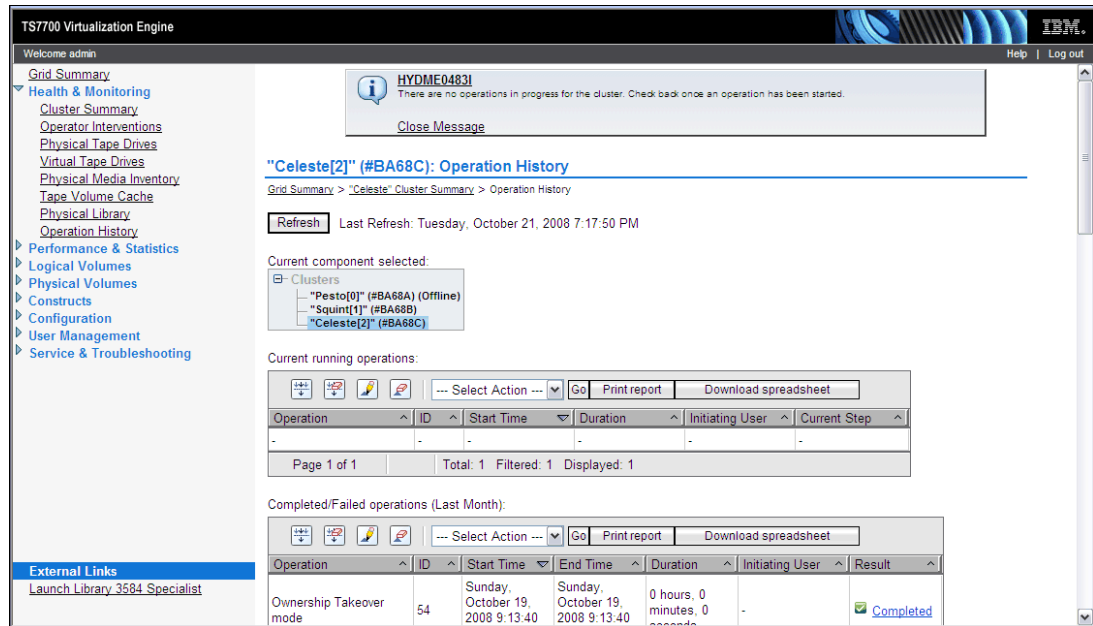


Figure 8-16 MI Operation History

The first table displays information for the currently running tasks on the TS7700 Virtualization Engine Cluster. The second table displays information about completed or failed tasks. You can print the table data by clicking **Print Report** (next to the Select Action menu). A comma separated value (.csv) file of the table data can be downloaded by clicking **Download spreadsheet**.

Note: A stand-alone mount or demount can fail and return an Error Recovery Action (ERA) code and modifier. To determine what this ERA code means, select **Help** from the top left of the window and a link to an ERA code table is displayed.

The spreadsheet has the following values:

- ▶ Task: Name of the task.
- ▶ ID: ID number of the task.
- ▶ Start Time: The time stamp for the start of the task.
- ▶ End Time: The time stamp for when the task ends. This is only available for completed tasks.
- ▶ Duration (h:m:s): In the case of completed task, this value displays the time the task took to complete. For currently running tasks, this displays the time the task has taken up to this point.
- ▶ Initiating user: The user login that initiated the task.
- ▶ Current step: Additional contextual information about the task's current progress and a link to additional information about the status of the task. This is only available for running tasks.
- ▶ Result: The result of the task. This information is only available for completed or failed tasks. A failed task or a task that has completed will have a link to additional information about the failure.

8.2.4 Performance and Statistics windows

This section shows the list of Performance and Statistics windows available. The graphical information available for monitoring your TS7700 Virtualization Engine and how it can be used to maximize subsystem resources is presented and discussed in Chapter 9, “Performance and monitoring” on page 589. The graphical views are from the last 15 minutes of activity, as shown in Figure 8-17.

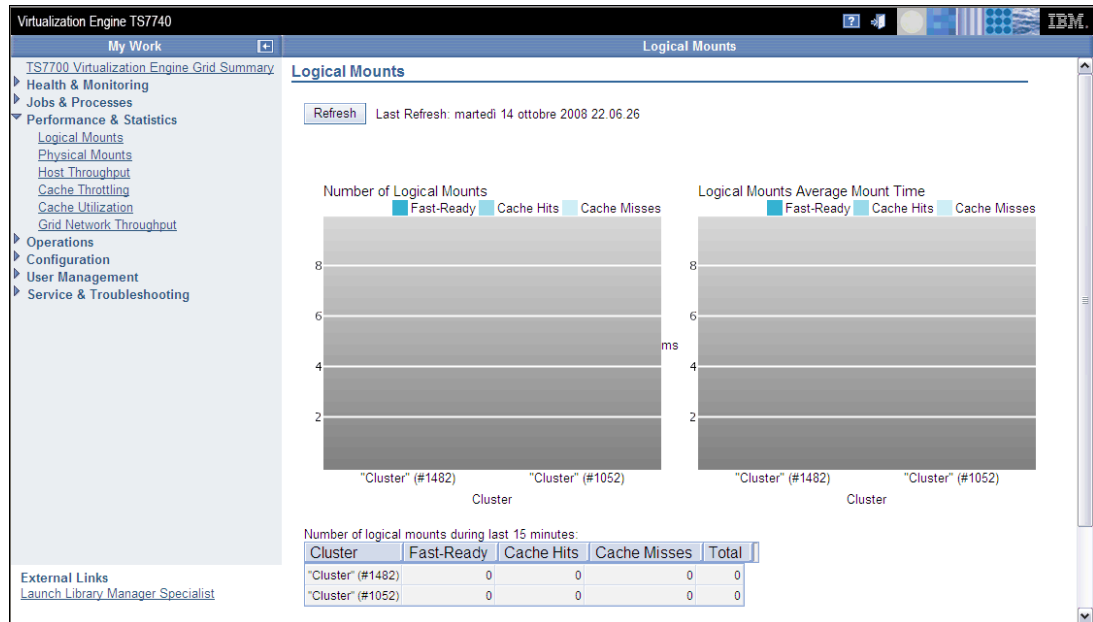


Figure 8-17 Logical Mounts statistics

Note: For more detailed information, see Chapter 9, “Performance and monitoring” on page 589.

8.2.5 Logical volumes

The topics in this section present information related to monitoring and manipulating logical volumes in the TS7700 Virtualization Engine. The topics are:

- ▶ Incoming Copy Queue
- ▶ Recall Queue
- ▶ Logical Volume Details
- ▶ Insert Logical Volumes
- ▶ Delete Logical Volumes
- ▶ Modify Logical Volumes
- ▶ Move Logical Volumes
- ▶ Logical Volume Search
- ▶ Fast Ready Categories

Incoming Copy Queue window

Use this window (Figure 8-18) to view the logical volume Incoming Copy Queue for a TS7700 Virtualization Engine Cluster. This view can only be selected if the cluster is part of a multi-cluster grid configuration.

The screenshot shows the 'Incoming Copy Queue' window for the 'Celeste[2]' cluster. The left sidebar contains a navigation menu with options like 'Grid Summary', 'Health & Monitoring', 'Performance & Statistics', 'Logical Volumes', 'Physical Volumes', 'Constructs', 'Configuration', 'User Management', and 'Service & Troubleshooting'. The main content area displays the cluster name, a 'Refresh' button, and the last refresh time. Below this, a tree view shows the cluster hierarchy with 'Celeste[2]' selected. A 'Download' button is available for generating a report. A table shows the number of logical volumes queued for copy to the cluster.

Copy Type	Quantity	Size (GiBs)
Rewind unload (RUN)	0	0
Deferred	699	268

Figure 8-18 TS7700 Virtualization Engine management interface Incoming Copy Queue

To obtain a report of each individual queued logical volume copy in a comma-separated file format (CSV), click the **Download** button.

Excel removes leading zeros from downloaded .csv (comma-separated value) files, which can lead to errors when data is displayed in a spreadsheet format. For example, if a VOLSER in a downloaded list has a value of 012345, it will appear in Excel as 12345. To preserve the leading zeros in your downloaded .csv files, perform the following steps:

1. When downloading a report in CSV format, click the **Download** button and then click **Save**.
2. When the file name appears, change the file extension from .csv to .txt. For instance, if the default name of the downloading file is `copies.csv`, change the name to `copies.txt`. The downloaded file will retain the comma-delimited format.

To open the downloaded file in Excel and retain any leading zeros, perform the following steps:

1. Open Excel.
2. Select **File** → **Open** and browse to the .txt file you downloaded. Highlight the .txt file and click **Open**. The Text Import Wizard will open.
3. In Step 1 of the Text Import Wizard, select **Delimited** as the original data type and click **Next**.
4. In Step 2 of the Text Import Wizard, select **Comma** as the delimiter and click **Next**.

5. In Step 3 of the Text Import Wizard, select **Text** as the column data format and click **Finish**.

You will now be able to save the downloaded file with a .txt or .xls extension while retaining any leading zeros.

The following information is available about the logical volume copy queue:

- ▶ Copy Type: Possible values are:
 - RUN: The copy will occur during rewind-unload and before the rewind-unload operation completes at the host.
 - Deferred: The copy will occur some time after the rewind-unload operation at the host.
- ▶ Quantity: The total number of queued logical volume copies of the indicated copy type.
- ▶ Size: The sum of the size of all volumes in the copy queue, displayed in gibibytes (GiB).

Downloadable report fields for queued logical volumes

The following information about individual queued logical volume copies is available in a downloadable comma separated value format (.csv) report that can be obtained by clicking the **Download** button:

- ▶ VOLSER: VOLSER of the logical volume being copied.
- ▶ Queue Type: The type of queue as reported by the cluster. Possible values are:
 - Rewind unload (RUN): The copy occurs before the rewind-unload operation completes at the host.
 - Deferred: The copy occurs some time after the rewind-unload operation completes at the host.
 - Immediate Deferred: A RUN copy that has been moved to the deferred state because of copy timeouts or TS7700 Virtualization Engine Grid states.
- ▶ Copy Activity: Status information about the copy activity of the logical volume copy. Possible values are:
 - Complete: A consistent copy exists at this location.
 - In Progress: A copy is required and currently in progress.
 - Queued: A copy is required at this location and is waiting to be processed.
 - Required: A copy is required at this location but has not started or completed.
 - Not Required: A copy is not required at this location.
 - Reconcile: Pending updates exist against this location's volume. The copy activity will be updated after the pending updates are resolved.
- ▶ Age in Queue: Elapsed time that the copy has been in the queue, displayed in hours (H), minutes (M), and seconds (S) in the following format: HH:MM:SS.
- ▶ Volume Size: Bytes that will be written during the logical volume copy.

Recall Queue window for TS7740 Virtualization Engine

Use this window (Figure 8-19) to view the logical volume Recall Queue and to promote logical volumes to the front of the queue in the TS7740 Virtualization Engine.

Note: If the cluster does not possess a physical library, the Recall Queue page will not be visible on the TS7700 Virtualization Engine management interface.

A recall of a logical volume retrieves the logical volume from physical tape and places it in the cache. A queue is used to process these requests.

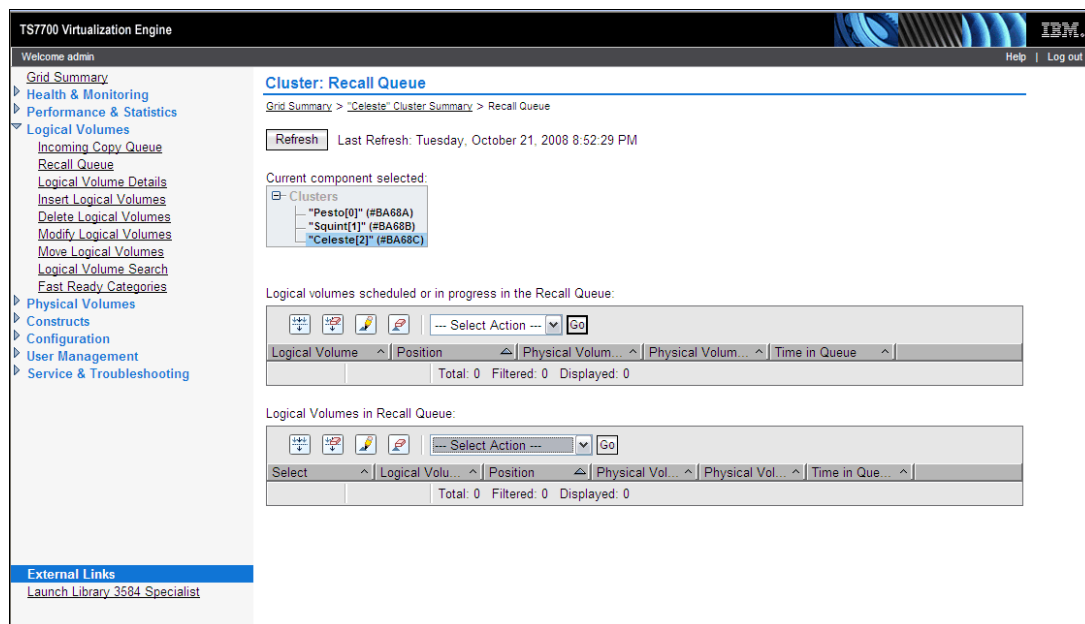


Figure 8-19 TS7700 Virtualization Engine management interface Recall Queue

This window displays recall queue information in two tables:

- ▶ Logical volumes scheduled or in progress in the Recall Queue window, which displays logical volumes scheduled for recall or in the process of being recalled.
- ▶ Logical volumes in the Recall Queue window, which displays the queue of volumes waiting to be scheduled for a recall.

Both tables display the following logical volume information:

- ▶ Logical Volume: The logical volume to be recalled.
- ▶ Position: The position of the logical volume in the recall queue. Possible values are:
 - In Progress: A recall is in progress for the volume.
 - Scheduled: The volume is scheduled to be recalled. If optimization is enabled, the TS7700 Virtualization Engine will schedule recalls to be processed from the same physical volume.
 - Number in Queue: A three-digit number that represents the volume's current position in the list of volumes that have not yet been scheduled.

- ▶ **Physical Volume 1:** The physical volume on which the logical volume resides.
- ▶ **Physical Volume 2:** A second physical volume the logical volume resides on if the logical volume spans a physical volume. Spanning will only occur with migrated data.
- ▶ **Time in Queue:** the length of time the logical volume has been in the queue, displayed in hours (H), minutes (M), and seconds (S) in the following format: HH:MM:SS.

However, only the logical volumes in the Recall Queue table will permit you to modify the Recall Queue by moving one or more logical volumes to the top of the queue. Move a logical volume to the top of the queue by selecting it from the Logical volumes in Recall Queue table, then select **Move To Head Of Queue** from the Select Action drop-down menu and click **Go**. You will be asked to confirm your decision to move a logical volume to the top of the queue. Click **Yes** to move the selected logical volume to the top of the queue. Click **No** to abandon the Move to Head of Queue operation.

Logical Volume Details window

Use the Logical Volume Details window, shown in Figure 8-20, to obtain detailed information about the state of a logical volume in the TS7700 Virtualization Engine Grid.

To obtain details about a logical volume, enter the logical volume's VOLSER in the available text field, as shown in Figure 8-20, and then select the **Get Details** button. The VOLSER must be six alphanumeric characters.

You can view details and the current status of a logical volume in a cluster by using the Logical volume summary image or the Logical volume details table. You can also use the Physical locations of requesting logical volume table to view information about requesting logical volumes on each cluster.

Note: If the cluster does not possess a physical library, physical resources are not shown in the Logical volume summary, Logical volume details table, or the Physical locations of requesting logical volume table.

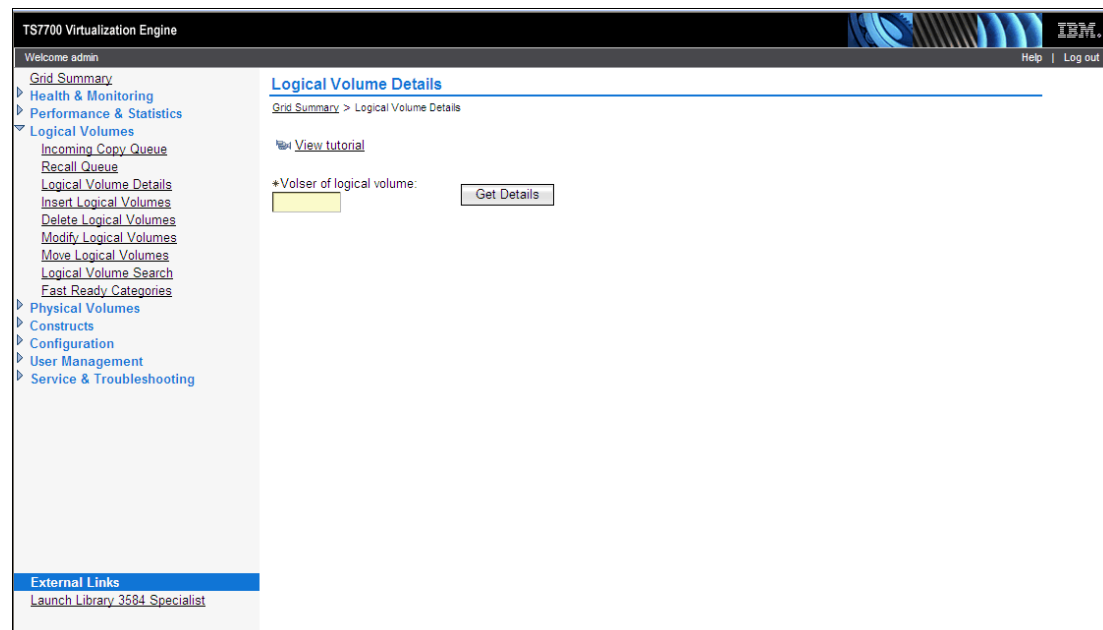


Figure 8-20 MI Logical Volume Details

Figure 8-21 shows an example of a Logical Volume Details result.

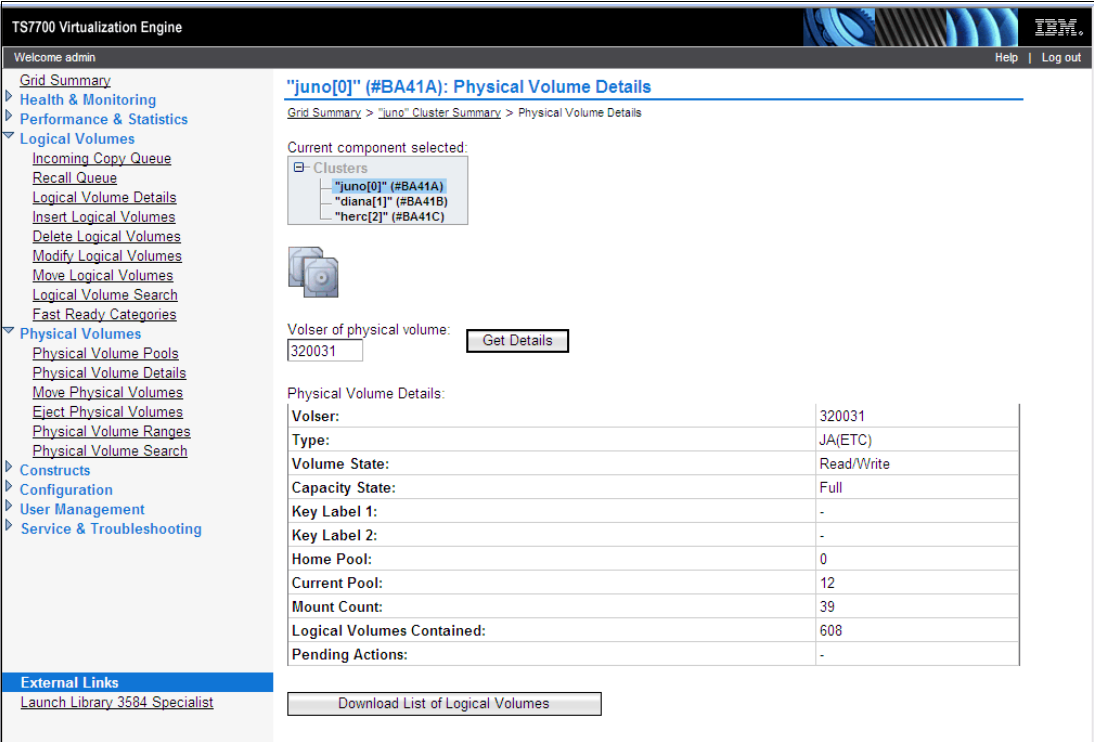


Figure 8-21 MI Logical Volume Details

The logical volume summary image is divided into virtual and physical sections.

The cluster that currently owns the logical volume is highlighted in this image. The status of each logical volume shown in the Logical volume summary image is indicated by the icon used to represent that logical volume; the Legend explains the meaning of each icon. To view additional details about the status of a logical volume or to find links to other relevant management interface pages, position your cursor over that logical volume's icon in the Logical volume summary image.

Figure 8-22 shows the logical volume details legend.

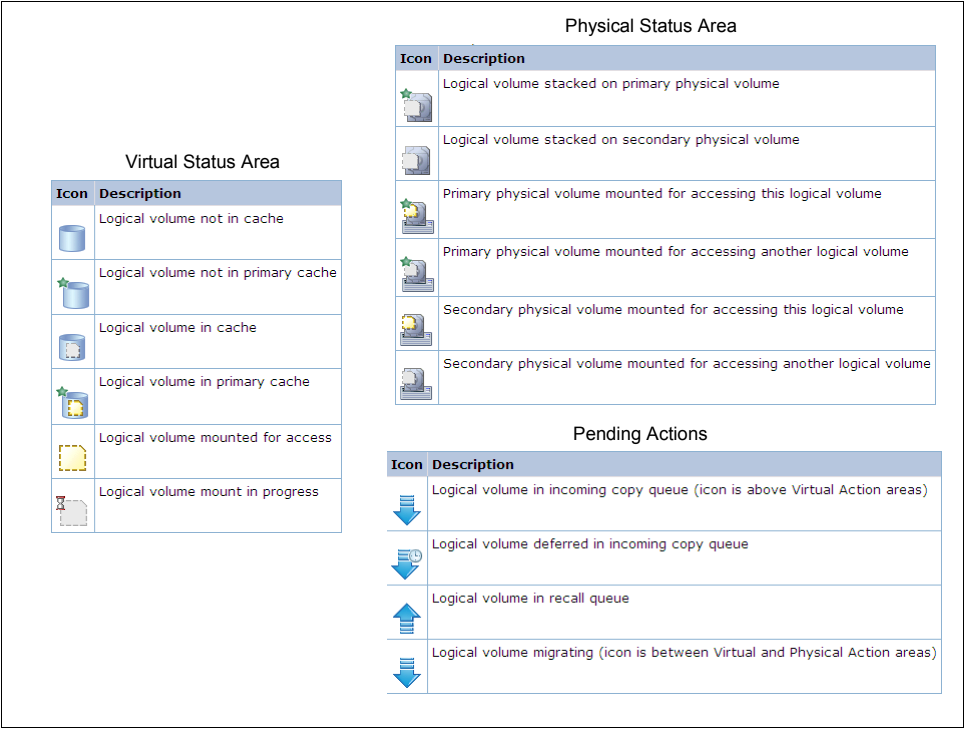


Figure 8-22 MI Logical Volume Details legend

The figure shows a full legend for the graphical representation of logical volume details in the TS7700 Virtualization Engine.

Figure 8-23 shows the virtual section of the summary window.

LEGEND	
Logical volume details:	
Volser:	V90000
Media Type:	Enhanced Capacity Cartridge System Tape (800 MiB)
Compressed Data Size:	1.2 MiB
Maximum Volume Capacity:	800 MiB
Current Owner:	"Cluster[0]" (#BA19A)
Currently Mounted:	No
vNode:	-
Virtual Drive:	-
Cache Copy Used for Mount:	"Cluster[0]" (#BA19A)
Mount State:	-
Cache Management Preference Group:	1
Last Accessed by a Host:	Thursday, August 14, 2008 2:10:40 AM
Last Modified:	Tuesday, July 22, 2008 1:18:43 AM
Category:	000F
Storage Group:	SGB19P03
Management Class:	MCP00L00
Storage Class:	CACHEIN
Data Class:	HSMDCLS7
Volume Data State:	Active
Earliest Deletion On:	-

Figure 8-23 Logical Volume virtual section

The Logical volume details table shows details and the current status of a logical volume in a cluster. These are logical volume properties that are consistent across the Grid.

The virtual section of Logical Volume Details in Figure 8-23 includes:

- ▶ **Volser:** The VOLSER of the logical volume. This is a six-character number that uniquely represents the logical volume in the virtual library.
- ▶ **Media Type:** The media type of the logical volume. Possible values are:
 - Cartridge System Tape (400 MiB).
 - Enhanced Capacity Cartridge System Tape (800 MiB).
- ▶ **Compressed Data Size:** The compressed file size of the logical volume in cache expressed in bytes (B), kibibytes (KiB), mebibytes (MiB), or gibibytes (GiB).
- ▶ **Maximum Volume Capacity:** The maximum size of the logical volume expressed in bytes (B), kibibytes (KiB), mebibytes (MiB), or gibibytes (GiB). This capacity is set by the volume's Data Class.
- ▶ **Current owner:** The name of the cluster that currently owns the latest version of the logical volume.
- ▶ **Currently Mounted:** Whether the logical volume is currently mounted in a virtual drive.
- ▶ **vNode:** The name of the vNode on which the logical volume is mounted.
- ▶ **Virtual Drive:** The ID of the virtual drive on which the logical volume is mounted.
- ▶ **Cache Copy Used for Mount:** The cluster name of the cache chosen for I/O operations for mount based on consistency policy, volume validity, residency, performance, and cluster mode.

- ▶ **Mount State:** The current mount state of the logical volume. Possible values are:
 - Mounted: The volume is mounted.
 - Mount Pending: A mount request has been received and is in progress.
 - Recall Queued: A mount request has been received and requires a recall.
 - Recalling: A mount request has been received and a recall from physical tape is in progress.
- ▶ **Cache Management Preference Group:** The Preference Group (PG) for the logical volume. The PG determines how soon volumes are removed from cache following their copy to tape. This information is only displayed if a physical library exists in the Grid. Possible values are:
 - 0: Volumes in this group are preferred to be removed from cache over other volumes. The preference is to remove volumes by size, largest first.
 - 1: Volumes in this group are preferred to be retained in cache over other volumes. A “least recently used” algorithm is used to select volumes for removal from cache if there are no volumes to remove in preference group 0.
 - Unknown: The preference group cannot be determined.
- ▶ **Last Accessed by a Host:** The date and time the logical volume was last accessed by a host. The time recorded reflects the time zone in which the user's browser is located.
- ▶ **Last Modified:** The date and time the logical volume was last modified by a host. The time recorded reflects the time zone in which the user's browser is located.
- ▶ **Category:** The number of the Fast Ready category to which the logical volume belongs. A Fast Ready category groups logical volumes for quick, non-specific use.
- ▶ **Storage Group:** The name of the Storage Group that defines the primary pool for the logical volume's premigration.
- ▶ **Management Class:** The name of the Management Class applied to the cluster. This policy defines the copy process for volume redundancy.
- ▶ **Storage Class:** The name of the Storage Class applied to the cluster. This policy classifies logical volumes to automate storage management.
- ▶ **Data Class:** The name of the Data Class applied to the cluster. This policy classifies logical volumes to automate storage management.
- ▶ **Volume Data State:** The state of the data on the logical volume. Possible values are:
 - New: The logical volume is in the insert category or a non-fast-ready category and data has never been written to it.
 - Active: The logical volume is currently located within a private category and contains data.
 - Scratched: The logical volume is currently located within a fast ready (scratch) category and its data is not scheduled to be automatically deleted.
 - Pending Deletion: The volume is currently located within a fast ready (scratch) category and its contents will become a candidate for automatic deletion when the earliest deletion time has passed. Automatic deletion will occur sometime thereafter.

Note: The volume can be accessed for mount or category change prior to the automatic deletion and therefore the deletion might not be completed.

- Pending Deletion with Hold: The volume is currently located within a fast ready (scratch) category configured with hold and the earliest deletion time has not yet passed. The volume is not accessible by any host operation until the volume has left the hold state. After the earliest deletion time has passed, the volume will become a candidate for deletion and move to the Pending Deletion state. While in this state, the volume is accessible by all legal host operations.
- Deleted: The volume is either currently within a fast ready category or has previously been in a fast ready category where it become a candidate for automatic deletion and the deletion was carried out successfully. Any mount operation to this volume will be treated as a fast ready (scratch) mount, because no data is present.
- ▶ Earliest Delete On: The date and time when the logical volume will be deleted. The time recorded reflects the time zone in which the user's browser is located. If there is no expiration date set, this value displays as “-”.
- ▶ Logical WORM: Whether the logical volume is formatted as a write-once, read many (WORM) volume. Possible values are Yes and No.

Figure 8-24 shows the Cluster-specific Logical Volume Properties part of the window.

Cluster-specific Logical Volume Properties:				
	"Archie[0]" (#BA96A)	"Veronica[1]" (#BA96B)	"Dohnut[2]" (#BA96C)	"Veronca1[3]" (#BA96D)
In Cache	-	-	-	-
Primary Physical Volume	None	None	None	None
Secondary Physical Volume	None	None	None	None
Copy Activity	Not Required	Not Required	Not Required	Not Required
Queue Type	-	-	-	-
Copy Mode	No copy	No copy	No copy	No copy
Deleted	-	-	-	-

Figure 8-24 Cluster-specific Logical Volume Properties

The Cluster-specific Logical Volume Properties table displays information about the specified logical volumes on each cluster. These are properties that are specific to each cluster. The logical volume details and status displayed include:

- ▶ Cluster: The cluster location of the logical volume copy.
- ▶ In Cache: Whether the logical volume is in cache for this cluster.
- ▶ Primary Physical: The physical volume that contains the specified logical volume. Click the VOLSER hyperlink to open the Physical Stacked Volume Details page for this physical volume. A value of None means that no primary physical copy is to be made. This column is only visible if a physical library is present in the Grid. If there is at least one physical library in the Grid, the value in this column is shown as “-” for those clusters not attached to a physical library.
- ▶ Secondary Physical: Secondary physical volume that contains the specified logical volume. Click the VOLSER hyperlink to open the Physical Stacked Volume Details page for this physical volume. A value of None means that no secondary physical copy is to be made. This column is only visible if a physical library is present in the Grid. If there is at least one physical library in the Grid, the value in this column is shown as “-” for those clusters not attached to a physical library.
- ▶ Copy Activity: Status information about the copy activity of the logical volume copy. Possible values are:
 - Complete: A consistent copy exists at this location.
 - In Progress: A copy is required and currently in progress.
 - Required: A copy is required at this location but has not started or completed.

- Not Required: A copy is not required at this location.
- Reconcile: Pending updates exist against this location's volume. The copy activity updates after the pending updates get resolved.
- ▶ Queue Type: The type of queue as reported by the cluster. Possible values are:
 - Rewind unload (RUN): The copy occurs before the rewind-unload operation completes at the host.
 - Deferred: The copy occurs some time after the rewind-unload operation completes at the host.
 - Immediate Deferred: A RUN copy that has been moved to the deferred state because of copy timeouts or TS7700 Virtualization Engine Grid states.
- ▶ Copy Mode: The copy behavior of the logical volume copy. Possible values are:
 - Rewind unload (RUN): The copy occurs before the rewind-unload operation completes at the host.
 - Deferred: The copy occurs some time after the rewind-unload operation at the host.
 - No copy: No copy will be made.
 - Exist: A consistent copy exists at this location although No Copy is intended. A consistent copy existed at this location at the time the logical volume was mounted. After the volume is modified, the Copy Mode of this location changes to No Copy.
- ▶ Deleted: The date and time when the logical volume on the cluster was deleted. The time recorded reflects the time zone in which the user's browser is located. If the volume has not been deleted, this value displays as “-”.
- ▶ Removal Residency: The residency state of the logical volume. Values are only displayed in this field if the cluster is part of a hybrid grid. Possible values are:
 - Dash (-): Removal Residency does not apply to the cluster.
 - Removed: The logical volume has been removed from the cluster.
 - Resident: The logical volume is a candidate for removal, but the removal has not yet occurred.
 - Retained: An attempt to remove the logical volume occurred, but the operation failed. The copy on this cluster cannot be removed based on the configured copy policy and the total number of configured clusters. Removal of this copy lowers the total number of consistent copies within the Grid to a value below the required threshold. If a removal is expected at this location, verify that the copy policy is configured appropriately and that copies are being replicated to other peer clusters. This copy can be removed only after a sufficient number of replicas exist on other peer clusters.
 - Deferred: An attempt to remove the logical volume occurred, but the operation failed. This state can result from a cluster outage or any state within the Grid that disables or prevents replication. The copy on this cluster cannot be removed based on the configured copy policy and the total number of available clusters capable of replication. Removal of this copy lowers the total number of consistent copies within the Grid to a value below the required threshold. This copy can be removed only after a sufficient number of replicas exist on other available peer clusters. A subsequent attempt to remove this volume occurs after no outage exists and replication is allowed to continue.
- ▶ Removed Time: The date and time the logical volume was removed from the cluster. The time recorded reflects the time zone in which the user's browser is located.

Insert Logical Volumes window

Use this window (Figure 8-25) to insert a range of logical volumes in the TS7700 Virtualization Engine Grid. Logical volumes inserted on an individual cluster will be available to all clusters within a Grid configuration.

During cartridge entry processing, even if the library is online and operational to a given host, at least one device needs to be online (or have been online) to that host for the library to be able to send the cartridge entry attention interrupt to that host. If the library is online and operational, yet there are no online devices to a given host, that host will not receive the attention interrupt from the library unless a device had previously been varied online. To get around this limitation, ensure that at least one device is online (or had been online) to each host or use the LIBRARY RESET,CBRUXENT command to initiate cartridge entry processing from the host. This is especially important if you only have one host attached to the library that owns the volumes being entered. In general, after you have entered volumes into the library, if you do not see the expected CBR36xxI cartridge entry messages being issued, you can use the LIBRARY RESET,CBRUXENT command from MVS to initiate cartridge entry processing. The LIBRARY RESET,CBRUXENT command causes the host to ask for any volumes in the insert category.

TS7700 Virtualization Engine

Welcome admin [Help](#) [Log out](#)

[Grid Summary](#)
▶ [Health & Monitoring](#)
▶ [Performance & Statistics](#)
▼ [Logical Volumes](#)
 [Incoming Copy Queue](#)
 [Recall Queue](#)
 [Logical Volume Details](#)
 [Insert Logical Volumes](#)
 [Delete Logical Volumes](#)
 [Modify Logical Volumes](#)
 [Move Logical Volumes](#)
 [Logical Volume Search](#)
 [Fast Ready Categories](#)
▶ [Physical Volumes](#)
▶ [Constructs](#)
▶ [Configuration](#)
▶ [User Management](#)
▶ [Service & Troubleshooting](#)

[External Links](#)
[Launch Library 3584 Specialist](#)

Insert Logical Volumes

[Grid Summary](#) > Insert Logical Volumes

[Refresh](#) Last Refresh: Wednesday, October 22, 2008 12:47:21 AM

Current availability across entire grid:

Currently Inserted:	170,282
Maximum Allowed:	1,000,000
Available Slots:	829,718

Show inserted volume ranges between and [Show](#)

Insert a new logical volume range:

+Starting volser: +

☐ Quantity ☐ Ending volser

Initially owned by:
 ▼

Media type:
☐ Cartridge System Tape (400 MiB)
☒ Enhanced Capacity Cartridge System Tape (800 MiB)

☒ Set Constructs

Storage Groups: ▼ Storage Classes: ▼ Data Classes: ▼ Management Classes: ▼

[Insert](#)

Figure 8-25 TS7700 Virtualization Engine MI Insert Logical Volumes window

The table shown at the top of Figure 8-25 on page 446 shows current information about the number of logical volumes in the TS7700 Virtualization Engine:

Currently Inserted	The total number of logical volumes inserted into the composite library
Maximum Allowed	The total maximum number of logical volumes that can be inserted into the composite library
Available Slots	The available slots remaining for logical volumes to be inserted, which are obtained by subtracting the Currently Inserted logical volumes from the Maximum Allowed

To view the current list of logical volume ranges in the TS7700 Virtualization Engine Grid, enter a logical volume range and click **Show**.

The following information is useful if you want to perform an Insert a new logical volume range function:

- ▶ **Starting VOLSER:** The first logical volume to be inserted. The range for inserting logical volumes begins with this VOLSER number.
- ▶ **Quantity:** Select this option to insert a set amount of logical volumes beginning with the Starting VOLSER. Enter the quantity of logical volumes to be inserted in the adjacent text field. You can insert up to 10,000 logical volumes at one time.

Note: Each logical volume is inserted sequentially. Several factors, such as current subsystem workload and hardware configuration, can increase the time required for each individual logical volume insert process.

- ▶ **Ending VOLSER:** Select this option to insert a range of logical volumes. Enter the ending VOLSER number in the adjacent text field.
- ▶ **Initially owned by:** The name of the cluster that will own the new logical volumes. Select a cluster from the drop-down menu.
- ▶ **Media type:** Media type of the logical volume (or volumes). Possible values are:
 - Cartridge System Tape (400 MiB)
 - Enhanced Capacity Cartridge System Tape (800 MiB)
- ▶ **Set Constructs:** Select this check box to specify constructs for the new logical volume (or volumes). Then use the drop-down menu below each construct to select a predefined construct name. You can specify any or all of the following constructs:
 - Storage Group
 - Storage Class
 - Data Class
 - Management Class

To insert a range of logical volumes, complete the fields listed and click **Insert**. You are prompted to confirm your decision to insert logical volumes. To continue with the insert operation, click **Yes**. To abandon the insert operation without inserting any new logical volumes, click **No**.

Delete Logical Volumes window

Use this window (Figure 8-26) to delete unused logical volumes from the TS7700 Virtualization Engine.

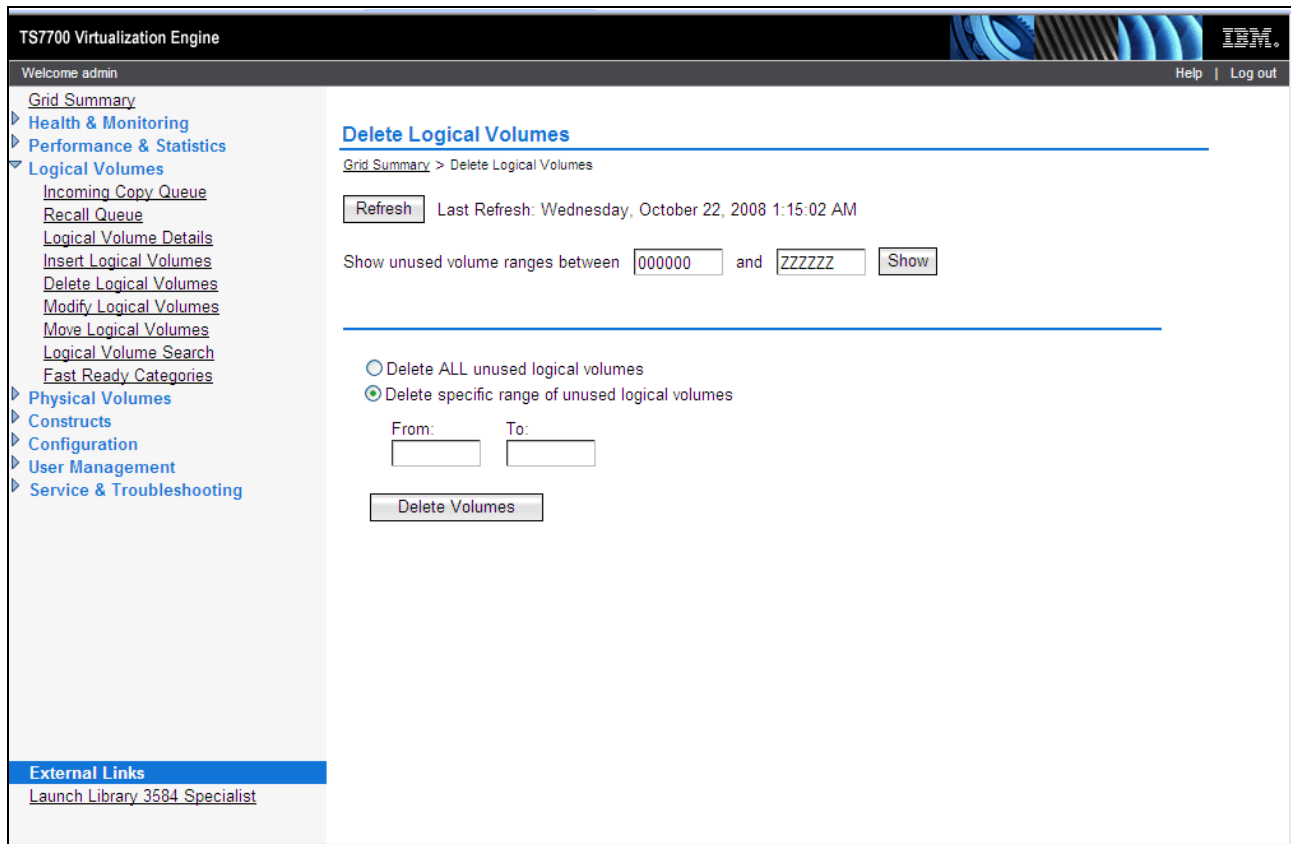


Figure 8-26 TS7700 Virtualization Engine MI Delete Logical Volumes window

To delete unused logical volumes, select one of the options described below and click the **Delete Volumes** button. A confirmation window will be displayed; click **Yes** to delete or **No** to cancel. To view the current list of unused logical volume ranges in the TS7700 Virtualization Engine Grid, enter a logical volume range at the bottom of the window and click **Show**. A logical volume range deletion can be cancelled while in progress at the Cluster Operation History window.

Important: A volume can only be ejected from the insert category if the volume has never been moved out of the insert category after initial insert and has never received write data from a host.

This window has the following options:

- ▶ Delete ALL unused logical volumes: Deletes all unused logical volumes across all VOLSER ranges.
- ▶ Delete specific range of unused logical volumes: All unused logical volumes in the entered VOLSER range will be deleted.
- ▶ From: The start of the VOLSER range to be deleted if **Delete specific range of unused logical volumes** is selected.
- ▶ To: The end of the VOLSER range to be deleted if **Delete specific range of unused logical volumes** is selected.

Modify Logical Volumes window

Use this window (see Figure 8-27) to modify the constructs associated with existing logical volumes in the TS7700 Virtualization Engine composite library.

TS7700 Virtualization Engine

Welcome admin

Grid Summary

Health & Monitoring

Performance & Statistics

Logical Volumes

Incoming Copy Queue

Recall Queue

Logical Volume Details

Insert Logical Volumes

Delete Logical Volumes

Modify Logical Volumes

Move Logical Volumes

Logical Volume Search

Fast Ready Categories

Physical Volumes

Constructs

Configuration

User Management

Service & Troubleshooting

External Links

Launch Library 3584 Specialist

Modify Logical Volumes

Grid Summary > Modify Logical Volumes

Refresh Last Refresh: Wednesday, October 22, 2008 1:25:24 AM

The modify logical volumes function is only valid for logical volumes used by non-MVS hosts. MVS hosts will automatically assign constructs for logical volumes.

Show inserted volume ranges between 000000 and ZZZZZZ Show

Volume Range:

*From: *To:

Storage Groups: Storage Classes: Data Classes: Management Classes:

No Change No Change No Change No Change

Modify

Figure 8-27 TS7700 Virtualization Engine MI Modify Logical Volumes window

To display a range of existing logical volumes, enter the starting and ending VOLSERS in the fields at the top of the page and click **Show**.

To modify constructs for a range of logical volumes, identify a Volume Range, then use the Constructs drop-down menus described as follows to select construct values and click **Modify**:

- ▶ Volume Range: The range of logical volumes to be modified.
 - From: The first VOLSER in the range.
 - To: The last VOLSER in the range.
- ▶ Constructs: Use the following drop-down menus to change one or more constructs for the identified Volume Range. (From each drop-down menu, you can select a predefined construct to apply to the Volume Range, No Change to retain the current construct value, or dashes (-----) to restore the default construct value.)
 - Storage Groups: Use this drop-down menu to change the Storage Group for the identified Volume Range.
 - Storage Classes: Use this drop-down menu to change the Storage Class for the identified Volume Range.

- Data Classes: Use this drop-down menu to change the Data Class for the identified Volume Range.
- Management Classes: Use this drop-down menu to change the Management Class for the identified Volume Range.

You are asked to confirm your decision to modify logical volume constructs. To continue with the operation, click **Yes**. To abandon the operation without modifying any logical volume constructs, click **No**.

Move Logical Volumes window

Use this window (Figure 8-28) to move a range of logical volumes used by the TS7740 Virtualization Engine to a target pool, or cancel a move request already in progress.

If a move operation is already in progress, a warning message will be displayed. You can view move operations already in progress from the Operation History window.

TS7700 Virtualization Engine

Welcome admin

Grid Summary

Health & Monitoring

Performance & Statistics

Logical Volumes

Incoming Copy Queue

Recall Queue

Logical Volume Details

Insert Logical Volumes

Delete Logical Volumes

Modify Logical Volumes

Move Logical Volumes

Logical Volume Search

Fast Ready Categories

Physical Volumes

Constructs

Configuration

User Management

Service & Troubleshooting

External Links

Launch Library 3584 Specialist

Help | Log out

"Juno[0]" (#BA41A): Move Logical Volumes

Grid Summary > "Juno" Cluster Summary > Move Logical Volumes

Refresh Last Refresh: Wednesday, October 22, 2008 1:37:22 AM

This page allows you to move all logical volumes from a physical volume or range of physical volumes, to a target pool.

Current component selected:

- Clusters
 - "Juno[0]" (#BA41A)
 - "diana[1]" (#BA41B)
 - "herc[2]" (#BA41C)

Cancel Move Requests>>

Physical Volume Range:

From: To:

Existing Ranges:

--Select An Existing Range--

*Target Pool: Media Type:

--Select a Pool-- --Any 3592--

*Move Type:

Deferred Move

Priority Move

☐ Honor Inhibit Reclaim Schedule

Move

Figure 8-28 MI Move Logical Volumes

To cancel a move request, select the **Cancel Move Requests** link. The available options to cancel a move request are:

- ▶ Cancel All Moves: Cancels all move requests.
- ▶ Cancel Priority Moves Only: Cancels only priority move requests.
- ▶ Cancel Deferred Moves Only: Cancels only deferred move requests.
- ▶ Select a Pool: Cancels move requests from the designated source pool (1-32), or from all source pools.

If you want to move logical volumes, you must define a volume range or select an existing range, select a target pool, and identify a move type:

- ▶ **Physical Volume Range:** The range of physical volumes to move. You may use either this option or Existing Ranges to define the range of volumes to move, but not both.
From: VOLSER of the last physical volume in the range to move.
To: VOLSER of the first physical volume in the range to move.
- ▶ **Existing Ranges:** The list of existing physical volume ranges. You may use either this option or Volume Range to define the range of volumes to move, but not both.
- ▶ **Media Type:** The media type of the physical volumes in the range to move. If no available physical stacked volume of the given media type is in the range specified, no logical volume is moved.
- ▶ **Target Pool:** The number (1-32) of the source pool from which logical volumes are moved.
- ▶ **Move Type:** Used to determine when the move operation will occur. Possible values are:
 - **Deferred:** Move operation will occur in the future as schedules permit.
 - **Priority:** Move operation will occur as soon as possible.
 - **Honor Inhibit Reclaim Schedule:** An option of the Priority Move Type, it specifies that the move schedule will occur in conjunction with Inhibit Reclaim schedule. If this option is selected, the move operation will not occur when Reclaim is inhibited.

After you define your move operation parameters and click the **Move** button, you will be asked to confirm your request to move physical volumes. If you select **No**, you will return to the Move Logical Volumes window.

Logical Volumes Search window

Use this window (see Figure 8-29) to search for logical volumes in a given TS7700 Virtualization Engine Cluster according to VOLSER, category, media type, expiration date, or inclusion in a group or class.

IBM System Storage™ TS7700 Virtualization Engine

Welcome admin

[Grid Summary](#)
▶ [Health & Monitoring](#)
▶ [Performance & Statistics](#)
▶ [Logical Volumes](#)
 [Incoming Copy Queue](#)
 [Recall Queue](#)
 [Logical Volume Details](#)
 [Insert Logical Volumes](#)
 [Delete Logical Volumes](#)
 [Modify Logical Volumes](#)
 [Move Logical Volumes](#)
 [Logical Volume Search](#)
 [Fast Ready Categories](#)
▶ [Physical Volumes](#)
▶ [Constructs](#)
▶ [Configuration](#)
▶ [User Management](#)
▶ [Service & Troubleshooting](#)

[External Links](#)
[Launch Library 3584 Specialist](#)

Make selections below to define the database search criteria. The more search criteria used, the more restrictive the search.

Current cluster selected:

- Clusters
 - Archie[0]* (#BA96A)
 - Veronica[1]* (#BA96B)
 - Dohnnut[2]* (#BA96C)
 - Veronica1[3]* (#BA96D)

Previous Searches

New Search Name:

Search Options

Volser: Category: Media Type:

Equal to Expire Time: Time Units:

Removal Residency: Removed Time:

Storage Group: or --Storage Groups in Archie--

Management Class: or --Management Classes in Archie--

Storage Class: or --Storage Classes in Archie--

Data Class: or --Data Classes in Archie--

Volser Flags:

	Yes	No	Ignore
Mounted	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Logical WORM	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>

Search Results Options

Figure 8-29 MI Logical Volume Search entry window

You can view the results of a previous query, or create a new query to search for logical volumes.

Note: Only one search can be executed at a time. If a search is in progress, an information message will occur at the top of the Logical volumes search window. You can cancel a search in progress by clicking the **Cancel Search** button.

To view the results of a previous search query, select the **Previous Searches** hyperlink to see a table containing a list of previous queries. Click a query name to display a list of logical volumes that match the search criteria.

Up to 10 previously named search queries can be saved. To clear the list of saved queries, check the check box adjacent to one or more queries to be removed, select **Clear** from the drop-down menu, and click **Go**. This operation will not clear a search query already in progress. You will be asked to confirm your decision to clear the query list. Select **Yes** to clear the list of saved queries, or **No** to retain the list of queries.

To create a new search query, enter a name for the new query. Enter a value for any of the fields and select **Search** to initiate a new logical volume search. The query name, criteria, start time, and end time are saved along with the search results.

Enter your parameters in the New Search Name and Volser fields and then click **Search**, as shown in Figure 8-30, where we prepare to search for logical volume 640104.

IBM System Storage™ TS7700 Virtualization Engine

Welcome admin

Grid Summary

▶ Health & Monitoring

▶ Performance & Statistics

▼ Logical Volumes

Incoming Copy Queue

Recall Queue

Logical Volume Details

Insert Logical Volumes

Delete Logical Volumes

Modify Logical Volumes

Move Logical Volumes

Logical Volume Search

Fast Ready Categories

▶ Physical Volumes

▶ Constructs

▶ Configuration

▶ User Management

▶ Service & Troubleshooting

External Links

Launch Library 3584 Specialist

Make selections below to define the database search criteria. The more search criteria used, the more restrictive the search.

Current cluster selected:

Clusters

- Archie[0] (#BA96A)
- Veronica[1] (#BA96B)
- Dohnnut[2] (#BA96C)
- Veronca[3] (#BA96D)

Previous Searches

New Search Name: 640104

▼ Search Options

Volser: 640104

Category:

Media Type:

Equal to

Expire Time:

Time Units: Not Set

Removal Residency:

Removed Time:

Storage Group:

or

--Storage Groups in Archie--

Management Class:

or

--Management Classes in Archie--

Storage Class:

or

--Storage Classes in Archie--

Data Class:

or

--Data Classes in Archie--

Volser Flags:

	Yes	No	Ignore
Mounted	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Logical WORM	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>

▶ Search Results Options

Search

Figure 8-30 MI Entering Data in the Logical Volume Search window

If you are looking for the results of earlier searches, click **Previous Searches** on the Logical Volume Search window, shown in Figure 8-29 on page 452, and the window shown in Figure 8-31 opens.

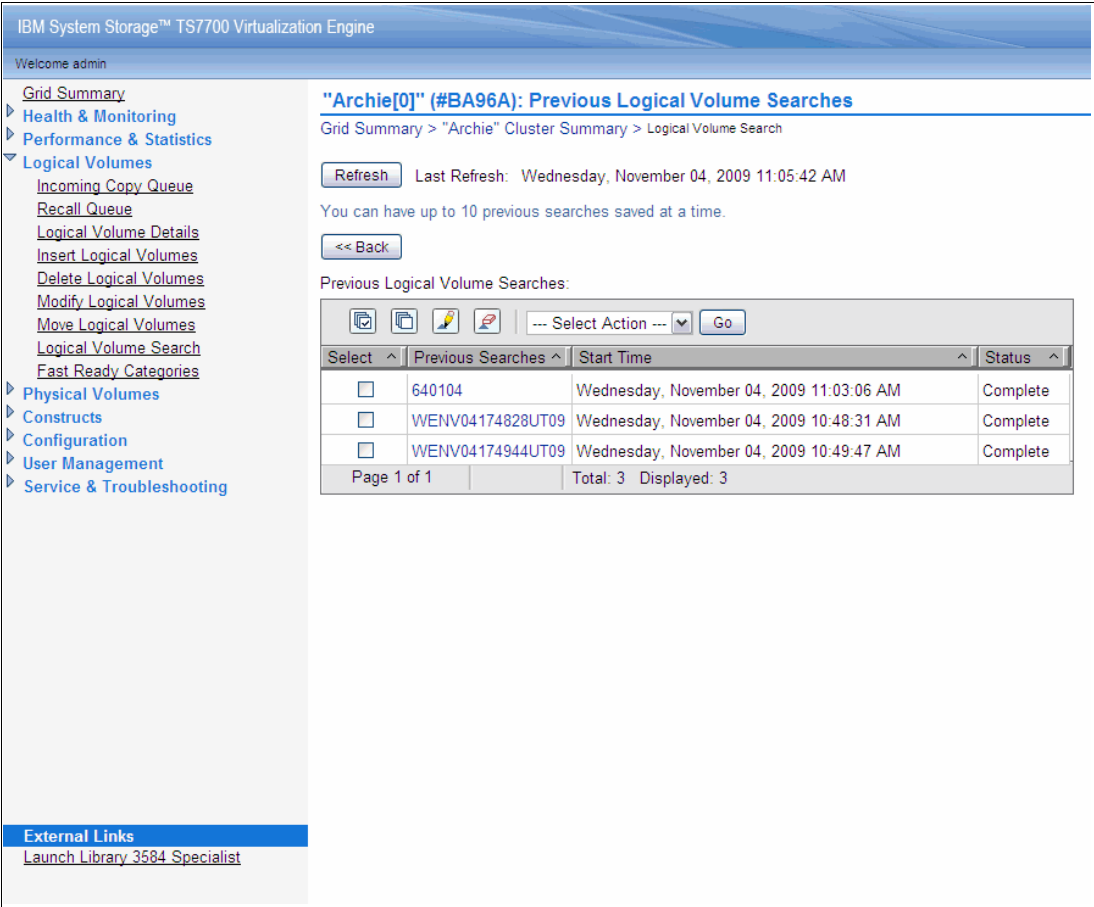


Figure 8-31 MI Logical Volume Search previous searches

In the table of the last ten searches (Figure 8-31 on page 454), click your search name and the search result window opens, as shown in Figure 8-32.

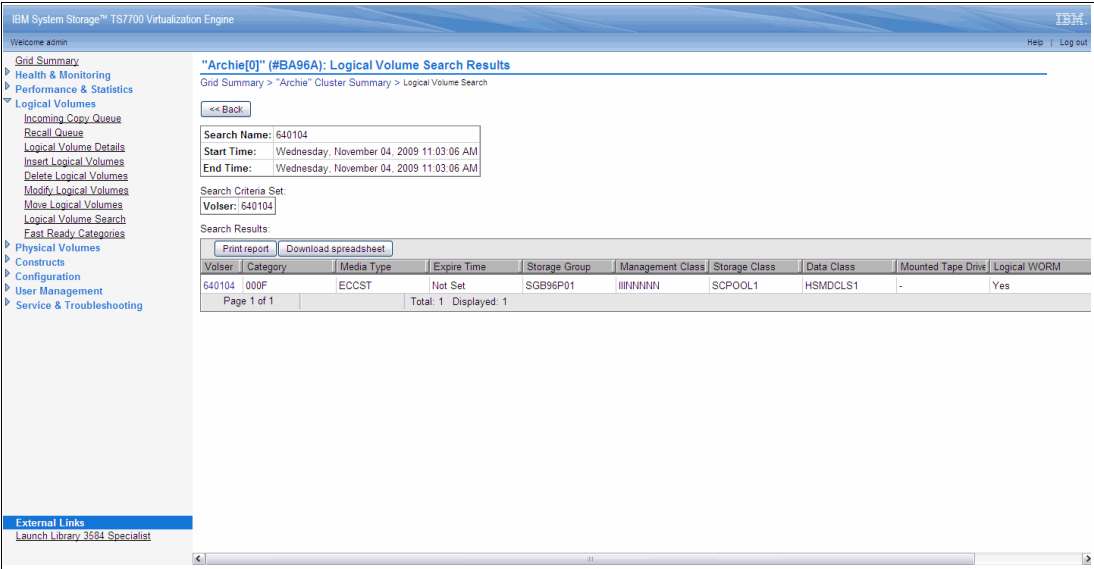


Figure 8-32 MI Logical Volume Search Results window

The meanings of the entry fields in Figure 8-29 on page 452, where you can enter your Logical Volume search, are:

- ▶ **VOLSER:** The volume's serial number. This field can be left blank. You can also use the following wild card characters in this field:
 - % (percent): Represents zero or more characters.
 - * (asterisk): Translated to % (percent). Represents zero or more characters.
 - . (period): Translated to _ (single underscore). Represents one character.
 - _ (single underscore): Represents one character.
 - ? (question mark): Translated to _ (single underscore). Represents one character.

- **Category:** The name of the category to which the logical volume belongs. This value is a four-character hexadecimal string. This field can be left blank. See Table 8-2 for possible values.

Table 8-2 Category values

Hexadecimal value	Category
0001-FEFF	General programming use
FF00	Insert
FF01	Insert for VTS stacked volumes
FF03	Scratch for VTS stacked volumes
FF04	Private for VTS stacked volumes
FF05	Disaster recovery for VTS stacked volumes
FF06	Disaster recovery for VTS stacked volumes
FF08	VTS stacked volume internal label is unreadable
FF09	Temporary category used by the VTS
FF10	Convenience eject
FF11	High-capacity eject
FF12	Export pending volumes
FF13	Exported volumes
FF14	Import volumes
FF15	Import pending volumes
FF16	Unassigned volumes (VTS import)
FF17	Export-hold volumes
FF20	Volume with corrupted tokens (Peer-to-peer VTS usage only)
FFF4	Cleaner volume (ETC)
FFF5	Service volume (ETC)
FFF6	Service volume (HPCT/EHPCT)
FFF9	Service volume (CST/ECCST)
FFFA	Manually ejected
FFFD	Cleaner volume (HPCT/EHPCT)
FFFE	Cleaner volume (CST/ECCST)
FFFF	VOLSER specific

- **Media Type:** The type of media on which the volume resides. Use the drop-down menu to select from the available media types. This field can be left blank.
- **Expire Time:** The amount of time in which logical volume data will expire. Enter a number. This field is qualified by the values Equal to, Less than, or Greater than in the preceding drop-down menu and defined by the succeeding drop-down menu under the heading Time Units. This field may be left blank.

- ▶ **Residency:** The residency state of the logical volume. Possible values are:
 - Resident: The search includes only logical volumes classified as resident.
 - Ignore: The search ignores any values in the Residency field. This is the default selection.
 - Removed Before: The search includes only logical volumes removed prior to a given date and time. If you select this value, you must also complete the fields for Removed Date and Time.
 - Removed After: The search includes only logical volumes removed after a given date and time. If you select this value, you must also complete the fields for Removed Date and Time.
 - Retained: The search includes only logical volumes classified as retained.
 - Deferred: The search includes only logical volumes classified as deferred.
- ▶ **Removed Date and Time:** The date and time that informs the Removed Before or Removed After search queries. These values reflect the time zone in which your browser is located.
 - Date: The calendar date according to month (M), day (D), and year (Y); it uses the format MM/DD/YYYY. This field is accompanied by a date chooser calendar icon. You can enter the month, day, and year manually, or you can use the calendar chooser to select a specific date. The default for this field is blank.
 - Time: The Coordinated Universal Time (UTC) in hours (H), minutes (M), and seconds (S). The values in this field must take the form HH:MM:SS. Possible values for this field include 00:00:00 through 23:59:59. This field is accompanied by a time chooser clock icon. You can enter hours and minutes manually using 24 hour time designations, or you can use the time chooser to select a start time based on a 12 hour (a.m./p.m.) clock. The default for this field is midnight (12:00:00).
- ▶ **Time Units:** Units of time used to define Expire Time. This field is required if the Expire Time field is completed. Possible values include:
 - Seconds.
 - Minutes.
 - Hours.
 - Days.
 - Weeks.
 - Not Set: The default entry. This value is only accepted if the Expire Time field is blank.
 - Expired: This value is only accepted if the Expire Time field is blank.
- ▶ **Storage Group:** The name of the storage group in which the logical volume resides. You can enter a name in the empty field, or select a name from the adjacent drop-down menu. This field can be left blank.
- ▶ **Management Class:** The name of the Management Class to which the logical volume belongs. You can enter a name in the empty field, or select a name from the adjacent drop-down menu. This field can be left blank.
- ▶ **Storage Class:** The name of the storage class to which the logical volume belongs. You can enter a name in the empty field, or select a name from the adjacent drop-down menu. This field can be left blank.
- ▶ **Data Class:** The name of the data class to which the logical volume belongs. You can enter a name in the empty field, or select a name from the adjacent drop-down menu. This field can be left blank.

- ▶ **Mounted:** Whether the logical volume is mounted. Possible values are:
 - Ignore: The search ignores any values in the Mounted field. This is the default selection.
 - Yes: The search includes only mounted logical volumes.
 - No: The search includes only unmounted logical volumes.
- ▶ **Logical WORM:** Whether the logical volume is defined as write-once, read-many (WORM). Possible values are:
 - Ignore: The search ignores any values in the Logical WORM field. This is the default selection.
 - Yes: The search includes only WORM logical volumes.
 - No: The search does not include any WORM logical volumes.

Note: You can print or download the results of a search query using the Print Report or Download Spreadsheet buttons on the Volumes found: table at the end of Search Results window, as shown in Figure 8-32 on page 455.

Fast Ready Categories window

Use this window (Figure 8-33) to add, modify, or delete a Fast Ready category of physical volumes.

TS7700 Virtualization Engine

Welcome admin

Grid Summary

Health & Monitoring

Performance & Statistics

Logical Volumes

Physical Volumes

Constructs

Configuration

User Management

Service & Troubleshooting

External Links

Launch Library 3584 Specialist

Fast Ready Categories

Grid Summary > Fast Ready Categories

Last Refresh: Wednesday, October 22, 2008 10:43:27 AM

Fast-Ready Categories:

Select	Category	Expire Time	Time Units	Expire Hold
<input type="radio"/>	0001		2 Days	No
<input type="radio"/>	0002		3 Days	Yes

Page 1 of 1 Total: 2 Displayed: 2

Figure 8-33 MI Fast Ready Categories window

A category is a logical grouping of physical volumes for a predefined use. A Fast Ready category groups logical volumes for quick, non-specific use. This grouping enables quick mount times because the TS7700 Virtualization Engine can order category mounts without recalling data from a stacked volume.

The Fast Ready Categories Table lists all Fast Ready categories and their attributes. If no Fast Ready categories are defined, this table will not be visible.

Note: Fast Ready categories and Data Classes work at the system level and are unique for all clusters in a Grid. This means that if you modify them in one cluster it will apply to all clusters on that Grid.

You can use the Fast Ready Categories table to create a new or modify or delete an existing Fast Ready category.

Referring to Figure 8-33 on page 458, information displayed in the Fast Ready Categories Table includes:

- ▶ **Category:** The Fast Ready category number. This is a four-digit hexadecimal number created by a user when he or she adds a new category. Valid characters for this field are:
 - A-F, 0-9

Note: You may not use category name 0000 or “FFxx”, where xx equals 0–9 or A-F. 0000 represents a null value, and “FFxx” is reserved for hardware.

- ▶ **Expire Time:** The number of hours or days in which logical volume data categorized as Fast Ready will expire. Used in conjunction with Time Units. If this field is set to 0, the categorized data will never expire. The minimum Expire Time is 24 hours and the maximum Expire Time is 195 weeks, 1,365 days, or 32,767 hours.
- ▶ **Time Units:** The unit of time used to define Expire Time. Possible values are:
 - Hours: Select if your Expire Time will be defined by the number of hours.
 - Days: Select if your Expire Time will be defined by the number of days.
 - Weeks: Select if your Expire Time will be defined by the number of weeks.
- ▶ **Expire Hold:** Whether a logical volume in a given category may be mounted or moved to another category before its data expires. Possible values are:
 - Yes: The logical volume may *not* be mounted or moved to another category before its data expires.
 - No: The logical volume may be mounted or moved to another category before its data expires.

Use the drop-down menu on the Fast Ready Categories Table to add a new category or modify or delete an existing category.

To add a new category, select **Add** from the drop-down menu. Complete the fields for the information that will be displayed in the Fast Ready Categories Table, as defined previously.

Note: You must choose either the **No expiration** or **Set expiration** radio button before the new category can be created. If you select **Set expiration**, you must complete the Expire Time and Time Units fields.

To modify an existing Fast Ready category, click the radio button from the **Select** column that appears adjacent to the number of the category you want to modify. Select **Modify** from the drop-down menu. You will be able to change any field displayed on the Fast Ready Categories Table.

To delete an existing Fast Ready category, click the radio button from the **Select** column that appears adjacent to the number of the category you want to delete. Select **Delete** from the drop-down menu.

8.2.6 Physical volumes

The topics in this section present information related to monitoring and manipulating physical volumes in the TS7740 Virtualization Engine.

Use this window (Figure 8-34) to view or modify settings for physical volume pools to manage the physical volumes used by the TS7740 Virtualization Engine. If the Grid does not possess a physical library, this page is not visible on the TS7700 Virtualization Engine management interface.

To watch a tutorial showing how to modify pool encryption settings, click the **View tutorial** link.

The screenshot shows the TS7700 Virtualization Engine management interface. The left sidebar contains a navigation menu with options like Grid Summary, Health & Monitoring, Performance & Statistics, Logical Volumes, Physical Volumes, Constructs, Configuration, User Management, and Service & Troubleshooting. The main content area is titled '"juno[0]" (#BA41A): Physical Volume Pools'. It includes a 'Refresh' button and a 'Last Refresh' timestamp. Below this is a 'View tutorial' link. A tree view shows the current component selected: Clusters > "juno[0]" (#BA41A). The 'Physical volume pool properties' section has two tabs: 'Pool Properties' (active) and 'Encryption Settings'. The 'Pool Properties' tab displays a table with columns for Select, Pool, Me..., Fir..., Se..., Bor..., Re..., Ma..., Da..., Da..., Ag..., Da..., Ma... The table lists 9 physical volume pools.

Select	Pool	Me...	Fir...	Se...	Bor...	Re...	Ma...	Da...	Da...	Ag...	Da...	Ma...
<input type="checkbox"/>	1	3592	Any...	None	Borr...	11	All ...	0	0	1	0	0
<input type="checkbox"/>	2	3592	JA...	JJ-E...	Borr...	12	All ...	0	1	0	0	0
<input type="checkbox"/>	3	3592	Any...	None	Borr...	13	All ...	0	0	0	1	95
<input type="checkbox"/>	4	3592	Any...	None	Borr...	4	All ...	0	0	0	0	0
<input type="checkbox"/>	5	3592	Any...	None	Borr...	5	All ...	0	0	0	0	0
<input type="checkbox"/>	6	3592	Any...	None	Borr...	6	All ...	0	0	0	0	0
<input type="checkbox"/>	7	3592	Any...	None	Borr...	7	All ...	0	0	0	0	0
<input type="checkbox"/>	8	3592	Any...	None	Borr...	8	All ...	0	0	0	0	0
<input type="checkbox"/>	9	3592	Any...	None	Borr...	9	All ...	0	0	0	0	0

Figure 8-34 MI Physical Volume Pools

Physical Volume Pools windows

The Physical Volume Pool Properties table displays the media properties and encryption settings and for every physical volume pool defined for a given cluster in the Grid. This table contains two tabs:

- Pool Properties
- Encryption Settings

Note: Pools 1-32 are preinstalled. Pool 1 functions as the default pool and will be used if no other pool is selected. All other pools must be defined before they can be selected.

Pool Properties tab

The Pool Properties tab displays properties of physical volume pools. Pool information displayed on this tab includes:

- ▶ Pool: The pool number. This is a whole number between 1 and 32, inclusive.
- ▶ Encryption: The encryption state of the pool. Possible values are:
 - Enabled: Encryption is enabled on the pool.
 - Disabled: Encryption is not enabled on the pool.
- ▶ Media Class: The supported media class of the storage pool. The only possible value is 3592.
- ▶ First Media (Primary): The primary media type that the pool can borrow or return to the common scratch pool (Pool 0). Possible values are:
 - Any 3592: Any available 3592 cartridge can be used.
 - JA: Enterprise Tape Cartridge (ETC).
 - JB: Enterprise Extended-Length Tape Cartridge (ETCL).
 - JJ: Enterprise Economy Tape Cartridge (EETC).

To modify pool properties, check the check box next to one or more pools listed in the Physical Volume Pool Properties table and select **Properties** from the drop-down menu. The Pool Properties table will be displayed. You will be able to modify the fields Media Class and First Media, defined previously, and the following fields:

- ▶ Second Media (Secondary): The second choice of media type that the pool can borrow from. The options shown will exclude the media type chosen for First Media. Possible values are:
 - Any 3592: Any available 3592 cartridge can be used.
 - JA: Enterprise Tape Cartridge (ETC).
 - JB: Enterprise Extended-Length Tape Cartridge (ETCL).
 - JJ: Enterprise Economy Tape Cartridge (EETC).
 - None: The only option available if the Primary Media type is Any 3592.
- ▶ Borrow Indicator: Defines how the pool is populated with scratch cartridges. Possible values are:
 - Borrow, Return: A cartridge is borrowed from the Common Scratch Pool (CSP) and returned when emptied.
 - Borrow, Keep: A cartridge is borrowed from the CSP and retained, even after being emptied.
 - No Borrow, Return: A cartridge may not be borrowed from CSP, but an emptied cartridge will be placed in CSP. This setting is used for an empty pool.
 - No Borrow, Keep: A cartridge may not be borrowed from CSP, and an emptied cartridge will be retained.
- ▶ Reclaim Pool: The pool to which logical volumes are assigned when reclamation occurs for the stacked volume on the selected pool.
- ▶ Maximum Devices: The maximum number of physical tape drives that the pool can use for premigration.

- ▶ **Export Pool:** The type of export supported if the pool is defined as an Export Pool (the pool from which physical volumes are exported). Possible values are:
 - **Not Defined:** The pool is not defined as an Export pool
 - **Copy Export:** The pool is defined as a Copy Export pool.
- ▶ **Days Before Secure Data Erase:** The number of days a physical volume that is a candidate for Secure Data Erase can remain in the pool without access to a physical stacked volume. Each stacked physical volume possesses a timer for this purpose, which is reset when a logical volume on the stacked physical volume is accessed. Secure Data Erase occurs at a later time, based on an internal schedule. Secure Data Erase renders all data on a physical stacked volume inaccessible. The valid range of possible values is 1-365. Clearing the check box deactivates this function.
- ▶ **Days Without Access:** The number of days the pool can persist without access to set a physical stacked volume. Each physical stacked volume has a timer for this purpose, which is reset when a logical volume is accessed. The reclamation occurs at a later time, based on an internal schedule. The valid range of possible values is 1-365. Clearing the check box deactivates this function.
- ▶ **Age of Last Data Written:** The number of days the pool has persisted without write access to set a physical stacked volume as a candidate for reclamation. Each physical stacked volume has a timer for this purpose, which is reset when a logical volume is accessed. The reclamation occurs at a later time, based on an internal schedule. The valid range of possible values is 1-365. Clearing the check box deactivates this function.
- ▶ **Days Without Data Inactivation;** The number of sequential days the pool's data ratio has been higher than the Maximum Active Data to set a physical stacked volume as a candidate for reclamation. Each physical stacked volume has a timer for this purpose, which is reset when data inactivation occurs. The reclamation occurs at a later time, based on an internal schedule. The valid range of possible values is 1-365. Clearing the check box deactivates this function. If deactivated, this field will not be used as a criteria for reclaim.
- ▶ **Maximum Active Data:** The ratio of the amount of active data in the entire physical stacked volume capacity. This field is used in conjunction with Days Without Data Inactivation. The valid range of possible values is 5 - 95 (%). This function is disabled if Days Without Data Inactivation is not checked.
- ▶ **Reclaim Threshold:** The percentage used to determine when to perform reclamation of free storage on a stacked volume. When the amount of active data on a physical stacked volume drops below this percentage, a reclamation operation will be performed on the stacked volume. The valid range of possible values is 5 - 95 (%). Clearing the check box deactivates this function; 10% is the default value.

To modify encryption settings for one or more physical volume pools, select the **Encryption Settings** tab, then select the check box next to one or more volume pools. From the drop-down menu, select **Modify Encryption Settings** and the Modify Encryption Settings window will open, as shown in Figure 8-35.

IBM System Storage™ TS7700 Virtualization Engine

Welcome admin

Grid Summary
 Health & Monitoring
 Performance & Statistics
 Logical Volumes
 Physical Volumes
 Physical Volume Pools
 Physical Volume Details
 Move Physical Volumes
 Eject Physical Volumes
 Physical Volume Ranges
 Physical Volume Search
 Constructs
 Configuration
 User Management
 Service & Troubleshooting

External Links
 Launch TS3500 Library Specialist

"Pesto[0]" (#BA68A): Modify Encryption Settings

Refresh Last Refresh: Monday, September 27, 2010 2:47:34 PM

Selected pools:
 Pool
 23

Encryption Settings:
☒ Enabled ☐ Disabled
☒ Use Encryption Key Manager default key
 Key Mode 1:
 Hash Label
 Key Label 1:
 HYDRA_KEY1
☐ Use Encryption Key Manager default key
 Key Mode 2:
 Hash Label
 Key Label 2:
 HYDRA_KEY1
 OK Cancel

Figure 8-35 MI Modify Encryption Settings

Encryption Settings tab

The Encryption Settings tab displays the encryption key settings for physical volume pools. Encryption information displayed on this tab includes:

- ▶ Pool: The pool number. This number is a whole number 1 - 32, inclusive.
- ▶ Encryption: The encryption state of the pool. Possible values are:
 - Enabled: Encryption is enabled on the pool.
 - Disabled: Encryption is not enabled on the pool.
- ▶ Use Encryption Key Manager default key: Select this check box to populate the Key Label 1 field using a default key provided by the encryption key manager.

Note: The Use Encryption Key Manager default key check box occurs prior to both Key Label 1 and Key Label 2 fields. You must select this check box for each label to be defined using the default key.

- ▶ Key Label 1: The pool's current encryption key Label 1. The label must consist of ASCII characters and cannot exceed 64 characters. No heading or trailing blank is permitted, though an internal space is allowed. Lower-case characters are internally converted to upper case upon storage, so key labels are reported using upper-case characters.

Note: If the encryption state is disabled, this field is blank.

- ▶ Key Mode 1: Encryption Mode used with Key Label 1. When Key Label 1 is disabled, this field is invalidated. Possible values for this field are:
 - Clear Label: The data key is specified by the key label in clear text.
 - Hash Label: The data key is referenced by a computed value corresponding to its associated public key.
- ▶ Use Encryption Key Manager default key: Check this box to populate the Key Label 2 field using a default key provided by the encryption key manager.

Note: The Use Encryption Key Manager default key check box occurs prior to both Key Label 1 and Key Label 2 fields. You must check this box for each label to be defined using the default key.

- ▶ Key Label 2: The pool's current encryption key Label 2. The label must consist of ASCII characters and cannot exceed 64 characters. No heading or trailing blank is permitted, although an internal space is allowed. Lower-case characters are internally converted to upper case upon storage, so key labels are reported using upper-case characters.

Note: If the encryption state is disabled, this field is blank.

- ▶ Key Mode 2: Encryption Mode used with Key Label 2. When Key Label 2 is disabled, this field is invalidated. Possible values for this field are:
 - Clear Label: The data key is specified by the key label in clear text.
 - Hash Label: The data key is referenced by a computed value corresponding to its associated public key.

Physical Volume Details window

Use this window (Figure 8-36 on page 465) to obtain detailed information about a physical stacked volume in the TS7740 Virtualization Engine.

To obtain details about a physical stacked volume, enter the volume's VOLSER number in the available text field and select the **Get Details** button. The VOLSER number must be six characters in length.

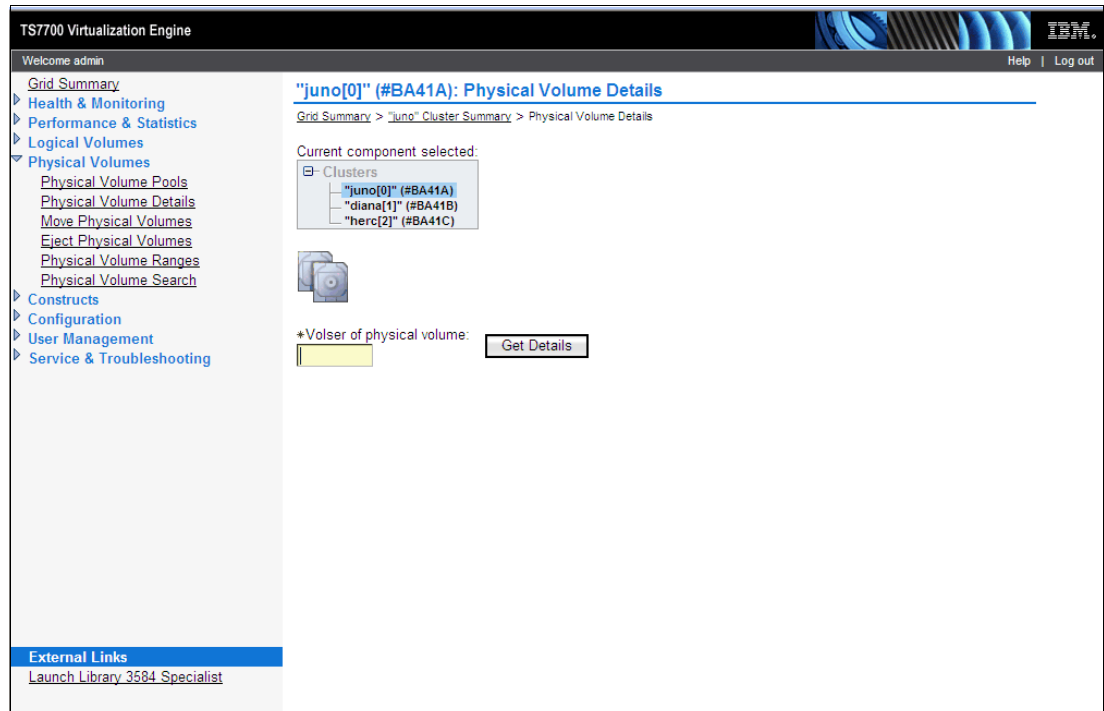


Figure 8-36 MI Physical Volume Details

The window shown in Figure 8-37 is displayed when details for a physical stacked volume are retrieved. The example is for volume 320031.

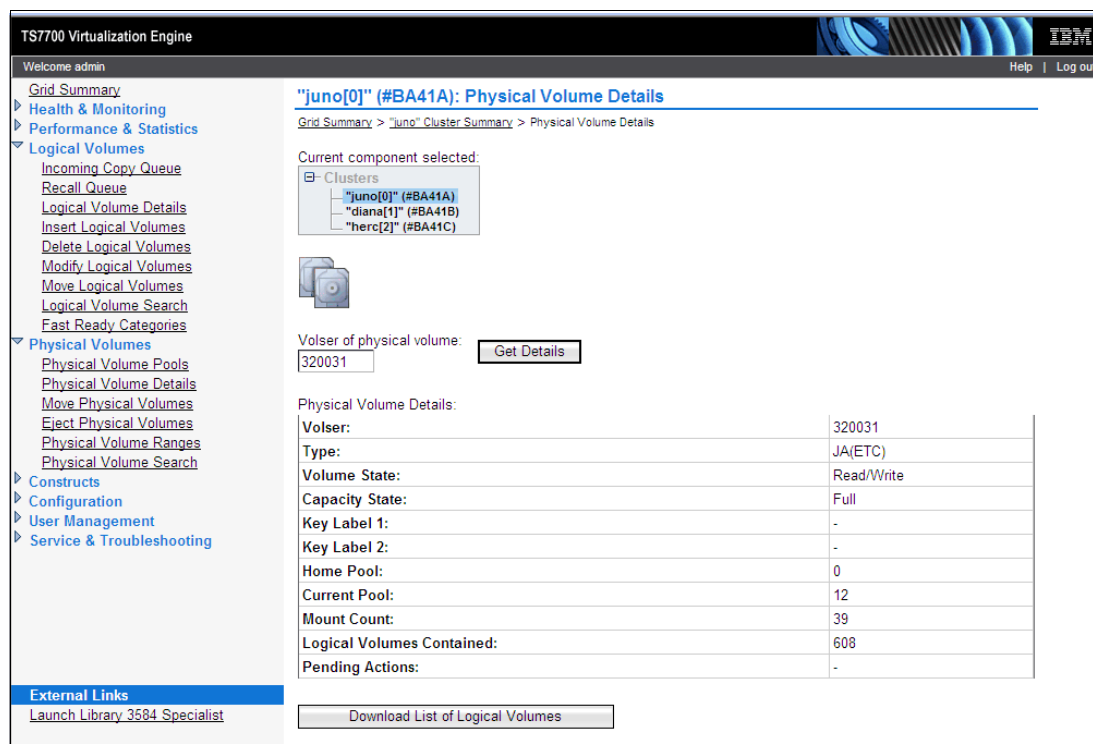


Figure 8-37 MI Physical Volume Details result

To obtain details about the specific logical volumes on the physical stacked volume, click the **Download List of Logical Volumes** button below the table. If the window for the download does not appear and the machine attempting the download is running Windows® XP SP2, click the **Downloading Comma Separated Value (.csv) Files** link at the bottom of this help page for a solution to this problem.

Note: Excel will remove leading zeros from downloaded CSV files. To preserve leading zeroes in files your download, see "Incoming Copy Queue window" on page 436.

You can see information about the following items:

- ▶ **VOLSER:** Six character VOLSER number of the physical stacked volume.
- ▶ **Type:** The media type of the physical stacked volume. Possible values are:
 - JA(ETC): Enterprise Tape Cartridge.
 - JB(ETCL): Enterprise Extended-Length Tape Cartridge.
 - JJ(EETC): Enterprise Economy Tape Cartridge.
- ▶ **Volume state:** Possible values are:
 - Read-Only: The volume is in a read only state.
 - Read/Write: The volume is in a read/write state.
 - Unavailable: The volume is in use by another task or is in a pending eject state.
 - Destroyed: The volume is damaged and unusable for mounting.
- ▶ **Capacity state:** Possible values are empty, filling, and full.

- ▶ Key Label 1 / Key Label 2: The encryption key label that is associated with a physical volume. Up to two key labels might be present. If there are no labels present, then the volume is not encrypted.
- ▶ Home pool: The pool number the physical volume was assigned to when it was inserted into the library, or the pool it was moved to through the management interface Move/Eject Stacked Volumes function.
- ▶ Current pool: The current storage pool in which the physical volume resides.
- ▶ Mount count: The number of times the physical volume has been mounted since being inserted into the library.
- ▶ Logical Volumes Contained: Number of logical volumes contained in this physical stacked volume.
- ▶ Pending Actions: Whether a move or eject operation is pending. Possible values are:
 - Pending eject.
 - Pending priority eject.
 - Pending deferred eject.
 - Pending move to pool #, where # represents the destination pool.
 - Pending priority move to pool #, where # represents the destination pool.
 - Pending deferred move to pool #, where # represents the destination pool.

Move Physical Volumes window

Use this window (Figure 8-38) to move a range or quantity of physical volumes used by the TS7740 Virtualization Engine to a target pool, or cancel a previous move request.

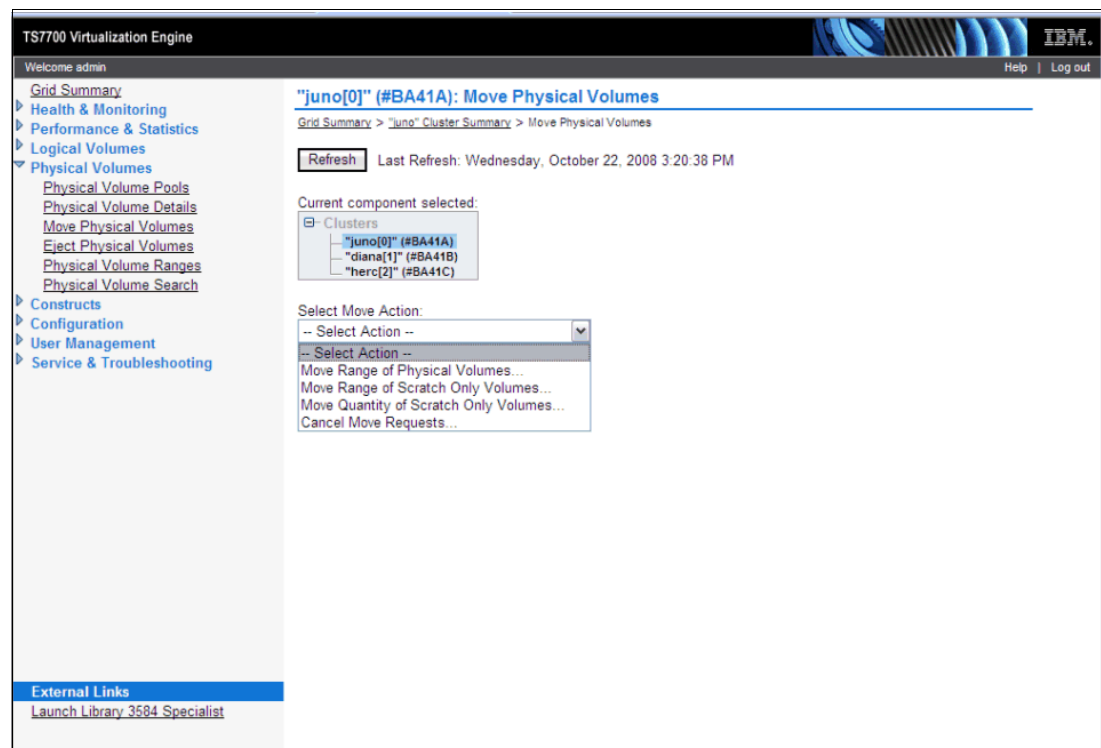


Figure 8-38 MI Move Physical Volumes window

The Select Move Action Menu provides options for moving physical volumes to a target pool. The options available to move physical volumes to a target pool are:

- ▶ **Move Range of Physical Volumes:** Moves physical volumes in the range specified to the target pool.
- ▶ **Move Range of Scratch Only Volumes:** Moves scratch volumes in the range specified to the target pool.
- ▶ **Move Quantity of Scratch Only Volumes:** Moves a specified quantity of physical volumes from the source pool to the target pool.
- ▶ **Cancel Move Requests:** Cancels any previous move request.

If you select **Move Range of Physical Volumes** from the Select Move Action Menu, you will be asked to define a volume range or select an existing range, select a target pool, and identify a move type:

- ▶ **Volume Range:** The range of physical volumes to move. You may use either this option or Existing Ranges to define the range of volumes to move, but not both.
 - **To:** VOLSER of the first physical volume in the range to move.
 - **From:** VOLSER of the last physical volume in the range to move.
- ▶ **Existing Ranges:** The list of existing physical volume ranges. You may use either this option or Volume Range to define the range of volumes to move, but not both.
- ▶ **Target Pool:** The number (0-32) of the source pool from which physical volumes are moved.
- ▶ **Move Type:** Used to determine when the move operation will occur. Possible values are:
 - **Deferred:** The move operation will occur in the future as schedules permit.
 - **Priority:** The move operation will occur as soon as possible.
- ▶ **Honor Inhibit Reclaim Schedule:** An option of the Priority Move Type, it specifies that the move schedule will occur in conjunction with Inhibit Reclaim schedule. If this option is selected, the move operation will not occur when reclamation is inhibited.

If you select **Move Range of Scratch Only Volumes** from the Select Move Action Menu, you will be asked to define a volume range or select an existing range, and select a target pool, as defined previously.

If you select **Move Quantity of Scratch Only Volumes** from the Select Move Action Menu, you will be asked to define the number of volumes to be moved, identify a source and target pool, and select a media type:

- ▶ **Number of Volumes:** The number of physical volumes to be moved.
- ▶ **Source Pool:** The number (0-32) of the source pool from which physical volumes are moved.
- ▶ **Target Pool:** The number (0-32) of the target pool to which physical volumes are moved.
- ▶ **Media Type:** Specifies the media type of the physical volumes in the range to be moved. The physical volumes in the range specified to move must be of the media type designated by this field, or else the move operation will fail.

After you define your move operation parameters and click the **Move** button, you will be asked to confirm your request to move physical volumes. If you select **No**, you will return to the Move Physical Volumes window. To cancel a previous move request, select **Cancel Move Requests** from the Select Move Action Menu. The available options to cancel a move request are:

- ▶ Cancel All Moves: Cancels all move requests.
- ▶ Cancel Priority Moves Only: Cancels only priority move requests.
- ▶ Cancel Deferred Moves Only: Cancels only deferred move requests.
- ▶ Select a Pool: Cancels move requests from the designated source pool (0-32), or from all source pools.

Eject Physical Volumes window

Use this window (Figure 8-39) to eject a range or quantity of physical volumes used by the TS7740 Virtualization Engine or cancel a previous eject request.

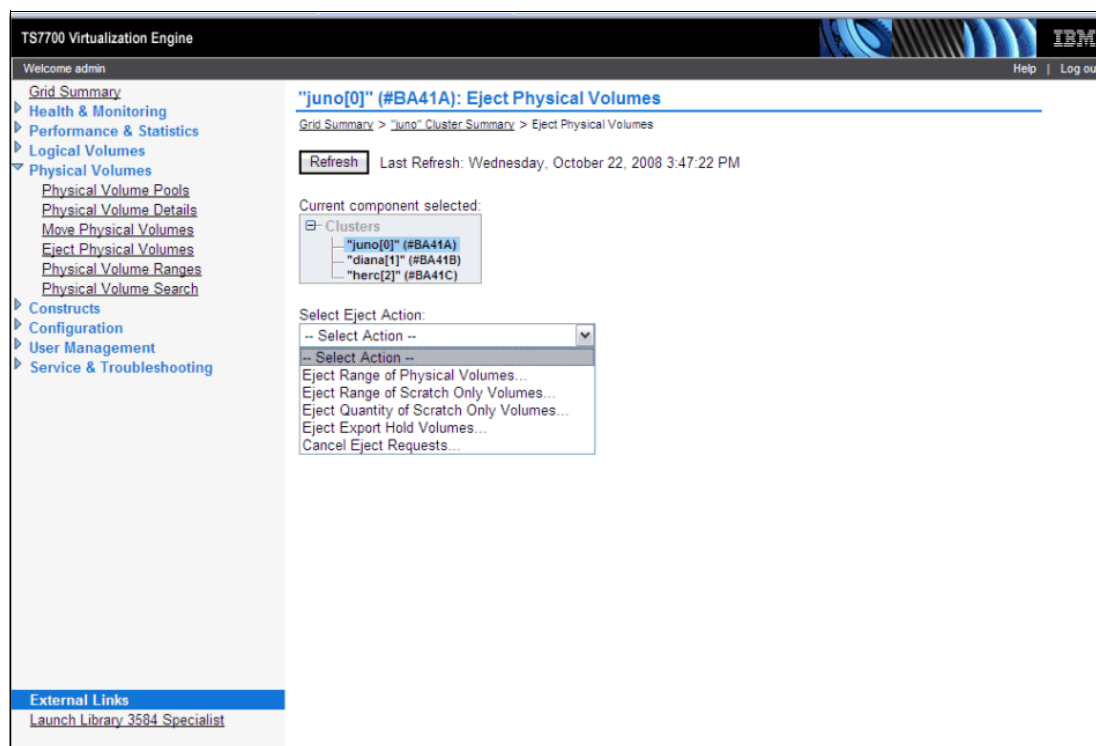


Figure 8-39 MI Eject Physical Volumes window

The Select Eject Action menu provides options for ejecting physical volumes. The options available to eject physical volumes are:

- ▶ Eject Range of Physical Volumes: Ejects physical volumes in the range specified.
- ▶ Eject Range of Scratch Only Volumes: Ejects scratch volumes in the range specified.
- ▶ Eject Quantity of Scratch Only Volumes: Ejects a specified quantity of physical volumes.
- ▶ Eject Export Hold Volumes: Ejects a subset of the volumes in the Export/Hold Category.
- ▶ Cancel Eject Requests: Cancels any previous eject request.

If you select **Eject Range of Physical Volumes** from the Select Eject Action Menu, you will be asked to define a volume range or select an existing range and identify an eject type:

- ▶ **Volume Range:** The range of physical volumes to eject. You may use either this option or Existing Ranges to define the range of volumes to eject, but not both.
 - **To:** VOLSER of the first physical volume in the range to eject.
 - **From:** VOLSER of the last physical volume in the range to eject.
- ▶ **Existing Ranges:** The list of existing physical volume ranges. You may use either this option or Volume Range to define the range of volumes to eject, but not both.
- ▶ **Eject Type:** Used to determine when the eject operation will occur. Possible values are:
 - **Deferred:** The eject operation will occur in the future as schedules permit.
 - **Priority:** The eject operation will occur as soon as possible.
 - **Honor Inhibit Reclaim Schedule:** An option of the Priority Eject Type, it specifies that the eject schedule will occur in conjunction with Inhibit Reclaim schedule. If this option is selected, the eject operation will not occur when reclamation is inhibited.

If you select **Eject Range of Scratch Only Volumes** from the Select Eject Action Menu, you will be asked to define a volume range or select an existing range, as defined previously.

If you select **Eject Quantity of Scratch Only Volumes** from the Select Eject Action Menu, you will be asked to define the number of volumes to be ejected, identify a source pool, and select a media type:

- ▶ **Number of Volumes:** The number of physical volumes to be ejected.
- ▶ **Source Pool:** The number (0-32) of the source pool from which physical volumes are ejected.
- ▶ **Media Type:** Specifies the media type of the physical volumes in the range to be ejected. The physical volumes in the range specified to eject must be of the media type designated by this field, or else the eject operation will fail.

If you select **Eject Export Hold Volumes** from the Select Eject Action Menu, you will be asked to select the VOLSER (or VOLSERs) of the volumes to be ejected. To select all VOLSERs in the Export Hold category, select **Select All** from the drop-down menu.

After you define your eject operation parameters and click the **Eject** button, you will be asked to confirm your request to eject physical volumes. If you select **No**, you will return to the Eject Physical Volumes window.

To cancel a previous eject request, select **Cancel Eject Requests** from the Select Eject Action Menu. The available options to cancel an eject request are:

- ▶ **Cancel All Ejects:** Cancels all eject requests.
- ▶ **Cancel Priority Ejects Only:** Cancels only priority eject requests.
- ▶ **Cancel Deferred Ejects Only:** Cancels only deferred eject requests.

Physical Volume Ranges window

Use this window (Figure 8-40) to add, modify, or delete physical volume ranges in, or eject unassigned physical volumes from, a library attached to a TS7740 Virtualization Engine.

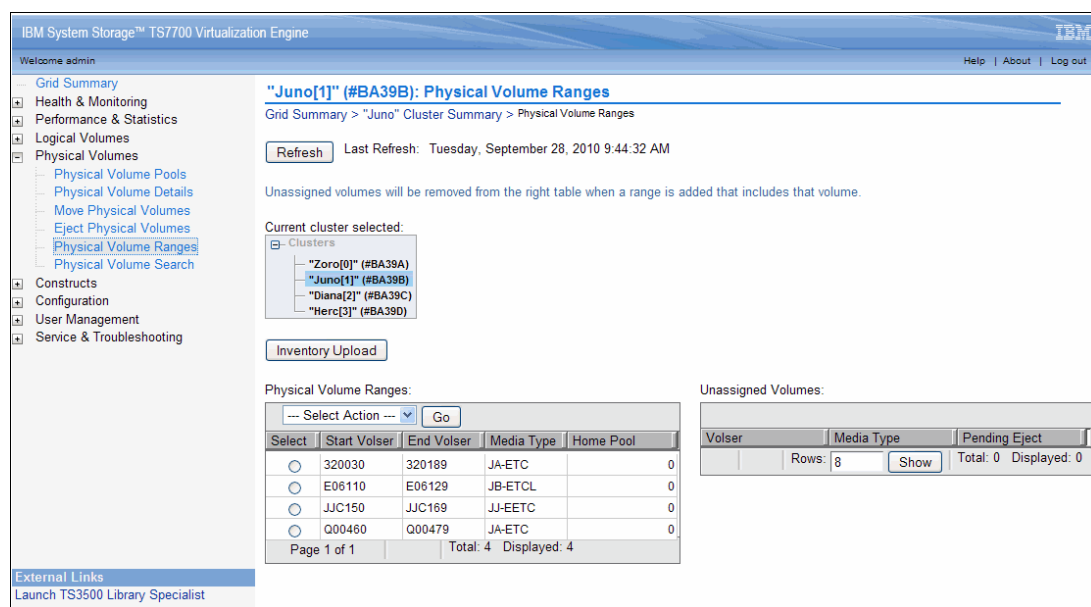


Figure 8-40 MI Physical Volume Ranges window

Click the **Inventory Upload** button to synchronize the physical cartridge inventory from the attached tape library with the TS7740 Virtualization Engine database.

The VOLSER Ranges Table displays the list of defined VOLSER ranges for a given component.

You can use the VOLSER Ranges Table to create a new VOLSER range or modify or delete a predefined VOLSER range.

Note: If a physical volume's VOLSER is defined within a TS3500 Tape Library Cartridge Assignment Policy (CAP) range, upon being inserted, that volume can be inventoried according to its VOLSER definition. If the physical volume's VOLSER occurs outside a defined range, operator intervention will be required to dedicate the inserted physical volume to TS7740 Virtualization Engine resources.

In Figure 8-40, status information displayed in the VOLSER Ranges Table includes:

- From: The first VOLSER in a defined range.
- To: The last VOLSER in a defined range.
- Media Type: The media type for all volumes in a given VOLSER range. Possible values are:
 - JB-ETCL: Enterprise Extended-Length Tape Cartridge
 - JA-ETC: Enterprise Tape Cartridge
 - JJ-EETC: Enterprise Economy Tape Cartridge
- Home Pool: The home pool to which the VOLSER range is assigned.

Use the drop-down menu in the VOLSER Ranges Table to add a new VOLSER range or modify or delete a predefined range.

To add a new VOLSER range, select **Add** from the drop-down menu. Complete the fields for information that will be displayed in the VOLSER Ranges Table, as defined previously.

To modify a predefined VOLSER range, select the radio button from the Select column that appears in the same row as the name of the VOLSER range you want to modify. Select **Modify** from the drop-down menu and make your desired changes to the information that will be displayed in the VOLSER Ranges Table, as defined previously.

To delete a predefined VOLSER range, select the radio button from the Select column that appears in the same row as the name of the VOLSER range you want to delete. Select **Delete** from the drop-down menu. A confirmation window will appear; select **Yes** to delete the indicated VOLSER range, or select **No** to cancel the request.

Referring to Figure 8-40 on page 471, the Unassigned Volsers table displays the list of unassigned physical volumes for a given component that are pending ejection. A VOLSER is removed from this table when a new range that contains the VOLSER is added.

You can use the Unassigned Volsers table to eject one or more physical volumes from a library attached to a TS7740 Virtualization Engine.

Status information displayed in the Unassigned Volsers table includes:

- ▶ Volser: The VOLSER associated with a given physical volume.
- ▶ Media Type: The media type for all volumes in a given VOLSER range. Possible values are:
 - JB-ETCL: Enterprise Extended-Length Tape Cartridge
 - JA-ETC: Enterprise Tape Cartridge
 - JJ-EETC: Enterprise Economy Tape Cartridge
- ▶ Pending Eject: Whether the physical volume associated with the given VOLSER is awaiting ejection.

To eject a physical volume, check the check box adjacent to the VOLSER associated with the volume and then select **Eject** from the drop-down menu on the Unassigned Volsers table. To eject multiple volumes, check multiple check boxes before selecting **Eject** from the drop-down menu. Buttons on the tool bar of the Unassigned Volsers table will select all or de-select all VOLSERS in the table. A confirmation page will appear after **Eject** is selected; select **Yes** to eject the indicated volume (or volumes), or **No** to cancel the request.

Physical Volumes Search window

Use this window (Figure 8-41) to search for physical volumes in a given TS7740 Virtualization Engine cluster according to one or more identifying features.

IBM System Storage™ TS7700 Virtualization Engine

Welcome admin Help | About | Log out

Grid Summary

Health & Monitoring

Performance & Statistics

Logical Volumes

Physical Volumes

Physical Volume Pools

Physical Volume Details

Move Physical Volumes

Eject Physical Volumes

Physical Volume Ranges

Physical Volume Search

Constructs

Configuration

User Management

Service & Troubleshooting

"Pesto[0]" (#BA68A): Physical Volume Search

Grid Summary > "Pesto" Cluster Summary > Physical Volume Search

Make selections below to define the database search criteria. The more search criteria used, the more restrictive the search.

Current cluster selected:

Clusters

- "Pesto[0]" (#BA68A)
- "Squint[1]" (#BA68B)
- "Celeste[2]" (#BA68C)

Previous Searches

New Search Name:

Search Options

Volser: Media Type:

Home Pool: Current Pool:

Encryption key: ☐ Search for the default key

☐ Pending All Ejects

☐ Pending All Moves to All Pools

Volser Flags:

	Yes	No	Ignore
Misplaced	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Mounted	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Inaccessible	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Encrypted	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Export Hold	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Read Only Recovery	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Unavailable	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Pending Secure Data Erase	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Copy Exported	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>

External Links

Launch TS3500 Library Specialist

Search Results Options

Search

Figure 8-41 MI Physical Volume Search entry window

You can view the results of a previous query, or create a new query to search for physical volumes.

Note: Only one search can be executed at a time. If a search is in progress, an information message will occur at the top of the Physical Volumes Search window. You can cancel a search in progress by clicking the **Cancel Search** button within this message.

To create a new search query, enter a name for the new query. Enter a value for any of the fields defined below and select **Search** to initiate a new physical volume search. The query name, criteria, start time, and end time are saved along with the search results.

- ▶ **VOLSER:** The volume's serial number. This field can be left blank. You can also use the following wild card characters in this field:
 - % (percent): Represents zero or more characters.
 - * (asterisk): Translated to % (percent). Represents zero or more characters.
 - . (period): Represents one character.
 - _ (single underscore): Translated to . (period). Represents one character.
 - ? (question mark): Translated to . (period). Represents one character.
- ▶ **Media Type:** The type of media on which the volume resides. Use the drop-down menu to select from available media types. This field can be left blank.
- ▶ **Home Pool:** The pool number (0-32) the physical volume was assigned to when it was inserted into the library, or the pool it was moved to through the management interface Move/Eject Stacked Volumes function. This field may be left blank.
- ▶ **Current Pool:** The number of the storage pool (0-32) in which the physical volume currently resides. This field may be left blank.
- ▶ **Encryption Key:** The encryption key label designated when the volume was encrypted. This is a text field. Possible values include:
 - A name identical to the first or second key label on a physical volume. Any physical volume encrypted using the designated key label is included in the search.
 - Search for the default key. Select this check box to search for all physical volumes encrypted using the default key label.
- ▶ **Pending Eject:** This option is used to decide whether to include physical volumes pending an eject in the search query. Possible values include:
 - All Ejects: All physical volumes pending eject are included in the search.
 - Priority Ejects: Only physical volumes classified as priority eject are included in the search.
 - Deferred Ejects: Only physical volumes classified as deferred eject are included in the search.
- ▶ **Pending Move to Pool:** Whether to include physical volumes pending a move in the search query. Possible values are:
 - All Moves: All physical volumes pending a move are included in the search.
 - Priority Moves: Only physical volumes classified as priority move are included in the search.
 - Deferred Moves: Only physical volumes classified as deferred move are included in the search.

Any of these values can be modified using the adjacent To drop-down menu. Use this drop-down menu to narrow your search to a specific pool set to receive physical volumes. Possible values are:

- ▶ **All Pools:** All pools will be included in the search.
- ▶ **0–32:** The number of the pool to which the selected physical volumes will be moved.

- **VOLSER Flags:** This option is used to decide whether to include misplaced, mounted, or inaccessible physical volumes in the search query. Possible values include:
 - **Yes:** Selected physical volume types are included in the search.
 - **No:** Selected physical volume types are excluded from the search.
 - **Ignore:** The selected status is ignored during the search operation.

Figure 8-42 shows a sample Physical Volume Search query for a VOLSER named 320031.

IBM System Storage™ TS7700 Virtualization Engine

Welcome admin

Grid Summary > "Juno" Cluster Summary > Physical Volume Search

Make selections below to define the database search criteria. The more search criteria used, the more restrictive the search.

Current cluster selected:

- Clusters
 - "Zoro[0]" (#BA39A)
 - "Juno[1]" (#BA39B)**
 - "Dane[2]" (#BA39C)
 - "Herc[3]" (#BA39D)

Previous Searches

New Search Name: Luciano

Search Options

Volser: 320031 Media Type: [dropdown]

Home Pool: [dropdown] Current Pool: [dropdown]

Encryption key: [text] ☐ Search for the default key

☐ Pending All Ejects [dropdown]

☐ Pending All Moves [dropdown] to All Pools [dropdown]

Volser Flags:

	Yes	No	Ignore
Misplaced	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Mounted	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Inaccessible	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Encrypted	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Export Hold	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Read Only Recovery	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Unavailable	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Pending Secure Data Erase	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Copy Exported	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>

Search Results Options

Search

Figure 8-42 MI Physical Volume Search selection

Clicking the **Search** button opens the window shown in Figure 8-43.

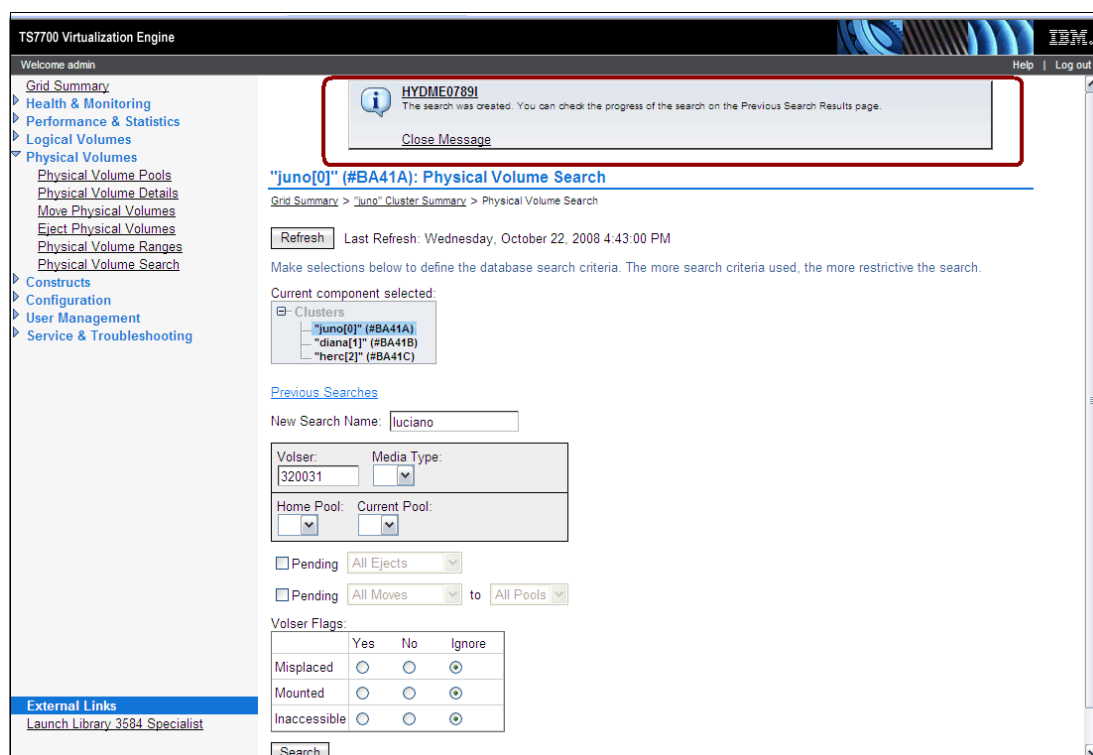


Figure 8-43 MI Physical Volume Search query message

Note: A message identifier, HYDME0789I, is produced to indicate that a search was created; the progress can be viewed in the Previous Search Results window.

To view the results of your query or one of the last ten previous search queries, select **Previous Searches** in Figure 8-43 and you will get a list of the last ten search queries, as shown in Figure 8-44 on page 477.

Up to 10 previously named search queries can be saved. To clear the list of saved queries, check the check box adjacent to one or more queries to be removed, select **Clear** from the drop-down menu, and click **Go**. This operation will not clear a search query already in progress. You will be asked to confirm your decision to clear the query list. Select **Yes** to clear the list of saved queries, or **No** to retain the list of queries.

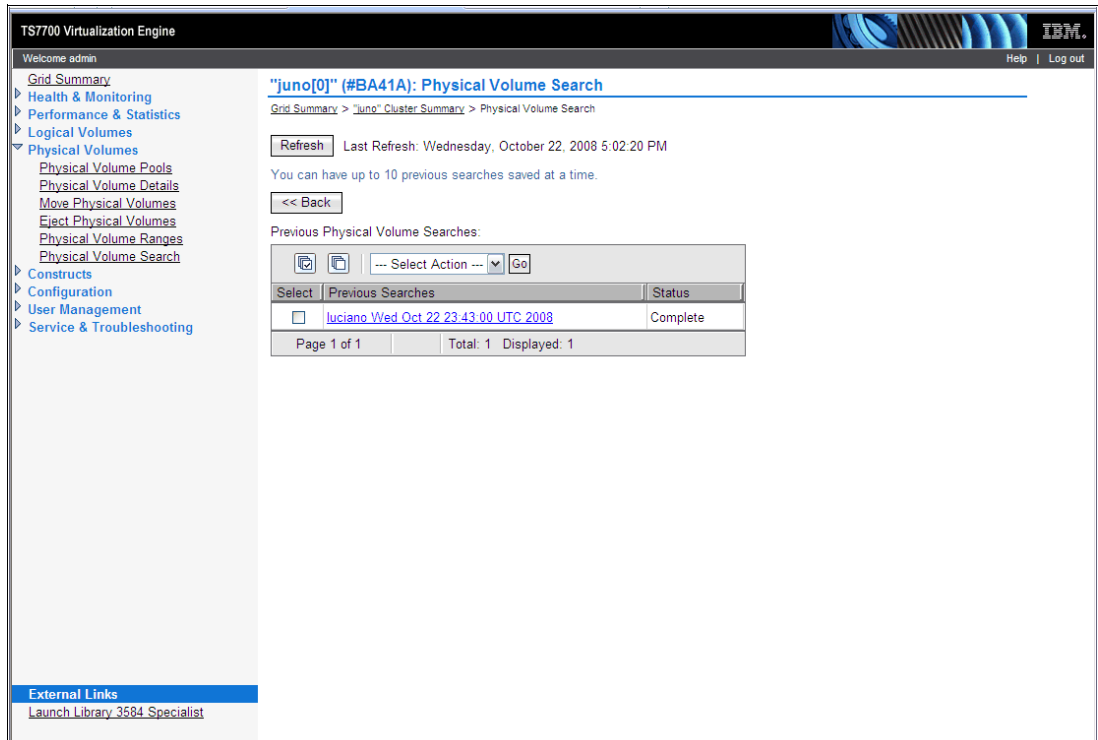


Figure 8-44 MI Physical Volume Search up to previous ten searches

In Figure 8-44, click a query name to display a list of physical volumes that match the search criteria. See Figure 8-45 for a description of the previous Physical Volume Search results.

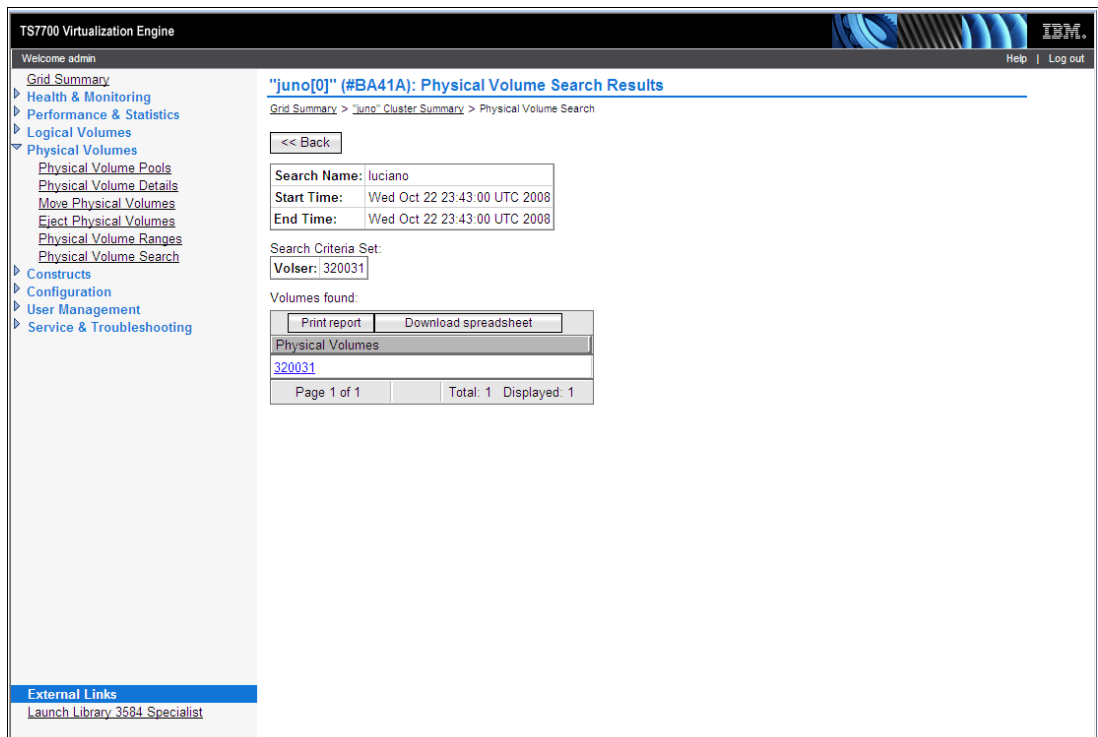


Figure 8-45 MI results of a previously submitted Physical Volume Search

The Physical Volume Search Results window displays a list of physical volumes that meet the criteria of an executed search query.

The table at the top of the Physical Volume Search Results window displays the name of the completed search query and its start and end times. The table following the heading Search Criteria Set displays the criteria used to generate the search query.

The Search Results table displays a list of VOLSERS belonging to the physical volumes that meet the search criteria. Each VOLSER listed under the Physical Volumes heading is linked to its physical volume details window. Click the VOLSER to open the details page and view additional information about a specific physical volume, as shown in Figure 8-46 on page 479.

The Search Results table displays up to 10 VOLSERS per page. To view the results contained on another window, select the right- or left-pointing arrow at the bottom of the table, or enter a page number in the adjacent field and click **Go**. You can print or download the results of a search query using the **Print Report** or **Download Spreadsheet** buttons at the top of the Volumes found table.

Figure 8-46 shows the Physical Volume Details for VOLSER 320031.

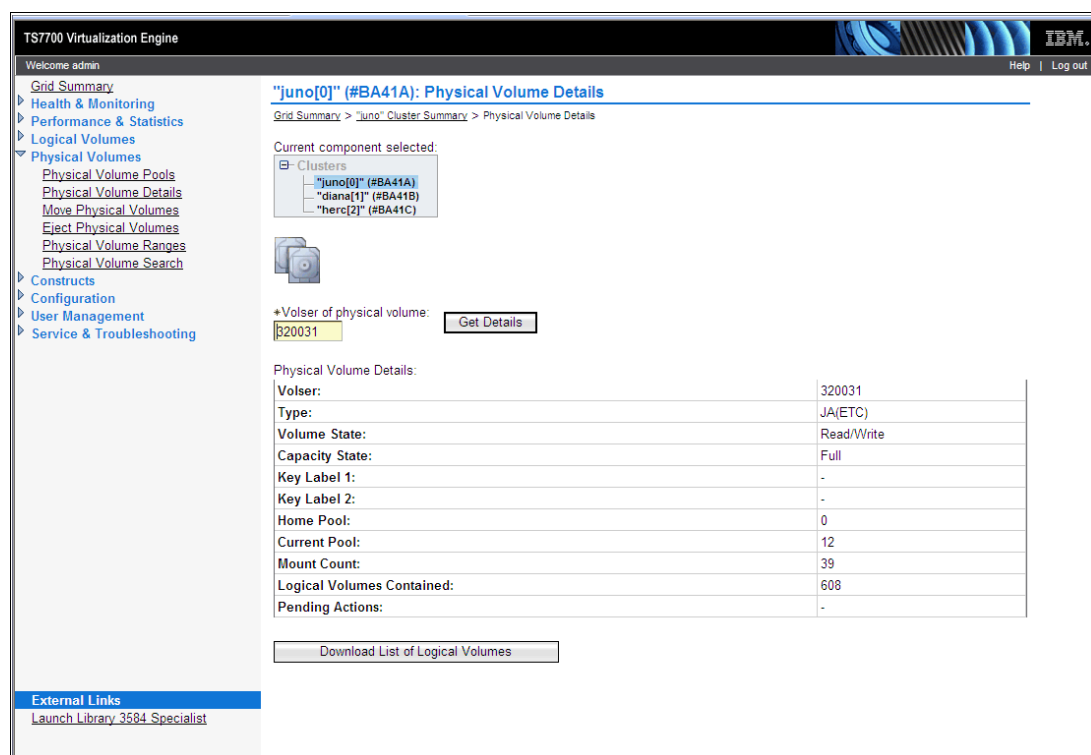


Figure 8-46 MI Physical Volume Details window

8.2.7 Constructs

The topics in this section present information that is related to TS7700 Virtualization Engine storage constructs.

Storage Groups window

Use this window (Figure 8-47 on page 480) to add, modify, or delete a storage group used to define a primary pool for logical volume premigration.

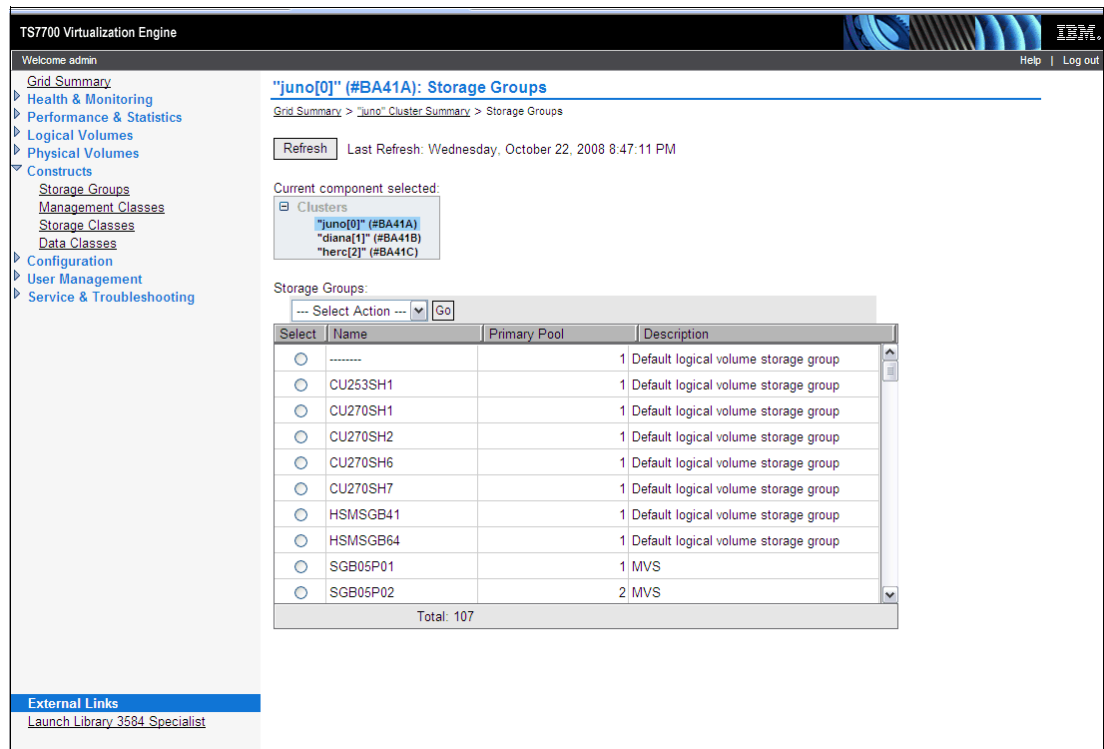


Figure 8-47 MI Storage Groups window

The Storage Groups table displays all existing storage groups available for a given cluster.

You can use the Storage Groups table to create a new storage group, modify an existing storage group, or delete a storage group.

The status information displayed in the Storage Groups table includes:

- ▶ **Name:** The name of the storage group. Each storage group within a cluster must have a unique name. Valid characters for this field are A-Z, 0-9, \$, @, *, #, and %.
- ▶ **Primary Pool:** The primary pool for premigration. Only validated physical primary pools may be selected. If the cluster does not possess a physical library, this column will not be visible and the management interface categorizes newly created storage groups using pool 1.
- ▶ **Description:** A description of the storage group.

Use the drop-down menu in the Storage Groups table to add a new storage group or modify or delete an existing storage group.

To add a new storage group, select **Add** from the drop-down menu. Complete the fields for information that will be displayed in the Storage Groups table, as defined previously.

Note: If the cluster does not possess a physical library, the Primary Pool field will not be available in the Add or Modify options.

To modify an existing storage group, select the radio button from the Select column that appears adjacent to the name of the storage group you want to modify. Select **Modify** from the drop-down menu. Complete the fields for information that will be displayed in the Storage Groups table, as defined previously.

To delete an existing storage group, select the radio button from the Select column that appears adjacent to the name of the storage group you want to delete. Select **Delete** from the drop-down menu. You will be prompted to confirm your decision to delete a storage group. If you select **Yes**, the storage group will be deleted. If you select **No**, your request to delete will be cancelled.

Management Classes window

Use this window (Figure 8-48) to define, modify, copy, or delete the Management Class that defines the TS7700 Virtualization Engine copy policy for volume redundancy.

The Current Copy Policy Table displays the copy policy in force for each component of the Grid. If no Management Class is selected, this table will not be visible. You must select a Management Class from the Management Classes Table to view copy policy details.

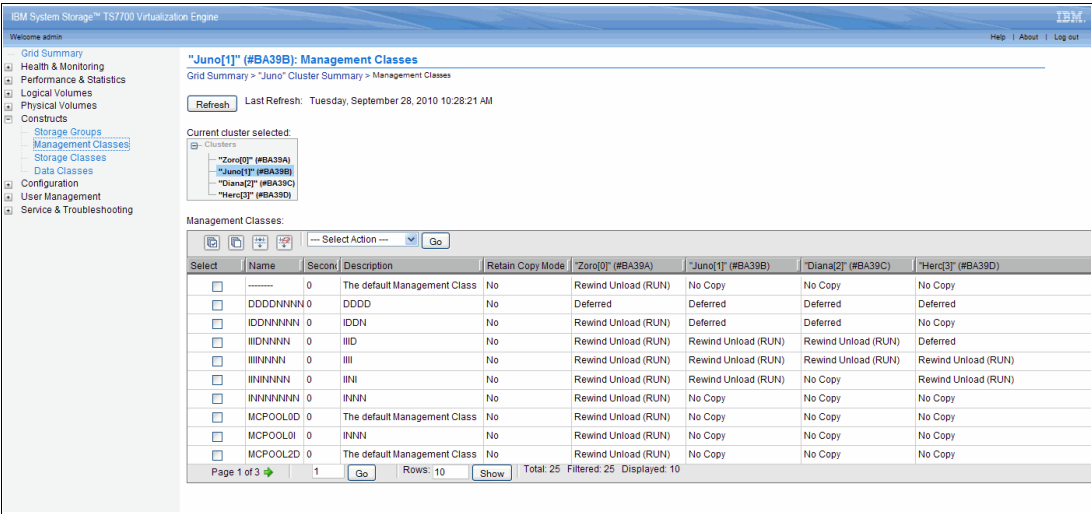


Figure 8-48 MI Management Classes window

The Management Classes Table (Figure 8-48) displays defined Management Class copy policies that may be applied to a cluster.

You can use the Management Classes Table to create a new Management Class, modify an existing Management Class, or delete one or more existing Management Classes. The default Management Class may be modified, but cannot be deleted. The default Management Class uses dashes (-----) for the symbolic name.

Status information displayed in the Management Classes Table includes:

- ▶ **Name:** The name of the Management Class. Valid characters for this field are A-Z, 0-9, \$, @, *, #, and %. The first character of this field may not be a number. This is the only field that may not be modified after it is added.
- ▶ **Secondary Pool:** The target pool in the volume duplication. If the cluster does not possess a physical library, this column will not be visible and the management interface categorizes newly created storage groups using pool 0.
- ▶ **Description:** A description of the Management Class definition. The value in this field must be from 1 to 70 characters in length.
- ▶ **Retain Copy Mode:** Whether previous copy policy settings on non-fast-ready logical volume mounts are retained.

Retain Copy Mode prevents the copy modes of a logical volume from being refreshed by an accessing host device if the accessing cluster is not the same cluster that created the volume. When Retain Copy Mode is enabled through the management interface, previously assigned copy modes are retained and subsequent read or modify access does not refresh, update, or merge copy modes. This allows the original number of copies to be maintained.

Use the drop-down menu on the Management Classes table to add a new Management Class, modify an existing Management Class, or delete one or more existing Management Classes.

To add a new Management Class, select **Add** from the drop-down menu and click **Go**. Complete the fields for information that will be displayed in the Management Classes Table, as defined previously. You can create up to 256 Management Classes per TS7700 Virtualization Engine Grid.

Note: If the cluster does not possess a physical library, the Secondary Pool field will not be available in the Add option.

The Copy Action drop-down menu appears adjacent to each cluster in the TS7700 Virtualization Engine Grid. The Copy Action menu allows you to select, for each component, the copy mode to be used in volume duplication.

The following actions are available in this menu:

- ▶ **No Copy:** No volume duplication will occur if this action is selected.
- ▶ **Rewind Unload (RUN):** Volume duplication occurs when the Rewind Unload command is received. The command will only return after the volume duplication completes successfully.
- ▶ **Deferred:** Volume duplication occurs at a later time based on the internal schedule of the copy engine.

To modify an existing Management Class, check the check box from the Select column that appears in the same row as the name of the Management Class you want to modify. You can modify only one Management Class at a time. Select **Modify** from the drop-down menu and click **Go**. Of the fields listed previously in the Management Classes Table, or available from

the Copy Action drop-down menu, you will be able to change all of them except the Management Class name.

Note: If the cluster does not possess a physical library, the Secondary Pool field will not be available in the Modify option.

To delete one or more existing Management Classes, check the check box from the **Select** column that appears in the same row as the name of the Management Class you want to delete. Check multiple check boxes to delete multiple Management Classes. Select **Delete** from the drop-down menu and click **Go**.

Note: You will not be permitted to delete the default Management Class.

Storage Classes window

Use this window (Figure 8-49) to define, modify, or delete a storage class used by the TS7700 Virtualization Engine to automate storage management through classification of data sets and objects.

The Storage Classes table displays the list of storage classes defined for each component of the Grid.

Note: Storage Class constructs do not apply to TS7700 Virtualization Engine Clusters that do not possess a physical library.

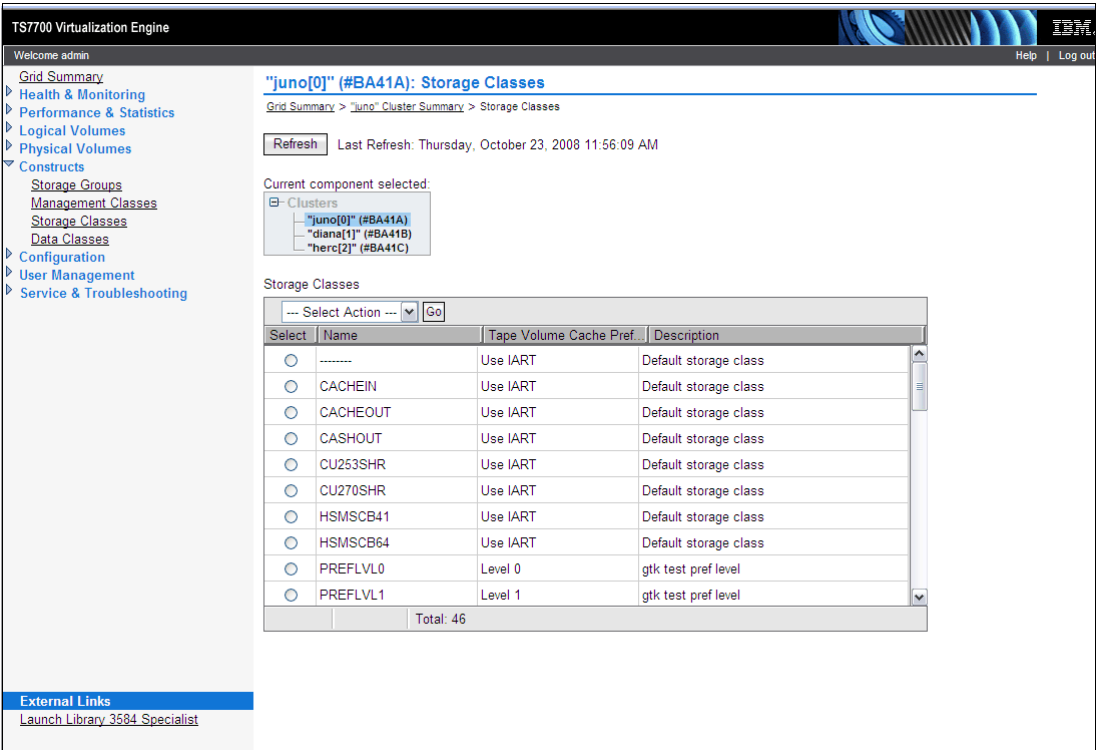


Figure 8-49 MI Storage Classes window

The Storage Classes table (Figure 8-49) lists defined storage classes that are available to control data sets and objects within a cluster.

You can use the Storage Classes table to create a new storage class or modify or delete an existing storage class. The default storage class may be modified, but cannot be deleted. The default Storage Class has dashes (-----) as the symbolic name.

The Storage Classes table displays the following status information:

- ▶ **Name:** The name of the storage class. Valid characters for this field are A-Z, 0-9, \$, @, *, #, and %. The first character of this field may not be a number. The value in this field must be between 1 and 8 characters in length.
- ▶ **Tape Volume Cache Preference:** The preference level for the storage class; it determines how soon volumes are removed from cache following their copy to tape. Possible values are:
 - **Use IART:** Volumes are removed according to the TS7700 Virtualization Engine's Initial Access Response Time (IART).
 - **Level 0:** Volumes are removed from the tape volume cache as soon as they are copied to tape
 - **Level 1:** Copied volumes remain in the tape volume cache until additional space is required, then the first volumes are removed to free up space in the cache.
- ▶ **Description:** A description of the storage class definition. The value in this field must be between 1 and 70 characters in length.

Use the drop-down menu in the Storage Classes Table to add a new storage class or modify or delete an existing storage class.

To add a new storage class, select **Add** from the drop-down menu. Complete the fields for the information that will be displayed in the Storage Classes Table, as defined previously.

Note: You can create up to 256 storage classes per TS7700 Virtualization Engine Grid.

To modify an existing storage class, select the radio button from the Select column that appears in the same row as the name of the storage class you want to modify. Select **Modify** from the drop-down menu. Of the fields listed in the Storage Classes Table, you will be able to change all of them except the storage class name.

To delete an existing storage class, select the radio button from the Select column that appears in the same row as the name of the storage class you want to delete. Select **Delete** from the drop-down menu. A window will appear to confirm the storage class deletion. Select **Yes** to delete the storage class or **No** to cancel the delete request.

Data Classes window

Use this window (Figure 8-50) to define, modify, or delete a TS7700 Virtualization Engine Data Class used to automate storage management through classification of data sets.

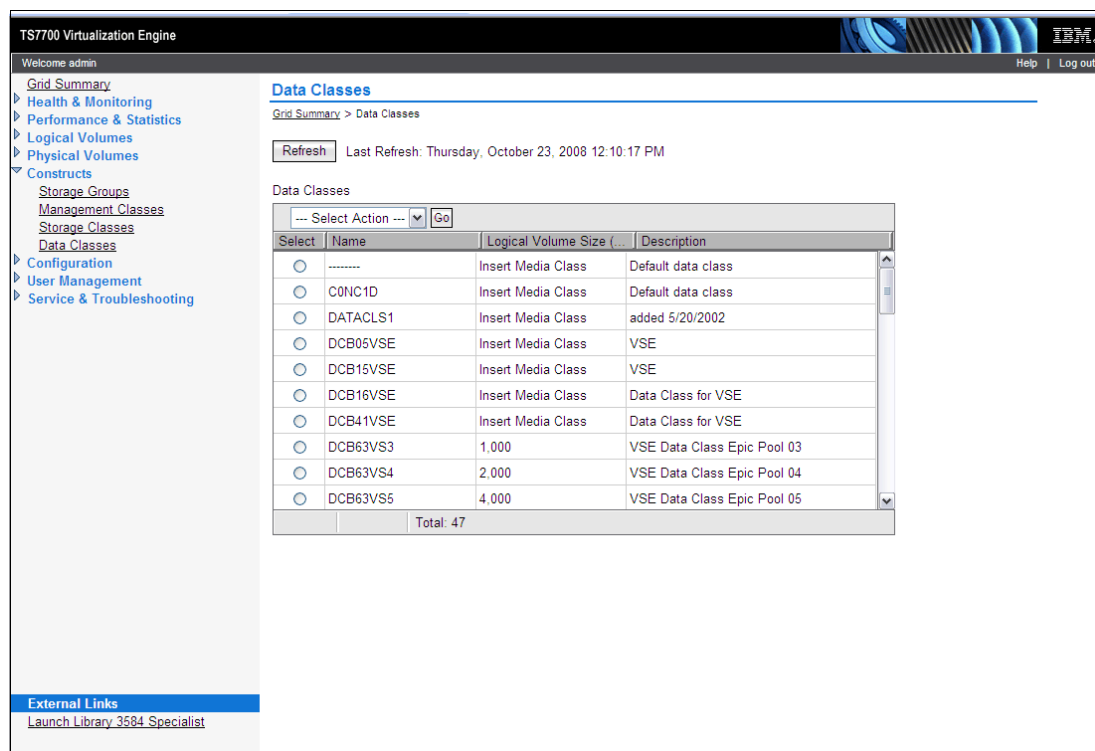


Figure 8-50 MI Data Classes window

Note: Fast Ready Categories and Data Classes work at the system level and are unique for all clusters in a Grid. This means that if you modify them on one cluster that they will be applied to all clusters in the Grid.

The Data Classes Table (Figure 8-50) displays the list of Data Classes defined for each cluster of the Grid.

You can use the Data Classes Table to create a new data class or modify or delete an existing Data Class. The default Data Class may be modified, but cannot be deleted. The default Data Class has dashes (-----) as the symbolic name.

The Data Classes table lists the following status information:

- ▶ **Name:** The name of the Data Class. Valid characters for this field are A-Z, 0-9, \$, @, *, #, and %. The first character of this field may not be a number. The value in this field must be between 1 and 8 characters in length.
- ▶ **Logical Volume Size (MiB):** The logical volume size of the Data Class, which determines the maximum number of MiB for each logical volume in a defined class. Possible values are Insert Media Class, where the logical volume size is not defined (the Data Class will not be defined by a maximum logical volume size). Possible values are 1,000 MiB, 2,000 MiB, or 4,000 MiB.

- **Description:** A description of the Data Class definition. The value in this field must be at least 0 and no greater than 70 characters in length.
- **Logical WORM:** Whether logical WORM (write-once, read-many) is set for the data class. Logical WORM is the virtual equivalent of WORM tape media, achieved through software emulation. This setting is available only when all clusters in a Grid operate at R1.6 and later.

Possible values are as follows:

- **Yes:** Logical WORM is set for the data class. Volumes belonging to the data class are defined as logical WORM.
- **No:** Logical WORM is not set. Volumes belonging to the data class are not defined as logical WORM. This is the default value for a new data class.

Use the drop-down menu in the Data Classes Table to add a new Data Class or modify or delete an existing Data Class.

To add a new Data Class, select **Add** from the drop-down menu and click **Go**. Complete the fields for information that will be displayed in the Data Classes Table.

Note: You can create up to 256 Data Classes per TS7700 Virtualization Engine Grid.

To modify an existing Data Class, select the radio button from the Select column that appears in the same row as the name of the Data Class you want to modify. Select **Modify** from the drop-down menu and click **Go**. Of the fields listed in the Data Classes Table, you will be able to change all of them except the default Data Class name.

To delete an existing Data Class, select the radio button from the Select column that appears in the same row as the name of the Data Class you want to delete. Select **Delete** from the drop-down menu and click **Go**. A window will appear to confirm the Data Class deletion. Select **Yes** to delete the Data Class or **No** to cancel the delete request.

8.2.8 Configuration

In this section, we describe the functions of the TS7700 Virtualization Engine MI that are related to configuring the TS7700 Virtualization Engine.

Grid Identification Properties window

Use this window (Figure 8-51 on page 487) to view and alter identification properties for the TS7700 Virtualization Engine Grid. It can be used to distinguish the composite library.

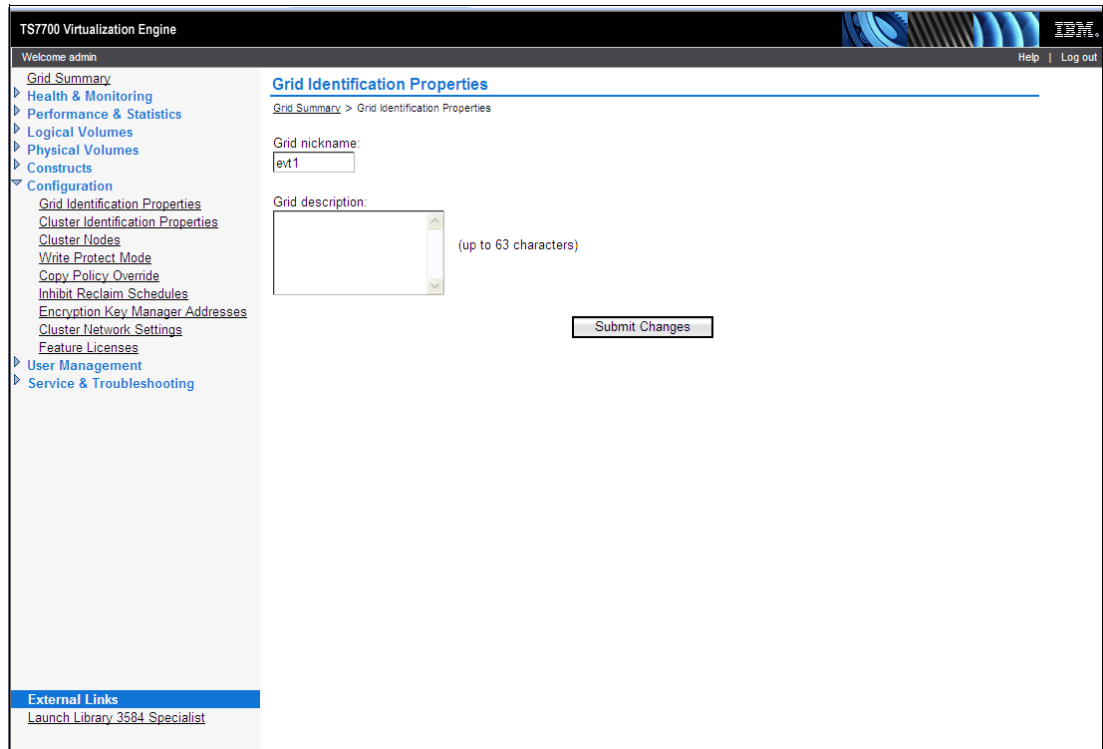


Figure 8-51 MI Grid Identification Properties window

The following information, related to Grid identification, is displayed. To change the Grid identification properties, edit the available fields and click **Submit Changes**. The available fields are as follows:

- ▶ Grid nickname: The Grid nickname must be one to eight characters in length and composed of alphanumeric characters with no spaces. The characters @, ., -, and + are also allowed.
- ▶ Grid description: A short description of the Grid. You can use up to 63 characters.

Cluster Identification Properties window

Use this window (Figure 8-52) to view and alter cluster identification properties for the TS7700 Virtualization Engine. This can be used to distinguish this distributed library.

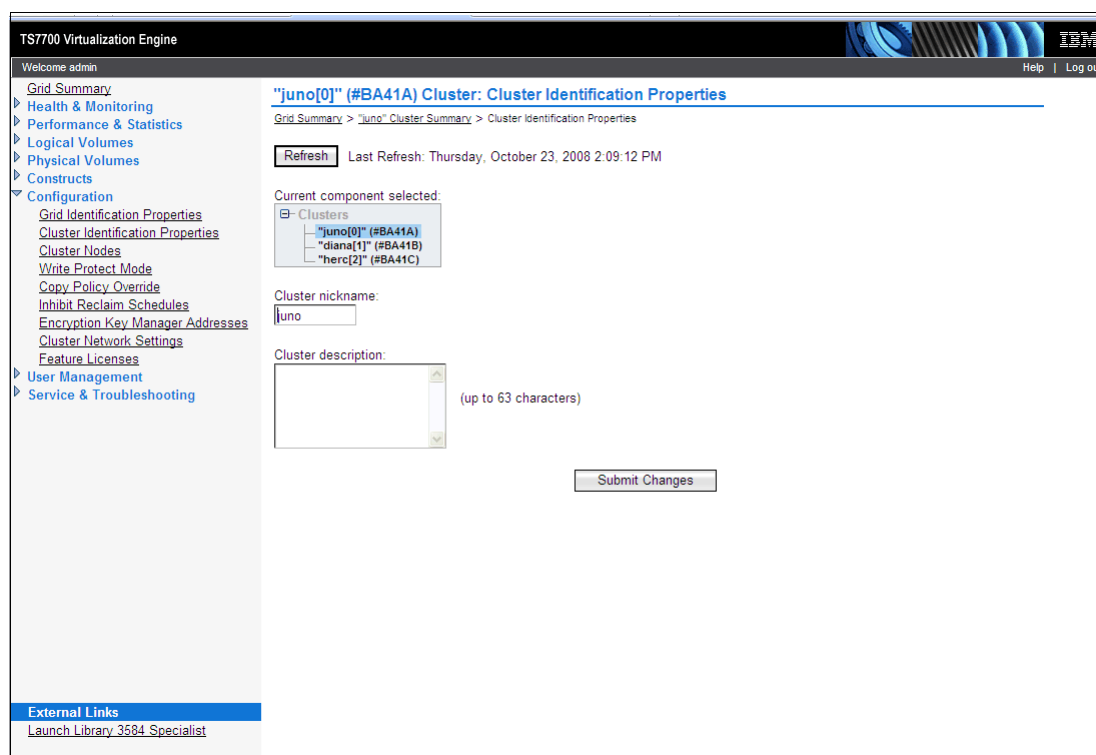


Figure 8-52 MI Cluster Identification Properties window

The following information related to cluster identification is displayed. To change the cluster identification properties, edit the available fields and click **Submit Changes**. The available fields are:

- ▶ **Cluster nickname:** The cluster nickname must be one to eight characters in length and composed of alphanumeric characters. Blank spaces and the characters @, ., -, and + are also allowed. Blank spaces cannot be used in the first or last character positions.
- ▶ **Cluster description:** A short description of the cluster. You can use up to 63 characters.

Cluster Families window

Use the Cluster Families window (Figure 8-53) to view information and perform actions related to TS7700 Virtualization Engine Cluster families, a feature only available on clusters operating at Release 1.6 and later.

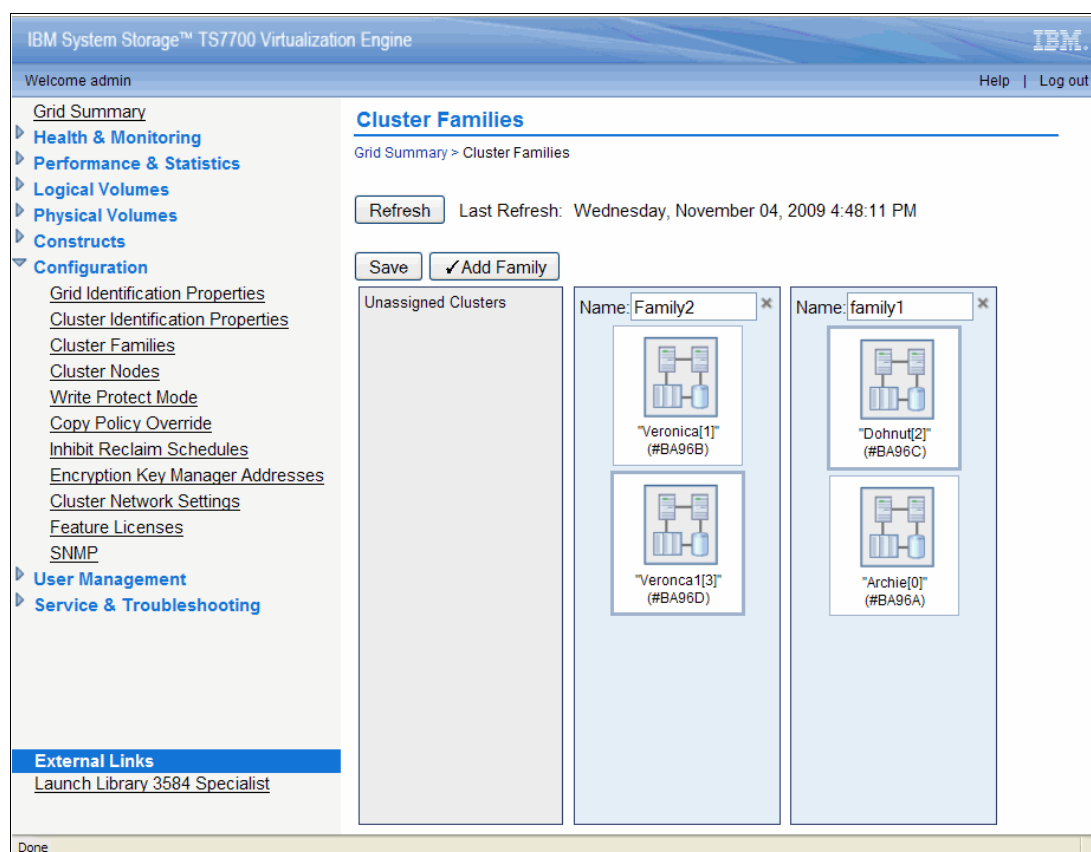


Figure 8-53 MI Cluster Families window

Data transfer speeds between TS7700 Virtualization Engine Clusters sometimes vary. The cluster family configuration groups clusters so that microcode can optimize Grid connection performance between the grouped clusters.

To view or modify cluster family settings, first verify that these permissions are granted to your assigned user role. If your user role includes cluster family permissions, select the **Modify** button to perform the following actions:

- Add a family: Click the **Add** button to create a new cluster family. A new cluster family placeholder is created to the right of any existing cluster families. Enter the name of the new cluster family in the active Name text box. Cluster family names must be one to eight characters in length and composed of Unicode characters. Each family name must be unique. Clusters are added to the new cluster family by relocating a cluster from the Unassigned Clusters area using the method described in the *Move a cluster* function, described next.

- ▶ Move a cluster: You can move one or more clusters, by drag and drop, between existing cluster families, to a new cluster family from the Unassigned Clusters area, or to the Unassigned Clusters area from an existing cluster family.
 - Select a cluster: A selected cluster is identified by its highlighted border. Select a cluster from its resident cluster family or the Unassigned Clusters area by:
 - Clicking the cluster with your mouse.
 - Using the Spacebar on your keyboard.
 - Press and hold the Shift key while selecting clusters to select multiple clusters at one time.
 - Press the Tab key on your keyboard to switch between clusters before selecting one.
 - Move the selected cluster or clusters by:
 - Click and hold the mouse button on the cluster and drag the selected cluster to the destination cluster family or the Unassigned Clusters area.
 - Using the arrow keys on your keyboard to move the selected cluster or clusters right or left.

Note: An existing cluster family cannot be moved within the Cluster Families window.

- ▶ Delete a family: You can delete an existing cluster family. Click the X in the top right corner of the cluster family you want to delete. If the cluster family you attempt to delete contains any clusters, a warning message is displayed. Click **OK** to delete the cluster family and return its clusters to the Unassigned Clusters area. Click **Cancel** to abandon the delete action and retain the selected cluster family.
- ▶ Save changes: Click the **Save** button to save any changes made to the Cluster Families window and return it to read-only mode.

Note: Each cluster family should contain at least one cluster. If you attempt to save changes and a cluster family does not contain any clusters, an error message displays and the Cluster Families window remains in edit mode.

Cluster Nodes window

Use this window (Figure 8-54) to view information and perform actions related to cluster nodes on the TS7700 Virtualization Engine.

To watch a tutorial that shows how to work with clusters, click the **View tutorial** link.

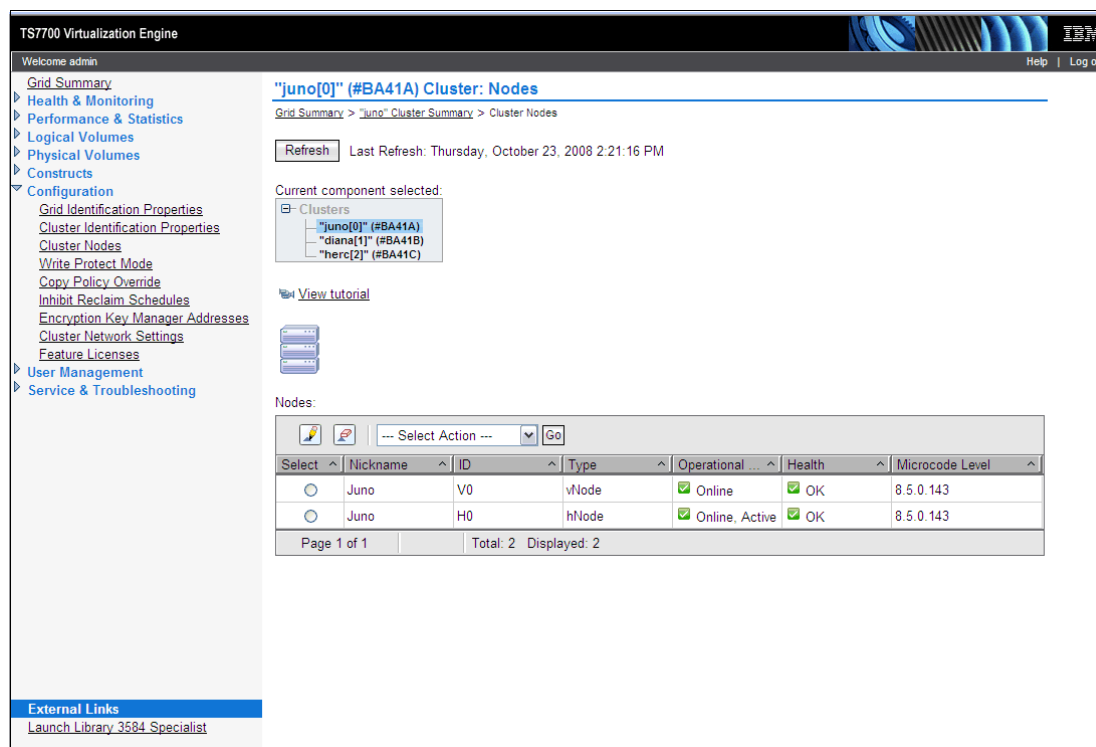


Figure 8-54 MI Cluster Nodes window

Table fields

The following fields are available:

- ▶ Nickname: Nickname of the node followed by the library sequence number.
- ▶ ID: Node ID.
- ▶ Type: Type of the node. Possible values are vNode or hNode.
- ▶ Operational State: Possible values are Online, Going Online, Going Offline, and Offline.

Note: If the node type is hNode, the Operational State will also display as Active.

- ▶ Microcode Level: Current Licensed Internal Code level of the node.

Table actions

There are several actions available for cluster nodes accessed through the Select Action drop-down menu:

- ▶ Details & Health: View details about the selected node. Details shown include general node information, node state and health information, and networking settings for the selected node.
- ▶ Modify Node Nickname: Modify the nickname of the selected node.

Write Protect Mode window

Use the page shown in Figure 8-55 to view Write Protect Mode settings for a TS7700 Virtualization Engine Cluster

When Write Protect Mode is enabled on a cluster, host commands fail if they are issued to logical devices in that cluster and attempt to modify a volume's data or attributes. Meanwhile, host commands that are issued to logical devices in peer clusters are allowed to continue with full read and write access to all volumes in the library. Write Protect Mode is used primarily for customer-initiated disaster recovery testing. In this scenario, a recovery host that is connected to a non-production cluster must access and validate production data without any risk of modifying it.

A cluster can be placed into Write Protect Mode only if the cluster is online. After it is set, the mode is retained through intentional and unintentional outages and can only be disabled through the same management interface panel used to enable the function. When a cluster within a grid configuration has Write Protect Mode enabled, standard grid functions such as logical volume replication and logical volume ownership transfer are unaffected.

Logical volume categories can be excluded from Write Protect Mode. Up to 16 categories can be identified and set to include or exclude from Write Protect Mode using the Category Write Protect Properties table. Additionally, volumes in any Fast Ready category can be written to if the **Ignore fast ready characteristics of write protected categories** check box is selected.

The screenshot shows the IBM System Storage™ TS7700 Virtualization Engine management interface. The left sidebar contains a navigation menu with categories like Grid Summary, Health & Monitoring, Performance & Statistics, Logical Volumes, Physical Volumes, Constructs, Configuration, User Management, and Service & Troubleshooting. The 'Configuration' section is expanded, showing sub-items like Grid Identification Properties, Cluster Identification Properties, Cluster Families, Cluster Nodes, Write Protect Mode (selected), Copy Policy Override, Inhibit Reclaim Schedules, Encryption Key Manager Addresses, Cluster Network Settings, Feature Licenses, SNMP, and Copy Export Settings.

The main content area is titled '"veronica[1]" (#BA96B): Write Protect Mode'. Below the title is a breadcrumb trail: Grid Summary > "veronica" Cluster Summary > Write Protect Mode. A 'Refresh' button and 'Last Refresh: Monday, September 27, 2010 9:10:02 PM' are displayed. A tree view shows the current cluster selected: Clusters, with sub-items "archie[0]" (#BA96A), "veronica[1]" (#BA96B) (selected), "dohnut[2]" (#BA96C), and "veronca1[3]" (#BA96D).

The 'Current State' is 'Disabled'. There are two radio buttons: 'Disable Write Protect Mode' (selected) and 'Enable Write Protect Mode'. A checked checkbox labeled 'Ignore fast ready characteristics of write protected categories' is also present. A 'Submit Changes' button is at the bottom of this section.

Below this is the 'Category Write Protect Properties' section. It includes a dropdown menu '--- Select Action ---' and a 'Go' button. A table with columns 'Select', 'Category', 'Excluded from Write Protect', and 'Description' is shown. The table is currently empty, with a 'Total: 0' at the bottom right.

At the bottom left, there is an 'External Links' section with a link to 'Launch TS3500 Library Specialist'.

Figure 8-55 MI Write Protect Mode window

The current Write Protect mode setting for the active cluster is displayed; buttons for enabling and disabling Write Protect mode are available, as shown in Figure 8-55 on page 492. The following fields are available:

- ▶ **Current State:** Either Enabled or Disabled.
- ▶ **Enable Write Protect Mode:** Only available if Write Protect Mode is currently disabled. Clicking this button will enable Write Protect Mode.
- ▶ **Disable Write Protect Mode:** Only available if Write Protect Mode is currently enabled. Clicking this button will disable Write Protect Mode.
- ▶ **Ignore fast ready characteristics of write protected categories:** If this check box is selected, any volume that is defined by a Fast Ready category becomes writable.

Use the Category Write Protect Properties table to add, modify, or delete categories to be selectively excluded from Write Protect Mode. When Write Protect Mode is enabled, any categories added to this table must display a value of Yes in the Excluded from Write Protect field before the volumes in that category can be modified by an accessing host.

The following Category fields are listed in the Category Write Protect Properties table:

- ▶ **Category:** The identifier for a defined category. This is an alphanumeric hexadecimal value between 0x0001 and 0xFEFF (0x0000 and 0xFFxx cannot be used). Letters used in the category value must be capitalized.
- ▶ **Excluded from Write Protect:** Whether the category is excluded from Write Protect Mode. Possible values are:
 - **Yes:** The category is excluded from Write Protect Mode. When Write Protect is enabled volumes in this category can be modified when accessed by a host.
 - **No:** The category is not excluded from Write Protect Mode. When Write Protect is enabled volumes in this category cannot be modified when accessed by a host.
- ▶ **Description:** A descriptive definition of the category and its purpose.
- ▶ Use the drop-down menu on the Category Write Protect Properties table to add a new category or modify or delete an existing category. You must click the Submit Changes button to save any changes made to Write Protect Mode settings.

Note: You can add up to 16 categories per cluster. However, you can add, modify, or delete categories only if all cluster in the Grid are at R1.7 or later.

Copy Policy Override window

Use this window (Figure 8-56) to set local copy policy and I/O overrides for a TS7700 Virtualization Engine Cluster.

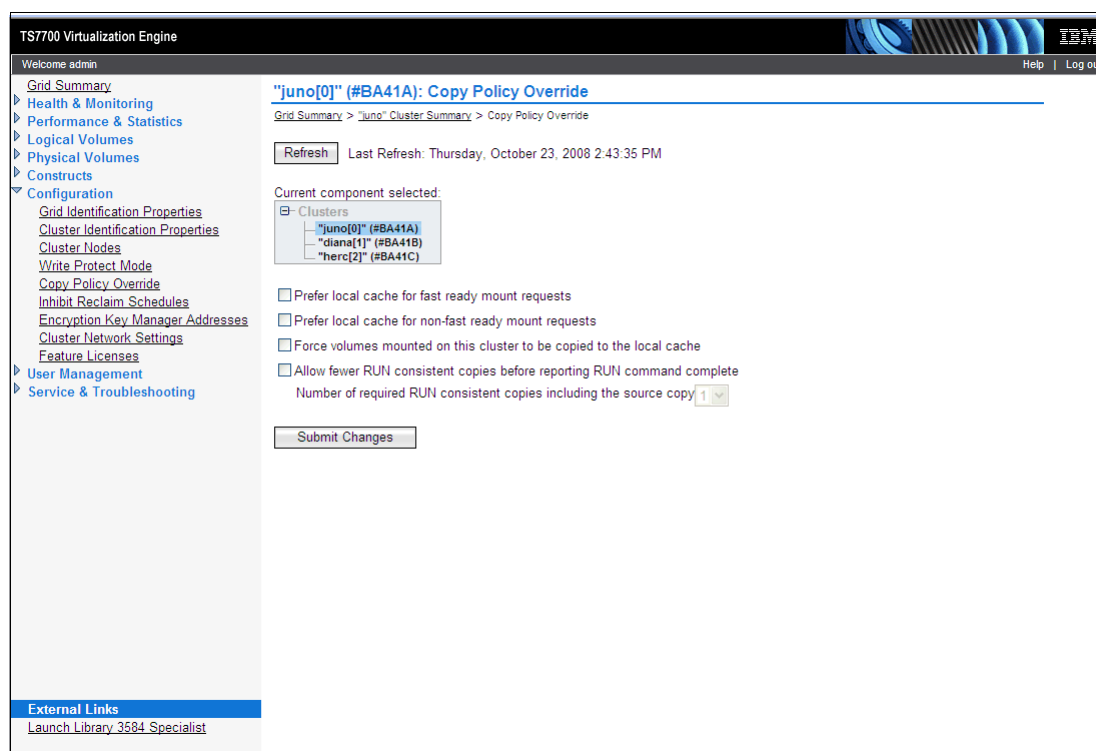


Figure 8-56 MI Copy Policy Override window

The following settings allow a user to specify cluster overrides for certain I/O and copy operations. These settings override the default TS7700 Virtualization Engine behavior and can be different for every cluster in a Grid.

- ▶ **Prefer local cache for fast ready mount requests:** A fast ready mount will select a local copy as long as a copy is available and a cluster copy consistency point is not specified as No Copy in the Management Class for the mount. The cluster is not required to have a valid copy of the data.
- ▶ **Prefer local cache for non-fast ready mount requests:** This override will cause the local cluster to satisfy the mount request as long as the cluster is available and the cluster has a valid copy of the data, even if that data is only resident on physical tape. If the local cluster does not have a valid copy of the data, then default cluster selection criteria applies.
- ▶ **Force volumes mounted on this cluster to be copied to the local cache:** For a non-fast ready mount, this override causes a copy to be performed on to the local cluster as part of the mount processing. For a fast ready mount, this setting has the effect of overriding the specified Management Class with a copy consistency point of Rewind/Unload for the cluster. This does not change the definition of the Management Class, but serves to influence the replication policy.
- ▶ **Allow fewer RUN consistent copies before reporting RUN command complete:** If selected, the value entered at Number of required RUN consistent copies, including the source copy, will be used to determine the maximum number of RUN copies, including the source, which must be consistent before the Rewind/Unload operation completes. If this option is not selected, the Management Class definitions are to be used explicitly. Thus, the number

of RUN copies can be from one to the number of clusters in the Grid configuration or the total number of clusters configured with a RUN copy consistency point.

Inhibit Reclaim Schedules window

Use this window (Figure 8-57) to add, modify, or delete Inhibit Reclaim Schedules used to postpone tape reclamation in a TS7740 Virtualization Engine cluster.

If the Grid possesses a physical library, but the selected cluster does not, this page is visible but disabled on the TS7700 Virtualization Engine management interface. If the Grid does not possess a physical library, this page is not visible on the TS7700 Virtualization Engine management interface.

Reclamation can improve tape utilization by consolidating data on some physical volumes, but it consumes system resources and can affect host access performance. The Inhibit Reclaim Schedules function can be used to disable reclamation in anticipation of increased host access to physical volumes.

The Schedules table displays the list of Inhibit Reclaim schedules defined for each partition of the grid. It displays the day, time, and duration of any scheduled reclamation interruption. All inhibit reclaim dates and times are displayed first in Coordinated Universal Time (UTC) and then in local time.

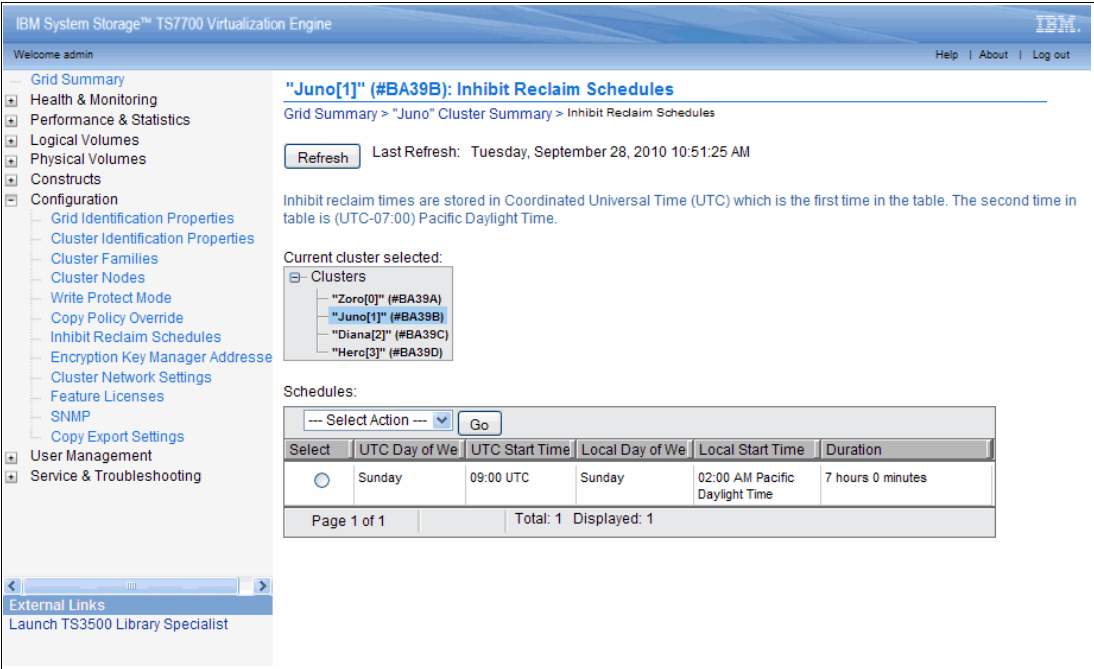


Figure 8-57 MI Inhibit Reclaim Schedules window

In Figure 8-57, the Schedules table displays the list of Inhibit Reclaim schedules defined for each partition of the grid. It displays the day, time, and duration of any scheduled reclamation interruption. All inhibit reclaim dates and times are displayed first in Coordinated Universal Time (UTC) and then in the local time of the browser. The conversion is automatic.

The following status information is listed in the Schedules table:

- ▶ **UTC Day of Week:** The UTC day of the week on which the reclamation will be inhibited. Possible values for this field include Every Day, Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.
- ▶ **UTC Start Time:** The UTC time at which reclamation will begin to be inhibited. The values in this field must take the form HH:MM. Possible values for this field include 00:00 through 23:59. The Start Time field is accompanied by a time chooser clock icon. You can enter hours and minutes manually using 24-hour time designations, or you can use the time chooser to select a start time based on a 12-hour (AM/PM) clock.
- ▶ **Local Start Time:** The local time at which reclamation will begin to be inhibited. The values in this field must take the form HH:MM. The time recorded reflects the local time of the browser.
- ▶ **Duration:** The number of days (D) hours (H) and minutes (M) that the reclamation will be inhibited. The values in this field must take the form DD days HH hours MM minutes. Possible values for this field include 0 day 0 hour 1 minute through 1 day 0 hour 0 minute if the day of the week is Every Day. Otherwise, possible values for this field are 0 day 0 hour 1 minute through 7 days 0 hour 0 minute.

Use the drop-down menu on the Schedules table to add a new Reclaim Inhibit schedule or modify or delete an existing schedule.

To add a new schedule, select **Add** from the Select Action drop-down menu and click **Go**. Complete the fields for information that will be displayed in the Schedules table. The Start Time field is accompanied by a time chooser clock icon. You can enter hours and minutes manually using 24 hour time designations, or you can use the time chooser to select a start time based on a 12 hour (AM/PM) clock.

Note: Inhibit Reclaim Schedules are *not* permitted to overlap.

To modify an existing schedule, select the radio button from the Select column that appears in the same row as the schedule you want to modify. Select **Modify** from the Select Action drop-down menu and click **Go**. You may modify any field listed in the Schedules table.

To delete an existing schedule, select the radio button from the Select column that appears in the same row as the schedule you want to delete. Select **Delete** from the Select Action drop-down menu and click **Go**. You are prompted to confirm the schedule deletion. Click **Yes** to delete the Inhibit Reclaim schedule, or click **No** to cancel the delete request.

Note: At Release 1.7 and later, all times are entered in reference to the local time of the browser and automatically converted to UTC for use by the TS7700 Virtualization Engine.

Encryption Key Manager Addresses window

Use this window (Figure 8-58 on page 497) to set the encryption key manager addresses in a TS7700 Virtualization Engine.

To watch a tutorial showing the properties of the Encryption Key Manager, click the **View tutorial** link.

If the cluster does not possess a physical library, the Encryption Key Manager addresses page will not be visible on the TS7700 Virtualization Engine management interface.

The Encryption Key Manager assists encryption-enabled tape drives in generating, protecting, storing, and maintaining encryption keys that are used to encrypt information

being written to and decrypt information being read from tape media (tape and cartridge formats).

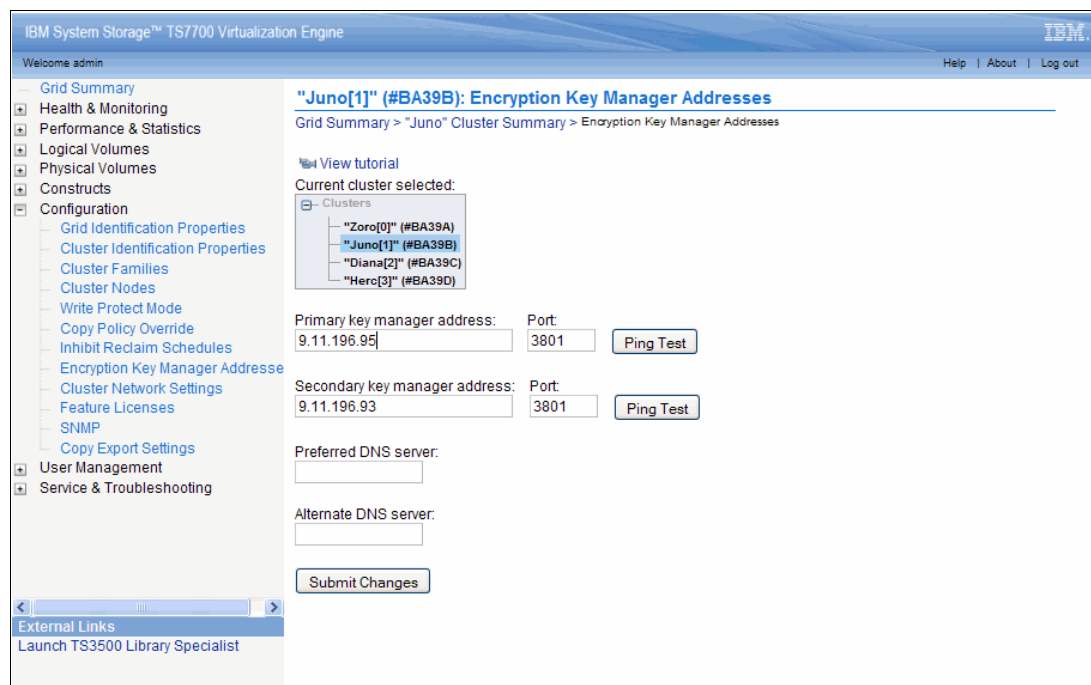


Figure 8-58 MI Encryption Key Manager Addresses window

The following settings are used to configure the TS7700 Virtualization Engine connection to an Encryption Key Manager:

- ▶ Primary key manager address: The key manager server name or IP address that is primarily used.
- ▶ Primary key manager port: The port number of the primary key manager. The default value is 3801. This field is only required if a primary key address is used.
- ▶ Secondary key manager address: The key manager server name or IP address that is used when the primary key manager is unavailable.
- ▶ Secondary key manager port: The port number of the secondary key manager. The default value is 3801. This field is only required if a secondary key address is used.
- ▶ Preferred DNS server: The Domain Name Server (DNS) that is primarily used. DNS addresses are only needed if you specify a symbolic domain name for one of the key manager addresses rather than a numeric IP address. If you need to specify a DNS, be sure to specify both a primary and an alternate so you do not lose access to your Encryption Key Manager because of one of the DNS servers being down or inaccessible. This address can be in IPv4 format.
- ▶ Alternate DNS server: The Domain Name Server that is used in case the preferred DNS server is unavailable. If a Preferred DNS server is specified, be sure to also specify an alternate DNS. This address can be in IPv4 format.
- ▶ Using the Ping Test: Use the Ping Test buttons to check component network connection to a key manager after changing a component's address or port. If you change a key manager address or port and do not submit the change before using the Ping Test button, you receive the following warning:

In order to perform a ping test, you must first submit your address or port changes.

After the Ping Test has been issued, you will receive one of two following messages:

- The ping test against the address “<key manager address>” on port “<port>” was successful.
- The ping test against the address “<key manager address>” on port “<port>” from “<cluster>” has failed.

Additional failure detail is provided with the error text to assist in problem determination.

Click the **Submit** button to save changes to any of the settings. To discard changes and return the field settings to their current values, click the **Reset** button.

Cluster Network Settings window

Use this window (Figure 8-59) to view and alter cluster network settings for a TS7700 Virtualization Engine.

Figure 8-59 MI Cluster Network Settings window

The following information related to cluster network settings is displayed. To change cluster network settings, edit the available fields and click **Submit Changes**.

- Management IP address: The IP address used for connecting to the management interface.

Note: Changing this setting can result in loss of connection to the management interface.

- Primary router address: The IP address of the primary router contained in the TS7700 Virtualization Engine.
- Alternate router address: The IP address of the secondary router contained in the TS7700 Virtualization Engine.

- ▶ Subnet mask: The subnet mask of the network on which the cluster resides
- ▶ Default gateway: The gateway IP address of the network on which the cluster resides.

Feature Licenses window

Use this window (Figure 8-60) to view information about, activate, or remove feature licenses from a TS7700 Virtualization Engine Cluster.

The screenshot shows the 'Juno[0]' cluster configuration page. The 'Current Available Resources' section displays the following data:

Cluster common resources:	Peak data throughput:
Cluster-Wide Disk Cache Enabled: 6 TB	vNode Peak Data Throughput
Cross-Cluster Communication (Grid): Enabled	"Juno" (v0) 600 MB/s

The 'Currently activated feature licenses' table is as follows:

Select	Feature C...	Feature D...	License K...	Node	Node Seri...	Activated	Expires
<input type="radio"/>	5267	1 TB Cache Enablement	e480520a3e4...	"Juno" (h0) hNode	78-H1156	Thursday, May 8, 2008 7:39:47 AM	Never
<input type="radio"/>	5267	1 TB Cache Enablement	e480520a3e4...	"Juno" (h0) hNode	78-H1156	Thursday, May 8, 2008 7:39:48 AM	Never
<input type="radio"/>	5267	1 TB Cache Enablement	e480520a3e4...	"Juno" (h0) hNode	78-H1156	Thursday, May 8, 2008 7:39:49 AM	Never
<input type="radio"/>	5267	1 TB Cache Enablement	e480520a3e4...	"Juno" (h0) hNode	78-H1156	Thursday, May 8, 2008 7:39:50 AM	Never
<input type="radio"/>	5267	1 TB Cache Enablement	e480520a3e4...	"Juno" (h0) hNode	78-H1156	Thursday, May 8, 2008 7:39:51 AM	Never

Figure 8-60 MI Feature Licenses window

In the Feature Licenses window, you have the following fields:

- ▶ Cluster common resources: The Cluster common resources table displays a summary of resources affected by activated features.
 - Cluster Wide Disk Cache Enabled: The amount of disk cache enabled for the entire cluster, in tebibytes (TiB).
 - Cross-Cluster Communication (Grid): Whether cross-cluster communication is enabled on the Grid. If this is enabled, multiple clusters can form a Grid. Possible values are Enabled and Disabled.
- ▶ Peak data throughput: The Peak data throughput table displays, for each vNode, the peak data throughput in mebibytes (MiB).
 - vNode: Name of the vNode.
 - Peak data throughput: The upper limit of the data transfer speed between the vNode and the host, displayed in MiB.
- ▶ Currently activated feature licences: The Currently activated feature licenses table displays a summary of features installed on each cluster.
 - Feature Code: The feature code number of the installed feature.
 - Feature Description: A description of the feature installed by the feature license.

- License Key: The feature's 32-character license key.
- Node: The name and type of the node on which the feature is installed.
- Node Serial Number: The serial number of the node on which the feature is installed.
- Activated: The date and time the feature license was activated.
- Expires: The expiration status of the feature license. Possible values are:
 - Day/Date: The day and date on which the feature license is set to expire.
 - Never: The feature is permanently active and never expires.
 - One-time use: The feature can be used once and has not yet been used.

Use the drop-down menu on the Currently activated feature licenses table to activate or remove a feature license. You can also use this drop-down menu to sort and filter feature license details.

Activate New Feature License

Use this procedure for activating feature licenses in a TS7700 Virtualization Engine.

In the Feature Licenses window (Figure 8-60 on page 499), perform the following steps:

1. Select **Activate New Feature License** from the Select Action drop-down menu and click **Go**.
2. The Activate New Feature License window opens. Enter the 32-character feature license in the available text field and click **Activate** to continue or **Cancel** to cancel the activation.
3. The Confirm Feature Activation window opens and shows the details of the feature license. To activate the feature license, click **Yes**; to cancel activation, click **No**.

Remove Selected Feature License

Use this procedure to remove feature licenses in a TS7700 Virtualization Engine.

From the Feature Licenses page, perform the following steps:

1. Select the radio button from the Select column that appears adjacent to the feature license you want to remove.
2. Select **Remove Selected Feature License** from the Select Action drop-down menu and click **Go**. You will be prompted to confirm your decision to remove a feature license. If you select **Yes**, the license will be removed. If you select **No**, your request to remove will be cancelled.

8.2.9 User management

The topics in this section present information that is related to managing user access in a TS7700 Virtualization Engine. Release 1.6 and later provide a series of enhancements in user access management and changed considerably from the previous release levels.

Two methods of managing user authentication policy are available at TS7700 Virtualization Engine R1.6 and later. They provide the option to use a locally administered authentication policy or the centralized Role-Based Access Control (RBAC) method using an externally managed Storage Authentication Service policy.

Local Authentication Policy is managed and applied within the cluster or clusters participating a Grid. In a multi-cluster Grid configuration user IDs and their associated roles are defined through the management interface on one of the clusters. The user IDs and roles are then propagated through the Grid to all participating clusters.

Storage Authentication Service policy provides the capability to centrally manage user ID's and roles. Storage Authentication Service policies store user and group data on a separate server and map relationships between users, groups, and authorization roles when a user signs into a cluster. Network connectivity to an external System Storage Productivity Center (SSPC) server is required. Each cluster in a Grid can operate its own Storage Authentication Service policy

You may access the following options through the User Management link:

- ▶ **Security Settings:** Use this window to view security settings for a TS7700 Virtualization Engine Grid. From this page, you can also access panels to add, modify, assign, test, and delete security settings.
- ▶ **Roles and Permissions:** Use this window to set and control user roles and permissions for a TS7700 Virtualization Engine Grid.
- ▶ **Change Password:** Use this window to change your password for a TS7700 Virtualization Engine Grid.
- ▶ **SSL Certificates:** Use this window to view, import, or delete SSL certificates to support connection to a Storage Authentication Service server from a TS7700 Virtualization Engine Cluster.
- ▶ **InfoCenter Settings:** Use this page to upload a new TS7700 Virtualization Engine Information Center to the cluster's management interface.

The options for User Management in the management interface for TS7700 Virtualization Engine are shown in Figure 8-61.

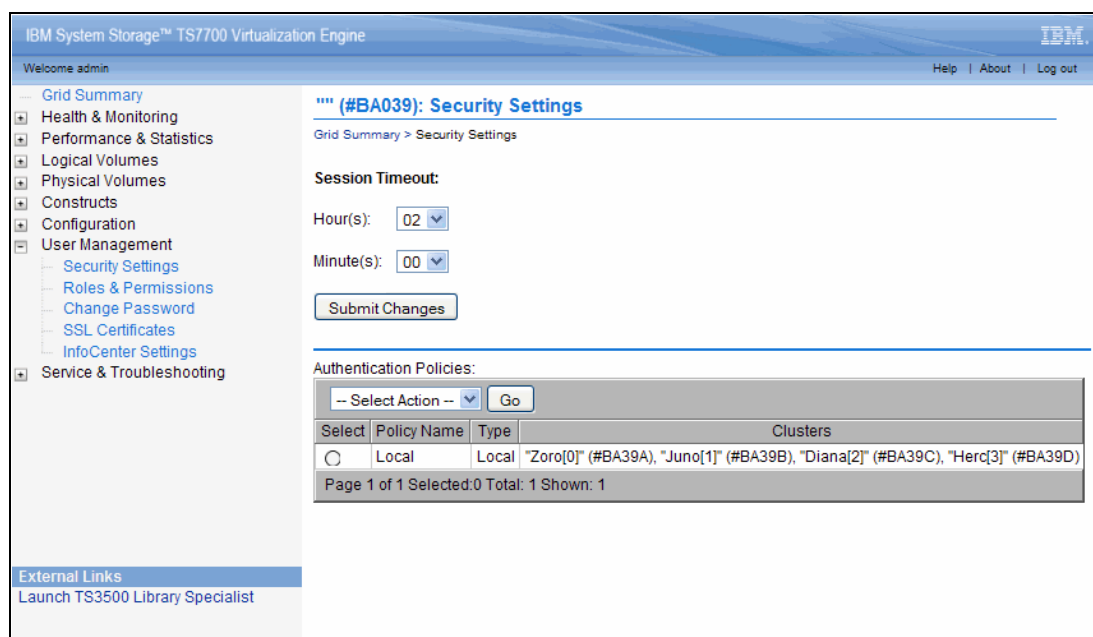


Figure 8-61 R1.7 TS7700 Virtualization Engine User Management options

Figure 8-61 shows the Security Settings summary panel. It is our entry point to enabling security policies.

Use the Session Timeout policy to specify the number of hours and minutes that the management interface can be idle before the current session expires and the user is redirected to the login page. This setting is propagated to all participating clusters in the Grid.

To modify the maximum idle time, select values from the Hours and Minutes drop-down menus and click **Submit Changes**. The parameters for Hours and Minutes are:

- ▶ **Hours:** The number of hours the management interface can be idle before the current session expires. Possible values for this field are 00 through 23.
- ▶ **Minutes:** The number of minutes the management interface can be idle before the current session expires. Possible values for this field are 00 through 55, selected in five-minute increments.

The Authentication Policies table shown in Figure 8-61 on page 501 lists the following information:

- ▶ **Policy Name:** The name of the policy that defines the authentication settings. The policy name is a unique value composed of one to 50 Unicode characters. Heading and trailing blank spaces are trimmed, although internal blank spaces are permitted. After a new authentication policy has been created its policy name cannot be modified.

Note: The Local Policy name is Local and cannot be modified.

- ▶ **Type:** The policy type which can be one of the following values:
 - **Local:** A policy that replicates authorization based on user accounts and assigned roles. It is the default authentication policy. When enabled, it is enforced for all clusters in the Grid. If Storage Authentication Service is enabled, the Local policy is disabled. This policy can be modified to add, change, or delete individual accounts, but the policy itself cannot be deleted.
 - **Storage Authentication Service:** A policy that maps user, group, and role relationships upon user login. Each cluster in a Grid can operate its own Storage Authentication Service policy by means of assignment. However, only one authentication policy can be enabled on any particular cluster within the Grid, even if the same policy is used within other clusters of the same Grid domain. A Storage Authentication Service policy can be modified, but can only be deleted if it is not in use on any cluster in the Grid.
 - **Clusters:** The clusters for which the authentication policy is in force.

Adding a new user to the Local Authentication Policy

A Local Authentication Policy replicates authorization based on user accounts and assigned roles. It is the default authentication policy. In this section we will look at the various panels required to manage the Local Authentication Policy.

To add a user to the Local Authentication Policy for a TS7700 Virtualization Engine Grid, perform the following steps:

1. On the TS7700 Virtualization Engine management interface select **User Management** → **Security Settings** from the left navigation panel.
2. Click the **Select** button next to the **Local policy** name on the **Authentication Policies** table.
3. Select **Modify** from the Select Action drop-down menu and click **Go**.
4. On the **Local Accounts** table, select **Add** from the Select Action drop-down menu and click **Go**.

5. In the Add User window, enter values for the following required fields:
 - Username: The new user's login name. This value must be 1 - 128 characters in length and composed of Unicode characters. Spaces and tabs are not allowed.
 - Role: The role assigned to the user account. The role can be a predefined role or a user-defined role. Possible values are:
 - Operator: The operator has access to monitoring information, but is restricted from changing settings for performance, network configuration, feature licenses, user accounts, or custom roles. The operator is also restricted from inserting and deleting logical volumes.
 - Lead Operator: The lead operator has access to monitoring information and can perform actions for a volume operation. The lead operator has nearly identical permissions to the administrator, but may not change network configuration, feature licenses, user accounts, or custom roles.
 - Administrator: The administrator has the highest level of authority and may view all pages and perform any action, including addition and removal of user accounts. The administrator has access to all service functions and TS7700 Virtualization Engine resources.
 - Manager: The manager has access to monitoring information and performance data and functions, and may perform actions for users, including adding, modifying, or deleting user accounts. The manager is restricted from changing most other settings, including those for logical volume management, network configuration, feature licenses, or custom roles.
 - Custom roles: The administrator can name and define two custom roles by selecting the individual tasks permitted to each custom role. Tasks may be assigned to a custom role in the Roles and assigned permissions table in the Roles & Permissions Properties window.
 - Cluster Access: The clusters to which the user has access. A user can have access to multiple clusters.
6. To complete the operation click **OK**. To abandon the operation and return to the Modify Local Accounts page, click **Cancel**.

Figure 8-62 and Figure 8-63 show specific panels that are used for managing users with the Local Authentication Policy method.

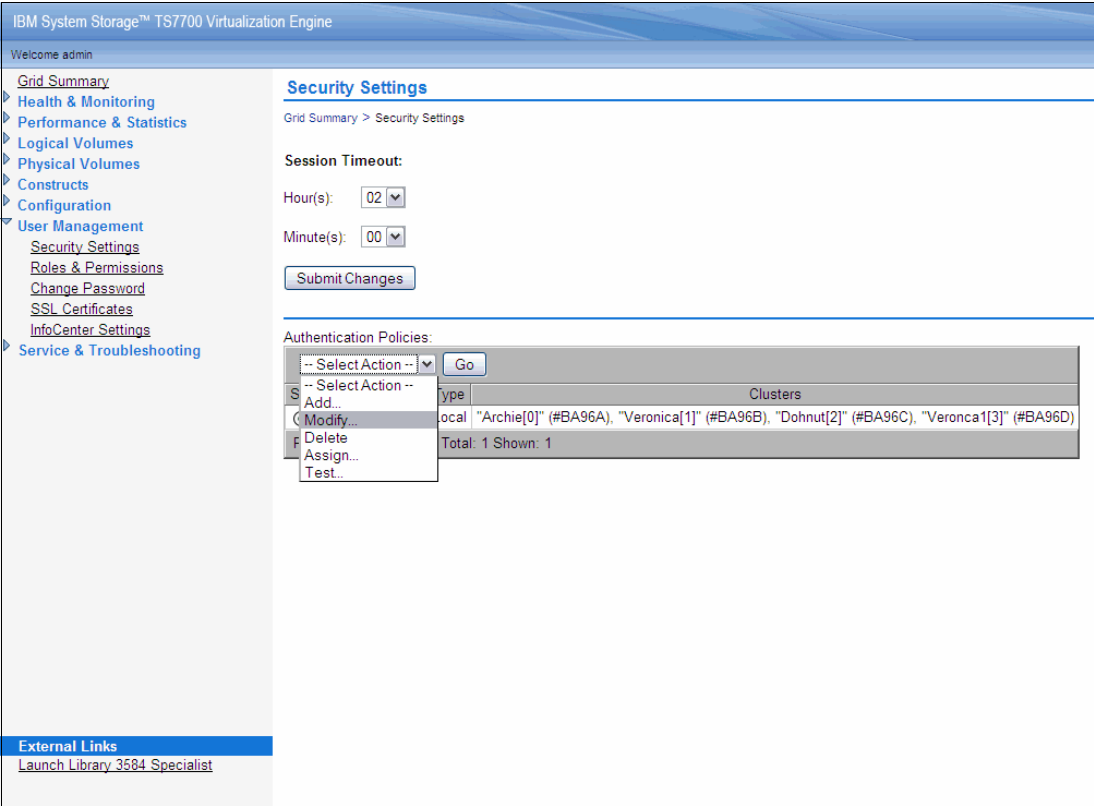


Figure 8-62 Creating a new user: Part 1

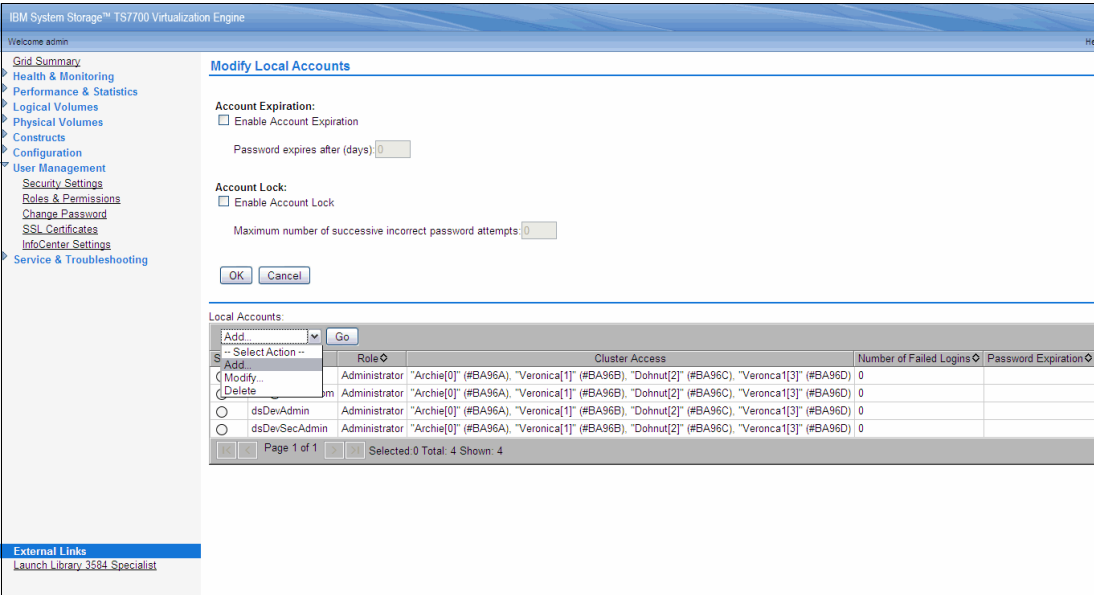


Figure 8-63 Creating a new user: Part 2

Modifying user or group of the Local Authentication Policy

Use this window (Figure 8-64 on page 506) to modify a user or group property for a TS7700 Virtualization Engine Grid.

Note: Passwords for the users are changed from this window also.

To modify a user account belonging to the Local Authentication Policy, perform these steps:

1. On the TS7700 Virtualization Engine management interface select **User Management** → **Security Settings** from the left navigation pane.
2. Click the **Select** button next to the **Local policy** name on the **authentication policies** table.
3. Select **Modify** from the Select Action drop-down menu and click **Go**.
4. On the **Local Accounts** table, click the **Select** radio button next to the username of the policy you want to modify.
5. Select **Modify** from the Select Action drop-down menu and click **Go**.
6. Modify the values for the any of the following fields:
 - Role: The role assigned to the user account. Possible values are as follows:
 - Operator: The operator has access to monitoring information, but is restricted from changing settings for performance, network configuration, feature licenses, user accounts, or custom roles. The operator is also restricted from inserting and deleting logical volumes.
 - Lead Operator: The lead operator has access to monitoring information and can perform actions for volume operation. The lead operator has nearly identical permissions to the administrator, but may not change network configuration, feature licenses, user accounts, or custom roles.
 - Administrator: The administrator has the highest level of authority and may view all pages and perform any action, including addition and removal of user accounts. The administrator has access to all service functions and TS7700 Virtualization Engine resources.
 - Manager: The manager has access to monitoring information and performance data and functions, and may perform actions for users, including adding, modifying, or deleting user accounts. The manager is restricted from changing most other settings, including those for logical volume management, network configuration, feature licenses, or custom roles.
 - Custom roles: The administrator can name and define two custom roles by selecting the individual tasks permitted to each custom role. Tasks may be assigned to a custom role in the Roles and assigned permissions table from the Roles & Permissions Properties window.
 - Cluster Access: The clusters to which the user has access. A user can have access to multiple clusters.
7. To complete the operation, click **OK**. To abandon the operation and return to the Modify Local Accounts page, click **Cancel**.

Note: You cannot modify the User Name or Group Name. Only the role and the clusters to which it is applied can be modified.

The Modify Local Account panel is shown in Figure 8-64. In the Cluster Access table you can use the **Select** check box to toggle all the cluster check boxes on and off.

IBM System Storage™ TS7700 Virtualization Engine

Welcome admin

[Grid Summary](#)

▶ [Health & Monitoring](#)

▶ [Performance & Statistics](#)

▶ [Logical Volumes](#)

▶ [Physical Volumes](#)

▶ [Constructs](#)

▶ [Configuration](#)

▶ **User Management**

[Security Settings](#)

[Roles & Permissions](#)

[Change Password](#)

[SSL Certificates](#)

[InfoCenter Settings](#)

▶ [Service & Troubleshooting](#)

Modify Local Account

Username:

Password:

Confirm password:

Role:

Clusters Access:

<input type="checkbox"/> Select	Cluster
<input checked="" type="checkbox"/>	"Archie[0]" (#BA96A)
<input checked="" type="checkbox"/>	"Veronica[1]" (#BA96B)
<input checked="" type="checkbox"/>	"Dohnut[2]" (#BA96C)
<input checked="" type="checkbox"/>	"Veronca1[3]" (#BA96D)

External Links

[Launch Library 3584 Specialist](#)

Figure 8-64 Modify user or group

Adding a Storage Authentication Service policy

A Storage Authentication Service Policy maps user, group, and role relationships upon user login with the assistance of a System Storage Productivity Center (SSPC). In this section we will look at the various panels required to manage the Storage Authentication Service Policy.

To add a new Storage Authentication Service Policy for a TS7700 Virtualization Engine Grid, perform the following steps:

1. On the TS7700 Virtualization Engine management interface select **User Management** → **Security Settings** from the left navigation panel.
2. On the **Authentication Policies** table select **Add** from the **Select Action** drop-down menu.
3. Click **Go** to open the **Add Storage Authentication Service Policy** panel shown in Figure 8-65 on page 507. The following fields are available for completion:
 - a. Policy Name: The name of the policy that defines the authentication settings. The policy name is a unique value composed of one to 50 Unicode characters. Heading

and trailing blank spaces are trimmed, though internal blank spaces are permitted. After a new authentication policy has been created its policy name cannot be modified.

- b. Primary Server URL: The primary URL for the Storage Authentication Service. The value in this field is composed of one to 256 Unicode characters and takes the following format:

`https://<server_IP_address>:secure_port/TokenService/services/Trust`

- c. Alternate Server URL: The alternate URL for the Storage Authentication Service if the primary URL cannot be accessed. The value in this field is composed of one to 256 Unicode characters and takes the following format:

`https://<server_IP_address>:secure_port/TokenService/services/Trust`

Note: If the Primary or Alternate Server URL uses the HTTPS protocol, a certificate for that address must be defined on the SSL Certificates page.

- Server Authentication: Values in the following fields are required if WebSphere® Application Server security is enabled on the WebSphere Application Server that is hosting the Authentication Service. If WebSphere Application Server security is disabled, the following fields are optional:
 - Client User Name: The user name used with HTTP basic authentication for authenticating to the Storage Authentication Service.
 - Client Password: The password used with HTTP basic authentication for authenticating to the Storage Authentication Service.
4. To complete the operation click **OK**. To abandon the operation and return to the Security Settings page, click **Cancel**.

The Figure 8-65 shows an example of a completed Add Storage Authentication Service Policy panel.

IBM System Storage™ TS7700 Virtualization Engine

Welcome admin

Grid Summary

Health & Monitoring

- Cluster Summary
- Operator Interventions
- Physical Tape Drives
- Virtual Tape Drives
- Physical Media Inventory
- Tape Volume Cache
- Physical Library
- Operation History

Performance & Statistics

Logical Volumes

Physical Volumes

Constructs

Configuration

User Management

- Security Settings
- Roles & Permissions
- Change Password
- SSL Certificates
- InfoCenter Settings

Service & Troubleshooting

External Links

Launch TS3500 Library Specialist

#BA039: Add Storage Authentication Service Policy

Server Settings

* Policy Name:
Bali

* Primary Server URL:
https://112.40.29.1:16311/TokenService/services/Trust Retrieve Signer Information

Alternate Server URL:
https://112.40.29.2:16311/TokenService/services/Trust Retrieve Signer Information

Server Authentication

A user ID and password are required if authentication is enabled on the server hosting the authentication service.

User ID:
bashas

Password:

OK Cancel

Figure 8-65 Add Storage Authentication Service Policy

After you click **OK** to confirm the creation of the new Storage Authentication Policy, the panel shown in Figure 8-66 on page 508, opens. In the Authentication Policies table, no clusters are assigned to the newly created policy therefore the Local Authentication Policy is enforced.

When in this state the newly created policy can be deleted as it is not applied to any of the clusters.

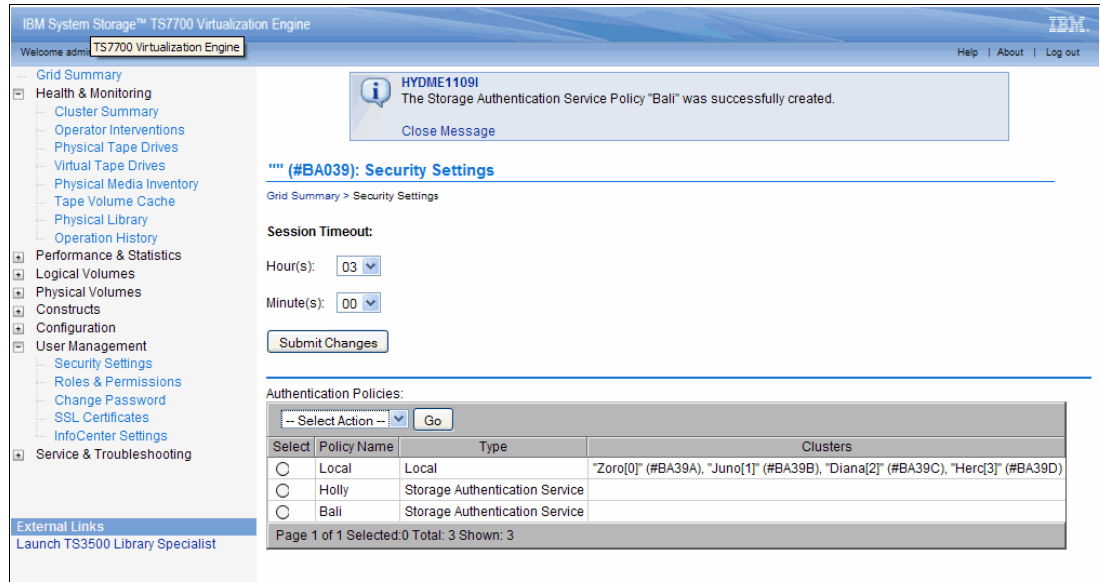


Figure 8-66 Addition of Storage Authentication Service Policy completed

Adding a User to a Storage Authentication Policy

To add a new user to a Storage Authentication Service Policy for a TS7700 Virtualization Engine Grid, perform the following steps:

- ▶ On the TS7700 Virtualization Engine management interface select **User Management** → **Security Settings** from the left navigation panel.
- ▶ In the Authentication Policies table select **Modify** from the Select Action drop-down menu as shown in Figure 8-67.

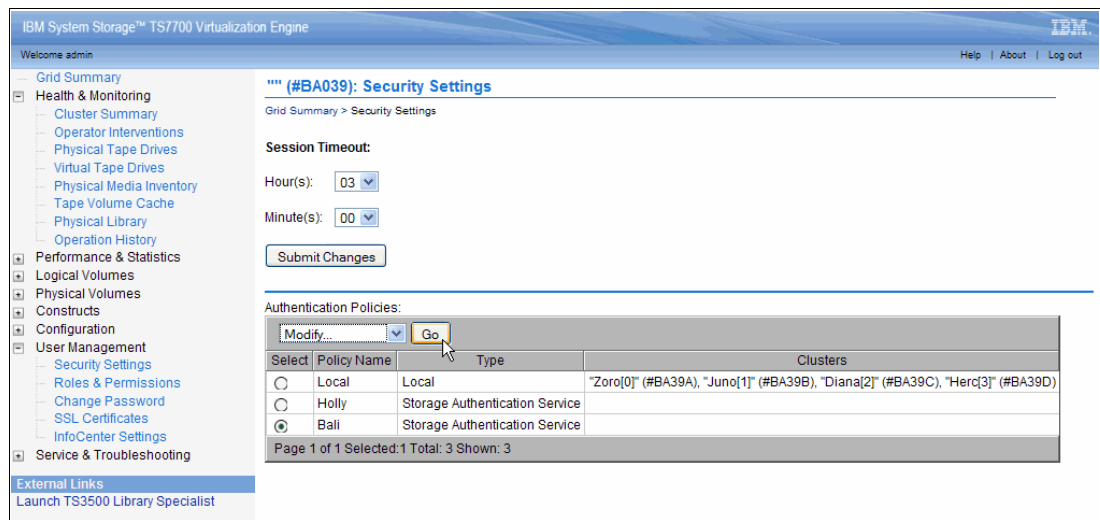


Figure 8-67 Selecting Modify Security Settings

- ▶ Click **Go** to open the Modify Storage Authentication Service Policy panel shown in Figure 8-68 on page 509.

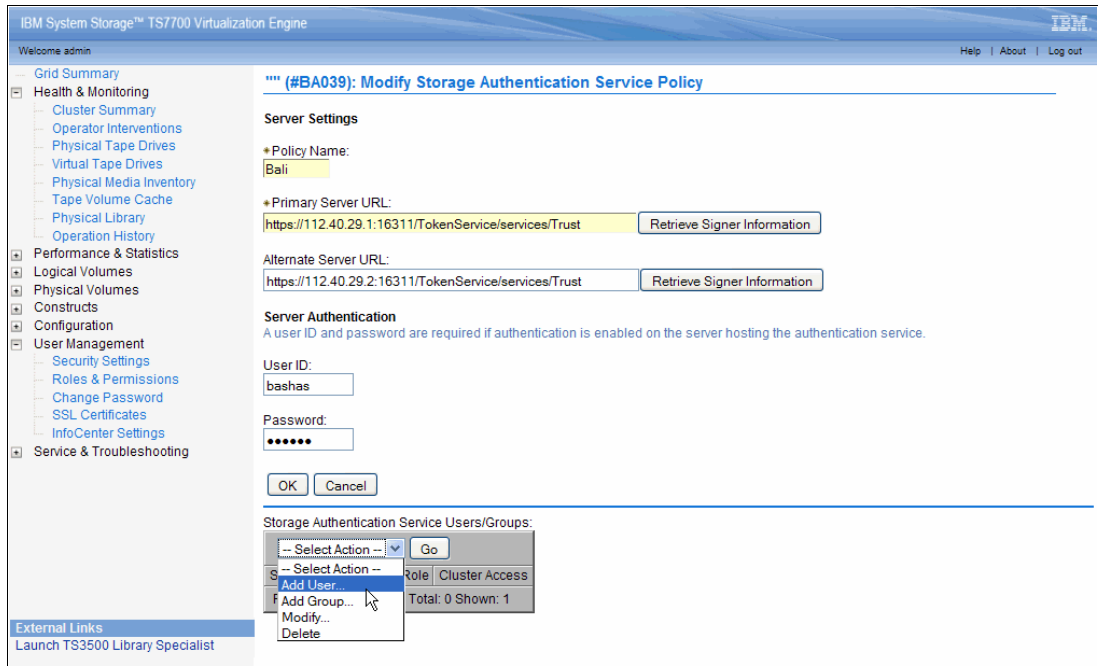


Figure 8-68 Adding a User to the Storage Authentication Service Policy

5. In the Modify Storage Authentication Service Policy panel shown in Figure 8-68, navigate to the **Storage Authentication Service Users/Groups** table at the bottom:
 - a. Select **Add User** from the Select Action drop-down menu.
 - b. Click **Go** to open the Add Storage Authentication Service User panel shown in Figure 8-69.

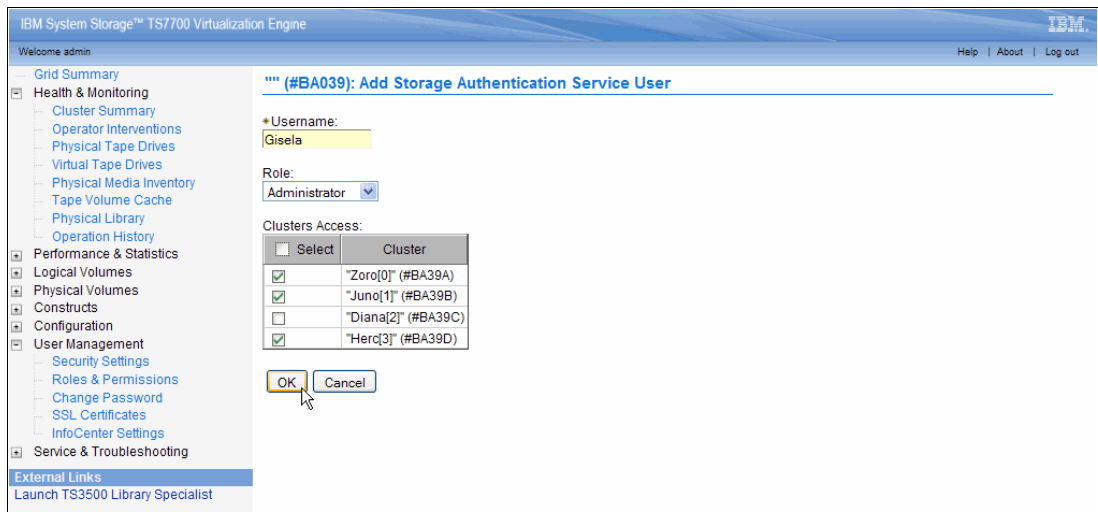


Figure 8-69 Defining Username, Role and Cluster Access Permissions

6. In the Add Storage Authentication Service User window, enter values for the following required fields:
 - Username: The new user's login name. This value must be 1 to 128 characters in length and composed of Unicode characters. Spaces and tabs are not allowed.
 - Role: The role assigned to the user account. The role can be a predefined role or a user-defined role. Possible values are as follows:
 - Operator: The operator has access to monitoring information, but is restricted from changing settings for performance, network configuration, feature licenses, user accounts, or custom roles. The operator is also restricted from inserting and deleting logical volumes.
 - Lead Operator: The lead operator has access to monitoring information and can perform actions for a volume operation. The lead operator has nearly identical permissions to the administrator, but may not change network configuration, feature licenses, user accounts, or custom roles.
 - Administrator: The administrator has the highest level of authority and may view all pages and perform any action, including addition and removal of user accounts. The administrator has access to all service functions and TS7700 Virtualization Engine resources.
 - Manager: The manager has access to monitoring information and performance data and functions, and may perform actions for users, including adding, modifying, or deleting user accounts. The manager is restricted from changing most other settings, including those for logical volume management, network configuration, feature licenses, or custom roles.
 - Custom roles: The administrator can name and define two custom roles by selecting the individual tasks permitted to each custom role. Tasks may be assigned to a custom role in the Roles and assigned permissions table in the Roles & Permissions Properties window.
 - Cluster Access: The clusters to which the user has access. A user can have access to multiple clusters.
7. To complete the operation click **OK**. To abandon the operation and return to the Modify Local Accounts page, click **Cancel**.
8. Click **OK** after the fields are complete; a user is added to the Storage Authentication Service Policy as shown in Figure 8-70 on page 511.

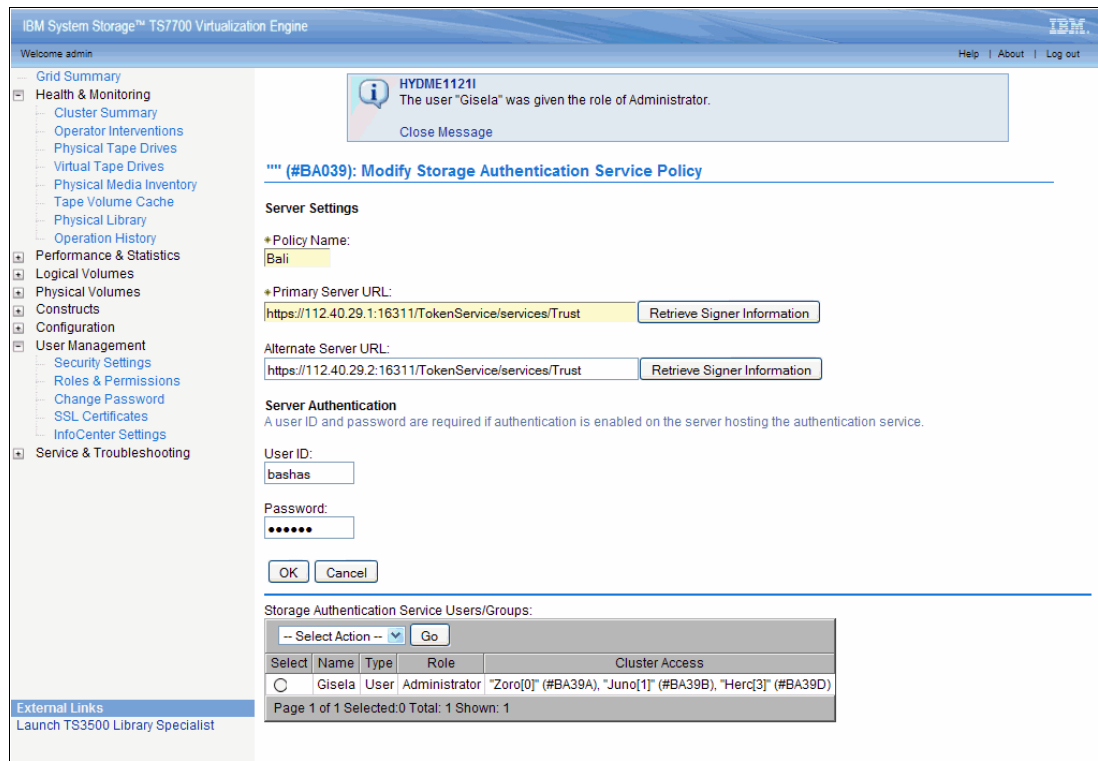


Figure 8-70 Successful Addition of Username to Storage Authentication Service Policy

Assigning Clusters to a Storage Authentication Policy

Clusters participating in a multi-cluster grid can have unique Storage Authentication policies active. To assign an authentication policy to one or more clusters you must have authorization to modify authentication privileges under the new policy. To verify that you have sufficient privileges with the new policy, you must enter a user name and password recognized by the new authentication policy.

To add a new user to a Storage Authentication Service Policy for a TS7700 Virtualization Engine Grid, perform the following steps:

- ▶ On the TS7700 Virtualization Engine management interface select **User Management** → **Security Settings** from the left navigation panel.
- ▶ In the Authentication Policies table, select **Assign** from the Select Action drop-down menu, as shown in Figure 8-71 on page 512.

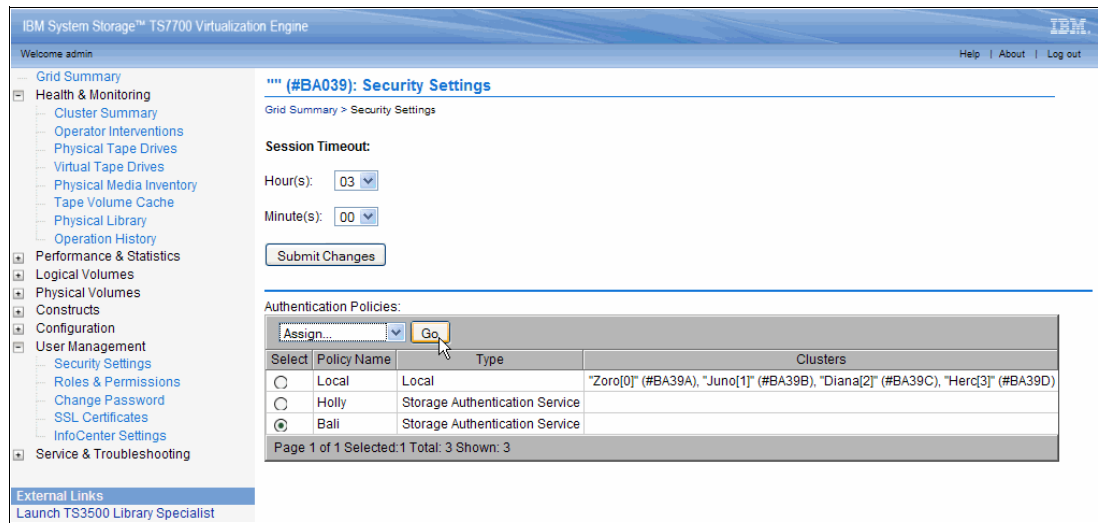


Figure 8-71 Assigning Storage Authentication Service Policy to Grid Resources

9. Click **Go** to open the Assign Authentication Policy panel, shown in Figure 8-72.

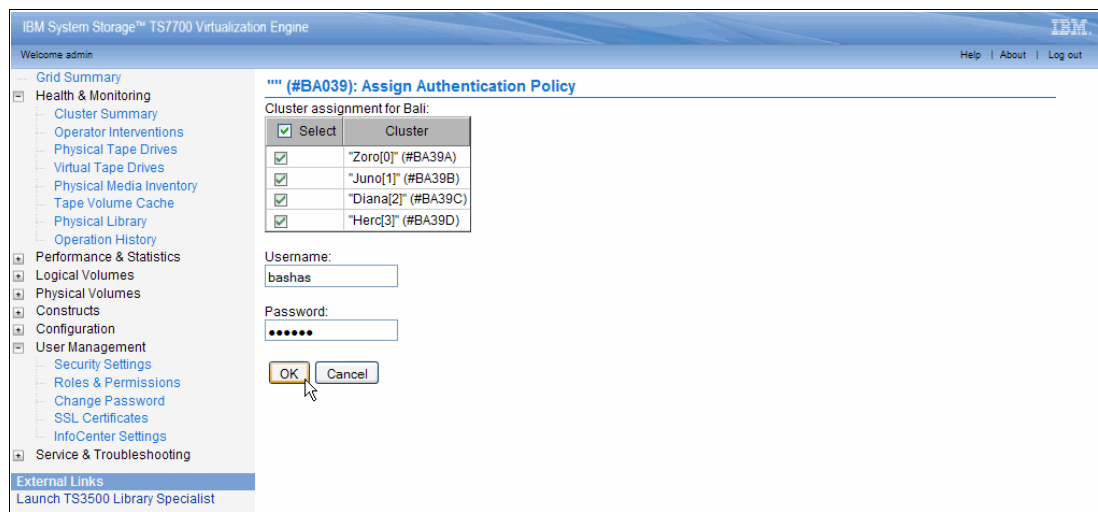


Figure 8-72 Cluster Assignment Selection for Storage Authentication Service Policy

10. To apply the authentication policy to a cluster, select the check box next to the cluster's name.

To assign an authentication policy to one or more clusters you must have authorization to modify authentication privileges under the new policy. To verify that you have sufficient privileges with the new policy you must enter a user name and password recognized by the new authentication policy. Enter values for the following fields:

- User Name: Your user name for the TS7700 Virtualization Engine management interface.
- Password: Your password for the TS7700 Virtualization Engine management interface.

11. To complete the operation click **OK**. To abandon the operation and return to the *Security Settings* panel, click **Cancel**.

Figure 8-73 shows successful assignment of the clusters to the Storage Authentication Service Policy.

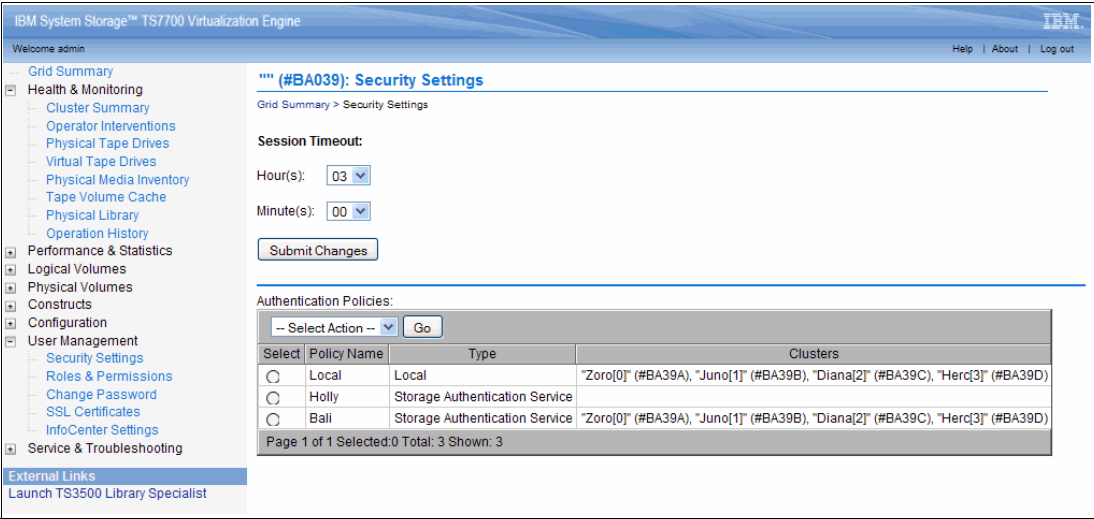


Figure 8-73 Successful Cluster Assignment to Storage Authentication Service Policy

Deleting a Storage Authentication Policy

You can delete a Storage Authentication Service policy if it is not in force on any cluster. You cannot delete the Local policy. In the Authentication Policies table in Figure 8-74 no clusters are assigned to the policy to be deleted, therefore the policy can be deleted. If clusters are assigned to the policy, use **Modify** from the Select Action drop down to remove the assigned clusters.

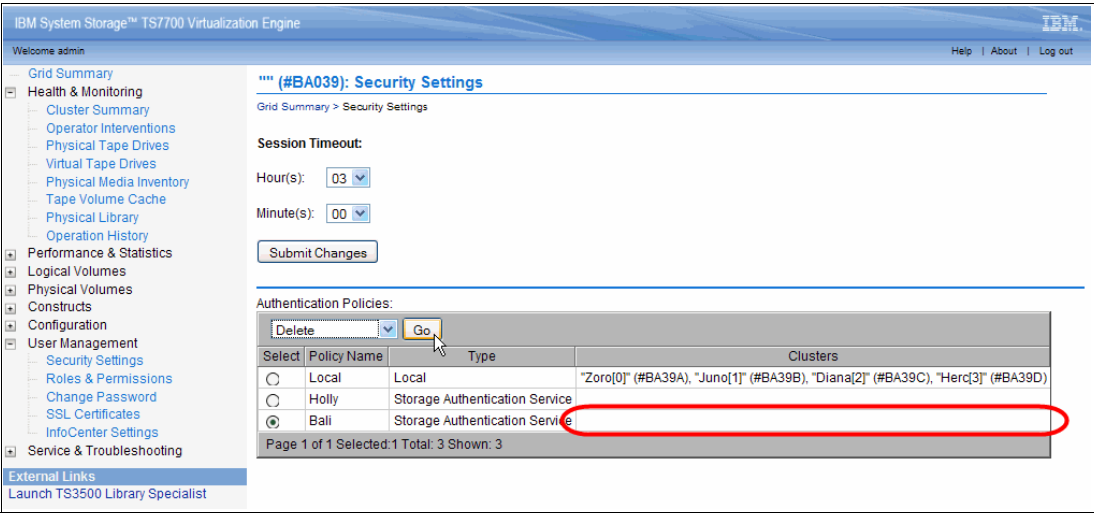


Figure 8-74 Deleting a Storage Authentication Service Policy

To delete a Storage Authentication Service policy from a TS7700 Virtualization Engine Grid, perform the following steps:

1. On the TS7700 Virtualization Engine management interface select **User Management** → **Security Settings** from the left navigation panel.
2. In the Authentication Policies table, select **Delete** from the Select Action drop-down menu as shown Figure 8-74 on page 513. From the Security Settings page, navigate to the Authentication Policies table and perform the following steps:
 - a. Select the radio button next to the policy you want to delete.
 - b. Select **Delete** from the Select Action drop-down menu.
 - c. Click **Go** to open the Confirm Delete Storage Authentication Service policy page.
 - d. Click **OK** to delete the policy and return to the *Security Settings* panel, or click **Cancel** to abandon the delete operation and return to the *Security Settings* panel.

The delete operation completed successfully as shown in Figure 8-75.

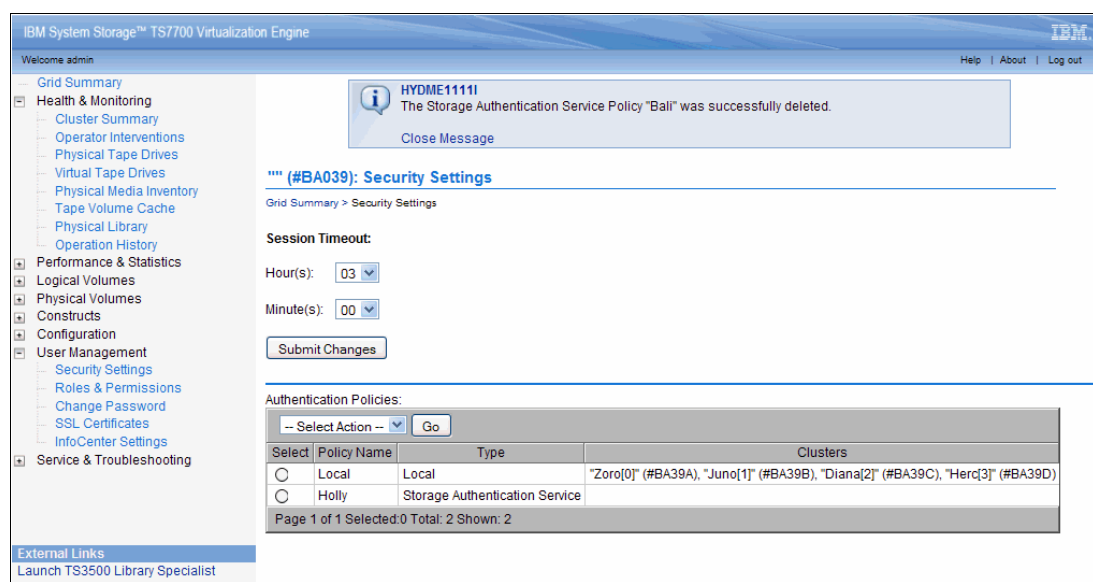


Figure 8-75 Successful Deletion of a Storage Authentication Service Policy

Roles & Permissions window

You can use this window (Figure 8-76) to set and control user roles and permissions for a TS7700 Virtualization Engine Grid.

To watch a tutorial showing how to set up user access, click the **View tutorial** link.

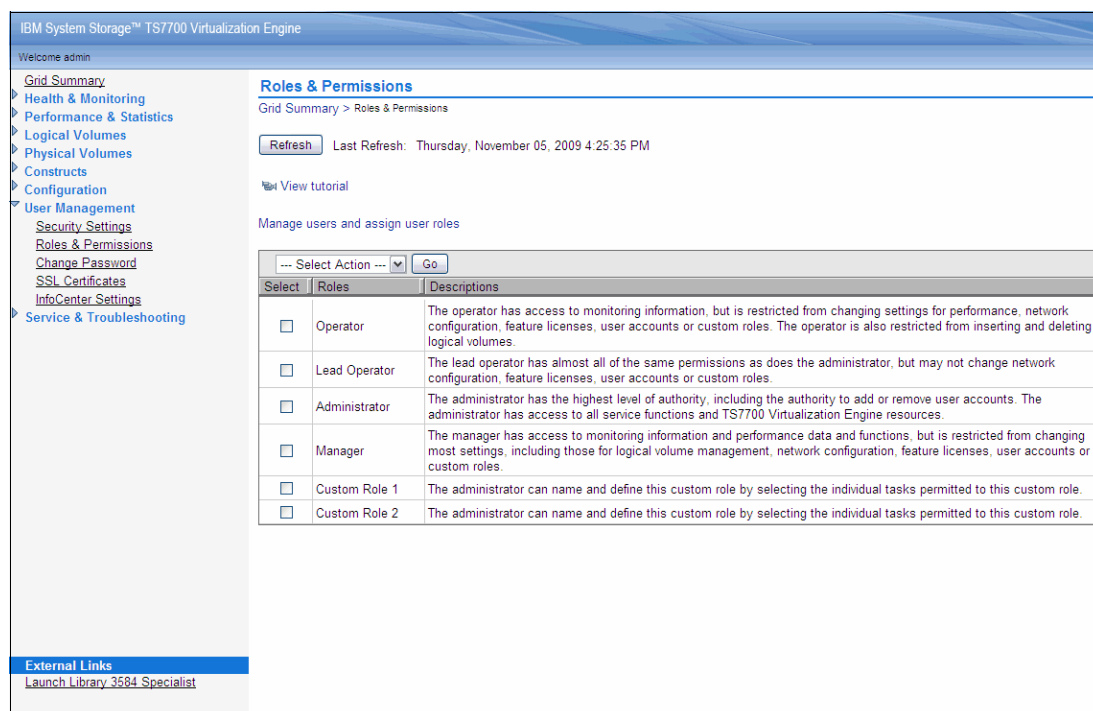


Figure 8-76 TS7700 Virtualization Engine MI Roles & Permissions window

Figure 8-76 shows the list of user roles and a summary of each role:

- ▶ **Operator:** The operator has access to monitoring information, but is restricted from changing settings for performance, network configuration, feature licenses, user accounts, or custom roles. The operator is also restricted from inserting and deleting logical volumes.
- ▶ **Lead Operator:** The lead operator has access to monitoring information and can perform actions for volume operation. The lead operator has nearly identical permissions to the administrator, but may not change network configuration, feature licenses, user accounts, or custom roles.
- ▶ **Administrator:** The administrator has the highest level of authority and may view all windows and perform any action, including addition and removal of user accounts. The administrator has access to all service functions and TS7700 Virtualization Engine resources.
- ▶ **Manager:** The manager has access to monitoring information and performance data and functions, and may perform actions for users. The manager is restricted from changing most settings, including those for logical volume management, network configuration, feature licenses, user accounts, or custom roles.
- ▶ **Custom roles:** The administrator can name and define two custom roles by selecting the individual tasks permitted to each custom role. Tasks may be assigned to a custom role in the Roles and Assigned Permissions view.

Roles and Assigned Permissions table

The Roles and Assigned Permissions table is a dynamic table that displays the complete list of TS7700 Virtualization Engine Grid tasks and the permissions that are assigned to selected user roles.

To view the Roles and Assigned Permissions table, perform the following steps:

1. Select the check box to the left of the role to be displayed. You may select more than one role to display a comparison of permissions.
2. Select **Properties** from the Select Action menu.
3. Click **Go**.

The first column of the Roles and Assigned Permissions table lists all the tasks available to users of the TS7700 Virtualization Engine. Subsequent columns show the assigned permissions for selected role (or roles). A checkmark denotes permitted tasks for a user role. A null dash (-) denotes prohibited tasks for a user role. Permissions for predefined user roles may not be modified. You may modify permissions for custom roles in the Roles and Assigned Permissions table. You can modify only one custom role at a time.

To modify a custom role, perform the following steps:

1. Enter a unique name for the custom role in the Name of Custom Role field.
2. Modify the custom role to fit your requirements by selecting (permitting) or deselecting (prohibiting) tasks. Selecting or deselecting a parent task similarly affects any child tasks. However, a child task may be selected or deselected independently of a parent task. You may apply the permissions of a predefined role to a custom role by selecting a role from the Role Template drop-down menu and clicking **Apply**. You may then customize the permissions by selecting or deselecting tasks.
3. After all tasks for the custom role have been selected, click **Submit Changes** to activate the new custom role.

Note: You may apply the permissions of a predefined role to a custom role by selecting a role from the Role Template drop-down menu and clicking **Apply**. You may then customize the permissions by selecting or deselecting tasks.

Change Password window

Use the window shown in Figure 8-77 on page 517 to change your password for a TS7700 Virtualization Engine Grid.

Note: This window is visible only when the current security policy is the Local policy.

Password rules

Adhere to the following rules when you set the password:

- ▶ Passwords must be at least six alphanumeric characters in length, but no more than 16 alphanumeric characters in length.
- ▶ Passwords must contain at least one number.
- ▶ The first and last characters of a password cannot be numbers.
- ▶ Your password cannot contain your user name.

Setting the password

To set the password, complete the following fields:

- ▶ Current Password: Your current password.
- ▶ New Password: Your new password, which must be compliant with the password rules.
- ▶ Re-enter New Password: Your new password, exactly as you entered it into the New Password field.

After you complete these fields, click **Submit**.

Figure 8-77 shows an overview of the Change Password window.

The screenshot shows the 'Change Password' window within the IBM System Storage™ TS7700 Virtualization Engine interface. The window has a blue header bar with the title 'IBM System Storage™ TS7700 Virtualization Engine' and a sub-header 'Welcome admin'. On the left side, there is a navigation menu with the following items: 'Grid Summary', 'Health & Monitoring', 'Performance & Statistics', 'Logical Volumes', 'Physical Volumes', 'Constructs', 'Configuration', 'User Management' (expanded), 'Security Settings', 'Roles & Permissions', 'Change Password' (selected), 'SSL Certificates', 'InfoCenter Settings', and 'Service & Troubleshooting'. The main content area on the right is titled 'Change Password' and contains the following fields: 'Current Password:', 'New Password:', and 'Re-enter New Password:'. Each field has a corresponding text input box. Below these fields is a 'Submit' button. At the bottom of the window, there is an 'External Links' section with a link to 'Launch Library 3584 Specialist'.

Figure 8-77 Change Password window

SSL Certificates window

Use this window (Figure 8-78) to view, import, or delete SSL certificates that support secure connections to a Storage Authentication Service server from a TS7700 Virtualization Engine Cluster. If a Primary or Alternate Server URL, defined by a Storage Authentication Service policy, uses the HTTPS protocol, a certificate for that address must be defined in this window

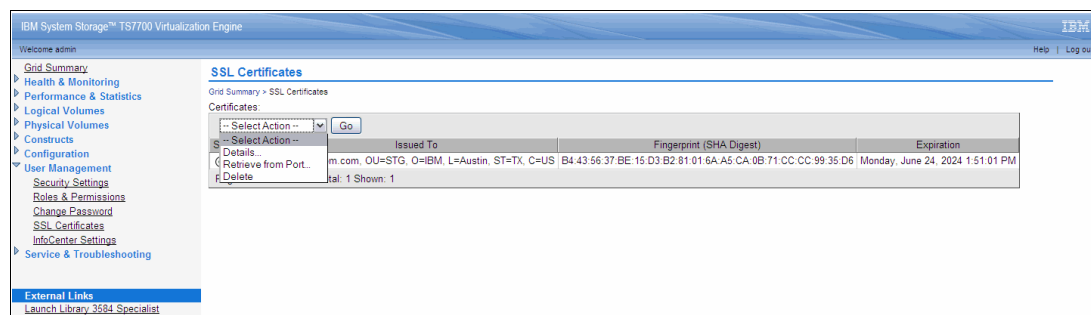


Figure 8-78 SSL Certificates window

The Certificates table displays the following identifying information for SSL certificates on the cluster:

- ▶ **Alias:** A unique name to identify the certificate on the machine.
- ▶ **Issued To:** The distinguished name of the entity requesting the certificate.
- ▶ **Fingerprint:** A number that specifies the Secure Hash Algorithm (SHA hash) of the certificate. This number can be used to verify the hash for the certificate at another location, such as the client side of a connection.
- ▶ **Expiration:** The expiration date of the signer certificate for validation purposes.

To import a new SSL certificate, perform the following steps:

1. Select **Retrieve from port** from the Select Action drop-down menu and click **Go**. The Retrieve from Port window opens.
2. Enter the host and port from which the certificate is retrieved and a unique value for the alias.
3. Click **Retrieve Signer Information**. To import the certificate, click **OK**. To abandon the operation and return to the SSL Certificates window, click **Cancel**.

To delete an existing SSL certificate, perform the following steps:

1. Select the radio button next to the certificate you want to delete, select **Delete** from the Select Action drop-down menu, and click **Go**.

The Confirm Delete SSL Certificate window opens and prompts you to confirm your decision to delete the SSL certificate.

2. Click **OK** to delete the certificate and return to the SSL Certificates window. Click **Cancel** to abandon the delete operation and return to the SSL Certificates window.

InfoCenter Settings window

Use this window (Figure 8-79) to upload a new TS7700 Virtualization Engine Information Center to the cluster's management interface.

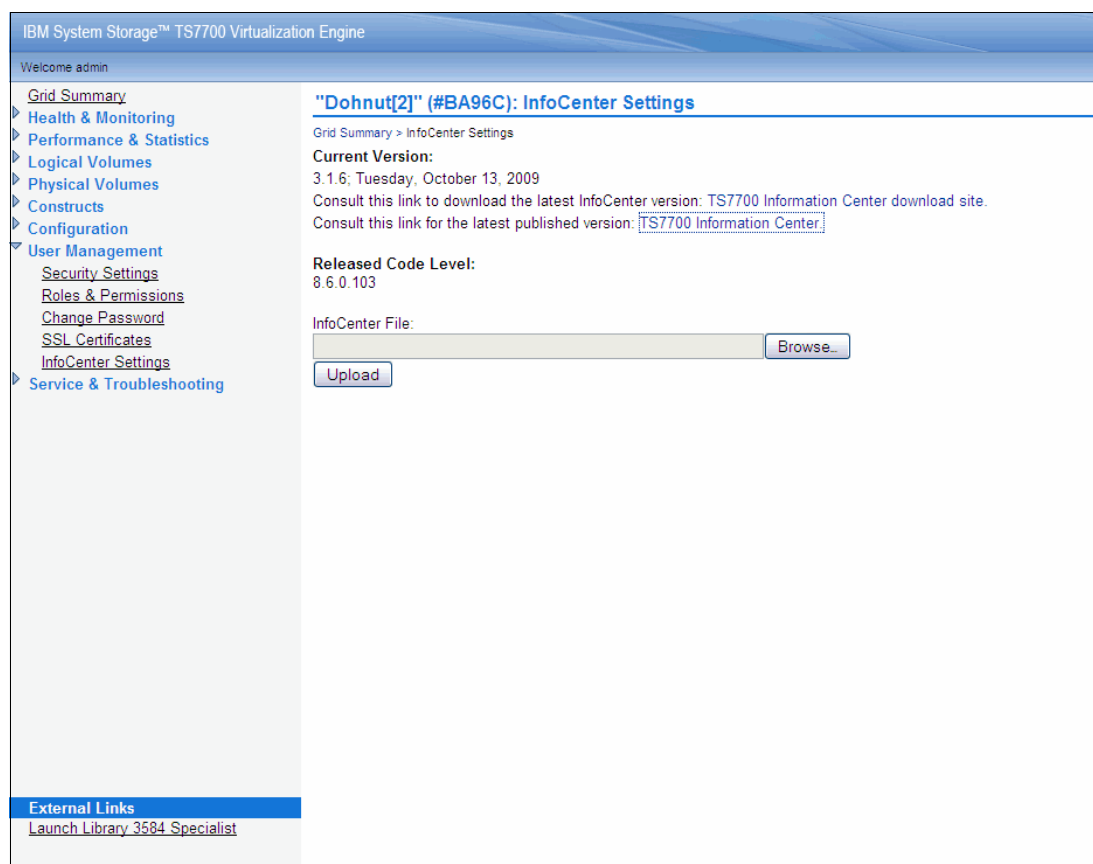


Figure 8-79 InfoCenter Settings window

This window has the following items:

- ▶ Current Version section, where you can identify or access the following items:
 - Identify the version level and date of the Information Center installed on the cluster.
 - Access a product database where you can download a JAR file containing a newer version of the Information Center.
 - Access an external site displaying the most recent published version of the Information Center.
- ▶ The TS7700 Information Center download site link.

Click this link to open the Fix Central product database so you can download a new version of the TS7700 Virtualization Engine Information Center as a .jar file, as follows:

 - a. Select **System Storage** from the Product Group menu and click **Continue**.
 - b. Select **Tape Systems** from the Product Family menu.
 - c. Select **TS7700 Virtualization Engine** from the Product menu.
 - d. Click **Go**.
 - e. Select **Download** from the Support & downloads navigation box.
 - f. Select **Downloadable files** from the Recommended fix navigation box.

- g. Select **Information Centers** from the Results heading.
- h. Select your desired Information Center .jar file.
- i. Follow the instructions on the Fix Central download page to download the JAR file.
- ▶ **TS7700 Information Center link**
 To verify that the installed version of the Information Center is the most recent version, compare the version level and date of the installed Information Center with that of the latest published Information Center. You can open and view the latest published Information Center by clicking the **TS7700 Information Center** link.

 This link opens the TS7700 Virtualization Engine Customer Information Center, an external site. This site is best viewed in a new tab or window. After you open the new tab or window, you can see the version number and date of the TS7700 Virtualization Engine Customer Information Center in the right-hand pane of the new tab or window under the Version heading. If this version is newer than the Current Version displayed in the InfoCenter Settings window, contact your SSR to obtain an updated Information Center file.
- ▶ **Released Code Level**
 This section displays the level of microcode installed on the accessing cluster. Compare this value to the Released Code Level shown in the latest published TS7700 Virtualization Engine Customer Information Center.
- ▶ **InfoCenter File**
 Use this field to locate and upload the new Information Center .jar file on a local drive. Updated versions of the Information Center are distributed as they are made available. The Information Center is distributed as a single .jar file with a file name similar to `com.ibm.storage.ts7700.doc.3.1.7.v20091013.jar`.

 After you receive a new .jar file containing the updated Information Center (either from the Fix Central database or from an SSR), save the .jar file to a local directory.

 To upload and install the new Information Center, perform the following steps:
 - a. Click the **Browse** button to open the File Upload window.
 - b. Navigate to the folder containing the new .jar file.
 - c. Highlight the new .jar file name and click **Open**.
 - d. Click **Upload** to install the new Information Center on the cluster's management interface.

8.2.10 Service and troubleshooting

The following topics in this section present information about performing service operations and troubleshooting problems for the TS7700 Virtualization Engine:

- ▶ Service mode
- ▶ Ownership Takeover mode
- ▶ Repair Logical Volumes
- ▶ Copy Export Settings
- ▶ Copy Export Recovery
- ▶ Copy Export Recovery Status
- ▶ Cluster Nodes Online/Offline
- ▶ Backup Settings
- ▶ Restore Settings
- ▶ Cluster Shutdown

Service Mode window

Use this window (Figure 8-80) to put a TS7700 Virtualization Engine Cluster into Service mode.

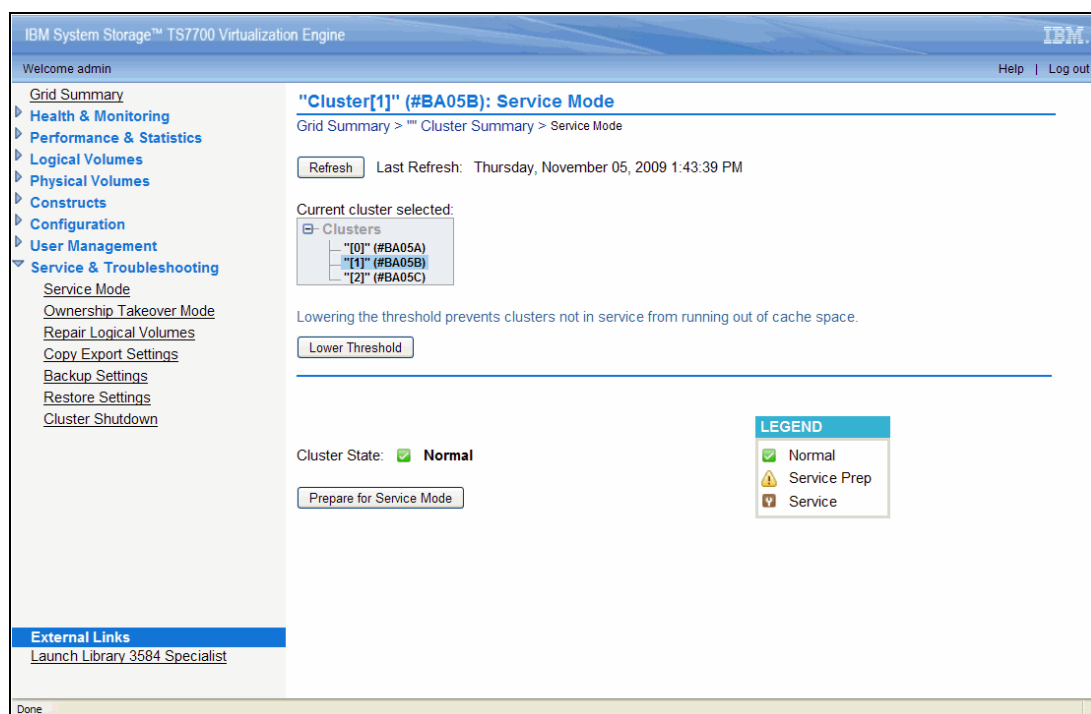


Figure 8-80 MI Service Mode window for a hybrid TS7700 Virtualization Engine configuration

Release 1.6 and later provide support for a hybrid TS7700 Virtualization Engine Grid. A hybrid TS7700 Virtualization Engine Grid is defined as a grid that combines TS7700 Virtualization Engine Clusters that both do and do not attach to a physical library. Additional functions are provided for cache management to reduce the performance impact of scheduled maintenance events to the TS7700 Virtualization Engine Grid.

For a TS7720 Virtualization Engine cluster in a hybrid grid, you may use the Lower Threshold button to lower the required threshold at which logical volumes are removed from cache. When the threshold is lowered, additional logical volumes already copied to another cluster are removed, creating additional cache space for host operations. This step is necessary to ensure that logical volume copies can be made and validated prior to a Service mode event. The default Removal Threshold is equal to 95% of the cache Used Size minus 2 TB (see the Used Size field in the Tape Volume Cache window). You can lower the threshold to any value between 4 TB and the remainder of the Used Size minus 2 TB.

In a TS7700 Virtualization Engine Grid, Service Prep can occur on only one cluster at any one time. If Service Prep is attempted on a second cluster at the same time, the attempt fails. After Service Prep has completed for one cluster and that cluster is in Service mode, then another cluster can be placed in Service Prep. A cluster in Service Prep automatically cancels Service Prep if its peer in the Grid experiences an unexpected outage while the Service Prep process is still active.

Note: Although you may place all clusters except one in Service mode, the best approach is if only one cluster is in Service mode at a time. If more than one cluster is in Service mode, and you cancel Service mode on one of them, then that cluster does not return to normal operation until Service mode is cancelled on *all* the clusters.

The following items are available when viewing the current operational mode of a cluster; also, you may set the cache's Removal Threshold in preparation for maintenance events in a TS7700 Virtualization Engine hybrid grid configuration. You may use the Service mode window to put the TS7700 Virtualization Engine into Service mode or back into Normal mode.

- ▶ **Lower Threshold:** For a TS7720 Virtualization Engine cluster in a hybrid grid, use this option to lower the required threshold at which logical volumes are removed from cache.

Note: The Lower Threshold button is visible only in a hybrid grid (one that contains both a TS7720 Virtualization Engine and TS7740 Virtualization Engine).

The Lower Threshold button is disabled if the selected cluster is the only TS7720 Virtualization Engine cluster (not connected to a physical library) in the Grid.

- ▶ **Current State:** Can be any of the following states:
 - **Normal:** The cluster is in a normal operation state; Service Prep may be initiated on this cluster.
 - **Service Prep:** The cluster is preparing to go into Service mode. The cluster is completing operations (that is, copies owed to other clusters, ownership transfers, and lengthy tasks such as inserts and token reconciliation) that require all clusters to be synchronized.
 - **Service:** The cluster is in Service mode. The cluster is normally taken offline in this mode for service actions or to activate new code levels.

Depending on what mode the cluster is in, a different action is presented by the button below the Cluster State display. You may use this button to place the TS7700 Virtualization Engine into Service mode or back into Normal mode:

- ▶ **Prepare for Service Mode:** This option puts the cluster into Service Prep mode and allows the cluster to finish all current operations. If allowed to finish Service Prep, the cluster enters Service mode. This option is only available while the cluster is in Normal mode. To cancel Service Prep mode, click the **Return to Normal Mode** button.
- ▶ **Return to Normal Mode:** Returns the cluster to Normal mode. This option is available if the cluster is in Service Prep or Service mode. A cluster in Service Prep mode or Service mode returns to Normal mode if the Return to Normal Mode button is selected.

You are prompted to confirm your decision to change the Cluster State. Click **OK** to change the Cluster State, or **Cancel** to abandon the change operation.

Ownership Takeover Mode window

Use this window (Figure 8-81) to enable or disable Ownership Takeover mode for a failed cluster in a TS7700 Virtualization Engine.

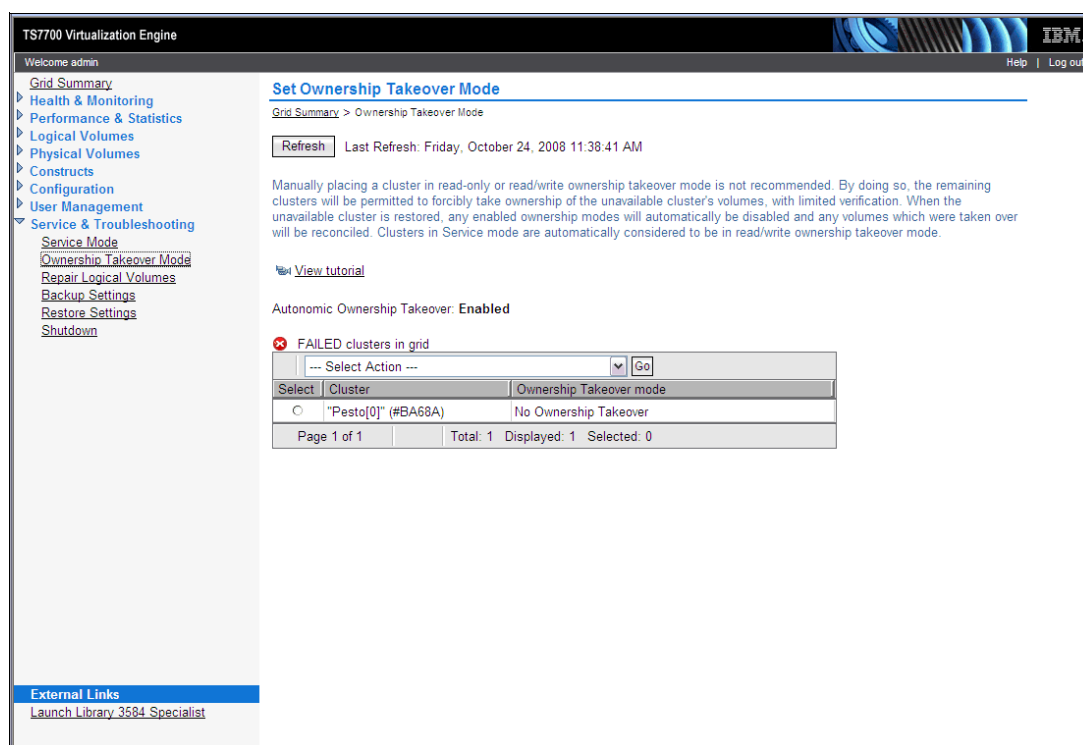


Figure 8-81 MI Set Ownership Takeover Mode window

For a cluster that is in a failed state, enabling Ownership Takeover mode allows other clusters in the Grid to obtain ownership of logical volumes that are owned by the failed cluster. Normally, ownership is transferred from one cluster to another through communication between the clusters. When a cluster fails or the communication path between clusters fails, the normal means of transferring ownership is not available. Enabling a read/write or read-only takeover mode must not be done if only the communication path between clusters has failed. A mode must only be enabled for a cluster that is no longer operational. The integrity of logical volumes in the Grid can be compromised if a takeover mode is enabled for a cluster that was not actually in a failed state.

The types of Ownership Takeover mode are as follows:

- **Read/Write Ownership Takeover mode:** Allows other non-failed clusters in the Grid to perform read and write operations and change to private or scratch status on logical volumes owned by the failed cluster. Ownership of the logical volume can only be obtained if there is a consistent copy of the logical volume available on the Grid or if the logical volume is in a scratch category. If there was no cluster failure, there is the possibility that two sites could be writing data to the same logical volume.
- **Read-Only Ownership Takeover Mode:** Allows other non-failed clusters in the Grid to perform read operations on the logical volumes that are owned by the failed cluster. The status of the volume, private/scratch, cannot be changed in this mode. If there was no cluster failure, there is the possibility a logical volume accessed by another cluster in Read-Only Ownership Takeover Mode contains data that is older than what is on the owning cluster. This can occur if the logical volume was modified on the owning cluster while the communication path between the clusters was down.

- ▶ **Autonomic Ownership Takeover Manager:** The Autonomic Ownership Takeover Manager (AOTM) can be either Enabled or Disabled. If AOTM is enabled, the TS7700 Virtualization Engine Grid will initiate an ownership takeover for a failed cluster. This will occur after a configured grace period and independently from user input. If the Autonomic Ownership Takeover Manager (AOTM) is installed and configured, it will perform an additional check to determine if the unavailable cluster is in a failed state when an ownership takeover is initiated. If AOTM determines that the cluster is not in a failed state, the ownership takeover attempt will fail.

The following Ownership Takeover mode information is available for failed clusters:

- ▶ **Cluster:** The failed cluster's name.
- ▶ **Ownership Takeover Mode.** Possible values are as follows:
 - **No Ownership Takeover:** The failed cluster is not in any Ownership Takeover mode.
 - **Read/Write Ownership Takeover:** Logical volumes owned by this cluster can be obtained for read and write operations.
 - **Read-only Ownership Takeover:** Logical volumes owned by this cluster are available for read and mount operations.

Figure 8-82 shows the three possible options you have in the Select Action drop-down menu.

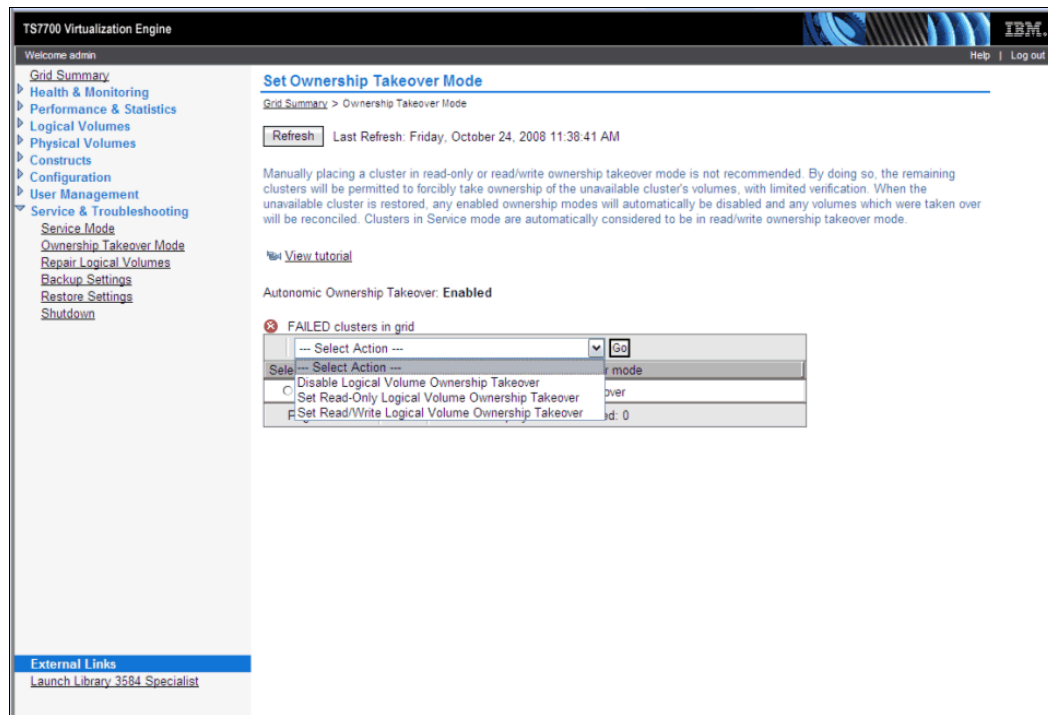


Figure 8-82 MI Set Ownership Takeover Mode window

For a Grid that does not have a cluster in a failed or offline state, you cannot enter a Ownership Takeover mode selection from any of the cluster’s MI connections, as shown in Figure 8-83.

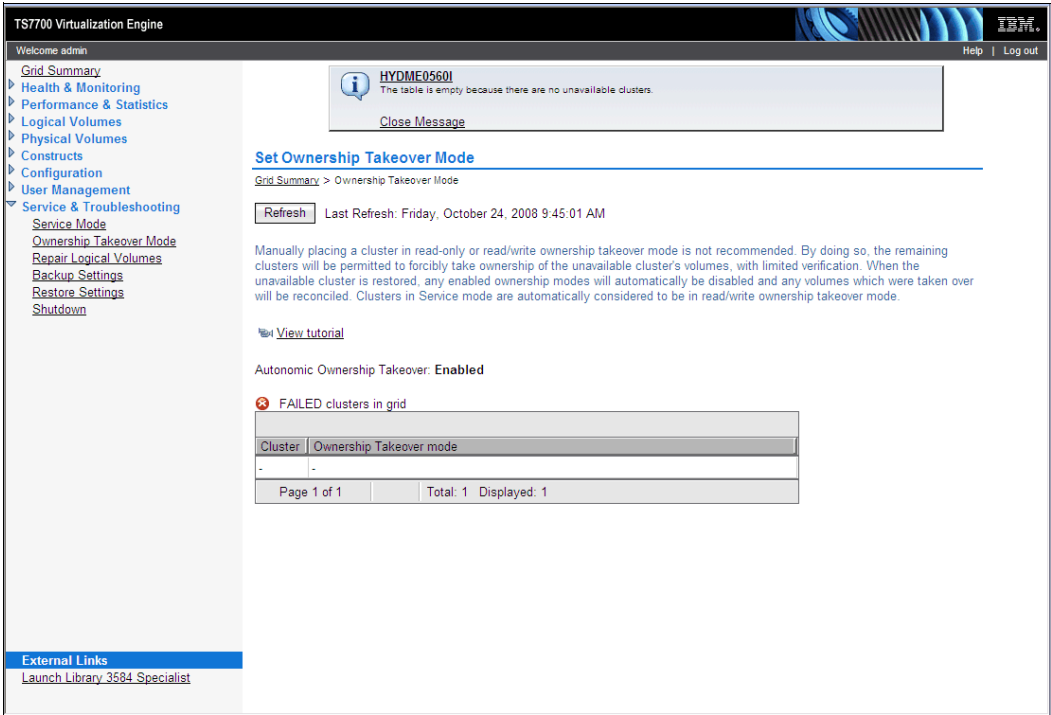


Figure 8-83 MI Ownership Takeover failed

The message HYDME0560I at the top of Figure 8-83 tells you that there are no unavailable clusters, as confirmed by the FAILED clusters in grid table in the middle of window.

Repair Logical Volumes window

Use this window (Figure 8-84 on page 526) to repair logical volumes in the damaged category for a TS7700 Virtualization Engine Grid.

To print the table data, click **Print Report** next to the Select Action menu. To download a comma-separated value (.csv) file of the table data, click **Download spreadsheet**.

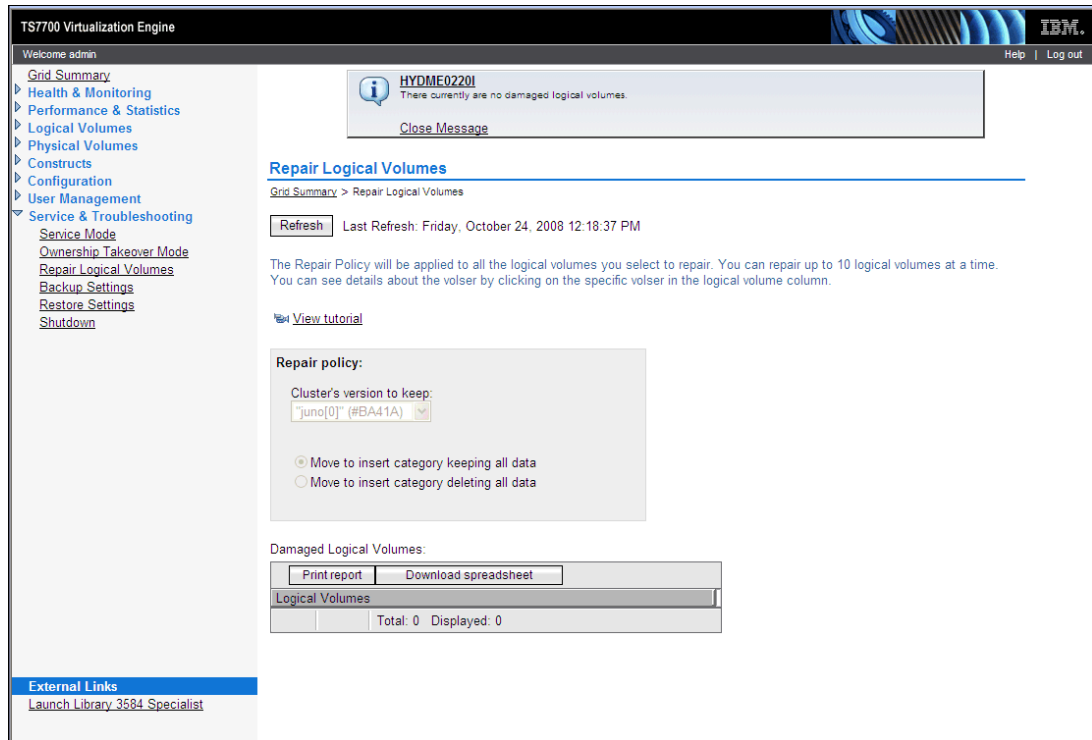


Figure 8-84 MI Repair Logical Volumes window

You have the following options:

- ▶ **Repair Policy:** The repair policy is applied to all damaged logical volumes selected. The options include:
 - **Cluster's version to keep:** The selected cluster will obtain ownership of the logical volume when the repair is complete. The cluster's version of the logical volume will be the basis for repair if the **Move to insert category keeping all data** option is selected.
 - **Move to insert category keeping all data:** Use this option if the data on the logical volume is intact and still relevant. If data loss has occurred, do not use this option. If the cluster chosen in the repair policy has no data for the logical volume to be repaired, choosing this option will be the same as choosing **Move to insert category deleting all data**.
 - **Move to insert category deleting all data:** The repaired logical volumes will be moved to the insert category and all data will be erased. Use this option if the volume has been returned to scratch or if data loss has rendered the volume obsolete. If the volume has been returned to scratch, then the data on the volume is no longer needed. If data loss has occurred on the volume, then data integrity issues might occur if the data on the volume is not erased.
- ▶ **Damaged Logical Volumes:** In this table, the damaged logical volumes are displayed. Multiple logical volumes can be selected by checking the check box in the Select column. The Repair Policy will be applied to all logical volumes selected.
- ▶ **Logical Volume:** The VOLSER of the damaged logical volume. If you click this field, which is also a hyperlink, the Damaged Logical Volumes details window opens. It provides more information.
- ▶ **Damaged logical volume details:** Use this window to view details about a damaged logical volume in a TS7700 Virtualization Engine.

The following information is presented in a table about the damaged logical volume and should assist in selecting a cluster for a repair action if data is to be retained. The information is presented for each cluster:

- ▶ **Cluster:** The cluster name.
- ▶ **Last Modified:** The date and time the logical volume was last modified.
- ▶ **Last Mounted:** The date and time the logical volume was last mounted.
- ▶ **Data Exists:** Possible values are as follows:
 - **Exists:** The data exists and is in the local cache or was migrated to physical tape.
 - **Does not exist:** The data for this logical volume does not exist on this cluster.
- ▶ **Size/MiBs:** The size of data on the logical volume in mibibytes.
- ▶ **Category:** Category attribute of the logical volume. This is used to denote a grouping of logical volumes.
- ▶ **Media Type:** The media type of the logical volume.
- ▶ **Ownership Takeover Time:** The date the last time an ownership takeover occurred for this logical volume.
- ▶ **Data Level:** Every time the logical volume is written to, this value increases. Inserted logical volumes start with a data level of 100. This factor is secondary to Insert Level when choosing a cluster for the repair policy.
- ▶ **Insert Level:** When a group of logical volumes are inserted, they are assigned an insert level. Later inserts are given a higher insert level. This factor is the most important when choosing a cluster in the repair policy if data is will be retained. A higher value means higher consistency for the logical volume's data.
- ▶ **Data Consistent:** If the cluster's logical volume copy's data level is considered the latest data level, then data is consistent.

Backup Settings window

Use this window (Figure 8-85) to back up the settings from a TS7700 Virtualization Engine Cluster.

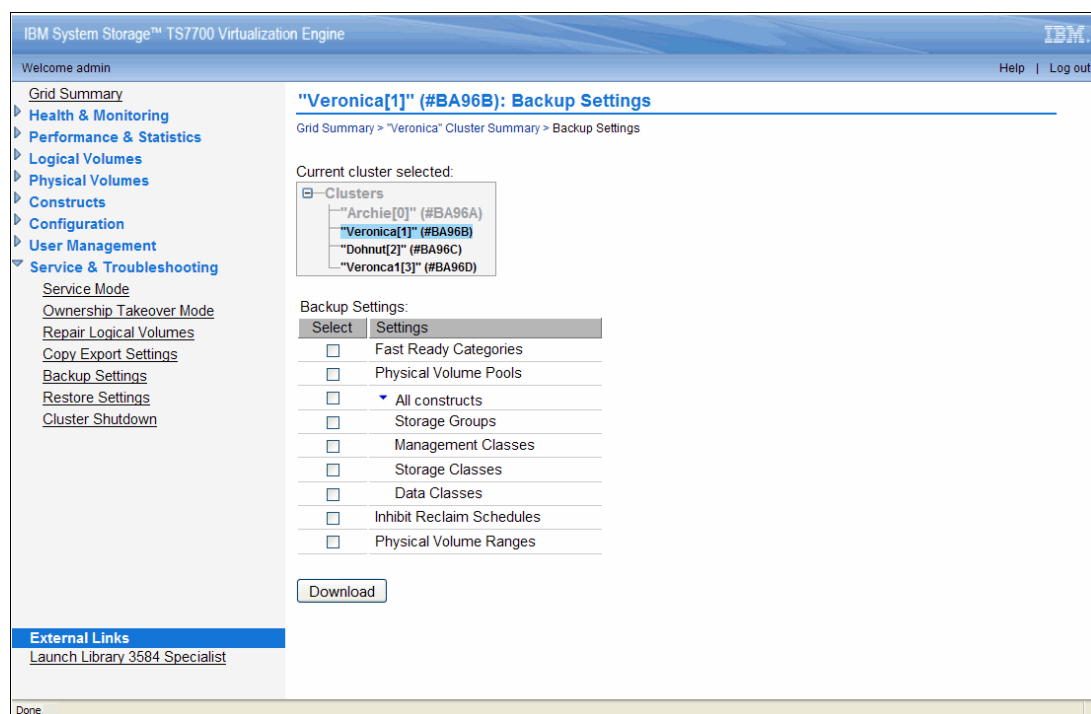


Figure 8-85 MI Backup Settings window

Figure 8-85 shows a list of cluster settings that are available for backup:

- **Fast Ready Categories:** Check the check box adjacent to this setting to back up Fast Ready Categories that are used to group logical volumes.
- **Physical Volume Pools:** Check the check box adjacent to this setting to back up physical volume pool definitions.

Note: If the cluster does not possess a physical library, physical volume pools will not be available.

- **All constructs:** To select all of the constructs for backup, select the check box adjacent to this setting.
- **Storage Groups:** Select the check box adjacent to this setting to back up defined storage groups.
- **Management Classes:** Check the check box adjacent to this setting to back up defined Management Classes.
- **Storage Classes:** Select the check box adjacent to this setting to back up defined storage classes.
- **Data Classes:** Select the check box adjacent to this setting to back up defined data classes.

- ▶ **Inhibit Reclaim Schedules:** Select the check box adjacent to this setting to back up Inhibit Reclaim Schedules used to postpone tape reclamation.
- ▶ **Physical Volume Ranges:** Select the check box adjacent to this setting to back up defined physical volume ranges.

Note: If the cluster does not possess a physical library, the option to Inhibit Reclaim Schedules will not be available.

To back up cluster settings, select a check box adjacent to any of the settings and then click the **Backup** button. The backup file `ts7700_cluster<cluster ID>.xml` is created. This file is an XML Meta Interchange file. You are prompted to open the backup file or save it to a directory. Save the file. Modify the file name before saving if you want to retain this backup file for subsequent backup operations.

Restore Settings window

Use this window (Figure 8-86) to restore the settings from a TS7700 Virtualization Engine Cluster to a recovered or new cluster.

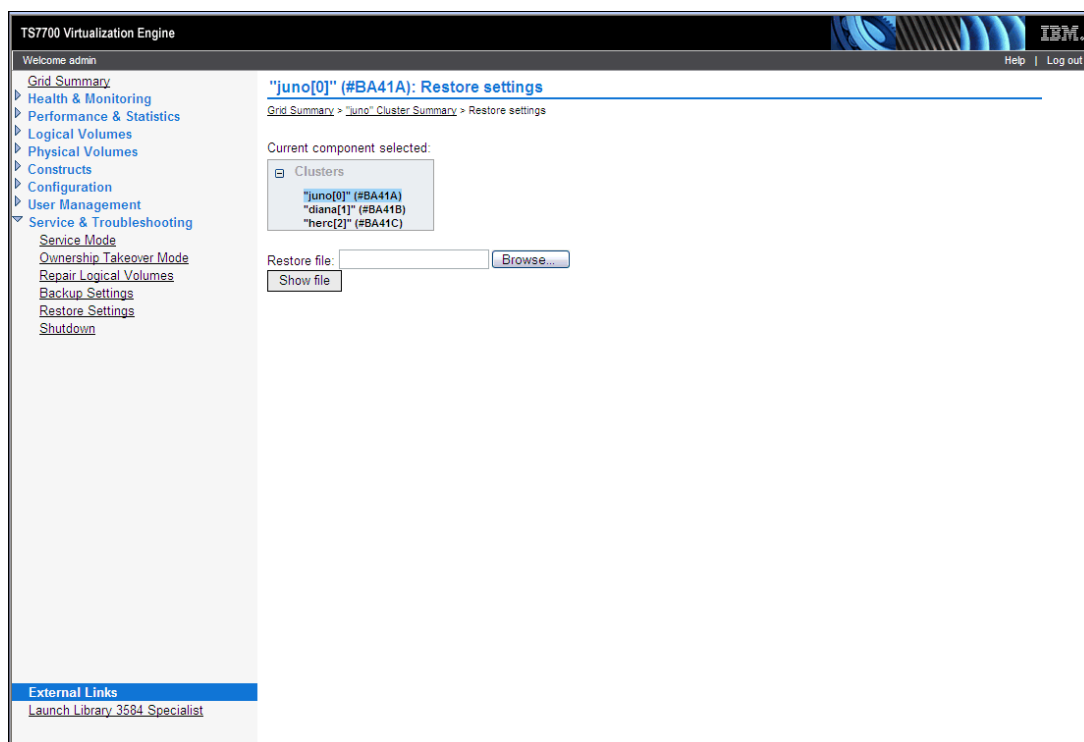


Figure 8-86 MI Restore Settings window

Click **Browse** to locate and upload the backup file you will use to restore cluster settings. Click **Show File** to review the cluster settings contained in the backup file. The backup file can contain any of the following settings, but only those settings that are defined by the backup file will be shown:

- ▶ **Constructs:** To restore all of the displayed constructs, check the check box adjacent to this setting.
- ▶ **Storage Groups:** Select the check box adjacent to this setting to restore defined storage groups.

- **Management Classes:** Select the check box adjacent to this setting to restore defined Management Classes.

Important: Management Class settings are related to the number and order of clusters in a Grid. Take special care when restoring this setting. If a Management Class is restored to a Grid having more clusters than the Grid had when the backup was performed, the copy policy for the new cluster or clusters are set as No Copy. If a Management Class is restored to a Grid having fewer clusters than the Grid had when the backup was performed, the copy policy for the now-nonexistent clusters are changed to No Copy. The copy policy for the first cluster will be changed to RUN to ensure one copy exists in the cluster.

- **Storage Classes:** Select the check box adjacent to this setting to restore defined storage classes.
- **Data Classes:** Select the check box adjacent to this setting to restore defined data classes.

Note: If this setting is selected and the cluster does not support logical WORM, the Logical WORM setting is disabled for all data classes on the cluster.

Logical WORM support is provided by TS7700 Virtualization Engine R1.6 and later.

- **Inhibit Reclaim Schedule:** Select the check box adjacent to this setting to restore Inhibit Reclaim Schedules used to postpone tape reclamation.

Note: A current Inhibit Reclaim Schedule is not overwritten by older settings. An earlier Inhibit Reclaim Schedule is not restored if it conflicts with an Inhibit Reclaim Schedule that currently exists.

If the backup file was created by a cluster that did not possess a physical library, the Inhibit Reclaim Schedules settings will be reset to their defaults.

- **Fast Ready Categories:** Select the check box adjacent to this setting to restore the Fast Ready Categories used to group logical volumes.
- **Physical Volume Pools:** Select the check box adjacent to this setting to restore physical volume pool definitions.

Note: If the backup file was created by a cluster that did not possess a physical library, Physical Volume Pool settings will be reset to their defaults.

- **Physical Volume Ranges:** Select the check box adjacent to this setting to back up defined physical volume ranges.

Note: If the backup file was created by a cluster that did not possess a physical library, the Physical Volume Range settings are reset to their defaults.

After clicking **Show File**, the name of the cluster from which the backup file was created is displayed at the top of the window, along with the date and time the backup occurred. Select the check box next to each setting to be restored and click **Restore**.

Note: The restore operation overwrites existing settings on the cluster.

A warning window opens and prompts you to confirm your decision to restore settings. Click **OK** to restore the settings or **Cancel** to cancel the restore operation.

The restore cluster settings operation can take five minutes or longer and its progress can be tracked in the Operation History window.

Restoring to or from a cluster without a physical library

If either the cluster that created the backup file or the cluster performing the restore operation does not possess a physical library, upon completion of the restore operation all physical tape library settings are reset to their defaults. One of the following warning messages is displayed in the confirmation window:

- ▶ The file was backed up from a system with a physical tape library attached but this cluster does not have a physical tape library attached. If you restore the file to this cluster, all the settings for physical tape library will have default values.
- ▶ The file was backed up from a cluster without a physical tape library attached but this cluster has a physical tape library attached. If you restore the file to this cluster, all the settings for physical tape library will have default values.

If a change in the cluster configuration is detected, the affected settings are as follows:

- ▶ Inhibit Reclaim Schedule
- ▶ Physical Volume Pools
- ▶ Physical Volume Ranges

Cluster Shutdown window

Use this window (Figure 8-87) to remotely shut down a TS7700 Virtualization Engine Cluster for a planned power outage or in an emergency.

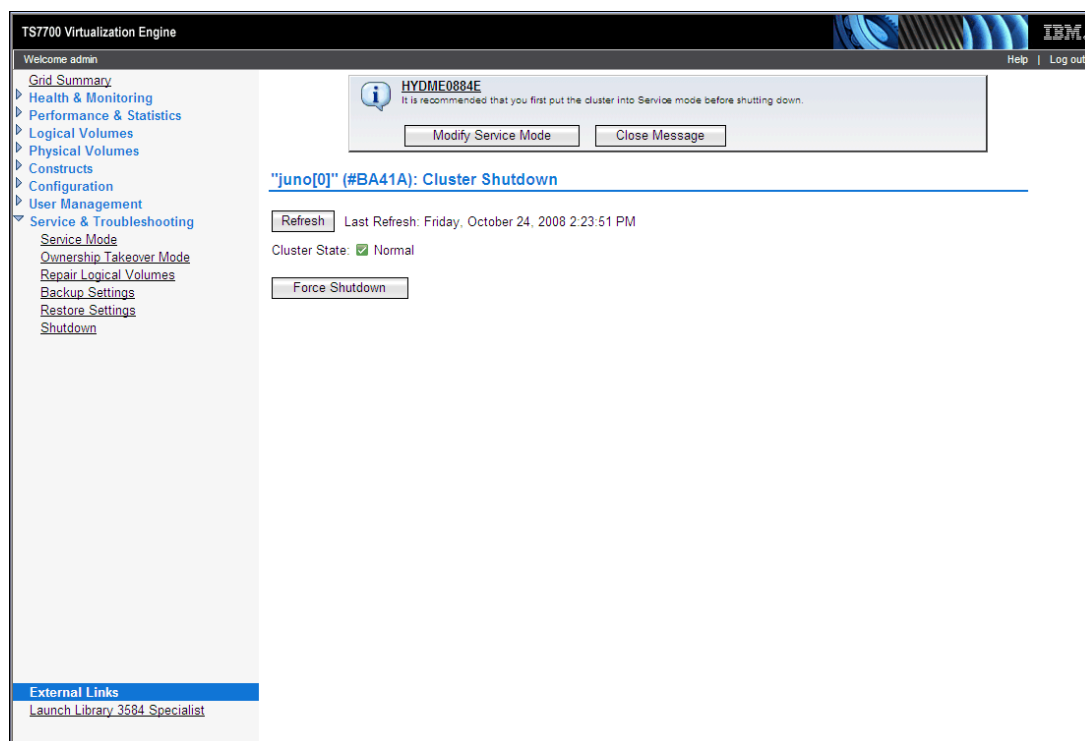


Figure 8-87 MI Cluster Shutdown window

This window is visible from the TS7700 Virtualization Engine management interface whether the TS7700 Virtualization Engine is online or offline.

Note: You can shut down only the cluster to which you are logged in. To shut down another cluster, you must log out of the current cluster and log into the cluster you want to shut down.

Before you shut down the TS7700 Virtualization Engine Cluster, you must decide whether your circumstances provide adequate time to perform a clean shutdown. A clean shutdown is not required, but is good practice to do for a TS7700 Virtualization Engine Grid configuration. A clean shutdown requires you to first place the cluster in service mode to ensure that no jobs are being processed during a shutdown operation. If you cannot place the cluster in service mode, you can use this window to force a shutdown of the cluster.

Attention: A forced shutdown can result in lost access to data and to job failure.

A cluster shutdown operation initiated from the TS7700 Virtualization Engine management interface also shuts down the cache. The cache must be restarted before any attempt is made to restart the TS7700 Virtualization Engine.

If you select **Shutdown** from the left side menu for a cluster that is still online, as shown at the top of the page in Figure 8-87, a message alerts you to first put the cluster in Service mode before shutting down.

In Figure 8-87 on page 532, the Cluster State field shows the operational status of the TS7700 Virtualization Engine and appears above the button used to force its shutdown. You have the following options;

- ▶ Cluster State. Possible values are as follows:
 - Normal: The cluster is in an online, operational state and is part of a TS7700 Virtualization Engine Grid.
 - Service: The cluster is in Service mode or is a stand-alone machine.
 - Offline: The cluster is offline. It might be shutting down in preparation for Service mode.
- ▶ Shutdown: This button initiates a shutdown operation.

Important: After a shutdown operation is initiated, it cannot be canceled.

- Clicking the Shutdown button in Normal mode: If you click **Shutdown** while in Normal mode, you receive a warning message recommending that you place the cluster in Service mode before proceeding. To place the cluster in Service mode, select **Modify Service Mode**. To continue with the force shutdown operation, select **Close Message**. If you opt to continue with the force shutdown operation, you are prompted to provide your password to proceed. Enter your password and select **Yes** to continue or select **No** to abandon the shutdown operation.
- Clicking the Shutdown button in Service mode: If you select the **Shutdown** button while in Service mode, you will be asked to confirm your decision. Click **OK** to continue or, click **Cancel** to abandon the shutdown operation.

When a shutdown operation is in progress, the Shutdown button is disabled and the status of the operation is displayed in an information message. The shutdown sequence is as follows:

1. Going offline
2. Shutting down
3. Powering off
4. Shutdown completes

Important: Verify that power to the TS7700 Virtualization Engine and to the cache is shut down before attempting to restart the system.

A cluster shutdown operation initiated from the TS7700 Virtualization Engine management interface also shuts down the cache. The cache must be restarted first and allowed to power up to a operational state before any attempt is made to restart the TS7700 Virtualization Engine.

Copy Export Recovery window

Use this window (Figure 8-88) to test a Copy Export recovery or perform an actual Copy Export recovery on a TS7700 Virtualization Engine.

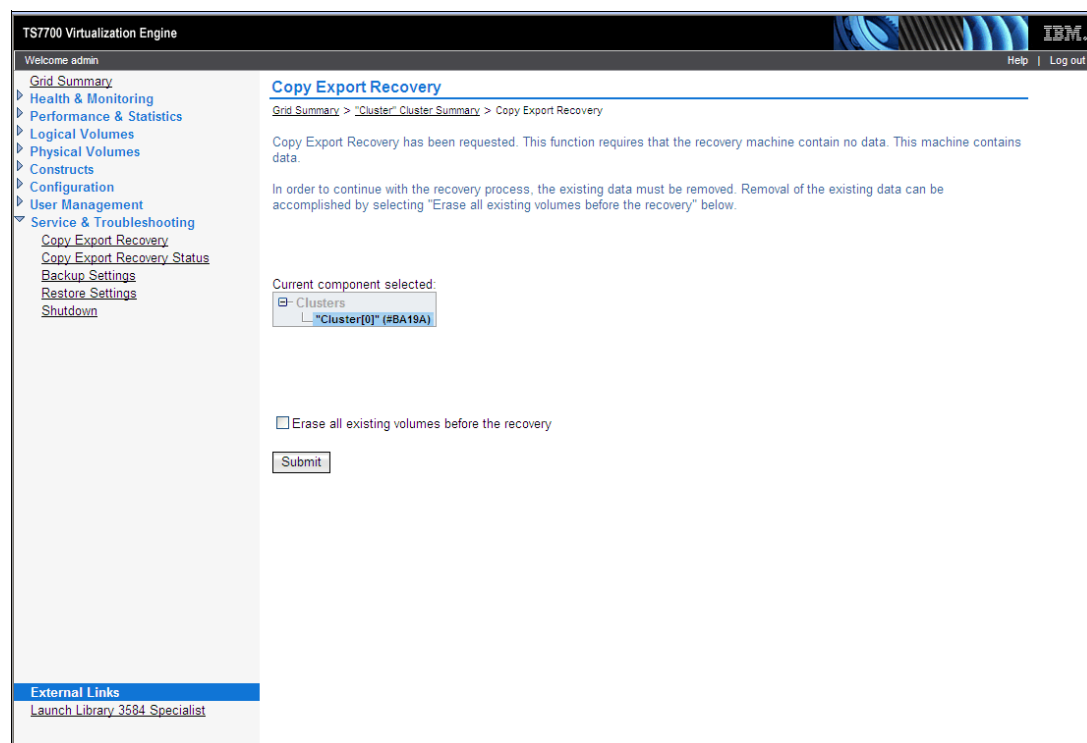


Figure 8-88 MI Copy Export Recovery window

If the Grid possesses a physical library, but the selected cluster does not, this page is visible but disabled on the TS7700 Virtualization Engine management interface, and the following message is displayed:

The cluster is not attached to a physical tape library.

If the Grid does not possess a physical library, this window is not visible on the TS7700 Virtualization Engine management interface.

Copy Export permits the export of all logical volumes and the logical volume database to physical volumes, which can then be ejected and saved as part of a data retention policy for disaster recovery. You can also use this function to test system recovery.

Note: Copy Export recovery is described in more detail in *IBM Virtualization Engine TS7700 Series Copy Export Function User's Guide*. This white paper and the most recently published white papers about this topic are available at the Techdocs website by searching for "TS7700" at the following address:

<http://www.ibm.com/support/techdocs/atmsastr.nsf/Web/Techdocs>

During the Copy Export recovery, all current logical volumes and physical volumes will be erased from the database and logical volumes are erased from the cache. Do *not* attempt this operation on a cluster where current data will be saved.

Before this operation is attempted, all physical media involved in the Copy Export recovery must be inserted. Two confirmation windows open after you select to erase all current data.

The second confirmation window requires reentry of the user's password before attempting the Copy Export recovery.

Important: Copy Export recovery can only target a Single Cluster configuration for TS7700 Virtualization Engine Grid volumes. The recovered cluster can be merged into an existing TS7700 Virtualization Engine Grid with all participating clusters at R1.7 or later.

If you create a new secondary copy, the original secondary copy is deleted because it becomes inactive data. For example, if you modify constructs for logical volumes that have already been exported and the logical volumes are remounted, a new secondary physical volume is created. The original physical volume copy is deleted without overwriting the logical volumes. When the copy export operation is rerun, the new, active version of the data is used. If this condition occurs, know that this is working as intended. In other words, the secondary copy, when recreated, expires the copy on the exported volume because it is no longer active data. This is how any volume works that has a secondary copy. When rewritten, the copy on the PVOL is no longer active and is no longer associated with the original physical volume.

The following fields and options are presented to the user to assist in testing a recovery or performing a recovery:

- ▶ **VOLSER of physical stacked volume for Recovery Test:** The physical volume from which the Copy Export recovery will attempt to recover the database.
- ▶ **Disaster Recovery Test Mode:** This option determines whether a Copy Export recovery will be run as a test or to recover a machine that has suffered a disaster. If this check box is selected (contains a check mark, which is the default status), the Copy Export recovery runs as a test. If the box is clear, the recovery process runs in "normal" mode, as when recovering from an actual disaster.

When the recovery is run as a test, the content of exported tapes remains unchanged. Additionally, primary physical copies remain unrestored and reclaim processing is disabled to halt any movement of data from the exported tapes. Any new volumes written to the machine are written to newly added scratch tapes and will not exist on the previously exported volumes. This ensures that the data on the Copy Export tapes remains unchanged during the test.

In contrast to a test recovery, a recovery performed in “normal” mode rewrites logical volumes to physical storage if the constructs change, so that the logical volume’s data can be put in the correct pools. With this type of recovery, reclaim processing remains enabled and primary physical copies are restored, requiring the addition of scratch physical volumes. A recovery run in this mode allows the data on the Copy Export tapes to expire in a normal manner and those physical volumes to be reclaimed.

- Erase all existing logical volumes during recovery: This check box is visible if logical volume or physical volume data is present in the database. A Copy Export Recovery operation erases any existing data. No option exists to retain existing data while performing the recovery. You must select this check box to proceed with the Copy Export Recovery operation.

You can verify the status of your Copy Export recovery by clicking **Copy Export Recovery Status**, as shown in Figure 8-89.

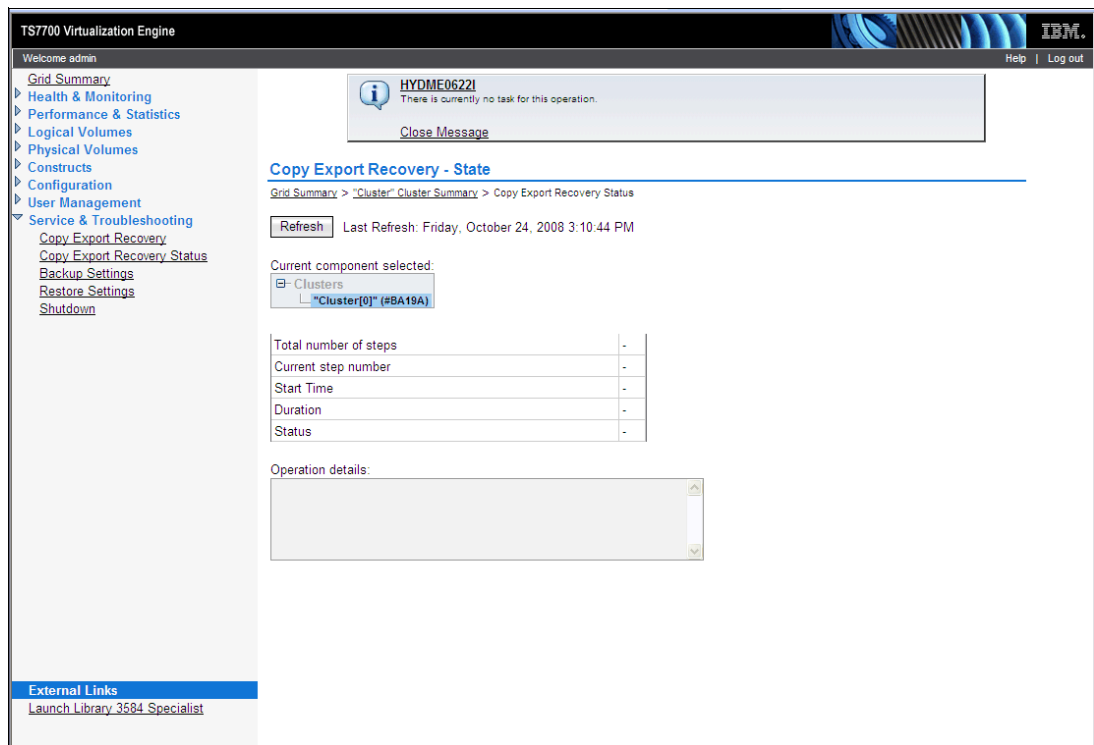


Figure 8-89 MI Copy Export Recovery - State window

Cluster nodes online/offline window

Note: The Cluster nodes online/offline window is visible only on a stand-alone cluster configuration and when the cluster is not online.

Use this window to verify the online state of a TS7700 Virtualization Engine Cluster and to skip a merge pending update operation if it is preventing an offline TS7700 Virtualization Engine Cluster from coming online.

This window displays information about a cluster's online status. Information or action messages displayed at the top of the page indicate the cluster's status.

A cluster can take a long time to come online, especially if a merge operation is pending. If a pending merge operation is preventing the cluster from coming online, you have the option to skip the merge step to reduce the time needed for the cluster to come online. Click **Skip Step** to skip the merge operation. This button is available only if the cluster is in a blocked state, waiting to share pending updates with one or more unavailable clusters.

Note: If you click the **Skip Step** button, pending updates against the local cluster might remain undetected until the unavailable cluster or clusters become available.

Click **Logout** to log off from the local cluster.

A cluster might be stuck in a pending online state because it is not able to communicate with its peers. While in a pending online state, the cluster is trying to discover if ownership takeovers have occurred. Without communicating with its peers, another cluster might have obtained ownership of a logical volume through an ownership takeover.

Important:

- ▶ If ownership takeover was set at any of the peers, the old data might possibly be sent the host if the cluster is forced online. Therefore, be sure that you know whether any peer clusters have ever enabled Ownership Takeover mode against this cluster while it was unavailable before attempting to force this cluster online. In addition, if this cluster is currently in service, automatic ownership takeover from unavailable peers is also likely and should be considered before attempting to force this cluster online.
- ▶ If Autonomic Ownership Takeover Manager (AOTM) is installed and configured, it attempts to determine if all unavailable peer clusters are actually in a failed state. If it determines that the unavailable cluster is not in a failed state, it blocks an attempt to force the cluster online. If the unavailable cluster is not actually in a failed state, the forced online cluster might be taking ownership of volumes that it should not be. If AOTM discovers that all unavailable peers have failed and network issues are not to blame, this cluster is then forced into an online state. After online, AOTM might further allow ownership takeover against the unavailable clusters if the AOTM option is enabled. Additionally, manual ownership takeover can be enabled if necessary.

8.2.11 Drive cleaning

The 3592 J1A, TS1120, and TS1130 Tape Drives need periodic cleaning and request cleaning from the TS3500 Tape Library. Automatic cleaning enables the TS3500 Tape Library to respond to any tape drive's request for cleaning and to begin the cleaning process without operator or TS7700 Virtualization Engine involvement. Automatic cleaning applies to all logical libraries that are configured in the TS3500 Tape Library. Automatic cleaning is required and cannot be disabled when the Advanced Library Management System (ALMS) is enabled. See *IBM System Storage TS3500 Tape Library with ALMS Operator Guide*, GA32-0594 for more information.

Important: ALMS is a requirement for IBM System z attachment. ALMS is always installed and enabled in a TS7700 Virtualization Engine z/OS environment; therefore, automatic cleaning is enabled.

Inserting cleaning cartridges into a TS3500 Tape Library

This section introduces two available methods to insert a cleaning cartridge into the TS3500 Tape Library. The process to insert cleaning cartridges varies depending on the setup of the TS3500 Tape Library. You can use the TS3500 Tape Library Specialist or the library's operator window to insert a cleaning cartridge.

Note: If virtual I/O slots are enabled, your library automatically imports cleaning cartridges.

Using the Tape Library Specialist Web interface to insert a cleaning cartridge

To use the Tape Library Specialist Web interface to insert a cleaning cartridge into the TS3500 Tape Library, perform the following steps:

1. Open the door of the I/O station and insert the cartridge so that the bar code label faces the interior of the library and the write-protect switch is on the right.
2. Close the door of the I/O station.
3. Type the Ethernet IP address on the URL line of the browser and press Enter. The System Summary window opens.
4. Select **Manage Cartridges** → **I/O Station**. The I/O Station window appears.
5. Follow the instructions in the window.

Using the operator panel to insert a cleaning cartridge

To use the operator panel to insert a cleaning cartridge into the 3584 Tape Library, perform the following steps:

1. From the library's Activity touchscreen, press **MENU** → **Manual Operations** → **Insert Cleaning Cartridge** → **ENTER**. The library displays the following message:

Insert Cleaning Cartridge into I/O station before you continue. Do you want to continue?

2. Open the door of the I/O station and insert the cartridge so that the bar code label faces the interior of the library and the write-protect switch is on the right.
3. Close the door of the I/O station.
4. Press **YES**. The following message is displayed:

Moving cleaning cartridge

This message is displayed while the library scans for one or more cleaning cartridges in the I/O stations:

- If one or more cleaning cartridges are present, the library moves the cleaning cartridges (one by one) to the lowest empty slots. If the library uses both LTO and 3592 tape cartridges, the accessor moves each cleaning cartridge to a storage location that contains similar media (using a separate move operation for each type of media). The library displays the following message:

Insertion of Cleaning Cartridges has completed.

- If no cleaning cartridges are in the I/O stations, the library displays the message "No cleaning cartridge found in the I/O station."
5. Press **ENTER** to return to the Manual Operations menu.
 6. Press **BACK** until you return to the Activity panel.

Removing cleaning cartridges from a TS3500 Tape Library

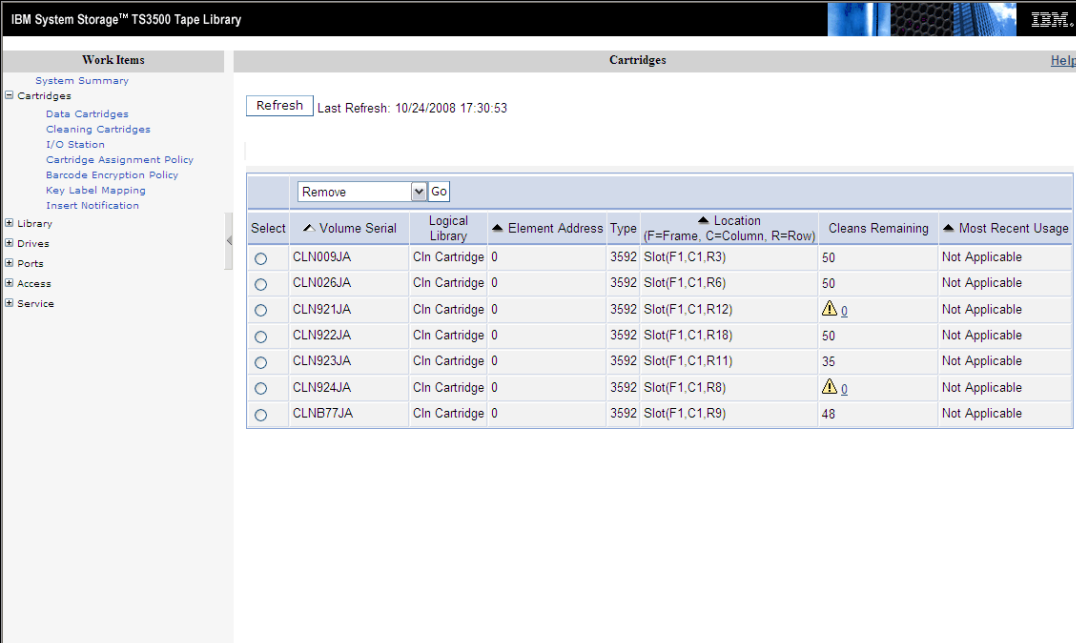
In this section, we describe how to remove a cleaning cartridge by using the TS3500 Tape Library Specialist. You can also use the operator panel. See *IBM System Storage TS3500 Tape Library with ALMS Operator Guide*, GA32-0594 for more information.

To use the TS3500 Tape Library Specialist Web interface to remove a cleaning cartridge from the tape library, perform the following steps:

1. Type the Ethernet IP address on the URL line of the browser and press Enter. The System Summary window opens.
2. Select **Cartridges** → **Cleaning Cartridges**. The Cleaning Cartridges window opens, as shown in Figure 8-90.
3. Select a cleaning cartridge, then from the Select Action drop-down list, select **Remove**, and then click **Go**.
4. Look at the Activity pane in the operator window to determine whether the I/O station that you want to use is locked or unlocked. If the station is locked, use your application software to unlock it.
5. Open the door of the I/O station and remove the cleaning cartridge.
6. Close the door of the I/O station.

Determine the cleaning cartridge usage in the TS3500 Tape Library

You can determine the usage of the cleaning cartridge in the same window that is used for the removal of the cleaning cartridges. See the Cleans Remaining column shown in Figure 8-90.



Select	Volume Serial	Logical Library	Element Address	Type	Location (F=Frame, C=Column, R=Row)	Cleans Remaining	Most Recent Usage
<input type="radio"/>	CLN009JA	Cln Cartridge 0		3592	Slot(F1,C1,R3)	50	Not Applicable
<input type="radio"/>	CLN026JA	Cln Cartridge 0		3592	Slot(F1,C1,R6)	50	Not Applicable
<input type="radio"/>	CLN921JA	Cln Cartridge 0		3592	Slot(F1,C1,R12)	⚠ 0	Not Applicable
<input type="radio"/>	CLN922JA	Cln Cartridge 0		3592	Slot(F1,C1,R18)	50	Not Applicable
<input type="radio"/>	CLN923JA	Cln Cartridge 0		3592	Slot(F1,C1,R11)	35	Not Applicable
<input type="radio"/>	CLN924JA	Cln Cartridge 0		3592	Slot(F1,C1,R8)	⚠ 0	Not Applicable
<input type="radio"/>	CLNB77JA	Cln Cartridge 0		3592	Slot(F1,C1,R9)	48	Not Applicable

Figure 8-90 TS3500 Tape Library cleaning cartridges

8.3 System-managed tape

This section describes the commands that are used to operate a tape library in a z/OS and system-managed tape environment. It is not intended to replace the full operational procedures in the product documentation. It is a quick reference for the needed DFSMS and MVS commands.

8.3.1 DFSMS operator commands

Some of the commands contain “libname” as a variable. In this case, the SMS-defined library name is required. Depending on whether you refer to a TS7700 Virtualization Engine composite library, TS7700 Virtualization Engine distributed library, or your native drives’ partition, the output will be slightly different for some of these commands. For more information about DFSMS commands, see *z/OS DFSMS Object Access Method Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC35-0427.

Information from the TS3500 Tape Library itself is contained in some of the outputs. However, you cannot switch the operational mode of the TS3500 Tape Library with z/OS commands.

Note: DFSMS and MVS commands apply only to SMS-defined libraries. The library name defined during the definition of a library in ISMF is required for “libname” in the DFSMS commands.

The following DFSMS operator commands support the tape library:

► **LIBRARY EJECT,volser{,PURGE|KEEP|LOCATION}{,BULK}**

This command is used to request the ejection of a volume from a tape library. The options available for this command are:

- Eject to the convenience I/O station for physical volumes (no additional specification). Delete from the tape library for logical volumes.
- Eject to the bulk output station (BULK or B) for physical volumes. Delete from the tape library for logical volumes.
- Remove the volume record from the TCDB (PURGE or P).
- Keep the volume record in the TCDB and update it to indicate that the cartridge has been ejected (KEEP or K). If the record contains information in the SHELF location field, it is not changed. If the SHELF location field is empty, the operator must enter information about the new location as a reply to WTOR. The reply can be up to 32 characters long.
- Keep the volume record in the TCDB and update it, including updating the SHELF location even if there is some information in this field (LOCATION or L). The operator has to enter the new information as a reply to WTOR.

If none of the variations (PURGE, KEEP, or LOCATION) are indicated in the command, a default decides whether the record is kept or purged. This default can be set separately for each library through the ISMF Library Definition window.

This command is available for the operator to eject single cartridges. Mass ejection of cartridges is usually performed through program interfaces such as ISMF, a tape management system, or a batch job.

► **LIBRARY SETCL, device-number, media-type**

This command allows the setting of the media type of the scratch volume that is to be loaded into the ICL of the specified tape drive. You must issue the command on the system on which the drive is online. The other hosts are notified when the drive is varied

online on the system. With the TS7700 Virtualization Engine, cartridge loaders are simulated and can be set to a media type. However, this approach can influence MVS allocation so care must be taken before using this command.

If the media assignment by this command differs from the current assignment, the ICL is emptied, and the proper cartridges are loaded.

► **VARY SMS,LIBRARY(libname),OFFLINE**

This command acts on the SMS library, which is referred by libname, that is, it stops tape library actions and gradually makes all of the tape units within this logical library unavailable. The units are varied offline “for library reasons”, which means that they are not accessible because the whole SMS-defined library is offline.

This simple form is a single-system form. The status of the library remains unaffected in other MVS systems.

Note: This command does not change the operational mode of the TS3500 Tape Library itself. It only applies to the SMS-defined logical libraries.

► **VARY SMS,LIBRARY(libname),ONLINE**

This command is required to bring the SMS-defined library back to operation after it has been offline.

The logical library does not necessarily go offline as a result of an error in some component of the physical library. Therefore, some messages for error situations request that the operator first vary the library offline and then back online. This usually clears all error indications and returns the library back into operation. Of course, this is only the MVS part of error recovery. You must clear the hardware, software, or operational error within the physical library and TS7700 Virtualization Engine before you bring the library back to work with MVS.

► **VARY SMS,LIBRARY(libname,sysname,...),ON/OFF and VARY SMS,LIBRARY(libname,ALL),ON/OFF**

This extended form of the VARY command can affect more than one system. The first form affects one or more named MVS systems. The second form performs the VARY action on all systems within the SMSplex.

The VARY SMS command allows the short forms ON and OFF as abbreviations for ONLINE and OFFLINE, respectively.

► **DISPLAY SMS,OAM**

This command gives a single line of information about all tape libraries (if present), their tape units, storage cells, and scratch cartridges.

This is the view of the single system where the command was executed. The number of unallocated, online drives is given under the heading AVL DRV (available drives).

If both optical libraries and tape libraries are defined in the SMS configuration, two multiline WTOs are displayed. The first multiline display produced by the library control system (LCS) is the display of optical library information. The second multiline display contains tape library information.

► **DISPLAY SMS,LIBRARY(libname|ALL),STATUS**

The library status display shows the SMS view of either one SMS-defined library or all SMS-defined libraries. The result contains one line of information for each library. This is a multihost view, which basically indicates whether the SMS-defined library is online, offline, or pending offline.

STATUS is the default parameter.

► **DISPLAY SMS,LIBRARY(ALL),DETAIL**

The DETAIL display, although a single-system view, gives slightly more information. The display is similar to the result of DISPLAY SMS,OAM, but each library gets its own line of information.

► **DISPLAY SMS,LIBRARY(libname),DETAIL**

This command provides details about the status of a single library. It is the only command that displays the library state (auto, pause, or manual mode). Reasons for the mode and indications of inoperative parts of the library are given in additional status lines. Examples of special situations are:

- Safety enclosure interlock open
- Vision system not operational
- Convenience output station full
- Out of cleaner volumes

► **DISPLAY SMS,STORGRP(grpname|ALL)**

There are no new parameters in the Storage Group display command because the optical library request formats are adequate here.

This display command is a general form of a request and gives the total SMS multihost view of the situation. The result is a display of the status of either all Storage Groups (DASD, optical, and tape) or a single Storage Group. There is no format to display one category only.

► **DISPLAY SMS,STORGRP(grpname|ALL),DETAIL**

The DETAIL display is not much more detailed than the general display. Only the library names of this Storage Group are indicated. This display is, in fact, more restricted than the general display. It gives the view of only one system, the view of its OAM, as the header line indicates.

The LISTVOL parameter of DISPLAY SMS,STORGRP is not used for tape Storage Groups. Although you can view a volume list through ISMF, a similar listing on the console is too long to be meaningful.

► **DISPLAY SMS,VOLUME(volser)**

This command displays all information that is stored about the volume in the TCDB (the VOLCAT) and some nonpermanent state information, such as “volume mounted on library-resident drive”.

8.3.2 Library LMPOLICY command

Use the LIBRARY LMPOLICY command to assign or change a volume's policy names outboard at the library. You can use this command only for private, library-resident volumes that reside in a library that supports outboard policy management.

The processing for the LIBRARY LMPOLICY command invokes the LCS external services FUNC=CUA function. Any errors that the CUA interface returns can also be returned for the LIBRARY LMPOLICY command. If the change use attribute installation exit (CBRUXCUA) is enabled, the CUA function calls the installation exit. This can override the policy names that you set using the LIBRARY LMPOLICY command.

The results of this command are specified in the text section of message CBR1086I. To verify the policy name settings and to see whether the CBRUXCUA installation exit changed the policy names you set, display the status of the volume.

The syntax of the LIBRARY LMPOLICY command to assign or change volume policy names is shown in Example 8-1.

Example 8-1 LIBRARY LMPOLICY command syntax

```
LIBRARY|LI LMPOLICY|LP , volser ,SG= storage group name |*RESET*  
                        ,SC= storage class name |*RESET*  
                        ,MC= Management Class name |*RESET*  
                        ,DC= data class name |*RESET*
```

The following parameters are required:

► LMPOLICY | LP

Specifies a request to set one or more of a private volume's policy names outboard in the library in which the volume resides. The library must support outboard policy management.

► Volser

Volser specifies the volume serial number of a private volume that resides in a library with outboard policy management support.

► You must specify *at least one of the following optional parameters*. These parameters can be specified in any order.

– SG={storage group name | *RESET*}

Specifies a construct name for the SG parameter. If the request is successful, the construct name becomes the Storage Group for the volume in the TCDB and the Storage Group policy name in the library. If you specify the *RESET* keyword, you are requesting that OAM set the volume's Storage Group name to blanks in the TCDB, and to the default Storage Group policy in the library, which are also blanks.

– SC={storage class name | *RESET*}

Specifies a construct name for the SC parameter. If the request is successful, the construct name becomes the Storage Class policy name for the volume in the library. If you specify the *RESET* keyword, you are requesting that OAM set the volume's Storage Class name to the default Storage Class policy in the library, which are blanks.

– MC={Management Class name | *RESET*}

Specifies a construct name for the MC parameter. If the request is successful, the construct name becomes the Management Class policy name for the volume in the library. If you specify the *RESET* keyword, you are requesting that OAM set the volume's Management Class name to the default Management Class policy in the library (blanks).

– DC={data class name | *RESET*}

Specifies a construct name for the DC parameter. If the request is successful, the construct name becomes the Data Class policy name for the volume in the library. If you specify the *RESET* keyword, you are requesting that OAM set the volume's Data Class name to the default Data Class policy in the library, which are blanks.

The values you specify for the SG, SC, MC, and DC policy names must meet the Storage Management Subsystem (SMS) naming convention standards:

- Alphanumeric and national characters only
- Name must begin with an alphabetic or national character (\$, *, @, #, or %)
- No leading or embedded blanks
- Eight characters or less

8.3.3 Host Console Request function

The Library Request host console command provides a simple way for an operator to determine the status of the TS7700, to obtain information about the resources of the TS7700 or to perform an operation in the TS7700. It could also be used with automation software to obtain and analyze operational information that could then be used to alert a storage administrator that there is something that should be examined further. Specify the following information for the command:

- ▶ A library name, which can be a composite or a distributed library
- ▶ It also allows one to four keywords, with each keyword being a maximum of eight characters.

The specified keywords are passed to the TS7700 identified by the library name to instruct it on what type of information is being requested or which operation is to be performed. Based on the operation requested through the command, the TS7700 then returns information to the host that will be displayed as a multiline write to operator (WTO) message.

A TS7700 library is made up of a composite and one or more distributed libraries. The composite library represents the logical view of the aggregate of the underlying physical aspects of the library. In essence it is virtualizing the underlying physical TS7700s so they look like one library. From a host job processing perspective, the virtual tape device addresses and logical volumes it uses are part of the composite library although the actual resources used is on one of the underlying distributed libraries. A distributed library represents the view of the resources owned and managed by a specific TS7700 and associated physical library. For a single TS7700, it has a composite library view and a distributed library view. In a Grid configuration, there is a composite view for the entire configuration and individual distributed library views for each of the TS7700s. Most of the host requests are related to the physical resources of a TS7700 and as such are directed to the distributed library name for a specific TS7700. Logical volumes, however, have a view that spans all of the distributed libraries and is directed to the composite library name.

The Library Request command for Host Console Request is supported in z/OS V1R6 and later. Refer to OAM APAR OA20065 and device services APARs OA20066, OA20067, and OA20313. A detailed description of the Host Console Request functions and responses is available in *IBM Virtualization Engine TS7700 Series z/OS Host Command Line Request User's Guide*, which is available at the Techdocs website (search for the term "TS7700"):

<http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs>

Command syntax for the Host Console Request function

The Host Console Request is also referred to as the Library Request command. The syntax of the command is shown in Example 8-2.

Example 8-2 Host Console Request function syntax

```
LIBRARYILI REQUESTIREQ,library_name,  
,keyword1l,keyword2l,keyword3l,keyword4l  
,L=alname lname-a
```

The following parameters are required:

REQUEST REQ	Specifies a request to obtain information from the TS7700 Virtualization Engine or to perform an outboard operation.
library_name	Specifies the library name associated with the TS7700 Virtualization Engine to which the request should be directed. The library name specified may be a composite or a distributed library and which library is applicable depends on the other keywords specified.
keyword1	Specifies which operation is to be performed at the TS7700 Virtualization Engine.

The following parameters are optional. The optional parameters depend on the first keyword specified. Based on the first keyword specified, zero or more of the additional keywords might be appropriate.

keyword2	Specifies additional information in support of the operation specified with the first keyword.
keyword3	Specifies additional information in support of the operation specified with the first keyword.
keyword4	Specifies additional information in support of the operation specified with the first keyword. Keyword4 is prepared for future use.
L={a name name-a}	Specifies where to display the results of the inquiry: the display area (L=a), the console name (L=name), or both the console name and the display area (L=name-a). The name parameter can be an alphanumeric character string.

Note the following information:

- ▶ If the request is specific to the composite library, then the composite library name must be specified.
- ▶ If the request is specific to a distributed library, then the distributed library name must be used.
- ▶ If a request for a distributed library is received on a virtual drive address on a TS7700 cluster of a separate distributed library, the request is routed to the appropriate cluster for handling and the response is routed back through the requesting device address.

Note: The specified v must be from one to eight characters in length and can consist of alphanumeric (A-Z and 0-9) and the national character set (\$* @#%). A keyword cannot contain any blanks. The only checking performed by the host is to verify that the specified keywords conform to the supported character set. The validity of the keywords themselves and the keywords in combination with each other is verified by the TS7700 Virtualization Engine processing the command.

Overview of the Host Console Request functions

Table 8-3 lists all the commands and a short description of each of them.

Table 8-3 Overview of Host Console Request functions

Keyword1	Keyword2	Keyword3	Keyword4	Description	Comp. Lib.	Dist. Lib.	TS7720 TS7740
CACHE				Requests information about the current state of the cache and the data managed within it associated with the specific distributed library.	N/A	Y	Both
COPYEXP	<i>volser</i>	RECLAIM		Requests that the specified physical volume that has been exported previously in a Copy Export operation be made eligible for priority reclaim.	N/A	Y	TS7740
COPYEXP	<i>volser</i>	DELETE		Requests that the specified physical volume that has been exported previously in a Copy Export operation be deleted from the TS7700 Virtualization Engine database. The volume must be empty.	N/A	Y	TS7740
GRIDCNTL	COPY	DISABLE		Requests that copy operations for the specified distributed library be disabled. Copies that are in progress are allowed to complete, but no new copies using the specified distributed library as the source or target are initiated.	N/A	Y	Both
GRIDCNTL	COPY	ENABLE		Requests that copy operations for the specified distributed library be enabled. Copy operations can again use the specified distributed library as the source or target for copies.	N/A	Y	Both
LVOL	<i>volser</i>			Requests information about a specific logical volume.	Y	N/A	Both
LVOL	<i>volser</i>	PREFER MIGRATE	PROMOTE INFO	Requests a change in the cache management for a logical volume.	N/A	Y	Both TS7740
		REMOVE			N/A	Y	
		REMOVE			N/A	Y	TS7720
		REMOVE			N/A	Y	TS7720

Keyword1	Keyword2	Keyword3	Keyword4	Description	Comp. Lib.	Dist. Lib.	TS7720 TS7740
PDRIVE				Requests information about the physical drives and their current usage associated with the specified distributed library.	N/A	Y	TS7740
POOLCNT	00-32			Requests information about the media types and counts, associated with a specified distributed library, for volume pools beginning with the value in keyword2.	N/A	Y	TS7740
PVOL	volser			Requests information about a specific physical volume.	N/A	Y	TS7740
RECALLQ	volser			Requests the content of the recall queue starting with the specified logical volume. Keyword2 could be blank.	N/A	Y	TS7740
RECALLQ	volser	PROMOTE		Requests that the specified logical volume be promoted to the top of the recall queue.	N/A	Y	TS7740
SETTING	ALERT, CACHE, THROTTLE	See "SETTING"	See "SETTING"	Request to current cache control and throttling values.	N/A	Y	See "SETTING"
STATUS	GRID			Requests information about the copy, reconcile, and ownership takeover status of the libraries in a Grid configuration.	Y	N/A	Both
STATUS	GRIDLINK			Requests information about the status and performance of the links between the TS7700 Virtualization Engines in the Grid configuration.	N/A	Y	Both

SETTING

The SETTING request provides information about many of the current workflow and management settings of the cluster specified in the request and the ability to modify the settings. It also allows alerts to be set for many of the resources managed by the cluster.

In response to the SETTING request, the cluster associated with the distributed library in the request will modify its settings based on the additional keywords specified. If no additional keywords are specified, the request will just return the current settings. See Example 8-12 on page 565 for an example of the data returned after the rest of the keyword descriptions.

Figure 8-91 shows the alert thresholds available for various resources that are managed by the cluster.

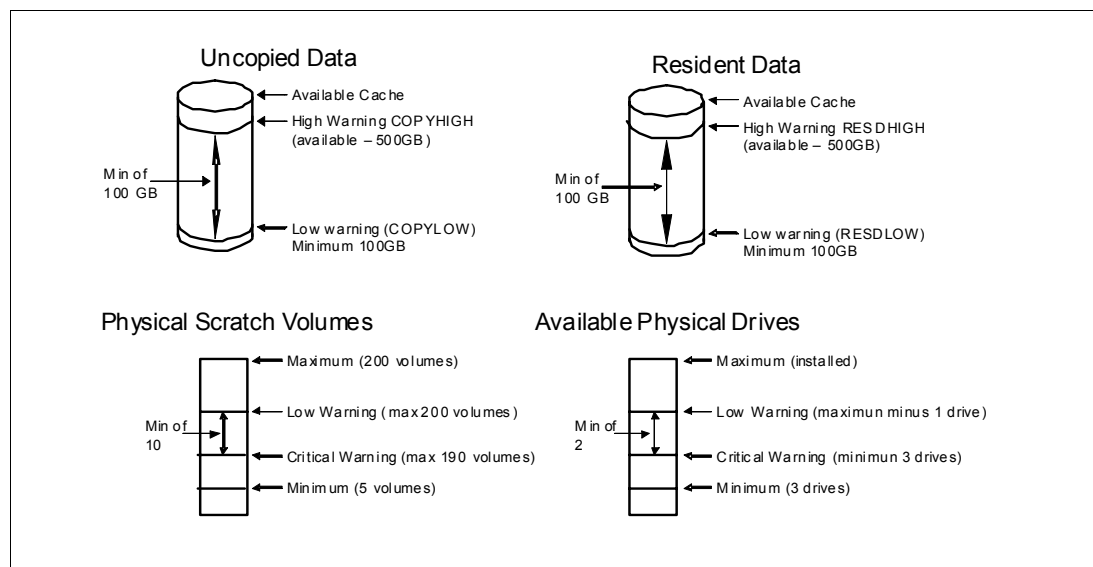


Figure 8-91 Setting Alert Threshold

When a value is specified, lead blanks or zeroes are ignored.

Note: All settings are persistent across machine restarts, service actions, or code updates. The settings are not carried forward as part of disaster recovery from Copy Exported tapes or the recovery of a system.

If a second keyword or alert is specified, the cluster will set different thresholds at which a message is sent to all hosts attached to the cluster and, in a Grid configuration, all hosts attached to all clusters. The third keyword specifies which alert threshold will be set and the fourth specifies the threshold value. All messages refer to the distributed library and will result in the following z/OS host console message:

CBR3750I Message from library distributed library name: Message Text

Thresholds can be set for many of the resources managed by the cluster. For each resource, two settings are provided. One warns that the resource is approaching a value that might result in an impact to the operations of the attached hosts. A second provides a warning that the resource has exceeded a value that might result in an impact to the operations of the attached hosts. When the second warning is reached, the warning message is repeated every 15 minutes.

The message text includes a message identifier that may be used to automate the capture and routing of these messages. To assist in routing messages to the appropriate individuals for handling, the messages that indicate that a resource is approaching an impact value will use message identifiers in a range of AL0000-AL4999. Message identifiers in a range of AL5000-AL9999 will be used for messages that indicate that a resource has exceeded an impact value.

Note: For messages where a variable is included (the setting), the value returned is left-justified without leading zeroes or right padding. For example:

AL5000 Uncopied data of 1450 GB above high warning limit of 1050 GB.

ALERT settings

Table 8-4 shows what ALERT thresholds are supported.

Table 8-4 ALERT threshold

Keyword 3	Keyword4	Description
COPYHIGH	value	<p>Uncopied Data High Warning Limit</p> <p>This is the threshold, in GBs of data in cache, that needs to be copied to other TS7700 Virtualization Engines in a Grid configuration, at which point the TS7700 Virtualization Engine will generate a message indicating that the amount of uncopied data has exceeded a high warning limit. For this message to be generated, the amount of uncopied data must be above the value specified in keyword4 for 150 seconds. As long as the amount of uncopied data remains above the threshold, the above warning limit message is repeated every 15 minutes. If the message had been generated and the amount of uncopied data is at or falls below that threshold for 150 seconds, a message is generated indicating that the amount of uncopied data is below the high warning limit.</p> <p>The default value is 0, which indicates no warning limit is set and no messages are to be generated.</p> <p>A value greater than the enabled cache size less 500 GB cannot be set. If a value is specified that is not 500 GB lower than the enabled cache size, the threshold is set to the enabled cache size less 500 GB.</p> <p>A value less than 100 GB cannot be set. If a value is specified that is not 100 GB or greater, the threshold is set to 100 GB.</p> <p>If the uncopied data low warning limit is not zero and a value is specified that is not 100 GB greater than the uncopied data low warning limit, the uncopied data low warning limit will be changed so that it is 100 GB less than the value specified (but never lower than 0).</p> <p>Message Text (xxxxxxx is the amount of uncopied data, yyyyyyy is the threshold):</p> <ul style="list-style-type: none">▶ When above the threshold: AL5000 Uncopied data of xxxxxxxx GB above high warning limit of yyyyyyy GB.▶ When below the threshold: AL5001 No longer above uncopied data high warning limit of yyyyyyy GB.

Keyword 3	Keyword4	Description
COPYLOW	value	<p>Uncopied Data Low Warning Limit</p> <p>This is the threshold, in GBs of data, in cache that needs to be copied to other TS7700 Virtualization Engines in a Grid configuration, at which the TS7700 Virtualization Engine will generate a message indicating that the amount of uncopied data has exceeded a low warning limit. For this message to be generated, the amount of uncopied data must be above the value specified in keyword 4 for 150 seconds. If the message had been generated and the amount of uncopied data is at or falls below that threshold for 150 seconds, a message is generated indicating that the amount of uncopied data is below the low warning limit.</p> <p>The default value is 0, which indicates no warning limit is set and no messages are to be generated.</p> <p>If the uncopied data high warning limit is set to 0, then the uncopied data low warning limit cannot be set to a value greater than the enabled cache size less 500 GB or to a value less than 100 GB. If a value is specified that is not 500 GB lower than the enabled cache size, the threshold is set to the enabled cache size less 500 GB. If a value is specified that is less than 100 GB, the threshold is set to 100 GB.</p> <p>If the uncopied data high warning limit is not zero, the uncopied data low warning limit cannot be set to a value greater than the uncopied data high warning limit less 100 GB. If a value is specified that is not 100 GB lower than the uncopied data high warning limit, the threshold is set to the uncopied data high warning limit less 100 GB (but never lower than 0).</p> <p>Message Text (xxxxxxx is the amount of uncopied data, yyyyyyy is the threshold):</p> <ul style="list-style-type: none"> ▶ When above the threshold: AL0000 Uncopied data of xxxxxxxx GB above low warning limit of yyyyyyy GB. ▶ When below the threshold: AL0001 No longer above uncopied data low warning limit of yyyyyyy GB.

Keyword 3	Keyword4	Description
PDRVCRIT	value	<p>Available Physical Drive Critical Warning Limit</p> <p>This is the threshold, in number of physical drives, at which the TS7700 Virtualization Engine will generate a message indicating that the number of available physical drives has fallen below the critical warning limit. For this message to be generated, the number of available physical drives must be below the value specified in keyword4 for 15 minutes. As long as the number of available physical drives is below the threshold, the below the threshold message is repeated every 15 minutes. If the message had been generated and the number of available physical drives is at or has risen above that threshold for 15 minutes, a message is generated indicating that the number of available physical drives is above the critical warning limit.</p> <p>The default value is 0, which indicates no warning limit is set and no messages are to be generated.</p> <p>A value greater than the number of installed physical drives minus 1 cannot be set. If a value is specified that is greater than the number of installed physical drives minus 1, the threshold is set to the number of installed physical drives minus 1. A value less than 3 cannot be set. If a value is specified that is less than 3, the threshold is set to 3.</p> <p>If the available physical drive low warning limit is not zero and a value is specified that is not 1 less than the available physical drive low warning limit, the available physical drive low warning limit will be changed so that it is 1 more than the value specified (but never more than the number of installed physical drives).</p> <p>Message Text (xx is the number of available drives, yy is the threshold):</p> <ul style="list-style-type: none"> ▶ When fallen below the threshold: AL5004 Available physical drives of xx is below critical limit of yy. ▶ When risen above the threshold: AL5005 Available physical drives no longer below critical limit of yy.

Keyword 3	Keyword4	Description
PDRVLOW	value	<p>Available Physical Drive Low Warning Limit</p> <p>This is the threshold, in number of physical drives, at which the TS7700 Virtualization Engine will generate a message indicating that the number of available physical drives has fallen below the low warning limit. For this message to be generated, the number of available physical drives must be below the value specified in keyword4 for 15 minutes. If the message had been generated and the number of available physical drives is at or has risen above that threshold for 15 minutes, a message is generated indicating that the number of available physical drives is above the low warning limit.</p> <p>The default value is 0, which indicates no warning limit is set and no messages will be generated.</p> <p>If the available physical drive critical warning limit is set to 0, then the available physical drive low warning limit cannot be set to a value greater than the number of installed physical drives or less than 3. If a value is specified that is greater than the number of installed physical drives, the threshold is set to the number of installed physical drives. If a value is specified that is less than 3, the threshold is set to 3.</p> <p>If the available physical drive critical warning limit is not zero and a value is specified that is not 1 greater than the available physical drive critical warning limit, the available physical drive low warning limit will be changed so that it is 1 more than the available physical drive critical warning limit (but never greater than the number of installed physical drives).</p> <p>Message Text (xx is the number of available drives, yy is the threshold):</p> <ul style="list-style-type: none"> ▶ When fallen below the threshold: AL0004 Available physical drives of xx is below low limit of yy. ▶ When risen above the threshold: AL0005 Available physical drives no longer below low limit of yy.

Keyword 3	Keyword4	Description
PSCRCRIT	value	<p>Physical Scratch Volume Critical Warning Limit This is the threshold, in number of scratch physical volumes, at which the TS7700 Virtualization Engine will generate a message indicating that the number of available scratch physical volumes has fallen below the critical warning limit. For this message to be generated, the number of available physical scratch volumes for one of the active general pools (pools 1-32) must be below the value specified in keyword4 for 16 minutes. Available volumes include those in pool 0 if borrowing is allowed for the pool. All media types allowed for borrowing are considered. As long as the number of scratch physical volumes is below the threshold, the below the threshold message is repeated every 16 minutes. If the message had been generated and the number of available physical scratch volumes for the pool is at or has risen above that threshold for 16 minutes, a message is generated indicating that the number of available physical scratch volumes is above the critical warning limit.</p> <p>The default value is 0, which indicates no warning limit is set and no messages will be generated.</p> <p>A value greater than 190 cannot be set. If a value is specified that is greater than 190, the threshold is set to 190. A value less than 5 cannot be set. If a value is specified that is less than 5, the threshold is set to 5. If the physical scratch volume low warning limit is not zero and a value is specified that is not 10 less than the physical scratch volume low warning limit, the physical scratch volume low warning limit will be changed so that it is 10 more than the value specified (but never greater than 200).</p> <p>Message Text (xxx is the number of available physical scratch volumes, yyy is the threshold, zz is the pool number):</p> <ul style="list-style-type: none"> ▶ When fallen below the threshold: AL5006 Available physical scratch volumes of xxx below critical limit of yyy for pool zz. ▶ When risen above the threshold: AL5007 Available physical scratch volumes no longer below critical limit of yyy for pool zz. <p>Note: The TS7700 Virtualization Engine will enter panic reclaim if the number of scratch volumes available to a defined pool is less than 2, including ones that it could borrow from pool 0.</p>

Keyword 3	Keyword4	Description
PSCRLOW	value	<p>Physical Scratch Volume Low Warning Limit This is the threshold, in number of scratch physical volumes, at which the TS7700 Virtualization Engine will generate a message indicating that the number of available scratch physical volumes has fallen below the low warning limit. For this message to be generated, the number of available physical scratch volumes for one of the active general pools (pools 1-32) must be below the value specified in keyword4 for 16 minutes. Available volumes include those in pool 0 if borrowing is allowed for the pool. All media types allowed for borrowing are considered. If the message had been generated and the number of available physical scratch volumes for the pool and media type is at or has risen above that threshold for 16 minutes, a message is generated indicating that the number of available physical scratch volumes is above the low warning limit.</p> <p>The default value is 0, which indicates no warning limit is set and no messages are to be generated.</p> <p>If the physical scratch volume critical warning limit is set to 0, then the physical scratch volume low warning limit cannot be set to a value greater than 200 or less than 5. If a value is specified that is greater than 200, the threshold is set to 200. If a value is specified that is less than 5, the threshold is set to 5.</p> <p>If the physical scratch volume critical warning limit is not zero and a value is specified that is not 10 greater than the physical scratch volume critical warning limit, the physical scratch volume low warning limit will be changed so that it is 10 more than the physical scratch volume critical warning limit (but never greater than 200).</p> <p>Message Text (xxx is the number of available physical scratch volumes, yyy is the threshold, zz is the pool number):</p> <ul style="list-style-type: none"> ▶ When fallen below the threshold: AL0006 Available physical scratch volumes of xxx below low limit of yyy for pool zz. ▶ When risen above the threshold: AL0007 Available physical scratch volumes no longer below low limit of yyy for pool zz. <p>Note: The TS7700 Virtualization Engine will enter panic reclaim if the number of scratch volumes available to a defined pool is less than 2, including ones that it could borrow from pool 0.</p>

Keyword 3	Keyword4	Description
REDHIGH	value	<p>Resident Data High Warning Limit</p> <p>This is the threshold, in gigabytes of resident data, at which the TS7700 Virtualization Engine will generate a message indicating that the amount of resident data has exceeded a high warning limit. For this message to be generated, the amount of resident data must be above the value specified in keyword 4 for 150 seconds. As long as the amount of resident data remains above the threshold, the above warning limit message is repeated every 15 minutes. If the message had been generated and the amount of resident data is at or falls below that threshold for 150 seconds, a message is generated indicating that the amount of resident data is below the high warning limit.</p> <p>The default value is 0, which indicates no warning limit is set and no messages are to be generated.</p> <p>A value greater than the enabled cache size less 500 GB cannot be set. If a value is specified that is not 500 GB lower than the enabled cache size, the threshold is set to the enabled cache size less 500 GB. A value less than 100 GB cannot be set. If a value is specified that is not 100 GB or greater, the threshold is set to 100 GB.</p> <p>If the resident data low warning limit is not zero and a value is specified that is not 100 GB greater than the resident data low warning limit, the resident data low warning limit will be changed so that it is 100 GB less than the value specified (but never lower than 0).</p> <p>Message Text (xxxxxxx is the amount of resident data, yyyyyyy is the threshold):</p> <ul style="list-style-type: none"> ▶ When above the threshold: AL5008 Resident data of xxxxxxxx GB above high warning limit of yyyyyyy GB. ▶ When below the threshold: AL5009 No longer above resident data high warning limit of yyyyyyy GB. <p>Note: For a TS7740 Virtualization Engine, resident data is the data that has not been copied to back-end physical volumes. For a TS7720 Virtualization Engine, resident data is the data that occupies the cache.</p>

Keyword 3	Keyword4	Description
RESLOW	Value	<p>Resident Data Low Warning Limit This is the threshold, in gigabytes of resident data, at which the TS7700 Virtualization Engine will generate a message indicating that the amount of resident data has exceeded a low warning limit. For this message to be generated, the amount of resident data must be above the value specified in keyword4 for 150 seconds. If the message had been generated and the amount of resident data is at or falls below that threshold for 150 seconds, a message is generated indicating that the amount of resident data is below the low warning limit.</p> <p>The default value is 0, which indicates no warning limit is set and no messages are to be generated.</p> <p>If the resident data high warning limit is set to 0, then the resident data low warning limit cannot be set to a value greater than the enabled cache size less 500 GB or to a value less than 100 GB. If a value is specified that is not 500 GB lower than the enabled cache size, the threshold is set to the enabled cache size less 500 GB. If a value is specified that is less than 100 GB, the threshold is set to 100 GB.</p> <p>If the resident data high warning limit is not zero, the resident data low warning limit cannot be set to a value greater than the resident data high warning limit less 100 GB. If a value is specified that is not 100 GB lower than the resident data high warning limit, the threshold is set to the resident data high warning limit less 100 GB (but never lower than 0).</p> <p>Message Text (xxxxxxx is the amount of resident data, yyyyyyy is the threshold):</p> <ul style="list-style-type: none"> ▶ When above the threshold: AL0008 Resident data of xxxxxxxx GB above low warning limit of yyyyyyyy GB. ▶ When below the threshold: AL0009 No longer above resident data low warning limit of yyyyyyyy GB. <p>Note: For a TS7740 Virtualization Engine, resident data is the data that has not been copied to back-end physical volumes. For a TS7720 Virtualization Engine, resident data is the data that occupies the cache</p>

CACHE settings

If second keyword of CACHE is specified, the cluster will modify how it controls the workflow and content of the tape volume cache.

The CACHE settings shown in Table 8-5 are supported.

Table 8-5 CACHE settings

Keyword 3	Keyword4	Description
COPYFSC	ENABLE, DISABLE	<p>Copies To Follow Storage Class Preference When the ENABLE keyword is specified, the logical volumes copied into the cache from a peer TS7700 Virtualization Engine are managed using the actions defined for the Storage Class construct associated with the volume as defined at the TS7700 Virtualization Engine receiving the copy.</p> <p>When the DISABLE keyword is specified, the logical volume copied into the cache from a peer TS7700 Virtualization Engine are managed as PG0 (prefer to be removed from cache).</p> <p>The default is disabled.</p>

Keyword 3	Keyword4	Description
PMPRIOR	value	<p>Premigration Priority Threshold This is the threshold, in gigabytes of unpemigrated data, at which the TS7700 Virtualization Engine will begin increasing the number of premigration tasks that will be allowed to compete with host I/O for cache and processor cycles. The amount of unpemigrated data must be above the value specified in keyword4 for 150 seconds before the additional premigration tasks are added. As the amount of data to premigrate continues to grow above this threshold setting, so do the number of enabled premigration tasks until the maximum is reached. The maximum is the number of available physical drives minus 1 (the TS7700 Virtualization Engine always keeps a minimum of one drive available for recalls). If the amount of unpemigrated data subsequently falls below this threshold for at least 150 seconds, the number of premigration tasks might be reduced depending on host I/O demand. If I/O host demand is high, the number of premigration tasks will eventually be reduced to a minimum of one. The default value is 800. If a value is specified that is higher than the premigration throttling threshold value, it is set to the premigration throttling threshold value. Note: Do not change this setting from the default unless you understand the impact that the change will have on the operation of the TS7700 Virtualization Engine. Raising the value might increase the length of time a peak write rate might be accepted, but also means that more data is solely resident in the cache and delays copying that data to physical tape.</p>
PMTHLVL	value	<p>Premigration Throttling Threshold This is the threshold, in gigabytes of unpemigrated data, at which the TS7700 Virtualization Engine will begin introducing a delay in responding to host write operations on all virtual tape device addresses of the TS7700 Virtualization Engine. The amount of unpemigrated data must be above the value specified in keyword4 for 150 seconds before the delay is introduced. The amount of delay begins with a few milliseconds and increases to 2000 milliseconds as the amount of unpemigrated data approaches the total cache capacity. The amount of unpemigrated data must fall below this threshold for 150 seconds for all delay to be removed. The default value is 1000. A value greater than the enabled cache size less 500 GB cannot be set. If a value is specified that is not 500 GB lower than the enabled cache size, the threshold is set to the enabled cache size less 500 GB. Note: Do not change this setting from the default unless you understand the impact that the change will have on the operation of the TS7700 Virtualization Engine. Raising the value might increase the length of time a peak write rate might be accepted, but also means that more data is solely resident in the cache and delays copying that data to physical tape.</p>
RECLPG0	ENABLE, DISABLE	<p>Recalls Preferred to be Removed from Cache When the ENABLE keyword is specified, logical volumes that are recalled into cache are managed as PG0 (prefer to be removed from cache). This overrides the actions defined for the Storage Class associated with the recalled volume. When the DISABLE keyword is specified, logical volume that are recalled into cache are managed using the actions defined for the Storage Class construct associated with the volume as defined at the TS7700 Virtualization Engine. The default is disabled.</p>

THROTTLE settings

If a second keyword of THROTTLE is specified, the cluster will modify how it controls the different data flow rates into and out to the cluster.

The THROTTLE settings shown in Table 8-6 are supported.

Table 8-6 THROTTLE settings

Keyword 3	Keyword4	Description
COPYFT	ENABLE, DISABLE	Full Cache Copy Throttling When the ENABLE keyword is specified, throttling when the cache is full of uncopied data to other TS7700 Virtualization Engines is enabled. When the DISABLE keyword is specified, throttling when the cache is full of uncopied data to other TS7700 Virtualization Engines is disabled. The default is enabled. Note: Full Cache Copy Throttling is also disabled for 24 hours after one of the other TS7700 Virtualization Engines has been in service mode. This is to prevent immediate host throttling when the TS7700 Virtualization Engine being service is returned to use.
DCOPYT	value	Deferred Copy Throttle This is a delay, in milliseconds, that the TS7700 Virtualization Engine will introduce for every 256 KB copy transfer for all deferred mode copy tasks in which it is the source. The delay is introduced when the TS7700 Virtualization Engine detects that it is in a period of high host write workload. A value of 0 to 250 milliseconds can be set. The default value is 125 milliseconds.
ICOPYT	ENABLE, DISABLE	Immediate Copy Throttling When the ENABLE keyword is specified, immediate copy throttling is enabled. When the DISABLE keyword is specified, immediate copy throttling is disabled. The default is enabled.

Examples of the Host Console Request functions

Let us examine some examples of the commands and the responses retrieved.

CACHE command

Example 8-3 shows the CACHE command.

Example 8-3 Library request command CACHE

```
LI REQ,BARR68A,CACHE
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR68A,CACHE.
CBR1280I LIBRARY BARR68A REQUEST. 149
KEYWORDS: CACHE
```

```
-----
TAPE VOLUME CACHE STATE V1
INSTALLED/ENABLED GBS 6000/ 2000
PARTITION  ALLOC  USED  PG0  PG1  PMIGR  COPY  PMT  CPYT
          0    2000  1750   47  1703    0    0    0    0
          1      0    0    0    0    0    0    0    0
          2      0    0    0    0    0    0    0    0
          3      0    0    0    0    0    0    0    0
          4      0    0    0    0    0    0    0    0
          5      0    0    0    0    0    0    0    0
          6      0    0    0    0    0    0    0    0
          7      0    0    0    0    0    0    0    0
```

The response shows the following information:

- ▶ Number of gigabytes (GB) installed and enabled
- ▶ One line per partition for future partition support
- ▶ Allocated GB and used GB in the cache and how they are split between PG0 and PG1
- ▶ Number of GB that must be copied to another cluster
- ▶ Premigration and copy throttling values in msec

GRIDCNTL command

Example 8-4 shows the GRIDCNTL DISABLE command.

Example 8-4 Library request command GRIDCNTL DISABLE

```
LI REQ,BARR68,GRIDCNTL,DISABLE
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR68,GRIDCNTL,DISABLE
CBR1280I LIBRARY BARR68 REQUEST. 509
KEYWORDS: GRIDCNTL,DISABLE
```

```
-----
GRID COPY CAPABILITIES V1
DISABLED FOR SOURCE AND TARGET COPIES
```

The response from GRIDCNTL DISABLE shows that copies have been stopped.

Example 8-5 shows the GRIDCNTL ENABLE command.

Example 8-5 Library request command GRIDCNTL ENABLE

```
LI REQ,BARR68,GRIDCNTL,ENABLE
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR68,GRIDCNTL,ENABLE
CBR1280I LIBRARY BARR68 REQUEST. 509
KEYWORDS: GRIDCNTL,ENABLE
```

```
-----
GRID COPY CAPABILITIES V1
  ENABLED FOR SOURCE AND TARGET COPIES
-----
```

The response from GRIDCNTL DISABLE Example 8-5 shows that copies have been restarted.

LVOL command

Example 8-6 shows the LVOL command.

Example 8-6 Library request command LVOL

```
LI REQ,BARR68,LVOL,693023
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR68,LVOL,693023.
CBR1280I LIBRARY BARR68 REQUEST. 509
KEYWORDS: LVOL,693023
```

```
-----
LOGICAL VOLUME INFORMATION V1
LOGICAL VOLUME:          693023
MEDIA TYPE:              CST
COMPRESSED SIZE (MB):    392
MAXIMUM VOLUME CAPACITY (MB): 0
CURRENT OWNER:           Pesto
MOUNTED LIBRARY:
MOUNTED VNODE:
MOUNTED DEVICE:
TVC LIBRARY:             Pesto
MOUNT STATE:
CACHE PREFERENCE:        PG1
CATEGORY:                300F
LAST MOUNTED (UTC):      2007-08-05 07:59:14
LAST MODIFIED (UTC):     2007-08-05 07:19:25
LAST MODIFIED VNODE:     00
LAST MODIFIED DEVICE:    1D
TOTAL REQUIRED COPIES:    3
KNOWN CONSISTENT COPIES: 3
IMMEDIATE-DEFERRED:      N
DELETE EXPIRED:          N
RECONCILIATION REQUIRED:  N
-----
```

LIBRARY	RQ	CACHE	PRI	PVOL	SEC	PVOL	COPY	ST	COPY	Q	COPY	CP
Pesto	N	N	JA6321	-----			CMPT		-			RUN
Squint	N	N	JA4238	-----			CMPT		-			RUN
Celeste	N	N	JB0450	-----			CMPT		-			RUN

The response from LVOL shows detailed information about the logical volume. The response shows the following information:

- ▶ Whether a logical volume is ECCST or CST and the size of the volume
- ▶ Number of copies and VOLSER of physical volumes where the logical volume resides
- ▶ Copy policy used

PVOL command

Example 8-7 shows the PVOL command.

Example 8-7 Library request command PVOL

```
LI REQ,BARR68A,PVOL,JA5313
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR68A,PVOL,JA5313.
CBR1280I LIBRARY BARR68A REQUEST. 225
KEYWORDS: PVOL,JA5313
```

```
-----
PHYSICAL VOLUME INFORMATION V1
  PHYSICAL VOLUME:  JA5313
  MEDIA TYPE:       JA
  DRIVE MODE:       E05
  FORMAT:           TS7700
  VOLUME STATE:     READ-WRITE
  CAPACITY STATE:   FILLING
  CURRENT POOL:     2
  MBYTES WRITTEN:   140718
  % ACTIVE DATA:   100.0
  LAST INSERTED:    2007-03-16 20:06:06
  WHEN EXPORTED:    N/A
  MOUNTS:           2
  LOGICAL VOLUMES:  374
  ENCRYPTED:         N
```

The response from PVOL shows detailed information of the logical volume. The response shows the following information:

- ▶ Media type, drive mode, format of the volume, and Volume State (read-write)
- ▶ Capacity in MB, valid data in percent, and number of logical volumes
- ▶ Shows whether the physical volume is exported and whether it is encrypted

POOLCNT command

Example 8-8 shows the POOLCNT command.

Example 8-8 Library request command PVOL

```
LI REQ,BARR68A,POOLCNT
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR68A,POOLCNT.
  CBRXLCS SCRATCH PROCESSING FAILED FOR: V66654 RC = 0004 RSN = 0004
CBR1280I LIBRARY BARR68A REQUEST. 441
KEYWORDS: POOLCNT
```

```
-----
PHYSICAL MEDIA COUNTS V1
  POOL      MEDIA  EMPTY  FILLING  FULL  ERASE  ROR    UNAVAIL  CXPT
    0        JA      1618
    1        JA        3      1      576    0      0        0        0
    2        JA        2      8      712    7      0        0
```

The response from POOLCNT shows detailed information about each physical volume pool. The response shows the following information:

- ▶ Shows detailed info from each pool
- ▶ Details about media type, and volumes that are empty, filling, or full
- ▶ Volumes eligible for erase
- ▶ Volumes that are in the Read-only state, unavailable, or in Copy Export state

RECALLQ command

Example 8-9 shows the RECALLQ command.

Example 8-9 Library request command RECALLQ

```
LI REQ,BARR68A,RECALLQ
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR68A,RECALLQ.
CBR1280I LIBRARY BARR68A REQUEST. 820
KEYWORDS: RECALLQ
```

```
RECALL QUEUE V1
POS  LVOL  PVOL1  PVOL2          TIME
IP  L00121 AB0456              175
IP  L99356 AA0350              201
SC  Y30458 AB0456              148
SC  L54019 AA0350              145
  1  L67304 AC0101              135
  2  T09356 AD5901  P00167      102
```

The response from RECALLQ on distributed library shows detailed information about the logical volume recall queue. The response shows the following information:

- ▶ The recall is in progress for volume L00121 and L99356.
- ▶ Volumes Y30458 and L54019 have a recall scheduled, which means a RECALLQ,volser,PROMOTE has been issued.
- ▶ Volume L67304 is in position 1 for recall and has been in the recall queue for 135 seconds.
- ▶ Volume T09365 spans from physical volume AD5901 to P00167.

STATUS command

Example 8-10 shows the STATUS command.

Example 8-10 Library request command STATUS,GRID

```
LI REQ,BARR68,STATUS,GRID
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR68,STATUS,GRID.
CBR1280I LIBRARY BARR68 REQUEST. 373
KEYWORDS: STATUS,GRID
```

GRID STATUS V1

COMPOSITE LIBRARY VIEW

LIBRARY	STATE	IMMED-DEFERRED		OWNERSHIP-T/O		RECONCILE
		NUM	MB	MODE	NUM	NUM
Pesto	ON	0	0	-	0	8
Squint	SVC	12	8713	SOT	24	12
Celeste	ON	0	0	-	0	0

DISTRIBUTED LIBRARY VIEW

LIBRARY	STATE	RUN-COPY-QUEUE		DEF-COPY-QUEUE		LINK STATE
		NUM	MB	NUM	MB	01234567
Pesto	ON	40	15691	0	0	-AA
Squint	UN	-	-	-	-	---
Celeste	ON	61	23929	0	0	AU-

The response from STATUS,GRID on Composite Library shows detailed information about the multi-cluster grid. The response shows the following information:

- ▶ As seen from the Composite Library, Squint is in service mode and there is a queue of data that needs to be copied. Squint is in “service ownership takeover” mode, meaning that the other two TS7700 Virtualization Engines must do recalls even though a logical volume resides in cache in Squint.
- ▶ As seen from the Distributed Library View, Squint is unknown, and Celeste has an unavailable link.

STATUS GRIDLINK command

Example 8-11 shows the STATUS GRIDLINK command.

Example 8-11 Library request command STATUS,GRIDLINK

```
LI REQ,BARR68,STATUS,GRIDLINK
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR68,STATUS,GRIDLINK.
CBR1280I LIBRARY BARR68 REQUEST. 373
KEYWORDS: STATUS,GRIDLINK
GRIDLINK STATUS V1
CAPTURE TIMESTAMP: 2008-02-08 12:45:32
LINK VIEW
  LINK NUM      CFG      NEG      READ  WRITE  TOTAL ERR  LINK STATE
                                MB/S    MB/S    MB/S
01234567
      0              1000  1000      87.2  102.4  189.6      0  -AA
      1              1000  1000      74.9  104.6  179.5      0  -AA
      2              0      0      0.0    0.0    0.0      0
      3              0      0      0.0    0.0    0.0      0
-----
LINK PATH LATENCY VIEW
  LIBRARY      LINK 0      LINK 1      LINK 2      LINK 3
          LATENCY IN MSEC
      TS001B          6          7          0          0
      TS001C         19         20          0          0
-----
CLUSTER VIEW
DATA PACKETS SENT:          103948956
DATA PACKETS RETRANSMITTED:  496782
PERCENT RETRANSMITTED:      0.4778
-----
LOCAL LINK IP ADDRESS
LINK 0 IP ADDR:  9.11.200.60
LINK 1 IP ADDR:  9.11.200.61
LINK 2 IP ADDR:
LINK 3 IP ADDR:
```

The response from STATUS,GRIDLINK on a Composite Library shows detailed information about the links between clusters. The response shows the following information:

- ▶ Detailed information about link throughput and their status.
- ▶ Latency seen by each cluster. This value is gathered every five minutes.
- ▶ Amount of data sent by this cluster.
- ▶ Percent of packets retransmitted. This means the amount of sent operations failed. In general, any value below 1.5% is acceptable.

SETTING command

Example 8-12 shows the SETTING command.

Example 8-12 Library request command SETTING

```
LI REQ,BARR70B,SETTING
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR70B,SETTING.
CBR1280I LIBRARY BARR70B REQUEST. 109
KEYWORDS: SETTING
```

```
-----
SETTING V1
ALERTS
COPYLOW  =          0 COPYHIGH =          0
LSCRLow  =          0 LSCRCRIT =          0
PDRVLOW  =          0 PDRVCRIT =          0
PSCRLow  =          0 PSCRCRIT =          0
RESDLow  =          0 RESDHIGH =          0
-----
```

```
-----
CACHE CONTROLS
COPYFSC  = DISABLED
RECLPGO  = DISABLED
PMPRIOR  =          800 PMTHLVL =          1000
-----
```

```
-----
THROTTLE CONTROLS
COPYFT   = ENABLED
ICOPYT   = ENABLED
DCOPYT   =          125
-----
```

The response from SETTING on distributed library shows detailed information about the ALERTS, CACHE, and THROTTLE controls.

Many of these commands could be subject to automation based on your own automation products. You could create your own actions to be taken by periodically issuing the Library Request commands to react to responses automatically and without operator interference. This could be used for proactive handling.

8.3.4 MVS System commands

The following commands are described in detail in *z/OS MVS System Commands*, SA22-7627.

► DS QT,devnum,1,RDC

This command displays identification, status, and diagnostic information about tape devices. You can use the command to display the LIBRARY-ID and the LIBPORT-ID that are stored for a device in a TS7700 Virtualization Engine.

Example 8-13 shows the sample output of a DS QT system command.

Example 8-13 Sample output of a DS QT system command

```
DS QT,1699,1,RDC
IEE459I 12.30.05 DEVSERV QTAPE 970
UNIT DTYPE DSTATUS CUTYPE DEVTYPE CU-SERIAL DEV-SERIAL ACL LIBID
1699 3490L ON-NRD 3490A20 3490B40 0177-10619 0177-10619 I 10007
  READ DEVICE CHARACTERISTIC
3490203490400000 1FF8808000000000 0000000000000000 0000000000000000
0100070100000000 4281000000000000 0000000000000000 0000000000000000
-----
      |  --
      |  |----->          4. Byte = LIBPORT-ID
      |  |----->  1.-3. Byte = LIBRARY-ID (omit first half byte)
LIBRARY-ID=10007
LIBPORT-ID=01
```

► DS QT,devnum,MED,nnn

This command displays information about the device type, media type, and the cartridge volume serial number. *devnum* is the device address in hexadecimal. *nnn* is the number of devices to query.

Example 8-14 shows the sample output of a DS QT system command.

Example 8-14 Sample output of a DS QT system command

```
IEE459I 11.32.31 DEVSERV QMEDIUM 608
UNIT RDTYPE EDTYPE EXVLSR INVLSR RMEDIA EMEDIA
0940 3590-E 3590-1 003700          3
  UNIT, the device address
RDTYPE, the real device type (physical)
EDTYPE, emulated device type
EXVLSR, external volume label
INVLSR, internal volume label
RMEDIA, real media type
EMEDIA, emulated media type
```

► VARY unit,ONLINE/OFFLINE

The VARY unit command in itself is no different from what it used to be. However, new situations are seen when the affected unit is attached to a library.

When the library is offline, the tape units cannot be used. This is internally indicated in a new status (offline for library reasons), which is separate from the normal unit offline status. A unit can be offline for both library and single-unit reasons.

A unit that is offline for library reasons only cannot be taken online by running VARY unit,ONLINE. Only VARY SMS,LIBRARY(...),ONLINE helps.

You can bring a unit online that was individually varied offline and was offline for library reasons by varying it online individually and varying its library online. The order of these activities is not important, but both are necessary.

Currently, no display directly gives the reason why the unit is offline, and there is no display that gives the name of the library to which this unit belongs.

► **DISPLAY U**

The DISPLAY U command displays the status of the requested unit. If the unit is part of a tape library (either manual or automated), device type 348X is replaced by 348L. An IBM 3490E is shown as 349L, and a 3590 or 3592 as 359L.

For a manual tape library, this might create a situation where it is no longer possible to see from the console response whether a particular tape unit supports IDRC, because this information is overlaid by the L, which indicates that the unit belongs to a library.

The output of DEVSERV is not changed in this way.

► **MOUNT devnum, VOL=(NL/SL/AL,serial)**

The processing of MOUNT has been modified to accommodate automated tape libraries and the requirement to verify that the correct volume has been mounted.

► **UNLOAD devnum**

The UNLOAD command allows you to unload a drive, if the rewind/unload process was not successful in the first place.

8.4 Basic operations

In the next few sections, we explain some of the basic tasks that might be needed during the operation of a TS7700 Virtualization Engine.

8.4.1 Clock and time setting

The TS7700 Virtualization Engine time can be set from a Network Time Protocol (NTP) server or by the SSR. It is set to UTC, also called GMT. See “Time coordination” on page 32 for more details about time coordination.

The TS3500 Tape Library time can be set from specialist work items by selecting **Library** → **Date and Time**, as shown in Figure 8-92.

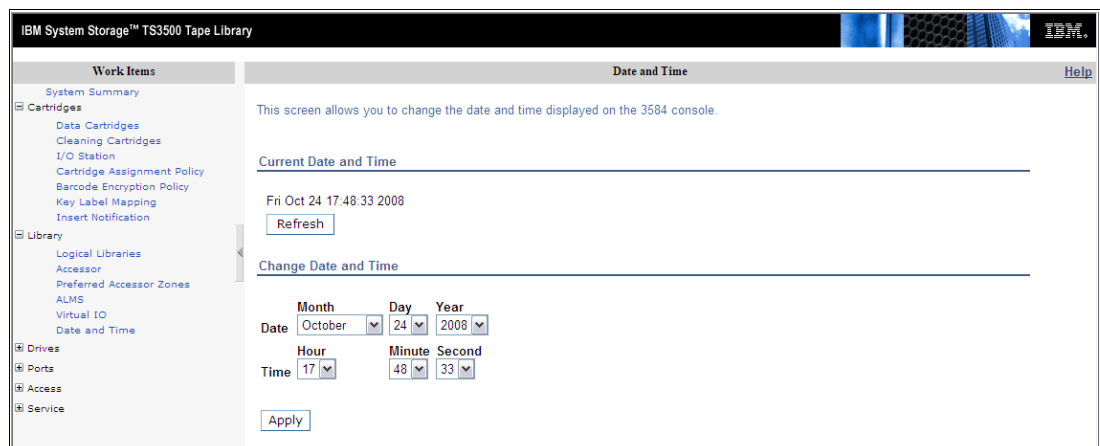


Figure 8-92 TS3500 Tape Library Specialist Date and Time

8.4.2 Library online/offline to host

The vary online/offline command for a TS7700 Virtualization Engine always uses the Composite Library ID either in a Single Cluster or in a multi-cluster grid environment.

8.4.3 Library in Pause mode

In a multi-cluster grid environment, one distributed library can enter the Pause mode, just as it is possible for a stand-alone VTS. The reasons for the pause can include an enclosure door being opened for clearing a device after a load/unload failure or removing cartridges from the high capacity I/O station. The following message is displayed at the host when a library is in Pause or manual mode:

```
CBR3757E Library library-name in {paused | manual mode} operational state
```

During Pause mode, all recalls and physical mounts are held up and queued by the TS7700 Virtualization Engine for later processing when the library leaves the Pause mode.

Because both scratch mounts and private mounts with data in the cache are allowed to execute, but not physical mounts, no more data can be moved out of the cache after the currently mounted stacked volumes are completely filled. The cache is filling up with data that has not been copied to stacked volumes. This results in significant throttling and finally in the stopping of any mount activity in the library. For this reason, it is important to minimize the amount of time spent with the library in Pause mode.

8.4.4 Preparing for service

When an element of the TS7700 Virtualization Engine needs to be serviced, it must be prepared prior to taking it offline; otherwise, continued host access to data might not be possible. The service preparation task is an administrative responsibility, and will remove the TS7700 Virtualization Engine Cluster from Grid participation. More details about service preparation can be found in Chapter 2, “Architecture, components, and functional characteristics” on page 15.

Note: Before invoking service preparation at the TS7700 Virtualization Engine, all virtual devices must be varied offline from the host. All logical volumes must be dismounted, all devices associated with the cluster varied offline, and all jobs moved to other clusters in the Grid before service preparation is invoked. After service is complete, and when the TS7700 Virtualization Engine is ready for operation, you must vary the devices online at the host.

Preparing a TS7700 Virtualization Engine for service

When an operational TS7700 Virtualization Engine needs to be taken offline for service, the TS7700 Virtualization Engine Grid must first be prepared for the loss of the resources involved, to provide continued access to customer data. The controls to prepare a TS7700 Virtualization Engine for service (Service Prep) are provided through the MI. This menu is described in Figure 8-80 on page 521.

Here is the message posted to all hosts when the TS7700 Virtualization Engine Grid is in this state:

```
CBR3788E Service preparation occurring in library library-name.
```

Preparing the tape library for service

If the TS3500 Tape Library in a TS7700 Virtualization Engine Grid must be serviced, the TS7740 Virtualization Engine associated with it must first be prepared for service. After the TS7740 Virtualization Engine has completed service preparation, the normal procedures for servicing the tape library can continue. See 8.2.10, “Service and troubleshooting” on page 520 for information about how to set the TS7700 Virtualization Engine in service preparation mode.

8.4.5 TS3500 Tape Library inventory

Use this window (Figure 8-93) from the TS3500 Tape Library specialist to perform Inventory/Audit.

You can choose **All Frames** or a selected frame from the drop-down menu.

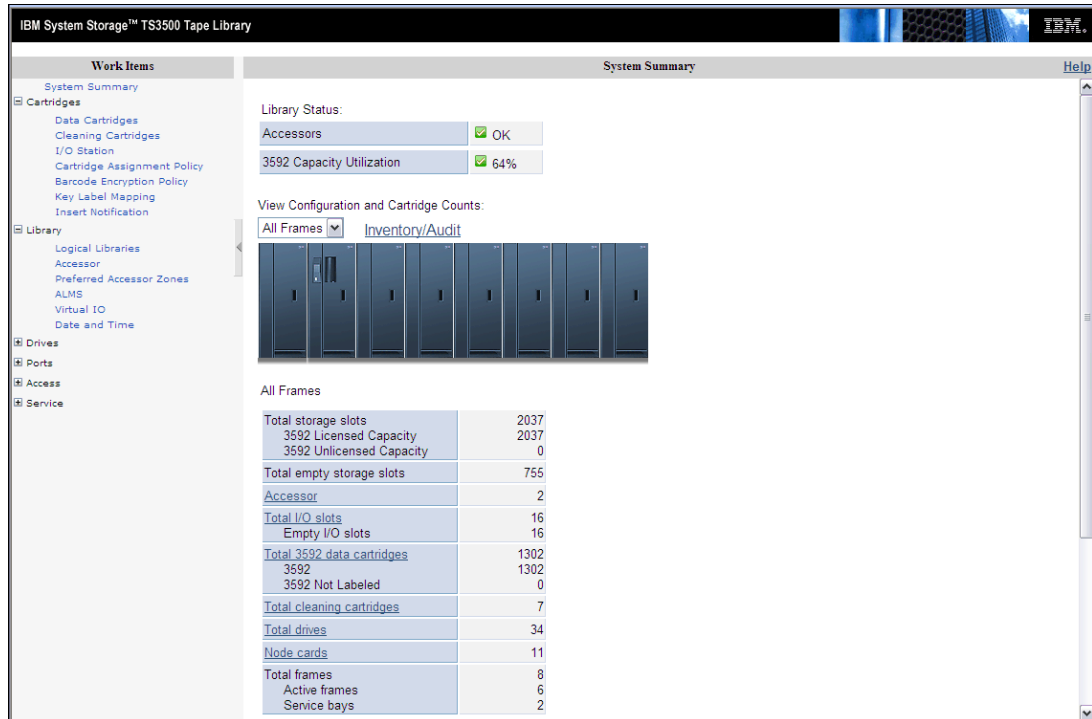


Figure 8-93 TS3500 Tape Library Inventory

After clicking the **Inventory/Audit** tab, you will receive the message shown in Figure 8-94.

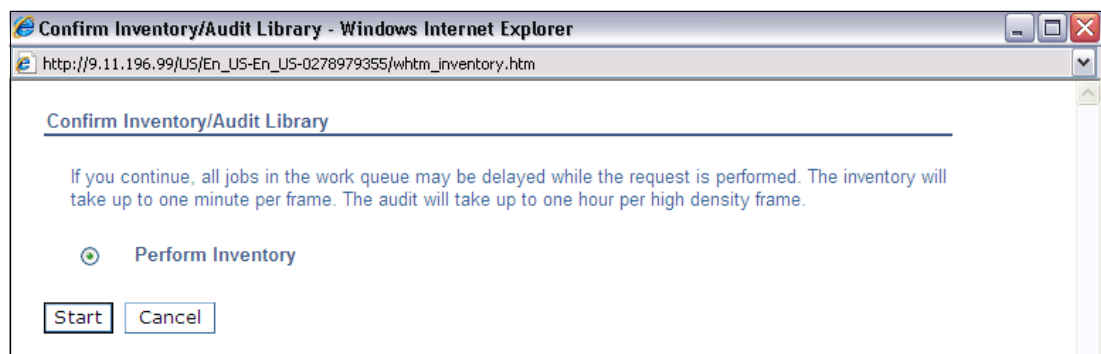


Figure 8-94 TS3500 Tape Library Inventory message

Important: As stated on the confirmation window (Figure 8-94), if you continue, all jobs in the work queue might be delayed while the request is performed. The inventory will take up to one minute per frame. The audit will take up to one hour per high density frame.

8.4.6 Inventory upload

See Figure 8-40 on page 471 for information about an inventory upload.

Click the **Inventory Upload** button to synchronize the physical cartridge inventory from the attached tape library with the TS7740 Virtualization Engine database.

8.5 Tape cartridge management

Most of the tape management operations are described in 8.1, “User interfaces” on page 410.

Here we give information about tape cartridges and labels, inserting and ejecting stacked volumes, and exception conditions.

8.5.1 3592 tape cartridges and labels

The data tape cartridge used in a 3592 contains the following items:

- ▶ A single reel of magnetic tape
- ▶ Leader pin (1)
- ▶ Clutch mechanism (2)
- ▶ Cartridge write-protect mechanism (3)
- ▶ Internal cartridge memory (CM)

Figure 8-95 shows a J-type data cartridge, and Table 8-7 on page 571 provides information about the various 3592 J-type data cartridges.




Figure 8-95 Tape cartridge

Table 8-7 Types of IBM 3592 tape cartridges

Text on product label; media type	Native capacity		Color of case	Color of label, door, and write-protect switch
	TS1120	TS1130		
Data; JB ^a	700 GB	1 TB	Black	Dark green
Cleaning; CLNxxxJA	N/A	N/A	Black	Gray
Data; JA ^b	300/500 GB	640 GB	Black	Dark blue
Economy; JJ ^c	60/100 GB	128 GB	Black	Light blue

- a. Use of JB tape cartridges is supported only with TS1120 Tape Drives operating in native capacity mode. If your TS7740 Virtualization Engine shipped prior to January 26, 2007, and you intend to use JB tape cartridges, you should order FC0521, Functional enhancement field. TS1120 Tape Drives/3592 E05 Tape Drives will operate in native E05 mode or in 3592 J1A emulation mode. However, all 3592 Tape Drives associated with the TS7740 Virtualization Engine must be TS1120 Tape Drives in order for 3592 E05 Tape Drives to operate in native E05 mode. To use 3592 E05 Tape Drives in native E05 mode, the drives must be set to E05 native mode. If you intend to use TS1120 Tape Drives in native E05 Mode, the minimum microcode levels are 8.0.1.xx for the TS7740 Virtualization Engine and 534.2x for the Library Manager. Use FC0521, Functional enhancement field, to receive the most current levels of TS7700 Virtualization Engine code.
- b. The native capacity of a JA tape cartridge used in a 3592 J1A drive or a 3592 E05 Tape Drive in J1A emulation mode is 300 GB. The native capacity of a JA tape cartridge used in a 3592 E05 Tape Drive drive in native mode is 500 GB.
- c. The native capacity of a JJ tape cartridge used in a 3592 J1A drive or 3592 E05 Tape Drive in native mode is 60 GB. The native capacity of a JJ tape cartridge used in a 3592 E05 Tape Drive in native mode is 100 GB.

Note: When 3592 J1A drives (or 3592 E05 Tape Drives in J1A emulation) are replaced with 3592 E06 Tape Drives, the TS7740 Virtualization Engine marks the J1A formatted tapes with active data FULL. By marking these tapes full, the TS7740 Virtualization Engine does not append more data, because the 3592 E06 Tape Drive cannot append data to a J1A formatted tape. As the active data on the J1A gets reclaimed or expired, the tape goes back to the scratch pool, and then eventually gets reformatted to the E06 data format.

If you have a tape that is written in E06 format, the capacity is 1 TB for JB/JX media, 640 GB for JA/JW media, and 128 GB for JJ/JR media.

Labels

The cartridges use a media label to describe the cartridge type, as shown in Figure 8-96 (JA example). In tape libraries, the library vision system identifies the types of cartridges during an inventory operation. The vision system reads a volume serial number (VOLSER), which appears on the label on the edge of the cartridge. The VOLSER contains from one to six characters, which are left-aligned on the label. If fewer than six characters are used, spaces are added. The media type is indicated by the seventh and eighth characters.




Figure 8-96 Cartridge label

8.5.2 Manual insertion of stacked cartridges

There are two methods for physically inserting a stacked volume into the TS3500 Tape Library:

- ▶ Opening the library doors and directly inserting the tape into the tape library storage cells
- ▶ Using the tape library I/O station

Inserting directly into storage cells

Open the front door of a frame and bulk load the cartridges directly into empty storage slots. This method takes the TS3500 Tape Library out of operation; therefore, use it only to add or remove large quantities of tape cartridges.

The TS3500 Tape Library Cartridge Assignment Policy (CAP) defines which volumes are assigned to which logical library partition. If the VOLSER is included in the TS7740 Virtualization Engine's range, then it will be assigned to the associated TS3500 Tape Library logical library partition.

After the doors on the library are closed and the tape library has performed inventory, the upload of the inventory to the TS7700 Virtualization Engine will be processed before the TS3500 Tape Library reaches the READY state. The TS7700 Virtualization Engine updates its database accordingly.

Notes:

- ▶ The inventory is performed only on the frame where the door is opened and not on the frames to either side. If you insert cartridges into a frame adjacent to the frame that you opened, then you must perform a manual inventory of the adjacent frame, using the operator panel on the TS3500 Tape Library itself.
- ▶ For a TS7740 Virtualization Engine it is important to note that the external cartridge bar code label and the internal VOLID label match or, as is the case for a new cartridge, the internal VOLID label is blank. If the external label and the internal label do not meet the aforementioned criteria the cartridge will be rejected.

Inserting cartridges using the I/O station

The TS3500 Tape Library detects volumes in the I/O station, and then moves the volumes to empty cells. The TS3500 Tape Library Cartridge Assignment Policy defines which volumes are assigned to which logical library. If the VOLSER is included in the System z range, it will be assigned to the TS3500 Tape Library logical library partition. If any VOLSER is not in a range defined by the CAP, then the operator identifies a System z logical library as the destination using the Insert Notification process.

Under certain conditions, cartridges are not assigned to a logical library partition in the TS3500 Tape Library. With TS7700 Virtualization Engine R1.5 and later, the TS3500 must have a dedicated logical partition for the cluster. Therefore, in a library with more than one partition, be sure that the Cartridge Assignment Policy is kept up to date with the cartridge volume range (or ranges) in use. This way minimizes conflicts by ensuring the cartridge is accessible only by the intended partition.

Note: Unassigned cartridges can exist in the TS3500 Tape Library, but “unassigned” can have different meanings and needs different actions from the operator. See *IBM System Storage TS3500 Tape Library with ALMS Operator Guide*, GA32-0594 for more information.

8.5.3 Exception conditions

On a physical library, one important exception condition is the Out of Physical Volumes condition.

Out of Physical Volumes

When a distributed library associated with a cluster runs out of scratch stacked physical volumes, operations of the TS7740 Virtualization Engine are impacted.

As part of normal processing, data is copied from cache to physical volumes in a primary pool managed by the Virtualization Engine. A copy might also be made to a physical volume in a secondary pool if the dual copy function is specified using Management Class. Empty physical volumes are needed in a pool, or if a pool is enabled for borrowing, in the common scratch pool, for operations to continue. If a pool runs out of empty physical volumes and there are no volumes that can be borrowed or borrowing is not enabled, operations that might use that pool on the distributed library must be suspended. If one or more pools run out of empty physical volumes, the distributed library enters the Out of Physical Scratch state. The Out of Physical Scratch state is reported to all hosts attached to the cluster associated with the distributed library and, if included in a Grid configuration, to the other clusters in the Grid. The following MVS console message is generated to inform an operator of this condition:

```
CBR3789E VTS library library-name is out of empty stacked volumes.
```

Library-name is the name of the distributed library in the state. The CBR3789E message will remain on the MVS console until empty physical volumes have been added to the library or the pool that is out has been enabled to borrow from the common scratch pool and there are empty physical volumes to borrow. Intervention required conditions are also generated for the out of empty stacked volume state and for the pool that is out of empty physical volumes.

If the option to send intervention conditions to attached hosts is set on the TS7700 Virtualization Engine that is associated with the distributed library, the following console

messages are also generated to provide specifics about the pool that is out of empty physical volumes:

CBR3750I Message from library library-name: 0P0138 The Common Scratch Pool (Pool 00) is out of empty media volumes.

CBR3750I Message from library library-name: 0P0139 Storage pool xx is out of scratch volumes.

The 0P0138 message indicates the media type that is out in the common scratch pool. These messages do not remain on the MVS console. The intervention conditions can be viewed through the TS7700 Virtualization Engine management interface.

If the TS7740 Virtualization Engine is in a Grid configuration, and if its associated distributed library enters the out-of-empty-stacked-volume state, operations are affected in other ways:

- ▶ All copy operations are immediately suspended in the cluster (regardless of which pool has become empty).
- ▶ If the cluster has a copy consistency point of RUN, the Grid enters the Immediate Mode Copy Operations Deferred state, and an MVS console message is generated:

CBR3787E One or more immediate mode copy operations deferred in library library-name.
- ▶ If another cluster attempts to copy a logical volume that is not resident in the cache, the copy attempt fails.
- ▶ In choosing a TVC cluster, the Grid prefers clusters that are not in the out-of-empty-stacked-volume state, but could still select a remote TVC whose cluster is in that state. If the data needed is not in the remote cluster's TVC, the recall of the data will fail. If data is being written to the remote cluster's TVC, the writes will be allowed, but because there might not be any empty physical volumes available to copy the data to, the cache might become full of data that cannot be copied and all host I/O using that cluster's TVC will become throttled to prevent a cache overrun.

Note: Because having a distributed library in the out-of-empty-stacked-volume state impacts operations in a TS7740 Virtualization Engine, it is something that should be avoided if at all possible.

Monitor the number of empty stacked volumes in a library. If the library is close to running out of a physical volume media type, action should be taken to either expedite the reclamation of physical stacked volumes or add additional ones. You can use the Bulk Volume Information Retrieval function to obtain the physical media counts for each library. The information obtained includes the empty physical volume counts by media type for the common scratch pool and each defined pool.

Above Threshold Warning state

The TS7740 Virtualization Engine enters the Above Threshold Warning state when the amount of data to copy exceeds the threshold for the installed cache capacity for five consecutive sample periods (the amount of data to copy is sampled every 30 seconds). The TS7740 Virtualization Engine leaves the Above Threshold Warning state when the amount of data to premigrate is below the threshold capacity for 30 consecutive sample periods. The consecutive sampling criteria is to prevent excessive messages from being created. This state produces the following message:

CBR3750I Message from library library-name:0P0160 Above threshold for uncopied data in cache, throttling possible

8.6 Managing logical volumes

In addition to the tasks described in Chapter 4, “Hardware implementation” on page 175 and in 8.2, “TS7740 Virtualization Engine management interface” on page 415, we cover a few more management tasks and considerations for logical volumes in the following sections.

8.6.1 Scratch volume recovery for logical volumes

The advantage of this method of managing data is that if you determine that a volume was mistakenly returned to scratch, you only have to return the volume to private status to recover from the mistake, as long as you have not reused the volume or the “grace period” has not expired. The method to recover depends on the tape management system used. In general, change the status volumes from scratch to private and change the expiration date by adding at least one week to prevent the Tape Management System from returning the volume to scratch during the next few days. For example, for DFSMSrmm, the following command will return the volume to private status and increase its retention period, including communicating the change to the TS7700 Virtualization Engine (see *z/OS DFSMSrmm Guide and Reference*, SC26-7404 for details about the command):

```
RMM CHANGEVOLUME yyyyyy STATUS(USER) RETPD(days) OWNER(userid)
```

In the command, *yyyyyy* is the VOLSER.

8.6.2 Ejecting logical volumes

Logical volumes are not physical entities that can be individually removed from the library. They reside on stacked volumes with many other logical volumes.

Because of the permanent nature of the EJECT, the TS7700 Virtualization Engine only allows you to EJECT a logical volume that is in either the INSERT or SCRATCH (defined with fast-ready attribute) categories. If a logical volume is in any other status, the EJECT fails. So, if we eject a scratch volume and later we need to recover the data, it will not be possible to do that action, so keep that in mind before you eject logical volumes.

Note: This fact has proven to be cumbersome for volumes that happen to be in the ERROR category (000E). An easy way to eject such volumes is to use ISMF screens to set these volumes to the PRIVATE status. The volume status is propagated to DFSMSrmm. You can use DFSMSrmm to subsequently assign the volume to the Pending Release status, and the next RMM HSKP run will return it to a SCRATCH status, allowing you to eject it.

Ejecting large numbers of logical volumes can have a performance impact on the host and the library.

Tapes that are in INSERT status can be ejected by the resetting of the return code through the CBRUXENT exit. This exit is usually provided by your tape management system vendor. Another way to EJECT cartridges in the INSERT category is by using the MI. For more information, see “Delete Logical Volumes window” on page 448.

After the tape is in SCRATCH status, follow the procedure for EJECT processing based on whether your environment is system-managed tape or BTLS. You also must follow the procedure that is specified by your tape management system vendor. For DFSMSrmm, issue the **RMM CHANGEVOLUME volser EJECT** command.

If your tape management system vendor does not specify how to do this, you can use one of the following commands:

- ▶ z/OS command `LIBRARY EJECT,volser`
- ▶ IDCAMS command `LIBRARY EJECT,volser` (for BTLS)
- ▶ ISMF EJECT line operator for the tape volume

The eject process fails if the tape is in another status or category. For libraries managed under DFSMS system-managed tape, the system command `LIBRARY EJECT,volser` issued to a logical volume in `PRIVATE` status fails with this message:

CBR3726I Function incompatible error code 6 from library <library-name> for volume <volser>.

Note: In a DFSMS system-managed tape environment, if you try to eject a logical volume and get this error, OAM notifies the tape management system. This is done through the OAM eject exit `CBRUXEJC` command before the eject request is sent to the tape library. The Integrated Library Manager will eventually fail the eject, but the tape management system has already marked the volume as ejected. Prior to APAR OW54054, there was no notification back that the eject failed.

Failed Eject Notification was added to OAM with APAR OW54054 and is currently in all supported releases of DFSMS. Any tape management system supporting this notification can use this function.

If your tape management system is DFSMSrmm, you can use the following commands to clean up the RMM CDS for failed logical volume ejects and to resynchronize the TCDB and RMM CDS:

```
RMM SEARCHVOLUME VOL(*) OWN(*) LIM(*) INTRANSIT(Y) LOCATION(vts) -  
    CLIST('RMM CHANGEVOLUME ',' LOC(vts)')  
  
EXEC EXEC.RMM
```

The first RMM command asks for a list of volumes that RMM thinks it has ejected and writes a record for each in a sequential data set called *prefix.EXEC.RMM.CLIST*. The CLIST then checks that the volume is really still resident in the VTS library and, if so, it corrects the RMM CDS.

Note: Limiting the number of outstanding ejects to a couple of thousand total per system will limit exposure to performance problems.

There are considerations to be aware of when ejecting large numbers of logical volumes. APAR OW42068 introduced Peer-To-Peer toleration support. This in effect treats all libraries as though they had the same eject processing restrictions as a Peer-to-Peer. This support is currently in all supported releases of DFSMS. There is a logical volume eject limit of 1000 that the Peer-to-Peer VTS can handle at any one time. Because libraries can be shared among systems, this limit can be reached quickly if many ejects are issued from multiple hosts.

OAM helps mitigate this situation by restricting the number of ejects sent to each library at a given time and manages all the outstanding requests. This management requires storage on the host, and a large number of ejects can force OAM to reserve large amounts of storage. Additionally, there is a restriction on the number of eject requests on the device services' queue. All of these conditions can have an impact on the host's performance.

Therefore, a good limit for the number of outstanding eject requests is no more than two thousand per system. Additional ejects can be initiated when others complete. For additional information, refer to APAR OW42068. The following commands can be used on the System z hosts to list the outstanding and the active requests:

```
F OAM,QUERY,WAITING,SUM,ALL
F OAM,QUERY,ACTIVE,SUM,ALL
```

8.7 Messages and displays

This section describes the enhanced message support and some relevant messages related to the TS7700 Virtualization Engine.

8.7.1 Console name message routing

Today, with console name message routing support, many of the library-specific messages are only issued to the specified library console (if defined) and not to the specified routing codes.

Although this is not specific to a TS7700 Virtualization Engine, the following critical, action-related messages are now issued using the specified library console and routing codes, providing maximum visibility:

- ▶ CBR3759E Library x safety enclosure interlock open.
- ▶ CBR3764E Library x all storage cells full.
- ▶ CBR3765E No cleaner volumes available in library x.
- ▶ CBR3753E All convenience output stations in library x are full.
- ▶ CBR3754E High capacity output station in library x is full.
- ▶ CBR3755E {Input|Output} door open in library x.
- ▶ CBR3660A Enter {list of media inserts} scratch volumes into x.

8.7.2 Alert setting messages

With the new SETTING function, we have a new set of messages that we need to know. These messages are described in 8.3.3, “Host Console Request function” on page 544, but we prefer to summarize here all the messages.

The message format is as follows:

CBR3750I Message from library *lib-id*: ALxxxx *message description*

The ALxxxx codes are as follows:

- ▶ AL0000 Uncopied data of xxxxxx GB above low warning limit of yyyyyy GB
- ▶ AL0001 No longer above uncopied data low warning limit of yyyyyy GB
- ▶ AL5000 Uncopied data of xxxxxx GB above high warning limit of yyyyyy GB
- ▶ AL5001 No longer above uncopied data high warning limit of yyyyyy GB
- ▶ AL0004 Available physical drives of xx is below low limit of yy
- ▶ AL0005 Available physical drives no longer below low limit of yy
- ▶ AL5004 Available physical drives of xx is below critical limit of yy
- ▶ AL5005 Available physical drives no longer below critical limit of yy
- ▶ AL0006 Available physical scratch volumes no longer below critical limit of yy for pool zz

- ▶ AL0007 Available physical scratch volumes of *xxx* below critical limit of *yyy* for pool *zz*
- ▶ AL5006 Available physical scratch volumes no longer below critical limit of *yyy* for pool *zz*
- ▶ AL5007 Available physical scratch volumes of *xxx* below critical limit of *yyy* for pool *zz*
- ▶ AL0008 Resident data of *xxxxxx* GB above low warning limit of *yyyyyy* GB
- ▶ AL0009 No longer above resident data low warning limit of *xxxxxx* GB
- ▶ AL5008 Resident data of *xxxxxx* GB above high warning limit of *yyyyyy* GB
- ▶ AL5009 No longer above resident data high warning limit of *xxxxxx* GB

8.7.3 Grid messages

This section lists some of the TS7700 Virtualization Engine Grid-specific messages that you might see. For a complete and current list, see the appropriate volume of *z/OS MVS System Messages*.

Incompatibility error message

In the case of an incompatible function error, you might see the message CBR3726I:

CBR3726I Function incompatible error code *error-code* from library *library-name* for volume *volser*.

In this message, an error has occurred during the processing of volume *volser* in library *library-name*. The library returned a unit check with an error code *error-code*, which indicates that an incompatible function has been requested. A command has been issued that requests an operation that is understood by the subsystem microcode, but cannot be performed because of one of the following errors:

- | | |
|--------------|--|
| X'00' | The function requested is not supported by the subsystem to which the order was issued. |
| X'01' | The library attachment facility is not installed and allowed. |
| X'02' | Not currently used. |
| X'03' | The high capacity input/output facility is not configured. |
| X'04' | Reserved. |
| X'05' | The volume requested to be mounted is not compatible with the device allocated. |
| X'06' | The logical volume can only be ejected if it is in the insert category and has a mount count of zero, or is assigned to a category that has the Fast-Ready attribute set. |
| X'07' | There is no pending import or export operation to cancel. |
| X'08' | There are not enough (four are needed) physical drives available to initiate the import or export operation. |
| X'09' | Reserved. |
| X'0A' | Reserved. |
| X'0B' | Reserved. |
| X'0C' | Reserved. |
| X'0D' | The TS7700 Virtualization Engine Grid subsystem is either in service preparation mode or has an unavailable component within the subsystem, such as an unavailable distributed library. Audit, eject, or entry-related commands are not being accepted at this time. |

X'0E'	The TS7700 Virtualization Engine Grid subsystem already has one thousand eject requests queued and is not accepting any more eject requests now.
X'0F'	An inappropriate library function was issued to the TS7700 Virtualization Engine Grid subsystem.
X'10'	The VTC in the Peer-to-Peer VTS subsystem or the distributed library in a TS7700 Virtualization Engine Grid configuration that the command was issued to is in read-only or write-protect mode and is not accepting requests that change the category or attributes of a volume. This mode of operation is provided to support disaster recovery operations in a configuration where the configuration is split between two physical sites.
X'12'	The volume specified has a non-zero expiration time associated with it. A volume in this state cannot be mounted, moved, or have its attributes modified until the expiration time has elapsed.
X'30'	The TS7700 Virtualization Engine Cluster that the command was received from, or has an available path to the cluster that currently owns the volume and ownership takeover is not enabled, cannot obtain ownership of the volume from the cluster that owns it and ownership takeover is not enabled.
X'31'	A non-recoverable internal microcode error was detected by the TS7700 Virtualization Engine.
X'32'	There is more than one valid copy of the specified export list volume in the TS7700 Virtualization Engine Grid configuration.
X'33'	An export operation was issued to a TS7700 Virtualization Engine that is performing a global operation. Global operations include volume inserts, volume deletions through the management interface, damaged volume recovery, and disaster recovery. Export operations are not being accepted at this time.
X'38'	An export operation was issued to a TS7700 Virtualization Engine and the export list volume specified is a logical WORM volume. The export list volume cannot be WORM.

8.7.4 Display Grid status

The following messages can be issued for the TS7700 Virtualization Engine Grid.

CBR1100I OAM status

This message is issued in response to the following operator command:

DISPLAY SMS,OAM

Example 8-15 shows the complete message text.

Example 8-15 DISPLAY SMS,OAM command

```

CBR1100I OAM status: 618
  TAPE TOT  ONL  TOT  TOT  TOT  TOT  TOT  ONL  AVL  TOTAL
    LIB  LIB  AL  VL  VCL  ML  DRV  DRV  DRV  SCRTCH
      5    2    0    0    2    0  192  192  127  46079
There are also 2 VTS distributed libraries defined.
CBRUXCUA processing ENABLED.
CBRUXEJC processing ENABLED.
CBRUXENT processing ENABLED.
CBRUXVNL processing ENABLED.

```

If any TS7700 Virtualization Engine subsystems are defined to the system, the following status line is displayed, reflecting the number of distributed libraries that are associated with the composite libraries:

There are also numvdl-lib VTS distributed libraries defined.

CBR1110I OAM library status

The CBR1110I message is issued in response to the following command:

DISPLAY SMS,LIBRARY(library-name),DETAIL

Example 8-16 shows the complete message text.

Example 8-16 DISPLAY SMS, LIBRARY command

CBR1110I OAM library status: 334										
TAPE	LIB	DEVICE	TOT	ONL	AVL	TOTAL	EMPTY	SCRATCH	ON OP	
LIBRARY	TYP	TYPE	DRV	DRV	DRV	SLOTS	SLOTS	VOLS		
HYDRA0	VCL	3957-V06	128	128	64	0	0	44365	Y	Y

MEDIA	SCRATCH		SCRATCH		SCRATCH					
TYPE	COUNT		THRESHOLD		CATEGORY					
MEDIA2	44365		100		0002					

DISTRIBUTED LIBRARIES: HYDRAD										
LIBRARY ID: 10001										
OPERATIONAL STATE: AUTOMATED										
ERROR CATEGORY SCRATCH COUNT:						0				
CORRUPTED TOKEN VOLUME COUNT:						0				

Note: Library type VCL indicates the Composite Library, as opposed to VDL for the Distributed Library.

CBR1180I OAM tape volume status

This CBR1180I message is issued in response to the following operator command:

DISPLAY SMS,VOLUME(volser)

Example 8-17 lists the complete message text for logical volume A01878.

Example 8-17 DISPLAY SMS,VOLUME command

```
CBR1180I OAM tape volume status: 933
VOLUME MEDIA STORAGE LIBRARY USE W C SOFTWARE LIBRARY
TYPE GROUP NAME ATR P P ERR STAT
CATEGORY
A01878 MEDIA2 SGHYP02 HYDRA0 P N N NOERROR PRIVATE
-----
RECORDING TECH: 36 TRACK COMPACTION: YES
SPECIAL ATTRIBUTE: NONE ENTER/EJECT DATE: 2006-08-03
CREATION DATE: 2006-08-03 EXPIRATION DATE: 2006-09-23
LAST MOUNTED DATE: 2006-09-20 LAST WRITTEN DATE: 2006-09-20
SHELF LOCATION:
OWNER: BHAEUSS
LM SG: SGHYP02 LM SC: SCHYPG1 LM MC: MCHYNCOP LM DC: DCHY1GB
-----
Logical volume.
Volume is cache resident.
```

8.7.5 Warning link status degraded messages

With Dynamic Link Load Balancing function in the TS7700 Virtualization Engine, host console messaging is provided to assist with GRID performance problems. Host console messages are issued when the subsystem determines that the GRID links have crossed a preset performance balance threshold.

There are two network links from each cluster participating in a GRID configuration. Every five minutes, the Dynamic Link Load Balancing function evaluates the capabilities of each link between the clusters. If performance in one of the links is 60% less than the other, a warning message is displayed at the System Console in the following format:

```
CBR3796E Grid links degrade in library library_name
```

When the Grid links are unbalanced, we can use the Host Console Request function STATUS GRIDLINK (see 8.3.3, “Host Console Request function” on page 544) to get additional information that can be useful with problem determination. Identifying which of the two links is being affected can be helpful in resolving the unbalanced condition. Because we cannot change the TS7700 Virtualization Engine link parameters in any way, we are limited to investigating the problem from the perspective of the GRID networking equipment or cabling.

A detailed description of the Host Console Request functions and responses is available in the *IBM Virtualization Engine TS7700 Series z/OS Host Command Line Request User's Guide* white paper. The most recently published white papers are available at the Techdocs website; search for TS7700 Virtualization Engine at the following address:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs>

When the Grid network performance issue is resolved and the links are balanced, a message is presented at the System Console in the following format:

```
CBR3797E Grid links in library_name are no longer degraded
```

8.7.6 Warning VTS operation degraded messages

When a VTS is operating in a degraded state, the following message is issued:

CBR3786E VTS operation degraded in library *library-name*

When the degradation is resolved, you see this message:

CBR3768I VTS operations in library *library-name* no longer degraded

8.7.7 Warning cache use capacity (TS7720 Virtualization Engine)

For the TS7720 Virtualization Engine, warning and critical cache free space messages are displayed:

CBR3792E Library *library-name* has entered the limited cache free space warning state

CBR3794E Library *library-name* has entered the out of cache resources critical state.

When the cache situation is resolved, the following messages are shown:

CBR3793I Library *library-name* has left the limited cache free space warning state

CBR3795I Library *library-name* has left the out of cache resources critical state

8.8 Recovery scenarios

In this section, we discuss some of the potential recovery scenarios you might have to perform. You are notified about most of the errors that require operator attention through an Intervention Required message on the TS7700 Virtualization Engines management interface and the hosts.

8.8.1 Hardware conditions

Here we describe potential hardware failure scenarios.

IBM 3592 Tape Drive failure (TS7740 Virtualization Engine with library attachment only)

When the TS7700 Virtualization Engine determines that one of its tape drives is not operating correctly and requires service (it is likely that the drive has excessive read or write errors), the drive is marked offline and an SSR must be engaged. The following intervention required message is displayed on the Library Manager Console:

CBR3750I MESSAGE FROM LIBRARY *lib*: Device xxx made unavailable by a VTS. (VTS z)

Operation of the TS7700 Virtualization Engine continues with a reduced number of drives until the repair action on the drive is complete. To recover, the SSR repairs the failed tape drive and makes it available for the TS7700 Virtualization Engine to use it again.

Power failure

User data is protected in the event of a power failure, as it is stored on the TVC. Any host jobs reading or writing to virtual tapes will fail as they would with a real IBM 3490E and will need to be restarted after the TS7700 Virtualization Engine is available again. When power is restored and stable, the TS7700 Virtualization Engine must be powered up manually. The TS7700 Virtualization Engine will recover access to the TVC, using information available from the TS7700 Virtualization Engine database and logs.

TS7700 Virtualization Engine database corruption

To recover from corruption of the TS7700 Virtualization Engine database, you have to perform a disaster recovery. This is a TS7700 Virtualization Engine process further described in Chapter 10, “Failover and disaster recovery scenarios” on page 725.

The TS7700 Virtualization Engine maintains a database of information about the location and status of logical volumes on the stacked volumes it manages. When a stacked volume has been filled with logical volumes, a backup of the entire database is placed at the end of the filled stacked volume. The database contains a time and date stamp that identifies when the backup was performed.

When the database copy operation is complete, a message is sent to the attached hosts:

```
CBR3750I MESSAGE FROM LIBRARY lib: VTS Database Backup written to Physical Tape  
xxxxxx.
```

The disaster recovery process causes the TS7740 Virtualization Engine to load the stacked volumes, locate the latest version of the database, and restore it. Any logical volumes written after the last backup of the TS7740 Virtualization Engine database are lost. When the restoration is complete, a message is displayed on the Library Manager console informing you of the date and time when the TS7700 Virtualization Engine database was restored.

Accessor failure and manual mode (TS7740 Virtualization Engine with Tape Library attachment only)

If the TS3500 Tape Library does not have the dual accessors installed, failure of the accessor results in the library being unable to automatically mount physical volumes. If TS3500 Tape Library dual accessors are installed, the second accessor takes over. Then, you may call your SSR to repair the failed accessor.

Gripper failure (TS7740 Virtualization Engine only with Tape Library attachment only)

The TS3500 Tape Library has dual grippers; if a gripper fails, library operations continue with the other one. While the gripper is being repaired, the accessor is not available until the repair is complete. If the dual accessors are installed, the second accessor is used until the gripper is repaired. For detailed information about operating the TS3500 Tape Library, see *IBM System Storage TS3500 Tape Library with ALMS Operator Guide*, GA32-0594.

Out of stacked volumes (TS7740 Virtualization Engine with Tape Library attachment only)

If the Tape Library runs out of stacked volumes, then copying to the 3592 Tape Drives will fail and an Intervention Required message is sent to the host and the TS7700 Virtualization Engine management interface. All further logical mount requests are delayed by the Library Manager until more stacked volumes are added to the TS3500 Tape Library connected to the TS7740 Virtualization Engine. To recover, insert more stacked volumes. Copy processing can then continue.

Damaged cartridge pin

The 3592 has a metal pin that is grabbed by the feeding mechanism in the 3592 tape drive to load the tape on to the take-up spool inside the drive. If this pin gets dislodged or damaged, then follow the instructions in *IBM TotalStorage Enterprise Tape System 3592 Operators Guide*, GA32-0465 to correct the problem.

Important: Repairing a 3592 tape should only be done for data recovery. After the data has been moved to a new volume, replace the repaired cartridge.

Broken tape

If a 3592 tape cartridge is physically damaged and unusable (the tape is crushed, or the media is physically broken, for example), the TS7740 Virtualization Engine cannot recover the contents. This is the same for any tape drive media cartridges. You can generate a list of logical volumes that are on that stacked volume. Consult with your SSR to determine if IBM services are available to attempt data recovery from a broken tape.

Logical mount failure

When a mount request is received for a logical volume, the TS7700 Virtualization Engine determines whether the mount request can be satisfied and if so, tells the host that it will process the request. Unless an error condition is encountered in the attempt to mount the logical volume, the mount operation completes and the host is notified that the mount was successful. With the TS7700 Virtualization Engine, the way that a mount error condition is handled is different than with the prior generations of VTSs. With the prior generation of VTSs, the VTS always indicated to the host that the mount completed, even if a problem had occurred. When the first I/O command was issued, the VTS would then fail that I/O because of the error. This would result in a failure of the job, without the opportunity to try and correct the problem and retry the mount.

With the TS7700 Virtualization Engine subsystem, if an error condition is encountered during the execution of the mount, instead of indicating that the mount was successful, the TS7700 Virtualization Engine returns completion and reason codes to the host indicating that a problem was encountered. With DFSMS, the logical mount failure completion code results in the console messages shown in Example 8-18:

Example 8-18 Unsuccessful mount completion and reason codes

```
CBR4195I LACS RETRY POSSIBLE FOR JOB job-name
CBR4171I MOUNT FAILED. LVOL=logical-volser, LIB=library-name,
PVOL=physical-volser, RSN=reason-code
...
CBR4196D JOB job-name, DRIVE device-number, VOLSER volser, ERROR CODE error-code.
REPLY 'R' TO RETRY OR 'C' TO CANCEL
```

Reason codes provide information about the condition that caused the mount to fail. The reason codes that might be presented are as follows:

X'10'	Internal Error Detected
X'11'	Resend Special Case
X'20'	Specific Volume In Use On Another Cluster
X'21'	Scratch Volume Selected In Use On Another Cluster
X'22'	Valid Volume Inaccessible
X'23'	Local Cluster Path to Volume's Data No Longer Available
X'24'	Remote Cluster Path To Volume's Data No Longer Available
X'25'	Copy Required, but Cluster Copying Inhibited
X'30'	Local Cluster Recall Failed, Stacked Volume Misplaced
X'31'	Local Cluster Recall Failed, Stacked Volume Inaccessible
X'32'	Local Cluster Recall Failed, Stacked Volume Unavailable
X'33'	Local Cluster Recall Failed, Stacked Volume No Longer In Library
X'34'	Local Cluster Recall Failed, Stacked Volume Load Failure
X'35'	Local Cluster Recall Failed, Stacked Volume Access Error

X'38'	Remote Cluster Recall Failed, Stacked Volume Misplaced
X'39'	Remote Cluster Recall Failed, Stacked Volume Inaccessible
X'3A'	Remote Cluster Recall Failed, Stacked Volume Unavailable
X'3B'	Remote Cluster Recall Failed, Stacked Volume No Longer In Library
X'3C'	Remote Cluster Recall Failed, Stacked Volume Load Failure
X'3D'	Remote Cluster Recall Failed, Stacked Volume Access Error

When a mount is failed on a TS7700 Virtualization Engine, you may attempt to resolve the underlying problem indicated by the reason code and then have the mount retried. For example, if the failure was because a recall was required and the stacked volume was unavailable because it was accidentally removed from the library, recovery involves returning the volume to the library and then replying to the outstanding message with RETRY.

Orphaned logical volume

This situation occurs when the TS7700 Virtualization Engine database has a reference to a logical volume but no reference to its physical location. This could result from hardware or internal software errors. When it does occur, any data on the logical volume is lost. When this error occurs, contact your SSR.

Internal-external label mismatch

If a label mismatch occurs, the stacked volume is ejected to the convenience Input/Output station and the Intervention Required condition is posted at the TS7700 Virtualization Engine management interface and sent to the host console (Example 8-19).

Example 8-19 Label mismatch

```
CBR3750I MESSAGE FROM LIBRARY lib: A stacked volume has a label mismatch and has
been ejected to the Convenience Input/Output Station.
Internal: xxxxxx, External: yyyyyy
```

The host is notified that intervention-required conditions exist. Investigate the reason for the mismatch. If possible, relabel the volume to use it again.

Failure during reclamation

If there is a failure during the reclamation process, the process is managed by the TS7700 Virtualization Engine microcode. No user action is needed; recovery is managed internally.

Excessive temporary errors on stacked volume

When a stacked volume is determined to have an excessive number of temporary data errors, to reduce the possibility of a permanent data error, the stacked volume is placed in read-only status.

8.8.2 TS7700 Virtualization Engine software failure

If a problem develops with the TS7700 Virtualization Engine software, the TS7700 Virtualization Engine issues an Intervention Required message to the TS7700 Virtualization Engine management interface and host console and attempts to recover. In the worst case, this can involve a reboot of the TS7700 Virtualization Engine itself. If the problem persists, you need to contact your SSR. The Intervention Required message shown in Example 8-20 on page 586 is sent to the host console.

Example 8-20 VTS software failure

CBR3750I MESSAGE FROM LIBRARY *lib*: Virtual Tape System z has a CHECK-1 (xxxx) failure

The TS7700 Virtualization Engine internal recovery procedures handle this situation and restarts the TS7700 Virtualization Engine. See Chapter 10, “Failover and disaster recovery scenarios” on page 725 for more details.

8.9 TS7720 Virtualization Engine considerations

With a homogenous TS7720 Virtualization Engine configuration at Release 1.5 and later, certain operations are unavailable or disabled on the management interface, which indicates that these functions are not applicable to the TS7720 Virtualization Engine, such as any functions that are related to the management of physical libraries, drives, and cartridges.

8.9.1 Management interface (MI)

In this section, we discuss using the TS7700 Virtualization Engine management interface to operate the TS7720 Virtualization Engine.

Health & Monitoring

Figure 8-97 shows you the available Health & Monitoring options in a TS7720 Virtualization Engine configuration.

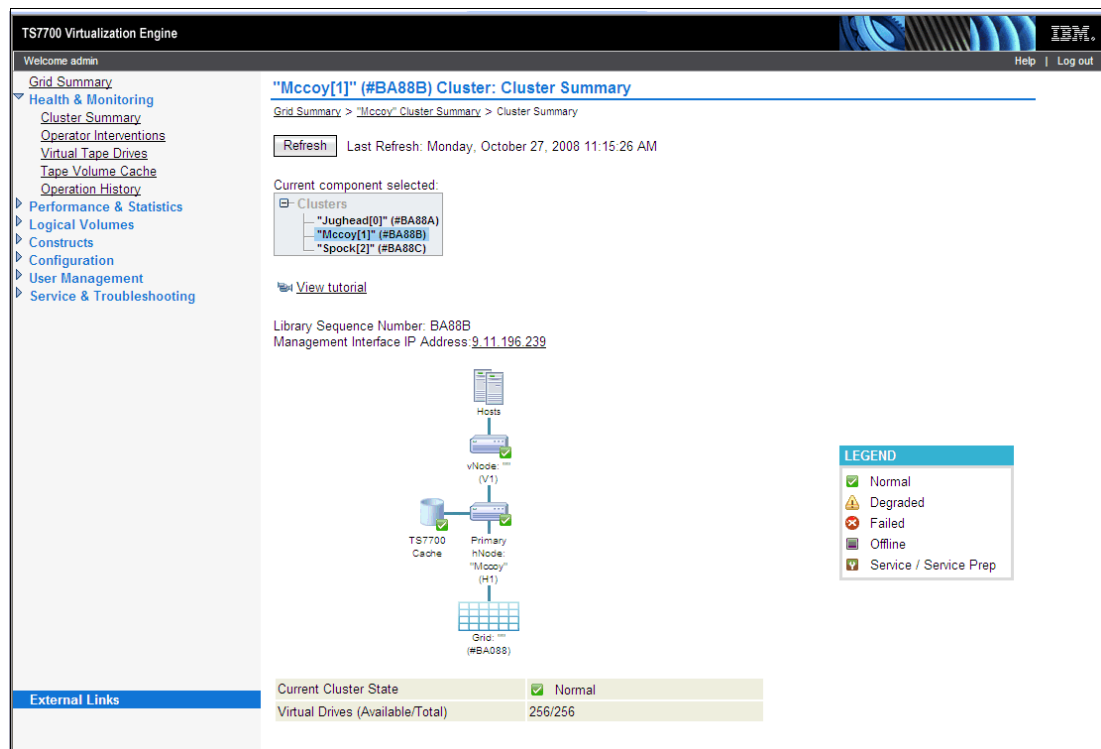


Figure 8-97 TS7720 Virtualization Engine MI Health & Monitoring options

Comparing Figure 8-97 on page 586 to 8.2.3, “Health & Monitoring (TS7740 Virtualization Engine with Library attachment)” on page 419, you can easily see that all selections

regarding Physical Tape are missing (Physical Tape Drives, Physical Media Inventory, and Physical Library).

For all the remaining and possible options, see 8.2, “TS7740 Virtualization Engine management interface” on page 415.

Logical Volumes

Figure 8-98 shows the available options under the Logical Volumes selection menu.

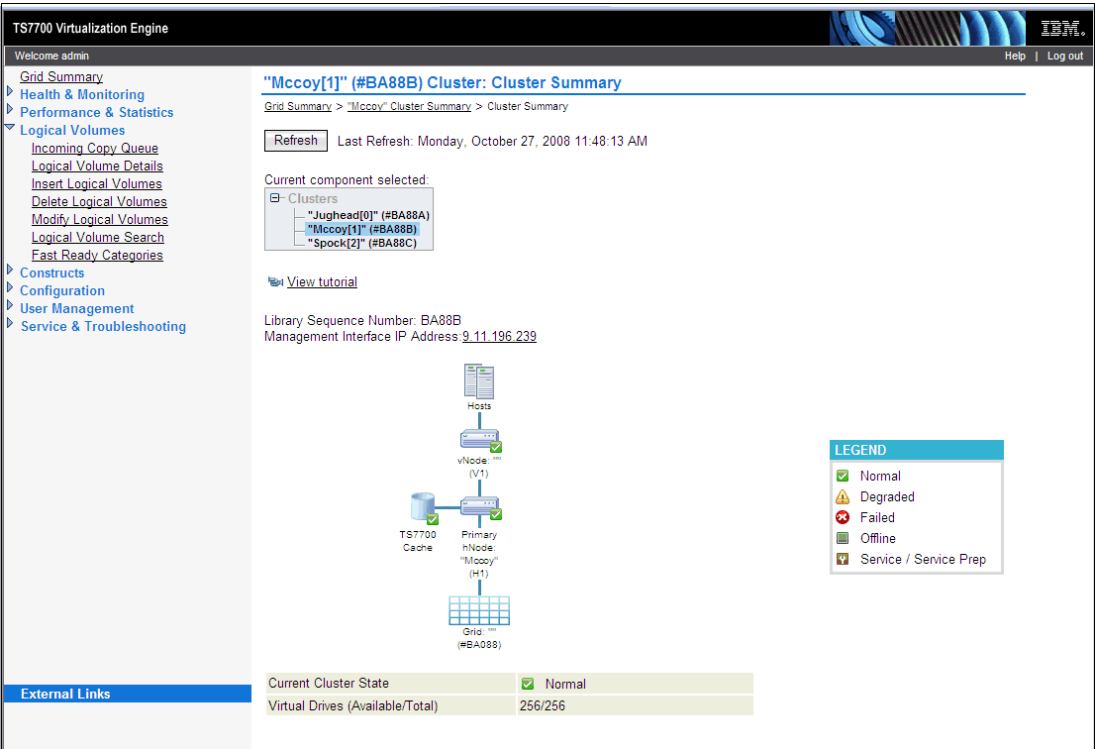


Figure 8-98 TS7720 Virtualization Engine MI Logical Volumes options

Comparing Figure 8-98 to 8.2.5, “Logical volumes” on page 435, you can easily see that all selections regarding Physical Tape are missing (Recall Queue and Move Logical Volumes). Notice also that the entire Physical Volumes selection is removed from the work area on the left.

For all the remaining possible options, see 8.2, “TS7740 Virtualization Engine management interface” on page 415.

Configuration

Figure 8-99 shows the available options under the **Configuration** selection menu.

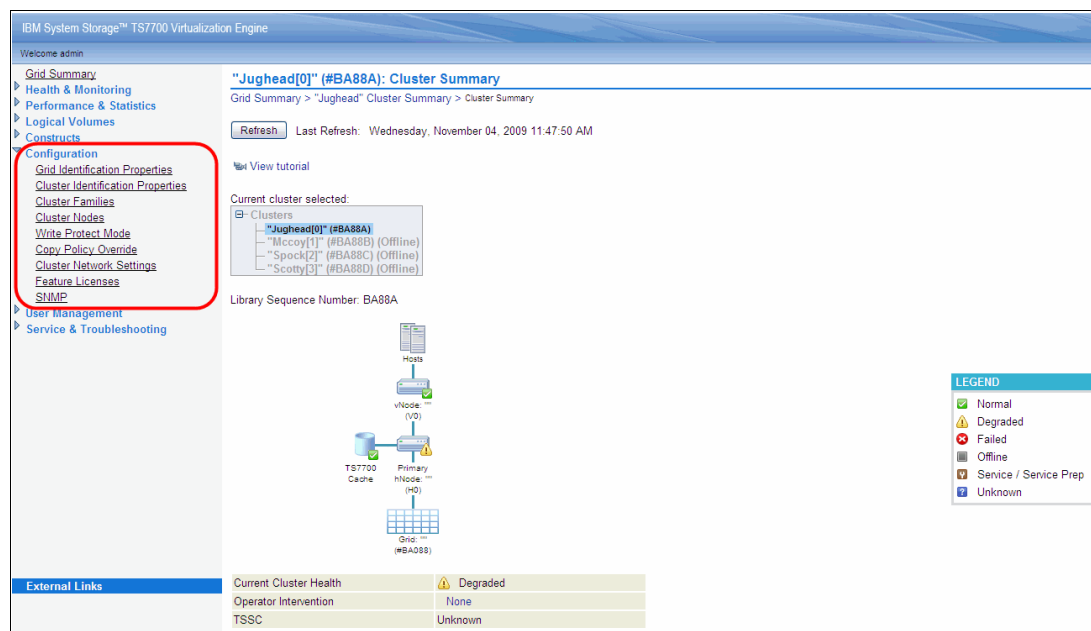


Figure 8-99 TS7720 Virtualization Engine MI Configuration options

Comparing Figure 8-99 to 8.2.8, "Configuration" on page 486, you can easily see that all selections regarding Physical Tape are missing (Inhibit Reclaim Schedules and Encryption Key Manager Addresses).

For all the remaining possible options, see 8.2, "TS7740 Virtualization Engine management interface" on page 415.



Performance and monitoring

In this chapter, we describe the factors that determine and influence the performance of the IBM Virtualization Engine TS7700. We also describe what actions to take, when necessary, to improve the TS7700 Virtualization Engine's performance.

This chapter includes the following information:

- ▶ An overview of the shared tasks that are running in the TS7700 Virtualization Engine server
- ▶ A description of a TS7700 Virtualization Engine monitoring and performance evaluation methodology
- ▶ A walkthrough of a TS7700 Virtualization Engine capacity planning case study
- ▶ A review of BVIR and VEHSTATS reporting

We discuss the TS7700 Virtualization Engine shared resources so that you can understand the impact that contention for these resources has on the performance of the TS7700 Virtualization Engine.

The monitoring section can help you understand the performance related data recorded in the TS7700 Virtualization Engine. We also discuss the performance issues that might arise with the TS7700 Virtualization Engine. This chapter can also help you recognize the symptoms that indicate that the TS7700 Virtualization Engine configuration is at or near its maximum performance capability. The information provided can help you evaluate the options available to improve the throughput and performance of the TS7700 Virtualization Engine.

The capacity planning case study illustrates some guidelines and techniques for the management of virtual and stacked volumes associated with the TS7700 Virtualization Engine.

9.1 TS7700 Virtualization Engine performance characteristics

The TS7700 Virtualization Engine can provide significant benefits to a tape processing environment. In general, performance depends on such factors as total system configuration, Tape Volume Cache (TVC) capacity, the number of physical tape drives available to the TS7700 Virtualization Engine, the number of channels, the read/write ratio, and data characteristics, such as blocksize and mount pattern. These performance factors and numbers for different configurations are shown in the following pages.

Based on initial modeling and measurements, and assuming a 2.66:1 compression ratio, Figure 9-1 shows the evolution in the write performance with TS7700 family, which is also described in more detail in the *IBM System Storage TS7700 Virtualization Engine TS7720 and TS7740 Performance* white paper. This paper and the most recently published performance white papers are available on the Techdocs website by searching for TS7700 Virtualization Engine at the following address:

<http://w3.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101674>




Figure 9-1 VTS and TS7700 maximum write bandwidth

Figure 9-1 shows the evolution of performance in the TS7700 IBM Virtualization Engine family compared with the previous member of IBM Tape Virtualization, the IBM Virtual Tape Server (VTS). The numbers were obtained from runs with 128 concurrent jobs, each job writing 800 MB (uncompressed) using 32 KB blocks, number of buffers (BUFNO) used by QSAM is twenty (QSAM BUFNO = 20).

Figure 9-2 shows the read performance numbers in a similar fashion.




Figure 9-2 VTS and TS7700 maximum read hit bandwidth

The numbers shown in Figure 9-2 were obtained with 128 concurrent jobs in all runs, each job reading 800 MB (uncompressed) using 32 KB blocks, QSAM BUFNO = 20.

Compared with the VTS, the TS7700 has introduced faster performing hardware components (such as faster FICON channel adapters, the more powerful TS7700 Virtualization Engine controller, and faster disk cache) along with the new TS7700 Virtualization Engine architecture, providing for improved performance and throughput characteristics of the TS7700 Virtualization Engine. From a performance aspect, important characteristics of the architecture are as follows:

- ▶ With the selection of DB2 as the central repository, the TS7700 Virtualization Engine provides a standard SQL interface to the data and all data is stored and managed in one place. DB2 also allows for more control over back-end performance.
- ▶ The cluster design with vNode and hNode provides increased configuration flexibility over the monolithic design of the Virtual Tape Server.
- ▶ The use of TCP/IP instead of FICON for site-to-site communication eliminates the requirement to use channel extenders.

9.1.1 Performance overview

This section describes several performance attributes of the TS7700.

Performance workloads and metrics

Performance charts in this chapter have been derived from measurements that generally attempt to simulate common user environments, namely a large number of jobs writing or reading multiple tape volumes simultaneously. Unless otherwise noted, all measurements

were made with 128 simultaneously active virtual tape jobs. Each tape job was writing or reading 800 MB of uncompressed data using 32 KB blocks and QSAM BUFNO = 20, where KB = 1024 bytes, using data that compresses at the TS7700 at 2.66:1. Stand-alone measurements were made with four 4-gigabit (Gb) FICON channels on a z10 host. Grid measurements were made with four 2-Gb FICON channels on a z990 host. All runs are begun with the virtual tape subsystem inactive. Unless otherwise stated, all runs were made with default tuning values (DCOPYT = 125, PMPIOR = 800, PMTHLV L= 1000, ICOPYT = ENABLED, Reclaim disabled, Number of premigration drives = Number of installed drives -1). See the *IBM Virtualization Engine TS7700 Series Best Practices - Understanding, Monitoring and Tuning the TS7700 Performance* white paper for detailed description of different tuning settings:

<http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101465>

Types of throughput

Because the TS7720 is a disk-cache-only cluster, its read and write data rates have been found to be fairly consistent throughout a given workload. Because the TS7740 contains physical tapes to which the cache data is periodically written, the TS7740 has been found to exhibit four basic throughput rates:

- ▶ Peak write
- ▶ Sustained write
- ▶ Read-hit see , “Read-hit and recall throughput” on page 592
- ▶ Recall

These four rates are described next.

Peak and sustained write throughput

For the TS7740, a measurement is not begun until all previously written data has been copied, or premigrated, from the disk cache to physical tape. Starting with this initial condition, data from the host is first written into the TS7740 disk cache with little if any premigration activity taking place. This approach allows for a higher initial data rate, and is termed the *peak* data rate. After a pre-established threshold is reached of non-premigrated data, the amount of premigration is increased, which can reduce the host write data rate. This threshold is called the premigration priority threshold, and has default value of 800 GB. When a further threshold of non-premigrated data is reached, the incoming host activity is actively throttled to allow for increased premigration activity. This throttling mechanism operates to achieve a balance between the amount of data coming in from the host and the amount of data being copied to physical tape. The resulting data rate for this mode of behavior is called the “sustained” data rate, and could theoretically continue on forever, given a constant supply of logical and physical scratch tapes. This second threshold is called the *premigration throttling threshold*, and has a default value of one terabyte (TB). These two thresholds can be used in conjunction with the peak data rate to project the duration of the peak period. Note that both the priority and throttling thresholds can be increased through a host command line request, described later in this paper.

Read-hit and recall throughput

Similar to write activity, there are two types of TS7740 read performance:

- ▶ *Read-hit* (also referred to as *peak*) occurs when the data requested by the host is currently in the disk cache.
- ▶ *Recall* (also referred to as *read-miss*) occurs when the data requested is no longer in the disk cache and must be first read in from physical tape.

Read-hit data rates are typically higher than recall data rates.

Summary

The two read performance metrics, along with peak and sustained write performance are sometimes referred to as the *four corners* of virtual tape performance. The charts in the book show three of the four corners:

- ▶ Peak write
- ▶ Sustained write
- ▶ Read hit

Recall performance is dependent on several factors that can vary greatly from installation to installation, such as number of physical tape drives, spread of requested logical volumes over physical volumes, location of the logical volumes on the physical volumes, and length of the physical media. Laboratory measurements have indicated that recall performance with TS7740 is similar to that for VTS.

Grid considerations

Up to four TS7700 clusters can be linked together to form a grid configuration. The connection between these clusters is provided by two 1-Gb TCP/IP links. Data written to one TS7700 cluster can be optionally copied to the other cluster (or clusters). Up to four TS7700 clusters can be host-driven, depending on individual customer requirements.

Data can be copied between the clusters in either deferred or RUN (also known as *Immediate*) copy mode. In RUN copy mode the rewind-unload at job end is held up until the received data is copied to the other cluster (or clusters). In deferred copy mode data is queued for copying, but the copy does not have to occur prior to job end. Deferred copy mode allows for a temporarily higher data rate than RUN copy mode, and can be useful for meeting peak workload demands. Be sure, however, that there is sufficient recovery time for deferred copy mode so that the deferred copies can be completed prior to the next peak demand.

Deferred copies are controlled during heavy host I/O with a default setting of 125 ms for the Deferred Copy Throttling (DCT). More priority can be given to deferred copying by lowering the DCT value.

Performance hardware configuration

Table 9-1 lists the hardware that was used in performance measurements. Except stated otherwise, all measurements were made using the maximum resources available in terms of cache size and tape drives.

Table 9-1 Performance measurements setup

TS7700	Model	Cache drawers	Disk drive	Tape drives	IBM host
TS7740 (Stand alone)	3956 CC8/CX7	4	600 GB	12 TS1130	z10 with four 4-GB FICON
TS7720 (Stand alone)	3956 CS8/XS7	7	2 TB	N/A	z10 with four 4-GB FICON
TS7740 (Grid)	3956 CC8/CX7	4	600 GB	8 TS1120	z990 with four 2-GB FICON
TS7720 (Grid)	3956 CS8/XS7	7	2 TB	N/A	z990 with four 2-GB FICON

9.1.2 Single cluster performance

Besides the CPU processor, two important factors that affect performance are workloads and cache capacity. The following graphs show TS7700 single cluster, or stand-alone, performing under various workloads and using various cache configurations.

Figure 9-3 shows the stand-alone bandwidth achieved by TS7720 Virtualization Engine with release 1.7. The chart compares performance with various cache sizes, using the configurations introduced in R1.7.




Figure 9-3 TS7720 3956 CS8/XS7 stand-alone write or read bandwidth

Figure 9-4 on page 595 shows the numbers for the TS7740 stand-alone cluster. The performance measurements were run with various workloads and cache sizes, using the cache configurations introduced in R1.7.




Figure 9-4 TS7740 3956 CC8/CX7 stand-alone write or read bandwidth

All runs in the two previous charts were performed with 128 concurrent jobs, each job writing 800 MB (uncompressed) using 32 KB blocks, QSAM BUFNO = 20 and four 4-GB FICON channels from a z10 LPAR.

for more details about single cluster performance characteristics, see the most recently published white paper, *IBM System Storage TS7700 Virtualization Engine TS7720 and TS7740 Performance*, which is available at the Techdocs website; search for TS7700 Virtualization Engine at the following address:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs>

9.1.3 Multi-cluster grid performance

The graphs in this section show the performance by hybrid two-cluster grids, which are used in two ways:

- ▶ Grid H1 (see Figure 9-5 on page 596) has a TS7720 in the production site, connected to the remote TS7740 in a disaster recovery configuration. Only TS7720 has host connections and active logical tape drives.
- ▶ Grid H2 (see Figure 9-5 on page 596) is used in a high availability mode, both TS7720 and TS7740 have active host connections and logical tape drives varied online.

Figure 9-5 shows the grids used in the performance test for the two-cluster hybrid grid.




Figure 9-5 TS7700 Hybrid two-cluster grids H1 and H2

Figure 9-6 shows the resultant numbers.




Figure 9-6 Two-way TS7700 hybrid grid (TS7720 3956 CS8/XS7 and TS7740 3956 CC8/CX7) maximum bandwidth

All runs were made with 128 concurrent jobs. Each job writing or reading was 800 MB (uncompressed) using 32-KB block size, data compressed at the TS7700 at 2.66:1, QSAM BUFNO=20, and z990 host with four 2-Gb FICON channels. Clusters are located at zero or near zero distance to each other in the laboratory configuration.

Note: At the time of writing this book, no performance numbers were available for three-cluster and four-cluster grids with R1.7 Licensed Microcode. We kept the graphs for three-cluster and four-cluster from previous R1.6 only as a reference. Some numbers might have changed with the new code. See the *IBM System Storage TS7700 Virtualization Engine TS7720 and TS7740 Performance* paper, available at the Techdocs website for the latest information. Search for TS7700 Virtualization Engine at the following address:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs>

Figure 9-7 shows the TS7740 Virtualization Engine three-cluster grid maximum bandwidth for dual-active clusters (R1.6).




Figure 9-7 R1.6 TS7700 three-clusters hybrid grid maximum bandwidth for dual active clusters

The TS7740s used were TS7740s 3956 CC7/CX7, which had higher throughput than the older model of TS7740 3956 CC6/CX6. All runs were made with 128 concurrent jobs. Each job writing or reading was 800 MB (uncompressed) using 32-KB block size, data compressed at the TS7700 at 2.66:1, QSAM BUFNO=20, and z990 host with four 2-Gb FICON channels. Clusters are located at zero or near zero distance to each other in laboratory setup.

Figure 9-8 shows the numbers for a four-cluster grid with dual active clusters (R1.6).




Figure 9-8 R1.6 TS7700 four-cluster grid Maximum Bandwidth for Dual Active Cluster

The TS7740s in the grid were TS7740s 3956 CC7/CX7. All runs were made with 128 concurrent jobs. Each job writing or reading was 800 MB (uncompressed) using 32 KB block size, data compressed at the TS7700 at 2.66:1, QSAM BUFNO=20, and z990 host with four 2-Gb FICON channels. Clusters are located at zero or near zero distance to each other in laboratory setup.

Figure 9-9 shows the TS7700 four-cluster grid with quadruple active clusters (R1.6).




Figure 9-9 R1.6 TS7700 four-cluster hybrid grid maximum bandwidth for quadruple active clusters

The TS7740s used in the Grid were TS7740s 3956 CC7/CX7. All runs were made with 128 concurrent jobs. Each job writing or reading was 800 MB (uncompressed) using 32-KB block size, data compressed at the TS7700 at 2.66:1, QSAM BUFNO=20, and z990 host with four 2-Gb FICON channels. Clusters are located at zero or near zero distance to each other in laboratory setup.

9.1.4 Distance and performance

Several factors, among others, can affect Grid performance:

- ▶ Distance
- ▶ Quality of links
- ▶ Number of utilized TCP/IP links (1 or 2)
- ▶ Percentage of link bandwidth available,
- ▶ Type of I/O activity (write, read, or copy)

The testing procedures use simulated delays to show the impact of placing the two clusters many miles from each other. Figure 9-10 on page 601 shows the remote copy performance simulating different distances, charted together with the packet loss effect over performance. Data packet loss is directly related to the quality of the Grid links.




Figure 9-10 Single direction copy at simulated distances and percents of package loss (R1.6)

To better understand the chart, a rough principle is that every millisecond of time delay corresponds to 100 miles (160 km) of distance. For example, the light blue bar corresponds to 500 miles in the chart (800 km).

Figure 9-11 shows the same correlation for remote read.




Figure 9-11 Single direction remote read performance at simulated distances and percents of packet loss (R1.6)

Similarly, Figure 9-12 shows the remote write performance under the same conditions.




Figure 9-12 Single direction remote write performance at simulated distances and percents of packet loss (R1.6)

9.1.5 TS7740 tape drives and overall cluster performance

The physical tape drives are managed by the TS7740 Virtualization Engine internal management software and cannot be accessed from any other attached host. These drives are used exclusively by the TS7740 Virtualization Engine for the mounts required for copying virtual volumes to stacked volumes, recalling virtual volumes into the cache, and reclaiming stacked volume space.

The availability of TS7740 Virtualization Engine physical tape drives for certain functions can significantly affect TS7740 Virtualization Engine performance. The TS7740 Virtualization Engine manages the internal allocation of these drives as required for various functions, but it always reserves at least one physical drive for recall and one drive for premigration.

Tape Volume Cache management algorithms also influence the allocation of back-end physical tape drives, as described in the following examples:

- ▶ Cache freespace low: The TS7740 Virtualization Engine increases the number of drives available to the pr-migration function and reduces the number of drives available for recalls.
- ▶ Premigration threshold crossed: The TS7740 Virtualization Engine reduces the number of drives available for recall down to a minimum of one drive to make drives available for the premigration function.

The number of drives available for recall or copy are also reduced during reclamation.

If the number of drives available for premigration is restricted, this can lead to limiting the number of virtual volumes in the cache that are eligible to be migrated, and which can then lead to free space or copy queue throttling being applied.

If the number of drives for recall is restricted, this can lead to elongated virtual mount times for logical volumes being recalled.

Recall performance is highly dependent on both the placement of the recalled logical volumes on stacked volumes and the order in which the logical volumes are recalled. To minimize the effects of volume pooling on sustained write performance, volumes are premigrated using a different distribution algorithm.

This algorithm chains several volumes together on the same stacked volume for the same pool. This can change recall performance, sometimes making it better, sometimes making it worse. Other than variations in performance because of differences in distribution over the stacked volumes, recall performance should be constant.

Reclaim policies must be set in the management interface (MI) for each volume pool. Reclamation occupies drives and can affect performance. The Inhibit Reclaim Schedule is also set from the MI and can prevent reclamation from running during specified time frames during the week. If Secure Data Erase is used, fewer physical tape drives might be available even during times when you use inhibited reclamation. If used, limit it to a specific group of data. Inhibit Reclaim specifications only partially apply to Secure Data Erase. Secure Data Erase does not honor your settings and thus can run erasure operations as long as there are physical volumes to be erased.

The use of Copy Export and Selective Dual Copy also increases the use of physical tape drives. Both are used to create two copies of a logical volume in a TS7740 Virtualization Engine.

Figure 9-13 on page 603 shows how the sustained write data rate can be affected by the back-end physical tape drives. The chart shows sustained write data rates achieved in the laboratory for a stand-alone TS7740 with different numbers of TS1130 backend physical tape drives. The data for the chart was measured with no TS7740 activity other than the sustained writes and premigration (host write balanced with premigration to tape).




Figure 9-13 TS7740 (3956 CC7/CC8) stand-alone sustained write versus number of online drives

All runs were made with 128 concurrent jobs. Each job wrote 800 MB (uncompressed) using 32KB blocks, data compression 2.66 to 1, QSAM BUFNO = 20, and four 4-Gb FICON channels from a z10 LPAR.

Figure 9-14 on page 604 shows the TS7740 premigration rates, so that the rates at which cache-resident data is copied to physical tapes, and how it depends on the number of drives available for premigration.

Note: No TS7740 activity other than premigration is measured in this chart.




Figure 9-14 TS7740 (3956 CC7/CX7) stand-alone premigration rate versus premigration drives

All runs were made with 128 concurrent jobs. Each job wrote 800 MB (uncompressed) using 32-KB blocks, data compression 2.66:1, QSAM BUFNO = 20, and four 4-Gb FICON channels from a z10 LPAR.

9.2 TS7700 components and tasks distribution

In the process of writing scratch volumes or premigrating and recalling virtual volumes from physical stacked volumes, hardware components are shared by tasks running on the TS7700 Virtualization Engine. Some of these tasks represent users' work (such as scratch mounts) and other tasks are associated with the internal operations of the TS7700 Virtualization Engine (such as reclamation in a TS7740 Virtualization Engine). All these tasks will be sharing the same resources especially the TS7700 Virtualization Engine Server processor, the Tape Volume Cache (TVC), and the physical tape drives attached to a TS7740 Virtualization Engine. Contention might occur for these resources when high workload

demands are placed on the TS7700 Virtualization Engine. To manage the use of shared resources, the TS7700 Virtualization Engine uses various resource management algorithms, which can have a significant impact on the level of performance achieved for a specific workload.

In this section, we discuss the effects on performance of the following shared resources:

- ▶ TS7700 Virtualization Engine processor
- ▶ Tape Volume Cache (TVC)
- ▶ TS7740 Virtualization Engine physical tape drives
- ▶ TS7740 Virtualization Engine physical stacked volumes

See Figure 9-15 on page 605 for this discussion. We look at the tasks that TS7700 has to perform, and correlate the tasks to the components involved. Also, we look at the various available tuning points that can be used to favor certain tasks over others.

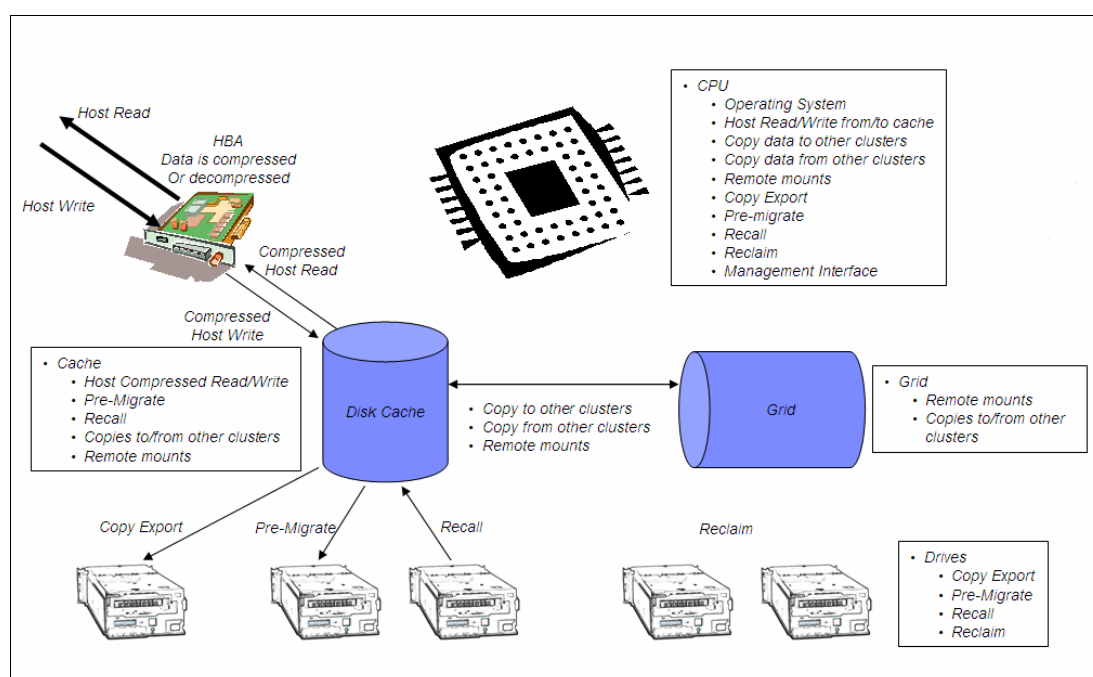


Figure 9-15 TS7700 components and Tasks

9.2.1 Tasks performed by CPU (processor cycles)

All tasks running in the TS7700 Virtualization Engine server require a portion of processor cycles. From one cluster's perspective, these tasks includes:

- ▶ Operating system: TS7700 runs on a AIX® operating system.
- ▶ Host read and write data: Customer data being written to or being read from TS7700 cache. This data is compressed coming into the TS7700 and decompressed going back to host by the host bus adapters (HBAs). Internally, data passes from HBAs through buffers in the CPU memory, and from the memory to the Fibre Channel Protocol adapter (FCP) in its way to the cache. The data rate is limited by the number of Performance Increments installed. Also, the Host Write Throttle affects the read/write rate.
- ▶ Copy data to other clusters: This task is present in a multi-cluster grid, when data is copied from this cluster's cache to other clusters as either immediate (RUN) or deferred copies. Deferred copies are influenced by the deferred copy throttle (DCT), also called deferred copy read throttle. Immediate copies are not throttled by the originating cluster.

- ▶ Copy data from others clusters: this is the data copied into this cluster's cache from other clusters as either immediate (RUN) or deferred copies. Both types of copies are regulated by the copy throttle.
- ▶ Remote mounts: Either from another cluster accessing one volume in this cluster's cache or this cluster accessing remote volumes in others cluster's cache. Host write throttle has control over remote mount for write from another cluster to this cluster. Remote mounts do not pass data through the cache of the cluster whose virtual logical device was allocated for the mount.
- ▶ Copy export (TS7740 only): When a copy export occurs, data to be exported that resides only in cache is premigrated to physical tape. Also, each physical tape to be exported is mounted and a snapshot of the TS7740 database is written to it. Avoid to request copy export within your heavy production period.
- ▶ Premigrate (TS7740 only): Logical volumes that resides only in cache are written to physical tape. Premigration includes both primary and secondary copies. There are several algorithms in premigration:
 - Idle Premigration
 - Fast Host Write Premigration
 - Somewhat Busy Premigration Ramp Up
 - Preferred Premigration
 - Limit number of drives for premigration in physical volume pools definitions
- ▶ Recall (TS7740 only): Bringing data from a physical tape to cache to satisfy a mount.
- ▶ Reclaim (TS7740 only): Transferring logical volumes from a physical tape to another. This data passes through the CPU's memory; cache is not used in this operation. Reclaim is influenced by the number of available drives, by the Reclaim Threshold and the Inhibit Reclaim Schedule.
- ▶ Management interface: The management interface (MI) is a task that consumes CPU power. The MI is used to configure, operate and monitor the TS7700.

9.2.2 Tasks performed by the TS7700 cache

The TS7700 cache is the focal point for most data transfers, everything except reclaim activity. When evaluating performance, remember that the cache has a finite size and I/O bandwidth to satisfy a variety of tasks. Equally important to remember is that all data moved within the TS7700 is compressed host data. The Host Bus Adapters compress and decompress the data as it is written and read by the host.

- ▶ Host read/write
- ▶ Premigrate
- ▶ Recall
- ▶ Copies to/from other clusters
- ▶ Remote mounts

Note: Remote mounts to other clusters do not move data through local cache. Also Reclaim data does not move through the cache.

9.2.3 Tasks performed by the Grid

The Grid is not a physical component that can be clearly identified, as we can do with the TS7700 Virtual Engine (CPU) or the TS7700 Cache. A functional module runs within the TS7700, using internal resources such as CPU and memory, and external resources like the gigabit link infrastructure. That functional module is responsible for the following items:

- ▶ Copies to/from other clusters
 - Immediate copies
 - Deferred copies
- ▶ Remote Mounts
 - Using another cluster's cache
 - Another cluster using this cluster's cache
- ▶ Cluster Coordination Traffic
 - Ownership takeover
 - Volume Attribute changes
 - Insert of logical volumes
 - Configuration

9.2.4 Tasks performed by the TS7740 Tape Drives

The TS7740 physical back-end tape drives read and write data for a variety of reasons. The back-end drives are shared among these tasks. Balancing their usage is controlled by default algorithms that can be adjusted somewhat by the user. It is important to monitor the use of the physical drives ensuring there are enough back-end drives to handle the workload:

- ▶ Copy Export
 - Volumes that have not being premigrated are copied to physical tape
 - Writes the TS7740 database to each exported volume
- ▶ Premigrate
 - Logical volumes are copied from cache to physical tape
 - Primary and secondary copies
 - Secondary copies can be copy export volumes
- ▶ Recall
 - Non-fast-ready mount
 - Whole logical volume read into cache
- ▶ Reclaim
 - Two drives per task
 - Runs to completion (source empty)

9.3 Throttling, tasks, and knobs

The main goal of this section is to help you understand how data flows within the TS7700, how the internal resources are used by the regular tasks, where the limits are and how to adjust the available tuning points, metaphorically knobs, to get the maximum performance from your TS7700, considering your variables and particular needs.

To make a good use of the panoply of tunable parameters available to you, you should have a good understanding of the types of throttling within the TS7700, and the underlying mechanisms.

Throttling is the mechanism adopted to control and balance the several tasks which run at the same time within the TS7700, prioritizing some tasks over others. These mechanisms are called upon only when the system reaches higher levels of utilization, when the components are used near to the fullest of its capacity and the bottleneck starts to show up. The criteria is based in what is required from the TS7700 by the user in one side balanced to what is vital to

the TS7700 functioning in the other side. This control is accomplished by holding the launch of new tasks, prioritizing the more important tasks over the others. After the tasks are dispatched and running, some control over the execution is accomplished by slowing down an specific functional area, introducing calculated amounts of delay in the operations, to give some alleviation to an overloaded component or leave some extra CPU cycles to another needed function, or yet await for a slower operation to finish.

In this section, we explain the throttling algorithms and the control knobs that can be customized by you to best suit the TS7700 behavior to your particular needs.

9.3.1 Throttling in the TS7700

The TS7700 subsystem aggregated I/O rate is the sum of all data transfers (reads or writes) that are performed by all active logical drives mounted at a certain time in a TS7700. Generally speaking, each logical drive that is mounted and transferring data gets a proportional share of the total data throughput. For instance, if there are 50 mounted and active virtual drives, which one will receive, in average, two percent of the current host data transfer. Important factors that can impact host data transfer bandwidth are the so-called housekeeping tasks, such as premigrating, immediate and deferred copies, copy export tasks, reclaim activities. Also, in the other side, how much data the host is attempting to drive into and out of the TS7700.

The subsystem has a series of self regulatory mechanisms that tries to optimize the shared resources usage. Subsystem resources, such as CPU, cache bandwidth, cache size, host channel bandwidth, grid network bandwidth, physical drives, and so forth are limited, and must be shared by all tasks moving data throughout the subsystem.

The resources will implicitly throttle by themselves when reaching their limits. The TS7700 introduces a variety of explicitly throttling methods to give the higher priority tasks more of the shared resources. The following list prioritizes normally running tasks that move data:

1. Immediate copies
2. Recalls
3. Copy export
4. Host I/O
5. Reclaims
6. Premigration
7. Deferred copies

TS7700 can get into certain situations that might require special treatment, being granted a higher priority. Examples are as follows:

- ▶ Panic reclamation: The TS7740 detects the number of empty physical volumes has dropped below the minimum value and reclaims need to be done immediately to increase the count.
- ▶ Cache fills with copy data: To protect from having un-copied volumes removed from cache the TS7740 throttles data coming into the cache.
- ▶ Cache overfills: If no more data can be placed into the cache before data is removed then other tasks trying to add to the cache are heavily throttled.

9.3.2 Host Write Throttle

This mechanism is applied to limit the amount of data written into cache from the host. It acts upon host incoming write operations through the host bus adapters, and write operations coming over the Grid interface because of mounts from other clusters. This throttle is triggered by the following items:

- ▶ Full Cache: A full cache of data is to be copied
- ▶ Immediate Copy: Large amounts of immediate copy data needs to be moved
- ▶ Premigrate: Large Amounts of data in cache has not been written to tape yet
- ▶ Free Space: Cache is nearly full of any kind of data.

Figure 9-16 is a visual representation of the Host Write Throttle mechanism and where it applies.

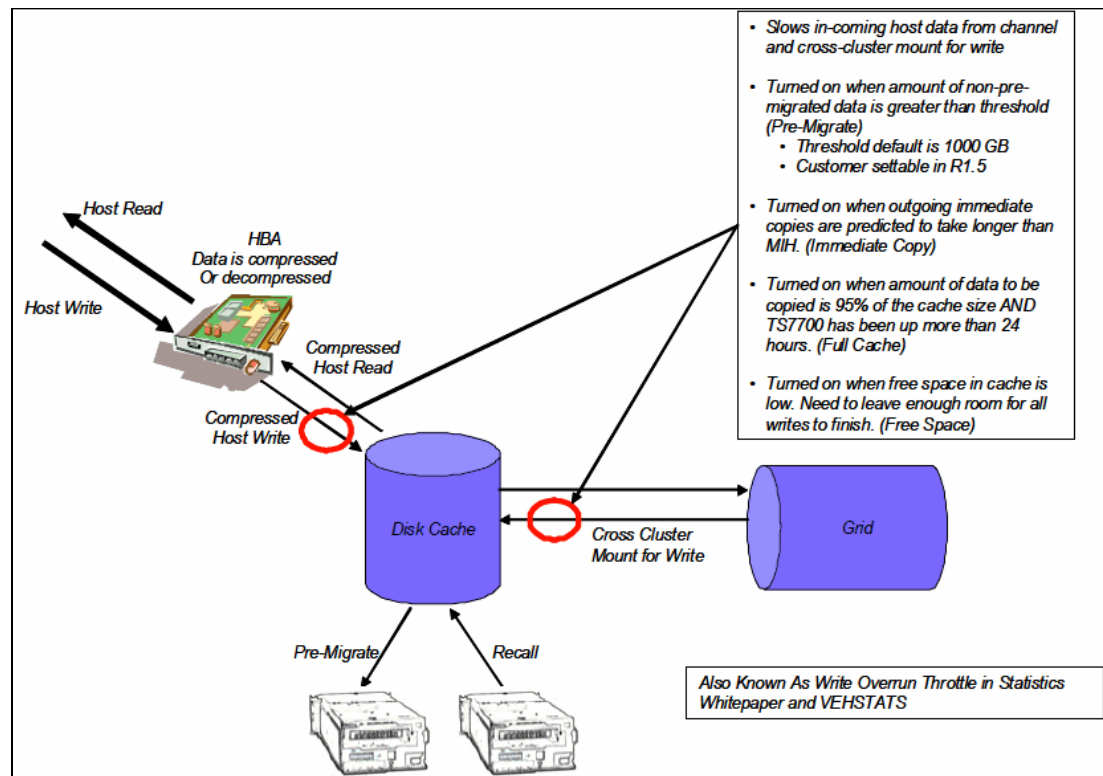


Figure 9-16 Host Write Throttle

9.3.3 Copy Throttle

Copy Throttle is used to limit the amount of data being written into cache from other clusters's copy data. This mechanism throttles incoming copies, both immediate and deferred, that come from other clusters in the Grid.

Copy throttling is triggered by the following items:

- ▶ Full Cache
- ▶ Amount of data to be premigrated

Figure 9-17 shows the mechanism and where it applies.

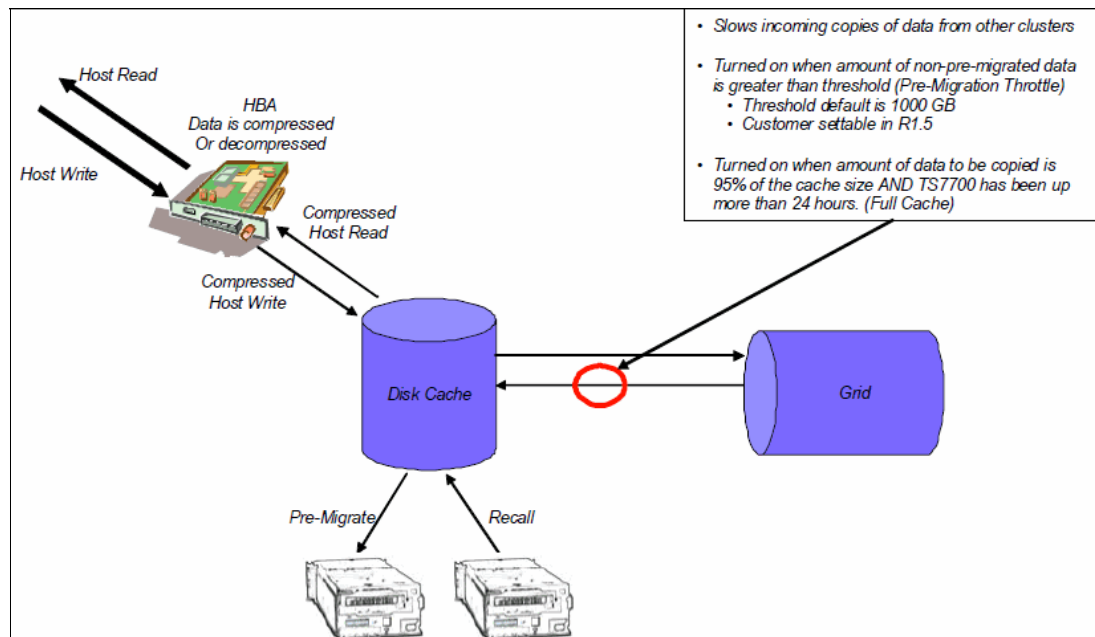


Figure 9-17 Copy Throttle

Host write throttle and copy throttle are triggered by the same factors, as follows:

- ▶ Full cache: Cache is full of data that needs to be copied to another cluster.
 - Amount of data to be copied to another cluster is greater than 95% of cache size *and* the TS7700 has been up more than 24 hours.
 - Full Cache is reported as Write Throttle and Copy Throttle in VEHSTATS.
- ▶ Immediate copy: Immediate copies to other clusters, where this cluster is the source, are taking too long or are predicted to take too long.
 - The TS7700 evaluates the need for this throttle every two minutes.
 - The TS7700 examines the depth of the immediate copy queue and the amount of time that the copies have been in the queue to determine if the throttle should be applied, as follows:

The algorithm looks at age of oldest immediate copy in the queue:

- If oldest is 10 - 30 minutes old, the TS7700 sets the throttle in the range of 0.00166 seconds to two seconds, and linear ramp in the range of 10 - 30 minutes.
- The maximum throttle (2 seconds) is applied immediately if an immediate copy has been in the queue for 30 minutes or longer.

Look at quantity of data, and calculate how long the transfer will take.

- If greater than 35 minutes, the TS7700 sets throttle to max (2 seconds).
- If 5 - 35 minutes, the TS7700 sets throttle from .01111 seconds to 2 seconds, linear ramp from 5 to 35 minutes.

- Immediate copy is reported as Write Throttle in VEHSTATS.

Note: The time required for a 4000 MB immediate copy is five times longer than an 800 MB immediate copy.

- Host Write Throttle, because of Immediate Copies taking too long, can be disabled by using the Host Console Request.
- ▶ Premigrate: Amount of data to be premigrated is above the threshold (default 1000 GB)
 - Premigrate is reported as Write Throttle and Copy Throttle in VEHSTATS.
 - These throttle values will be equal if Premigrate is the sole reason for throttling.
- ▶ Free Space: Invoked when cache is near full of any data.
 - Used to make sure there is enough cache to handle the currently mounted volumes.
 - Free Space is reported as Write Throttle in VEHSTATS.

9.3.4 Deferred Copy Throttle

Deferred Copy Throttle (DCT) acts on the rate of deferred copies going to other clusters. This throttle is applied with the intention of saving some CPU cycles to favor host I/O data transfer. Here, you are trading deferred copies done by host I/O data throughput rate. This throttle is also known as Deferred Copy Read Throttle, because other clusters are “reading” data from this cluster. See Figure 9-18.

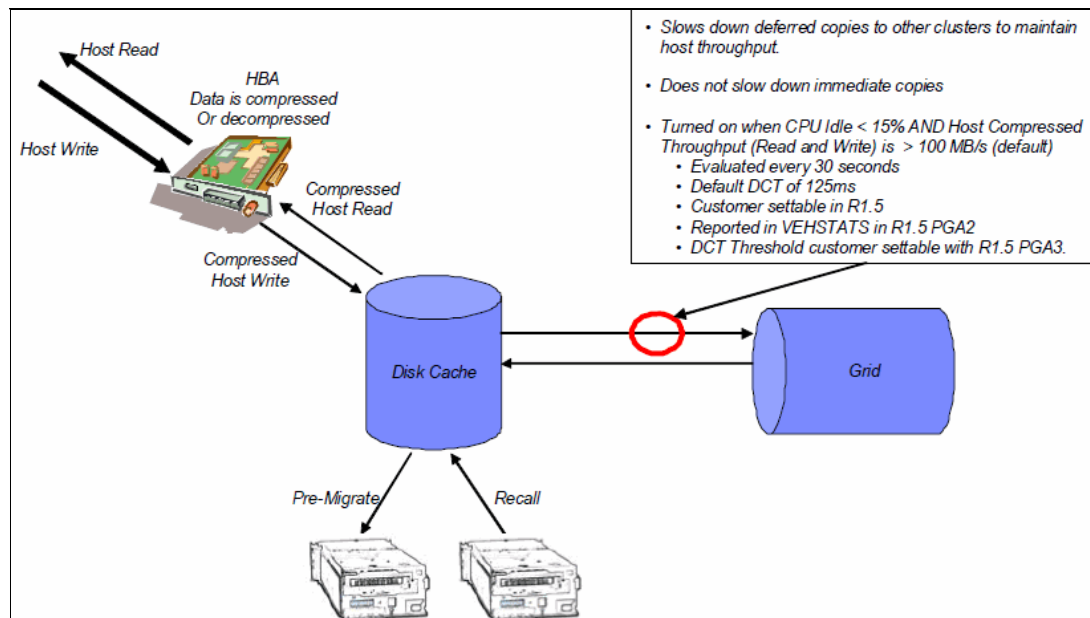


Figure 9-18 Deferred Copy Throttle (Deferred Read Copy Throttle)

Deferred Copy Throttle is triggered by CPU usage. Usage and the host compressed throughput are monitored and evaluated every 30 seconds. DCT is invoked whenever CPU usage goes higher than 85% and the compressed host throughput is bigger than 100 MBps.

Note: The 100 MBps threshold is the default and can be changed by the customer through the Host Console Request

Deferred Copy Throttle remains in effect for the subsequent 30 second interval, after which the TS7700 will reevaluate the scenario.

Note: Default Deferred Copy Throttle value is 125ms.

This 125 ms of throttling if applied severely slow down deferred copy activity, translating to 125 ms being added between each block of 256 KB of data sent through the replication grid for a volume copy.

The DCT can be set by using Host Console Request function. For details about setting DCT, see *IBM Virtualization Engine TS7700 Series z/OS Host Command Line Request User's Guide* which is available on Techdocs website. Use the SETTING, THROTTLE, DCOPYT keywords.

The DCT Threshold can be set by using Host Console Request function. The setting of the DCT is discussed in details in *IBM Virtualization Engine TS7700 Series z/OS Host Command Line Request User's Guide* which is available on Techdocs website. Use the SETTING, THROTTLE, DCTAVGTD keywords.

9.3.5 Managing tasks within TS7700

This section describes more command management tasks with the TS7700.

Prioritizing copies within the grid

In a multi-cluster grid, copies of volumes are made from one cluster to another. Copies are made in either an immediate or deferred manner. There are three classifications of copies: Immediate, Immediate-Deferred, and Deferred. Immediate-deferred are volumes that were originally defined as immediate copies but were changed to deferred copies. An immediate copy can be changed to an immediate-deferred copy if the cluster that is to receive the copy is not online. The TS7700 processes immediate copies as top priority, immediate-deferred as second priority, and deferred as third priority. Volumes in the source cluster are defined as either PG0 or PG1 volumes through the Management Class. Within the three copy classifications, the TS7700 gives priority to PG0 volumes over PG1 volumes. The reason for prioritizing PG0 over PG1 is to allow the PG0 volumes to be removed from cache as soon as possible thus allowing for more PG1 volumes to reside in cache.

The order in which copies are processed is as follows:

1. Immediate PG0
2. Immediate PG1
3. Immediate-Deferred PG0
4. Immediate-Deferred PG1
5. Deferred PG0
6. Deferred PG1

Premigration tasks

The TS7740 uses a variety of criteria to manage the number of premigration tasks. The TS7700 looks at these criteria every five seconds to determine if one more premigration task should be added. Adding a premigration task is based on the following factors and others:

- ▶ Host-compressed write rate
- ▶ CPU activity
- ▶ How much data needs to be premigrated per pool
- ▶ How much data needs to be premigrated in total

Figure 9-19 on page 614 shows migration demonstration. A premigration task does not preempt a recall, reclaim or copy export task. Four algorithms are working in concert to determine whether another premigration task should be started. General details are described below. The actual algorithm has several nuances that are not described here.

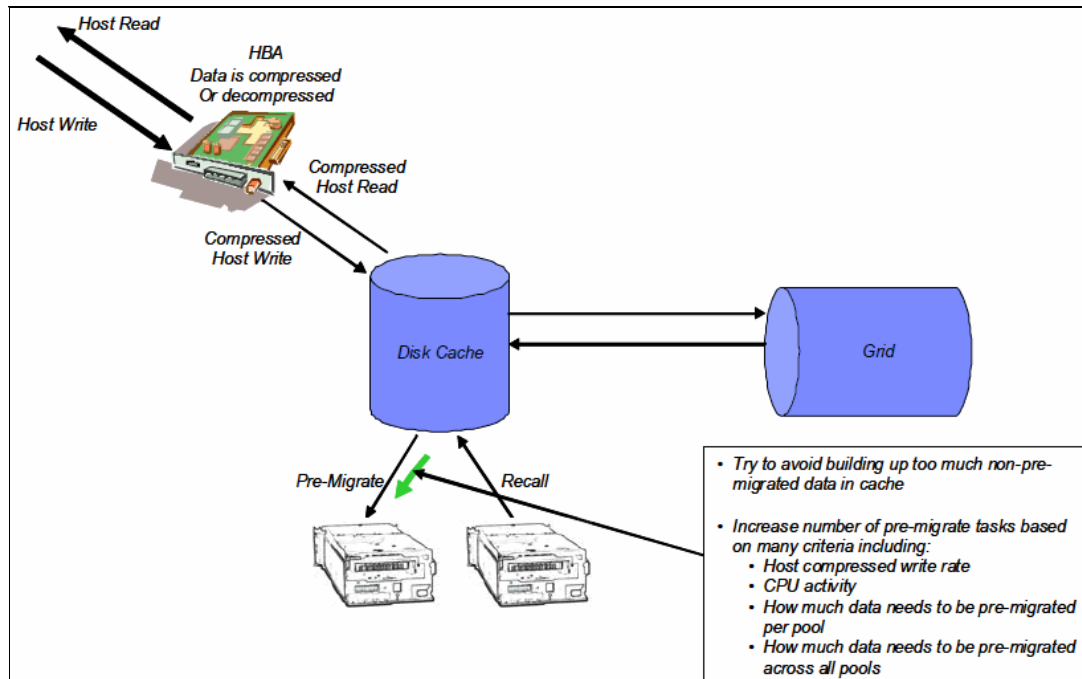


Figure 9-19 Migration demonstration

Note the following factors:

- ▶ Idle premigration
 - If the CPU usage is idle more than 5% of the time, a premigrate task is started, if appropriate.
 - The number of tasks is limited to six or the maximum premigration drives defined by pool properties, whichever is less.
- ▶ Fast host write premigration mode
 - Compressed host write rate is higher than 30 MBps *and* CPU idle time lesser than 1%.
 - Premigration tasks are limited to two, for all pools, and lowered to one or zero if mode continues.
- ▶ Premigration ramp up
 - Compressed host write is less than 30 MBps *and* CPU idle time less than 5%.
 - Indicates the TS7700 is somewhat busy.
 - Is limited to available back-end drives minus one. Also limited by the maximum number of premigrate drives setting in pool properties.
- ▶ Preferred premigration
 - Amount of non premigrated data exceeds the Preferred Premigration threshold. The default threshold is 800 GB.
 - Limited to available back-end drives minus one. Also limited by the maximum number of premigrate drives setting in pool properties.
 - Preferred Premigration takes precedence over the other algorithms.

Immediate-copy set to immediate-deferred state

The goal of an immediate copy is to complete one or more RUN consistency point copies of a logical volume prior to surfacing status of the RUN command to the mounting host. If one or more of these copies cannot complete, the replication state of the targeted volume will enter the immediate-deferred state. The volume will remain in an immediate-deferred state until all of the requested RUN consistency points contain a valid copy. The immediate-deferred volume will replicate with a priority greater than standard deferred copies, but lower than non-deferred immediate copies. There are numerous reasons why a volume might enter the immediate-deferred state. For example, it might not complete within 40 minutes, or, one or more clusters targeted to receive an immediate copy are not available. Independent of why, the host application or job that is associated with the volume is not aware that its previously written data has entered the immediate-deferred state.

From Release 1.6 about the reason why a volume moves to the immediate-deferred state is contained in the ERA 35 sense data. The codes are divided into unexpected and expected reasons. New failure content is introduced into the CCW(RUN) ERA35 sense data:

- ▶ Byte 14 FSM Error: If set to 0x1C (Immediate Copy Failure), the additional new fields are populated.
- ▶ Byte 18 Bits 0:3 – Copies Expected: Indicates how many RUN copies were expected for this volume.
- ▶ Byte 18 Bits 4:7 – Copies Completed: Indicates how many RUN copies were actually verified as successful before surfacing SNS.
- ▶ Byte 19 – Immediate Copy Reason Code:
 - Unexpected – 0x00 to 0x7F: The reasons are based on unexpected failures:
 - 0x01 – A valid source to copy was unavailable
 - 0x02 – Cluster targeted for a RUN copy is not available (unexpected outage).
 - 0x03 – Forty minutes has passed and one or more copies have timed out.
 - 0x04 – Is downgraded to immediate-deferred because of health/state of RUN target clusters.
 - 0x05 – Reason is unknown.
 - Expected – 0x80 to 0xFF: The reasons are based on configuration or a result of planned outages:
 - 0x80 – One or more RUN target clusters are out of physical back end scratch.
 - 0x81 – One or more RUN target TS7720 clusters are low on available cache (95%+ full)
 - 0x82 – One or more RUN target clusters are in service-prep or service.
 - 0x83 – One or more clusters have copies explicitly disabled through the Library Request operation
 - 0x84 – The volume cannot be reconciled and is currently ‘Hot’ against peer clusters.

The additional data contained within the CCW(RUN) ERA35 sense data can be used within a z/OS custom user exit to act on a job moving to the immediate-deferred state. Because the requesting application that results in the mount has already received successful status prior to the issuing of the CCW(RUN), it cannot act on the failed status. However, future jobs can be suspended or other custom operator actions can be taken using the information provided within the sense data.

9.4 Data movement through cache

This section describes the data flow through the TS7700 cache in several configurations. Our goal is to help you understand the data traffic and the pieces of data movement that have to share the resources of the TS7700.

Note: Reclaim data in a TS7740 is transferred from one tape drive to another, not passing through the cache.

We look at TS7720 data flow, because this is the TS7700 basic configuration in terms of cache data flow.

For both the TS7720 and TS7740, the following data flows through the subsystem:

- *Uncompressed host data* is compressed by the Host Bus Adapters (HBAs) and the compressed data is written to cache.
- *Compressed data* is read from the cache and decompressed by the HBA, as shown in Figure 9-20.




Figure 9-20 TS7720 single-cluster data flow

Compare to a TS7740 single-cluster shown in Figure 9-21 on page 617:

- If a read is requested and the logical volume does not exist in the cache, a stacked physical tape is mounted and the logical volume is read into cache. The host then reads the logical volume from the TS7740 cache.
- Host data will be written from cache to the physical stacked volumes in a process called *preigrate*.

The back-end drives account for the difference. In a TS7740, the tape drives are used for recall data from the stacked volumes, and to preigrate data into physical tape.

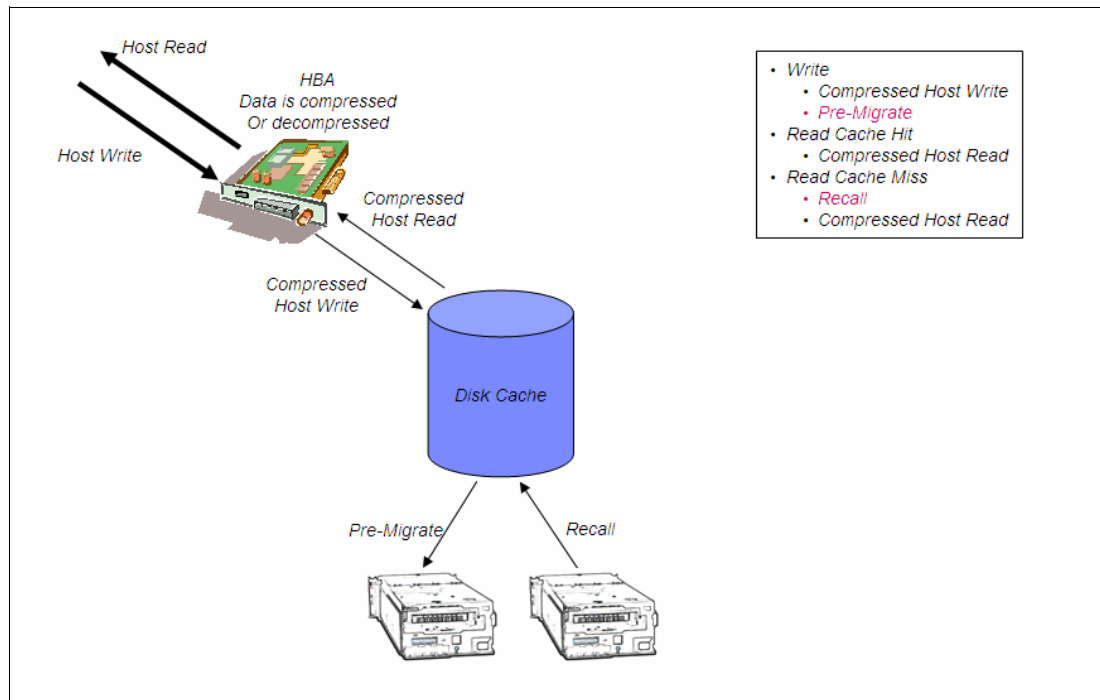


Figure 9-21 TS7740 single-cluster data flow

9.4.1 Cache Data Flow in a TS7700 two-cluster grid

Figure 9-22 shows a TS7720 two-cluster grid being used in balanced mode. host accesses are available to both clusters, and the customer operates with logical drives online in both clusters. For a scratch mount, host allocation can select a virtual device from either cluster. The figure shows data flow from a cache perspective in cluster 0.

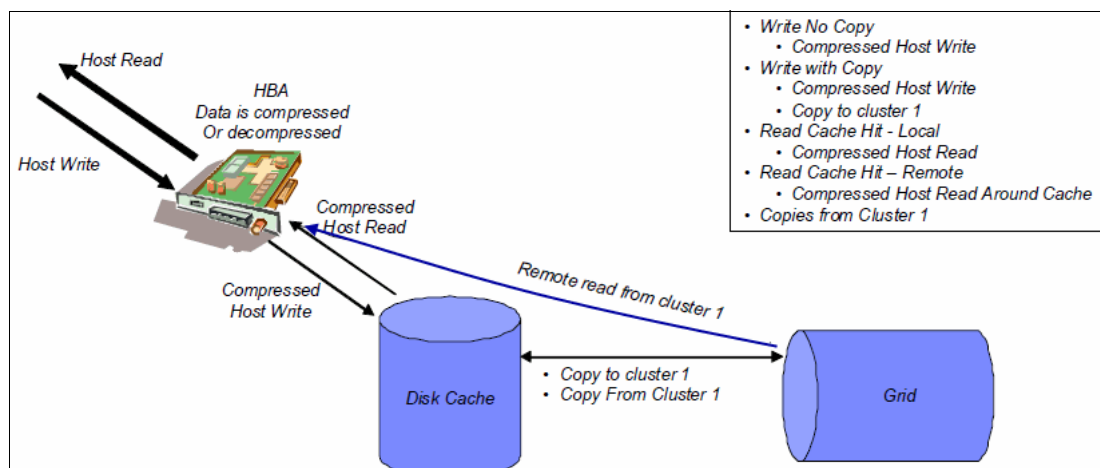


Figure 9-22 TS7720 two-cluster grid: balanced mode

For both the TS7720 and TS7740, the following data is moved through the subsystem:

- ▶ Local write with no remote copy, that is, copy consistency point (CCP) of run-none (RN) includes writing the compressed host data to cache.
- ▶ Local write with remote copy (CCP of DD or RR) includes writing the compressed host data to cache and to the grid:
 - For a CCP of RR the copy is immediate and must complete before Rewind/Unload (RUN) is complete. Copies are placed in the immediate copy queue.
 - For a CCP of DD the copy is deferred where the completion of the RUN is not tied to the completion of the copy operation. Copies are placed in the Deferred Copy Queue.
- ▶ Remote write with no local copy (CCP of NR) includes writing compressed host data to the grid. (Not shown in the picture).
- ▶ Local read with a local cache hit. Here, the compressed host data is read from the local cache.
- ▶ Local read with a remote cache hit. Here, the compressed host data is read from the remote cache through the grid link.
- ▶ Immediate and deferred copies from the remote cluster. Here compressed host data is received on the grid link and copied into the local cache.

For the TS7740, we have the back-end tape drives for recalls and premigrates.

Figure 9-23 shows a similar configuration with TS7740, and shows data flow from the cache perspective in cluster 0.

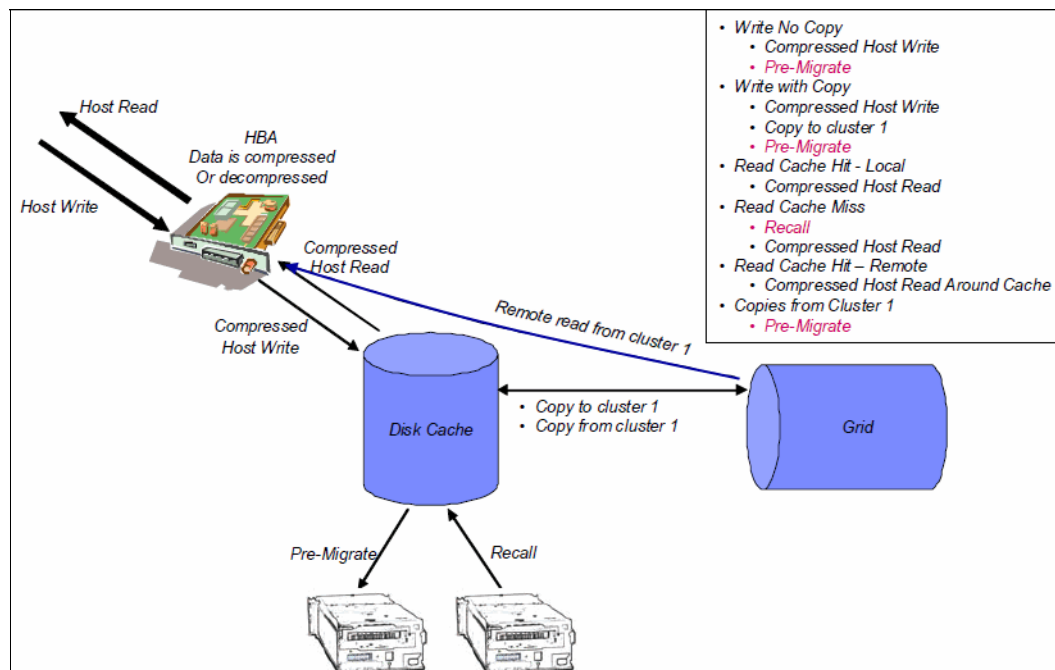


Figure 9-23 TS7740 two-cluster grid: balanced mode

Note the following information:

- ▶ A write with copy or no copy to another cluster includes the premigrate process.
- ▶ A read with local cache miss has one of the following results:
 - A remote mount without recall
 - A recall into local cache from local stacked volume
 - A remote mount requiring recall from remote stacked volume
- ▶ Host data is written from cache to the physical stacked volumes in a premigrate process, including data written as a result of a local mount, volumes copied from other clusters, and a remote mount to this cluster for write (not shown).

The Grid can be used in a preferred mode. Preferred mode means that only one cluster will have the logical drives varied online. Host allocation will select a virtual device only from the cluster with varied on virtual devices. Data movement through the cache in this mode is a subset of the Balanced Mode model.

9.4.2 Cache data flow in a TS7700 three-cluster grid

For this model, we consider a three-cluster grid configuration where two clusters are in the production site, being used in high availability (balanced mode), that is. both clusters in production site operates with their logical drives online, being optioned by the host allocation for mounts. The third cluster is in a remote site, kept in a vault for disaster recovery site, with its virtual devices varied offline.

We look at the data movement through the subsystem. Balanced mode means the host has virtual devices in both clusters varied online. Host allocation can select a virtual device from either cluster. For both the TS7720 and TS7740 the following data is moved through the subsystem:

- ▶ Local write with no remote copy (CCP of RNN) includes writing the compressed host data to cache.
- ▶ Local write with high availability (HA) copy (CCP of RDN or RRN) includes writing the compressed host data to cache and to the grid.
 - For a CCP of RRN the copy to the other HA cluster is immediate and must complete before rewind-unload (RUN) is complete. Copies are placed in the immediate copy queue.
 - For a CCP of DDN the copy to the other HA cluster is deferred where the completion of the RUN is not tied to the completion of the copy operation. Copies are placed in the Deferred Copy Queue.
- ▶ Local write with HA and remote copy (CCP of RRD or RDD/DDD) includes writing the compressed host data to cache and to the grid.
 - For a CCP of RRD, the copy to the HA cluster is immediate. The copy to the remote cluster is deferred. The immediate copy will be sourced from the mounting cluster. The remote copy will be sourced from either of the two HA clusters. The grid link performance and other factors are used to determine which cluster the remote cluster will source the deferred copy from.
 - For a CCP of RDD/DDD, the copies to the other HA cluster and remote cluster are deferred. Each cluster may source the copy from either of the other two clusters. Which cluster has a valid copy, the grid link performance and other factors are used to determine which cluster the two clusters will source the deferred copy from.

- ▶ Remote write with no local copy (CCP of NRN or NNR) includes writing compressed host data to the grid (Not shown in diagram below).
- ▶ Local read with a local cache hit. Here the compressed host data is read from the local cache.
- ▶ Local read with a remote cache hit. Here the compressed host data is read from one of the other two cluster's cache through the grid link.
- ▶ Immediate and deferred copies from the other HA cluster. Here compressed host data is received on the grid link and copied into the local cache.

See Figure 9-24 for the TS7720 representation.

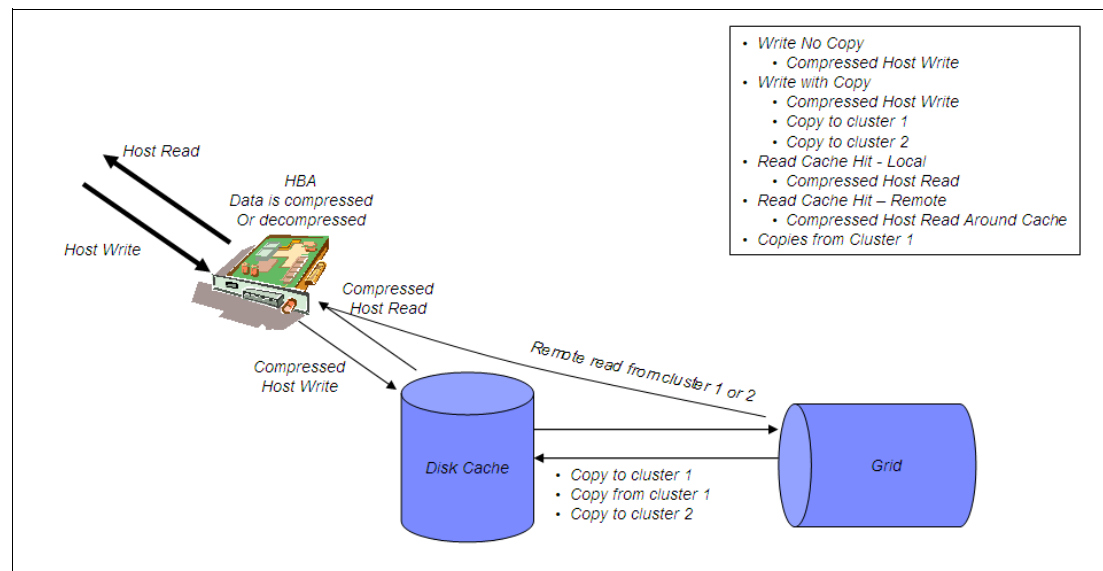


Figure 9-24 TS7720 three-cluster grid, HA, DR Mode

For the TS7740, we add back-end tape drives for recalls and premigrates:

- ▶ A write with no copy and a write with copy also includes the premigrate process.
- ▶ A read with local cache miss has one of the following results:
 - A remote mount without recall
 - A recall into local cache from local stacked volume
 - A remote mount requiring recall from remote stacked volume
- ▶ Host data will be written from cache to the physical stacked volumes in a premigrate process.

This addition includes data written as a result of a local mount for write, volumes copied from other clusters, and a mount for write from the HA cluster (not shown). Figure 9-25 for the TS7740's take on the data moved through cache in this Grid.

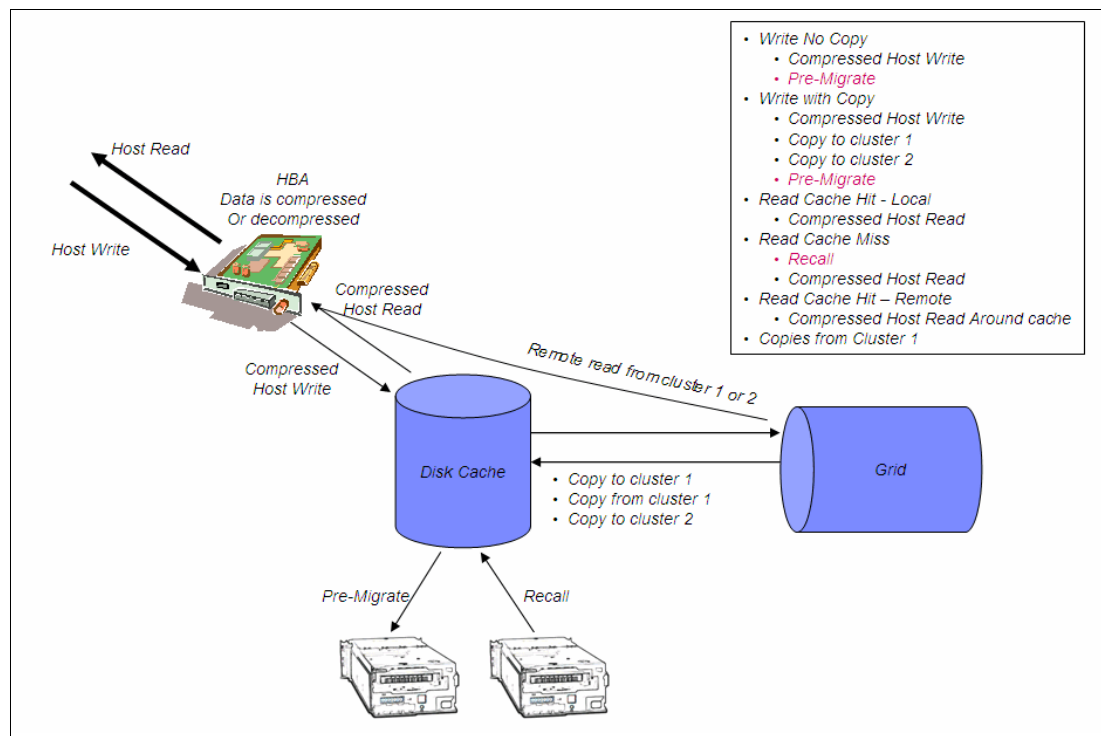


Figure 9-25 TS7740 three-cluster grid, HA, DR Mode

9.4.3 TS7700 Four-cluster grid considerations

With the four-cluster Grid, the copies going back and forth the fourth cluster (Cluster 3) must also be considered. See Figure 9-26 on page 622 for a TS7740 representation in a four-cluster grid.

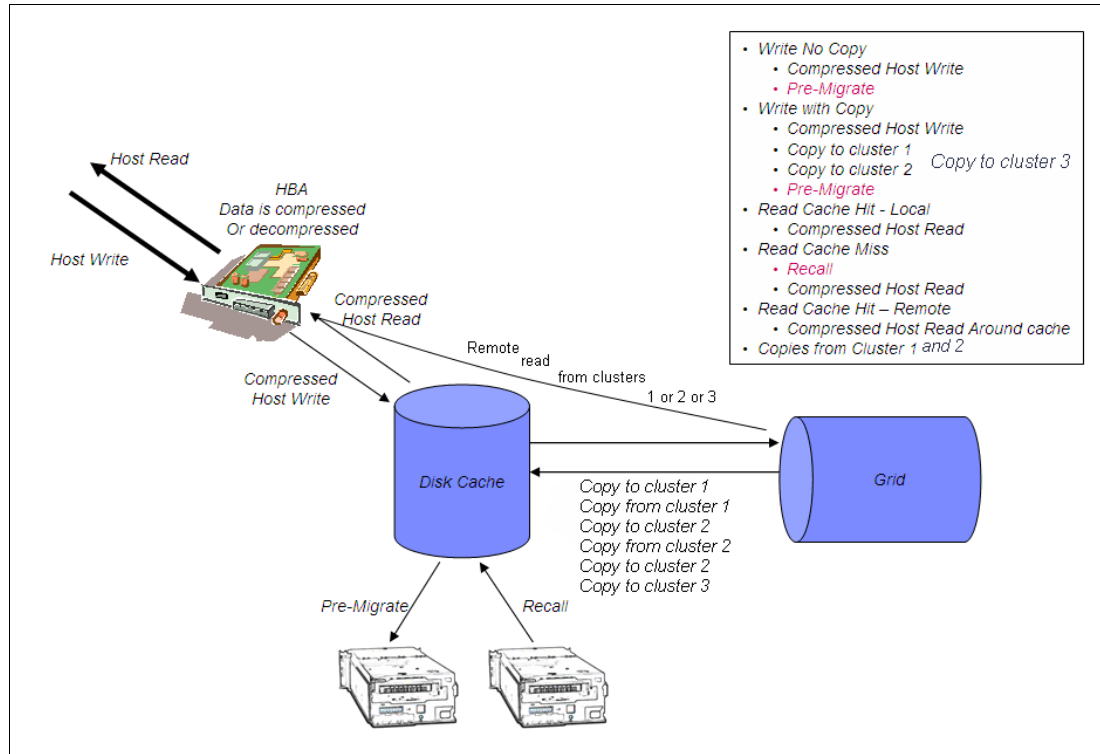


Figure 9-26 TS7740 four-cluster grid configuration, HA, DR mode

One possible configuration for a four-cluster grid is two local clusters (let's say cluster 0 and 1) and two remote clusters (cluster 2 and 3). We can configure the Copy Consistency Points (CCPs) in a way that cluster 0 is replicated to cluster 2, and data written to cluster 1 is replicated to cluster 3. In this way, you will have two two-cluster grid that is within this four-cluster grid. With this configuration, two copies of data are in the grid. All data is accessible by all clusters either within the cluster or through the grid, which means that all data is available when one of the clusters is not available. See Figure 9-27 for the four-cluster grid example.

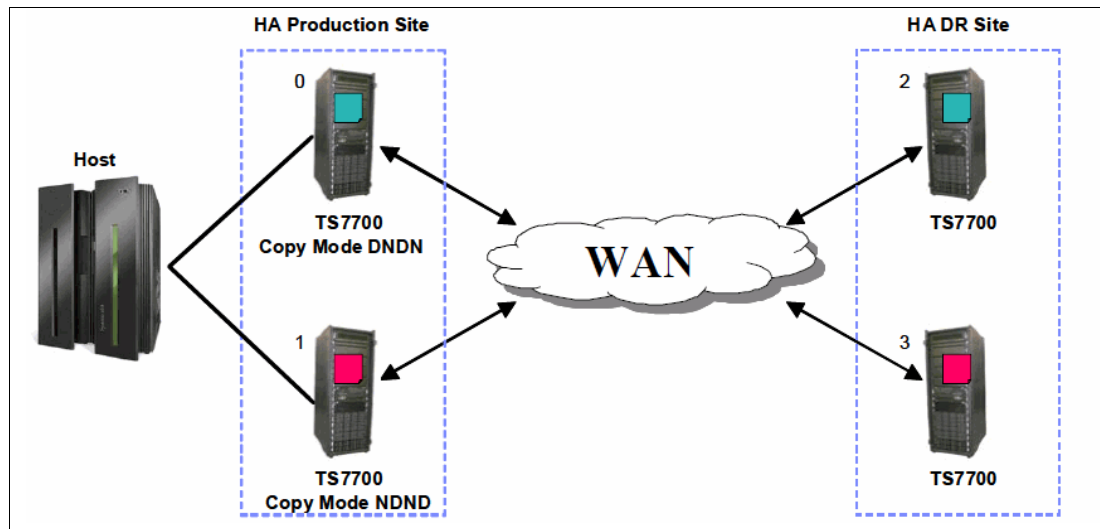


Figure 9-27 Two two-cluster grids in one

With Dynamic Allocation Assist, which is supported by JES2 but not JES3, the host can allocate a virtual device for a private mount on the best cluster. The best cluster is typically the cluster that contains the logical volume in its cache.

9.4.4 TS7700 hybrid grid considerations

Hybrid grids provide many possible configurations. Six basic combinations are available (one 2-way, two 3-way, and three 4-way possibilities). The same performance considerations apply to hybrids and homogeneous grids. An interesting hybrid grid, which is illustrated in Figure 9-28, is one where two or three TS7720 clusters are attached to the production hosts, and a single, perhaps remote, TS7740 DR cluster. The TS7720 clusters do not replicate to each other, but all of the TS7720 clusters replicate to the single TS7740. The advantages are a large front-end cache as presented by the TS7720 clusters, and a deep back-end for archiving and disaster recovery. However, the replication traffic from all of the TS7720 clusters is traveling across the same grid network. It is essential that adequate network bandwidth be provided to handle the traffic to the TS7740. Also, the network needs to have enough bandwidth to retrieve logical volumes that reside only in the TS7740 cluster. Figure 9-28 shows three TS7720 production clusters and one TS7740 DR cluster:

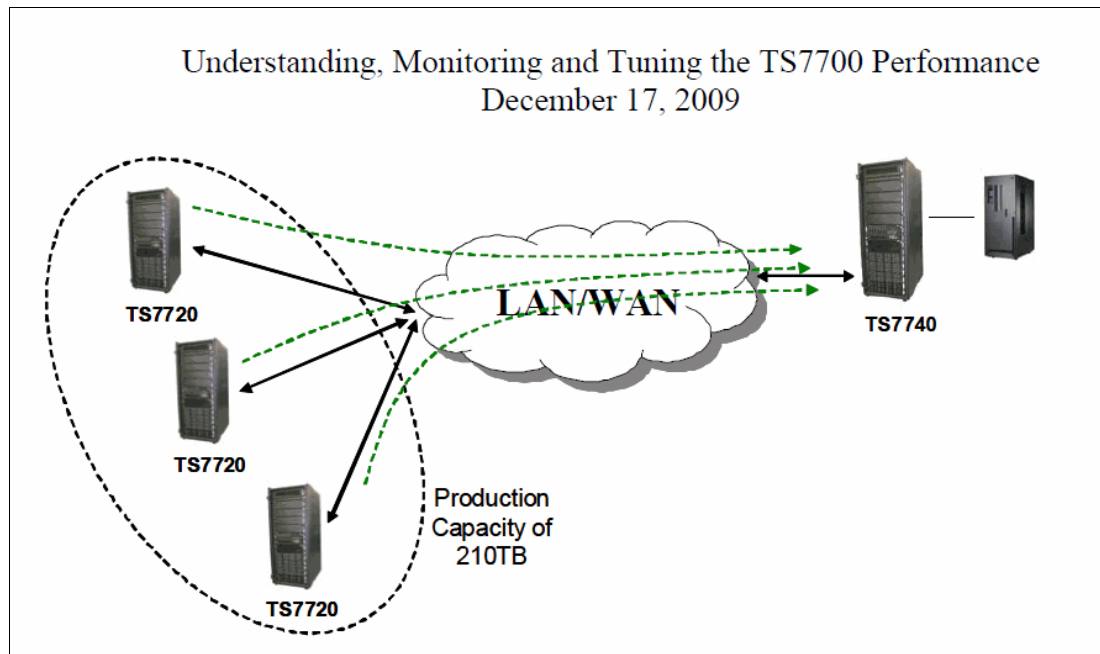


Figure 9-28 Hybrid grid (three TS7720 Production Clusters, one TS7740 DR Cluster)

9.4.5 Cluster families and cooperative replication

Prior to release 1.6 and cluster families, only the Copy Consistency Points could be used to direct who gets a copy of data and when they get the copy. Decisions as to where to source a volume from were left to each cluster in the grid. Two copies of the same data may be transmitted across the grid links for two distant clusters. With the introduction of cluster families in release 1.6, you can make the copying of data to other clusters more efficient by influencing where a copy of data is sourced. This becomes important with three- and four-cluster grids where the clusters may be geographically separated.

For example, when two clusters are at one site and the other two are at a remote site, and the two remote clusters need a copy of the data, cluster families make it so only one copy of the

data is sent across the long grid link. Also, when deciding where to source a volume, a cluster will give higher priority to a cluster in its family over another family. A cluster family establishes a special relationship between clusters. Typically families are grouped by geographic proximity to optimize the use of grid bandwidth. Family members are given higher weight when deciding which cluster to prefer for TVC selection.

Figure 9-29 illustrates how cooperative replication occurs with cluster families. Cooperative replication is used for deferred copies only. When a cluster needs to pull a copy of a volume it will prefer a cluster within its family. The example uses CCPs of RRDD. With cooperative replication one of the family B clusters at the DR site pulls a copy from one of the clusters in production family A. The second cluster in family B waits for the other cluster in family B to finish getting its copy then pulls it from its family member. This way the volume travels only once across the long grid distance.

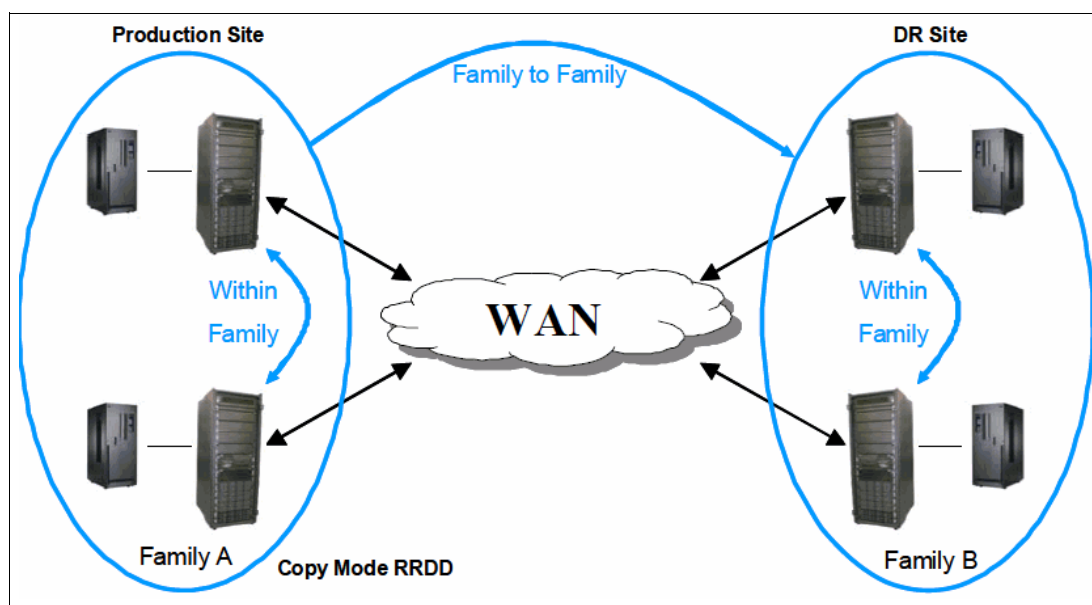


Figure 9-29 Cluster families

Cooperative replication includes another layer of consistency. A family is considered consistent when only one member of the family has a copy of a volume. Because only one copy is required to be transferred to a family, the family is consistent after the one copy is complete. Because a family member prefers to get its copy from another family member instead of getting the volume across the long grid link, the copy time is much shorter for the family member. Because each family member is pulling a copy of a different volume, this will make a consistent copy of all volumes to the family quicker. With cooperative replication, a family will prefer retrieving a new volume that the family does not have a copy of yet, over copying a volume within a family. When fewer than 20 new copies are to be made from other families, the family clusters will copy among themselves. This means second copies of volumes within a family are deferred in preference to new volume copies into the family. Without families, a source cluster attempts to keep the volume in its cache until all clusters needing a copy have gotten their copy. With families, a cluster's responsibility to keep the volume in cache is released after all families needing a copy have it. This way allows PGO volumes in the source cluster to be removed from cache sooner.

See the following resources for details about cluster families:

- ▶ *IBM Virtualization Engine TS7700 Series Best Practices - Hybrid Grid:*
<http://w3.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101656>
- ▶ *TS7700 Technical Update (R1.7) and Hybrid Grid Best Practices:*
<http://w3.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS4165>

9.4.6 Retain Copy Mode and Device Allocation Assist

This section covers several of the modes that are available with the TS7700.

Retain Copy Mode is an optional setting where a volume's existing Copy Consistency Points (CCPs) are honored instead of applying the CCPs defined at the mounting cluster. This setting applies to private volume mounts for reads or write appends. It is used to prevent more copies of a volume in the grid than desired. The example in Figure 9-30 is a four-cluster grid where cluster 0 replicates to cluster 2, and cluster 1 replicates to cluster 3. The desired result is that only two copies of data remain in the grid after the volume is accessed. Later, the host wants to mount the volume written to cluster 0. On a JES2 system *Device Allocation Assist* is used to determine which cluster is the best to request the mount from. Device Allocation Assist asks the grid which cluster to allocate a virtual drive from. The host will then attempt to allocate a device from the best cluster, in this case, cluster 0.

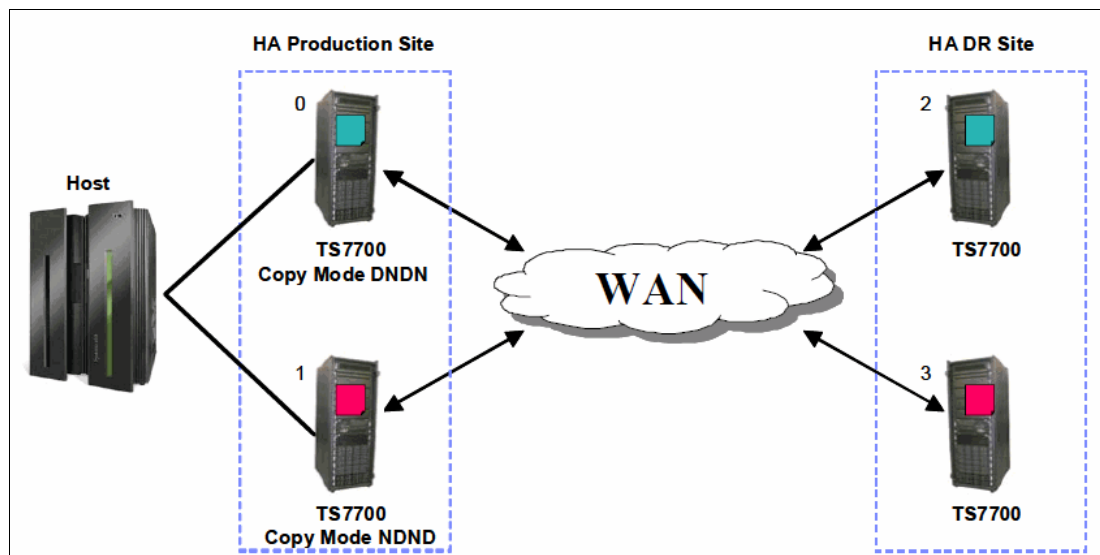


Figure 9-30 Four-cluster grid with Device Allocation Assist

Note: JES3 does not support Device Allocation Assist.

JES3 does not support Device Allocation Assist, so 50% of the time the host allocates to the cluster that does not have a copy in its cache. Without Retain Copy Mode, three or four copies of a volume will exist in the grid after the dismount instead of the desired two. In the case where host allocation picks the cluster that does not have the volume in cache, one or two additional copies are created on clusters 1 and 3 because the CCPs indicate the copies should be made to clusters 1 and 3. For a read operation, four copies remain. For a write append, three copies are created. This is illustrated in Figure 9-31 on page 626.

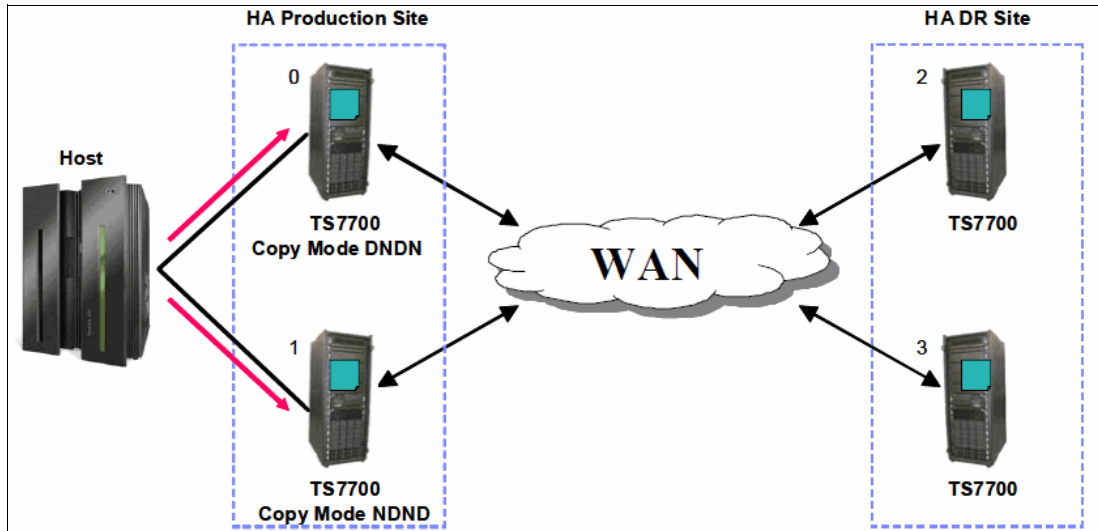


Figure 9-31 Four-cluster grid without Device Allocation Assist, Retain Copy Mode Disabled

With the Retain Copy Mode option set, the original CCPs of a volume are honored instead of applying the CCPs of the mounting cluster. A mount of a volume to the cluster that does not have a copy in its cache results in a cross cluster (remote) mount instead. The cross cluster mount uses the cache of the cluster that contains the volume. The CCPs of the original mount are used. In this case, the result is that cluster 0 and 2 have the copies and clusters 1 and 3 do not. This is illustrated in Figure 9-32.

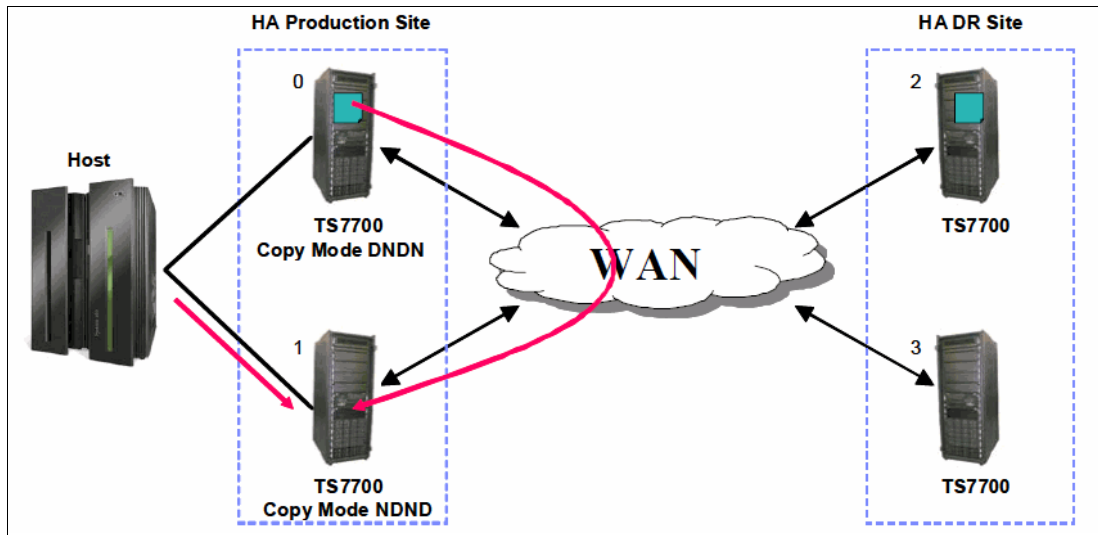


Figure 9-32 Four-cluster grid without Device Allocation Assist, Retain Copy Mode enabled

Another example of the need for Retain Copy Mode is when one of the production clusters is not available. All allocations are made to the remaining production cluster. When the volume only exists in clusters 0 and 2, the mount to cluster 1 will result in a total of three or four copies. This applies to both JES2 and JES3 without Retain Copy Mode enabled. See Figure 9-33 on page 627.




Figure 9-33 Four-cluster grid, One Production Cluster Down, Retain Copy Mode Disabled

The example in Figure 9-34 shows that the Retain Copy Mode is enabled and one of the production clusters is down. In the scenario where the cluster containing the volume to be mounted is down, the host will allocate to a device on the other cluster, in this case, cluster 1. A cross-cluster mount using Cluster 2's cache occurs, the original two copies remain. If the volume is appended to it is changed on cluster 2 only. Cluster 0 will get a copy of the altered volume when it rejoins the grid.

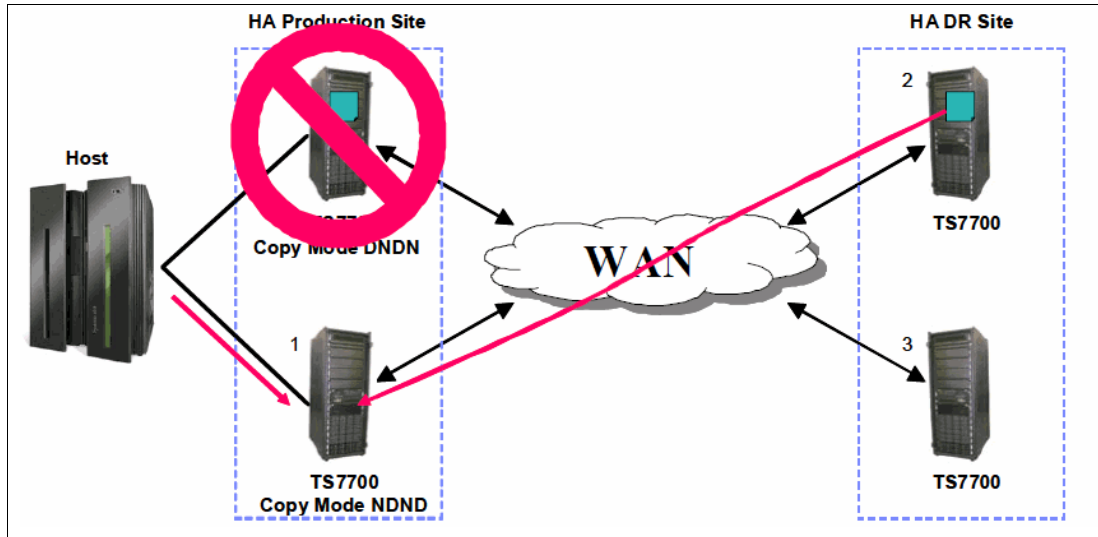


Figure 9-34 Four-cluster grid, One Production Cluster Down, Retain Copy Enabled

See the *IBM Virtualization Engine TS7700 Series Best Practices - Hybrid Grid* white paper at the Techdocs website:

<http://w3.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS4165>

9.5 Monitoring TS7700 Virtualization Engine performance

The IBM Virtualization Engine TS7700 Series is the latest in the line of tape virtualization products that has revolutionized the way mainframe customers utilize their tape resources. As the capability of tape virtualization has grown, so has the need to efficiently manage the large number of logical volumes the system supports. Internal to the TS7700 Virtualization Engine, a large amount of information is captured and maintained about the state and operational aspects of the resources within the TS7700 Virtualization Engine. The TS7700 Virtualization Engine provides a management interface based on open standards through which a storage management application can request specific information the TS7700 Virtualization Engine maintains. The open standards are not currently supported for applications running under z/OS, so an alternative method is needed to provide the information to mainframe applications.

You may use the following interfaces, tools, and methods to monitor the TS7700 Virtualization Engine:

- ▶ IBM System Storage TS3500 Tape Library Specialist
- ▶ TS7700 Virtualization Engine management interface
- ▶ Bulk Volume Information Retrieval function (BVIR)
- ▶ VEHSTATS
- ▶ VEPSTATS

The Specialist and management interface are web-based. With the BVIR function, various types of monitoring and performance-related information can be requested through a host logical volume in the TS7700 Virtualization Engine. Finally, the VEHSTATS and VEPSTATS tools can be used to format the BVIR responses, which are in a binary format, to create usable statistical reports. All interfaces, tools, and methods to monitor the TS7700 Virtualization Engine are explained in detail in the following sections. An overview of these interfaces, and tools, and methods are shown in Figure 9-35.

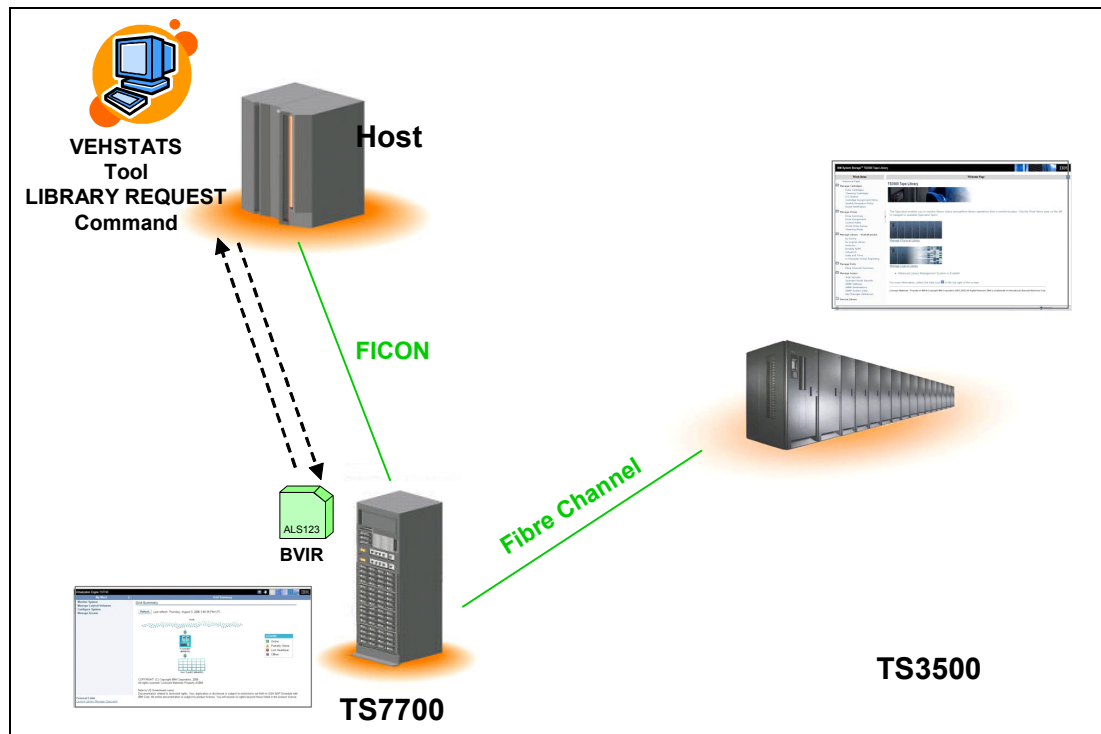


Figure 9-35 Interfaces, tools, and methods to monitor the TS7700 Virtualization Engine

9.6 Using the TS3500 Tape Library Specialist for monitoring

The Tape Library Specialist (TS3500 Tape Library Specialist) allows management and monitoring of items related to the TS3500 Tape Library. Initially, the web user interface to the TS3500 Tape Library only supported a single user at any given time. Now, each Ethernet-capable frame on the TS3500 Tape Library allows five simultaneous users of the web user interface so that multiple users can access the TS3500 Tape Library Specialist interface at the same time.

Figure 9-36 shows the TS3500 Tape Library System Summary window.

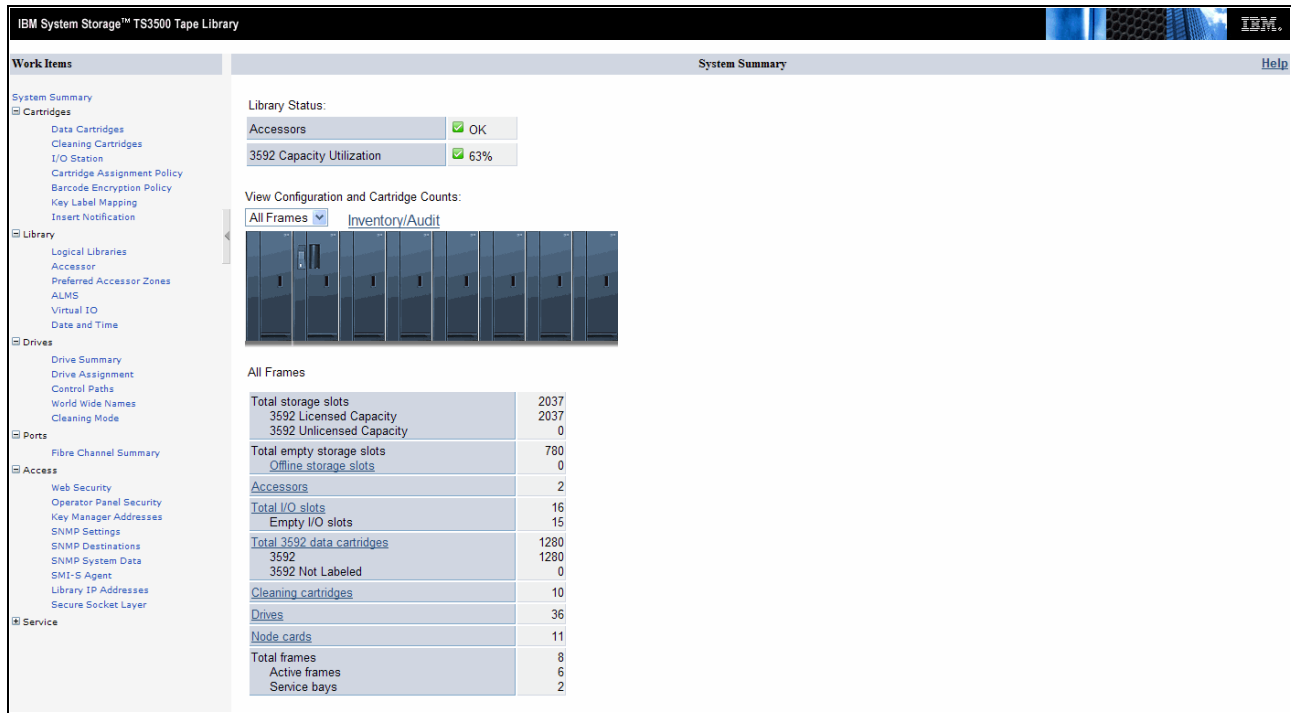


Figure 9-36 TS3500 Tape Library Specialist System Summary window

The TS3500 Tape Library Specialist session will time out after a default setting of ten minutes. This is different from the TS7700 Virtualization Engine management interface. You can change the default values through the TS3500 Tape Library Specialist by selecting **Manage Access** → **Web Security**, which opens the window shown in Figure 9-37 on page 630.

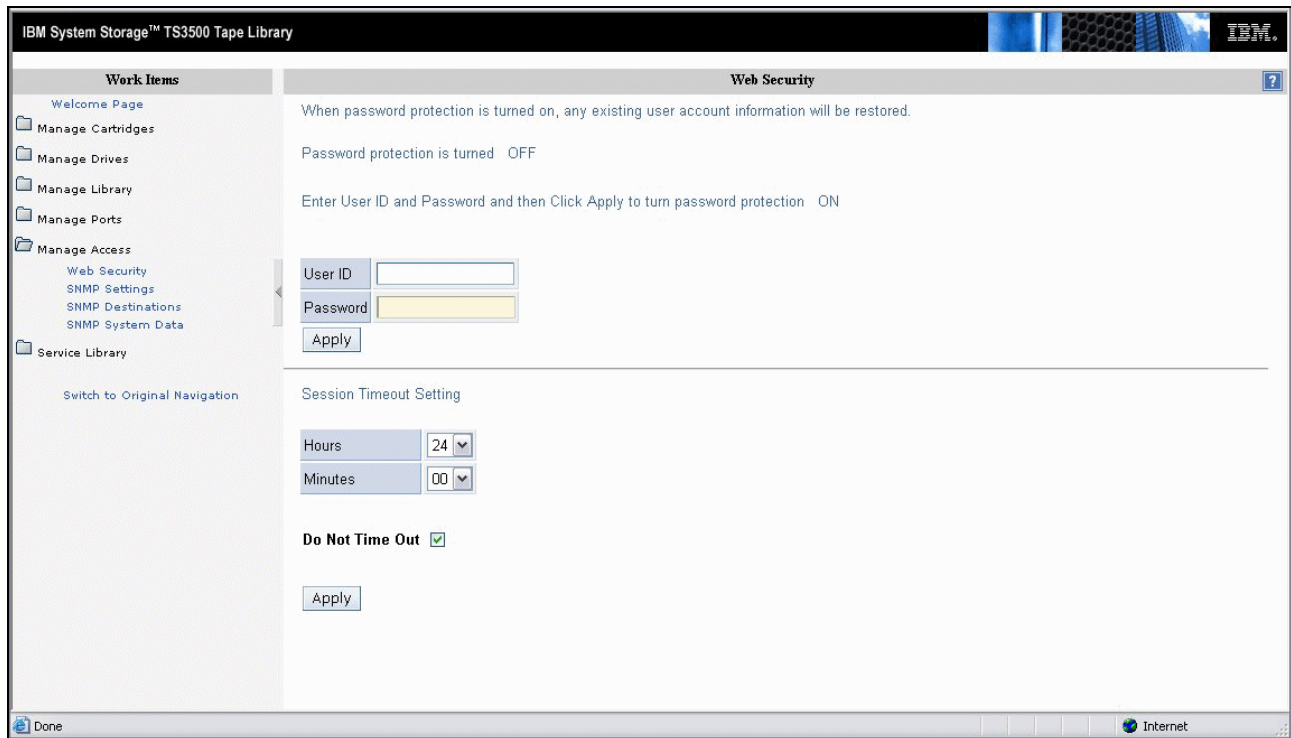


Figure 9-37 TS3500 Tape Library Specialist Manage Access: Web Security window

Some information provided by the TS3500 Tape Library Specialist is in a display-only format and there is no option to download data. Other windows provide a link for data that is available only when downloaded to a workstation. The data, in comma-separated value (CSV) format, can be downloaded directly to a PC and then used as input for snapshot analysis for the TS3500. This information refers to the TS3500 and its physical drives usage statistics from a TS3500 standpoint only.

For further information, including how to request and use this data, see *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789.

The following information is available:

- ▶ Accessor Usage: display only
 - Activity of each Accessor and gripper
 - Travel meters of Accessors
- ▶ Drive Status and Activity: display only
- ▶ Drive Statistics: download only
 - Last VOLSER on this drive
 - Write and Read MB per drive
 - Write and Read errors corrected per drive
 - Write and Read errors uncorrected per drive
- ▶ Mount History for cartridges: download only
 - Last Tape Alert
 - Number of Mounts of a specific cartridge
 - Number of Write and Read retries of a specific cartridge in the life cycle
 - Number of Write and Read permanent Errors of a specific cartridge in the life cycle

- Fibre Port statistics: download only

The Fibre Port statistics include fiber errors, aborts, resets, and recoveries between the TS7700 Virtualization Engine and the physical tape drives in the TS3500 Tape Library.

Restriction: This statistic does not provide information from the host to the TS7700 Virtualization Engine or host to controller.

- Library statistics, on a hourly basis: download only

- Total Mounts
- Total Ejects
- Total Inserts
- Average and Maximum amount of time a drive was mounted on a drive (residency)
- Average and Maximum amount of time was needed to perform a single mount
- Average and Maximum amount of time was needed to perform an eject

These statistics can be downloaded to a work station for further analysis. These statistics are not included in the Bulk Volume Information Retrieval (BVIR) records processed by the TS7700 Virtualization Engine.

Consideration: The information is available only on the TS3500 Tape Library.

9.7 Using the TS7700 Virtualization Engine management interface for monitoring

The TS7700 Virtualization Engine management interface (MI) belongs to the family of tools used for reporting and monitoring IBM storage products. These tools do not provide reports, but can be used for online queries about the status of the TS7700 Virtualization Engine, its components, and the distributed libraries. They also provide information about the copies that have not been completed yet and on the amount of data to be copied.

The TS7700 Virtualization Engine MI is based on a web server that is installed in each TS7700 Virtualization Engine. You can access this interface with any standard web browser.

The TS7700 Virtualization Engine MI is a Storage Management Initiative - Specification (SMI-S) compliant interface providing the user a single access point to manage resources without requiring their physical presence.

The TS7700 Virtualization Engine MI is a Storage Management Initiative - Specification (SMI-S) compliant interface that provides to the user a single access point to remotely manage resources through a standard web browser. The MI is required for implementation and operational purposes. In a TS7700 Virtualization Engine configuration, two possible web interfaces are available:

- The TS3500 Tape Library Specialist
- The TS7700 Virtualization Engine management interface

A link is available to the TS3500 Tape Library Specialist from the TS7700 Virtualization Engine management interface, as shown at the lower left corner the page of Figure 9-38 on page 632. This link might not be available if not configured during TS7740 installation or for a TS7720 Virtualization Engine.

In this section, we introduce the Performance and Statistics panels of the TS7700 Virtualization Engine management interface.

9.7.1 Performance and statistics

This section presents information that is related to viewing performance information and statistics for the TS7700 Virtualization Engine for Single and multi-cluster grid configurations. The graphical views display snapshots of the processing activities from the last 15 minutes. You can access the following selections by navigating to the Performance & Statistics section in the TS7700 Virtualization Engine management interface (MI). The examples are taken from different cluster configurations.

The navigation pane is available on the left side of the MI, as shown in the Grid Summary window shown in Figure 9-38.

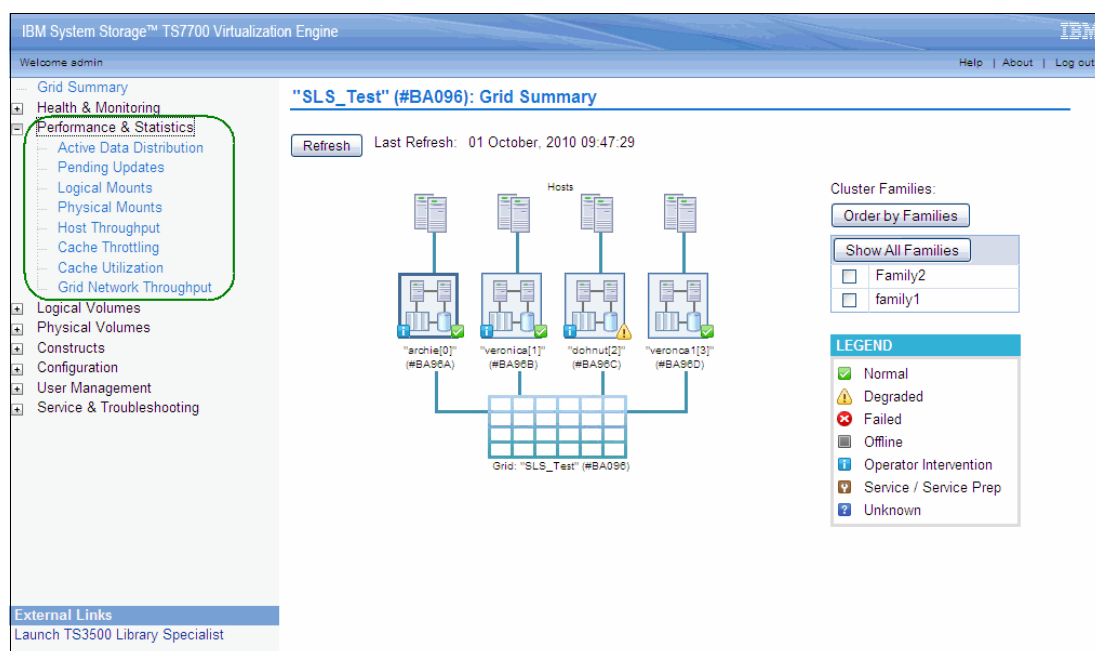


Figure 9-38 TS7700 Virtualization Engine management interface Grid summary

Active Data Distribution

This window (Figure 9-39) shows the active pools and correspondent data distribution (number of cartridges by a occupancy percentage range).

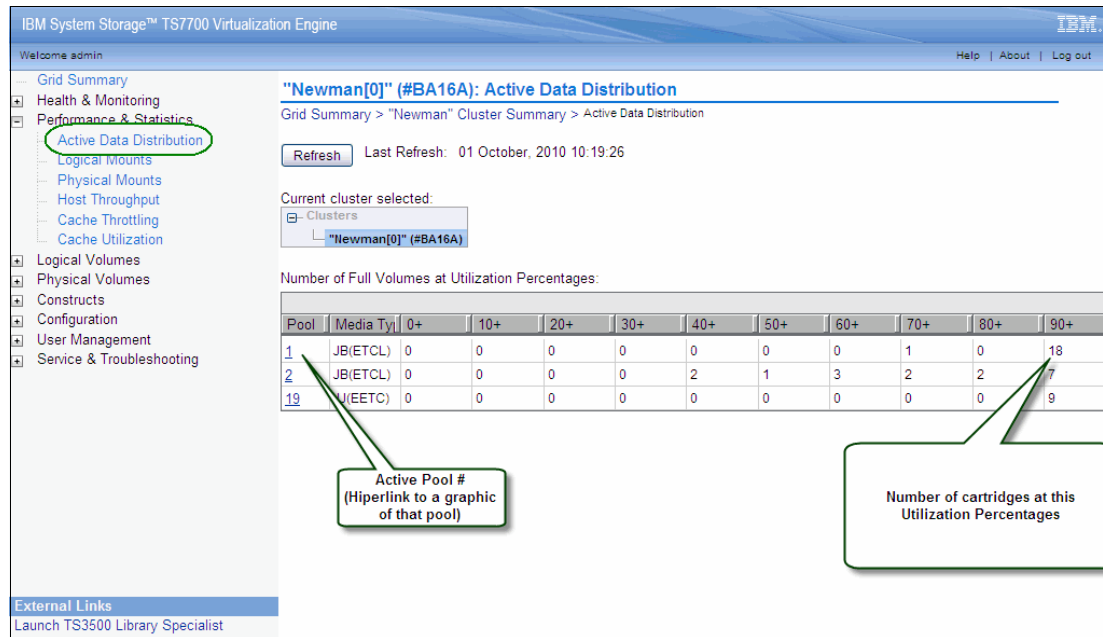


Figure 9-39 TS7740 MI Active Data Distribution

Click a pool link. Information for the pool is displayed, as shown in Figure 9-40

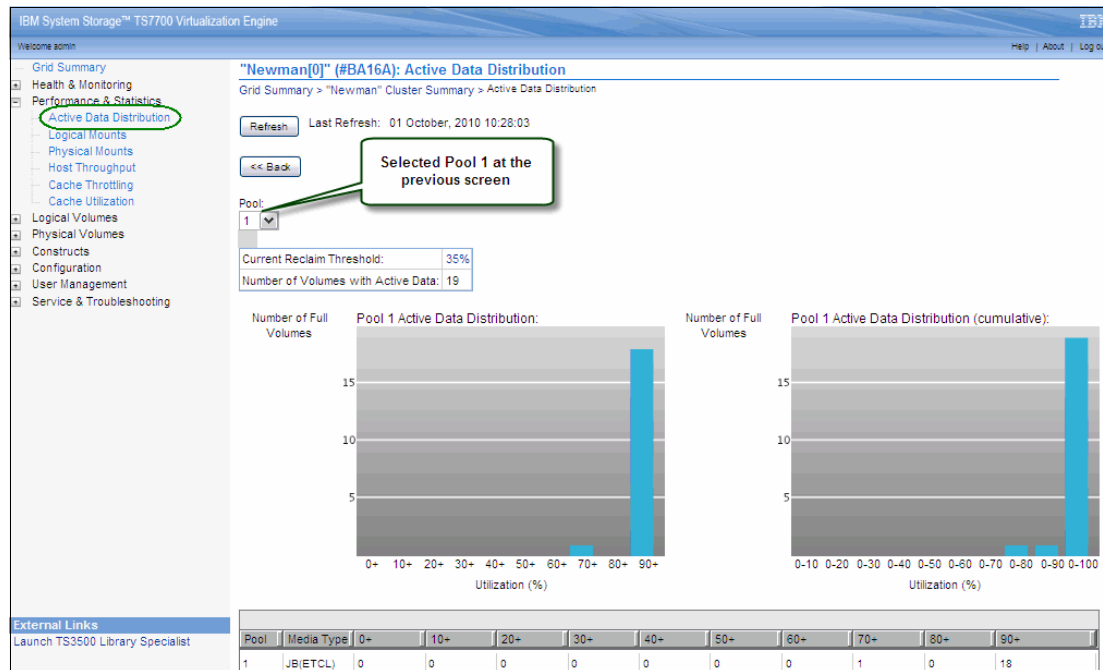


Figure 9-40 Pool Active Data Distribution

Pending Updates

Use this page (Figure 9-41) to view the pending updates for the IBM Virtualization Engine TS7700 Grid. The existence of pending updates indicates that updates occurred while a cluster was offline, in Service Prep mode, or in Service mode. Before any existing pending updates can take effect, all clusters must be online.

This page provides a summary of the number of outstanding updates for each cluster in an IBM Virtualization Engine TS7700 Grid. You can also use this page to monitor the progress of pending immediate-deferred copies, which, like pending updates, normally result from changes made while a cluster is Offline, in Service Prep mode, or in Service mode.

Note: Pending immediate-deferred copies may also result from grid network problems.

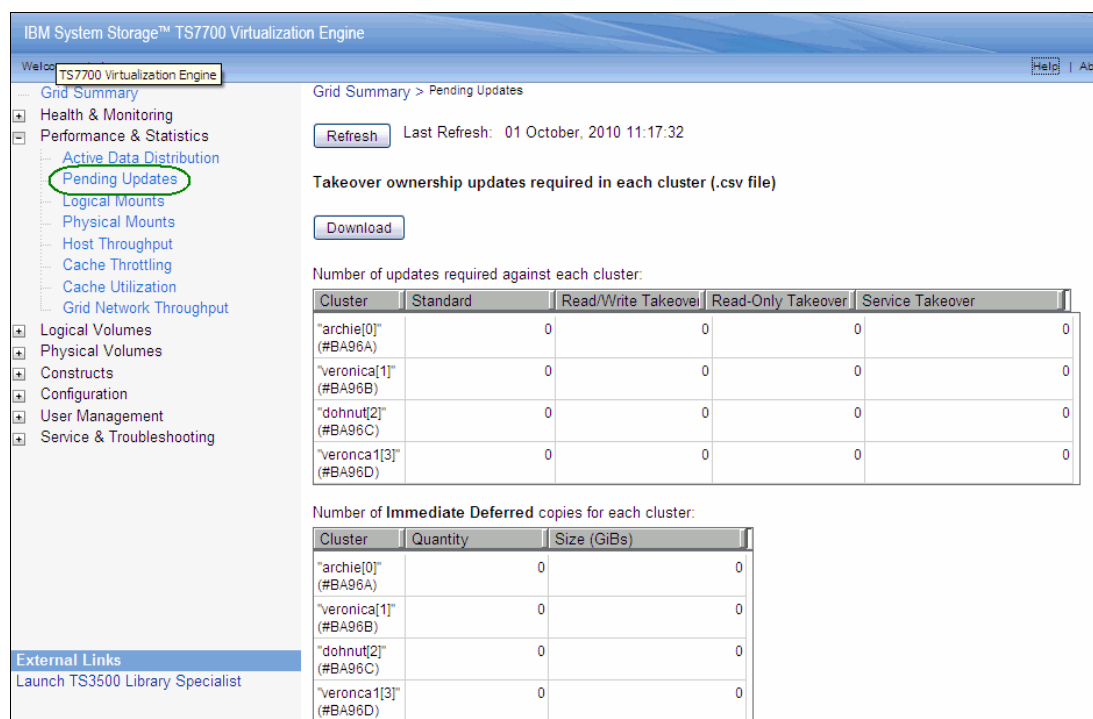


Figure 9-41 Grid Pending Updates

Logical Mounts window

Use this window (Figure 9-42) to view logical mount statistics for the TS7700 Virtualization Engine. The logical mount statistics for each cluster are displayed in two bar graphs and tables: one for the number of mounts and one for average mount time. The example in Figure 9-42 is from a TS7700 Virtualization Engine Cluster that is part of a multi-cluster grid configuration (four-cluster grid).

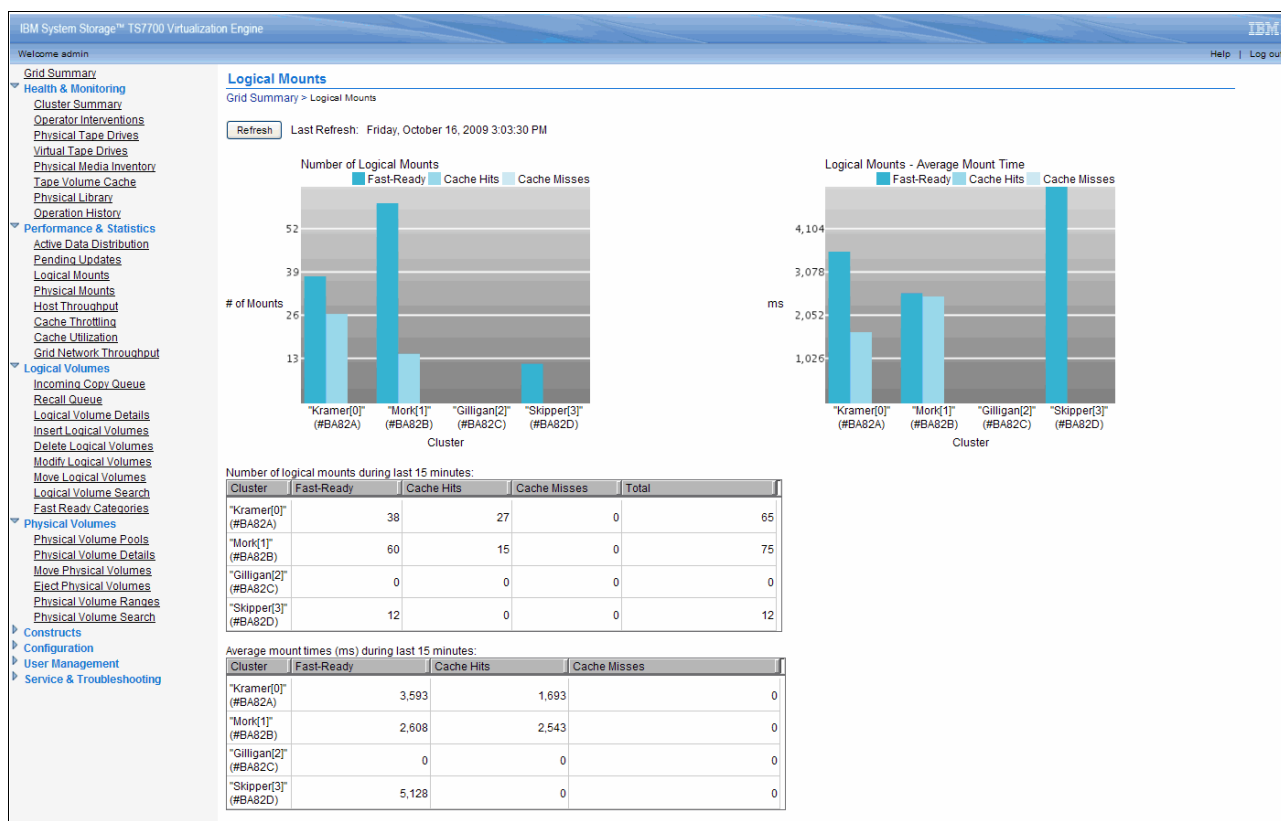


Figure 9-42 TS7700 Virtualization Engine management interface Logical Mounts window

"Number of logical mounts during last 15 minutes" table has the following information:

Cluster	The cluster name
Fast-Ready	Number of logical mounts completed using the Fast-Ready method
Cache Hits	Number of logical mounts completed from cache
Cache Misses	Number of mount requests that were not fulfilled from cache
Total	Total number of logical mounts

"Average mount times (ms) during last 15 minutes" table has the following information:

Cluster	The cluster name
Fast-Ready	Average mount time for Fast-Ready scratch logical mounts
Cache Hits	Average mount time for logical mounts completed from cache
Cache Misses	Average mount time for requests that are not fulfilled from cache

Physical Mounts window

Use this window (Figure 9-43) to view physical mount statistics for the TS7740 Virtualization Engine. The physical mount statistics for each cluster are displayed in two bar graphs and tables: one for the number of mounts by category and one for average mount time per cluster. The example in Figure 9-43 is from a TS7740 Virtualization Engine Cluster that is part of a multi-cluster grid configuration (four-cluster grid).

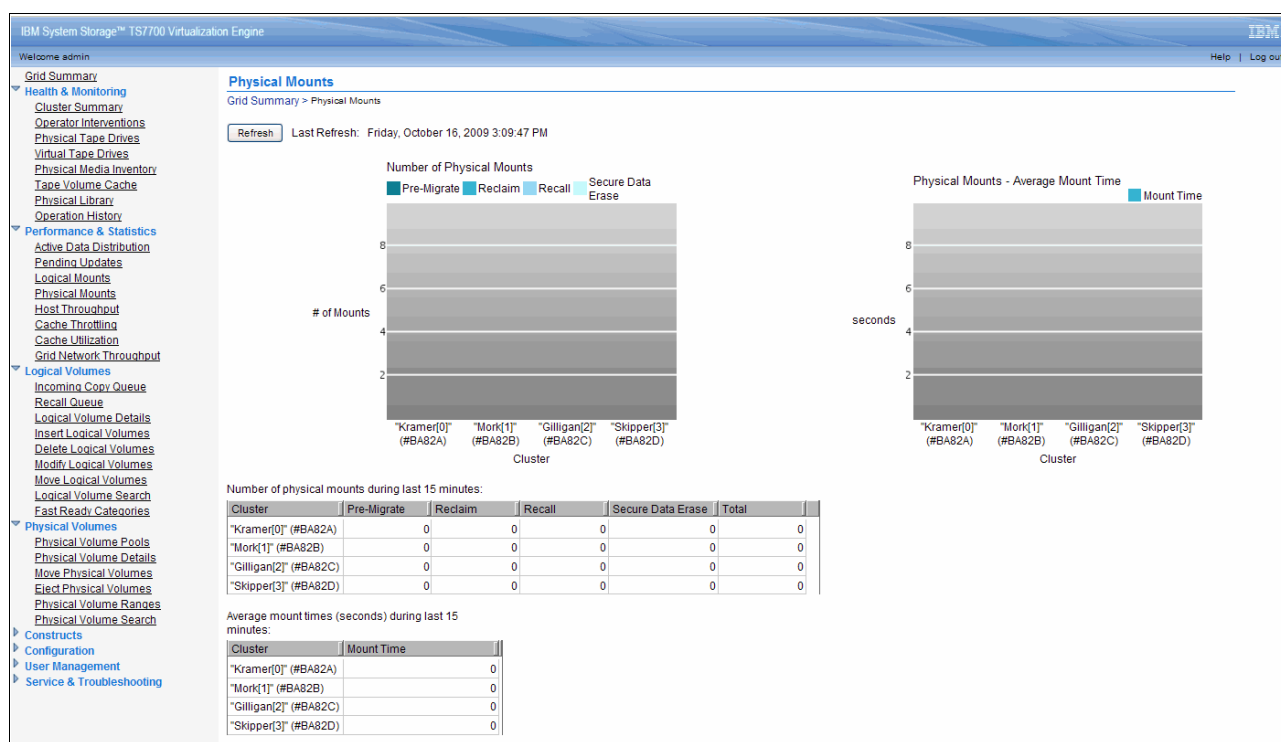


Figure 9-43 TS7740 Virtualization Engine management interface Physical Mounts window

The table cells show the following items:

Cluster	The cluster name
Pre-migrate	Number of premigrate mounts
Reclaim	Number of reclaim mounts
Recall	Number of recall mounts
Secure Data Erase	Number of Secure Data Erase mounts
Total	Total physical mounts
Mount Time	Average mount time for physical mounts

Host Throughput window

You can use this window (Figure 9-44 on page 637) to view host throughput statistics for the TS7700 Virtualization Engine. The information is provided in 15-second intervals, unlike the 15-minute intervals of other performance data.

Use this window to view statistics for each cluster, vNode, host adapter, and host adapter port in the Grid. At the top of the window is a collapsible tree where you view statistics for a specific level of the Grid and cluster. Click the Grid to view information for each cluster. Click the cluster link to view information for each vNode. Click the vNode link to view information for each host adapter. Click a host adapter link to view information for each port.

The example in Figure 9-44 is from a TS7700 Virtualization Engine Cluster that is part of a multi-cluster grid configuration (four-cluster grid).

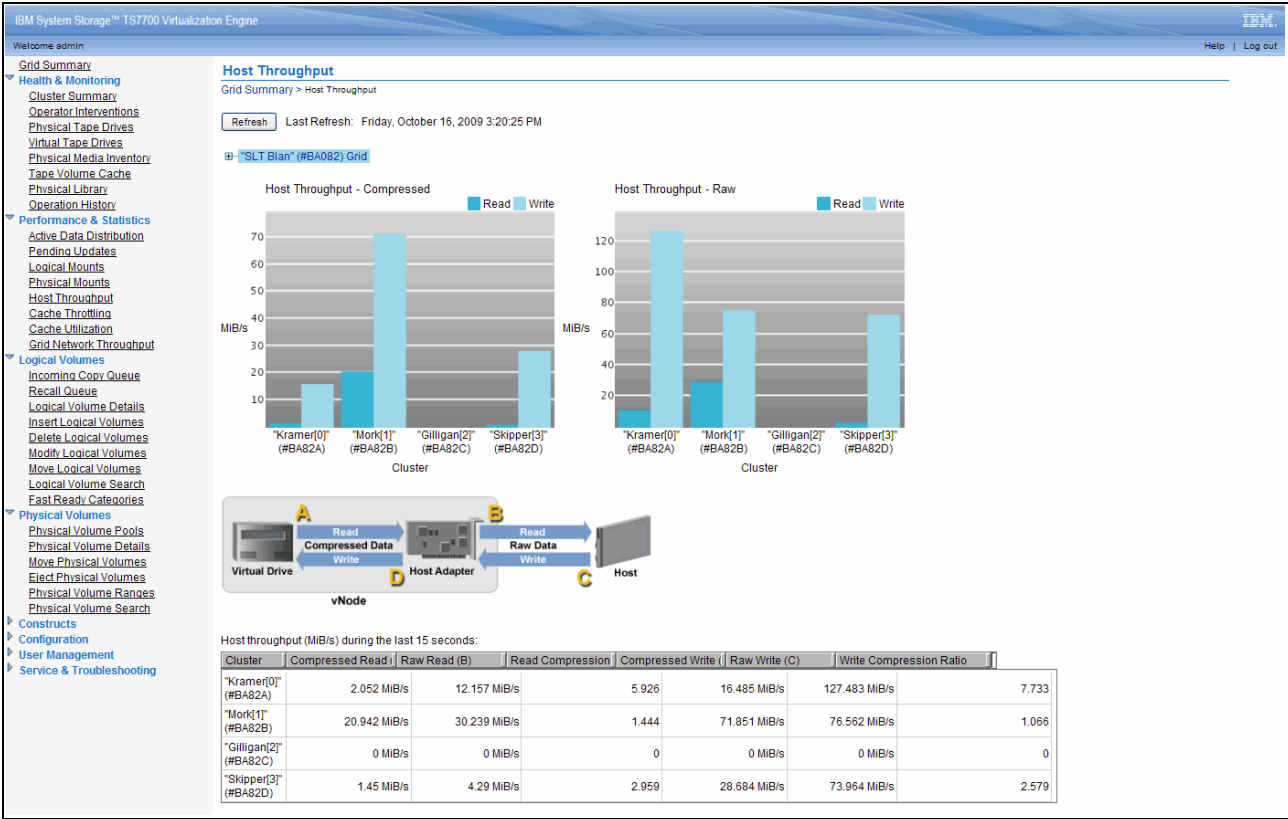


Figure 9-44 TS7700 Virtualization Engine management interface Host Throughput window

The host throughput data is displayed in two bar graphs and one table. The bar graphs are for raw data coming from the host to the host bus adapter (HBA) and for compressed data going from the HBA to the virtual drive on the vNode.

The letter next to the table heading corresponds with the letter in the diagram above the table. Data is available for a cluster, vNode, host adapter or host adapter port. The table cells include the following items:

- Cluster** (vNode, Host Adapter, Host Adapter Port)
The cluster or cluster component for which data is being displayed
- Compressed Read (A)** Amount of data read between the virtual drive and HBA
- Raw Read (B)** Amount of data read between the HBA and host
- Read Compression Ratio** Ratio of compressed read data to raw read data
- Compressed Write (D)** Amount of data written from the HBA to the virtual drive
- Raw Write (C)** Amount of data written from the host to the HBA
- Write Compression Ratio** Ratio of compressed written data to raw written data

Cache Throttling window

You can use this window (Figure 9-45) to view cache throttling statistics for the TS7700 Virtualization Engine. The example in Figure 9-45 is from a TS7700 Virtualization Engine Cluster that is part of a multi-cluster grid configuration (four-cluster grid).



Figure 9-45 TS7700 Virtualization Engine management interface Cache Throttling window

Cache throttling is a time interval applied to TS7700 Virtualization Engine internal functions to improve throughput performance to the host. The cache throttling statistics for each cluster in regards to copy and write are displayed both in a bar graph form and in a table. The table shows the following items:

Cluster	The cluster name
Copy	The amount of time inserted between internal copy operations
Write	The amount of time inserted between internal write operations

Cache Utilization window

You can use this window (Figure 9-46) to view cache utilization statistics for the TS7700 Virtualization Engine. The example in Figure 9-46 is from a TS7740 Virtualization Engine cluster that is part of a multi-cluster grid configuration (four-cluster grid).

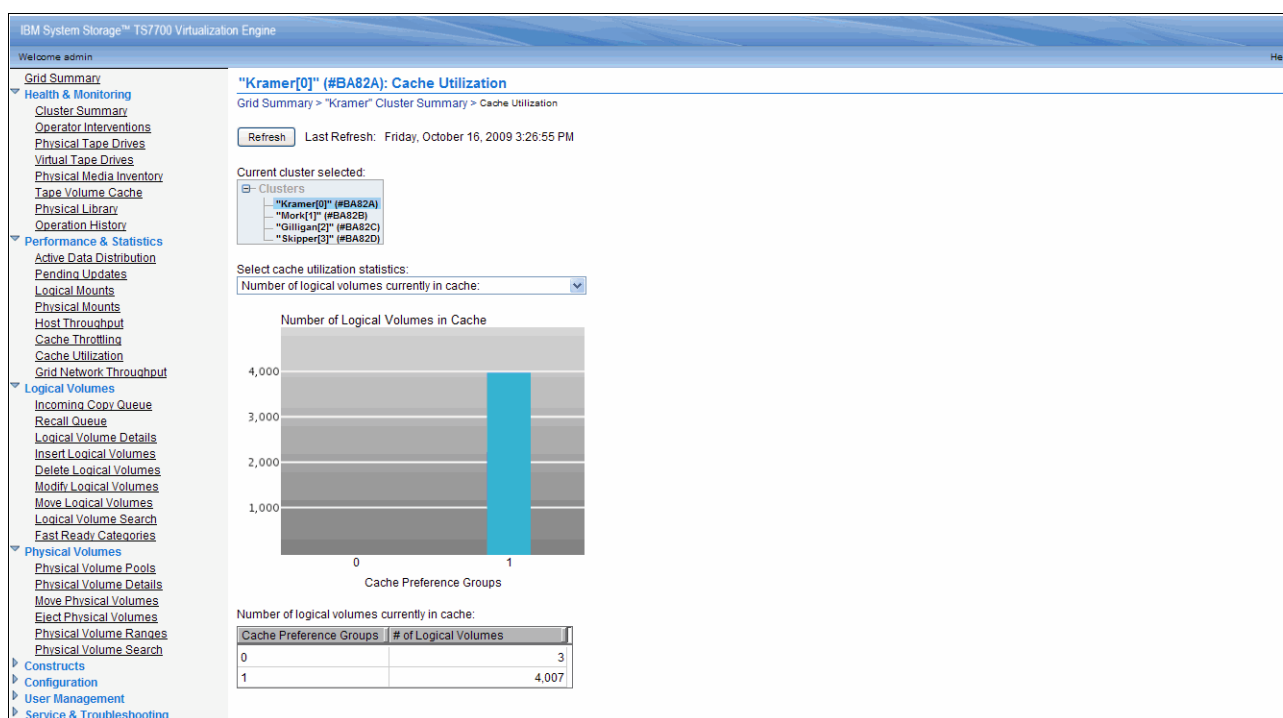


Figure 9-46 TS7740 Virtualization Engine management interface Cache Utilization window

The cache utilization statistics can be selected for each individual cluster. Various aspects of cache performance are displayed for each cluster. Select them from the “Select cache utilizations statistics” drop-down menu. The data is displayed in both bar graph and table form and further by preference groups 0 and 1.

The available cache utilization statistics are as follows:

- ▶ Cache Preference Group. Possible values are as follows:
 - 0: Volumes in this group have a preference for removal from cache over other volumes.
 - 1: Volumes in this group have a preference to be retained in cache over other volumes.
- ▶ Number of logical volumes currently in cache: The number of logical volumes present in the cache preference group
- ▶ Total amount of data currently in cache: Total amount of data present in volumes assigned to the cache preference group
- ▶ Median duration that volumes have remained in cache: Rolling average of the cache age of volumes migrated out of this cache preference group for the specified amount of time (last 4 hours, 48 hours, 35 days)
- ▶ Number of logical volumes migrated: Rolling average of the number of volumes migrated to this cache preference group (4 hours, 48 hours, 35 days). Bar graph is not used.

Note: Median Duration in Cache and Number of Logical Volumes Migrated statistics have a table column for each of the time periods mentioned in parentheses.

Grid Network Throughput window

Use this window (Figure 9-47) to view network path utilization (Grid Network Throughput) statistics for the TS7700 Virtualization Engine Cluster.

Note: Grid Network Throughput option is not available in a stand-alone Cluster.

This window presents information about cross-cluster data transfer rates. This selection will be present only in a multi-cluster grid configuration. If the TS7700 Virtualization Engine Grid only has one cluster, there is no cross-cluster data transfer through the Ethernet adapters.

The example in Figure 9-47 is from a TS7700 Virtualization Engine Cluster that is part of a multi-cluster grid configuration (four-cluster grid).

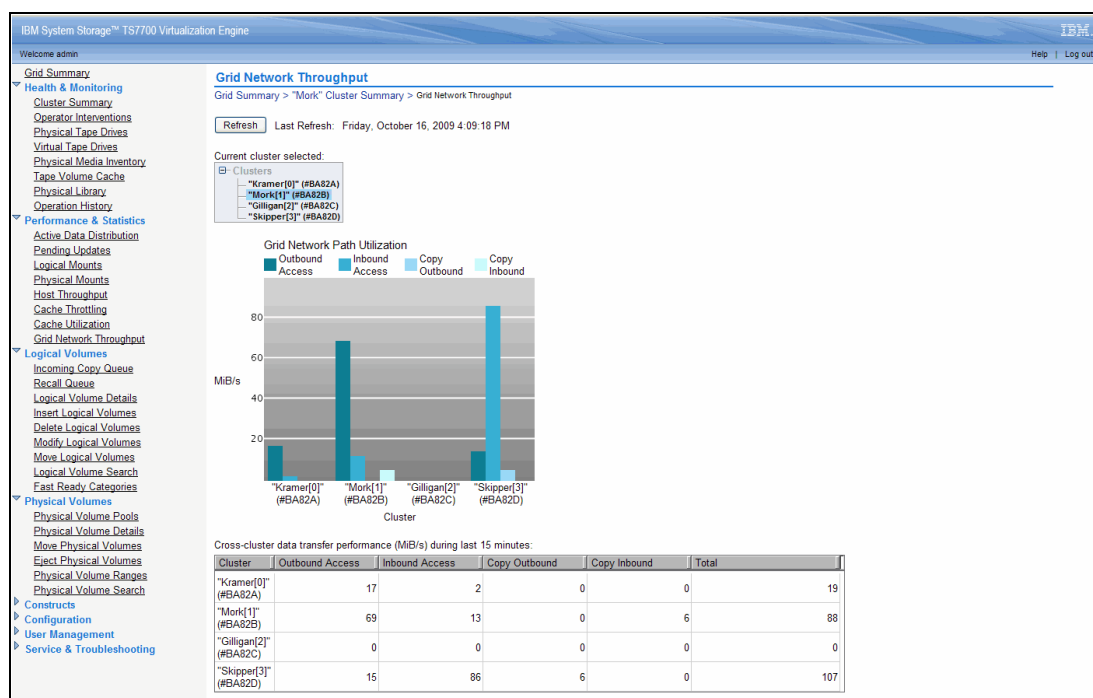


Figure 9-47 TS7700 Virtualization Engine management interface Grid Network Throughput in a four-cluster grid

The table displays data for cross-cluster data transfer performance (MBps) during the last 15 minutes. The table cells show the following items:

Cluster	The cluster name
Outbound Access	Data transfer rate for host operations that move data from the specified cluster into one or more remote clusters
Inbound Access	Data transfer rate for host operations that move data into the specified cluster from one or more remote clusters
Copy Outbound	Data transfer rate for copy operations that pull data out of the specified cluster into one or more remote clusters
Copy Inbound	Data transfer rate for copy operations that pull data into the specified cluster from one or more remote clusters

9.8 Tuning the TS7700

To change the “knobs” settings, altering the behavior of the TS7700 you must be able to collect the Statistics data from your Clusters, use the available tools to format and plot the binary data, and understand the resulting graphs.

As an working example, see the Figure 9-48. This is cache throughput plotted from VEHSTATS.

Plotting Cache Throughput from VEHSTATS

When evaluating performance, a graph that reveals a lot of information in a small space is the *cache throughput* for a cluster. The example here is from a three-cluster grid in an HA and DR configuration. Clusters 0 and 1 are the production clusters and replicate data to each other. cluster 2 is the remote DR cluster. It receives copies from both cluster 0 and 1.

The cache throughput graph stacks the data rates of data (MBps) that is moved through cache to show the total amount of data moved through cache. Remember, all data passed through cache is compressed data. The graph view helps you see the ebb and flow of various tasks in the TS7700 including host IO, housekeeping tasks (copy to other clusters, premigrate), and recall of data from tape into cache.

For consistency, be sure you use the same coloring scheme in the graphs so that they are more easily recognizable and familiar. It is best to plot 24 hours or less in a single chart so the data can be read easily. All data is available in VEHSTATS, in both the HOURFLAT form and in the formatted reports. The descriptions below describe the columns used in both the HOURFLAT file and the formatted reports.




Figure 9-48 TS7740 activity graph

Note the following information:

- Compressed Host IO: This is the MBps of the data written to and read from cache. This bar is bright green (Avg_Dev_MB_s).
- Data copied from this cluster to other clusters: This is the rate at which copies of data to other clusters are made. This cluster is the source of the data and includes copies to all other clusters in the grid. The DCT applied by this cluster applies to this data. For a two-cluster grid this is a single value. For a three-cluster grid there are two values, one for copies to each of the other clusters. For a four-cluster grid there are three values, one for

copies to each of the other clusters. The descriptions below are for plotting cluster 0's cache throughput. Use the appropriate columns when plotting other clusters. These bars are cyan or light blue (Avg_01_MB_s or Avg_02_MB_s).

- Data copied to this cluster from other clusters: This is the rate at which other clusters are copying data into this cluster. This cluster is the target of the data and includes copies from all other clusters in the grid. The DCT applied by other clusters applies to this data. For a two-cluster grid this is a single value. A three-cluster grid has two values: one for copies from each of the other clusters. A four-cluster grid has three values, one for copies from each of the other clusters. The following descriptions are for plotting cluster 0's cache throughput. Use the appropriate columns when plotting other clusters. These bars are light purple (Avg_10_MB_s).
- Data premigrated from cache to tape: This is the rate at which data is being read from cache and written to physical tape. This bar is bright yellow (Pre_Mig).
- Data recalled from tape to cache: This is the rate at which data is being read from tape into cache for a mount requiring a recall. This bar is dark blue (Recall).
- Uncompressed Host IO: You can optionally plot this statistic on the graph. It should not be part of the stacked bars showing cache throughput. This helps customers see what the host view of throughput is relative to the cache activity. This is a solid red line (Max_Chan_MB_s).

Interpreting Cache Throughput

The TS7700 cache has a finite bandwidth of just over 300 MBps (CC6/CX6 is 320 MBps; CC7/CX7 is 350 MBps). Typical cache IO rates during busy periods are in the range of 250 - 300 MBps. This bandwidth is shared between the Host IO (compressed), copy activity, premigration activity, and reclaim activity (blocked during heavy production periods using the Inhibit Reclaim Schedule). The TS7700 balances these tasks using various thresholds and controls in an effort to prefer host IO. The two major housekeeping tasks at work are the premigration of data from cache to tape, and deferred copies to/from other clusters.

The graph in Figure 9-49 on page 643 shows the cache throughput for cluster 0 which is part of a three-cluster grid. Clusters 0 and 1 are both attached to production hosts and make deferred copies to each other. Cluster 2 receives copies of data from both clusters 0 and 1. This is an HA and DR three-cluster grid. The rise and decline of data flow of the host IO can be seen in both the uncompressed Host IO rate (red line) and the compressed host IO rate (bright green bar). When cluster 0 turns on the DCT, to slow the copies to other clusters, both of the cyan/light blue bars are smaller. When cluster 1 turns on the DCT, the light purple bar becomes smaller in size. These cyan/light blue and light purple bars increase in size when the DCT is turned off or is only turned on sporadically during an interval. The premigration activity is shown with the yellow bars.




Figure 9-49 Premigration activity

Fast Host Write premigration algorithm

The control of premigration tasks is discussed in an earlier section. The Fast Host Write algorithm limits the number of premigration tasks to two, one, or zero. This limit occurs when the compressed host write rate is greater than 30 MBps and the CPU is greater than 99% busy. The portion encircled on the left of the graph (Figure 9-49) illustrates this algorithm in effect. During the 02:30 and 02:45 intervals, the amount of premigrate activity is limited. During the next 5 intervals the premigration activity is zero except for a small amount during the 03:15 interval. After this period of intense host activity and CPU usage, the premigrate tasks are allowed to start again.

False-end-of-batch phenomenon

The false-end-of-batch phenomenon (encircled at the right in Figure 9-49) is triggered when the compressed host I/O rate drops off and the DCT is turned off. In the graph the host I/O drops off during the 20:45 interval, after a high host I/O rate during the 20:15 and 20:30 intervals. The lower host I/O rate triggers the TS7700 to turn off the DCT thus unleashing the deferred copies. The cyan bars account for over 250 MBps of cache throughput at 20:45 through 21:15. When the host wants to push more I/O, the TS7700 is slow to react to reapply the DCT. This is a chicken and the egg scenario: The host wants to send more I/O but the TS7700 sees low activity so it concentrates on housekeeping. Eventually the Host I/O increases enough to allow the thresholds and controls to moderate the housekeeping tasks. You can see the host I/O rate ramp up starting about 21:45 and the deferred copies diminish. Again, at 22:45 the host I/O falls off and the DCT is turned off. This has been labeled the *false-end-of-batch* indication.

The R1.5 PGA 3 (September 2009) release adds a control that allows the customer to set the DCT threshold (default of 100 MBps). Also, the algorithm that decides when to turn off the DCT is enhanced so that it waits for a sustained period of lower Host IO activity before turning off the DCT. This causes the DCT to remain in effect during brief dips in host IO activity. Also, the DCT threshold adjustment allows the threshold to be lowered thus keeping the DCT applied longer. This should help to allow the false-end-of-batch to occur but keep the bandwidth open for host IO longer. The Host Console Request can be used to set the DCT Threshold using the SETTING, THROTTLE, DCTAVGTD keywords. Be careful when adjusting the DCT threshold because setting the threshold lower can cause a larger build-up in deferred copies. You might want to adjust the Preferred Pre-migration threshold and Pre-Migration Throttling threshold to a higher value. Also, setting the DCT threshold too low

could cause the deferred copy queue to never catch up during a 24-hour period. Carefully monitor the deferred copy queue depth after adjusting this threshold.

9.8.1 Tuning the TS7700

This section describes the various adjustments available to tune the TS7700.

Deferred Copy Throttle (DCT) value and threshold

The DCT is used to regulate outgoing deferred copies to other clusters to prefer host throughput. For some customers host throughput is more important than the deferred copies, but for others, deferred copies are just as important. Adjusting the DCT value and threshold can allow you to tune the performance of the deferred copies.

Deferred Copy Throttle (DCT) value

When the DCT threshold is reached, the TS7700 adds a delay to each block of deferred copy data sent across the grid links from a cluster. The larger the delay, the slower that the overall grid performance becomes. The performance of the grid links is also affected by the latency time of the connection. The one-way latency, which is half the round trip latency, is essentially an additional DCT. This means that a DCT value for one customer can have a different effect for another, if their grid latency times differ. For example, customer A sees a round-trip latency of 10 ms and customer B sees a latency of 44 ms. With a DCT value set to 70 ms, customer A has an effective DCT of $70 + 5 = 75$ ms, whereas customer B has an effective DCT of $70 + 22 = 92$ ms. The default DCT is 125ms. The effect on host throughput as the DCT is lowered is not linear. Lab testing suggests the knee of the curve is at approximately 30 ms. This is with a grid latency of 0ms. As the DCT value is lowered towards 30 ms the host throughput is affected somewhat and deferred copy performance improves somewhat. Around 30 ms the host throughput is affected more significantly as well as deferred copy performance. If the DCT needs to be adjusted from the default value, the initial recommended DCT plus network latency value is 70 - 85 ms. Favor the value to 70 ms if the customer is more concerned with deferred copy performance or towards 85 ms if the customer is concerned about sacrificing host throughput. The DCT value should take into account the grid one-way latency. For example, with a one-way latency of 5 ms, and the desired latency is 70 ms, set the DCT value to 65 ms. After adjusting the DCT, the host throughput and Deferred Copy Queue must be monitored to see if the desired balance of host throughput and deferred copy performance has been achieved. Lowering the DCT can improve deferred copy performance at the expense of host throughput. Depending on the code level of the TS7700, the DCT is set either by the SSR (in R1.4, R1.4a) or the customer (in R1.5):

- ▶ For R1.4 systems, the `vtd_dct` tool is used to set the DCT value in milliseconds. IBM Support Center has instructions for the use of this tool.
- ▶ For R1.4a systems, the DCT is set by creating a file with the DCT value in microseconds. Support has the instructions for the use of this tool.
- ▶ For R1.5 and later systems the DCT value can be set by the customer through the Host Console Request command. The setting of this throttle is discussed in detail in the *IBM Virtualization Engine TS7700 Series z/OS Host Command Line Request User's Guide* which is available from the Techdocs website. Use the SETTING, THROTTLE, DCOPYT keywords.

Deferred Copy Throttle (DCT) threshold

Prior to R1.5 PGA3, the DCT threshold is fixed at 100 MBps of compressed host I/O. With R1.5 PGA3, the DCT threshold can be changed by using the Host Console Request. This value is used to determine at what average host I/O rate to keep deferred copy throttling on. The average host I/O rate is a rolling average of the I/O rate over a 20-minute period. When

this average rate exceeds the DCT threshold, the deferred copies are delayed as specified by the DCOPYT value. The default DCT threshold value is 100 MBps. A value of 0 sets the DCT threshold to the default. The threshold can be set to a value of 0 (default of 100) or 1 - 500 (MBps). For each 30-second period, the host throughput is measured, and if it is greater than the DCT threshold (default of 100 MBps) *and* the CPU usage is >85%, deferred copy throttling is turned on. The 20- minute average (40 samples) of host throughput is used to determine whether deferred copy throttling should remain set or turned off. When the 20-minute average is greater than the DCT threshold, the deferred copy (and CPU usage is greater than 85%) throttling remains set. If both the current sample and the 20-minute average are below the DCT threshold, the DCT is turned off. For R1.5 PGA3 and later machines, the DCT threshold value can be set by the customer through the Host Console Request command.

Preferred Pre-migration and Pre-migration Throttling Thresholds

These two thresholds are triggered by the amount of non premigrated data in the cache. For releases prior to R1.5 these thresholds are fixed:

- ▶ 800 GB for the Preferred Pre-migration threshold
- ▶ 1000 GB for the Pre-migration Throttle threshold

For R1.5, the thresholds default to 800 GB and 1000 GB, respectively. You may modify the thresholds by using the Host Console Request. The amount of non premigrated data is available with the Host Console Request CACHE request.

For instance, in the following example, 750 GB of data has yet to be premigrated:

```
LI REQ,lib_name,CACHE
TAPE VOLUME CACHE STATE V1
INSTALLED/ENABLED GBS 6000/ 3000
PARTITION ALLOC USED PG0 PG1 PMIGR COPY PMT CPYT
          0      2000      1880      0 1880      750      0      14      14
```

With R1.5, the TS7700 historical statistics, which are available through BVIR and VEHSTATS, show the amount of non premigrated data at the end of each reporting interval. This value is also available on the TS7700 management interface as a point-in-time statistic. Two host warning messages, low and high, can be configured for the TS7700 using the Host Console Request function. See *IBM Virtualization Engine TS7700 Series z/OS Host Command Line Request User's Guide*, which is available on the Techdocs website. Use the following keywords:

- ▶ SETTING, ALERT, RESDHIGH
- ▶ SETTING, ALERT, RESDLOW

Preferred Pre-migration Threshold

When this threshold is crossed and the number of premigration processes increases, the host throughput will tend to decrease from the peak I/O rate. Lowering this value will decrease the peak throughput period. This will also delay the amount of time before premigration throttling may occur. The purpose is to hopefully cause data to be premigrated faster and avoid the Pre-Migration Throttling threshold. See "Premigration tasks" on page 613 for details about how the premigration tasks are added. You might want to adjust this threshold lower to provide a larger gap between this threshold and the Pre-migration Throttling threshold. Do this if you want the gap to be larger but do not want to raise the Premigration Throttling Threshold. This threshold can be raised, along with the Pre-Migration Throttling Threshold, to defer premigration until after a peak period. This can improve the host IO rate because the premigration tasks are not ramped up as soon with lower threshold. This trades off an increased amount of un-premigrated data for a higher host IO rate during heavy production periods. The Preferred Pre-migration Threshold is set using the Host Console Request

function. The setting of this threshold is discussed in detail in the *IBM Virtualization Engine TS7700 Series z/OS Host Command Line Request User's Guide*, which is available on the Techdocs website. Use the SETTING, CACHE, PMPRIOR keywords.

Preferred Pre-migration Threshold

When this threshold is crossed and the number of premigration processes increases, the host throughput tends to decrease from the peak I/O rate. Lowering this value decreases the peak throughput period and also delays the amount of time before premigration throttling can occur. The purpose is to cause data to be premigrated faster and avoid the Pre-Migration Throttling threshold. See “Premigration tasks” on page 613 for details about how the premigration tasks are added.

You might want to adjust this threshold lower to provide a larger gap between this threshold and the Premigration Throttling threshold. Do this adjustment if you want the gap to be larger but do not want to raise the Premigration Throttling Threshold. This threshold can be raised, along with the Pre-Migration Throttling Threshold, to defer premigration until after a peak period. This step can improve the host I/O rate because the premigration tasks are not ready as soon with lower threshold. This trades an increased amount of non premigrated data for a higher host IO rate during heavy production periods. The Preferred Pre-migration Threshold is set by using the Host Console Request function. The setting of this threshold is discussed in detail in the *IBM Virtualization Engine TS7700 Series z/OS Host Command Line Request User's Guide*, which is available on Techdocs. Use the SETTING, CACHE, PMPRIOR keywords.

Pre-migration Throttling Threshold

When this threshold is crossed, the Host Write Throttle and Copy Throttle are both invoked. The purpose is to slow incoming data to allow the amount of non premigrated data to be reduced and not rise above this threshold. See section 9.3.2, “Host Write Throttle” on page 609 and 9.3.3, “Copy Throttle” on page 610 for more details on these throttles. You may want to adjust this threshold if there are periods where the amount of data entering the subsystem increases for a period of time and the existing threshold is being crossed for a short period of time. Raising the threshold will avoid the application of the throttles and keep host and copy throughput higher. However, the exposure is there will be more non premigrated data in cache. The extra non premigrated data will take a longer period to be premigrated. The customer needs to determine the balance. The Pre-migration Throttling Threshold is set using the Host Console Request function. Details about this setting are described *IBM Virtualization Engine TS7700 Series z/OS Host Command Line Request User's Guide* which is available on Techdocs website. Use the SETTING, CACHE, PMTHLVL keywords.

Disabling Host Write Throttle because of Immediate Copy

Host Write Throttle can be turned on because of immediate copies taking too long to copy to other clusters in the grid. Host Write Throttling is applied for various reasons including when the oldest copy in the queue is 10 or more minutes old. The TS7700 changes an immediate copy to immediate-deferred if the immediate copy has not started after 40 minutes in the immediate copy queue. The reason for this approach is to avoid triggering the 45-minute missing interrupt handler (MIH) in the host. When a copy is changed to immediate-deferred, the rewind/unload task is completed and the immediate copy becomes a high priority deferred copy. See “Immediate-copy set to immediate-deferred state” on page 615 for more information. You might decide to turn off Host Write Throttling because of immediate copies taking too long (if having the immediate copies take longer is acceptable). However, avoid the 40-minute limit where the immediate copies are changed to immediate-deferred. In grids where a large portion of the copies are immediate, better overall performance has been seen when the Host Write Throttle because of immediate copies is turned off. You are trading off

host IO for length of time required to complete an immediate copy. The enabling and disabling of the host write throttle because of immediate copies is discussed in detail in the *IBM Virtualization Engine TS7700 Series z/OS Host Command Line Request User's Guide* which is available on Techdocs. Use the SETTING, THROTTLE, ICOPYT keywords.

Making your cache deeper

A deeper cache will improve the likelihood of a volume being in cache for a recall. A cache-hit for a recall improves performance when compared to a cache-miss which requires a recall from physical tape. The TS7700 statistics provide a cache hit ratio for read mounts that can be monitored to make sure the cache-hit rate is not too low. Generally you want to keep the cache-hit ratio above 80%. Your cache can be made deeper in several ways:

- ▶ Add more cache.
- ▶ Use the Storage Class construct to utilize Preference Group 0 (PG0). PG0 volumes are removed from cache soon after they are premigrated to physical tape. PG0 volumes are actively removed from cache and do not wait for the cache to fill before being removed. This approach leaves more room for the PG1 volumes, which remain in cache as long as possible, to be available for recalls. Many customers have effectively made their cache deeper by examining their jobs and identifying which of them are most likely not to be recalled. Customers use Storage Class to assign these jobs to PG0.

Back-End Drives

Ensuring enough back-end drives are available is important. Below are general guidelines for the number of back-end drives versus the number of performance increments installed in the TS7740. If there are insufficient back-end drives, the performance of the TS7740 will suffer. As a guideline, use the ranges, listed in Table 9-2, of back-end drives based on the host throughput configured for the TS7740. The lower number of drives in the ranges is for scenarios that have few recalls, whereas the upper number is for scenarios that have numerous recalls. Remember, these are guidelines not rules.

Table 9-2 Performance increments versus back-end drives

Throughput	Back-end drives
100 MBps	4-6
200 MBps	5-8
300 MBps	7-10
400 MBps	9-12
500 MBps	10-14
600 MBps	12-16

Installing the correct number of back-end drives is important; also important is that the drives be available for use. Available means they are operational and might be idle or might be in use. The Host Console Request function can be used to set up warning messages for when the number of available drives drops. The setting of the Available Physical Drive Low and High Warning levels is discussed in detail in the *IBM Virtualization Engine TS7700 Series z/OS Host Command Line Request User's Guide* which is available on Techdocs. Use the SETTING, ALERT, PDRVLOW and SETTING, ALERT, PDRVCRT keywords.

9.8.2 Grid links

This section describes the grid link performance and the setting of the performance warning threshold.

Provide sufficient bandwidth

The customer network between the TS7700s must have sufficient bandwidth to account for the total replication traffic. For customers sharing network switches among multiple TS7700 paths or with other network traffic, the sum total of bandwidth on that network needs to be sufficient to account for all of the network traffic.

The TS7700 uses the TCP/IP protocol for moving data between each cluster. In addition to the bandwidth, there are other key factors which impact the throughput which the TS7700 can achieve. Factors that directly affect performance are as follows:

- ▶ Latency between the TS7700s
- ▶ Network efficiency (packet loss, packet sequencing, and bit error rates)
- ▶ Network switch capabilities
- ▶ Flow control to pace the data from the TS7700s
- ▶ Inter-switch link (ISL) capabilities such as flow control, buffering, performance

The TS7700s attempt to drive the network links at the full 1-Gb rate, which might be much higher than the network infrastructure is able to handle. The TS7700 supports the IP flow control frames, to have the network pace the rate at which the TS7700 attempts to drive the network. The best performance is achieved when the TS7700 is able to match the capabilities of the underlying network, resulting in fewer dropped packets.

Note: When the system attempts to give the network more data than it can handle, it begins to throw away the packets it cannot handle. This process causes TCP to have to stop, re-synchronize, and re-send amounts of data, resulting in a much less efficient use of the network.

To maximize network throughput, you must ensure the following items regarding the underlying network:

- ▶ The underlying network must have sufficient bandwidth to account for all network traffic expected to be driven through the system - eliminate network contention
- ▶ The underlying network must be able to support flow control between the TS7700s and the switches, allowing the switch to pace the TS7700 to the WAN's capability
- ▶ Flow control between the switches is also a potential factor, to ensure that they themselves are able to pace with each other's rate.
- ▶ Be sure that performance of the switch is capable of handling the data rates expected from all of the network traffic

In short, latency between the sites is the primary factor. However packet loss, because of bit error rates or because of the network not being capable of handling the 1-Gb per link, causes TCP to resend data, which multiplies the effect of the latency.

Grid link performance monitoring

The TS7700 generates a host message when it detects the grid performance is degraded. If the degraded condition persists, a call-home link is generated. The performance of two grid links is monitored periodically and if one link is performing worse than the other link, by an SSR alterable value, a warning message is generated and sent to the host. The purpose of this warning is to alert the customer that an abnormal grid performance difference exists. The

value must be adjusted so that warning messages are not generated because of normal variations of the grid performance. The default value is different for different code levels. For example, a setting of 75% means that if one link's performance is 25% lower than the other link, a warning message is generated. The grid link performance is available with the Host Console Request function and is available on the TS7700 management interface. The monitoring of the grid link performance using the Host Console Request is described in detail in the *IBM Virtualization Engine TS7700 Series z/OS Host Command Line Request User's Guide*, which is available on Techdocs. Use the STATUS, GRIDLINK keywords. The Grid Link Degraded threshold also includes two other values that can be set by the SSR:

- ▶ Number of degraded iterations: The number of consecutive five-minute intervals link degradation was detected before reporting an attention message.
- ▶ Generate Call Home iterations: The number of consecutive 5 minute intervals link degradation was detected before generating a Call Home. Figure 9-50 lists default and recommended values.

Description	Code Level	Default	Recommendation Two Cluster Grid	Recommendation Three or more Cluster Grid
Grid Link Degraded Threshold	R1.4a	75%	50%	10%
	R1.5 pre-PGA2	60%	50%	10%
	R1.5 PGA2+	60%	50%	50%
	R1.6	50%	50%	50%
Grid Link Degraded Iterations	R1.4a and above	3	9	9
Generate Call Home Iterations	R1.4a and above	9	12	12

Figure 9-50 Grid Link Degraded threshold

Use the default values unless you are receiving intermittent warnings. If you receive intermittent warnings, change the threshold and have the SSR change the iterations values to the recommended values. For some customers, the default and recommended values still must be adjusted. For example, a customer's clusters in a two-cluster grid are 2000 miles apart with a round trip latency of approximately 45 ms. The normal variation seen by this customer is 20 - 40%. The customer set the threshold value at 50% and the iterations to 9 and 12.

Reclaim operations

Reclaim operations consume two back-end drives per reclaim task. Reclaim operations also consume CPU MIPS. If needed, the TS7740 can allocate pairs of idle drives for reclaim operations being sure to leave one drive available for recall. Reclaim operations affect host performance, especially during peak workload periods. To tune your reclaim tasks using both the Reclaim Threshold and Inhibit Reclaim Schedule.

5.6.1 Reclaim threshold

The reclaim threshold directly affects how much data is moved during each reclaim operation. The default setting is 10% for each pool. Customers tend to raise this threshold too high because they want to store more data on their stacked volumes. The result is that reclaim

operations have to move larger amounts of data and consume back-end drive resources that are needed for recalls and premigration. After a reclaim task is started, it does not free up its back-end drives until the volume being reclaimed is empty. Figure 9-51 shows the reclaim threshold and the amount of data that must be moved, depending on the stacked tape capacity and the reclaim percentage. When the threshold is reduced from 40% to 20% only half of the data needs to be reclaimed, thus cutting the time and resources needed for reclaim in half.

Cartridge Capacity	Reclaim Threshold			
	10%	20%	30%	40%
300 GB	30 GB	60 GB	90 GB	120 GB
500 GB	50 GB	100 GB	150 GB	200 GB
640 GB	64 GB	128 GB	192 GB	256 GB
700 GB	70 GB	140 GB	210 GB	280 GB
1000 GB	100 GB	200 GB	300 GB	400 GB

Figure 9-51 Reclaim threshold

Inhibit Reclaim Schedule

Use the Inhibit Reclaim Schedule to inhibit reclaims during the customer’s busy periods, leaving back-end drives to be available for recalls and premigrates tasks. Be sure you start the inhibit time 30 - 60 minutes prior to the heavy workload period. This setting allows any started reclaim tasks to complete before the heavy workload period. Adjusting the Maximum Number of Reclaim Tasks Reclaim operations consume back-end drives, two per task, and consume CPU cycles. For this reason, use the Inhibit Reclaim Schedule to turn off reclaim operations during heavy production periods. When reclaim operations are not inhibited, you might want to limit the number of reclaim tasks. Perhaps there is moderate host I/O during the uninhibited period and reclaim is consuming too many back-end drives, CPU cycles, or both. The Host Library Request command has been expanded, starting with R1.6, to limit the number of reclaim tasks in the TS7740. The second keyword RECLAIM is added along with the third keyword of RCLMMAX. This expansion only applied to the TS7740. Also, the Inhibit Reclaim Schedule is still honored. The limit is turned off by setting the value to -1 (minus 1).

The maximum number of reclaim tasks is limited by the TS7740, based on the number of available back-end drives, as listed in Table 9-3, “reclaim tasks” on page 651.

Table 9-3 reclaim tasks

Number of available drives	Maximum number of reclaim tasks
3	1
4	1
5	1
6	2
7	2
8	3
9	3
10	4
11	4
12	5
13	5
14	6
15	6
16	7

Limiting number of premigration drives (max drives)

Each storage pool allows you to define the maximum number of back-end drives to be used for premigration tasks. There are several triggers that cause the TS7740 to ramp up the number of premigration tasks. If a ramp-up of premigration tasks occurs, followed by the need for more than one recall, the recall will have to wait until a premigration task is complete for a back-end drive to free up. A single premigration task can move up to 30 GB at one time. Having to wait for a back-end drive delays a logical mount that requires a recall. If this ramping up is causing too many back-end drives to be used for premigration tasks, you can limit the number of premigration drives in the pool properties panel.

- For systems before R1.5, the limit setting is in the Library Manager.
- For R1.5 and later systems, the limit setting is in the TS7740 MI. For Copy Export pools it is advisable that the maximum number of premigration drives be set appropriately. If the customer is exporting a small amount of data each day (one or two cartridges worth of data) you should limit the premigration drives to two. If more data is being exported, set the maximum to four. This setting limits the number of partially filled export volumes. Look at MB/GB written to a pool, compute MBps, compute maximum and average, and compute the number of premigration drives per pool. Base the number of drives by using 50 - 70 MBps per drive.

Avoid Copy Export during Heavy Production Periods

Because a Copy Export operation requires each physical volume to be exported to be mounted, the best approach is to perform the operation during a slower workload time.

9.9 TS7700 Virtualization Engine statistical data

The Virtualization Engine TS7700 statistics content and processes differ significantly from the VTS predecessor. In addition to completely new statistical record formats, major changes have been made to the frequency of collection, data storage, host data retrieval, and reporting tools.

9.9.1 Types of statistical records

The TS7700 Virtualization Engine provides two types of statistics:

- ▶ **Point-in-time (PIT) statistics:** These statistics are performance-related. The PIT information is intended to supply information about what the system is doing the instant that the request is made to the system. This information is not persistent on the system; the TS7700 updates these statistics every 15 seconds, but does not retain them. This information focuses on the individual components of the system and their current activity. These statistics report operations over the last full 15-second interval. The user can retrieve the PIT statistics from the TS7700 at any time, using the Bulk Volume Information Retrieval (BVIR) facility. A subset of point-in-time statistics is also available on the TS7700 management interface.
- ▶ **Historical (HIS) statistics:** These statistics encompass a wide selection of performance and capacity planning information. They are intended to help with capacity planning and tracking system use over an extended period of time. The information focuses more on the system as a whole, and the movement of data through the system. These statistics are collected by the TS7700 every 15 minutes, and stored for 90 days in a TS7700 database. The user can also retrieve these records by using BVIR.

Both point-in-time (PIT) and historical statistics are recorded. The PIT records present data from the most recent interval, providing speedometer like statistics. The historical statistics provide data that can be used to observe historical trends.

These statistical records are available to a host through the BVIR facility. See 9.10, “Bulk Volume Information Retrieval” on page 657 for more information about how to retrieve the statistics records and 9.11, “IBM Tape Tools” on page 706 for more information about how to format and analyze these records.

Each cluster in a Grid has its own set of PIT and historical statistics for both the vNode and hNode.

Point-in-time statistics

The data provided by this type of statistic is a snapshot of the TS7700 Virtualization Engine activity over the last 15-second interval. Each new 15-second interval data overlays the prior interval's data. But not all data is updated every 15 seconds (primarily hNode data). Those statistics contain Single and multi-cluster grid information.

You can obtain the point-in-time statistics through the appropriate BVIR request. The response returns the last PIT snapshot from the TS7700 Virtualization Engine. The data returned is not in a human readable format; it is primarily binary data, so use the `DCB=(RECFM=U,BLKSIZE=24000)` parameter. The records lengths differ, depending on the record type.

See the following resources:

- ▶ Appendix E, “IBM Virtualization Engine TS7700 statistical record format” on page 839
- ▶ The *IBM Virtualization Engine TS7700 Series Bulk Volume Information Retrieval Function User's Guide* white paper

The most recently published white papers are available at the Techdocs website by searching for TS7700 Virtualization Engine at the following address:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs>

Basically, the point-in-time statistics provide the following record types:

- ▶ vNode: Virtual Device
- ▶ vNode: Host Adapter
- ▶ hNode: HSM
- ▶ hNode: Grid

A variable number of data records are returned depending on the number of vNodes and hNodes (if in a multi-cluster grid configuration).

Each record has a common header that contains the following information:

- ▶ Record length
- ▶ Record version
- ▶ Record type
- ▶ Node and distributed library ID
- ▶ time stamp
- ▶ Machine type, model, and hardware serial number
- ▶ Code level

Point-in-time statistics: vNode virtual device

The kinds of point-in-time statistics provide the state of the reporting vNode:

- ▶ Overall state (online, offline, going online, and so on)
- ▶ Maximum configured throughput
- ▶ Number of installed virtual devices

In addition, the state and usage of each virtual drive is presented:

- ▶ Volume mounted or last mounted
- ▶ Distributed library access point
- ▶ Mount state (unloaded/mount in progress, failed, mounted, and so on)
- ▶ Device state (ready, write mode, BOT, and so on)
- ▶ Buffer wait condition count (whether the subsystem is pacing the channel or vice versa)
- ▶ Data transferred

Point-in-time statistics: vNode host adapter

These kinds of point-in-time statistics provide the status for each of the four host adapter positions:

- ▶ Adapter type installed
- ▶ State (online, offline, reloading, and so on)
- ▶ Location (drawer and slot)

In addition, the state and usage of each adapter port is provided:

- ▶ Port interface ID
- ▶ Data transferred before and after compression

Point-in-time statistics: hNode HSM

These kinds of point-in-time statistics provide the status for the management tasks running in the hNode:

- ▶ Premigrate and recall queue counts
- ▶ Throttling values (write, copy, and deferred copy)
- ▶ Library sequence number
- ▶ State and usage of the physical tape drives
 - Device type
 - Physical volume mounted and pool ID
 - Device state (offline or online)
 - Device role (idle, premig, recall, reclaim, and so on)
 - Logical volume being processed
 - Data transferred

Note: hNode HSM does not establish a relationship to the Hierarchical Storage Manager (DFSMSHsm), the z/OS storage management software.

For TS7720 Virtualization Engine statistics, because physical tape drives are not utilized, there is no information provided for functions related to back-end drive activity such as migration, recall, reclamation, and so on.

Point-in-time statistics: hNode Grid

These point-in-time statistics provide the Grid status from this hNode's perspective:

- ▶ Status and usage for the distributed library
 - Run and deferred copy queue counts for this distributed library
 - Active copies for this distributed library
 - Link status
- ▶ Data transferred to and from each cluster

Historical statistics

The data provided by this type of statistic is captured over a 15-minute interval in the TS7700 Virtualization Engine. Each new 15-minute interval data does not overlay the prior interval's data. However, not all data is updated every 15 minutes. Those statistics contain single and multi-cluster grid information. Up to 96 intervals per day can be acquired, and each day a file is generated that contains the historical statistics for that day. The historical statistics are kept in the TS7700 Virtualization Engine for 90 rolling days.

You can obtain the complete set or a subset of these historical statistics through the appropriate BVIR request (for more details, see “Historical statistics” on page 688). The request will specify the day or the days of needed historical statistics. For the current day, records up to the last 15-minute interval are returned. The data returned is not in a human readable format; it is primarily binary data, therefore use the following parameter:

DCB=(RECFM=U,BLKSIZE=24000)

The records are of different lengths depending on the record type. See Appendix E, “IBM Virtualization Engine TS7700 statistical record format” on page 839 for more information about the format of the statistics records. VEHSTATS is a tool that is available for formatting these records into reports. For more information about using VEHSTATS, see 9.11, “IBM Tape Tools” on page 706.

Basically, the historical statistics provide the following record types:

- ▶ vNode: Virtual Device
- ▶ vNode: Host Adapter
- ▶ hNode: HSM
- ▶ hNode: Library
- ▶ hNode: Grid
- ▶ hNode: Import/Export

The number of records returned varies depending on the number of vNodes and hNodes (if in a multi-cluster grid configuration).

Each record has a common header that contains the following information:

- ▶ Record length
- ▶ Record type
- ▶ Node and distributed library ID
- ▶ Time stamp
- ▶ Machine type, model, and hardware serial number
- ▶ Code level

Historical statistics: vNode virtual device

These types of historical statistics will provide the following information about the usage for a vNode's virtual devices:

- ▶ Number of installed virtual devices
- ▶ Virtual device type
- ▶ Blocksizes being written
- ▶ Configured throughput
- ▶ Min, max, and average virtual devices mounted

Historical statistics: vNode host adapter

You can use these historical statistics to determine the status for each of the two or four host adapter positions. The provided statistics contain the following information:

- ▶ Adapter type installed
- ▶ State (online, offline, and so on)
- ▶ Location (drawer or slot)

Both the state and usage of each adapter port are provided:

- ▶ Port interface ID
- ▶ Interface data transfer rate setting (capable and actual)
- ▶ Data transferred before and after compression
- ▶ Selective and system reset counts

Historical statistics: hNode HSM

This data portion within the historical statistics will give you hNode HSM information. You will get the VOLSER of the physical volume with the latest snapshot of the database on it.

You also see the state and status of the Tape Volume Cache. The provided statistics contain the following information:

- ▶ Usable size in GB
- ▶ Throttling values
- ▶ State of each cache partition
 - Partition size
 - Number of fast-ready, cache hit, or cache miss mounts
 - Average fast-ready, cache hit, or cache miss mount times

- Number of volumes in cache by preference group
- Space occupied by the volumes in cache by preference group
- Volume aging by preference group
- Migration and replication

Import/Export historical record

Information imported and exported is as follows:

- ▶ Physical volumes imported and exported
- ▶ Logical volumes imported and exported
- ▶ Amount of data imported and export

Historical statistics: hNode library

This part of the historical statistics give you library information, such as:

- ▶ Information about the attached physical library
- ▶ Information about the physical tape devices in the library
- ▶ Information about the common scratch pool media in the library
- ▶ Information about each physical volume pool in the library

Note: The TS7720 Virtualization Engine does not use a physical tape library, therefore hNode physical library information is not available from a TS7720 Virtualization Engine.

With a TS7740 Virtualization Engine, the attached physical tape library might have multiple subsystems sharing the TS3500 Tape Library through ALMS partitioning. Only data related to the partition attached to this TS7740 Virtualization Engine is provided in the hNode library report.

Historical statistics: hNode Grid

This last set of historical statistics gives you information in a multi-cluster grid environment for the following items:

- ▶ Status and usage for the distributed library
 - Number of logical volumes and data to copy for this distributed library
 - Average age of the deferred and immediate copy queues for this distributed library
 - Number of distributed libraries in the Grid configuration
- ▶ Data transferred to and from each cluster

9.9.2 Collecting and analyzing TS7700 Virtualization Engine statistical records

Consider the following aspects when you work with the statistics and reporting mechanisms that are provided by TS7700 Virtualization Engine:

- ▶ Getting the statistics and reports: The main interface to access statistic and reports from the TS7700 Virtualization Engine is the Bulk Volume Information Retrieval (BVIR) facility. Depending on the request, you will receive readable output or, for the TS7700 Virtualization Engine point-in-time and historical statistics, binary data. For binary data, further documentation and tools are needed. For details about how to use BVIR functions, see 9.10, “Bulk Volume Information Retrieval” on page 657.
- ▶ Formatting and displaying the information: Some of the response data of the BVIR functions is already in a readable format. For the remaining binary format data provided by the point-in-time and historical statistics, you need a formatting tool. IBM provides a tool called VEHSTATS. Further information about where to download this tool and how to use it is in 9.11, “IBM Tape Tools” on page 706.

9.10 Bulk Volume Information Retrieval

With the potential to support hundreds of thousands of logical volumes in a TS7700 Virtualization engine subsystem, providing a set of information for all of those volumes through normal channel control type commands is not practical. Luckily, the functions of a TS7700 Virtualization Engine subsystem that allows it to virtualize a tape volume, also allows for a simple and effective method to transfer the information to a requesting application. The TS7700 Virtualization Engine converts the format and storage conventions of a tape volume into a standard file managed by a file system within the subsystem.

The Bulk Volume Information Retrieval (BVIR) facility uses an IBM standard labeled tape volume to both initiate a request for information and return the results. By using a standard tape volume, no special interfaces or access methods are needed for an application to use this facility. In practice, no specific applications are required, as standard IBM utilities, such as IEBGENER, provide the function needed to request and obtain the information.

With BVIR, you are able to obtain information about all of the logical volumes managed by a TS7700 Virtualization Engine. The following data is available from a TS7700 Virtualization Engine:

- ▶ Volume Status Information
- ▶ Cache Contents Information
- ▶ Physical Volume to Logical Volume Mapping Information
- ▶ Point-in-Time Statistics
- ▶ Historical Statistics
- ▶ Physical Media Pools

For more information, see the following resources:

- ▶ *IBM Virtualization Engine TS7700 Series Bulk Volume Information Retrieval Function User's Guide:*

<http://w3.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101094>

- ▶ *IBM Virtualization Engine TS7700 Series Statistical Data Format White Paper:*

<http://w3.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100829>

The most recently published white papers are available on the Techdocs website by searching for TS7700 Virtualization Engine at either of the web addresses in the previous list.

9.10.1 Overview of the BVIR function

With the potential to support hundreds of thousands of logical volumes in a TS7700 Virtualization Engine, providing a set of information for all of those volumes through normal channel control type commands is not practical. Luckily, the functions of a TS7700 Virtualization Engine that allow it to virtualize a tape volume also provide a simple and effective method to transfer the information to a requesting application. The TS7700 Virtualization Engine converts the format and storage conventions of a tape volume into a standard file managed by a file system within the subsystem.

The BVIR function uses an IBM standard labeled tape volume to both initiate a request for information and return the results. By using a standard tape volume, no special interfaces or access methods are needed for an application to use this facility. In practice, no specific applications are required, as standard IBM utilities, such as IEBGENER, provide the function needed to request and obtain the information.

The following steps obtain information by using this function:

1. A single data set with the information request is written to a logical volume. The logical volume can be any logical volume in the subsystem from which the information is to be obtained. Either a scratch or specific volume request can be used. The data set contains a minimum of two records and a maximum of three records that specify the type of data being requested. The records are in human readable form, that is, lines of character data. The data set can be cataloged or uncataloged (although cataloging the data set can make it easier for subsequent access to the data). On closing the volume, the TS7700 Virtualization Engine server recognizes it as a request volume and “primes” the subsystem for the next step.

Note: Some information obtained through this function is specific to the cluster on which the logical volume is written, for example, cache contents or a logical-physical volume map. In a TS7700 Virtualization Engine Grid configuration with multiple clusters, use a Management Class for the volume to obtain statistics for a specific cluster.

Historical statistics for a multi-cluster grid can be obtained from any of the clusters.

2. The request volume is again mounted, this time as a specific mount. Seeing that the volume was primed for a data request, the TS7700 Virtualization Engine appends the requested information to the data set. The process of obtaining the information and creating the records to append can take up to several minutes, depending on the request and, from a host's viewpoint, is part of the mount processing time. After the TS7700 Virtualization Engine has completed appending to the data set, the host is notified that the mount has completed. The requested data can then be accessed like any other tape data set.

In a JES2 environment, the JCL to perform the two steps can be combined into a single job. However, in a JES3 environment, they must be run in separate jobs. This is because the volume will not be demounted and remounted between job steps in a JES3 environment.

After the response data set has been written to the request logical volume, that logical volume functions identically to any other logical volume in the subsystem. Subsequent mount requests and read accesses to the logical volume should have no effect on its contents. Subsequent mount requests and write accesses to the logical volume will overwrite its contents. It can be returned to scratch status and reused by any application.

Figure 9-52 shows the process flow of BVIR.

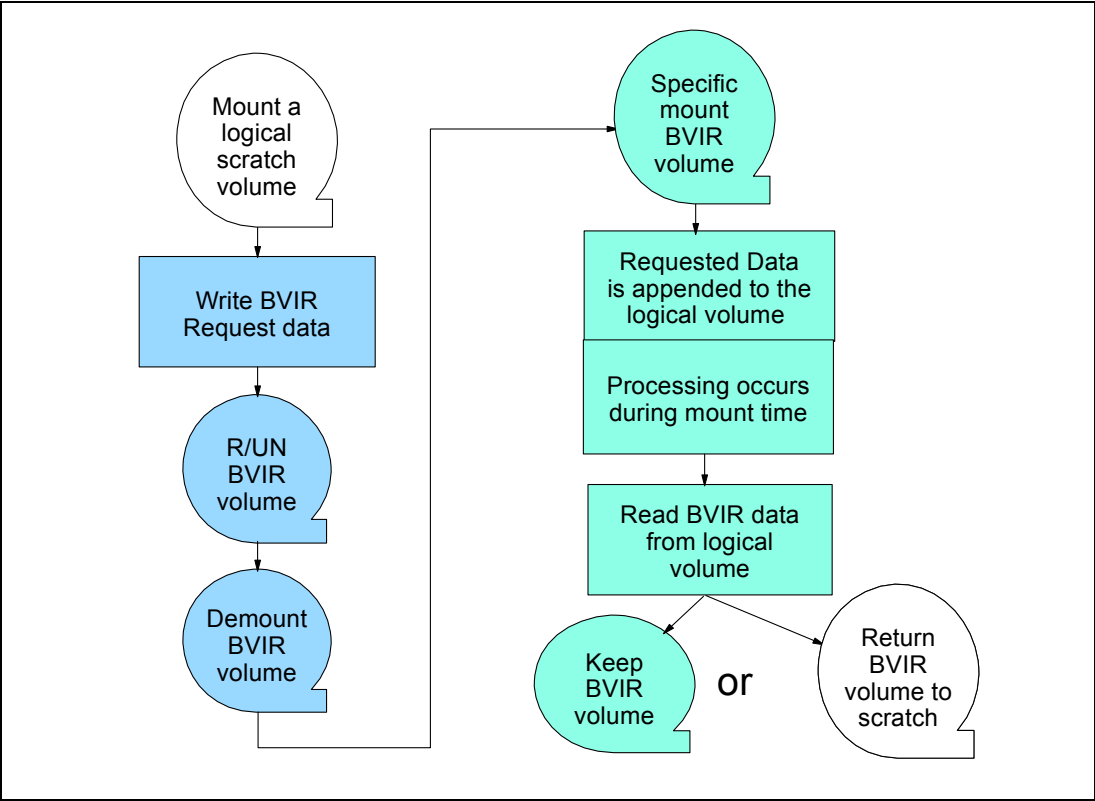


Figure 9-52 BVIR process flow

The building of the response information requires a small amount of resources from the TS7700 Virtualization Engine. Do not use the BVIR function to “poll” for a specific set of information and only issues one request at a time. Some requests, for example, the volume map, might take several minutes to complete, and to prevent “locking” out another request during that time, the TS7700 Virtualization Engine is designed to handle two concurrent requests. If more than two concurrent requests are issued, they will be processed as previous requests are completed.

While the requested data is always in a human readable format, depending on the request, the data returned from the TS7700 Virtualization Engine can be in human readable or binary form. See the response sections for the specifics of the returned data.

The general format for the request/response data set is shown in Example 9-1.

Example 9-1 BVIR output format

```
1234567890123456789012345678901234567890123456789012345
VTS BULK VOLUME DATA REQUEST
VOLUME MAP
11/20/2008 12:27:00 VERSION 02
S/N: 0F16F LIB ID: DA01A
```

PHYSICAL	LOGICAL	P/B	ORDER	PART	SIZE
P00024	GK0000	P	000001	1 OF 1	23.45 M
P00024	GK0020	P	000002	1 OF 1	76.50 M
P00024	GK0010	P	000003	1 OF 1	134.24 M

Note: When records are listed in this chapter, there will be an initial record showing “1234567890123...”. This record does not exist, but is provided to improve readability.

Record 0 is identical for all requests and not part of the output; it is for just supporting Records 1 through 5. Records 6 and above contain the requested output, which differs depending on the request:

- ▶ Records 1 and 2 contain the data request commands.
- ▶ Record 3 contains the date and time when the report was created and the version of BVIR.
- ▶ Record 4 contains the hardware serial number and the Distributed Library ID of the TS7700 Virtualization Engine.
- ▶ Record 5 contains all blanks.
- ▶ Record 6 contains the title of the output columns, as seen in Example 9-1 on page 659, for a Volume Map.
- ▶ Records 7 and above contain the requested data. The fields of these and the other reports are described in the next sections.

9.10.2 Prerequisites

Any logical volume defined to a TS7700 Virtualization Engine can be used as the request/response volume. Logical volumes in a TS7700 Virtualization Engine are formatted as IBM Standard Labeled volumes. Although a user can reformat a logical volume with an ANSI Standard Label or as an unlabeled tape volume, those formats are not supported for use as a request/response volume.

There are no restrictions regarding the prior use of a volume used as a request/response volume and no restrictions regarding its subsequent use for any other application. Use normal scratch allocation methods for each request (that is, use the DISP=(NEW,CATLG) parameter). In this way, any of the available scratch logical volumes in the TS7700 Virtualization Engine can be used. Likewise, when the response volume's data is no longer needed, the logical volume should be returned to scratch status through the normal methods (typically by deletion of the data set on the volume and a return-to-scratch policy based on data set deletion).

9.10.3 Request data format

Several types of data can be requested. The type of data requested is indicated in the request data set. The request data set must be the only data set on the volume and must be written with a record format of F and a logical record size of 80 bytes. Request information is in EBCDIC character form, beginning in the first character position of the record and padded with blank characters on the right to fill out the record.

Important: The request fields must be as shown. Not beginning with the first character position of the record or using extra blanks between words will result in a failed request.

Although the request data format uses fixed records, not all response records are fixed. For the point-in-time and historical statistics responses, the data records are of variable length and the record format used to read them is the Undefined (U) format. See Appendix F, “Sample JCLs” on page 889 for more information about this topic.

In a multi-site TS7700 Virtualization Engine Grid configuration, the request volume must be created on the cluster for which the data is being requested. The Management Class assigned to the volume needs to specify the particular cluster that is to have the copy of the request volume.

The format for the request data set records are listed in the following sections.

Record 1

Record 1 must contain the command exactly as shown in Example 9-2.

Example 9-2 BVIR Request Record 1

1234567890123456789012345678
VTB BULK VOLUME DATA REQUEST

The format for the request's data set records is shown inTable 9-4.

Table 9-4 BVIR Request Record 1

Record 1: Request Identifier		
Bytes	Name	Contents
1 - 28	Request Identifier	VTB BULK VOLUME DATA REQUEST
29 - 80	Blanks	Blank padding

Record 2

With Record 2, you can specify which data you want to obtain. The following options are available:

- ▶ VOLUME STATUS zzzzzz
- ▶ CACHE CONTENTS
- ▶ VOLUME MAP
- ▶ POINT IN TIME STATISTICS
- ▶ HISTORICAL STATISTICS FOR xxx
- ▶ HISTORICAL STATISTICS FOR xxx-yyy
- ▶ PHYSICAL MEDIA POOLS
- ▶ PHYSICAL VOLUME STATUS VOLUME zzzzzz
- ▶ PHYSICAL VOLUME STATUS POOL xx
- ▶ COPY AUDIT COPYMODE INCLUDE/EXCLUDE libids

The format for the request's data set records is shown in Table 9-5.

Table 9-5 BVIR Request Record 2

Record 2: Request Identifier		
Bytes	Name	Contents
1 - 80	Request	'VOLUME STATUS zzzzzz' or 'CACHE CONTENTS' or 'VOLUME MAP' or 'POINT IN TIME STATISTICS' or 'HISTORICAL STATISTICS FOR xxx-yyy' or 'PHYSICAL MEDIA POOLS' or 'PHYSICAL VOLUME STATUS VOLUME zzzzzz' or 'PHYSICAL VOLUME STATUS POOL xx' or 'COPY AUDIT COPYMODE INCLUDE/EXCLUDE libids' left justified, padded with blanks on the right

For the Volume Status and Physical Volume Status Volume requests, 'zzzzzz' specifies the volume serial number mask to be used. By using the mask, one to thousands of volume records can be retrieved for the request. The mask must be six characters in length, with the underscore character (_) representing a positional wildcard mask.

For example, assuming that volumes in the range of ABC000 through ABC999 have been defined to the cluster, a request of VOLUME STATUS ABC1_0 would return database records that exist for ABC100, ABC110, ABC120, ABC130, ABC140, ABC150, ABC160, ABC170, ABC180 and ABC190.

For the Historical Statistics request, 'xxx' specifies the Julian day being requested. Optionally, '-yyy' can also be specified and indicates that historical statistics from xxx through yyy are being requested. Valid days are 001 through 366 (to account for leap year). For leap years, February 29 is Julian day 060 and December 31 is Julian day 366, for other years, Julian day 060 is March 1, and December 31 is Julian day 365. If historical statistics do not exist for the day or days requested, that will be indicated in the response record (this can occur if a request is issued for a day prior to the day the system was installed, day or days the system was powered off, or after the current day before a rolling year has been accumulated). If a request spans the end of the year, for example a request that specified as HISTORICAL STATISTICS FOR 364-002, responses are provided for days 364, 365, 366, 001 and 002, regardless of whether the year was a leap year.

For Copy Audit, INCLUDE or EXCLUDE is specified to indicate which TS7700s clusters in a Grid configuration are to be included or excluded from the audit. COPYMODE is an option for taking a volume's copy mode for a cluster into consideration. If COPYMODE is specified, a single space must separate it from INCLUDE or EXCLUDE. The libid parameter specifies the library sequence numbers of the distributed libraries associated with each of the TS7700 clusters either to include or exclude in the audit. The parameters are separated by a comma. At least one libid parameter must be specified. For the Physical Volume Status Pool request, 'xx' specifies the pool for which the data is to be returned. If there are no physical volumes currently assigned to the specified pool, that will be indicated in the response record. Data might be requested for pools 0 through 32.

For point-in-time and historical statistics requests, any additional characters provided in the request record past the request itself are retained in the response data, but otherwise ignored. In a TS7700 Grid configuration, the request volume must be valid only on the specific cluster the data is to be obtained from. Use a specific Management Class that has a copy policy defined to indicate that only the desired cluster is to have a copy of the data. By ensuring that there is a sole copy of the request volume, any virtual device address on any of the clusters in the same Grid configuration can be used to request and access the data. You do not have to have host connectivity to the specific cluster. If a Management Class is used that indicates that more than one cluster is to have a valid copy of the request volume, unpredictable response data results can occur.

9.10.4 Response data format

When the request data set has been written to the volume and subsequently closed and demounted, when mounted again, the TS7700 Virtualization Engine validates the contents of the request volume and appends the requested data records to the data set.

Human readable appended records are 80 bytes in length. Binary data appended records can be variable in length of up to 32760 bytes. The data set is now a response data set. The appropriate block counts in the end of file (EOF) records will be updated to reflect the total number of records written to the volume.

After appending the records and updating the EOF records, the host that requested the mount is signaled that the mount is complete and can read the contents of the volume. If the contents of the request volume is not valid, either one or more error description records will be appended to the data set or the data set will be unmodified prior to signaling the host that the mount completed, depending on the problem encountered.

All human readable response records begin in the first character position of the record and are padded with blank characters on the right to fill out the record. All binary records are variable in length and are not padded.

Note: In the response records, the dates and times presented are all based on the internal clock of the TS7700 Virtualization Engine handling the request. The internal clock of a TS7700 Virtualization Engine is not synchronized to the host, but is synchronized with all other TS7700 Virtualization Engine.

The response data set contains both request records that we describe in 9.10.3, “Request data format” on page 660, and the response data set contains three explanatory records (Records 3 - 5) and, starting with Record 6, the actual response to the data request.

Records 1 - 2 and Record 3

The general format for the response data set is as follows:

- Records 1-2

Contains the contents of request records 1-2.

- Record 3

This record contains the date and time that the response data set was created and a format version number for the results, as shown in Example 9-3 and Table 9-6 on page 664.

Example 9-3 BVIR Response Record 3

123456789012345678901234567890
08/20/2009 12:27:00 VERSION 03

Table 9-6 BVIR Response Record 3

Bytes	Name	Description
1-10	Date	The date that the response was generated in the format of MM/DD/YYYY
11	Field Delimiter	Blank character
12-19	Time	The time the response was generated in the format HH:MM:SS
20	Field Delimiter	Blank character
21-27		"VERSION"
28	Field Delimiter	Blank character
29-30	Version Number	The Version Number of this response right-justified and zero-filled. Version number code level supported: 01 Original VTS 02 Hydra 1.5, Hydra 1.5PGA1–PGA3 03 Hydra 1.5 levels PGA4 and later, Hydra 1.6 04 Hydra 1.7
31-80	Blanks	Blank character padding

Record 4

This record contains both the five-character hardware serial number of the TS7700 Virtualization Engine and the five-character Distributed Library sequence number of the cluster that generated the response, as shown in Example 9-4 and Table 9-7.

Example 9-4 BVIR Response Record 4

123456789012345678901234
S/N: BA142 LIB ID: DA01A

Table 9-7 BVIR Response Record 4

Bytes	Name	Description
1-4		"S/N:"
5	Field Delimiter	Blank character
6-10	Serial Number	Hardware serial number of the TS7700 Virtualization Engine
11-12	Field Delimiter	Blank character
13-19		"LIB ID:"
20	Field Delimiter	Blank character
21-25	Distributed Library ID	Distributed library sequence number of the TS7700 Virtualization Engine
26-80	Blanks	Blank character padding

Bytes 6-10 contain a unique 5 character TS7700 Virtualization Engine hardware serial number that is configured during manufacturing of the TS7700 Virtualization Engine server node. The serial number is fixed in the subsystem microcode and reported as part of BVIR Response Record 4.

The TS7700 Virtualization Engine distributed library sequence number (Library ID) is a unique 5 character identifier assigned to each TS7700 Virtualization Engine Cluster during preinstallation planning. The IBM System Services Representative (SSR) uses the distributed library sequence number during installation configuration of the TS7700 Virtualization Engine. It is also labelled as the Library ID on the DFSMS Tape Library Define window or the Tape Library Display window.

Record 5

This record contains all blank characters.

Record 6-N

These records contain the specific response records based on the request. If the request could not be understood or was invalid, that will be indicated. See the following sections for the response record definitions.

9.10.5 Interpreting the BVIR response data

In this section, we explain how to interpret each BVIR Response Data Set for the specific request information, such as the following information:

- ▶ Volume Status Information
- ▶ Physical Volume Status Information
- ▶ Cache Contents Information
- ▶ Physical Volume to Logical Volume Mapping Information
- ▶ Point in Time Statistics
- ▶ Historical Statistics
- ▶ Physical Media Pools

Note: When records are listed in this chapter, an initial record shows “1234567890123...”. This record does not exist, but is provided to improve readability.

Volume Status Information

A database is maintained on each individual TS7740 Cluster that contains information related to the management of the logical volumes on the cluster and copy and resynchronization processes when the TS7700 Virtualization Engines are in a Grid configuration. Several returned database fields can be useful in handling operational exceptions at one or more clusters in a Grid configuration.

Data Inconsistent

This field indicates whether the cluster has a valid version of the data or not. If it indicates that the data on the logical volume is not valid, this means that the same volume on another TS7700 Virtualization Engine in the Grid has been modified and it has not yet been copied. For customers who use the deferred copy consistency point (which is typically when there is some significant distance between the TS7700 Virtualization Engines in the Grid configuration), there will be some number of volumes that are not consistent between the TS7700 Virtualization Engines at any point in time. If a situation occurs that renders the site inoperable where the source data resides, by issuing the Volume Status request to an operable TS7700 Virtualization Engine, this field can be used to identify the volumes that were not copied prior to the situation so that appropriate recovery steps can be performed for them.

MES Volume

This field indicates that the logical volume was created in the TS7700 Virtualization Engine Cluster or even created within a VTS, prior to being merged into a Grid configuration. Volumes that existed in a TS7700 Virtualization Engine Cluster prior to being included in a Grid configuration are not automatically copied to the other TS7700 Virtualization Engine Clusters in the configuration until they have been accessed and closed. This field could be used to determine which volumes in each TS7700 Virtualization Engine Cluster have not been copied, to build a set of jobs to access them, and force the copy. The PRESTAGE program from the TAPETOOL FTP site could support you in doing that job in an efficient way. The VEHSYNC job can be used to identify volumes needing copies.

Additional information about various tools available for monitoring your TS7700 Virtualization Engine is provided in 9.12.1, “VEHSTATS tool overview” on page 711. Access the TAPETOOL FTP site:

<ftp://ftp.software.ibm.com/storage/tapetool/>

Copy Required for Cluster n

This field indicates that a copy to another TS7700 Virtualization Engine Cluster in a Grid configuration is required. In cases where deferred mode copy is used, this field can be used to determine whether a critical set of volumes has completed their copy operations to specific clusters.

Volume Ownership and Volume Ownership Taken

At any point in time, a logical volume is owned by a specific cluster. If required, ownership is transferred as part of mount processing. If a logical volume is mounted on a virtual drive anywhere in the composite library, ownership will not be transferred until the volume is unloaded. Ownership can transfer in one of two ways:

- ▶ Through communication with the current owning cluster
- ▶ Through a recovery process called ownership takeover

Normally, if the cluster receiving a mount command does not have ownership of the volume, it requests the transfer of volume ownership from the current owning cluster. If the volume is not in use, then ownership is transferred.

However, if the cluster receiving the mount request cannot communicate with the owning cluster, that method does not work. In this case, the requesting clusters cannot determine whether the owning cluster has failed or just the Grid network links to it have failed. Operator intervention is required to indicate that the owning cluster has failed and that ownership takeover by the other clusters is allowed. The two types of ownership takeover are as follows:

- ▶ Write ownership takeover (WOT): The cluster taking over ownership of the volume has complete freedom to modify the contents of the volume or modify any of the properties associated with the volume. This includes scratch mounts.
- ▶ Read ownership takeover (ROT): The cluster taking over ownership of the volume is restricted to reading the volume's data only. Therefore, a cluster in ROT mode fails a scratch mount request for which it is unable to acquire volume ownership.

Current and Pending Category

One of the key properties associated with a volume is the *category* that it is assigned. The primary usage for category is to group scratch volumes together. A volume's category assignment changes as the volume is used. The current category field indicates the category the volume is assigned to within the TS7700 Virtualization Engine Integrated Library Manager function. The pending category field indicates that a new category assignment is in progress

for the volume. These fields can be used to determine whether the category assignments are in sync between the clusters and the host databases.

Data Deleted

As part of normal processing in a TS7700 Virtualization Engine Cluster, you can specify that after a certain period of time after being returned to scratch, the contents of a volume can be deleted. This field indicates whether or not the data associated with the volume has been deleted on the cluster.

Removal State

As part of normal processing in a TS7700 Virtualization Engine Grid configuration where a mixture of both TS7740 Virtualization Engine and TS7720 Virtualization Engine clusters exist, a data removal or migration process occurs where data is removed from TS7720 Virtualization Engine clusters to prevent TS7720 Virtualization Engine clusters from overrunning their tape volume cache. This field, and the removal time stamp, can be used to determine whether or not the data associated with the volume has been removed.

Hot

This field represents the cluster's view of which clusters have down level token or volume metadata information as a result of a cluster outage. When clusters are unavailable because of expected or unexpected outages, the remaining clusters mark the unavailable cluster for pending reconciliation by updating this hot mask. The field represents both Insert or Eject pending updates, or regular pending updates. Insert/Eject updates are related to volumes being inserted or ejected during the outage. Regular pending updates are for updates that occur to the volume during an outage as a result of normal operations such as host I/O. Each bit within the mask represents which clusters are viewed as needing reconciliation.

The volume status information returned represents the status of the volume on the cluster the requested volume was written. In a TS7700 Virtualization Engine Grid configuration, separate requests must be issued to each cluster to obtain the volume status information for the individual clusters. Using the volume serial number mask specified in the request, a response record is written for each matching logical volume that exists in the cluster. A response record consists of the database fields defined in Table 9-8 on page 668. Fields are presented in the order defined in the table and are comma-separated. From TS7700 Virtualization Engine R1.6 on, the overall length of each record was increased to 640 bytes with blank padding after the last field as needed. The first few fields of the record returned for VOLSER ABC123 would be as shown in Example 9-5.

Example 9-5 BVIR volume status information

```
123456789012345678901234567890123456789012345678901234567890123
ABC123,0,2009-04-22-11.56.45.871263,0,0,32,0,N,2548,N,8719,N...
```

Record 6-N

If the request is for a specific volume and it does not exist in the cluster's database, or the request is for a range of volumes and none of the volumes exist in the cluster's database, the record shown in Example 9-6 is returned.

Example 9-6 BVIR request for a non-existing specific volume

```
12345678901234567890123456789012345678901234567890123
NO VOLUME RECORD(S) EXIST FOR THE REQUESTED VOLUME(S)
```

For the requested volumes that do exist in the cluster's database, each of these records provides information for one logical volume, as shown in Example 9-5 on page 667 and according to Table 9-8.

Table 9-8 Decoding volume status information

Field name	Description
Volser	Six-character volume serial number (VOLSER).
domain_lock_cluster	The cluster identifier of the cluster that currently has ownership of the volume. Zero (0) is the identifier for the first cluster in a Grid.
domain_lock_timestamp	Timestamp of the last volume ownership change for the volume. The format of the time stamp is as follows: Year-Month-Day-Hour.Minute.Second.Microsecond, for example: 2006-05-23-19.34.23.876129 When a volume is inserted in to the library, this field is set to: 1970-01-01-
mounted_cluster	If the volume is mounted, this the cluster identifier of the cluster to which the mount request for the volume was issued. The identifier is 0 for the first cluster in a Grid. If the volume is not mounted, -1 is indicated.
mounted_vnode	If the volume is mounted, this is the vNode identifier of the cluster to which the mount request for the volume was issued. The identifier is 0 for the first VNode in a cluster. If the volume is not mounted, -1 is indicated.
mounted_device	If the volume is mounted, this is the virtual device number in the vNode of the cluster to which the mount request for the volume was issued. The identifier is 0 for the first virtual device in a VNode. If the volume is not mounted, -1 is indicated.
tvcluster	This is the cluster identifier associated with the tape volume cache being used for the mounted volume. With remote mounting support, this cluster may not be the same cluster as the mounted_cluster. The identifier is 0 for the first cluster in a Grid.
data_inconsistent	This field indicates whether the cluster has an inconsistent copy of the volume. Y indicates it does not have a valid copy, N indicates that it has a valid copy
data_level	This is the current data level value for the volume on the cluster. Every time a volume's data is modified, its data level is incremented. If the data level matches the corresponding data level current, this cluster's version of the volume is viewed as consistent
properties_inconsistent	This field indicates whether the cluster has inconsistent properties for the volume. Y indicates it has one or more properties that are down level, N indicates that it does not. The properties consist of the category and storage constructs assigned to the volume.
properties_level	The current properties level value for the volume on the cluster. Every time a volume's properties are modified, its data level is incremented.
cluster0_copy_required	This field indicates whether Cluster 0 needs to make a copy of the volume from this cluster. Y indicates it does, N indicates that either a copy is not required or one has already been made.

Field name	Description
cluster1_copy_required	This field indicates whether Cluster 1 needs to make a copy of the volume from this cluster. Y indicates it does, N indicates that either a copy is not required or one has already been made.
cluster2_copy_required	This field indicates whether Cluster 2 needs to make a copy of the volume from this cluster. Y indicates it does, N indicates that either a copy is not required or one has already been made.
cluster3_copy_required	This field indicates whether Cluster 3 needs to make a copy of the volume from this cluster. Y indicates it does, N indicates that either a copy is not required or one has already been made.
cluster4_copy_required	This field indicates whether cluster 4 needs to make a copy of the volume from this cluster. Y indicates it does, N indicates that either a copy is not required or one has already been made.
cluster5_copy_required	This field indicates whether cluster 5 needs to make a copy of the volume from this cluster. Y indicates it does, N indicates that either a copy is not required or one has already been made.
cluster6_copy_required	This field indicates whether cluster 6 needs to make a copy of the volume from this cluster. Y indicates it does, N indicates that either a copy is not required or one has already been made.
cluster7_copy_required	This field indicates whether cluster 7 needs to make a copy of the volume from this cluster. Y indicates it does, N indicates that either a copy is not required or one has already been made.
Cluster0_copy_mode	<p>This field indicates whether Cluster 0 is to have a copy of the volume and the copy consistency point defined for the volume. The values are as follows:</p> <ul style="list-style-type: none"> ▶ I: Rewind unload (RUN) copy consistency point ▶ D: Deferred copy consistency point ▶ N: No copy ▶ E: The volume was previously assigned a copy consistency point of rewind unload or deferred, but was changed to no copy and a private mount operation occurred against the volume.
Cluster1_copy_mode	<p>This field indicates whether Cluster 1 is to have a copy of the volume and the copy consistency point defined for the volume. The values are as follows:</p> <ul style="list-style-type: none"> ▶ I: Rewind unload (RUN) copy consistency point ▶ D: Deferred copy consistency point ▶ N: No copy ▶ E: The volume was previously assigned a copy consistency point of rewind unload or deferred, but was changed to no copy and a private mount operation occurred against the volume.
Cluster2_copy_mode	<p>This field indicates whether Cluster 2 is to have a copy of the volume and the copy consistency point defined for the volume. The values are as follows:</p> <ul style="list-style-type: none"> ▶ I: Rewind unload (RUN) copy consistency point ▶ D: Deferred copy consistency point ▶ N: No copy ▶ E: The volume was previously assigned a copy consistency point of rewind unload or deferred, but was changed to no copy and a private mount operation occurred against the volume.

Field name	Description
Cluster3_copy_mode	<p>This field indicates whether Cluster 3 is to have a copy of the volume and the copy consistency point defined for the volume. The values are as follows:</p> <ul style="list-style-type: none"> ▶ I: Rewind unload (RUN) copy consistency point ▶ D: Deferred copy consistency point ▶ N: No copy ▶ E: The volume was previously assigned a copy consistency point of rewind unload or deferred, but was changed to no copy and a private mount operation occurred against the volume.
Cluster4_copy_mode	<p>This field indicates whether Cluster 4 is to have a copy of the volume and the copy consistency point defined for the volume. The values are as follows:</p> <ul style="list-style-type: none"> ▶ I: Rewind unload (RUN) copy consistency point ▶ D: Deferred copy consistency point ▶ N: No copy ▶ E: The volume was previously assigned a copy consistency point of rewind unload or deferred, but was changed to no copy and a private mount operation occurred against the volume.
Cluster5_copy_mode	<p>This field indicates whether Cluster 5 is to have a copy of the volume and the copy consistency point defined for the volume. The values are as follows:</p> <ul style="list-style-type: none"> ▶ I: Rewind unload (RUN) copy consistency point ▶ D: Deferred copy consistency point ▶ N: No copy ▶ E: The volume was previously assigned a copy consistency point of rewind unload or deferred, but was changed to no copy and a private mount operation occurred against the volume.
Cluster6_copy_mode	<p>This field indicates whether Cluster 6 is to have a copy of the volume and the copy consistency point defined for the volume. The values are as follows:</p> <ul style="list-style-type: none"> ▶ I: Rewind unload (RUN) copy consistency point ▶ D: Deferred copy consistency point ▶ N: No copy ▶ E: The volume was previously assigned a copy consistency point of rewind unload or deferred, but was changed to no copy and a private mount operation occurred against the volume
Cluster7_copy_mode	<p>This field indicates whether Cluster 7 is to have a copy of the volume and the copy consistency point defined for the volume. The values are as follows:</p> <ul style="list-style-type: none"> ▶ I: Rewind unload (RUN) copy consistency point ▶ D: Deferred copy consistency point ▶ N: No copy ▶ E: The volume was previously assigned a copy consistency point of rewind unload or deferred, but was changed to no copy and a private mount operation occurred against the volume.

Field name	Description
mes_flag	<p>This field indicates whether the volume was part of a MES merge operation.</p> <ul style="list-style-type: none"> ▶ Y indicates the volume existed prior to merging the cluster into a Grid configuration. After code level 8.3 this value is no longer used and is replaced with W and M. ▶ W indicates the volume existed prior to merging the cluster into a Grid configuration and has not been accessed because the MES merge operation. This field is only set within the cluster that was previously not aware of the volume prior to the merge. ▶ M indicates the volume existed prior to merging the cluster into a Grid configuration and has been mounted/demounted without being modified. The volume will be copied to the clusters specified in the copy_mode fields. This field is only set within the cluster that was previously not aware of the volume prior to the merge. ▶ N indicates that the volume was not part of the MES merge operation or if it was, has since been modified and successfully copied to the clusters specified in the copy_mode fields.
recall_error	<p>This field indicates whether or not an error occurred while recalling this volume from a physical tape associated with the tvc_cluster. Y indicates that an error occurred the last time a recall was attempted with the volume. N indicates that the last recall of the volume was successful.</p>
read_error	<p>This field indicates whether the data associated with the volume is corrupted. Y indicates that the volume's data is corrupted. N indicates that the data is not corrupted. A volume's data is determined to be corrupted if an error is detected when reading the data.</p>
disaster_rec	<p>This field indicates the disaster recovery state for the volume. Y indicates that the volume has been through the disaster recovery process. N indicates that no disaster recovery processing has been done or is required for the volume.</p>
data_deleted	<p>This field indicates whether the data associated with the volume has been deleted by the delete expired volume data function. Y indicates that the data has been deleted. N indicates that the data has not been deleted.</p>
volume_damaged	<p>The volume/token is determined to be damaged or unrecoverable without human intervention. Y indicates that the volume/token has been damaged. N indicates that the volume/token has not been damaged.</p>
Mount_operation_time	<p>Time stamp of when the volume was last mounted or unloaded. The format of the time stamp is Year-Month-Day-Hour.Minute.Second.Microsecond, for example: 2009-09-14-10.27.48.545823 If the volume has not yet been mounted, this field is set to: 1970-01-01-00.00.00.000000</p>
Media_type	<p>The media type defined when the volume was inserted. The values are as follows:</p> <ul style="list-style-type: none"> ▶ 0: Cartridge System Tape (400 MB) ▶ 1: Enhanced Capacity Cartridge System Tape (800 MB)

Field name	Description
takeover_active	<p>This field indicates whether the volume's ownership has been taken over. Bits 22-23 indicates the ownership transfer method.</p> <ul style="list-style-type: none"> ▶ bits 0-7: Reserved ▶ bits 8-15: Steal ID ▶ bits 16-19: Transfer version ▶ bits 20-21: Reserved ▶ bits 22-23: Takeover mode <ul style="list-style-type: none"> 0 - Write Ownership Takeover (WOT) 1 - Read Ownership Takeover (ROT) 2 - Service Ownership Takeover (SOT) ▶ bits 24-27: Cluster ID which was the original owner ▶ bits 28-31: Cluster ID which took control of this volume
pending_category	This is the category the volume is to be set to within the Integrated Library Manager of the cluster.
current_category	<p>This is the category the volume is currently assigned to in the Integrated Library Manager of the cluster.</p> <p>When this value and the pending_category are equal, then the category is consistent</p>
category_time	<p>Prior to code R1.5, this time stamp represents when the pending_category field was last changed. For R1.5 and later, this time stamp represents when the volume's category was last changed within the composite library. The format of the time stamp is Year-Month-Day-Hour.Minute.Second.Microsecond, for example: 2009-09-14-10.27.48.545823</p> <p>Prior to the first change in category, this field is set to: 1970-01-01-00.00.00.000000</p>
takeover_timestamp	<p>This is the time stamp of when the volume's ownership was last changed because of write or read ownership takeover. The format of the time stamp is Year-Month-Day-Hour.Minute.Second.Microsecond, for example: 2009-09-14-14.27.48.545823</p> <p>If the volume has not had its ownership taken, this field is set to: 1970-01-01-00.00.00.000000</p>
subcmd_cluster	If an LSVC or LSVA is currently active against the volume, this field represents the cluster identifier of the cluster the LSVC or LSVA request for the volume was issued to. Zero is the identifier for the first cluster in a Grid. If the volume has no active LSVC or LSVA against it, -1 is indicated.
subcmd_vnode	If an LSVC or LSVA is currently active against the volume, this field represents the vnode identifier of the cluster to which the LSVC or LSVA request for the volume was issued. Zero is the identifier for the first vnode in a cluster. If the volume has no active LSVC or LSVA against it, -1 is indicated.
subcmd_device	If an LSVC or LSVA is currently active against the volume, this field represents the virtual device in the vNode of the cluster to which the LSVC or LSVA request for the volume was issued. Zero is the identifier for the first virtual device in a vNode. If the volume has no active LSVC or LSVA against it, -1 is indicated.
data_level_current	This is the highest known level in the Grid.

Field name	Description
data_level_last_modify	<p>Defines a mask to show what cluster, hnode, vnode, and device last wrote to the volume:</p> <ul style="list-style-type: none"> ▶ bit 0: Reserved ▶ bits 1-3: Cluster ▶ bits 4-7: hNode ▶ bits 8-15: vNode ▶ bits 16-31: Logical device
properties_level_current	This is the highest known level in the Grid.
properties_level_last_modify	<p>Defines a mask to show what cluster, vnode, and device last issued an LSVC/LSVA:</p> <ul style="list-style-type: none"> ▶ bit 0: Reserved ▶ bits 1-3: Cluster ▶ bits 4-7: hNode ▶ bits 8-15: vNode ▶ bits 16-31: Logical device
subcmd_operation_time	<p>This time stamp represents when an LSVC or LSVA was issued to this volume. The format of the time stamp is Year-Month-Day-Hour.Minute.Second.Microsecond, for example: 2009-09-14-14.27.48.545823</p> <p>Prior to the first LSVA/LSVA for the volume, this field is set to: 1970-01-01-00.00.00.000000</p>
hot	<p>Defines a hot mask that shows the definition of regular hot and insert hot:</p> <ul style="list-style-type: none"> ▶ bit 0-15: Reserved ▶ bits 16-23: Token Insert/Eject Hot From right to left, each bit position represents a cluster index starting with cluster0. If a cluster index is set to 1, that cluster is viewed as token hot because of an insert or eject operation. ▶ bits 24-31: Token Standard Hot From right to left, each bit position represents a cluster index starting with cluster0. If a cluster index is set to 1, that cluster is viewed as token hot due a standard operation.
hot_creation_time	<p>Indicates the time stamp of when the hot mask became non-zero. The format of the time stamp is: Year-Month-Day-Hour.Minute.Second.Microsecond, for example: 2009-09-14-14.27.48.545823</p> <p>Prior to the first change, this field is set to: 1970-01-01-00.00.00.000000</p>
hot_update_time	<p>The time stamp of when the hot mask had at least one bit set. The format of the time stamp is Year-Month-Day-Hour.Minute.Second.Microsecond, for example: 2009-09-14-14.27.48.545823</p> <p>Prior to the first change, this field is set to: 1970-01-01-00.00.00.000000</p>

Field name	Description
steal_conflict	<p>When two clusters communicate for the first time after an outage, an initial volume token update occurs where volumes that might have been taken over during the outage are flagged. This field determines the result of this initial token update that must occur prior to a cluster moving to the available state:</p> <ul style="list-style-type: none"> ▶ Y: This volume might have been taken over by a peer cluster. Token reconciliation must resolve the condition before host access is allowed. ▶ N: Ownership takeover is suspected for this volume. ▶ X: An illegal or double takeover was detected. The volume will be put into the volume damaged state.
close_inconsistent	The volume is currently mounted by a device within the composite library. If the mounted_device value is not -1, then this cluster has the volume mounted.
last_properties_modify_time	<p>This field represents a time stamp of when this cluster's volume last had its local properties updated either through a host operation or through a reconciliation process. The format of the time stamp is Year-Month-Day-Hour.Minute.Second.Microsecond, for example: 2009-09-14-14.27.48.545823</p> <p>Prior to the first change, this field is set to: 1970-01-01-00.00.00.000000</p>
last_data_modify_time	<p>This field represents a time stamp of when this cluster's volume became consistent. The consistent_type field determines how it became consistent. The format of the time stamp is Year-Month-Day-Hour.Minute.Second.Microsecond, for example: 2009-09-14-14.27.48.545823</p> <p>Prior to the first change, this field is set to: 1970-01-01-00.00.00.000000</p>
expire_time	<p>If delete-expired, this time represents when it was deleted. The format of the time stamp is Year-Month-Day-Hour.Minute.Second.Microsecond, for example: 2009-09-14-14.27.48.545823</p> <p>Prior to the first change, this field is set to: 1970-01-01-00.00.00.000000</p>
consistent_type	<p>If the volume is consistent, this field states what method was used to bring this volume up to consistency:</p> <ul style="list-style-type: none"> ▶ U: The method is unknown. ▶ T: This volume was the actual version written to directly by the last host write operation. ▶ C: Grid replication was used to bring the volume up to consistency. ▶ M: An MES process was used to bring the volume up to consistency. ▶ R: The volume was recovered as part of a disaster recovery event.

Field name	Description
removal_timestamp	In a TS7700 Virtualization Engine hybrid grid configuration, TS7720 Virtualization Engine clusters may remove volumes from tape volume cache after replicating to peer clusters. If the removal state shows that this volume was removed, this time stamp represents the time of removal. The format of the time stamp is Year-Month-Day-Hour.Minute.Second.Microsecond, for example: 2009-09-14-14.27.48.545823 Prior to the first change, this field is set to: 1970-01-01-00.00.00.000000
removal_state	In a TS7700 Virtualization Engine hybrid grid configuration, TS7720 Virtualization Engine clusters may remove volumes from tape volume cache after replicating to peer clusters. This field represents the current state of the volume within this cluster. TS7740 Virtualization Engine clusters will always state N. <ul style="list-style-type: none"> ▶ N: The local TS7720 Virtualization Engine cluster has not attempted to remove this volume. ▶ R: The local TS7720 Virtualization Engine cluster removed this volume. ▶ X: The local TS7720 Virtualization Engine cluster determined that removal is not possible for this volume. ▶ D: The local TS7720 Virtualization Engine attempted to remove this volume, but the state of the Grid did not allow it to complete. An attempt will be made at a later time.
world_id	This field represents an ASCII version of a 12 byte unique world wide identifier associated with the volume. Only volumes bound as logical worm volumes utilize this field.
write_mount_count	This field represents how many times this volume, under the current world ID instance, was mounted and at least one write operation occurred.
removal_policy	This field represents the value assigned to the logical volume through the constructs that were introduced with code level 8.7.x.x. The allowed values are as follows: <ul style="list-style-type: none"> ▶ 0 - Prefer remove If removal is enabled, volumes assigned to this removal policy will be removed first after the minimum retention time has passed. The order of which the volumes are removed is in least recently used order when the TS7720 is reaching full capacity. ▶ 1 - Prefer keep Group 1 (LRU) (Default value and only value allowed in code level 8.6.x.x) - If removal is enabled, volumes assigned to this removal policy will be removed second after all "Prefer Remove" candidates have been removed. The same minimum retention time must have elapsed and the same least recently used access is used to determine order. ▶ 4 - Pinned (Disabled) If removal is enabled, volumes assigned to this removal policy are never removed. A pinned volume is removed from cache. Note: Independent of being Pinned or having a Minimum Retention Timed defined, all volumes that have been returned to scratch become preferred candidates for removal over volumes which remain private.
Blanks	As needed to pad record to 640 bytes

Cache Contents Information

Volumes that are accessed by a host are maintained in the tape volume cache, which is managed by each cluster. The cache may be partitioned into up to eight partitions. The TS7700 Virtualization Engine controls the movement of logical volumes out of a cache partition as space is needed for newly created or recalled volumes for that partition. The primary goal of the cache management algorithms in the TS7700 Virtualization Engine is to maximize the utilization of its cache for volumes that have some likelihood to be accessed again. The cache management function of the TS7700 Virtualization Engine arranges the volumes in a cache partition in the anticipated order they are to be removed when space is needed. To remove a volume from cache it must first have been premigrated (which means copied to a physical tape). For this reason, it is possible that volumes with a higher order number are removed from cache first. As part of the Advanced Policy Management functions of the TS7700 Virtualization Engine, the Storage Class construct provides for customer control of the partition for a volume's data and cache preferencing policies for the management of the volume in cache.

Two preferencing policies are supported:

- Preference group 0 (PG0)

When space is needed in the cache, premigrated volumes assigned to PG0 are removed from cache before volumes assigned to preference group 1. Within PG0, the volumes are ordered for removal from cache by largest volumes first.

Note: Volumes that are assigned to PG0 may also be removed from the cache, independent of the need for cache space, as a background task within the TS7700 Virtualization Engine.

- Preference Group 1 (PG1)

When space is needed in the cache and there are no premigrated PG0 volumes to remove, premigrated volumes that are assigned to PG1 are removed. Within PG1, the volumes are ordered for removal from cache based on time since last access (LRU).

Note: The order of removal of a volume from cache may also be influenced by other storage constructs settings for a volume, so the order that is presented in the response data should not be relied on to be exact.

The contents of the cache that are associated with the specific cluster that the request volume is written to are returned in the response records. In a TS7700 Grid configuration, separate requests must be issued to each cluster to obtain the cache contents of all of the clusters. The response records are written in 80-byte fixed block format.

Note: The generation of the response might take several minutes to complete depending on the number of volumes in the cache and how busy the TS7700 cluster is at the time of the request.

Record 6

This record provides a heading for the data records to follow. See Table 9-9.

Table 9-9 Record 6 layout

Bytes	Name	Description
1-3		Blank characters
4-8		'ORDER' Heading for the volume order column
9	Field Delimiter	Blank Character
10-15		'VOLSER' Heading for the volume serial number column
16	Field Delimiter	Blank characters
17-34		'DATE/TIME IN CACHE' Heading for the data/time stamp columns
35-38	Field Delimiter	Blank characters
39-40		'PG' Heading for the preference group column
41-42	Field Delimiter	Blank character
43-46		'PART' Heading for the partition column
47-48	Field Delimiter	Blank character
49-52		'SIZE' Heading for the LVOL size column
53-80	Blanks	Blank character padding

Record 7-N

Each record provides information for one logical volume. The cache contents information is provided by partition, then in the order in which the volumes are anticipated to be removed from the cache partition:

- ▶ 1 VOL020 11/30/2005 11:57:00 0 0 799997952
- ▶ 2 VOL019 11/29/2005 03:00:00 0 0 799997952
- ▶ 3 VOL023 11/20/2005 09:57:00 1 0 799997952
- ▶ 4 VOL016 11/20/2005 10:01:00 1 0 799997952
- ▶ 5 ABC309 11/20/2005 17:31:00 0 1 799997952
- ▶ 6 ABC333 11/20/2005 11:44:00 1 1 799997952
- ▶ 7 ABC789 11/20/2005 10:51:00 1 1 799997952
- ▶ 8 ABC234 11/20/2005 04:27:00 1 1 983040
- ▶ 9 ABC045 11/19/2005 21:45:00 1 1 9994240

Table 9-10 shows Record 7-N.

Table 9-10 Record 7-N layout

Bytes	Name	Description
1-8	Order in Cache Partition	The order in which volumes are to be removed from the cache partition, right justified and blank filled.
9	Field Delimiter	Blank character
10-15	Logical Volser	Six character volume serial number
16	Field Delimiter	Blank character
17-26	Data in Cache	The date the volume was created or recalled into the cache in the format of MM/DD/YYYY.
27	Field Delimiter	Blank character
28-35	Time in Cache	The time the volume was created or recalled into the cache in the format of HH:MM:SS.
36-39	Field Delimiter	Blank characters
40	Preference Group	Contains the preference group the volume is assigned.
41-45	Field Delimiter	Blank characters
46	Partition	Contains the partition in which the volume is resident.
47-48	Field Delimiter	Blank characters
49-60	LVOL size	Size of the volume in cache after compression, right-justified and blank-filled. This data is specified in bytes as a decimal value.
61-80	Blanks	Blank character padding

Note: The contents of the cache typically are all private volumes; however, some might have been returned to scratch status soon after being written. The VTS does not filter the cache contents based on the private or scratch status of a volume.

Historical statistics

A TS7700 Virtualization Engine is continually logging information regarding the activities within it. The logged information is referred to as statistical information and is recorded in two forms, point-in-time (PIT) and historical. Historical statistics indicate the operational aspects of the TS7700 Virtualization Engine accumulated over a 15-minute interval of time. The data from each 15-minute interval is maintained and logged within the TS7700 Virtualization Engine. A request for historical statistics results in a response file that contains all of the data logged up to that point for the requested julian day. Other than an information header, Historical statistics are provided in character and binary format fields. The sizes and format of the statistical records are defined in the *IBM Virtualization Engine TS7700 Series Statistical Data Format White Paper*.

<http://w3.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100829>

The historical statistics for all clusters are returned in the response records. In a TS7700 Grid configuration, this way means that the request volume can be written to any cluster to obtain the information for the entire configuration.

Notes:

- ▶ If a cluster or node is not available at the time the historical statistics are recorded, except for the headers, all the data fields for that cluster or node will be zeroes
- ▶ The TS7700 Virtualization Engine retains 90 days worth of historical statistics. If you want to keep statistics for a longer period of time, be sure you retain the logical volumes that are used to obtain the statistics.
- ▶ The request records are written in fixed block (FB) format. To read the response records, use the undefined (U) format with a maximum blocksize of 24000. The response records are variable in length.

Record 6-N

Table 9-11 shows Record 6-N. The historical statistical response records for each day that is requested are preceded by the following human readable header:

HISTORICAL STATISTICS FOR xxx or NO HISTORICAL STATISTICS FOR xxx

If historical statistics are requested for a day which the TS7700 Virtualization Engine does not have historical data, the No Historical Statistics for xxx header is the only response record provided for the day.

Table 9-11 Record 6-N

Bytes	Name	Description
1-46		HISTORICAL STATISTICS for xxx or NO HISTORICAL STATISTICS FOR xxx, left-justified, padded with blanks on the right.
47-80	Blanks	Blank character padding

Data response records

Each data response record provides historical information for the nodes of the TS7740 Cluster for the day indicated in the header record (see Table 9-12). Character data is encoded using EBCDIC and non-character data is in binary format. Records vary in length. The first four bytes of the record identify the length, version, and data type of the record. Because the historical information is provided for all nodes in the configuration, including nodes in all clusters in a Grid configuration, the number of records returned depends on the configuration. There are two records for each vNode and from four to seven records for each hNode (depends on the number of physical libraries attached to the cluster). The application that processes the response volume records must not assume that the records for a vNode or hNode are presented in any specific order.

Table 9-12 Historical information

Bytes	Name	Description
1-2	Length	Total length of the response record in binary
3	Version	Binary 1
4	DataType	Identifies the type of response record data
5-length	Data	Response data

Data type

The statistics for a node in the subsystem are subdivided into various data types.

Physical volume information

A database is maintained on each individual TS7740 cluster that contains information related to the management of the physical volumes on the cluster. The physical volume status information that is returned represents the status of the volume or volumes on the cluster to which the request volume is written. In a TS7700 Grid configuration, separate requests must be issued to each cluster to obtain the physical volume status information for the individual clusters. A response record is written for each physical volume, selected based on the volume serial number mask or pool number specified in the request, that exists in the cluster. A response record consists of the database fields defined in the following table. Fields are presented in the order defined in the table and are comma-separated (.). The overall length of each record is 400 bytes with blank padding after the last field as needed. For example, the first few fields of the record returned for VOLSER A03599 would be as follows:

A03599,2,FULL,READ-WRITE,2007-05-05-06.40.08.030061,2007-05-04-13.45.15.918473,...

Note: The generation of the response might take several minutes to complete depending, on the number of volumes requested and how busy the TS7700 cluster is at the time of the request.

Record 6-N

If this is a TS7720 configuration, the response is returned, as in Table 9-13.

Table 9-13 Record 6-N response

Bytes	Name	Description
1-57		NOT SUPPORTED IN A DISK-ONLY TS7700 VIRTUALIZATION ENGINE
58-80	Blanks	Blank character padding

If the request is for a specific physical volume and it does not exist in the cluster's database, or the request is for a range of physical volumes and none of the volumes exist in the cluster's database, or the request is for a pool and there are no physical volumes currently resident in the pool, the following record is returned, as in Table 9-14:

NO VOLUME RECORD(S) EXIST FOR THE REQUESTED VOLUME(S)

Table 9-14 No volume record bytes

Bytes	Name	Description
1-53		'NO VOLUME RECORD(S) EXIST FOR REQUESTED VOLUME(S)'
	Blanks	Blank character padding

For the requested volumes that do exist in the cluster's database, each of these records provides information for each physical volume, as shown in Table 9-15.

Table 9-15 Decoding volume record Information

Field Name	Description
VOLSER	Six-character volume serial number
CURRENT_POOL	This field indicates the pool to which the volume is currently assigned. Pool 0 is the common scratch pool. Pools 1-32 are the specific data pools.
VOLUME_STATUS	<p>This field indicates the volume's current capacity state. The following values are the ones that can be indicated:</p> <ul style="list-style-type: none"> ▶ EMPTY: The volume contains no data and is available for use as a physical scratch volume. ▶ FILLING: The volume contains valid data, but is not yet full. It is available for additional data. ▶ FULL: The volume contains valid data. At some point, it was marked as full and additional data cannot be added to it. A volume can be marked full in some cases short of the volume capacity limit. ▶ UNKNOWN: The volume's capacity state is unknown.
VOLUME_ACCESS	<p>This field indicates the volume's current accessibility state. The following values are the ones that can be indicated:</p> <ul style="list-style-type: none"> ▶ READ-WRITE: The volume can be read from or written to. ▶ READ-ONLY: The volume contains valid data, but had encountered read or write errors that exceeded an allowable threshold. It is only accessible for reading. ▶ UNAVAILABLE: The volume is unavailable to the TS7740 Virtualization Engine. ▶ DAMAGED: The volume has been damaged physically or logically such that it cannot be mounted or read from. ▶ COPY-EXPORTED: The volume has been exported using the Copy Export function and is not currently resident in the library.
STARTED_EXPIRING_TIMESTAMP	<p>This is the time stamp of when the data first started expiring on the physical volume. The format of the time stamp is Year-Month-Day-Hour.Minute.Second.Microsecond, for example: 2009-05-23-19.34.23.876129</p> <p>This field is set to 1970-01-01-00.00.00.000000 when the volume becomes empty (returned to scratch) or is inserted into the library.</p>
BECAME_EMPTY_TIMESTAMP	<p>This is the time stamp of when all of the data on the volume last became invalid and the volume is empty. The format of the time stamp is Year-Month-Day-Hour.Minute.Second.Microsecond, for example, 2009-05-23-19.34.23.876129.</p> <p>When the volume is inserted into the library, this field is set to 1970-01-01-00.00.00.000000.</p>
BECAME_FULL_TIMESTAMP	<p>This is the time stamp of when the volume is last marked as full. The format of the time stamp is Year-Month-Day-Hour.Minute.Second.Microsecond, for example, 2009-05-23-19.34.23.876129.</p> <p>When the volume is inserted into the library, this field is set to 1970-01-01-00.00.00.000000.</p>

Field Name	Description
RECORDING_FORMAT	<p>This field indicates the format that was used in recording the data on the volume. The following values are those that can be reported:</p> <ul style="list-style-type: none"> ▶ 4: The volume is written in the 3592 J1A format. ▶ 5: The volume is written in the TS1120 E05 format. ▶ 6: The volume is written in the TS1120 E05E format (encrypted). ▶ 7: The volume is written in the TS1130 E06 format. ▶ 8: The volume is written in the TS1130 E06E format (encrypted). <p>Values not defined are reserved.</p>
PCT_UTILIZED	<p>This field indicates the percent of the volume that contains active data. The percentage is determined using the active_size and the total_bytes_written fields. The number is reported in 1/10th of a percent. This field is updated hourly (if needed). It is reset to 0 when the volume becomes empty (scratch).</p>
TOTAL_BYTES_WRITTEN	<p>This field indicates the number of bytes that were written to the volume when it was filled. It is reset to 0 when the volume becomes empty (scratch).</p>
ACTIVE_SIZE	<p>This field indicates the number of bytes of the active data on the volume. It is reset to 0 when the volume becomes empty (scratch).</p>
ERROR_TYPE_FLAG	<p>This field indicates whether the TS7740 Virtualization Engine has detected an error with the use of the volume. The error values are used by IBM Support.</p>
SCRATCH_COUNT	<p>This field indicates the number of times the volume has been returned to scratch status since it was inserted into the library.</p>
MOUNT_COUNT	<p>This field indicates the number of times the volume has been mounted since it was inserted into the library.</p>
MEDIA_TYPE	<p>This field indicates the media type for the volume. The following values can be reported:</p> <ul style="list-style-type: none"> ▶ 5: IBM JA ▶ 6: IBM JJ ▶ 7: IBM JB <p>Values not defined are reserved.</p>
BATE_ROR_REASON	<p>If the volume has a volume access value of Read-Only, this field provides additional information about the reason for the access value. A volume may be marked as Read-Only for reasons other than a media error. These values might be used by IBM Support.</p>
DVM_ROR_REASON	<p>If the volume has a volume access value of Read-Only, this field provides additional information about the reason for the access value. A volume may be marked as Read-Only for reasons other than a media error. These values might be used by IBM Support.</p>
ERASE_FLAG	<p>This field indicates whether or not the volume needs to be subjected to secure data erase.</p>
MOUNT_FAILURE_COUNT	<p>This field indicates the number of times a mount failure occurred with this volume since it was inserted into the library.</p>
LAST_DEVICE_MOUNTED	<p>This field contains the device number on which the volume was last mounted.</p>

Field Name	Description
LAST_WRITE_TIMESTAMP	The time stamp of when the volume was last written to. The format of the time stamp is Year-Month-Day-Hour.Minute.Second.Microsecond, for example, 2009-05-23-19.34.23.876129. When the volume is inserted into the library, this field is set to 1970-01-01-00.00.00.000000.
LAST_READ_TIMESTAMP	The time stamp of when the volume was last read from. The format of the time stamp is Year-Month-Day-Hour.Minute.Second.Microsecond, for example, 2009-05-23-19.34.23.876129. When the volume is inserted into the library, this field is set to 1970-01-01-00.00.00.000000.
KEKLABEL1	If the volume has been encrypted, this field contains the first key label. A key label can be up to 64 characters in length. If the volume is not encrypted, this field is null. It is set to null when the volume becomes empty (scratch).
KEKLABEL2	If the volume has been encrypted, this field contains the second key label. A key label can be up to 64 characters in length. If the volume is not encrypted, this field is null. It is set to null when the volume becomes empty (scratch).
INSERT_TIMESTAMP	The time stamp of when the volume was inserted into the library. The format of the time stamp is Year-Month-Day-Hour.Minute.Second.Microsecond, for example, 2009-05-23-19.34.23.876129.
ADSM_FORMAT_FLAG	This field indicates whether the physical volume is written using the pre-TS7740 Virtualization Engine (B10/B20 VTS) for TS7740 Virtualization Engine format. The following are the values that can be indicated: <ul style="list-style-type: none"> ▶ 0: The volume is written in the TS7740 Virtualization Engine format. ▶ 1: The volume is written in the pre-TS7740 Virtualization Engineformat.
Blanks	As needed to pad a record to 400 bytes.

Physical volume to logical volume mapping information

The TS7740 Virtualization Engine maintains the mapping between logical and physical volumes in a database on each cluster. It is possible that there are inconsistencies in the mapping information provided with this function. This situation occurs when a logical volume is being moved from one physical volume to another one. For a period of time, the volume is shown on more than one physical volume. This can result in a small number of logical volumes reported as being on physical volumes which they were located on in the past, but presently are not.

Even with some inconsistencies, the mapping data is useful to customers that want to design jobs that recall data efficiently off of physical volumes. If the logical volumes reported on a physical volume are recalled together, the efficiency of the recalls will be increased. If a logical volume with an inconsistent mapping relationship is recalled, it will recall correctly, but an additional amount of a different physical volume may be required.

The physical volume to logical volume mapping information is returned in the response records for those physical volumes managed by the cluster on which the request volume is

written. In a TS7740 Virtualization Engine multi-cluster grid configuration, separate requests must be issued to each cluster to obtain the mapping for all physical volumes.

The response records are written in 80-byte fixed format.

Note: The generation of the response can take several minutes to complete depending on the number of active logical volumes in the library and how busy the TS7700 cluster is at the time of the request.

The first few fields of the record returned for VOLSER A03599 would, for example, be as follows:

A03599,2,FULL,READ-WRITE,2009-05-05-06.40.08.030061,2009-05-04-13.45.15.918473,...

Record 6

When this request is issued to a TS7720 Virtualization Engine where there are no physical tape resources, the response has no relevant data to return. When this request is issued to a TS7720 Virtualization Engine, the record shown in Example 9-7 is returned. Use Table 9-16 to decode the response.

Example 9-7 BVIR message when Volume Mapping is issued to TS7720 Virtualization Engine

```
123456789012345678901234567890123456789012345678901234567
NOT SUPPORTED IN A DISK-ONLY TS7700 VIRTUALIZATION ENGINE
```

Table 9-16 Decoding of Volume Mapping issued to TS7720 Virtualization Engine

Bytes	Name	Description
1-57		NOT SUPPORTED IN A DISK-ONLY TS7700 VIRTUALIZATION ENGINE
58-80	Field Delimiter	Blank characters

This record provides a heading for the data records to follow (Example 9-8 and Table 9-17 on page 685).

Example 9-8 BVIR heading for the data records for Volume Mapping information

```
123456789012345678901234567890123456789012345678
PHYSICAL LOGICAL P/B ORDER PART SIZE
```

Table 9-17 Decoding of the data records for Volume Mapping information

Bytes	Name	Description
1-8		PHYSICAL Heading for the physical volume serial number column
9-10	Field Delimiter	Blank characters
11-17		LOGICAL Heading for the logical volume serial number column
18	Field Delimiter	Blank characters
19-21		P/B Heading for the primary or backup pool indicator column
22-23	Field Delimiter	Blank characters
24-28		ORDER Heading for the logical volume order column
29-31	Field Delimiter	Blank characters
32-35		PART Heading for the logical volume spanning indicator column
36-44	Blanks	Blank characters
45-48		SIZE Heading for the logical volume compresses size
49-80	Blanks	Blank character padding

The primary and backup pool indicator column indicates whether the volume being reported resides on a primary pool volume or the secondary pool volume, as defined using the Selective Dual Copy function that is part of the TS7740 Virtualization Engine's advanced policy management function.

The size field reports the number of MBs, rounded to two decimal places, that a logical volume occupies on the physical volume. This includes the effect of compression performed on the data by the TS7740 Virtualization Engine, but does not include any effect of the compression performed by the physical drive. Any volume with a size of less than 5 KB will report a size of 0.00. When a volume indicates that it spans, the size shown represents the amount of the logical volume data that is stored on each physical volume.

Note: Only data created on prior generation B10 or B20 VTSs and migrated under the control of a TS7740 Virtualization Engine will indicate that they span from one physical volume to another. Volumes created or recopied by the TS7740 Virtualization Engine *do not span* physical volumes.

Record 7-N

Each of these records provides information for a logical to physical volume association. The records are ordered alphanumerically (0-9, A-Z) by physical VOLSER, then by logical sequence of the active logical volumes on the physical volume. See Example 9-9; Table 9-18 shows the decoding.

Example 9-9 BVIR data records for Volume Mapping information

12345678901234567890123456789012345678901
P00024 GK0000 P 000001 1 OF 1 23.45 M
P00024 GK0020 P 000002 1 OF 1 76.50 M
P00024 GK0010 P 000003 1 OF 1 145.70 M
P00024 GK0030 P 000004 1 OF 1 670.32 M
P00024 GK0040 P 000005 1 OF 1 1934.12 M
P00024 GK0060 P 000006 1 OF 1 0.00 M
P00024 GK0050 P 000007 1 OF 2 540.12 M
P00467 GK0050 P 000001 2 OF 2 540.12 M

Table 9-18 Decoding of the data records for Volume Mapping information

Bytes	Name	Description
1-6	Physical Volser	The physical VOLSER on which the logical VOLSER is located, left justified and padded with blanks.
7-10	Field Delimiter	Blank character.
11-16	Logical Volser	Logical VOLSER, left justified and padded with blanks.
17-19	Field Delimiter	Blank characters.
20	Pool Indicator	P indicates that the logical volume is the primary copy, while B indicates that the volume is the backup copy.
21-22	Field Delimiter	Blank character.
23-28	Order	The relative order of the logical volume on the physical volume.
29	Field Delimiter	Blank characters.
30-35	Spanning	Indicates whether the logical volume spans to another physical volume or not, and if so, which part.
36-41	Blanks	Blank characters.
42-45	Size	Integer part of the size, right justified and padded with leading blanks.
46	Decimal Point	“.”
47-48	Hundredths	Decimal part of the size.
49	Field Delimiter	Blank character.
50	Unit Indicator	M indicates the size is in megabytes, which is 1024x1024 bytes.
51-80	Blanks	Blank character padding.

Point-in-time statistics

A TS7700 Virtualization Engine is continually logging information regarding its activities. The logged information is referred to as statistical information and is recorded in two forms:

- Point-in-time
- Historical

Point-in-time statistics indicate the state and operation aspects of the TS7700 Virtualization Engine over a short interval of time. The time interval is currently approximately 15 seconds. A request for point-in-time statistics will respond with the data accumulated in the interval completed just prior to the request being processed. Because of this situation, the state information reported might lag behind the actual state of the TS7700 Virtualization Engine.

Other than an information header, point-in-time statistics are provided in a mixture of character and binary format fields. The record sizes and format of the statistical records are defined in the *IBM Virtualization Engine TS7700 Series Statistical Data Format White Paper*.

The point-in-time statistics for all clusters are returned in the response records. In a TS7700 Virtualization Engine Grid configuration, this means that the request volume can be written to any cluster to obtain the information for the entire configuration.

Note: The request records are written in fixed block (FB) format. To read the response records, use the undefined (U) format with a maximum blocksize of 24000. The response records are variable in length.

Record 6

This record provides a human readable header for the response record (Example 9-10 and Table 9-19).

Example 9-10 BVIR heading for the data records for point-in-time statistics

123456789012345678901234
POINT IN TIME STATISTICS

Table 9-19 Decoding of the data records for point-in-time statistics

Bytes	Name	Description
1-24		POINT IN TIME STATISTICS
25-80	Blanks	Blank character padding

Record 7-N

Each of these records provides point-in-time information for the nodes of the TS7700 Virtualization Engine Cluster. Character data is encoded using EBCDIC and non-character data is in binary format. Records vary in length. The first four bytes of the record identify the length, version, and data type of the record.

Because the point-in-time information is provided for all nodes in the configuration, including nodes in all clusters in a Grid configuration, the number of records returned depends on the configuration. There are two records for each vNode and two records for each hNode. The application processing the response volume records should not assume that the records for a vNode or hNode are presented in any specific order.

Table 9-20 lists the decoding of the response command.

Table 9-20 Decoding the data records for point-in-time statistics

Bytes	Name	Description
1-2	Length	Total length of the response record in binary.
3	Version Number	Version Number of this response.
4	Data Type	Identifies the type of response record data.
5-length	Data	Response data.

Data type

The statistics for a node in the subsystem are subdivided into different data types. See the *Virtualization Engine TS7700 Series Statistical Data Format* white paper. The most recently published white papers are available at the Techdocs website by searching for TS7700 Virtualization Engine at the following address:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs>

Binary response data

This record contains the actual response data. See the *Virtualization Engine TS7700 Series Statistical Data Format* white paper for the format of the response data for each data type. The most recently published white papers are available at the Techdocs website by searching for TS7700 Virtualization Engine at the following address:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs>

Historical statistics

Historical statistics indicate the operational aspects of the TS7700 Virtualization Engine accumulated over a 15-minute interval. The data from each 15-minute interval is maintained and logged in the TS7700 Virtualization Engine. A request for historical statistics results in a response file that contains all of the data logged up to that point for the requested date.

Other than an information header, historical statistics are provided in character and binary format fields. The format of the statistical records are defined in the *Virtualization Engine TS7700 Series Statistical Data Format* white paper. The most recently published white papers are available at the Techdocs website by searching for TS7700 Virtualization Engine at the following address:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs>

The historical statistics for all clusters are returned in the response records. In a TS7700 Virtualization Engine Grid configuration, this means that the request volume can be written to any cluster to obtain the information for the entire configuration.

Note: The TS7700 Virtualization Engine retains 90 days worth of historical statistics. If you want to keep statistics for a longer period of time, retain the logical volumes used to obtain the statistics.

The request records are written in fixed block (FB) format. To read the response records, use the Undefined (U) format with a maximum blocksize of 24,000. The response records are variable in length.

Record 6-N

The historical statistical response records for each day requested are preceded by a human-readable header. If historical statistics are requested for a day for which the TS7700 Virtualization Engine does not have historical data, the NO HISTORICAL STATISTICS FOR xxx, (where xxx is the Distributed Library ID) header is the only response record provided for the day (Example 9-11 and Table 9-21).

Example 9-11 BVIR heading for the data records for historical statistics

```
1234567890123456789012345678901234567890123456789012345
HISTORICAL STATISTICS FOR xxx or NO HISTORICAL STATISTICS FOR xxx
```

Table 9-21 Decoding of the data records for historical statistics

Bytes	Name	Description
1-46		HISTORICAL STATISTICS FOR xxx or NO HISTORICAL STATISTICS FOR xxx, left justified, padded with blanks on the right
47-80	Blanks	Blank character padding

Data response records

Each of these records provides historical information for the nodes of the TS7700 Virtualization Engine Cluster for the day indicated in the header record. Character data is encoded using EBCDIC and non-character data is in binary format. Records vary in length. The first four bytes of the record identify the length, version, and data type of the record. Because the historical information is provided for all nodes in the configuration, including nodes in all clusters in a Grid configuration, the number of records returned depends on the configuration. There are two records for each vNode and from four to seven records for each hNode (depending on the number of physical libraries attached to the cluster). The application processing the response volume records should not assume that the records for a vNode or hNode are presented in any specific order. Table 9-22 provides the information for decoding the response command.

Table 9-22 Decoding the data response for historical statistics

Bytes	Name	Description
1-2	Length	Total length of the response record in binary.
3	Version Number	Version Number of this response.
4	Data Type	Identifies the type of response record data.
5-length	Data	Response data.

Data type

The statistics for a node in the subsystem are subdivided into different data types. See the *IBM Virtualization Engine TS7700 Series Statistical Data Format* white paper for the description of the data types.

Physical media pools

The TS7700 Virtualization Engine supports separating the physical volumes it manages into pools. The supported pools include a pool that contains scratch (empty) volumes that are common and up to 32 pools that may contain scratch (empty) and data (filling/full) volumes. Pools can borrow and return volumes from the common scratch pool. Each pool can contain several types of media.

For pool 00 (common scratch pool), because it only contains an empty volume, only the empty count is returned. Volumes that have been borrowed from the common pool are not included.

For pools 1-32, a count of the physical volumes that are empty, are empty and waiting for erasure, are in the process of being filled, or have been marked as full, is returned. The count for empty includes physical volumes that have been specifically assigned to the pool and volumes that were borrowed from the common scratch pool but have not yet been returned. The count of volumes that are marked as Read Only or Unavailable (including destroyed volumes) are returned. Also, the full data volumes contain a mixture of valid and invalid data. Response records are provided for the distribution of active data on the data volumes marked as full for a pool.

Information is returned for the common pool and all other pools that are defined and have physical volumes associated with them.

The physical media pool information managed by the specific cluster the request volume is written to is returned in the response records. In a TS7700 Virtualization Engine Grid configuration, separate requests must be issued to each cluster to obtain the physical media pool information for all clusters.

The response records are written in 80-byte fixed format. Counts are provided for each media type associated with the pool (up to a maximum of eight) and the response lines are formatted, as shown in “Record 6” on page 690 through “Record N+2-M” on page 694.

Record 6

When this request is issued to a TS7720 Virtualization Engine where there are no physical tape resources, then the response has no relevant data to return. When this request is issued to a TS7720 Virtualization Engine configuration, the record shown in Example 9-12 is returned. Use Table 9-23 to decode the response.

Example 9-12 BVIR message when Volume Mapping is issued to TS7720 Virtualization Engine

```
123456789012345678901234567890123456789012345678901234567
NOT SUPPORTED IN A DISK-ONLY TS7700 VIRTUALIZATION ENGINE
```

Table 9-23 Decoding of Physical Media Pooling request issued to TS7720 Virtualization Engine

Bytes	Name	Description
1-57		NOT SUPPORTED IN A DISK-ONLY TS7700 VIRTUALIZATION ENGINE
58-80	Field Delimiter	Blank characters

This record provides a heading for the data records to follow (Example 9-13 and Table 9-24 on page 691).

Example 9-13 BVIR heading for the data records for Physical Media Pooling

```
1234567890123456789012345678901234567890123456789012345678
R POOL MEDIA EMPTY FILLING FULL ERASE ROR UNAVAIL
```

Table 9-24 Decoding of data records for Physical Media Pooling

Bytes	Name	Description
1		R Heading for the record type column
2	Field Delimiter	Blank character
3-6		POOL Heading for the pool number column
7	Field Delimiter	Blank character
8-12		MEDIA Heading for the media type column
13-14	Field Delimiter	Blank character
15-19		EMPTY Heading for the empty count column
20-21	Field Delimiter	Blank character
22-28		FILLING heading for the filling count column
29-31	Field Delimiter	Blank character
32-35		FULL Heading for the full count column
36-37	Field Delimiter	Blank character
38-42		ERASE Heading for the waiting to erase count column
43-46	Field Delimiter	Blank characters
47-49		ROR Heading for the read only recovery count column
50-51	Field Delimiter	Blank characters
52-58		UNAVAIL Heading for the unavailable count column
59-80	Filed Delimiter	Blank character padding

Record 7-N

Each of these records provides the counts for a specific media type for each pool. Up to eight media types may be reported for a pool, as shown in Example 9-14.

Example 9-14 BVIR response for the data records for Physical Media Pooling

123456789012345678901234567890123456789012345678

C	00	J	00343					
C	00	K	00523					
C	00	JA	00083					
C	00	JB	00051					
C	00	JJ	00002					
C	01	JA	00134	00003	00609	00000	00000	00000
C	01	JB	00001	00002	00001	00006	00000	00000
C	02	K	00023	00002	00325	00000	00000	00000
C	03	J	00103	00001	00336	00004	00000	00000
C	03	K	00034	00000	00185	00002	00001	00000

Use Table 9-25 to decode the data records shown in Example 9-14 on page 691.

Table 9-25 Decoding of the data records for Physical Media Pooling

Bytes	Name	Description
1	Record Type	Record type identifier. C stands for media count.
2-4	Field Delimiter	Blank characters.
5-6	Pool Number	The pool number. The pool number is right justified and padded with a leading zero.
7-10	Field Delimiter	Blank characters.
11-12	Media Type	The media type defined for the pool. Media types with a single character are padded on the left with a blank.
13-14	Field Delimiter	Blank characters.
15-19	Empty Count	The count of the physical volumes that are empty for the media type. The count is right justified and padded with leading zeroes.
20-23	Field Delimiter	Blank characters.
24-28	Filling Count	The count of the physical volumes that are in the filling state for the media type. The count is right justified and padded with leading zeroes. This field is all blanks for pool 00.
29-30	Field Delimiter	Blank characters.
31-35	Full Count	The count of the physical volumes that have been marked full for the media type. The count is right justified and padded with leading zeroes. This field is all blanks for pool 00.
36-37	Field Delimiter	Blank characters.
38-42	Erase Count	The count of the physical volumes that have been reclaimed, but need to be erased before they become empty. The count is right justified and padded with leading zeroes. This field is all blanks for pool 00.
43-44	Field Delimiter	Blank characters.
45-49	ROR Count	The count of the physical volumes that are in the read only recovery state.
50-53	Field Delimiter	Blank characters.
54-58	Unavailable Count	The count of the physical volumes that are in the unavailable or destroyed state.
59-80	Blanks	Blank character padding.

Record N+1

For pools 1-32, this record provides a heading for the data records to follow, as shown in Example 9-15.

Example 9-15 BVIR response for the data records N+1 for Physical Media Pooling

12345678901234567890123456789012345678901234567890123456789012
R POOL MEDIA 0+ 10+ 20+ 30+ 40+ 50+ 60+ 70+ 80+ 90+

Table 9-26 describes the layout of the record.

Table 9-26 Decoding of the data record N+1 for Physical Media Pooling

Bytes	Name	Description
1		R Heading for the record type column
2	Field Delimiter	Blank character
3-6		POOL Heading for the pool number column
7	Field Delimiter	Blank characters
8-12		MEDIA Heading for the media type column
13-16	Field Delimiter	Blank characters
17-18		0+ Heading for the 0-10% count column
19-21	Field Delimiter	Blank characters
22-24		10+ Heading for the 10-20% count column
25-27	Field Delimiter	Blank characters
28-30		20+ Heading for the 20-30% count column
31-33	Field Delimiter	Blank characters
34-36		30+ Heading for the 30-40% count column
37-39	Field Delimiter	Blank characters
40-42		40+ Heading for the 40-50% count column
43-45	Field Delimiter	Blank characters
46-48		50+ Heading for the 50-60% count column
49-51	Field Delimiter	Blank characters
52-54		60+ Heading for the 60-70% count column
55-57	Field Delimiter	Blank characters
58-60		70+ Heading for the 70-80% count column
61-63	Field Delimiter	Blank characters
64-66		80+ Heading for the 80-90% count column
67-69	Field Delimiter	Blank characters
70-72		90+ Heading for the 90+% count column
73-80	Field Delimiter	Blank character padding

Record N+2-M

Each of these records provides the active data distribution counts for a specific media type. Up to eight media types may be reported for a pool. Example 9-16 lists the first three pools.

Example 9-16 BVIR response for the data records N+2-M for Physical Media Pooling

123456789012345678901234567890123456789012345678901234567890123456789012

```
D 01 JA 00000 00002 00034 00056 00092 00078 00084 00065 00195 00183
D 01 JB 00000 00000 00000 00000 00000 00000 00000 00002 00003 00001
D 02 K 00000 00000 00000 00000 00003 00023 00033 00087 00076 00103
D 03 J 00000 00000 00000 00002 00023 00046 00003 00056 00072 00134
D 03 K 00000 00000 00003 00007 00015 00012 00041 00018 00024 00065
```

Table 9-27 describes the layout of the records.

Table 9-27 Decoding of the data records N+2-M for Physical Media Pooling

Bytes	Name	Description
1	Record Type	Record type identifier. D stands for active media data distribution count.
2-4	Field Delimiter	Blank characters.
5-6	Pool Number	This is the pool number, right justified and padded with a leading zero.
7-10	Field Delimiter	Blank characters.
11-12	Media Type	The media type defined for the pool. Media types with a single character are padded on the left with a blank.
13	Field Delimiter	Blank character.
14-18	0-9 Count	This is the count of the physical volumes that are marked full and have an active data percentage greater than 0 and less than 10. The count is right justified and padded with leading zeroes.
19	Field Delimiter	Blank character.
20-24	10s Count	The count of the physical volumes that are marked full and have an active data percentage greater than 10 and less than 20. The count is right justified and padded with leading zeroes.
25	Field Delimiter	Blank character.
26-30	20s Count	The count of the physical volumes that are marked full and have an active data percentage greater than 20 and less than 30. The count is right justified and padded with leading zeroes.
31	Field Delimiter	Blank character.
32-36	30s Count	The count of the physical volumes that are marked full and have an active data percentage greater than 30 and less than 40. The count is right justified and padded with leading zeroes.
37	Field Delimiter	Blank character.
38-42	40s Count	The count of the physical volumes that are marked full and have an active data percentage greater than 40 and less than 50. The count is right justified and padded with leading zeroes.
43	Field Delimiter	Blank character.

Bytes	Name	Description
44-48	50s Count	This is the count of the physical volumes that are marked full and have an active data percentage greater than 50 and less than 60. The count is right justified and padded with leading zeroes.
49	Field Delimiter	Blank character.
50-54	60s Count	This is the count of the physical volumes that are marked full and have an active data percentage greater than 60 and less than 70. The count is right justified and padded with leading zeroes.
55	Field Delimiter	Blank character.
56-60	70s Count	This the count of the physical volumes that are marked full and have an active data percentage greater than 70 and less than 80. The count is right justified and padded with leading zeroes.
61	Field Delimiter	Blank character.
62-66	80s Count	This is the count of the physical volumes that are marked full and have an active data percentage greater than 80 and less than 90. The count is right justified and padded with leading zeroes.
67	Field Delimiter	Blank character.
68-72	90s Count	This is the count of the physical volumes that are marked full and have an active data percentage greater than 90. The count is right justified and padded with leading zeroes.
73-80	Blanks	Blank character padding.

Active Data Distribution

The Active Data Distribution information displays the distribution of the data on the stacked volumes. There is a number for each 10% increment in active data for the number of used stacked tapes. You use that output to identify the number of tapes that will be freed up if you were to change the reclaim threshold.

With prior VTSs, the Active Data Distribution displayed by the Library Manager was a combination of all media types and pools. It was difficult to determine if changing the reclamation thresholds would yield a specific result for the targeted media type. With the TS7740 Virtualization Engine, this information is not available through a Library Manager. You must use the BVIR function to retrieve and analyze the data.

This BVIR Function provides counts of physical media by media type, pool, and uses, for a specific distributed library, a human readable formatted response, as shown in Example 9-17.

Example 9-17 Bulk Volume Information Retrieval for physical media pools

```
12345678901234567890123456789012345678901234567890123456789012
VTS BULK VOLUME DATA REQUEST
PHYSICAL MEDIA POOLS
05/20/2009 08:45:29 VERSION 03
S/N: 1B245 LIB ID: BAR09
```

R	POOL	MEDIA	EMPTY	FILLING	FULL	ERASE	ROR	UNAVAIL					
C	00	JA	00083										
C	00	JB	00051										
C	00	JJ	00002										
C	01	JA	00134	00003	00609	00000	00000	00000					
C	01	JB	00001	00002	00001	00006	00000	00000					
R	POOL	MEDIA	0+	10+	20+	30+	40+	50+	60+	70+	80+	90+	
D	01	JA	00000	00002	00034	00056	00092	00078	00084	00065	00195	00183	
D	01	JB	00000	00000	00000	00000	00000	00000	00000	00002	00003	00001	

In the upper area of the BVIR response for the PHYSICAL MEDIA POOLS query, you can identify the number of empty and active volumes per pool, that is, the number of stacked volumes that have active data on them for each pool. This includes stacked volumes that are in the process of being filled with new data. Pool 0 is the Common Scratch Pool (CSP) and therefore it can only contain EMPTY stacked volumes.

In the lower section of the output, you can see the active data distribution per pool. You should not see any volumes below your Reclaim Threshold Percentage value; if you do, perhaps you should decrease, if possible, your Inhibit Reclaim Schedule values.

Copy audit request

A database is maintained on each individual TS7740 Virtualization Engine Cluster that contains status information about the logical volumes defined to the Grid. Two key pieces of information are whether or not the cluster contains a valid copy of a logical volume and whether the copy policy for the volume indicates that it should have a valid copy.

This request performs an audit of the databases on a set of specified TS7700 Virtualization Engine distributed libraries to determine if there are any volumes that do not have a valid copy on at least one of them. If the COPYMODE option is specified, whether or not the volume is supposed to have a copy on the distributed library is taken into account when determining whether that distributed library has a valid copy. If COPYMODE is specified and the copy policy for a volume on a specific cluster is “R” or “D”, then that cluster is considered during the audit. If COPYMODE is specified and the copy policy for a volume on a specific cluster is “N”, then the volume’s validity state is ignored because that cluster does not need to have a valid copy. The request then returns a list of any volumes that do not have a valid copy, subject to the copy mode if the COPYMODE option is specified, on the TS7700 Virtualization Engine clusters specified as part of the request.

The specified clusters might not have a copy for several reasons:

- ▶ The copy policy associated with the volume did not specify that any of the clusters specified in the request were to have a copy and the COPYMODE option was not specified. This might be because of a mistake in defining the copy policy or because it was intended. For example, volumes used in a disaster recovery test only need to reside on the disaster recovery TS7700 Virtualization Engine and not on the production TS7700 Virtualization Engines. If the request specified only the production TS7700 Virtualization Engines, all of the volumes used in the test would be returned in the list.
- ▶ The copies have not yet been made from a source TS7700 Virtualization Engine to one or more of the specified clusters. This could be because the source TS7700 Virtualization Engine or the links to it are unavailable or because a copy policy of deferred was specified and a copy had not been completed when the audit was performed. In addition, one or more of the specified clusters might have completed their copy and then had their copy automatically removed as part of the TS7700 Virtualization Engine hybrid automated removal policy function. Automatic removal can only take place on TS7720 Virtualization Engine clusters in a hybrid configuration.
- ▶ Each of the specified clusters contained a valid copy at one time but has since removed them as part of the TS7700 Virtualization Engine hybrid automated removal policy function. Automatic removal can only take place on TS7720 Virtualization Engine clusters in a hybrid configuration.

The Copy Audit request is intended to be used for the following situations:

- ▶ A TS7700 Virtualization Engine is to be removed from a Grid configuration. Prior to its removal, you want to ensure that the TS7700 Virtualization Engines that are to remain in the Grid configuration have a copy of all the important volumes that were created on the TS7700 Virtualization Engine that is to be removed.
- ▶ A condition has occurred (because of a site disaster or as part of a test procedure) where one of the TS7700 Virtualization Engines in a Grid configuration is no longer available and you want to determine which, if any, volumes on the remaining TS7700 Virtualization Engines do not have a valid copy.

In the Copy Audit request, you need to specify which TS7700 Virtualization Engine clusters are to be audited. The clusters are specified by using their associated distributed library ID (this is the unique 5 character library sequence number defined when the TS7700 Virtualization Engine Cluster was installed). If more than one distributed library ID is specified, they are separated by a comma. The following rules determine which TS7700 Virtualization Engine clusters are to be included in the audit:

- ▶ When the INCLUDE parameter is specified, all specified distributed library IDs will be included in the audit. All clusters associated with these IDs must be available or the audit will be failed.
- ▶ When the EXCLUDE parameter is specified, all specified distributed library IDs will be excluded from the audit. All other clusters in the Grid configuration must be available or the audit will fail.
- ▶ Distributed library IDs specified are checked for being valid in the Grid configuration. If one or more of the specified distributed library IDs are invalid, the Copy Audit is failed and the response will indicate the IDs that are considered invalid.
- ▶ Distributed library IDs must be specified or the Copy Audit fails.

Here are some examples of valid requests (assume a three-cluster grid configuration with distributed library IDs of DA01A, DA01B, and DA01C):

- ▶ COPY AUDIT INCLUDE DA01A: Audits the copy status of all volumes on only the cluster associated with distributed library ID DA01A.
- ▶ COPY AUDIT COPYMODE INCLUDE DA01A: Audits the copy status of volumes that also have a valid copy policy on only the cluster associated with distributed library ID DA01A.
- ▶ COPY AUDIT INCLUDE DA01B,DA01C: Audits the copy status of volumes on the clusters associated with distributed library IDs DA01B and DA01C.
- ▶ COPY AUDIT EXCLUDE DA01C: Audits the copy status of volumes on the clusters in the Grid configuration associated with distributed library IDs DA01A and DA01B.

On completion of the audit, a response record is written for each logical volume that did not have a valid copy on any of the specified clusters. Volumes that have never been used, have had their associated data deleted, or have been returned to scratch, are not included in the response records. The record includes the volume serial number and the copy policy definition for the volume. The VOLSER and the copy policy definitions are comma separated, as shown in Example 9-18.

Example 9-18 BVIR message when Copy Audit is requested

```
1234567890123456789012345678901234567890
L00001,R,D,D,N,N,N,N,N,N,R,N,R,N,N,N,N,N
```

Notes:

- ▶ The output for Copy Audit includes copy consistency points for up to eight TS7700 Virtualization Engine clusters. This is to provide for future expansion of the number of clusters supported in a TS7700 Virtualization Engine Grid to the architected maximum.
- ▶ Copy Audit might take more than an hour to complete depending on the number of logical volumes that have been defined, how many clusters are configured in the Grid configuration, and how busy the TS7700 Virtualization Engines are at the time of the request.

Record 6

This record lists the unique distributed library IDs by their 5 character library sequence number separated by a comma. The distributed library IDs listed are either the ones used for a successful Copy Audit operation or ones that the TS7700 Virtualization Engine had a problem with, as shown in Example 9-19.

Example 9-19 BVIR message when Distributed Library IDs are returned

```
12345678901
DA01B,DA01C
```

Record 7-N

When there are logical volumes that do not have a copy on any of the specified TS7700 Virtualization Engine clusters, each of these records provide information for each volume. See the field descriptions in Table 9-28.

Table 9-28 Decoding of the data record 7-N for Copy Audit request

Field name	Description
VOLSER	6 character volume serial number.
cluster0_copy_mode	<p>This field indicates whether cluster 0 should have a copy of the volume and the copy consistency point defined for the volume. The values are as follows:</p> <ul style="list-style-type: none">▶ R: Rewind unload (RUN) copy consistency point.▶ D: Deferred copy consistency point.▶ N: No copy.▶ E: The volume was previously assigned a copy consistency point of rewind unload or deferred, but was changed to no copy.
cluster1_copy_mode	<p>This field indicates whether cluster 1 should have a copy of the volume and the copy consistency point defined for the volume. The values are as follows:</p> <ul style="list-style-type: none">▶ R: Rewind unload (RUN) copy consistency point.▶ D: Deferred copy consistency point.▶ N: No copy.▶ E: The volume was previously assigned a copy consistency point of rewind unload or deferred, but was changed to no copy.
cluster2_copy_mode	<p>This field indicates whether cluster 2 should have a copy of the volume and the copy consistency point defined for the volume. The values are as follows:</p> <ul style="list-style-type: none">▶ R; Rewind unload (RUN) copy consistency point.▶ D: Deferred copy consistency point.▶ N: No copy.▶ E: The volume was previously assigned a copy consistency point of rewind unload or deferred, but was changed to no copy.
cluster3_copy_mode	<p>This field indicates whether cluster 3 should have a copy of the volume and the copy consistency point defined for the volume. The values are as follows:</p> <ul style="list-style-type: none">▶ R; Rewind unload (RUN) copy consistency point.▶ D: Deferred copy consistency point.▶ N: No copy.▶ E: The volume was previously assigned a copy consistency point of rewind unload or deferred, but was changed to no copy. <p>This is new for Version 3.</p>
cluster4_copy_mode	<p>This field indicates whether cluster 4 should have a copy of the volume and the copy consistency point defined for the volume. The values are as follows:</p> <ul style="list-style-type: none">▶ R; Rewind unload (RUN) copy consistency point.▶ D: Deferred copy consistency point.▶ N: No copy.▶ E: The volume was previously assigned a copy consistency point of rewind unload or deferred, but was changed to no copy. .

Field name	Description
cluster5_copy_mode	<p>This field indicates whether cluster 5 should have a copy of the volume and the copy consistency point defined for the volume. The values are as follows:</p> <ul style="list-style-type: none"> ▶ R; Rewind unload (RUN) copy consistency point. ▶ D: Deferred copy consistency point. ▶ N: No copy. ▶ E: The volume was previously assigned a copy consistency point of rewind unload or deferred, but was changed to no copy.
cluster6_copy_mode	<p>This field indicates whether cluster 6 should have a copy of the volume and the copy consistency point defined for the volume. The values are as follows:</p> <ul style="list-style-type: none"> ▶ R; Rewind unload (RUN) copy consistency point. ▶ D: Deferred copy consistency point. ▶ N: No copy. ▶ E: The volume was previously assigned a copy consistency point of rewind unload or deferred, but was changed to no copy.
cluster7_copy_mode	<p>This field indicates whether cluster 7 should have a copy of the volume and the copy consistency point defined for the volume. The values are as follows:</p> <ul style="list-style-type: none"> ▶ R; Rewind unload (RUN) copy consistency point. ▶ D: Deferred copy consistency point. ▶ N: No copy. ▶ E: The volume was previously assigned a copy consistency point of rewind unload or deferred, but was changed to no copy.
cluster0_removal_state	<p>This field indicates the removal state of the volume on cluster 0. The values are as follows:</p> <ul style="list-style-type: none"> ▶ R: This TS7720 Virtualization Engine cluster has automatically removed its consistent copy because of the volumes configured removal policy within a TS7700 Virtualization Engine hybrid configuration. ▶ N: This cluster has not removed its copy. A TS7740 Virtualization Engine or a TS7720 Virtualization Engine in a non-hybrid configuration will always contain this value. <p>Note: If this cluster is not included in the copy audit, the value will default to N even though its actual value might be R. This is new for Version 3.</p>
cluster1_removal_state	<p>This field indicates the removal state of the volume on cluster 1. The values are as follows:</p> <ul style="list-style-type: none"> ▶ R: This TS7720 Virtualization Engine cluster has automatically removed its consistent copy because of the volumes configured removal policy within a TS7700 Virtualization Engine hybrid configuration. ▶ N: This cluster has not removed its copy. A TS7740 Virtualization Engine or a TS7720 Virtualization Engine in a non-hybrid configuration will always contain this value. <p>Note: If this cluster is not included in the copy audit, the value will default to N even though its actual value might be R. This is new for Version 3.</p>

Field name	Description
cluster2_removal_state	<p>This field indicates the removal state of the volume on cluster 2. The values are as follows:</p> <ul style="list-style-type: none"> ► R: This TS7720 Virtualization Engine cluster has automatically removed its consistent copy because of the volumes configured removal policy within a TS7700 Virtualization Engine hybrid configuration. ► N: This cluster has not removed its copy. A TS7740 Virtualization Engine or a TS7720 Virtualization Engine in a non-hybrid configuration will always contain this value. <p>Note: If this cluster is not included in the copy audit, the value will default to N even though its actual value might be R. This is new for Version 3.</p>
cluster3_removal_state	<p>This field indicates the removal state of the volume on cluster 3. The values are as follows:</p> <ul style="list-style-type: none"> ► R: This TS7720 Virtualization Engine cluster has automatically removed its consistent copy because of the volumes configured removal policy within a TS7700 Virtualization Engine hybrid configuration. ► N: This cluster has not removed its copy. A TS7740 Virtualization Engine or a TS7720 Virtualization Engine in a non-hybrid configuration will always contain this value. <p>Note: If this cluster is not included in the copy audit, the value will default to N even though its actual value might be R. This is new for Version 3.</p>
cluster4_removal_state	<p>This field indicates the removal state of the volume on cluster 3. The values are as follows:</p> <ul style="list-style-type: none"> ► R: This TS7720 Virtualization Engine cluster has automatically removed its consistent copy because of the volumes configured removal policy within a TS7700 Virtualization Engine hybrid configuration. ► N: This cluster has not removed its copy. A TS7740 Virtualization Engine or a TS7720 Virtualization Engine in a non-hybrid configuration will always contain this value. <p>Note: If this cluster is not included in the copy audit, the value will default to N even though its actual value might be R. This is new for Version 3.</p>
cluster5_removal_state	<p>This field indicates the removal state of the volume on cluster 3. The values are as follows:</p> <ul style="list-style-type: none"> ► R: This TS7720 Virtualization Engine cluster has automatically removed its consistent copy because of the volumes configured removal policy within a TS7700 Virtualization Engine hybrid configuration. ► N: This cluster has not removed its copy. A TS7740 Virtualization Engine or a TS7720 Virtualization Engine in a non-hybrid configuration will always contain this value. <p>Note: If this cluster is not included in the copy audit, the value will default to N even though its actual value might be R. This is new for Version 3.</p>

Field name	Description
cluster6_removal_state	<p>This field indicates the removal state of the volume on cluster 3. The values are as follows:</p> <ul style="list-style-type: none"> ► R: This TS7720 Virtualization Engine cluster has automatically removed its consistent copy because of the volumes configured removal policy within a TS7700 Virtualization Engine hybrid configuration. ► N: This cluster has not removed its copy. A TS7740 Virtualization Engine or a TS7720 Virtualization Engine in a non-hybrid configuration will always contain this value. <p>Note: If this cluster is not included in the copy audit, the value will default to N even though its actual value might be R. This is new for Version 3.</p>
cluster7_removal_state	<p>This field indicates the removal state of the volume on cluster 3. The values are as follows:</p> <ul style="list-style-type: none"> ► R: This TS7720 Virtualization Engine cluster has automatically removed its consistent copy because of the volumes configured removal policy within a TS7700 Virtualization Engine hybrid configuration. ► N: This cluster has not removed its copy. A TS7740 Virtualization Engine or a TS7720 Virtualization Engine in a non-hybrid configuration will always contain this value. <p>Note: If this cluster is not included in the copy audit, the value will default to N even though its actual value might be R. This is new for Version 3.</p>
composite_removal_state	<p>This field indicates the overall removal state of the volume on the Grid with respect to those clusters that are specified as part of the copy audit consistency check. The values are as follows:</p> <ul style="list-style-type: none"> ► R: All clusters included in the consistency check have a copy mode of R or D, are TS7720s, and are not consistent because of the removal of data in a TS7700 hybrid configuration. Only TS7720 clusters can remove volumes in a hybrid configuration. ► N: At least one cluster in the consistency check is not consistent because of an invalid copy mode policy or it has not yet completed its copy. <p>This is new for Version 3.</p>

If the audit does not find any logical volumes that are not valid on the specified TS7700 Virtualization Engine clusters, the record shown in Example 9-20 is returned (see Table 9-29 for more details).

Example 9-20 BVIR message when no invalid volumes exist on specified clusters

```
1234567890123456789012345678901234567890
NO INVALID VOLUMES ON SPECIFIED CLUSTERS
```

Table 9-29 Decoding of the data record 7-N with no invalid volumes for a Copy Audit request

Bytes	Name	Description
1-27		NO INVALID VOLUMES ON SPECIFIED CLUSTERS
28-80	Blanks	Blank character padding

If errors or operational conditions are encountered by the TS7700 Virtualization Engine while processing the copy audit request, the audit cannot be performed and one of the following

records is returned, depending on the condition (also see the error responses listed in “Unknown or invalid request” on page 705):

- One or more of the clusters associated with the specified distributed library IDs are not available because they are in service or offline. Record 6 lists the distributed library IDs specified that are associated with the unavailable cluster(s). See Example 9-21 and Table 9-30 for more information.

Example 9-21 BVIR message when clusters are not in a valid state

```
123456789012345678901234567
CLUSTERS NOT IN VALID STATE
```

Table 9-30 Decoding of the data record 7-N invalid state for a Copy Audit request

Bytes	Name	Description
1-27		CLUSTERS NOT IN VALID STATE
28-80	Blanks	Blank character padding

- No distributed library IDs are included in the request. At least one is required. Record 6 will be all blanks. See Example 9-22 and Table 9-31 for more information.

Example 9-22 BVIR message when no Distributed Library IDs were specified

```
1234567890123456789012345678
NO DISTRIBUTED IDS SPECIFIED
```

Table 9-31 Decoding of the data record 7-N non distributed IDs for a Copy Audit request

Bytes	Name	Description
1-28		NO DISTRIBUTED IDS SPECIFIED
29-80	Blanks	Blank character padding

- One or more of the specified distributed library IDs are not defined to the Grid configuration. Record 6 lists the distributed library IDs specified that are not defined. See Example 9-23 and Table 9-32 for more information.

Example 9-23 BVIR message when Distributed Library ID is invalid

```
123456789012345678901234
DISTRIBUTED ID NOT VALID
```

Table 9-32 Decoding of the data record 7-N invalid distributed for a Copy Audit request

Bytes	Name	Description
1-22		DISTRIBUTED ID NOT VALID
23-80	Blanks	Blank character padding

- The EXCLUDE parameter was specified and the distributed library IDs specified includes all defined in the Grid configuration. Record 6 lists the distributed library IDs specified. See Example 9-24 and Table 9-33 for more information.

Example 9-24 BVIR message when all clusters were excluded

```
1234567890123456789012345678901
ALL CLUSTERS HAVE BEEN EXCLUDED
```

Table 9-33 Decoding of the data record 7-N all clusters excluded for a Copy Audit request

Bytes	Name	Description
1-31		ALL CLUSTERS HAVE BEEN EXCLUDED
32-80	Blanks	Blank character padding

- Neither the INCLUDE or the EXCLUDE parameter was specified. See Example 9-25 and Table 9-34 for more information.

Example 9-25 BVIR message when INCLUDE or EXCLUDE parameter was not specified

```
123456789012345678901234567890123456
INCLUDE OR EXCLUDE MUST BE SPECIFIED
```

Table 9-34 Decoding of the data record 7-N invalid syntax for a Copy Audit request

Bytes	Name	Description
1-36		INCLUDE OR EXCLUDE MUST BE SPECIFIED
37-80	Blanks	Blank character padding

- For copy audit to be performed, it requires that no other major database update functions are being performed. If the TS7700 Virtualization Engine Grid is in the process of inserting logical volumes, performing a Copy Export operation, or is already executing a Copy Export operation, the message in Example 9-26 issued.

Example 9-26 BVIR message when global resources are busy

```
12345678901234567890123456789012345678
GLOBAL RESOURCE CURRENTLY BUSY - TRY AGAIN LATER
```

See Table 9-35 for more information.

Table 9-35 Decoding of the data record 7-N currently busy for a Copy Audit request

Bytes	Name	Description
1-48		GLOBAL RESOURCE CURRENTLY BUSY - TRY AGAIN LATER
48-90	Blanks	Blank character padding

- An internal error in the Grid resulted in no response being generated. See Example 9-27 and Table 9-36 for more information.

Example 9-27 BVIR message when an internal error caused no response

```
123456789012345678901234567890123
INTERNAL ERROR CAUSED NO RESPONSE
```

Table 9-36 Decoding of the data record 7-N internal error for a Copy Audit request

Bytes	Name	Description
1-33		INTERNAL ERROR CAUSED NO RESPONSE
34-80	Blanks	Blank character padding

Unknown or invalid request

If the request file does not contain the correct number of records or the first record is incorrect, the request file on the volume is unchanged and no error is indicated.

If the request file contains the correct number of records and the first record is correct but the second is not, the response file will indicate in Record 6 that the request is unknown, as shown in Example 9-28.

Example 9-28 BVIR message when an unknown or invalid request is submitted

```
12345678901234567890
UNKNOWN REQUEST TYPE
```

Table 9-37 shows the record format details.

Table 9-37 Decoding a BVIR unknown or invalid request

Bytes	Name	Description
1-20		UNKNOWN REQUEST TYPE
21-80	Blanks	Blank character padding

If the request file contains the correct number of records, the first record is correct, the second is recognized but includes a variable that is not within the range supported for the request, and the response file will indicate in record 6 that the request is invalid, as shown in Example 9-29.

Example 9-29 BVIR message when an invalid variable is specified

```
12345678901234567890123456
INVALID VARIABLE SPECIFIED
```

Table 9-38 shows the exact format of the record.

Table 9-38 Decoding the BVIR INVALID VARIABLE SPECIFIED message

Bytes	Name	Description
1-26		3. INVALID VARIABLE SPECIFIED
27-80	Blanks	Blank character padding

9.11 IBM Tape Tools

A set of tape tools is available on an *as-is* basis to help you monitor your tape environment. Several of these tools are specific to the TS7700 and are based on BVIR data, such as VEHSTATS, for example. Before we describe VEHSTATS and the reports that you can obtain by running the VEHSTATS jobs, we give you a general overview over the IBM Tape Tools and guide you through the installation of these tools in a z/OS environment.

9.11.1 Introduction to IBM Tape Tools

All the TS7700 Virtualization Engine monitoring and evaluating tools described in this chapter can be found at the following FTP site:

<ftp://ftp.software.ibm.com/storage/tapetool/>

Example 9-30 lists the content of the `Readme.txt` file that provides basic information about the tape tools.

Example 9-30 Readme.txt from the IBM Tape Tools website

IMPORTANT	IMPORTANT
Program enhancements will be made to handle data format changes when they occur. If you try to run new data with old program versions, results will be unpredictable. To avoid this situation, you need to be informed of these enhancements so you can stay current. To be informed of major changes to any of the tools distributed via this ftp site, send an email message to:	

TAPETOOL@US.IBM.COM

In the subject, specify NOTIFY. Nothing else is required in the body of the note. This will add you to our change distribution list.

The UPDATES.TXT file will contain a chronological history of all changes made to the tools. You should review that file on a regular basis, at least monthly, perhaps weekly, so you can see if any changes apply to you.

Look in file, OVERVIEW.PDF, for an overview of all currently available tools. The JCL, CNTL, and LOAD libraries for all the tools are zipped into IBMTTOOLS.EXE. IBMTTOOLS.TXT explains the complete installation procedure. Most tools have their own xxxxxx.txt file with more detail. There are no formal documentation manuals. The intent is to have enough JCL comment to allow the user to run the jobs without difficulty and to have adequate column headings and footnotes to make report output obvious without needing additional documentation.

If you feel that the JCL or report output needs more explanation, please send an mail to the address above indicating the area needing attention.

Most of these tools are z/OS-based and included in the `ibmtools.exe` file. A complete list of all tools that are included in the `ibmtools.exe` file is available in the `overview.pdf` file. The start of the `overview.pdf` file is shown in Figure 9-53.

IBM Tape Analysis Tools Overview					
ftp://ftp.software.ibm.com/storage/tapetool			<=== as IBMTOOLS.EXE		
Tool	Major Use	Benefit	Inputs	Outputs	Product pre-reqs
BADBLKSZ	Identify small VTS blk sizes	Improve VTS performance, make jobs run faster	Logrec MDR & CA1, TLMS, RMM, ZARA, CTLT	Volser, Jobname, Dsname for VTS volumes with small blk sizes	none
BIRTHDST	Find old cartridges	Improve data reliability	CA1, TLMS, RMM, ZARA, CTLT	List of active data on old media and cartridge birth date distribution.	none
BVIRHSTx	Get historical stats from TS7700	Creates U, VB, SMF format	TS7700	Statistics file	TS7700
BVIRPOOL	Identify available scratch by pool	Reports all pools at once	BVIR file	Physical media by pool	TS7700
BVIRPRPT	Reclaim copy export volumes	Based on active GB, not %	BVIR file	Detail report of data on volumes	TS7700
BVIRRPT	Identify VTS virtual volumes by owner	Determine which applications or users have virtual volumes	BVIR data & CA1, TLMS, RMM, ZARA, CTLT	Logical volumes by jobname or dsname, logical to physical rpts	VTS 7.4 for size
BVPTRPT	Point in Time stats as WTO	Immediately available	TS7700	Point in Time stats as WTO	TS7700
CHGDATE	Change dates in SMF 14,15, 21,30,40 records	Allows analyst to simulate combining workloads from two different periods.	SMF 14,15,21,30,40	Modified SMF records	none
CHKDUPS	Identify duplicate volsers	Identify duplicate volsers across multiple tape management catalogs	CA1, TLMS, RMM, ZARA, CTLT	List of duplicate volsers	none
COPYVTS	Copy lvols from old VTS	Recall lvols based on selected application(s)	BVIR data & CA1, TLMS, RMM, ZARA, CTLT	IEBGENER to recall lvols and copy to new VTS	Need copy program


Figure 9-53 Tape Tools listed in `overview.pdf` file

9.11.2 Tools download and installation

Public access is provided to the IBM Tape Tools library, which contains various tools that can help you analyze your tape environment. This set of utilities also include the VEHSTATS and VEPSTATS tools, which use the Bulk Volume Information Retrieval (BVIR) reports for comprehensive performance analysis.

Figure 9-54 shows several tools that are available to customers from the FTP site.

Index of <ftp://ftp.software.ibm.com/storage/tapetool/>

 [Up to higher level directory](#)





















Name	Size	Last Modified
 .message	1 KB	15.06.2000 00:00:00
 TS7700.VEHSTATS.Decoder.V10.pdf	414 KB	16.04.2010 10:38:00
 TapeWise.PDF	325 KB	17.08.2005 00:00:00
 VEHGRXCL.exe	1122 KB	05.08.2009 00:00:00
 VEHGRXCL.txt	4 KB	05.08.2009 00:00:00
 VTSGRXCL.EXE	805 KB	28.02.2007 00:00:00
 VTSGRXCL.TXT	2 KB	04.12.2003 00:00:00
 badblksz.txt	2 KB	05.08.2004 00:00:00
 eximcalc.123	11 KB	02.03.2007 00:00:00
 export.txt	1 KB	05.08.2004 00:00:00
 findlrg.txt	2 KB	07.05.2002 00:00:00
 fsrtmm.txt	1 KB	07.05.2002 00:00:00
 ftpcust.txt	2 KB	21.03.2002 00:00:00
 grpdsn.txt	2 KB	05.08.2004 00:00:00
 ibmcntl.xmi	2432 KB	09.08.2010 13:41:00
 ibmjcl.xmi	1119 KB	17.09.2010 13:44:00
 ibmload.xmi	4394 KB	17.09.2010 13:44:00
 ibmpat.xmi	280 KB	09.08.2010 13:41:00
 ibmtools.exe	1922 KB	17.09.2010 13:44:00
 ibmtools.txt	7 KB	28.01.2010 00:00:00

Figure 9-54 Tape tools catalog

The index is as at the following web address:

<http://public.dhe.ibm.com/storage/tapetool/>

For most tools, a text file is available. In addition, each job to run a tool contains a detailed description on the function of the tool and parameters that need to be specified.

Note: For the IBM Tape Tools, there are no warranties, expressed or implied, including the warranties of merchantability and fitness for a particular purpose.

To obtain the tape tools, download the `ibmtools.exe` file to your computer or use FTP from Time Sharing Option (TSO) on your z/OS system to directly upload the files contained in the `ibmtools.exe` file.

The `ibmtools.exe` file is a self-extracting .zip file, that is expand into four separate files:

<code>IBMJCL.XMI</code>	Contains the execution JCL for current tape analysis tools.
<code>IBMCNTL.XMI</code>	Contains parameters needed for job execution, but that do not need to be modified by the user.
<code>IBMLOAD.XMI</code>	Contains the load library for executable load modules.
<code>IBMPAT.XMI</code>	Contains the data pattern library which is only needed if you will be running the <code>QSAMDRVR</code> utility.

The `ibmtools.txt` file contains detailed information about how to download and install the tools libraries.

After you have created the three or four libraries on the z/OS host, be sure you execute the following steps:

1. Copy, edit, and submit `userid.IBMT00LS.JCL($$CPYLIB)` to create a new JCL library that has a unique second node (&SITE symbolic). This step creates a private JCL library for you from which you can submit jobs while leaving the original as is. CNTL and LOAD can then be shared by multiple users running jobs from the same system.
2. Edit and submit `userid.SITENAME.IBMT00LS.JCL($$TAILOR)` to tailor the JCL according to your system requirements.

The `updates.txt` file contains all fixes and enhancements made to the tools. Review this file regularly to determine whether any of the programs you use have been modified.

To ensure that you are not working with outdated tools, the tools are controlled through an EXPIRE member. Every three months, a new EXPIRE value will be issued that is good for the next 12 months. When you download the latest tools package any time during the year, you have at least nine months remaining on the EXPIRE value. New values are issued middle of January, April, July, and October.

If your IBM tools jobs stop running because the expiration date has passed, download the `ibmtools.exe` file again to get the latest `IBMT00LS.JCL(EXPIRE)` member.

9.11.3 IBM Tape Tools for TS7700 monitoring

In this section, we describe several tape tools that can be used to help you better understand your tape processing with regard to TS7700 operation and migration.

IOSTATS

IOSTATS tool is part of the `ibmtools.exe` file, which is available at the following address:

<ftp://ftp.software.ibm.com/storage/tapetool/>

You can use IOSTATS tool to measure job execution times. For example, you might want to compare the TS7700 Virtualization Engine performance before and after configuration changes.

IOSTATS can be run for a subset of job names for a certain period of time before the hardware installation. SMF type 30 records are required as input. The reports list the number of disk and tape I/O operations that were done for each job step, and the elapsed job execution time.

With the TS7700 Virtualization Engine running in a multi-cluster grid configuration, IOSTATS can be used for different purposes:

- ▶ To evaluate the effect of the multi-cluster grid environment, compare job execution times before implementation of the multi-cluster grid to those after migration, especially if you are operating in immediate copy (RUN, RUN data consistency point) mode.
- ▶ To evaluate the effect of hardware upgrades, compare job execution times before and after upgrading components of the TS7700 Virtualization Engine. For example, you might want to verify the performance impact of a larger Tape Volume Cache (TVC) capacity or the number of TS1130/TS1120/3592 Tape Drives.
- ▶ To evaluate the effect of changing the copy mode of operation on elapsed job execution time.

TAPEWISE

As with IOSTATS, TAPEWISE tool is available from the IBM Tape Tools FTP site. TAPEWISE can, based on input parameters, generate several reports that can help with various items:

- ▶ Tape activity analysis
- ▶ Mounts and MBs processed by hour
- ▶ Input and output mounts by hour
- ▶ Mounts by SYSID during an hour
- ▶ Concurrent open drives used
- ▶ Long VTS mounts (recalls)

MOUNTMON and MOUNTRPT

As with IOSTATS, MOUNTMON is available from the IBM Tape Tools FTP site. MOUNTMON runs as a started task or batch job and monitors all tape activity on the system. The program must be APF-authorized and, if it runs continuously, it writes statistics for each tape volume allocation to SMF or to a flat file.

Based on data that is gathered from MOUNTMON, the MOUNTRPT program can report on the following information:

- ▶ How many tape mounts are necessary
- ▶ How many are scratch
- ▶ How many are private
- ▶ How many by host system
- ▶ How many by device type
- ▶ How much time is needed to mount a tape
- ▶ How long are tapes allocated
- ▶ How many drives are being used at any given time
- ▶ What is the most accurate report of concurrent drive usage
- ▶ Which jobs are allocating too many drives

9.12 Using VEHSTATS for monitoring and reporting

In this section, we show how to work with the binary reports for point-in-time and historical statistics. These would have been collected by using the BVIR functions that are described in 9.10, “Bulk Volume Information Retrieval” on page 657.

To convert the binary response record from BVIR data to address your requirements, you can use the IBM tool VEHSTATS when working with historical statistics. When working with point-in-time statistics, you can use the IBM tool VEPSTATS. See 9.11.2, “Tools download and installation” on page 707 for specifics about where to obtain these tools. Details about

using BVIR can also be found in the *IBM Virtualization Engine TS7700 Series Bulk Volume Information Retrieval Function User's Guide*. The most recently published white papers are available at the Techdocs website by searching for TS7700 Virtualization Engine at the following address:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs>

With the record layout of the binary BVIR response data, you can decode the binary file or you may use the record layout to program your own tool for creating statistical reports.

9.12.1 VEHSTATS tool overview

The TS7700 Virtualization Engine's activity is recorded in the subsystem. The two types of statistics are as follows:

- ▶ Point-in-time statistics: A snapshot of activity in the last 15 seconds
- ▶ Historical statistics: Up to 90 days in 15-minute increments

Both sets of statistics can be obtained through the BVIR functions (see 9.10, "Bulk Volume Information Retrieval" on page 657).

Because both types of statistical data are delivered in binary format from the BVIR functions, you must translate the content into a readable format. You can do this task either manually by using the information provided in the *Virtualization Engine TS7700 Series Statistical Data Format* white paper, as referenced in 9.12.1, "VEHSTATS tool overview" on page 711, or by using an existing automated tool.

IBM provides a historical statistics tool called VEHSTATS. Like the other IBM Tape Tools, the program is provided *as-is*, without official support, for the single purpose of showing how the data might be reported. There is no guarantee of its accuracy, and there is no additional documentation available for this tool. We provide some guidance for interpretation of the reports in 9.12.3, "VEHSTATS reports" on page 713.

You can use VEHSTATS to monitor TS7700 Virtualization Engine virtual and physical back-end tape drives, and Tape Volume Cache (TVC) activity to do trend analysis reports, based on BVIR binary response data. The tool summarizes TS7700 Virtualization Engine activity on a specified time basis, up to 90 days in time sample intervals of 15 minutes or one hour, depending upon the data reported.

Figure 9-54 on page 708 highlights three files that might be helpful in reading and interpreting VEHSTATS reports:

- ▶ The TS7700. VEHSTATS.Decoder.V10.pdf file contains a description of the fields listed in the various VEHSTATS reports.
- ▶ The VEHGRXCL.txt file contains the description for the graphical package contained in VEHGRXCL.EXE.
- ▶ The VEHGRXCL.EXE file contains VEHSTATS_Model.ppt and VEHSTATS_Model.xls. You can use these files to create graphs of cluster activity based on the flat files created with VEHSTATS. Follow the detailed instructions in the VEHSTATS_Model.xls file to create these graphs.

9.12.2 Running the VEHSTATS jobs

You have several output options for VEHSTATS, and must submit separate jobs depending on your requirements. The IBMTTOOLS.JCL member VEHSTATS (see Example 9-31) provides guidance on which job to chose.

Example 9-31 Member VEHSTATS

THERE ARE NOW DIFFERENT VERSIONS OF THE VEHSTATS JOB DEPENDING
ON HOW YOU WANT TO VIEW OR SAVE THE REPORTS.

1. **VEHSTSO** WRITES REPORTS DIRECTLY TO SYSOUT (THIS IS OLD VEHSTATS)
 2. **VEHSTPS** WRITES FINAL REPORTS TO A SINGLE PHYSICAL SEQUENTIAL FILE WHERE REPORTS ARE WRITTEN WITH DISP=MOD.
 3. **VEHSTPO** WRITES FINAL REPORTS TO A PDSE WHERE EACH REPORT IS A SEPARATE MEMBER.
-

In addition to the VEHSTATS tool, sample BVIR jobs are included in the IBMTTOOLS libraries. These jobs help you obtain the input data from the TS7700 Virtualization Engine. With these jobs, you can control where the historical statistics are accumulated for long term retention. The TS7700 Virtualization Engine still maintains historical statistics for the previous 90 days, but you can have the pulled statistics recorded directly to the SMF log file or continue to use the disk flat file method. The flat files can be recorded as either RECFM=U or RECFB=VB.

Three specific jobs in IBMT00LS.JCL are designed to fit your particular needs:

BVIRHSTS	To write statistics to the SMF log file
BVIRHSTU	To write statistics to a RECFM=U disk file
BVIRHSTV	To write statistics to a RECFM=VB disk file

The VEHSTATS reporting program accepts any or all of the various formats of BVIR input. You define which input is to be used through a DD statement in the VEHSTATS job. The three input DD statements are optional, but at least one of the statements shown in Example 9-32 must be specified.

Example 9-32 VEHSTATS Input DD statements

```
/* ACTIVATE ONE OR MORE OF THE FOLLOWING DD STATEMENTS FOR YOUR DATA
/*STATSU DD DISP=SHR,
/*          DSN=&USERHLQ..#&VTSID..BVIRHIST.D070205.D070205
/*STATSVB DD DISP=SHR,
/*          DSN=&USERHLQ..#&VTSID..BVIRHIST.D070206.D070206
/*STATSMF DD DISP=SHR,          RECORDS WILL BE SELECTED BASED ON SMFNUM
/*          DSN=&USERHLQ..#&VTSID..SMF194
```

The SMF input file can contain all SMF record types kept by the user. The SMFNUM parameter defines which record number is processed when you specify the STATSMF statement.

The fields shown in the various reports depend on which ORDER member in IBMTOOLS.JCL is being used. Use the following steps to ensure that the reports and the flat file contain the complete information that you want to be in the reports:

1. Review which member is defined in the **ORDER=** parameter in the VEHSTATS job member.
2. Verify that none of the fields that you want to see have been deactivated by an asterisk in the first column. Example 9-33 shows sample active and inactive definitions in the ORDERV12 member of IBMTOOL.JCL. The sample statements define whether you want the amount of data in cache be displayed in MB or in GB.

Example 9-33 Sample statements in the ORDERV12 member

*ORDER=' PG0 MB IN TVC';	PG0	MEGABYTES IN CACHE
*ORDER=' PG1 MB IN TVC';	PG1	MEGABYTES IN CACHE
ORDER=' PG0 GB IN TVC';	PG0	GIGABYTES IN CACHE
ORDER=' PG1 GB IN TVC';	PG1	GIGABYTES IN CACHE

If you are planning to create graphics from the flat file using the graphics package from the IBM Tape Tools FTP site, you should specify the ORDERV12 member, because it contains all the fields that are used when creating the graphics, and verify that all statements are activated for all clusters in your environment.

9.12.3 VEHSTATS reports

VEHSTATS can be used to monitor TS7700 Virtualization Engine drive and Tape Volume Cache (TVC) activity and do trend analysis to see where the performance bottlenecks are. Also, comparative analysis can be used to determine whether an upgrade, such as adding additional physical tape drives, might improve the overall performance of the TS7740 Virtualization Engine. VEHSTATS is not a projection tool, but it provides the basis for an overall health check of the TS7700 Virtualization Engine.

VEHSTATS gives us a huge amount of information. We describe the reports that we think are most important, and for every report we give some results and analysis that can help you understand the reports better.

The following reports are available for the TS7700 Virtualization Engine:

- ▶ H20VIRT: Virtual Device Historical Records
- ▶ H21ADP00: vNode Adapter Historical Activity
- ▶ H21ADPXX: vNode Adapter Historical Activity combined (by adapter)
- ▶ H21ADPSU: vNode Adapter Historical Activity combined (total)
- ▶ H30TVC1: hNode HSM Historical Cache Partition
- ▶ H31IMEX: hNode Export/Import Historical Activity
- ▶ H32TDU12: hNode Library Historical Drive Activity
- ▶ H32CSP: hNode Library Hist Scratch Pool Activity
- ▶ H32GUP01: General Use Pools 01/02 through General Use Pools 31/32
- ▶ H33GRID: hNode Historical Peer-to-Peer Activity
- ▶ AVGRDST: Hrs Interval Average Recall Mount Pending Distribution
- ▶ DAYMRY: Daily Summary
- ▶ MONMRY: Monthly Summary
- ▶ COMPARE: Interval Cluster Comparison
- ▶ HOURFLAT: Hourly
- ▶ DAYHSMRY: Daily flat file

Note: Be sure you have a copy of TS7700. VEHSTATS.Decoder.V10.pdf available when you familiarize yourself with the VEHSTATS reports.

Virtual Device Activity

Example 9-34 shows the report for Virtual Device Activity. This report gives you an overview, per 15-minute interval, of the relevant time frame and has the following information:

- ▶ The minimum, average, or maximum (MIN, AVG, or MAX) mounted virtual drives
- ▶ The amount of channel blocks written based on blocksize

Note: The report is provided per cluster in the Grid. The report title includes the cluster number in the DIST_LIB_ID field.

Example 9-34 VEHSTATS report for Virtual Drive Activity

```
(C) IBM   REPORT=H2OVIRT (10210)          VNODE VIRTUAL DEVICE HISTORICAL RECORDS N ON 19AUG2010
@ 10:39:15   PAGE   1
GRID#=CC001  DIST_LIB_ID= 0  VNODE_ID= 0  NODE_SERIAL=CLOABCDE  VE_CODE_LEVEL=008.006.000.0110
UTC NOT CHG
21JUL10WE -VIRTUAL_DRIVES-
RECORD      --MOUNTED--      MAX
-----CHANNEL_BLOCKS_WRITTEN_FOR_THese_BLOCKSIZEs-----
TIME INST MIN AVG MAX  THRPUT    <=2048    <=4096    <=8192    <=16384    <=32768    <=65536
>65536
 4:15:00   256 114 124 127   MAX      630         0         0         0      2485298         0
0
 4:30:00   256 117 125 127   MAX      631         0         0         0      2026062         0
0
 4:45:00   256 113 124 127   MAX      530         0         0         0      2620099         0
0
 5:00:00   256 117 125 127   MAX      474         0         0         0      3118714         0
0
```

The most important fields in this report are CHANNEL BLOCKS WRITTEN FOR BLOCKSIZES. In general, we should see that the largest amount of blocks are written at 32768 or higher blocksize, but this is not a fixed rule. For example, DFSMSHsm writes a 16384 blocksize and DB2 writes a 4096 blocksize. From an I/O point of view, analysis of blocksize on performance is outside the scope of this book.

vNode Host Adapter Activity

The next example report provides details about the vNode Host Adapter Activity. Although there is a large amount of information available (one report per distributed library per FICON adapter), we believe that the vNode Adaptor Historical Activity Combined report is sufficient in providing an overall view of the FICON channel performance. As always, one report exists for each distributed library. This report is on an hourly basis with the following information:

- ▶ Total throughput per distributed library every hour
- ▶ Read and write channel activity
- ▶ Read and write device activity with compression rate achieved

Example 9-35 shows a sample report for Adapter 3 of Cluster 0.

Example 9-35

C) IBM REPORT=H21ADP03(10210)										VNODE ADAPTOR HISTORICAL ACTIVITY										RUN ON 18AUG2010 @ 8:04:29										PAGE 30																													
GRID#=CC001 DIST_LIB_ID= 0 VNODE_ID= 0										NODE_SERIAL=CLOABCDE										VE_CODE_LEVEL=008.006.000.0110										UTC NOT CHG																													
ADAPTOR 3 FICON-2 (ONLINE)										R DRAWER SLOT# 6																																																	
19JUL10MO PORT 0										MiB is 1024 based, MB is 1000 based										PORT 1																																							
RECORD GBS MB/										-----CHANNEL-----										-----DEVICE-----										GBS MB/										-----CHANNEL-----										-----DEVICE-----									
TIME		RTE		sec		RD_MB		MB/s		WR_MB		MB/s		RD_MB		COMP		WR_MB		COMP		RTE		sec		RD_MB		MB/s		WR_MB		MB/s		RD_MB		COMP		WR_MB		COMP																			
01:00:00		4		20		25827		7		49676		13		7741		3.33		19634		2.53		0		0		0		0		0		0		0		0		0																					
02:00:00		4		7		9204		2		18030		5		2100		4.38		6480		2.78		0		0		0		0		0		0		0		0		0																					
03:00:00		4		1		2248		0		4550		1		699		3.21		1154		3.94		0		0		0		0		0		0		0		0		0																					
04:00:00		4		0		0		0		69		0		0				24		2.87		0		0		0		0		0		0		0		0		0																					
05:00:00		4		0		1696		0		1655		0		550		3.08		540		3.06		0		0		0		0		0		0		0		0		0																					
06:00:00		4		9		8645		2		24001		6		3653		2.36		13589		1.76		0		0		0		0		0		0		0		0		0																					
07:00:00		4		4		6371		1		10227		2		2283		2.79		3503		2.91		0		0		0		0		0		0		0		0		0																					
08:00:00		4		2		5128		1		4950		1		2048		2.50		1985		2.49		0		0		0		0		0		0		0		0		0																					
09:00:00		4		3		6270		1		7272		2		2530		2.47		3406		2.13		0		0		0		0		0		0		0		0		0																					

The most important fields in this report are as follows:

- ▶ **TOTAL_MB/S:** This is the total throughput of the cluster. In a Multi Cluster configuration, we need to combine every distributed library's hourly figures to calculate the total throughput, giving us a better perspective of performance. For example, at the sixteenth time frame, we can see 253 MBps in both distributed libraries. This gives a total throughput of 506 MBps during this hour.
- ▶ **DEVICE_COMP** (to the right of the WR-GB column): This is the real rate of compression achieved each hour. The compression ratio is *highly* dependent on the nature of the data.

The host adapter activity is summarized per adapter and as a total of all adapters. This result is also shown in the vNode Adaptor Throughput Distribution report which we show in Example 9-36.

Example 9-36 Adapter Throughput Distribution report

(C) IBM REPORT=H21ADPSU(10210) VNODE ADAPTOR THROUGHPUT DISTRIBUTION RUN ON 18AUG2010 @ 8:04:29					
GRID#=CC001 DIST_LIB_ID= 0 VNODE_ID= 0 NODE_SERIAL=CLOABCDE					
VE_CODE_LEVEL=008.006.000.0110 UTC NOT CHG					
MB/SEC_RANGE		#INTERVALS	PCT	ACCUM%	
1	- 50	477	64.4	64.4	
51	- 100	191	25.8	90.2	
101	- 150	52	7.0	97.2	
151	- 200	17	2.2	99.5	
201	- 250	1	0.1	99.7	
251	- 300	2	0.2	100.0	

This report summarizes the overall host throughput and shows how many one hour intervals have shown which throughput, for example, if you look at the second line of the report data:

- ▶ The throughput was 51 - 100 MBps in 191 intervals.
- ▶ 191 intervals are 25.8% of the entire measurement period.
- ▶ In 90.2% of the measurement intervals, the throughput was below 100 MBps.

This report provides details of Cache Partition Activity in the TS7700 Virtualization Engine. You can identify the following information for each 15-minute interval:

- ▶ The percentage of read, write, and deferred copy throttling
- ▶ The number of Fast Ready mounts, cache hits, and cache misses
- ▶ The capacity and number of logical volumes by preference group (0 or 1) in cache

The report also shows information about the Preference Groups.

The most important fields in this report are:

- ▶ The ratio between FAST_RDY_MOUNTS, CACHE_HIT_MOUNTS, and CACHE_MISS_MOUNTS. In general, a high number of CACHE_MISSES could mean that additional cache capacity is needed or cache management policies need to be adjusted.
- ▶ FAST_RDY_AVG_SECS and CACHE_HIT_AVG_SECS should show only a few seconds. CACHE_MISS_AVG_SECS can list values higher than a few seconds, but higher values (more than two or three minutes) could indicate a lack of back-end physical tape drives; see “Physical Drive Activity” on page 716 for more information.

Another important report to look at is the report for Physical Drive Activity, grouped by device type. We list the information for TS1130 drives from a sample report in Example 9-37. From there, you can identify the following items for each 15-minute interval taken for the report:

- ▶ How many physical tape drives were installed.
- ▶ How many physical tape drives were available.
- ▶ How many drives (min/avg/max) were mounted.
- ▶ How much time (min/avg/max in seconds) the mount took.
- ▶ The number of physical mounts sorted by purpose:
 - STG: Recalls of logical volume back into cache
 - MIG: Premigration of logical volume from cache to physical tape
 - RCM: Reclamation
 - SDE: Secure Data Erase

Example 9-37 VEHSTATS for Physical Drives Activity

08JUL10TH -----PHYSICAL_DRIVES_3592-E06-----																			
RECORD	--MOUNTED--					-MOUNT_SECS-			----MOUNTS_FOR----										
TIME	INST	AVL	MIN	AVG	MAX	MIN	AVG	MAX	STG	MIG	RCM	SDE	TOT						
01:00:00	16	16		2	9		16	20	32		53			3	15		0	0	18
02:00:00	16	16	3	8	16	20	25	39	6	4	0	0	10						
03:00:00	16	16	1	4	9	20	20	21	4	2	0	0	6						
04:00:00	16	16	1	2	3	19	21	23	0	2	0	0	2						

The most important fields in this report are as follows:

- **PHYSICAL_DRIVE_MOUNTED_AVG**: If this value is equal or close to the maximum drives available during several hours, this might mean that more physical tape drives are required.
- **MOUNT_FOR (RCL MIG RCM SDE)**: This field presents the reason for each physical mount. If the percentage value in the Recall (RCL) column is high compared to the total number of mounts, this might indicate a need to evaluate the cache size or cache management policies. However, this is not a fixed rule and further analysis is required. For example, if HSM migration is into a TS7740 Virtualization Engine, we might see high recall activity during the morning, which can be driven by temporary development or user activity. This is normal and not a problem in itself.

Common Scratch Pool

The report for Common Scratch Pool presented in Example 9-38 shows you the amount of scratch tapes per physical media type. Review this report often to avoid a shortage in scratch stacked volumes.

Example 9-38 VEHSTATS report for Common Scratch Pool

18JUL10	-----SCRATCH_STACKED_VOLUMES_AVAILABLE_BY_TYPE-----									
RECORD	3590J	3590K	3592JA	3592JJ	NONE	NONE	3592JB	NONE		
TIME	MEDIA0	MEDIA1	MEDIA2	MEDIA3	MEDIA4	MEDIA5	MEDIA6	MEDIA7		
4:15:00	0		0		42		0		0	0
4:30:00	0		0		42		0		0	0
4:45:00	0		0		42		0		0	0
5:00:00	0		0		41		0		0	0

General Pool Use

The General Pool Use report is shown in Example 9-39. A single report always shows two pools. In this example, the report shows Pool 01 and Pool 02. You can see the following details per pool for each recorded time frame:

- ▶ The number of active logical volumes
- ▶ The amount of active data in GB
- ▶ The amount of data written in MB
- ▶ The amount of data read in MB
- ▶ The current reclamation threshold and target pool

Example 9-39 VEHSTATS report for General Pool Use

(C) IBM REPORT=H32GUP01(10210)										HNODE LIBRARY HIST GUP/POOLING ACTIVITY										RUN ON 18JUL2010 @ 16:57:51										PAGE 03	
GRID#=CC001 DIST_LIB_ID= 0 VNODE_ID= 0										NODE_SERIAL=CLOABCDE VE_CODE_LEVEL=008.006.000.0110										-										UTC NOTCHG	
18JUL10 POOL 01 3592-E05 3592JA										READ UN- POOL 02 3592-E05																				READ UN-	
RECORD	ACTIVE	ACTIVE	MB	MB	VOL_COUNT	RECLAIM-	ONLY	AVAI	ACTIVE	ACTIVE	MB	MB	VOL_COUNT	RECLAIM-	ONLY	AVAI															
TIME	LVOLS	GB	WRITTN	READ	SCR PRIV	PCT POOL	VOLS VOLS		LVOLS	GB	WRITTN	READ	SCR PRIV	PCT POOL	VOLS VOLS																
UPD INT=>	-ON_THE_HOUR-				ON_THE_HR				-ON_THE_HOUR-																						
4:15:00	65079	18052	5412	0	2 56	25 01	00 00		0	0	0	0	0	0	25 02	00 00															
4:30:00	65079	18052	37888	0	2 56	25 01	00 00		0	0	0	0	0	0	25 02	00 00															
4:45:00	65079	18052	83895	0	2 56	25 01	00 00		0	0	0	0	0	0	25 02	00 00															
00																															
5:00:00	65630	18206	94721	0	2 57	25 01	00 00		0	0	0	0	0	0	25 02	00 00															
5:15:00	65630	18206	98630	0	2 57	25 01	00 00		0	0	0	0	0	0	25 02	00 00															
5:30:00	65630	18206	124490	0	2 57	25 01	00 00		0	0	0	0	0	0	25 02	00 00															
5:45:00	65630	18206	119979	0	2 57	25 01	00 00		0	0	0	0	0	0	25 02	00 00															
6:00:00	67069	18610	108854	0	2 57	25 01	00 00		0	0	0	0	0	0	25 02	00 00															
6:15:00	67069	18610	108854	0	2 57	25 01	00 00		0	0	0	0	0	0	25 02	00 00															
6:30:00	67069	18610	97126	0	2 57	25 01	00 00		0	0	0	0	0	0	25 02	00 00															

Peer to Peer Activity

The Peer to Peer Activity report shown in Example 9-40 on page 718 provides various performance metrics of Grid activity. This report can be useful for installations working in deferred copy mode. This report allows, for example, the analysis of subsystem performance during peak Grid network activity, such as determining the maximum delay during the batch window.

For the duration of the report you can identify, in 15-minute increments, the following items:

- ▶ The number of logical volumes to be copied (valid only for a multi-cluster grid configuration)
- ▶ The amount of data to be copied (MB)
- ▶ The average age of Copy Jobs on the deferred and immediate copy queue
- ▶ The amount of data (in MB) to and from the Tape Volume Cache driven by copy activity

- The amount of data (MB) copied from other clusters (inbound data) to the cluster on which the report was executed

Note: Analyzing the report shown in Example 9-40, you see three active clusters with write operations from a host. This might not be a common configuration, but it is an example of a scenario to show the possibility of having three copies of a logical volume in a multi-cluster grid.

Example 9-40 VEHSTATS report for Peer to Peer Activity

(C) IBM REPORT=H33GRID (10210) HNODE HISTORICAL PEER-TO-PEER ACTIVITY RUN ON 18AUG2010 @ 8:04:29 PAGE 37															
GRID#=CC001 DIST_LIB_ID= 1 VNODE_ID= 0 NODE_SERIAL=CL1FEDCB VE_CODE_LEVEL=008.006.000.0110 UTC NOT CHG															
25JUN10FR	LVOLS	TO	TO	AV_DEF	AV_RUN	MB_TO	CALC	V_MNTS	MB_XFR	MB_XFR	1-->0	CALC	1-->2	CALC	
		RECEIVE	RECEIVE	---MINUTES---		COPY	SEC	OTHR_DL	RMT_WR	RMT_RD	COPY	SEC	COPY	SEC	
01:00:00	1	13	1	0	139077	38.6	43	1	346	61355	17.0	746	0.2	156	0.0
02:00:00	6	1518	7	0	150440	41.7	84	462	11410	64536	17.9	4448	1.2	1175	0.3
03:00:00	2	3239	3	0	88799	24.6	38	8	44	57164	15.8	1114	0.3	166	0.0
04:00:00	2	574	4	0	241205	67.0	4	82	29	109850	30.5	1409	0.3	401	0.1
05:00:00	3	1055	2	0	70637	19.6	9	390	136	51464	14.2	2488	0.6	0	
06:00:00	16	9432	2	0	187776	52.1	33	1519	491	100580	27.9	2526	0.7	463	0.1
07:00:00	0	0	0	0	86624	24.0	19	63	12649	50139	13.9	6036	1.6	1988	0.5
08:00:00	1	484	0	0	46314	12.8	26	30	12292	23216	6.4	9563	2.6	1971	0.5

The most important fields in this report are as follows:

- **MB_TO_COPY:** The amount of data pending a copy function to other clusters (outbound).
- **MB_FR:** The amount of data (MB) copied from the cluster (inbound data) identified in the column heading. The column heading 1-->2 indicates Cluster 1 is the copy source and Cluster 2 is the target.
- **CALC_MB/SEC:** This number shows the true throughput achieved when replicating data between the clusters identified in the column heading.

Summary reports

In addition to daily and monthly summary reports per cluster, VEHSTATS also provides a side-by-side comparison of all clusters for the entire measurement interval. Examine this report for an overall view of the Grid and for significant or unexpected differences between the clusters.

9.13 z/OS commands for monitoring

In addition to the previously introduced methods and options for monitoring the TS7700 Virtualization Engine, we provide additional points for further subsystem monitoring.

9.13.1 Display SMS command

Several DISPLAY SMS commands exist to display the OAM status, the Composite and Distribution Library, and volume details.

Several of these commands (shown in **bold**) and their responses are listed in Example 9-41, separated by a dashed line.

Example 9-41 DISPLAY SMS Command responses

```
D SMS,OAM
F OAM,D,OAM,L=BHAEUSS-Z
CBR1100I OAM status: 770
TAPE TOT ONL TOT TOT TOT TOT ONL AVL TOTAL
LIB LIB AL VL VCL ML DRV DRV DRV SCRTCH
2 2 1 0 1 0 138 138 136 27826
There are also 1 VTS distributed libraries defined.
CBRUXCUA processing ENABLED.
CBRUXEJC processing ENABLED.
CBRUXENT processing ENABLED.
CBRUXVNL processing ENABLED.
-----
D SMS,LIBRARY(ALL),DETAIL
F OAM,D,LIB,L=BHAEUSS-Z
CBR1110I OAM library status: 866
TAPE LIB DEVICE TOT ONL AVL TOTAL EMPTY SCRTCH ON OP
LIBRARY TYP TYPE DRV DRV DRV SLOTS SLOTS VOLS
TLIB AL 3584-L22 10 10 8 6019 5319 316 Y Y
TVTS VCL 3957-V06 128 128 128 0 0 27510 Y Y
TVTSD VDL 3957-V06 0 0 0 6000 5700 0 N Y
-----
D SMS,LIBRARY(ALL),STATUS
IGD002I 00:20:45 DISPLAY SMS 944

LIBRARY CLASS SYSTEM= 1
TLIB TAPE +
TVTS TAPE +
TVTSD TAPE -
***** LEGEND *****
. THE LIBRARY IS NOT DEFINED TO THE SYSTEM
+ THE LIBRARY IS ONLINE
- THE LIBRARY IS OFFLINE
P THE LIBRARY IS PENDING OFFLINE
SYSTEM 1 = MCEVS1
-----
D SMS,LIBRARY(TVTS),DETAIL
F OAM,D,LIB,TVTS,L=BHAEUSS-Z
CBR1110I OAM library status: 074
TAPE LIB DEVICE TOT ONL AVL TOTAL EMPTY SCRTCH ON OP
LIBRARY TYP TYPE DRV DRV DRV SLOTS SLOTS VOLS
TVTS VCL 3957-V06 128 128 128 0 0 27510 Y Y
-----
MEDIA SCRATCH SCRATCH SCRATCH
TYPE COUNT THRESHOLD CATEGORY
MEDIA1 100 0 0001
MEDIA2 27410 500 0002
-----
DISTRIBUTED LIBRARIES: TVTSD
-----
LIBRARY ID: 01955
OPERATIONAL STATE: AUTOMATED
ERROR CATEGORY SCRATCH COUNT: 0
CORRUPTED TOKEN VOLUME COUNT: 0
```

 Library supports import/export.
 Library supports outboard policy management.
 Library supports logical WORM.

D SMS,LIBRARY(TLIB),DETAIL

F OAM,D,LIB,TLIB,L=BHAEUSS-Z

CBR1110I OAM library status: 233

TAPE	LIB	DEVICE	TOT	ONL	AVL	TOTAL	EMPTY	SCRATCH	ON	OP
LIBRARY	TYP	TYPE	DRV	DRV	DRV	SLOTS	SLOTS	VOLS		
TLIB	AL	3584-L22		10	10		9	6019	5319	316 Y Y

MEDIA	SCRATCH	SCRATCH	SCRATCH
TYPE	COUNT	THRESHOLD	CATEGORY
MEDIA5	312	20	0005
MEDIA9	4	0	0009

LIBRARY ID: 01513

OPERATIONAL STATE: AUTOMATED

ERROR CATEGORY SCRATCH COUNT: 0

 Convenience I/O station installed.
 Convenience I/O station in Input mode.
 Convenience I/O station Empty.
 Bulk input/output not configured.

D SMS,VOLUME(A00416)

F OAM,D,VOL,A00416,L=BHAEUSS-Z

CBR1180I OAM tape volume status: 403

VOLUME	MEDIA	STORAGE	LIBRARY	USE	W	C	SOFTWARE	LIBRARY
	TYPE	GROUP	NAME	ATR	P	P	ERR STAT	CATEGORY
A00416	MEDIA5	TLIB	TLIB	P	N	N	NOERROR	PRIVATE

RECORDING TECH: EFMT2 COMPACTION: YES
 SPECIAL ATTRIBUTE: NONE ENTER/EJECT DATE: 2007-02-09
 CREATION DATE: 2007-02-08 EXPIRATION DATE: 2011-09-26
 LAST MOUNTED DATE: 2010-09-30 LAST WRITTEN DATE: 2010-09-30
 SHELF LOCATION:
 OWNER: ADSM

09.18.16 **d sms,vol(hyd210)**

09.18.16 STC00098 CBR1180I OAM tape volume status: 870

VOLUME	MEDIA	STORAGE	LIBRARY	USE	W	C	SOFTWARE	LIBRARY
	TYPE	GROUP	NAME	ATR	P	P	ERR STAT	CATEGORY
HYD210	MEDIA2	*SCRATCH*	ATL3484F	S	N	N	NOERROR	SCRMED2

RECORDING TECH: 36 TRACK COMPACTION: NO
 SPECIAL ATTRIBUTE: NONE ENTER/EJECT DATE: 2010-06-11
 CREATION DATE: 2010-06-11 EXPIRATION DATE:
 LAST MOUNTED DATE: 2011-01-30 LAST WRITTEN DATE: 2011-01-30
 SHELF LOCATION:
 OWNER:
 LM SG: LM SC: LM MC: LM DC:

 Logical volume.
 Volume is logical WORM.

For more information, see Chapter 8, "Operation" on page 409 and *z/OS DFSMS Object Access Method Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC35-0427.

9.13.2 Library command

The LIBRARY command and the LIBRARY REQUEST command, also known as the Host Console Request function, can be used to check for missing virtual drives or for the status of the Grid links. Example 9-42 shows the output of the LIBRARY DD command, which you can use to verify whether all virtual drives are available.

Example 9-42 Sample response for LI DD,libname command

```
LI DD,ATVIGA
RESPONSE=EGZB
CBR1220I Tape drive status: 338
```

DRIVE NUM	DEVICE TYPE	LIBRARY NAME	ON	OFF	REAS	N	LM	ICL	ICL	MOUNT
				LI	OP	PT	AV	CATE	LOAD	VOLUME
5F00	3490	ATVIGA		N	N	Y	N	A	NONE	N
5F01	3490	ATVIGA		N	N	Y	N	A	NONE	N
5F02	3490	ATVIGA		N	N	Y	N	A	NONE	N
5F03	3490	ATVIGA		N	N	Y	N	A	NONE	N
5F04	3490	ATVIGA		N	N	Y	N	A	NONE	N
5F05	3490	ATVIGA		N	N	Y	N	A	NONE	N
5F06	3490	ATVIGA		N	N	Y	N	A	NONE	N
5F07	3490	ATVIGA		N	N	Y	N	A	NONE	N
5F08	3490	ATVIGA		N	N	Y	N	A	NONE	N
5F09	3490	ATVIGA		N	N	Y	N	A	NONE	N
5F0A	3490	ATVIGA		N	N	Y	N	A	NONE	N
5F0B	3490	ATVIGA		N	N	Y	N	A	NONE	N
5F0C	3490	ATVIGA		N	N	Y	N	A	NONE	N
5F0D	3490	ATVIGA		N	N	Y	N	A	NONE	N
5F0E	3490	ATVIGA		N	N	Y	N	A	NONE	N
5F0F	3490	ATVIGA		N	N	Y	N	A	NONE	N
5F10	3490	ATVIGA		N	N	Y	N	A	NONE	N
5F11	3490	ATVIGA		N	N	Y	N	A	NONE	N
5F12	3490	ATVIGA		N	N	Y	N	A	NONE	N
5F13	3490	ATVIGA		N	N	Y	N	A	NONE	N
5FFA	3490	ATVIGA		N	N	Y	N	A	NONE	N
5FFB	3490	ATVIGA		N	N	Y	N	A	NONE	N
5FFC	3490	ATVIGA		N	N	Y	N	A	NONE	N
5FFD	3490	ATVIGA		N	N	Y	N	A	NONE	N
5FFE	3490	ATVIGA		N	N	Y	N	A	NONE	N
5FFF	3490	ATVIGA		N	N	Y	N	A	NONE	N

For more information about the LIBRARY command, see Chapter 8, “Operation” on page 409 and *z/OS DFSMS Object Access Method Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC35-0427.

9.14 What to look for and where

In this chapter, we describe tools, provide considerations, and give you important information to help you monitor and understand the performance indicators of your TS7700 Grid. In Table 9-39, we summarize where you can find the important information and what observations you can make.

Most checks that you should make in each shift ensure that the TS7700 environment is operating as expected. The checks that are made daily or weekly are intended for tuning and longer-term trend analysis.

The information in this table are intended as a basis for monitoring; you may tailor them to best fit your needs.

Table 9-39 Monitoring summary

Information	Source	Tracking	Reporting interval	Observation
All virtual drives online	LI DD,libname	Display each Composite Library and each system	Each shift	Report or act on any missing drive
VE Health Check	TS7700 MI	Display each Composite Library	Each shift	Report any offline or degraded status
Library online and operational	D SMS, LIB(ALL), DETAIL	Display each Composite Library and each system	Each shift	Verify availability to systems
Exits enabled	D SMS,OAM	Display each system	Each shift	Report any disabled exits
Virtual scratch volumes	D SMS, LIB(ALL), DETAIL	Display each Composite Library	Each shift	Report each shift
Physical scratch tapes	D SMS, LIB(ALL), DETAIL	Display each Composite Library	Each shift	Report each shift
Interventions	D SMS, LIB(ALL), DETAIL	Display each Composite Library	Each shift	Report or act on any interventions
Grid Link Status	LI REQ,libname, STATUS, GRIDLINK	Display each Composite Library	Each shift	Report any errors or elevated Retransmit%
Number of volumes on the deferred copy queue	TS7700 MI → Logical Volumes → Incoming Copy Queue	Display for each Cluster in the Grid	Each shift	Report and watch for gradual or sudden increases
Copy queue depths	TS7700 MI	Display for each system	Each shift	Report if queue depth is higher than usual
Virtual Mounts per Day	VEHSTATS	Rolling weekly trend	daily	Increase over time indicates increased workload
MB Transferred per Day	VEHSTATS	Rolling weekly trend	daily	Increase over time indicates increased workload

Information	Source	Tracking	Reporting interval	Observation
Virtual volumes managed	VEHSTATS	Rolling weekly trend	daily	Capacity planning: maximum one million per Grid
MB Stored	VEHSTATS	Rolling weekly trend	daily	Capacity planning ä general awareness
Backend Drive Utilization	VEHSTATS	Rolling weekly trend	daily	Check for periods of 100%
Daily Throttle Indicators	VEHSTATS	Rolling weekly trend	daily	Key performance indicator
Average Virtual Mount Time	VEHSTATS	Rolling weekly trend	daily	Key performance indicator
Cache Hit Percentage	VEHSTATS	Rolling weekly trend	daily	Key performance indicator
Physical Scratch Count	VEHSTATS	Rolling weekly trend	daily	Capacity planning ä general awareness
Available Slot Count	D SMS, LIB(ALL), DETAIL	Rolling weekly trend	daily	Capacity planning ä general awareness
Available Virtual Scratch Volumes	D SMS, LIB(ALL), DETAIL	Rolling weekly trend	daily	Drive insert
Data Distribution	BVIRPOOL Job	Watch for healthy distribution	weekly	Use for reclaim tuning
Times in Cache	VEHSTATS	Watch for healthy distribution	weekly	Preference Group tuning indicator



Failover and disaster recovery scenarios

In this chapter, we cover IBM Virtualization Engine TS7700 failover scenarios, and disaster recovery (DR) planning and considerations, with or without a Geographical Dispersed Parallel Sysplex (GDPS).

We discuss the following topics:

- ▶ TS7700 failover principles and scenarios
- ▶ Geographically Dispersed Parallel Sysplex considerations for a TS7700 Virtualization Engine Grid configuration
- ▶ Disaster recovery testing
- ▶ Real recovery, including using the Copy Export function

10.1 TS7700 Virtualization Engine Grid failover principles

To better understand and plan for the actions to be performed with the TS7700 Virtualization Engine Grid configuration in the event of failures, this section describes key concepts for Grid operation and the many failure scenarios that the Grid has been designed to handle. A TS7700 Virtualization Engine Grid configuration provides the following data access and availability characteristics:

- ▶ Accessing the data on a particular cluster requires that a host mount request be issued on a virtual device address that is defined for that cluster. The virtual device addresses for each cluster are independent. In a prior generation, the PTP VTS mount request was issued on a virtual device address defined for a virtual tape controller and the virtual tape controller then decided which VTS to use for data access.
- ▶ All logical volumes are accessible through any of the virtual device addresses on the TS7700 Virtualization Engine Clusters in the Grid configuration. The preference is to access a copy of the volume in the tape volume cache that is associated with the TS7700 Virtualization Engine Cluster on which the mount request is received. If a recall is required to place the logical volume in the tape volume cache on that TS7700 Virtualization Engine Cluster, it will be done as part of the mount operation. If a copy of the logical volume is not available at that TS7700 Virtualization Engine Cluster (either because it does not have a copy or the copy it does have is inaccessible because of an error), and a copy is available at another TS7700 Virtualization Engine Cluster in the Grid, the volume is accessed through the tape volume cache at the TS7700 Virtualization Engine Cluster that has the available copy. If a recall is required to place the logical volume in the tape volume cache on the other TS7700 Virtualization Engine Cluster, it will be done as part of the mount operation.
- ▶ Whether a copy is available at another TS7700 Virtualization Engine Cluster in a multi-cluster grid depends on the copy consistency point that had been assigned to the logical volume when it was written. The copy consistency point is set through the Management Class storage construct. It specifies if and when a copy of the data is made between the TS7700 Virtualization Engine Clusters in the Grid configuration. The copy consistency policies that can be assigned are as follows:
 - Rewind Unload (RUN) Copy Consistency Point: If a data consistency point of RUN is specified, the data created on one TS7700 Virtualization Engine Cluster is copied to the other TS7700 Virtualization Engine Cluster as part of successful rewind unload command processing, meaning that for completed jobs, a copy of the volume will exist on both TS7700 Virtualization Engine Clusters. Access to data written by completed jobs (successful rewind/unload) prior to the failure is maintained through the other TS7700 Virtualization Engine Cluster. Access to data of incomplete jobs that were in process at the time of the failure is not provided.
 - Deferred Copy Consistency Point: If a data consistency point of Deferred is specified, the data created on one TS7700 Virtualization Engine Cluster is copied to the specified TS7700 Virtualization Engine Clusters after successful rewind unload command processing. Access to the data through the other TS7700 Virtualization Engine Cluster is dependent on when the copy completes. Because there will be some delay in performing the copy, access might or might not be available when a failure occurs.
 - No Copy Consistency Point: If a data consistency point of No Copy is specified, the data created on one TS7700 Virtualization Engine Cluster is not copied to the other TS7700 Virtualization Engine Cluster. If the TS7700 Virtualization Engine Cluster data was written to fails, the data for that logical volume is inaccessible until that TS7700 Virtualization Engine Cluster's operation is restored.

- With the introduction of the three-cluster grid, Logical Volume Copy Consistency Override has been enabled. Using settings on each library, you can override existing RUN consistency points. Figure 10-1 shows an example of the setting from the management interface (MI), where a third copy is forced to Deferred Copy, but only if three copies are requested from the MC definitions on all three clusters.

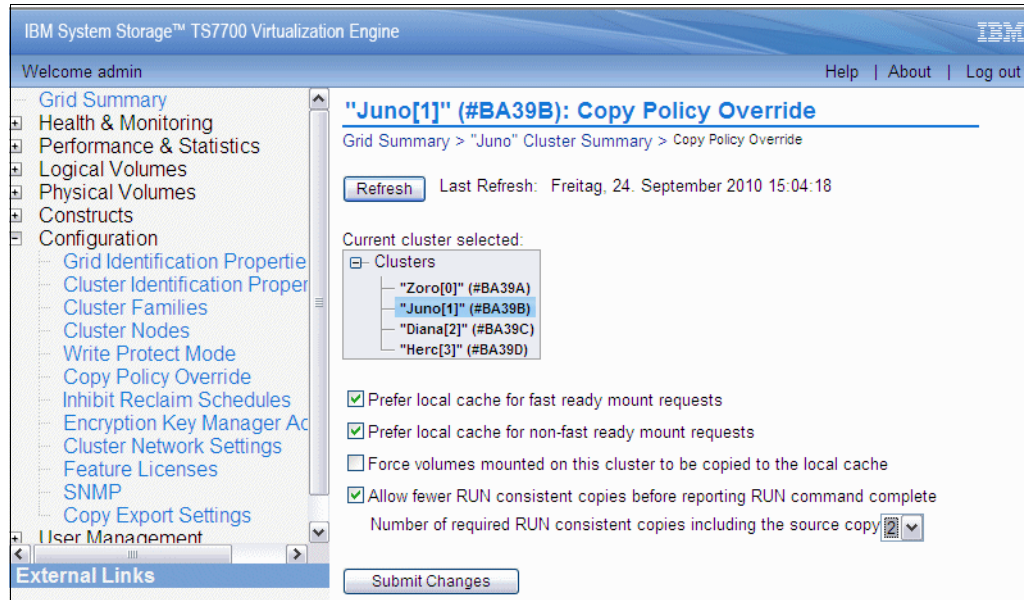


Figure 10-1 Logical Volume Copy Consistency Override

- The volume removal policy for hybrid Grid configurations was introduced in the R1.6 microcode level. Beginning with R1.7, it is available in any grid configuration that contains at least one TS7720 cluster. Because the TS7720 “Disk-Only” solution has a maximum storage capacity that is the size of its tape volume cache, after cache fills, this policy allows logical volumes to be automatically removed from cache while a copy is retained within one or more peer clusters in the Grid. When the auto-removal starts, all volumes in the fast-ready (scratch) category are removed first because these volumes are intended to hold temporary data. This mechanism could remove old volumes in a private category from the cache to meet pre-defined cache usage threshold if a copy of the volume is retained on one of the remaining clusters. A TS7740 cluster failure can affect the availability of old volumes (no logical volumes are removed from a TS7740 cluster).
- If a logical volume is written on one of the TS7700 Virtualization Engine Clusters in the Grid configuration and copied to the other TS7700 Virtualization Engine Cluster, the copy can be accessed through the other TS7700 Virtualization Engine Cluster. This is subject to the so-called *volume ownership*.

At any time, a logical volume is “owned” by a cluster. The owning cluster has control over access to the volume and changes to the attributes associated with the volume (such as category or storage constructs). The cluster that has ownership of a logical volume can surrender it dynamically to another cluster in the Grid configuration that is requesting a mount of the volume.

When a mount request is received on a virtual device address, the TS7700 Virtualization Engine Cluster for that virtual device must have ownership of the volume to be mounted or must obtain the ownership from the cluster that currently owns it. If the TS7700 Virtualization Engine Clusters in a Grid configuration and the communication paths between them are operational (Grid Network), the change of ownership and the processing of logical volume related commands are transparent in regards to the operation of the TS7700 Virtualization Engine Cluster.

However, if a TS7700 Virtualization Engine Cluster that owns a volume is unable to respond to requests from other clusters, the operation against that volume will fail, unless some additional direction is given. In other words, clusters will not automatically assume or take over ownership of a logical volume without being directed. This is done to prevent the failure of the Grid Network communication paths between the TS7700 Virtualization Engine Clusters resulting in both clusters thinking they have ownership of the volume. If more than one cluster has ownership of a volume, that could result in the volume's data or attributes being changed differently on each cluster, resulting in a data integrity issue with the volume.

If a TS7700 Virtualization Engine Cluster has failed or is known to be unavailable (for example, a power fault in IT center) or needs to be serviced, its ownership of logical volumes is transferred to the other TS7700 Virtualization Engine Cluster through one of the following modes. These modes are set through the management interface:

- Read Ownership Takeover: When Read Ownership Takeover (ROT) is enabled for a failed cluster, ownership of a volume is allowed to be taken from a TS7700 Virtualization Engine Cluster that has failed. Only read access to the volume is allowed through the other TS7700 Virtualization Engine Cluster in the Grid. After ownership for a volume has been taken in this mode, any operation attempting to modify data on that volume or change its attributes is failed. The mode for the failed cluster remains in place until a different mode is selected or the failed cluster has been restored.
 - Write Ownership Takeover: When Write Ownership Takeover (WOT) is enabled for a failed cluster, ownership of a volume is allowed to be taken from a cluster that has been marked as failed. Full access is allowed through the other TS7700 Virtualization Engine Cluster in the Grid. The mode for the failed cluster remains in place until a different mode is selected or the failed cluster has been restored.
 - Service Prep/Service Mode: When a TS7700 Virtualization Engine Cluster is placed in service preparation mode or is in service mode, ownership of its volumes is allowed to be taken by the other TS7700 Virtualization Engine Cluster. Full access is allowed. The mode for the cluster in service remains in place until it has been taken out of service mode.
- In addition to the manual setting of one of the ownership takeover modes, an optional automatic method named Autonomic Ownership Takeover Manager (AOTM) is available when each of the TS7700 Virtualization Engine Clusters is attached to a TotalStorage System Controller (TSSC) and there is a communication path provided between the TSSCs. AOTM is enabled and defined by the IBM SSR. If the clusters are in close proximity of each other, then multiple clusters in the same Grid can be attached to the same TSSC and the communication path is not required.

Guidance: The links between the TSSCs must not be the same physical links that are also used by cluster Grid Gigabit links. AOTM should have a different network to be able to detect that a missing cluster is actually down, and that the problem is not caused by a failure in the Grid Gigabit WAN links.

If enabled by the SSR, if a TS7700 Virtualization Engine Cluster cannot obtain ownership from the other TS7700 Virtualization Engine Cluster because it does not get a response to an ownership request, a check is made through the TSSCs to determine whether the owning TS7700 Virtualization Engine Cluster is inoperable or just that the communication paths to it are not functioning. If the TSSCs have determined that the owning TS7700 Virtualization Engine Cluster is inoperable, then they will enable either read or write ownership takeover, depending on what was set by the CE.

- AOTM enables an ownership takeover mode after a grace period and that can only be configured by an IBM SSR. Therefore, jobs can intermediately fail with an option to retry

until the AOTM enables the configured Takeover Mode. The grace period is set to 20 minutes, by default. The grace period starts when a TS7700 detects that a remote TS7700 has failed; it can take several minutes.

The following OAM messages can be displayed up until the point when AOTM enables the configured ownership takeover mode:

- CBR3758E Library Operations Degraded
 - CBR3785E Copy operations disabled in library
 - CBR3786E VTS operations degraded in library
 - CBR3750I Message from library libname: G0013 Library libname has experienced an unexpected outage with its peer library libname. Library libname may be unavailable or a communication issue may be present
 - CBR3750I Message from library libname: G0009 Autonomic ownership takeover manager within library libname has determined that library libname is unavailable. The Read/Write ownership takeover mode has been enabled.
 - CBR3750I Message from library libname: G0010 Autonomic ownership takeover manager within library libname has determined that library libname is unavailable. The Read-Only ownership takeover mode has been enabled.
- A failure of a TS7700 Virtualization Engine Cluster will cause the jobs using its virtual device addresses to abend. To re-run the jobs, host connectivity to the virtual device addresses in the other TS7700 Virtualization Engine Cluster must be enabled (if not already) and an appropriate ownership takeover mode selected. As long as the other TS7700 Virtualization Engine Cluster has a valid copy of a logical volume, the jobs can be retried.

If a logical volume is being accessed in a remote cache through a 1 Gb Ethernet link and that link fails, the job accessing that volume also fails. If the failed job is attempted again, the TS7700 Virtualization Engine will use the remaining 1 Gb Ethernet link. If both 1 Gb links fail, access to any data in a remote cache is not possible.

10.2 Failover scenarios

As part of a total systems design, you must develop business continuity procedures to instruct IT personnel in the actions that they need to take in the event of a failure. Test those procedures either during the initial installation of the system or at some other time.

The scenarios described in the following sections are from some of the scenarios described in the *IBM Virtualization Engine TS7700 Series Grid Failover Scenarios* white paper, which was written in an effort to assist IBM specialists and customers in developing such testing plans. The white paper is available at the following address:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100831>

The white paper documents a series of TS7700 Virtualization Engine Grid failover test scenarios for z/OS that were run in an IBM laboratory environment. Single failures of all major components and communication links and some multiple failures are simulated.

10.2.1 Test configuration

The hardware configuration used for the laboratory test scenarios is shown in Figure 10-2.




Figure 10-2 Grid test configuration for a two-cluster grid

For the Automatic Takeover scenarios, a TSSC attached to each of the TS7700 Virtualization Engine Clusters is required and an Ethernet connection between the TSSCs. Although all the components tested were local, the results of the tests should be similar, if not the same, for remote configurations. All FICON connections were direct, but again, the results should be valid for configurations utilizing FICON directors. Any supported level of z/OS software, and current levels of TS7700 Virtualization Engine, IBM 3953, and TS3500 Tape Library microcode, should all provide similar results. The test environment was MVS/JES2. Failover capabilities are the same for all supported host platforms, although host messages differ and host recovery capabilities might not be supported in all environments.

For the tests, all host jobs are routed to the virtual device addresses associated with TS7700 Virtualization Engine Cluster 0. The host connections to the virtual device addresses in TS7700 Virtualization Engine Cluster 1 are used in testing recovery for a failure of TS7700 Virtualization Engine Cluster 0.

An IBM Support team should be involved in the planning and execution of any failover tests. In certain scenarios, intervention by a service support representative (SSR) might be needed to initiate failures or restore “failed” components to operational status.

Test job mix

The test jobs running during each of the failover scenarios consist of 10 jobs that mount single specific logical volumes for input (read), and five jobs that mount single scratch logical volumes for output (write). The mix of work used in the tests is purely arbitrary, and any mix would be suitable, but in order for recovery to be successful, some logical drives must be available for a swap, and for that reason, less than the maximum number of virtual drives should be active during testing. Also, a large number of messages are generated during some scenarios, and fewer jobs will reduce the number of host console messages.

Note: The following scenarios were tested using TS7740 Virtualization Engine Clusters with attached TS3500 Tape Libraries. The scenarios also apply to the TS7720 Virtualization Engine as long as they are limited to virtual volume management and Grid communication.

10.2.2 Failover scenario 1

In this scenario, shown in Figure 10-3, we assume one host link to TS7740-0 fails. The failure might be the intermediate FICON directors, FICON channel extenders, or remote channel extenders.




Figure 10-3 Failure of a host link to a TS7740 Virtualization Engine

Effects of the failure

You see the following effects of the failure:

- ▶ All Grid components continue to operate.
- ▶ All channel activity on the failing host link is stopped.
- ▶ Host channel errors are reported or error information becomes available from the intermediate equipment.
- ▶ If alternate paths exist from the host to either TS7700, the host I/O operations may continue. Ownership takeover modes are not needed.
- ▶ All data remains available.

Recovery from failure

To recover from the failures, note the following information:

- ▶ Normal error recovery procedures apply for the host channel and the intermediate equipment
- ▶ You must contact your IBM SSR to repair the failed connection.

10.2.3 Failover scenario 2

In this scenario, shown in Figure 10-4, we assume a failure of both links between the TS7700 Virtualization Engine Clusters.




Figure 10-4 Failure of both links between the TS7700 Virtualization Engine Clusters

Effects of the failure

You will see the following effects of the failure:

- ▶ Jobs on virtual device addresses on TS7700 Cluster 0 continue to run because the logical volumes are using the TVC in Cluster 0.
- ▶ All scratch mounts to TS7700 Cluster 0 will succeed if it owns one or more volumes in the scratch category at the time of mount operation. After the scratch volumes owned by TS7700 Cluster 0 are exhausted, scratch mounts will begin to fail.
- ▶ The Grid enters the Grid Links Degraded state and the VTS Operations Degraded state.
- ▶ All copy operations are stopped.
- ▶ The Grid enters the Copy Operation Disabled state.
- ▶ If the RUN copy consistency point is being used, the Grid also enters the Immediate Mode Copy Completion's Deferred state.
- ▶ Call Home support is invoked.

Recovery from failure

Contact your IBM SSR for repair of the failed connection.

10.2.4 Failover scenario 3

In this scenario, shown in Figure 10-5, we assume a failure of a link between TS7700 Virtualization Engine Clusters with remote mounts.




Figure 10-5 Failure of a link between TS7700 Virtualization Engine Clusters with remote mounts

Effects of the failure

You will see the following effects of the failure:

- ▶ Any job in progress that is using the remote link between TS7700 Cluster 0 and TS7700 Cluster 1 that was disconnected will fail.
- ▶ If the job is resubmitted, it will succeed by using the other link.
- ▶ The Grid enters the Grid Links Degraded state and the VTS Operations Degraded state.
- ▶ Call Home support is invoked.

Recovery from failure

To recover from the failure, you must contact your IBM SSR to repair the failed connections.

10.2.5 Failover scenario 4

In this scenario, shown in Figure 10-6, we assume a failure of both links between TS7700 Virtualization Engine Clusters with remote mounts.




Figure 10-6 Failure of both links between TS7700 Virtualization Engine Clusters with remote mounts

Effects of the failure

You will see the following effects of the failure:

- ▶ Jobs on virtual device addresses on TS7700 Cluster 0 that are using TS7700 Cluster 1 as the TVC cluster will fail.
- ▶ Subsequent specific mount jobs that attempt to access the data through TS7700 Cluster 0 that exist only on TS7700 Cluster 1 will fail.
- ▶ All scratch mounts to TS7700 Cluster 0 will succeed if Cluster 0 owns one or more volumes in the scratch category at the time of mount operation. After the scratch volumes owned by TS7700 Cluster 0 are exhausted, scratch mounts will begin to fail.
- ▶ All copy operations are stopped.
- ▶ The Grid enters the Grid Links Degraded state, the VTS Operations Degraded state and the Grid enters the Copy Operation Disabled state.
- ▶ Call Home support is invoked.

Note: Although the data resides on TS7700-1, if it was mounted on TS7700-0 when the failure occurred, it is not accessible through the virtual device addresses on TS7700-1 because ownership transfer cannot occur.

Recovery from failure

To recover from the failures, you must contact your IBM SSR to repair the failed connections.

10.2.6 Failover scenario 5

In this scenario, shown in Figure 10-7, we assume a failure of the local TS7700 Virtualization Engine Cluster 0.




Figure 10-7 Failure of the local TS7700 Virtualization Engine Cluster 0

Effects of the failure

You will see the following effects of the failure:

- ▶ Virtual tape device addresses for TS7700 Cluster 0 will become unavailable.
- ▶ All channel activity on the failing host links are stopped.
- ▶ Host channel errors are reported or error information becomes available from the intermediate equipment.
- ▶ Jobs which were using the virtual device addresses of TS7700 Cluster 0 will fail.
- ▶ Scratch mounts that target volumes that are owned by the failed cluster will also fail until write ownership takeover mode is enabled. Scratch mounts that target pre-owned volumes will succeed. The Grid enters the Copy Operation Disabled and VTS Operations Degraded states.
- ▶ If the RUN copy consistency point is being used, the Grid also enters the Immediate Mode Copy Completion's Deferred state.
- ▶ All copied data can be made accessible through TS7700 Cluster 1 through one of the takeover modes. If a takeover mode for TS7700 Cluster 0 is not enabled, data will not be accessible through TS7700 Cluster 1 even if it has a valid copy of the data if the volume is owned by TS7700 Cluster 0.

Recovery from failure

To recover from the failures, you must perform the following steps:

- ▶ Enable write or read-only ownership takeover through the management interface.
- ▶ Rerun the failed jobs using the virtual device addresses associated with TS7700 Virtualization Engine Cluster 1.
- ▶ Normal error recovery procedures and repair apply for the host channels and the intermediate equipment.
- ▶ Contact your IBM SSR to repair the failed TS7700 Cluster.

10.2.7 Failover scenario 6

In this scenario, shown in Figure 10-8, we consider a failure of both links between TS7700 Virtualization Engine Clusters with Automatic Takeover.




Figure 10-8 Failure of both links between TS7700 Virtualization Engine Clusters with Automatic Takeover

Effects of the failure

You will see the following effects of the failure:

- ▶ Specific mount jobs subsequent to the failure using virtual device addresses on Cluster 0 that need to access volumes that are owned by Cluster 1 will fail (even if the data is local to Cluster 0). Jobs using virtual device addresses on Cluster 1 that need to access volumes that are owned by Cluster 0 will also fail.
- ▶ All scratch mounts to Cluster 0 will succeed so long as it owns one or more volumes in the scratch category at the time of mount operation. After the scratch volumes owned by Cluster 0 are exhausted, scratch mounts will begin to fail.
- ▶ All copy operations are stopped.
- ▶ The Grid enters the Grid Links Degraded state, the VTS Operations Degraded state and the Copy Operation Disabled state.
- ▶ If the RUN copy consistency point is being used, the Grid also enters the Immediate Mode Copy Completion's Deferred state.
- ▶ Call Home support is invoked.

Recovery from failure

To recover from the failures, contact your IBM SSR for repair of the failed connections.

10.2.8 Failover scenario 7

In this scenario we assume a production site with two TS7700 clusters (Cluster 0 and Cluster 1) active in production. The third TS7700 Cluster (Cluster 2) is located at a remote location without attachment to the production hosts. Cluster 3 is attached to a backup host, devices varied offline, and no active host, as shown in Figure 10-9.

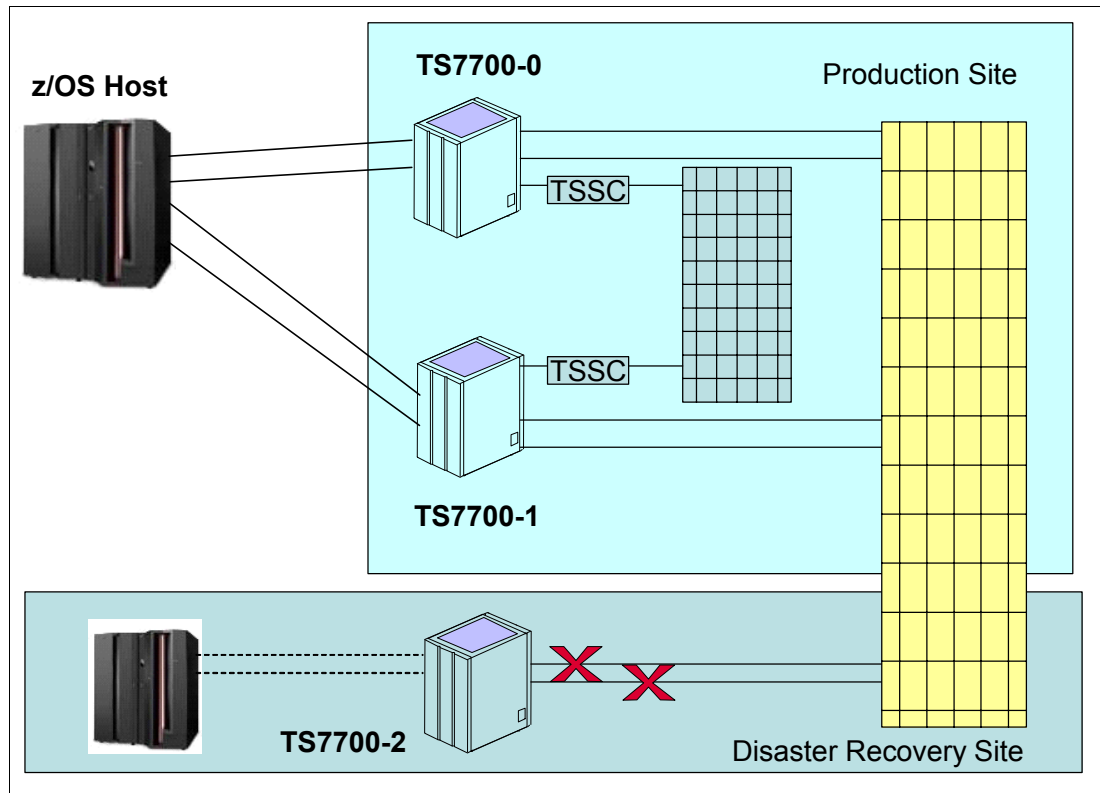


Figure 10-9 Three-cluster grid with failure on two links to Cluster 2

Different failures related to Cluster 0 and Cluster 1 are already described in the previous scenarios of this chapter. In Figure 10-9, we consider what to do when both links to Cluster 2 have failed and the only shared component from Cluster 0 and Cluster 1 to Cluster 2 is the network.

Effects of the failure

You will see the following effects of the failure:

- ▶ All copy operations between Cluster 2 and rest of the clusters are stopped.
- ▶ All copy operations between Cluster 0 and Cluster 1 continue.
- ▶ The Grid enters the Grid Links Degraded state, the VTS Operations Degraded state and the Copy Operations Disabled state.
- ▶ If the RUN copy consistency point is being used for Cluster 2, the Grid also enters the Immediate Mode Copy Completion's Deferred state.
- ▶ Call Home support is invoked.

Recovery from failure

Contact your IBM SSR for repair of the failed connections.

10.2.9 Failure scenario 9

In this scenario, we assume a four-cluster hybrid grid configuration with partitioned workload. At the production site, two TS7720 clusters are installed. At the remote site, two TS7740 clusters, attached to TS3500 tape libraries, are installed.

Virtual volumes are written on one cluster at the local site and copied to one cluster at the remote site, so that a copy of a volume will exist in Cluster 0 and Cluster 2 on in Cluster 1 and Cluster 3.

In the scenario, shown in Figure 10-10, we assume that the remote site fails. The Grid WAN is operational.

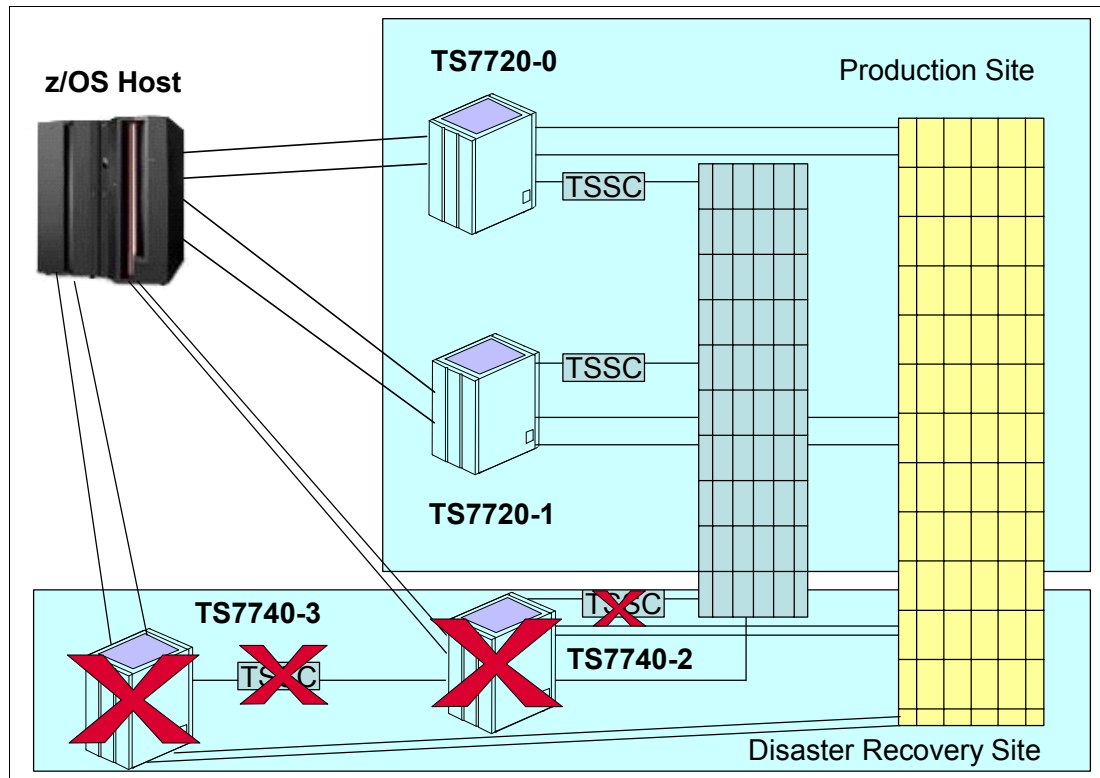


Figure 10-10 Four-cluster hybrid grid multiple failures

Effect of the failures

You will see the following effects of the failures:

- ▶ Jobs on virtual device addresses on Cluster 0 will continue to run because the logical volumes are in the TVC on Cluster 0 or Cluster 1.

Jobs that access old volumes, which the automatic removal mechanisms have already removed from the production clusters, will fail. Because TS7720s cannot copy to TS7740, they might eventually become full and all scratch mounts and specific mounts with modifications will fail.

- ▶ The Grid enters the Copy Operation Disabled and VTS Operations Degraded states.
- ▶ If the RUN copy consistency point is being used, the Grid also enters the Immediate Mode Copy Completion's Deferred state.
- ▶ All copy operations for Cluster 2 and Cluster 3 are stopped.
- ▶ Call Home support is invoked.

Recovery from failure

Normal error recovery procedures and repair apply for the host channels and the intermediate equipment. To recover from the failures, you must contact your IBM SSR to repair the failed connections.

10.3 Planning for disaster recovery

Although we can hope that a disaster does not happen, planning for such an event is important. This section provides information that can be used in developing a disaster recovery plan as it relates to a TS7700 Virtualization Engine.

Many aspects of disaster recovery planning must be considered:

- ▶ How critical is the data in the TS7700 Virtualization Engine?
- ▶ Can the loss of some of the data be tolerated?
- ▶ How much time can be tolerated before resuming operations after a disaster?
- ▶ What are the procedures for recovery and who will execute them?
- ▶ How will you test your procedures?

10.3.1 Grid configuration

With the TS7700 Virtualization Engine, two types of configurations can be installed:

- ▶ Single-cluster grid
- ▶ Multi-cluster grid

With a stand-alone system, a single TS7700 Virtualization Engine Cluster is installed. If the site at which that system is installed is destroyed, the data that is associated with the TS7700 Virtualization Engine might also have been destroyed. If a TS7700 Virtualization Engine is not usable, because of interruption of utility or communication services to the site, or significant physical damage to the site or the TS7700 Virtualization Engine itself, access to the data that is managed by the TS7700 Virtualization Engine is restored through automated processes designed into the product.

The recovery process assumes that the only elements available for recovery are the stacked volumes themselves. It further assumes that only a subset of the volumes is undamaged after the event. If the physical cartridges have been destroyed or irreparably damaged, recovery is not possible, as with any other cartridge types. It is important that you integrate the TS7700 Virtualization Engine recovery procedure into your current disaster recovery procedures.

Note: The disaster recovery process is a joint exercise that requires your involvement and your IBM SSR to make it as comprehensive as possible.

For many customers, the potential data loss or the recovery time required with a stand-alone TS7700 Virtualization Engine is not acceptable. For those customers, the TS7700 Virtualization Engine Grid provides a near-zero data loss and expedited recovery time solution. With a TS7700 Virtualization Engine multi-cluster grid configuration, two, three, or four TS7700 Virtualization Engine Clusters are installed, typically at two or three sites, and interconnected so that data is replicated between them. The way the two or three sites are used then differs, depending on customer requirements.

In a two-cluster grid, the typical use will be that one of the sites is the local production center and the other is a backup or disaster recovery center, separated by a distance dictated by your company's requirements for disaster recovery.

In a three-cluster grid, the typical use will be that two sites are connected to a host and the workload is spread evenly between them. The third site is strictly for disaster recovery and there will probably not be connections from the production host to the third site. Another use of a three-cluster grid might consist of three production sites, all interconnected and holding the backups of each other.

In a four-cluster grid, disaster recovery and high availability can be achieved, ensuring that two local clusters keep RUN volume copies and both clusters are attached to the host; the third and fourth remote clusters hold deferred volume copies for disaster recovery. This way can be configured in a crossed way, meaning that customer could run two production data centers, each serving as a backup for the other.

The only connection between the production sites and the disaster recovery site is the Grid interconnection. There is normally no host connectivity between the production hosts and the disaster recovery site's TS7700 Virtualization Engine. When client data is created at the production sites, it is replicated to the disaster recovery site as defined through outboard policy management definitions and SMS settings.

10.3.2 Planning guidelines

As part of planning a TS7700 Virtualization Engine Grid configuration to address this solution, you need to consider the following items:

- ▶ Plan for the necessary WAN infrastructure and bandwidth to meet the copy requirements that you need. You generally will need more bandwidth if you are primarily using a copy consistency point of RUN, because any delays in copy time caused by bandwidth limitations will result in an elongation of job run times. If you have limited bandwidth available between sites, use the deferred copy consistency point or only copy the data that is critical to the recovery of your key operations. The amount of data sent through the WAN can possibly justify the establishment of a separate, redundant, and dedicated network only for the multi-cluster grid.
- ▶ If you use a consistency point of deferred copy, and the bandwidth is the limiting factor, it is possible that some data has not been replicated between the sites, and the jobs that created that data must be rerun.
- ▶ Plan for host connectivity at your disaster recovery site with sufficient resources to perform your critical workloads. If the local TS7700 Virtualization Engine Cluster becomes unavailable, there is no local host access to the data in the disaster recovery site's TS7700 Virtualization Engine Cluster through the local cluster.
- ▶ Design and code the Data Facility System Management Subsystem (DFSMS) Automatic Class Selection routines to control what data gets copied and by which copy consistency point. You might need to consider management policies for testing your procedures at the disaster recovery site that are different from the production policies.
- ▶ Prepare procedures that your operators would execute in the event the local site becomes unusable. The procedures would include such tasks as bringing up the disaster recovery host, varying the virtual drives online, and placing the disaster recovery TS7700 Virtualization Engine Cluster in one of the ownership takeover modes.
- ▶ Do a periodic capacity planning of your tape setup to evaluate whether the disaster setup is still capable of handling the production in case of a disaster.

- ▶ If encryption is used in production, make sure that the disaster site supports encryption also. The Key Encrypting Keys (KEK), used with TS7700 R1.6 for production must be available at the disaster recovery site to enable the data key to be decrypted. With TS7700 R1.7, default keys are supported and enable key management without modifications required on the TS7740. On the tape setup, the TS1120/TS1130, the TS7700 Virtualization Engine, and the MI itself must support encryption. Validate that the TS7700 Virtualization Engine can communicate with the Encryption Key Manager (EKM) or Tivoli Key Lifecycle Manager (TKLM) and that the keystore itself is available.
- ▶ Consider how you will test your disaster recovery procedures. Many scenarios can be set up:
 - Will it be based on all data from an existing TS7700 Virtualization Engine?
 - Will it be based on using the Copy Export function and an empty TS7700 Virtualization Engine?
 - Will it be based on stopping production of one TS7700 Virtualization Engine and running production to the other during a period of time when one cluster is down for service?

10.3.3 Disaster recovery considerations with TS7740 and TS7720

Although we know the differences between a TS7740 Virtualization Engine and a TS7720 Virtualization Engine, from a DR point of view, these differences mean different DR capabilities depending on the configuration. Figure 10-11 shows three scenarios:

- ▶ A two-cluster grid with TS7720 Virtualization Engines only inside the dotted line
- ▶ A two-cluster hybrid grid with one TS7720 and one TS7740 inside the dashed line
- ▶ A two-cluster grid with TS7740 Virtualization Engines only inside the solid line




Figure 10-11 TS7700 Virtualization Engine two-cluster grid configuration

No differences exist between these configurations regarding failover scenarios (see 10.2, “Failover scenarios” on page 729) or GDPS (see 10.5, “Geographically Dispersed Parallel Sysplex” on page 752). However, because the TS7720 is not attached to physical tape drives, all DR testing or recovery actions that involve physical tape media are not supported in a Grid configuration consisting of TS7720 Virtualization Engines only. Copy Export Recovery (see 10.4, “Disaster recovery using Copy Export” on page 742) requires physical tape attachment and is therefore possible only in Grid configurations containing at least one TS7740 Cluster.

10.4 Disaster recovery using Copy Export

Copy Export provides a function to allow a copy of selected logical volumes written to the TS7700 Virtualization Engine to be removed and taken off site for disaster recovery purposes. All the original logical or physical volumes are intact and usable for the production system. In 5.3.2, “Implementing Copy Export” on page 288, we explain how to set up and use the Copy Export function. In the following sections, we provide an example of how to use a Copy Export Recovery process.

Note: You can only execute a Copy Export Recovery process in a single-cluster grid. After recovered, you can create a multi-cluster grid by joining the Grid with another single-cluster grid.

We offer instructions for how to implement and execute Copy Export Recovery. Note that these instructions also apply if you are running a DR test. If it is a test, it is specified in each step. For more details and error messages related to the Copy Export function, see the *IBM Virtualization Engine TS7700 Series Copy Export Function User's Guide* white paper, which is available at the following address:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101092>

Note: Copy Export is not applicable to the TS7720 Virtualization Engine. In a Hybrid Cluster, data can be moved from the TS7720 Virtualization Engine to the TS7740 Virtualization Engine Grid, and then you can run Copy Export on the data.

10.4.1 Planning and considerations for testing Copy Export Recovery

You must consider several factors when you prepare a recovery TS7700 Virtualization Engine for the Copy Export volumes. Copy Export Recovery can be executed in various ways. In this section, we list the planning considerations for Copy Export recovery.

TS7740 R1.7 Copy Export Recovery

Copy Export recovery can be used to restore previously created and copy exported tapes to a new, empty TS7740 Cluster. Beginning with R1.7, the same subset of tapes can be used to restore a TS7740 in an existing Grid as long as the new empty restore cluster will replace the no longer present source cluster.

This allows data that may have existed only within a TS7740 within a hybrid configuration to be restored while maintaining access to the still existing TS7720 clusters. This form of extended recovery must be carried out by IBM support personnel.

Customer-initiated Copy Export Recovery

Customer-initiated recovery restores copy-exported tapes to a stand-alone TS7740 for DR testing or as a recovery site. In this section, we describe the considerations for Copy Export

Recovery to a stand-alone TS7740 Cluster, which can be prepared in advance, for DR testing, for example. The TS7700 Virtualization Engine and associated library that is to be used for recovery of the copy exported logical volumes must meet the following requirements:

- ▶ The recovery TS7700 Virtualization Engine must have physical tape drives that match the capabilities of the source TS7700 Virtualization Engine, including encryption capability if the copy exported physical volume have been encrypted.
- ▶ If the source copy exported volumes have been encrypted, the recovery TS7700 Virtualization Engine must have access to a key manager that has the encryption keys for the data.
- ▶ There must be sufficient library storage slots in the library associated with the recovery TS7700 Virtualization Engine to hold all of the copy exported physical volumes from the source TS7700 Virtualization Engine.
- ▶ Only the copy exported volumes from a single source TS7700 Virtualization Engine can be used in the recovery process.
- ▶ The recovery TS7700 Virtualization Engine cannot be part of a Grid configuration.
- ▶ The recovery TS7700 Virtualization Engine must be configured as Cluster 0.
- ▶ The recover TS7700 Virtualization Engine and its associated management interface must be configured, have code loaded, and be in an online state to start the recovery.
- ▶ The code levels on the recovery TS7700 must be at the same or later code level than the source TS7700.
- ▶ If the recovery TS7700 Virtualization Engine is not empty of data (in the cache or the database), the Copy Export volumes must not be loaded into the attached library until the machine has been emptied of data. The procedures detailed in 10.4.3, “Performing Copy Export Recovery” on page 745 identifies when the physical volumes need to be loaded into the library and how to empty the recovery TS7700 Virtualization Engine of data.
- ▶ If another TS7700 Virtualization Engine or native drives are on another partition of the TS3500 Tape Library, the other partition must not have any VOLSERS that overlap with the VOLSERS to be recovered (including both logical and physical volumes). If any conflicts are encountered during the recovery process, the VOLSERS that conflict cannot be recovered, a warning message is displayed in the recovery status window on the recovery TS7700 Virtualization Engine management interface, and you cannot use the same library for both the source and recovery TS7700 Virtualization Engine.
- ▶ Other than the physical drive compatibility requirements listed, the source and recovery TS7700 Virtualization Engine can have a different configuration of features, that is, different cache capabilities, performance enablement features, and so on.
- ▶ You must add scratch physical volumes to the recovery TS7700 Virtualization Engine even if you are only going to be reading data. A minimum of two scratch volumes per defined pool in the recovery TS7740 is needed to prevent the recovery TS7740 from entering the out-of-scratch state. In the out-of-scratch state, logical volume mounts are not allowed. When adding scratch physical volumes to the recovery TS7740, do so only after the recovery has been performed and the recovery TS7740 is ready to be brought online to its attached hosts, otherwise their inventory records will have been erased during the recovery process. Physical volumes that are part of the copy export set and are now empty cannot be counted as scratch. After the Copy Export recovery is complete, and the recovery TS7740 Virtualization Engine is online to its hosts, you must insert logical volumes to be used as scratch volumes before you can write new data.
- ▶ If the recovery is for a real disaster (rather than only a test), verify that the actions defined for the storage management constructs that were restored during the recovery are what you want to continue to use.

10.4.2 Grid considerations for Copy Export

In a Grid configuration, a Copy Export operation is performed against an individual TS7740 Virtualization Engine, not across all TS7740 Virtualization Engines. In setting up to use Copy Export in a Grid, plan for the following items:

- ▶ Decide on which TS7740 Virtualization Engine in a Grid configuration will be used to export a specific set of data. Although you can set up more than one TS7740 can be used in the recovery process. You cannot merge copy exported volumes from more than one source TS7740 in the recovery TS7740.
- ▶ For each specific set of data to export, define a Management Class name. As described in “Setting up data management definitions” on page 292, on the TS7740 that will be used to export that data, define a secondary physical volume pool for that Management Class name and also make sure that you indicate that it is an export pool. Although you must define the Management Class name on all TS7700 Virtualization Engines in the Grid configuration, specify only the secondary physical volume pool on the TS7700 Virtualization Engine that is to perform the export operation. Specifying it on the other TS7700 Virtualization Engines in the Grid configuration does not interfere with the Copy Export operation, but it is a waste of physical volumes. The exception to this approach might be if you want one of the TS7700 Virtualization Engines in the Grid configuration to have a second physical copy of the data if the primary copies on other TS7740 Virtualization Engines are not accessible.
- ▶ As you are defining the Management Class name for the data, also make sure that the TS7700 Virtualization Engine to perform the export operation will have a copy policy specifying that it will have a copy.
- ▶ When the Copy Export operation is executed, the export list file volume must only be valid on the TS7700 Virtualization Engine performing the operation. You need to define a unique Management Class name to be used for the export list file volume. For that Management Class name, you need to define its copy policy so that a copy is only on the TS7700 Virtualization Engine that is to perform the export operation. If the VOLSER specified for the export list file volume when the export operation is initiated is resident on more than one TS7700 Virtualization Engine, the Copy Export operation will fail.

For more details and error messages related to the Copy Export function, see the *IBM Virtualization Engine TS7700 Series Copy Export Function User's Guide Version 1.4* white paper.

10.4.3 Performing Copy Export Recovery

Perform the following steps:

1. With the TS7740 and library in an online state, but with all virtual tape drives varied offline to any attached hosts, log in to the management interface and select **Service & Troubleshooting** → **Copy Export Recovery**, as shown in Figure 10-12.



Figure 10-12 Copy Export Recovery option in MI

You see only the Copy Export Recovery menu item, if you have been given access to that function by the overall system administrator on the TS7700. It is also not displayed if the TS7700 is configured in a Grid configuration. Contact your IBM SSR if you have to recover a TS7740 that is a member of a Grid.

2. The TS7740 determines that data exists in the cache and so do database entries. Copy Export Recovery cannot be performed until the TS7740 is empty. Figure 10-13 shows the window that opens, informing the operator that the TS7740 contains data that must be erased.

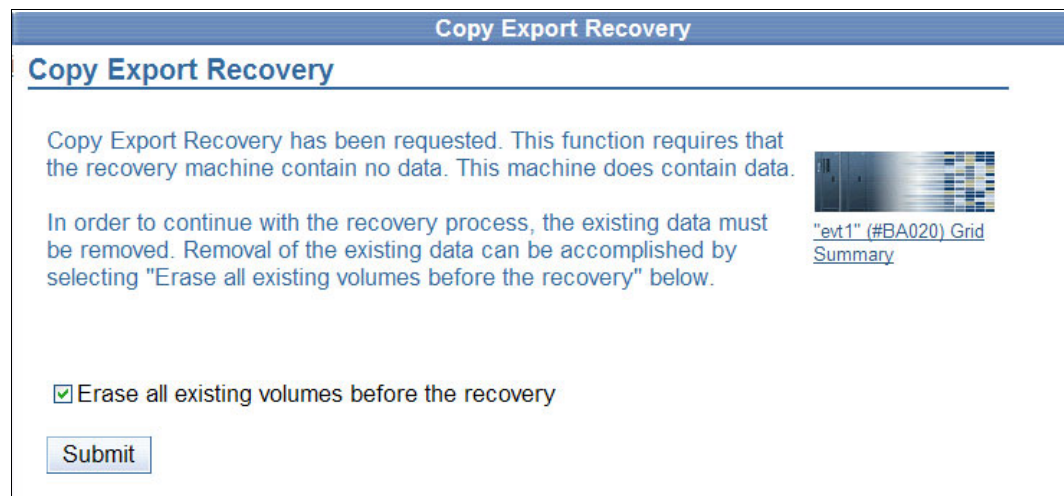


Figure 10-13 Copy Export Recovery window with erase volume option

3. Ensure that you are using the correct TS7740 (check to make sure you are logged in to the correct TS7740) and if so, select the **Erase all existing volumes before the recovery** check box and click **Submit**.

- The window shown in Figure 10-14 opens, providing you with the option to confirm and continue the erasure of data on the recovery TS7740 or to abandon the recovery process. It describes what data records are going to be erased and informs you of the next action to be taken. To erase the data, enter your login password and click **Yes**.

Confirm submission of Copy Export Recovery

"Erase all existing volumes before the recovery" has been selected. Continuing will erase the following:

- All logical volumes in the database
- All physical volumes in the database
- All logical volumes in the cache
- All logical and physical volume records associated with this TS7700 Virtualization Engine's Library Manager partition
- All database backups

If you are sure you want to continue, click the "Yes" button to log you out of the Management Interface. The TS7700 Virtualization Engine will be taken offline and all existing data will be erased. You should wait at least one minute before logging back in to the Management Interface. After you log in, the navigation will contain access to the Service & Troubleshooting -> Copy Export Recovery State page, which will allow you to monitor the erasure recovery process. At the conclusion of the erasure, the TS7700 Virtualization Engine will be placed online and the recovery process can be started. A click on the "No" button will abandon the erasure of existing data.

If you are sure you want to continue, re-enter your password in the box below.

*Password:

Figure 10-14 Confirm submission of erase data in TS7740

The TS7740 begins the process of erasing the data and all database records. As part of this step, you are logged off from the management interface.

- After waiting about one minute, log in to the management interface. Because the TS7740 is performing the erasure process, the only selection that is available through the Service & Troubleshooting menu is the **Copy Export Recovery State** window. Select that window to follow the progress of the erasure process. The window provides information about the process, including the total number of steps required, the current step, when the operation was initiated, the run duration, and overall status, as shown in Figure 10-15.

Copy Export Recovery State

Last refresh: Wednesday, June 6, 2007 1:32:06 PM

Total number of steps	5
Current step number	1
Start Time	Wednesday, June 6, 2007 1:28:46 PM
Duration	0 hours, 3 minutes, 21 seconds
Status	In Progress

Task details:

Taking the TS7700 offline.

Figure 10-15 Copy Export Recovery State window

6. To update the Copy Export Recovery State window, click **Refresh**. You may log off from the management interface by clicking **Logout**. Logging off from the MI does not terminate the erasure process. You can log back into the MI at a later time, and if the erasure process is still in progress, the window provides the latest status. If the process had completed, the Copy Export Recovery State menu item is not available.

The following tasks are listed in the task detail window as the erasure steps are being performed:

- Taking the TS7700 offline.
- The existing data in the TS7700 database is being removed.
- The existing data in the TS7700 cache is being removed.
- Cleanup (removal) of existing data.
- Requesting the TS7700 go online.
- Copy Export Recovery database cleanup is complete.

After the erasure process has completed, the TS7740 returns to its online state and you can continue with Copy Export Recovery.

7. Starting with an empty TS7740, you must perform several setup tasks by using the MI that is associated with the recovery TS7740 (for many of these tasks, you might only have to verify that the settings are correct; the settings are not deleted as part of the erasure step):
 - a. Verify or define the VOLSER range or ranges for the physical volumes that are to be used for and after the recovery. The recovery TS7740 must know what VOLSER ranges it owns. This step is done through the MI that is associated with the recovery TS7740.
 - b. If the copy exported physical volumes were encrypted, set up the recovery TS7740 for encryption support and have it connected to an external key manager that has access to the keys used to encrypt the physical volumes. If you will write data to the recovery TS7740, you must also define the pools to be encrypted and set up their key label or labels or define to use default keys.
 - c. If you are executing the Copy Export Recovery operations to be used as a test of your disaster recovery plans and have kept the Disaster Recovery Test Mode check box selected, the recovery TS7740 does not perform reclamation.

If you are running Copy Export Recovery because of a real disaster, verify or define the reclamation policies through the MI.

8. With the TS7740 in its online state, but with all virtual tape drives varied offline to any attached hosts, log in to the management interface and select **Service & Troubleshooting** → **Copy Export Recovery**, as shown in Figure 10-12 on page 745.

The TS7740 determines that it is empty and the window shown in Figure 10-16 opens.

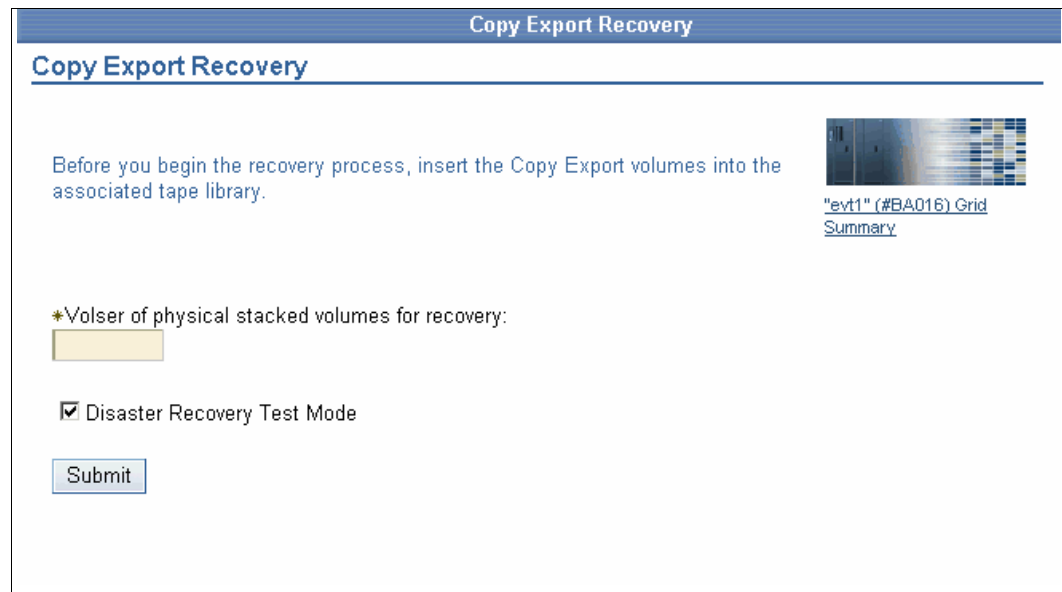


Figure 10-16 Copy Export Recovery window

At this time, load the copy exported physical volumes into the library. Multiple sets of physical volumes have likely been exported from the source TS7740 over time. All of the exported stacked volumes from the source TS7740 must be loaded into the library. If multiple pools were exported and you want to recover with the volumes from these pools, load all sets of the volumes from these pools. However, be sure that the VOLSER you provided is from the latest pool that was exported, because it will then have the latest overall database backup copy.

Notes:

- ▶ Before continuing the recovery process, be sure that all the copy-exported physical volumes have been added. Any volumes not known to the TS7740 when the recovery process continues will not be included and can lead to errors or problems. You can use the Physical Volume Search window from the MI to verify that all inserted physical volumes are known to the TS7740.
- ▶ Do not add any physical scratch cartridges at this time. You can do that after the Copy Export recovery operation has completed and you are ready to bring the recovery TS7740 online to the hosts.

9. After you have added all of the physical volumes into the library and they are now known to the TS7740, use the window shown in Figure 10-16 to enter the volume serial number of one of the copy exported volumes from the last set exported from the source TS7740. It contains the last database backup copy, which will be used to restore the recovery TS7740 Virtualization Engine's database. The easiest place to find a volume to enter is from the Export List File Volume Status file from the latest Copy Export operation.

If you are using Copy Export Recovery operation to perform a disaster recovery test, keep the **Disaster Recovery Test Mode** check box selected. The normal behavior of the TS7740 storage management function, when a logical volume in the cache is unloaded, is to examine the definitions of the storage management constructs associated with the volume. If the volume had been written to while it was mounted, the actions defined by the storage management constructs are taken. If the volume had not been modified, actions

are only taken if there has been a change in the definition of the storage management constructs since the last time the volume was unloaded. For example, if a logical volume is assigned to a storage group which had last had the volume written to pool 4 and either the storage group was not explicitly defined on the recovery TS7700 or specified a different pool, on the unload of the volume, a new copy of it will be written to the pool determined by the new storage group definition, even though the volume was only read. If you are just accessing the data on the recovery TS7700 for a test, you would not want the TS7700 to recopy the data. Keeping the check box selected causes the TS7700 to bypass its checking for a change in storage management constructs.

Another consideration with just running a test is reclamation. Running reclamation while performing a test will require scratch physical volumes and exposing the copy exported volumes to being reused after they are reclaimed. By keeping the **Disaster Recovery Test Mode** check box selected, the reclaim operation is not performed. With the **Disaster Recovery Test Mode** check box selected, the physical volumes used for recovery maintain their status of COPY-EXPORTED. This is so they cannot be re-used or used in a subsequent copy export operation. If you are using Copy Export because of a real disaster, clear the check box.

Enter the volume serial number, select the check box. and then click **Submit**.

10. The window shown in Figure 10-17 opens. It indicates the volume that will be used to restore the database. If you want to continue with the recovery process, click **Yes**. To abandon the recovery process, click **No**.

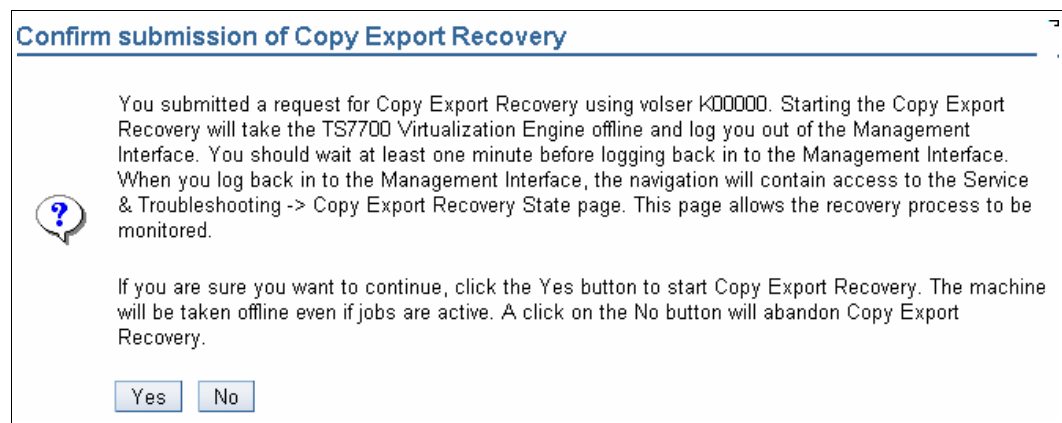


Figure 10-17 Confirm submission of Copy Export Recovery window

11. The TS7740 begins the recovery process. As part of this step, you are logged off from the management interface.
12. After waiting about 1 minute, log in to the management interface. Because the TS7740 is performing the recovery process, you can select only **Service & Troubleshooting → Copy Export Recovery State**. Select **Service & Troubleshooting → Copy Export Recovery State** to follow the progress of the recovery process.

The window provides information about the process, including the total number of steps required, the current step, when the operation was initiated, the run duration, and overall status, as shown in Figure 10-18.

Copy Export Recovery - State	
Refresh Last refresh: Wednesday, June 6, 2007 2:15:33 PM	
Total number of steps	12
Current step number	7
Start Time	Wednesday, June 6, 2007 2:03:24 PM
Duration	0 hours, 12 minutes, 10 seconds
Status	In Progress
Task details: Loading recovered data into the active database.	
Logout	

Figure 10-18 Copy Export Recovery - State window

The following tasks are listed in the task detail window as the Copy Export Recovery steps are performed:

- Taking the TS7700 offline.
- The requested recovery tape XXXXXX is being mounted on device YYY.
- The database backup is being retrieved from the specified recovery tape XXXXXX.
- The requested recovery tape is being demounted following retrieval of the database backup.
- The database backup retrieved from tape is being restored on the TS7700.
- The restored database is being updated for this hardware.
- The restored database volumes are being filtered to contain the set of logical volumes which were copy exported.
- Token ownership is being set to this cluster from previous cluster.
- The restored database is being reconciled with the contents of cache, XX of YY complete.
- Logical volumes are being restored on the Library Manager, XX of YY complete.
- Copy Export Recovery is complete.
- Copy Export Recovery from physical volume XXXXXX.
- Requesting the TS7700 go online.
- Loading recovered data into the active database.
- In progress.

After the Copy Export Recovery process completes successfully, the management interface returns to its full selection of tasks.

13. Now is the time to add scratch physical volumes to the library. Two scratch volumes are required for each active pool. Define the VOLSER range (or ranges) for the physical scratch volumes that are to be used for and after the recovery. The recovery TS7700 must know what VOLSER ranges it is to consider it owns. The steps are described in 4.3.3, “Defining VOLSER ranges for physical (stacked) volumes” on page 203.
14. If you had run Copy Export recovery because of a real disaster (you cleared the **Disaster Recovery Test Mode** check box), then verify that the defined storage management construct actions will manage the logical and physical volumes in the manner needed. During Copy Export Recovery, the storage management constructs and their actions will be restored to what was defined on the source TS7740. If you want the actions to be different, change them through the MI that is associated with the recovery TS7740.

If an error occurs, various possible error texts with detailed error descriptions can help you solve the problem. For more details and error messages related to the Copy Export Recovery function, see the *IBM Virtualization Engine TS7700 Series Copy Export Function User's Guide* white paper, which is available at the following address:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101092>

If everything is completed, you can vary the virtual devices online, continue the disaster recovery or DR testing, and the tapes are ready to read.

Note: For more general considerations about DR testing, see 10.6, “Disaster recovery testing considerations” on page 755.

10.4.4 Restoring the host and library environments

Before you can use the recovered logical volumes, you must restore the host environment also. We only list the minimum steps you need to continue the recovery process of your applications:

1. Restore the tape management system control data set.
2. Restore the DFSMS data catalogs, including the Tape Configuration Data Base (TCDB).
3. Define the I/O gen using the library ID of the recovery TS7740.
4. Update the library definitions in the SMS Control Data Set (SCDS) with the library IDs for the recovery TS7740 in the composite and distributed library definition windows.
5. Activate the I/O gen and the SCDS.

You might also want to update the library nicknames that are defined through the management interface for the Grid and cluster to match the library names defined to DFSMS. That way, the names shown on the management interface windows will match those used at the host for the composite and distributed libraries. To set up the composite name used by the host to be the Grid name, select **Configuration** → **Grid Identification Properties**. In the window that opens, enter the composite library name used by the host into the Grid nickname field. You can optionally provide a description. Likewise, to set up the distributed name, select **Configuration** → **Cluster Identification Properties**. In the window that opens, enter the composite library name used by the host into the Cluster nickname field. You can optionally provide a description. The names can be updated at any time.

10.5 Geographically Dispersed Parallel Sysplex

The IBM System z multi-site application availability solution, the Geographically Dispersed Parallel Sysplex (GDPS), integrates Parallel Sysplex technology and remote copy technology to enhance application availability and improve disaster recovery.

The GDPS topology is a Parallel Sysplex cluster spread across two sites, with all critical data mirrored between the sites. GDPS provides the capability to manage the remote copy configuration and storage subsystems, automates Parallel Sysplex operational tasks, and automates failure recovery from a single point of control, thereby improving application availability.

GDPS is a multi-site management facility incorporating a combination of system code and automation that utilizes the capabilities of Parallel Sysplex technology, storage subsystem mirroring, and databases to manage processors, storage, and network resources. It minimizes and potentially eliminates the impact of a disaster or planned site outage. The GDPS provides the ability to perform a controlled site switch for both planned and unplanned site outages, with no data loss, maintaining full data integrity across multiple volumes and storage subsystems, and the ability to perform a normal DBMS restart (not DBMS recovery) at the opposite site.

10.5.1 GDPS functions

GDPS provides the following functions:

- ▶ Remote Copy Management Facility (RCMF), which automates management of the remote copy infrastructure
- ▶ Planned reconfiguration support, which automates operational tasks from one single point of control
- ▶ Unplanned reconfiguration support, which recovers from a z/OS, processor, storage subsystem, or site failure

Remote copy management facility

RCMF was designed to simplify the storage administrator's remote copy management functions by managing the remote copy configuration rather than individual remote copy pairs. This includes the initialization and monitoring of the PPRC or XRC volume pairs based upon policy and performing routine operations on installed storage subsystems.

Planned reconfigurations

GDPS planned reconfiguration support automates procedures performed by an operations center. These include standard actions to:

- ▶ Quiesce a system's workload and remove the system from the Parallel Sysplex cluster (stop the system prior to a change window).
- ▶ IPL a system (start the system after a change window).
- ▶ Quiesce a system's workload, remove the system from the Parallel Sysplex cluster, and re-IPL the system (recycle a system to pick up software maintenance). The standard actions can be initiated against a single system or group of systems. Additionally, user-defined actions are supported (for example, a planned site switch in which the workload is switched from processors in site A to processors in site B).

Unplanned reconfigurations

GDPS was originally designed to minimize and potentially eliminate the amount of data loss and the duration of the recovery window in the event of a site failure; however, it will also minimize the impact and potentially mask an z/OS system or processor failure based upon GDPS policy. GDPS uses Point to Point Remote Copy (PPRC) or Extended Remote Copy (XRC) to help minimize or eliminate data loss. Parallel Sysplex cluster functions, along with automation, are used to detect z/OS system, processor, or site failures and to initiate recovery processing to help minimize the duration of the recovery window.

If a z/OS system fails, the failed system will automatically be removed from the Parallel Sysplex cluster, re-IPLed in the same location, and the workload restarted. If a processor fails, the failed system(s) will be removed from the Parallel Sysplex cluster, re-IPLed on another processor, and the workload restarted.

With PPRC, there will be limited or no data loss, based upon policy, because all critical data is being synchronously mirrored from site A to site B in the event of a site failure. Limited data loss occurs if the production systems continue to make updates to the primary copy of data after remote copy processing is suspended (any updates after a freeze will not be reflected in the secondary copy of data) and there is a subsequent disaster that destroys some or all of the primary copy of data. No data loss occurs if the production systems do not make any updates to the primary PPRC volumes after PPRC processing is suspended.

Depending on the type of application or recovery options selected by the enterprise, multiple freeze options are supported by GDPS (the freeze is always performed to allow the restart of the software subsystems).

10.5.2 GDPS considerations in a TS7740 Grid configuration

A key principle of GDPS is to have all I/O be local to the system running production. Another principle is to provide a simplified method to switch between the primary and secondary site if needed. The TS7740 Virtualization Engine in a Grid configuration provides a set of capabilities that can be tailored to allow it to operate efficiently in a GDPS environment. Those capabilities and how they may be used in a GDPS environment are described in this section.

Direct production data I/O to a specific TS7740

The hosts are directly attached to the TS7740 local to it, so that is your first consideration in directing I/O to a specific TS7740. Host channels from each site's GDPS hosts are also typically installed to connect to the TS7740 at the site remote to a host to cover recovery Scenario É on page 755, but during normal operation, the remote virtual devices are set offline in each GDPS host.

The default behavior of the TS7740 in selecting which TVC will be used for the I/O is to follow the Management Class definitions and considerations to provide the best overall job performance. It will, however, use a logical volume in a remote TS7740's TVC if required to perform a mount operation unless override settings on a cluster are used.

To direct the TS7740 to use its local TVC, perform the following steps:

1. For the Management Class used for production data, ensure that the local cluster has a copy consistency point. If it is important to know that the data has been replicated at job close time, specify a copy consistency point of Rewind/Unload (RUN). If some amount of data loss after a job closes can be tolerated, then a copy consistency point of Deferred can be used. You might have production data with different data loss tolerance; if that is the case, you might want to define more than one Management Class with different copy consistency points. In defining the copy consistency points for a Management Class, it is important that you define the same copy mode for each site, because in a site switch, the local cluster changes.
2. Set **Prefer Local Cache for Fast Ready Mounts** in the MI Copy Policy Override window. This override will select the TVC local to the TS7740 the mount was received on as long as it is available and a copy consistency point other than No Copy is specified for that Cluster in the Management Class specified with the mount. The Cluster does not have to have a valid copy of the data for it to be selected for the I/O TVC.
3. Set **Prefer Local Cache for Non-Fast Ready Mounts** in the MI Copy Policy Override window. This override will select the TVC local to the TS7740 the mount was received on as long as it is available and the cluster has a valid copy of the data, even if the data is only resident on a physical tape. Having an available, valid copy of the data overrides all other selection criteria. If the local cluster does not have a valid copy of the data, then without the next override, it is possible that the remote TVC would be selected.
4. Set **Force Volume Copy to Local**. This override has two effects, depending on the type of mount requested. For a non-fast ready mount, if a valid copy does not exist on the cluster, a copy is performed to the local TVC as part of the mount processing. For a fast ready mount, it has the effect of “ORing” the specified Management Class with a copy consistency point of Rewind/Unload for the cluster and this will force the local TVC to be used. The override does not change the definition of the Management Class; it serves only to influence the selection of the I/O TVC or force a local copy.
5. Make sure that these override settings are duplicated on both TS7740 Virtualization Engines.

Note: The MI Copy Policy Override window is shown in Figure 10-1 on page 727.

Switch site production from one TS7700 to the other

The way data is accessed by either TS7740 is based on the logical volume serial number. No changes are required in tape catalogs, JCL, or tape management systems. In case of a failure in a TS7740 Grid environment with GDPS, three scenarios can occur:

- ▶ GDPS switches the primary host to the remote location; the TS7740 Grid is still fully functional.
 - No manual intervention is required.
 - Logical volume ownership transfer is done automatically during each mount through the Grid.
- ▶ A disaster happens at the primary site; the GDPS host and TS7740 Cluster are down or inactive.
 - Automatic ownership takeover of volumes, which then will be accessed from the remote host, is not possible.
 - Manual intervention is required. Through the TS7740 management interface, the administrator has to invoke a manual ownership takeover. To do so, use the TS7740

management interface and select **Service and Troubleshooting** → **Ownership Takeover Mode**.

- ▶ Only the TS7740 Cluster at the GDPS primary site is down. In this case, two manual interventions are required:
 - Vary online remote TS7740 Cluster devices from the primary GDPS host.
 - Because the down cluster cannot automatically take ownership of volumes that will then be accessed from the remote host, manual intervention is required. Through the TS7740 management interface, the administrator has to invoke a manual ownership takeover. To do so, select **Service and Troubleshooting** → **Ownership Takeover Mode** in the TS7740 management interface.

10.6 Disaster recovery testing considerations

The TS7700 Virtualization Engine Grid configuration provides a solution for disaster recovery needs when data loss and the time for recovery are to be minimized. Although a real disaster is not something that can be anticipated, it is important to have tested procedures in place in case one occurs. As you design a test involving the TS7700 Virtualization Engine Grid configuration, there are several capabilities designed into the TS7700 Virtualization Engine that you should consider.

10.6.1 The test environment represents a point in time

The test environment is typically a point in time, meaning that at the beginning of the test, the catalog, TCDB, and tape management system control databases are all a snapshot of the production systems. Over the duration of the test, the production systems continue to run and make changes to the catalogs and TMS. Those changes are not reflected in the point-in-time snapshot. The main impact is that it is possible that a volume will be used in a test that has been returned to scratch status by the production system. The test system's catalogs and TMS will not reflect that change. If the links between the TS7700 Virtualization Engines remain connected, the TS7700 Virtualization Engine at the test location will be informed that a volume has been returned to scratch; it will not, however, prevent the test host from accessing the data on that volume. What is important is that a volume returned to scratch needed during the test is not re-used on the production system during the test. See 10.6.5, “Protecting production volumes with DFSMSrmm” on page 758 for further information about how to manage return-to-scratch handling during a test.

10.6.2 Breaking the interconnects between the TS7700 Virtualization Engines

The two approaches to conducting the test are as follows:

- ▶ The site-to-site links are broken.
- ▶ The links are left connected.

A test can be conducted with either approach, but each one has trade-offs. The main trade-offs for breaking the links are as follows:

- ▶ On the positive side:
 - You are sure that only the data that has been copied to the TS7700 Virtualization Engine connected to the test system is being accessed.
 - Logical volumes that are returned to scratch by the production system are not “seen” by the TS7700 Virtualization Engine under test.

- Test data that is created during the test is not copied to the other TS7700 Virtualization Engine.
- On the negative side:
 - If a disaster occurs while the test is in progress, data that was created by the production site after the links were broken is lost.
 - The TS7700 Virtualization Engine at the test site must be allowed to take over volume ownership (either read only or read write).
 - The TS7700 Virtualization Engine under test could select a volume for scratch that has already been used by the production system while the links were broken.

For most customers, the concern about losing data in the event of a disaster during a test is the major issue with using the link break method. The TS7700 Virtualization Engine has several design features that make valid testing possible without having to break the site-to-site links.

10.6.3 Writing data during the test

This test typically includes running a batch job cycle that creates new data volumes. This test can be handled in two ways:

- Have TS7700 Virtualization Engine available as the output target for the test jobs.
- Have a separate logical volume range that is defined for use only by the test system.

The second approach is the most practical in terms of cost. It would involve defining the VOLSER range to be used, defining a separate set of categories for scratch volumes in the DFSMS DEVSUP parmlib, and inserting the volume range into the test TS7700 Virtualization Engine prior to the start of the test. It is important that the test volumes inserted using the management interface are associated with the test system so that the TS7700 Virtualization Engine at the test site will have ownership of the inserted volumes.

If the links are to be kept connected during the time when the volumes are inserted, an important step is to make sure that the tape management system at the production site does *not* accept use of the inserted volume range, and that the tape management system at the test site does the following steps:

- Changes on production systems:
 - Use the RMM parameter REJECT ANYUSE(TST*), which means do not use VOLSERs named TST* here.
- Changes on the DR test systems:
 - Use the RMM parameter VLPOOL PREFIX(TST*) TYPE(S) to allow use of these volumes for default scratch mount processing.
 - Change DEVSUPxx to point to another scratch category. That would be the category of the TST* volumes.

Figure 10-19 might help you better understand what needs to be done to insert cartridges in a DR site to perform a DR test.

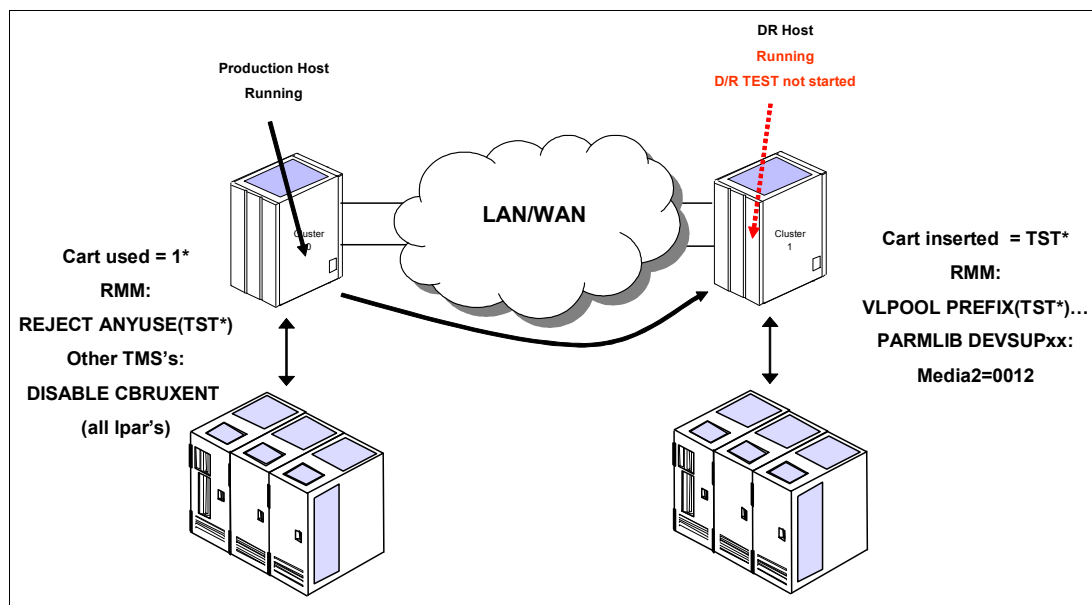


Figure 10-19 Insertion considerations in a DR test

After these settings are done, now is time to insert the new TST* logical volumes. Any new allocations that are performed by the DR test system will use only the logical volumes defined for the test. At the end of the test, the volumes can be returned to scratch status and left in the library, or deleted if desired.

Figure 10-19 shows that DR is in a running state, which means that the DR test itself is not started, but the DR system must be running before insertion can be done.

Note: Be aware that one logical unit must have been or is online on the test system before entering logical volumes. For more information, see 5.3.1, “z/OS and DFSMS/MVS System-managed tape” on page 286.

10.6.4 Protecting production volumes with TS7700 R1.7

While performing a test, you would not want the test system to inadvertently overwrite a production volume. You can put the TS7700 Virtualization Engine into the Write Protect mode (through the management interface) to prevent the test host from modifying production volumes. The Write Protect mode will prevent any host command issued to the test cluster from creating new data, modifying existing data, or changing volume attributes such as the volume category.

If you require that the test host be able to write new data, you may use the *selective write protect for DR testing* function that was introduced with TS7700 R1.7 which allows you to write to selective volumes during DR testing.

With Selective Write Protect, you may define a set of volume categories on the TS7700 that are excluded from the Write Protect Mode, thus enabling the test host to write data onto a separate set of logical volumes without jeopardizing normal production data, which remains write-protected. This way requires that the test host use a separate scratch category or categories from the production environment. If test volumes also must be updated, the test

host's private category must also be different from the production environment to separate the two environments.

You must determine the production categories that are being used and then define different, not yet used categories on the test host using the DEVSUPxx member. Be sure you define a minimum of four categories in the DEVSUPxx member: MEDIA1, MEDIA2, ERROR, and PRIVATE.

In addition to the host specification, you must also define on the TS7700 those volume categories that you are planning to use on the DR host and that need to be excluded from Write-Protect mode.

In 10.7.1, "Scenario 1 for a TS7700 two-cluster grid at R1.7" on page 764, instructions are provided regarding the necessary definitions for DR testing with a TS7700 Grid at R1.7.

The Selective Write Protect function enables you to read production volumes and to write new volumes from BOT while protecting production volumes from being modified by the DR host. Therefore, you cannot modify or append to volumes in the production hosts' PRIVATE categories, and DISP=MOD or DISP=OLD processing of those volumes is not possible.

10.6.5 Protecting production volumes with DFSMSrmm

As an alternative to the process described in 10.6.4, "Protecting production volumes with TS7700 R1.7" on page 757, if you are at a lower TS7700 microcode level and if you want to prevent overwriting production data, you can use the tape management system control to allow only read-access to the volumes in the production VOLSER ranges. However, the following process does not allow you to write data during the disaster recovery testing.

For example, with DFSMSrmm, you would insert these extra statements into the EDGRMMxx parmlib member:

- ▶ For production volumes in a range of A00000-A09999, add:
REJECT OUTPUT(A0*)
- ▶ For production volumes in a range of ABC000-ABC999, add:
REJECT OUTPUT(ABC*)

With REJECT OUTPUT in effect, products and applications that append data to an existing tape with DISP=MOD must be handled manually to function properly. If the product is DFSMSHsm, tapes that are filling (seen as not full) from the test system control data set must be modified to full by issuing commands. If DFSMSHsm then later needs to write data to tape, it will require a scratch volume related to the test system's logical volume range.

As a result of recent changes in DFSMSrmm, it should now be easier to manage this situation:

- ▶ In z/OS V1R10, the new rmm parmlib commands PRTITION and OPENRULE provides for flexible and simple control of mixed system environments as an alternative to the REJECT examples used here.
- ▶ In z/OS V1R9, you can specify additional EXPROC controls in the EDGHSKP SYSIN file to limit the return-to-scratch processing to specific subsets of volumes, so you could just EXPROC the DR system volumes on the DR system and the PROD volumes on the PROD system. You can still continue to run regular batch processing and also run expiration on the DR system.

Figure 10-20 helps us understand how we can protect our tapes in a DR test while our production system continues running.

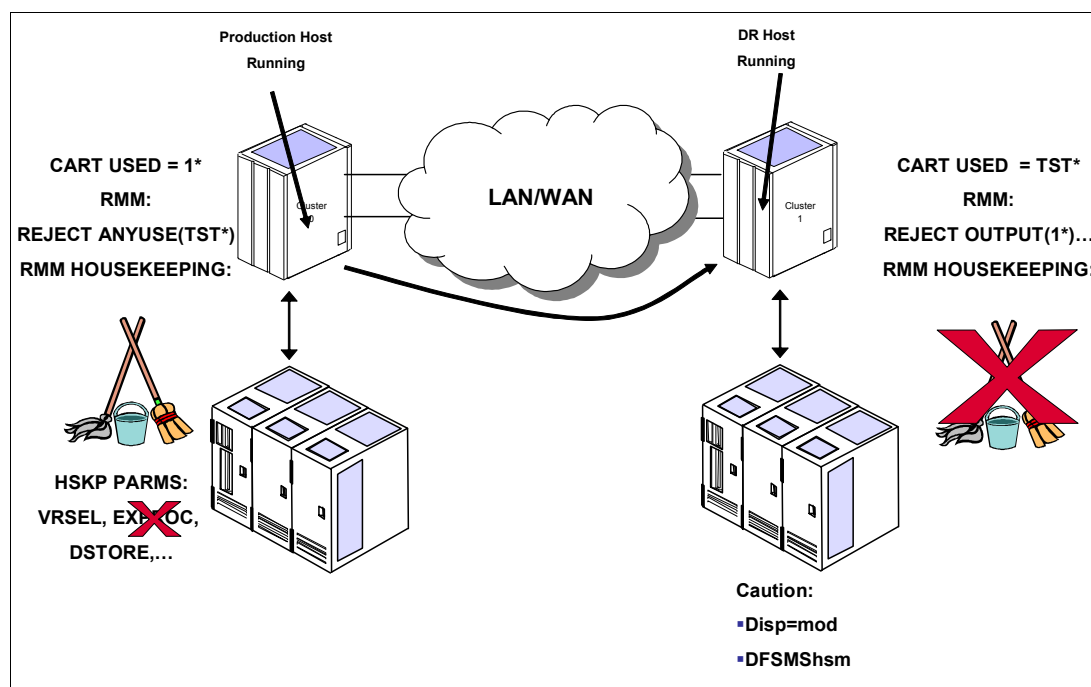


Figure 10-20 Work process in DR test

DR tests can last for days and many you may run return to scratch processing at this time. This is why we introduced the additional feature in Selective Write Protect to exclude fast ready characteristics.

Note: We refer to HSKP, because this is the name that we usually use in RMM. However, we are referring to the daily process that should be done with every other Tape Management System. From now on, when we refer to HSKP, which means the daily process on any other Tape Management System also.

This includes stopping any automatic short-on-scratch process, if enabled. For example, RMM has one emergency short-on-scratch procedure.

To explain what implications running the HSKP task in a DR test system has, we can use a example. In Table 10-1, we can see the status and definitions of one cartridge in a normal situation.

Table 10-1 VOLSER AAAAAA before returned to scratch from the DR site

Environment	DEVSUP	TCDB	RMM	MI	VOLSER
PROD	0002	Private	Master	000F	AAAAAA
DR	0012	Private	Master	000F	AAAAAA

Let us say cart AAAAAA is the master in both environments and that if there are any errors or mistakes, it is returned to scratch by the DR system. You can see its status in Table 10-2 on page 760.

Table 10-2 VOLSER AAAAAA after returned to scratch from the DR site

Environment	DEVSUP	TCDB	RMM	MI	VOLSER
PROD	0002	Private	Master	0012	AAAAAA
DR	0012	Scratch	Scratch	0012	AAAAAA

Now, we have cart AAAAAA in scratch category 0012. This presents two issues:

- ▶ If we need to access this volume from the Prod system, we need to change its status to master (000F) in the MI before we can access it. Otherwise, we lose the data in the cart, which can have serious consequences if you, for example, return to scratch 1000 volumes.
- ▶ Using DR RMM, reject using Prod carts to output activities. If this cart is mounted in response to a scratch mount, it will be rejected by RMM. Imagine that you have to mount 1000 scratch volumes because RMM rejected all of them before you get one validated by RMM.

To protect production volumes from unwanted return to scratch:

- ▶ Make sure that the RMM HSKP procedure is not running during the test window of the test system. There is a real risk of data loss if the test system returns production volumes to scratch and we have defined in TS7700 Virtualization Engine that the expiration time for virtual volumes is 24 hours. After this time, volumes can become unrecoverable.
- ▶ Make sure that the RMM short-on-scratch procedure does not start. The results could be the same as running a HSKP.

If you are going to perform the test with the site-to-site links broken, then you can use the Read Ownership Takeover mode to prevent the test system from modifying the production site's volumes. You can get further information about ownership takeover by referring to 10.6.9, "Ownership takeover" on page 762.

In addition to the protection options noted, you can also use the following RACF commands to protect the production volumes:

```
RDEFINE TAPEVOL x* UACC(READ) OWNER(SYS1)
SETR GENERIC(TAPEVOL) REFRESH
```

In the command, x is the first character of the VOLSER of the volumes to protect.

10.6.6 Control of copies

One of the issues with not breaking the links is that data being created as part of the test might be copied to the production site, wasting space and inter-site bandwidth. This situation can be avoided by defining the copy mode for the Management Classes differently at the test TS7700 Virtualization Engine than at the production TS7700 Virtualization Engine. Using a copy mode of No Copy for the production library site will prevent the test TS7700 Virtualization Engine from making a copy of the test data. It will not interfere with the copying of production data.

10.6.7 Return-to-scratch processing and test use of older volumes

In a test environment where the links between sites are not used, having the production system return logical volumes to scratch status that are to be used during the test for input is not an issue, because the TS7700 Virtualization Engine at the test site will not be aware of the change in status.

TS7700 R1.7 environments

With TS7700 R1.7, you can use the write protect definition **Ignore fast ready characteristics if write protected categories** to avoid conflicts. (Also see step 5 on page 766.) This approach only allows the DR host to read scratched volumes. It does not prevent the production host from using them. Either turning off return-to-scratch or configuring a long expire-hold time should be used too.

TS7700 pre-R1.7 environments

In a test environment where the links are maintained, care must be taken to ensure that logical volumes that are to be in the test are not returned to scratch status and used by production applications to write new data. There are several ways that to prevent conflicts between the return-to-scratch processing and the test use of older volumes:

1. Suspend all return-to-scratch processing at the production site. Unless the test is fairly short (hours, not days), this is not likely to be acceptable because of the risk of running out of scratch volumes, especially for native tape workloads. If all tape processing uses logical volumes, the risk of running out of scratch volumes can be eliminated by making sure that the number of scratch volumes available to the production system is enough to cover the duration of the test.

In z/OS V1R9 and later, you can specify additional EXPROC controls in the EDGHSKP SYSIN file to limit the return-to-scratch processing to specific subsets of volumes, so you could just EXPROC the DR system volumes on the DR system and the PROD volumes on the PROD system. Therefore, you can still continue to run regular batch processing and also run expiration on the DR system.

Keep in mind that if a volume is returned to scratch into a Fast-Ready category, even though DR knows that it is private (remember that TCDB and RMM is a snapshot of production data), mounting that volume through specific mount will not bring the data written to it (TS7700 Virtualization Engine will always mount a blank volume from a Fast-Ready category). It can be recovered by assigning the volume back to private or to a non-Fast-Ready category, or taking that category out of the Fast-Ready list and retrying the mount.

Even if the number of volumes in the list is larger than the number of volumes needed per day times the number of days of the test, you will still need to take steps to make it unlikely that a volume needed for test will be reused by production.

For more information, see the *IBM Virtualization Engine TS7700 Series Best Practices - Return-to-Scratch Considerations for Disaster Recovery Testing with a TS7700 Grid* white paper:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101281>

2. Suspend only the return-to-scratch processing for the production volume needed for the test. For RMM, this can be done using policy management through VRSs. A volume VRS can be set up that covers each production volume so that this will override any existing policies for data sets. For example, say the production logical volumes to be used in the test are in a VOLSER range of 990000-990999. To prevent them from being returned to scratch, the following subcommand would be run on the production system:

```
RMM AS VOLUME(990*) COUNT(99999) OWNER(VTSTEST) LOCATION(CURRENT) PRIORITY(1)
```

Then, EDGHSKP EXPROC can be run and not expire the data required for test.

After the test is finished, we have a set of tapes in TS7700 Virtualization Engine that belong to test activities. We need to decide what to do with these tapes. As a test ends, the RMM database and VOLCAT will probably be destaged (with all the data used in the test), but in the MI database, the tapes remain defined: one will be in master status and the others in scratch status.

Our guidance differs if the tapes will not be needed anymore, or if the tapes will be used in the future for future test activities. In the first case, manually release the volumes and then run EXPROC to return the volumes to scratch under RMM control. In the second case, you only have to manually release these volumes; the cartridges remain in the scratch status and ready for use.

Note: Although cartridges in MI remain ready to use, you must ensure that the next time you create the test environment that these cartridges are defined to RMM and VOLCAT. Otherwise, you will not be able to use them.

10.6.8 Copies flushed or kept as LRU

The default management of the data in the cache is different for host-created versus copied data. By default, data copied to a TS7700 Virtualization Engine from another TS7700 Virtualization Engine is preferred to be removed from the cache, largest first, but using the Last Recently Used algorithm.

To add some flexibility, we can use the cache management function (PG0 and PG1). In general, PG0 tapes are deleted first from cache. In lower activity periods, the smallest PG0 volumes are removed, but if the TS7700 Virtualization Engine is busy and immediately requires space, the largest PG0 volume is removed.

The default behavior for cache preference is done so that where a host is connected to both TS7700 Virtualization Engines in a Grid configuration, the effective cache size is the combination of both TS7700 Virtualization Engine TVCs. This way, more mount requests can be satisfied from the cache. These “cache hits” result in faster mount times because no physical tape mount is required. It does have the disadvantage in that in the case of a disaster, most of the recently copied logical volumes are not going to be resident in the cache at the recovery site, because they were copies and would have been managed to be removed from the cache.

You may modify cache behavior by using the SETTING Host Console command in two ways:

- ▶ COPYFSC enable/disable: When disabled, logical volumes copied into cache from a Peer TS7700 Virtualization Engine are managed as PG0 (prefer to be removed from cache).
- ▶ RECLPG0 enable/disable: When disabled, logical volumes that are recalled into cache are managed using the actions defined from the Storage Class construct associated with the volume as defined at the TS7700 Virtualization Engine.

10.6.9 Ownership takeover

If you will be performing the test with the links broken between sites, you must enable Read Ownership Takeover so that the test site can access the data on the production volumes owned by the production site. Because the production volumes are created by mounting them on the production site's TS7700 Virtualization Engine, that TS7700 Virtualization Engine will have volume ownership.

If you attempt to mount one of those volumes from the test system, without ownership takeover enabled, the mount will fail because the test site's TS7700 Virtualization Engine will not be able to request ownership transfer from the production site's TS7700 Virtualization Engine. By enabling Read Ownership Takeover, the test host will now be able to mount the production logical volumes and read their contents.

The test host will not be able to modify the production site-owned volumes or change their attributes. The volume looks to the test host as a write protected volume. Because the volumes that are going to be used by the test system for writing data were inserted through the management interface associated with the TS7700 Virtualization Engine at the test site, that TS7700 Virtualization Engine will already have ownership of those volumes and the test host will have complete read and write control of them.

Note: Never enable Write Ownership Takeover mode for a test! Write Ownership Takeover mode should only be enabled in the event of a loss or failure of the production TS7700 Virtualization Engine.

If you are not going to break the links between the sites, then normal ownership transfer will occur whenever the test system requests a mount of a production volume.

10.7 Disaster recovery testing detailed procedures

In 10.6, “Disaster recovery testing considerations” on page 755, we explain what to consider during DR test situations. We understand that the best we can do is give our best explanation, with detailed instructions that include all the steps necessary to do a DR test, such as pre-test task, post-test task, production host task, recovery site task, and so on.

We think that the best DR test is a “pseudo-real” DR test, which means stopping the production site and starting real production at the DR site. As we are aware of the difficulties of doing this task, in our scenarios we always assume that production must continue working during our DR test. The negative aspect of this approach is that DR test procedures and real disaster procedures can differ slightly.

Note: In a DR test on a pre-R1.7 TS7700 Grid with production systems running concurrently, be sure that no return-to-scratch or emergency short-on-scratch procedure is started in the test systems. We can return to scratch production tapes, as discussed in 10.6.5, “Protecting production volumes with DFSMSrmm” on page 758.

In a DR test on a R1.7 TS7700 Grid with production systems running concurrently, you can use the **Ignore fast ready characteristics of write-protected categories** option together with Selective Write Protect as described in 10.6.4, “Protecting production volumes with TS7700 R1.7” on page 757.

In the following sections, we describe procedures for four scenarios, depending on the TS7700 release level and on the Grid configuration and the connection status during the test. We cover the following scenarios:

1. Scenario 1 for a TS7700 two-cluster grid at R1.7

In this scenario, we describe the steps for performing a DR test by using the R1.7 DR testing enhancements. Whether the links between the clusters are broken or not is irrelevant, as has been explained before. See 10.7.1, “Scenario 1 for a TS7700 two-cluster grid at R1.7” on page 764.

2. Scenario 2 for a Pre-R1.7 two-cluster grid

In this scenario, we assume that the DR test is performed with production running in parallel on a TS7700 two-cluster grid. The links between both clusters are not broken, and we cannot use the DR enhancements that were introduced with TS7700 R1.7. See 10.7.2, “Scenario 2 for a TS7700 two-cluster grid Pre-R1.7” on page 767.

3. Scenario 3 for a Pre-R1.7 two-cluster grid

In this scenario, we assume that the DR test is performed on a TS7700 two-cluster on pre-R1.7 with the links broken between both clusters so the production cannot be affected by the DR test. See 10.7.3, “Scenario 3 for a TS7700 two-cluster grid Pre R1.7” on page 770.

4. Scenario 4 for a pre-R1.7 three-cluster grid

This scenario is similar to scenario 2, but we are running production in parallel on a three-cluster grid. The links between both clusters are not broken, and we cannot use the DR enhancements that were introduced with TS7700 R1.7. See 10.7.4, “Scenario 4 for a three-cluster grid” on page 773.

10.7.1 Scenario 1 for a TS7700 two-cluster grid at R1.7

Figure 10-21 shows a sample multi-cluster grid scenario using Selective Write Protect. In the remainder of this section, we refer to the left cluster as the *Production Cluster* and the right cluster as the *DR Cluster*.

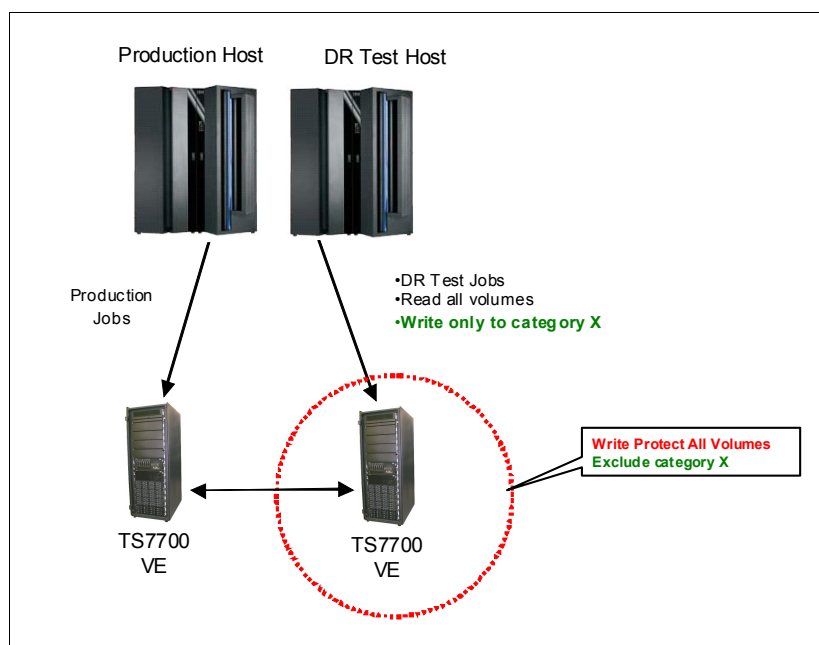


Figure 10-21 Sample DR testing scenario with TS7700 R1.7

Note: You can also use the steps described in the following procedure when performing DR testing on one cluster within a three- or four-cluster grid. To perform DR testing on more than one host or cluster, repeat the steps in the procedure on each of the DR hosts and clusters involved in the test.

Perform the following steps to prepare your DR environment accordingly:

1. Vary all virtual drives of the DR Cluster offline to the normal production hosts and to the DR hosts, and make sure that the production hosts have access to the Production Cluster so that normal tape processing can continue.
2. On the MI, select **Configuration** → **Write Protect Mode**. The window shown in Figure 10-22 on page 765 opens.

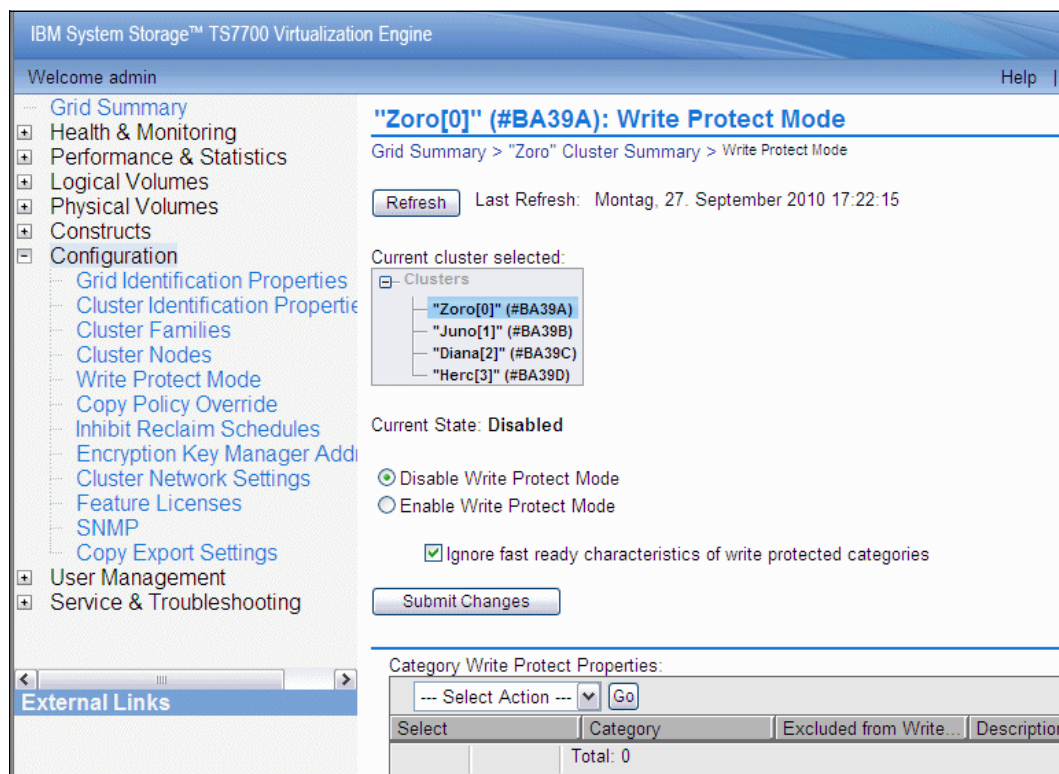


Figure 10-22 TS7700 Write Protect Mode window

3. Click **Enable Write Protect Mode** to set the cluster in Write Protect Mode.

Be sure to also leave the **Ignore fast ready characteristics if write protected categories** selected. This setting ensures that volumes in Production Fast Ready categories that are write-protected on the DR Cluster will be treated differently.

Normally, when a mount occurs to one of these volumes, the TS7700 assumes that the host starts writing at BOT and creates a stub. Also, when Expire Hold is enabled, the TS7700 will not allow any host access to these volumes until the hold period passes. Therefore, if the production host returns a volume to scratch "After" time zero, the DR host still believes within its catalog that the volume is private and the host will want to validate its contents. It cannot afford to allow the TS7700 to stub it or block access if the DR host attempts to mount it.

The **Ignore fast ready characteristics if write protected categories** option informs the DR Cluster that it should ignore these characteristics and treat the volume as a private volume. It will then surface the data versus a stub and will not prevent access because of expire hold states. However, it will still prevent write operations to these volumes.

Click **Submit Changes** to activate your selections.

4. Decide on which set of categories you want to use during DR testing on the DR hosts and confirm that no host system is using this set of categories, for example X'0030' to X'003F'. You define those categories to the DR host in a later step.

On the DR Cluster TS7700 management interface, define two Fast-Ready categories as described in 4.3.5, "Defining Fast Ready categories" on page 220. These two categories will be used on the DR host as scratch categories MEDIA1 and MEDIA2 (X'0031' and X'0032') and will be defined as excluded from Write-Protect mode.

5. In the DR Cluster MI, use the Write Protect Mode window (shown in Figure 10-22 on page 765) to define the entire set of categories to be excluded from Write-Protect Mode, including the Error and the Private Categories.

On the bottom of the window, click **Select Action** → **Add**, and then click **Go**. The next window opens (Figure 10-23).

Figure 10-23 Add Category window

Define the categories you have decided to use for DR testing, and make sure that **Excluded from Write Protect** is set to **Yes**. In our example, you would define volume categories X'0030' through X'003F', or, as a minimum X'0031' (MEDIA1), X'0032' (MEDIA2), X'003E' (ERROR), and X'003F' (PRIVATE).

6. On the DR Cluster, make sure that no copy is written to the Production Cluster that defines the CCP of **No Copy** in the Management Class definitions that are used by the DR host.
7. On the DR host, restore your DR system.
8. Change the DEVSUPxx member on the DR host to use the newly defined DR categories. DEVSUPxx controls installation-wide default tape device characteristics, for example:
 - MEDIA1 = 0031
 - MEDIA2 = 0032
 - ERROR = 003E
 - PRIVATE = 003F

Thus, the DR host is enabled to use these categories that have been excluded from Write-Protect Mode in Step 5.

9. On the DR host, define a new VOLSER range to your tape management system.

10. Insert that VOLSER range on the DR cluster and verify that Volume Insert Processing has assigned them to the correct Fast-ready categories.

11. On the DR host, vary online the virtual drives of the DR Cluster. Start DR testing.

10.7.2 Scenario 2 for a TS7700 two-cluster grid Pre-R1.7

In general, the most usual situation that we can face is a DR test in a DR site while real production occurs. In this situation, we have to assume that Grid links will not be broken, because the production site is working and it will need to continue copying cartridges to the DR site to be ready if a real disaster happens while we are running the test.

We assume the following information:

- ▶ That Grid links must not be broken
- ▶ That the production site will be running everyday jobs as usual
- ▶ That the DR site must not affect the production site in any way
- ▶ That the DR site is ready to start if a real disaster happens

Figure 10-24 shows the environment and the main tasks to perform in this DR situation.

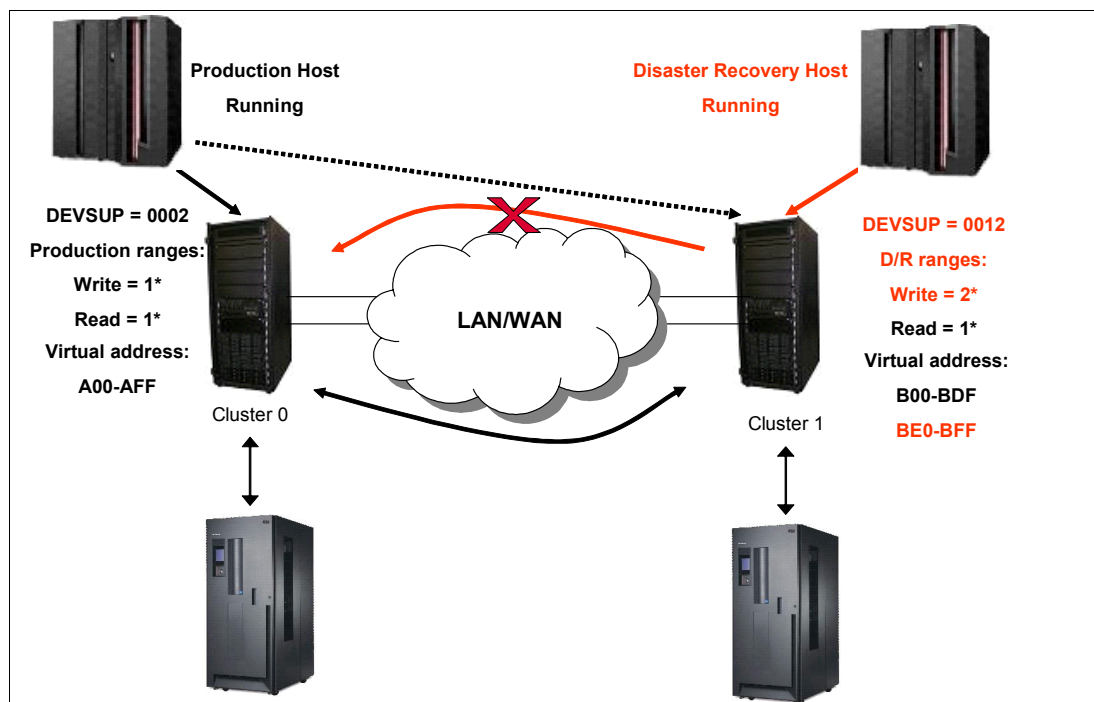


Figure 10-24 Disaster recovery environment: two clusters and links not broken

Note the following information about the figure:

- ▶ The production site can write and read its usual cartridges (in this case, 1*).
- ▶ The production site can write in any address in Cluster 0 or 1.
- ▶ The DR site can read production cartridges (1*), but cannot write on this range. We have to create a new range for this purpose (2*). This new range must not be accessible by the production site.
- ▶ We must ensure that no production tapes can be modified in any way by DR site systems.

- ▶ We must ensure that the production site does not rewrite tapes that will be needed during the DR test.
- ▶ We must not waste resources copying cartridges from the DR site to the production site.

Issues

Consider the following issues with TS7700 Pre-R1.7 environments:

- ▶ You should not run the HSKP process in production site unless you can run it without the EXPROC parameter in RMM. In z/OS V1R10, the new RMM parmlib commands PRTITION and OPENRULE provide for flexible and simple control of mixed system environments.

In z/OS V1R9 and later, you can specify additional EXPROC controls in the EDGHSKP SYSIN file to limit the return-to-scratch processing to specific subsets of volumes. So you could just use EXPROC on the DR system volumes on the DR system and use the PROD volumes on the PROD system. You can still continue to run regular batch processing and also run expiration on the DR system.
- ▶ With other TMSs, we need to stop the return-to-scratch process, if possible. If not, we must stop the whole daily process. To avoid problems with scratch shortage, you can add more logical volumes.
- ▶ If you run HSKP with the EXPROC (or daily processes in other TMSs) parameter in the production site, you cannot expire volumes that might be needed in the DR test. If you fail to do so, TS7700 Virtualization Engine sees these volumes as scratch and, with Fast Ready Category on, TS7700 Virtualization Engine presents the volume as a scratch volume, and you lose the data on the cartridge.
- ▶ Ensure that HSKP or short-on-scratch procedures are deactivated in the DR site.

Tasks before the DR test

Before performing the DR test of the TS7700 Virtualization Engine Grid, we must prepare the environment and do some tasks that will allow us to run the test without any problems or affecting our production systems.

Perform the following steps:

1. Plan and decide the scratch categories that are needed in the DR site (1*). See “Number of scratch volumes needed per day” on page 144.
2. Plan and decide the VOLSER ranges that will be used to write in the DR site (2*).
3. Modify the production site PARMLIB RMM member EDGRMMxx:
 - a. Include REJECT ANYUSE (2*) to avoid having the production system using or accepting the insertion of 2* cartridges.
 - b. If your Tape Management System is not RMM, disable CBRUXENT exit before inserting cartridges in the DR site.
4. Plan and decide the virtual address used in the DR site during the test (BE0-BFF).
5. Insert additional scratch virtual volumes to ensure that during the DR test that production cartridges can return to scratch but are not rewritten afterwards. This has to be done in the production site. For more information, see “Physical Tape Drives” on page 424.
6. Plan and define a new Management Class with copy policies on No Rewind (NR) using the MI at the DR site, for example, NOCOPY. For more information, see 8.2.7, “Constructs” on page 479.

Tasks during the DR test

After starting the DR system, but before the real DR test can start, you must change several things to be ready to use tapes from the DR site. Usually, the DR system is started using a “clone” image of the production system, so we need to alter some values and definitions to customize the image for the DR site.

The necessary steps are as follows:

1. Modify DEVSUPxx in SYS1.PARMLIB at the DR site and define the scratch category chosen for DR.
2. Use the command DEVSER QLIB,CATS at the DR site to change dynamically scratch categories. See “DEVSERV QUERY LIBRARY command” on page 282.
3. Modify the test PARMLIB RMM member EDGRMMxx at the DR site:
 - a. include REJECT OUTPUT (1*), to allow only read activity against production cartridges.
 - b. If you have another TMS product, ask your software provider how to use a similar function, if one exists.
4. Modify test PARMLIB RMM member EDGRMMxx at the DR site and delete REJECT ANYUSE(2*) to allow write and insertion activity of 2* cartridges.
5. Define a new SMS MC (NOCOPY) in SMS CDS at the DR site.
6. Modify the MC ACS routine at the DR site. All the writes must be directed to MC NOCOPY.
7. Restart the SMS configuration at the DR site.
8. Insert a new range (2*) of cartridges from the MI at the DR site. Ensure that all the cartridges are inserted in DR TS7700 Virtualization Engine so the ownership of this carts is TS7700 Virtualization Engine in DR site.
 - a. If you have RMM, your cartridges are defined automatically to TCDB and RMM.
 - b. If you have another TMS, check with the OEM software provider. In general, to add cartridges to other TMS, you need to stop them.
9. Do the next modification in DFSMSHsm at the DR site:
 - a. Mark all HSM ML2 cartridges as full by using the DELVOL MARKFULL HSM command.
 - b. Run HOLD HSM RECYCLE.
10. Again, make sure the following procedures do not run:
 - RMM housekeeping activity at the DR site
 - Short-on-scratch RMM procedures at the DR site

Tasks after the DR test

After the test is finished, we have a set of tapes in TS7700 Virtualization Engine that are used by the test activities. We need to decide what to do with these tapes. As the test ends, the RMM database and VOLCAT is destaged (because all the data was used in the test), but in the MI database, the tapes remain defined: One will be in master status and the others in scratch status.

Our guidance differs if the tapes are not needed anymore, or if the tapes will be used for future DR test activities.

In the first case (tapes are not needed anymore), perform the following steps:

1. Stop the RMM address space and subsystem and, using ISMF 2.3 (at the DR site), return to scratch all private cartridges.
2. After all of the cartridges are in the scratch status, use ISMF 2.3 again (at the DR site) to eject all the cartridges. Keep in mind that MI can only accept 1000 eject commands at one time. If you have to eject a high amount of cartridges, this process will be time consuming.

In the second case (tapes will be used in the future), run only step 1. Cartridges remain in the scratch status and ready for future use.

Note: Although cartridges in MI remain ready to use, you must ensure that the next time you create the test environment that these cartridges are defined to RMM and VOLCAT. Otherwise, you will not be able to use them.

10.7.3 Scenario 3 for a TS7700 two-cluster grid Pre R1.7

In other situations, we can choose to break Grid links, even if our production system is running during a DR test.

In this approach, we assume the following information:

- ▶ The Grid links are broken.
- ▶ The production site will be running everyday jobs as usual.
- ▶ The DR site cannot affect the production site.
- ▶ The DR site is ready for a real disaster.

Do not use logical drives in the DR site from the production site.

If you decide to “break” links during your DR test, you must review carefully your everyday work. For example, if you have 3 TB of cache and you write 4 TB of new data every day, you are a good candidate for a large amount of throttling, and probably during your batch window. To understand throttling, see “Host Write Throttle” on page 609.

After the test ends, we can have many virtual volumes in pending copy status. When TS7700 Virtualization Engine Grid links are restored, communication will be restarted, and the first task that TS7700 Virtualization Engine will do is make a copy of the volumes created during our “links broken” window. This can affect TS7700 Virtualization Engine performance.

If our DR test runs over several days, we can minimize the performance degradation by suspending copies using the GRIDCNTL Host Console Command. After our test is over, we can enable the copy again, during a low peak workload, to avoid or minimize performance degradation. See 8.3.3, “Host Console Request function” on page 544 for more information.

Figure 10-25 can help you understand the environment and the main tasks to perform in this DR situation.

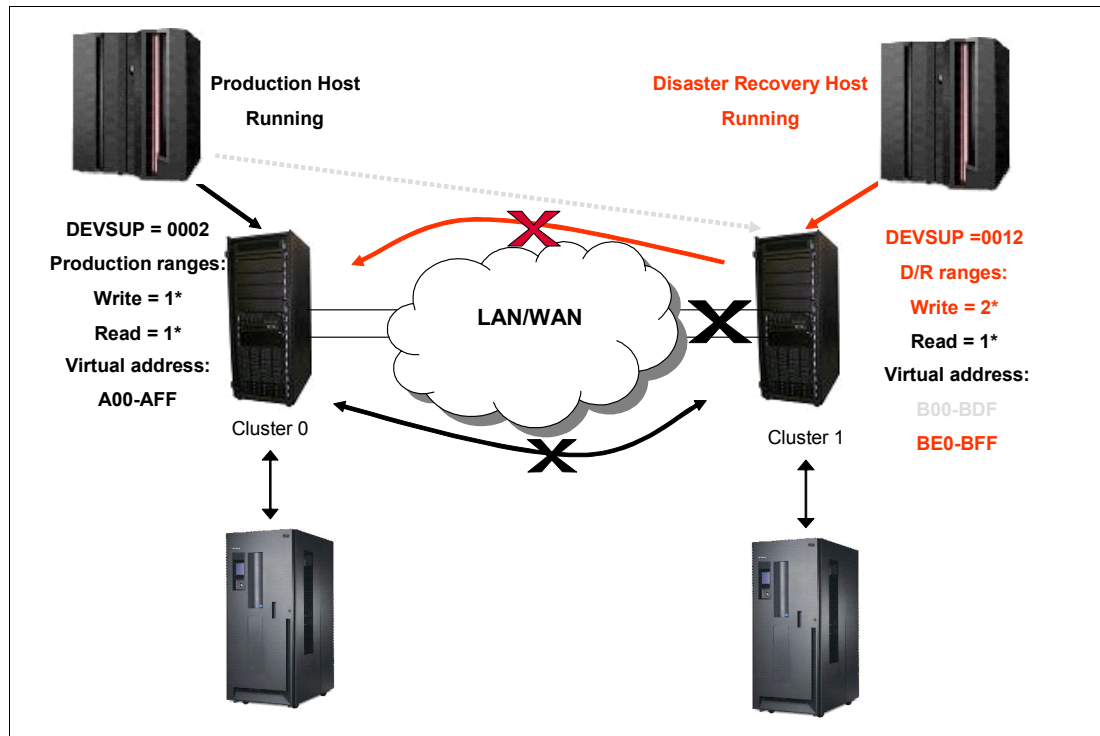


Figure 10-25 Disaster recovery environment: two clusters and links broken

Note the following information about the figure:

- ▶ The production site can write and read its usual cartridges (in this case, 1*).
- ▶ The production site only writes to virtual addresses associated with Cluster 0. The tapes remain, pending copy.
- ▶ The DR site can read production cartridges (1*) but cannot write on this range. We have to create a new one for this purpose (2*). This new range must not be accessible by the production site.
- ▶ We must ensure that no production tapes can be modified by the DR site systems.
- ▶ We must ensure that the production site does not rewrite tapes that will be needed during the DR test.
- ▶ We must not waste resources copying cartridges from the DR site to the production site.

Issues

Consider the following items:

- ▶ You can run the whole HSKP process at the production site. As communications are broken, the return-to-scratch process cannot be completed in the DR TS7700 Virtualization Engine, so your production tapes never return to scratch in the DR site.
- ▶ In this scenario, be sure that HSKP or short-on-scratch procedures are deactivated in the DR site.

Tasks before the DR test

Before you start the DR test for the TS7700 Virtualization Engine Grid, we must prepare the environment and perform several tasks so that we can run the test without any problems and without affecting our production site.

Perform the following steps:

1. Plan and decide on the scratch categories needed at the DR site (1*). See “Number of logical volumes” on page 145 for more information.
2. Plan and decide on the VOLSER ranges that will be used to write at the DR site (2*).
3. Plan and decide on the virtual address used at the DR site during the test (BE0-BFF).
4. Plan and define a new Management Class with copy policies on NR in the MI at the DR site, for example, NOCOPY. For more information, see 8.2.7, “Constructs” on page 479.

Tasks during the DR test

After starting the DR system, but before DR itself can start, you must change several things to be ready to use tapes from the DR site. Usually, the DR system is started using a “clone” image of the production system, so we need to alter some values and definitions to customize the DR site.

Perform the following steps:

1. Modify DEVSUPxx in SYS1.PARMLIB at the DR site and, when you define the scratch category, choose DR.
2. Use the DEVSER QLIB,CATS command at the DR site to change dynamically scratch categories. See “DEVSERV QUERY LIBRARY command” on page 282 for more information.
3. Modify test PARMLIB RMM member EDGRMMxx at the DR site
 - a. include REJECT OUTPUT (1*) to allow only read activity against production cartridges.
 - b. If you have another TMS product, ask your software provider for a similar function. We believe that there might not be similar functions in other TMSs.
4. Define a new SMS MC (NOCOPY) in SMS CDS at the DR site.
5. Modify the MC ACS routine at the DR site. All the writes must be directed to MC NOCOPY.
6. Restart the SMS configuration at the DR site.
7. Insert a new range of cartridges from the MI at the DR site. Ensure that all the cartridges are inserted in the DR TS7700 Virtualization Engine so the ownership of these cartridges are at the DR site.
 - a. If you have RMM, your cartridges will be defined automatically to TCDB and RMM.
 - b. If you have another TMS, check with the OEM software provider. In general, to add cartridges to other TMS, you need to stop them.
8. Now you can break the link connection between clusters. If you do this step before cartridge insertion, the insertion will fail.
9. If either of the following conditions apply, skip this step:
 - If you have the Autonomic Ownership Takeover function running.
 - if you usually write in the production site. See “Ownership Takeover Mode window” on page 523 for more information.

Otherwise, modify the ownership takeover mode in MI in the cluster at the Production site. Select **Write-only takeover mode**, which is needed only if you are working in balanced mode.

10. Modify ownership takeover mode in MI in the cluster at the DR site. Select **Read-only takeover mode**, because we only need to read production cartridges.
11. Do the next modification in DFSMSHsm at the DR site:
 - a. Mark all HSM ML2 cartridges as full by using the DELVOL MARKFULL, HSM command.
 - b. Run HOLD HSM RECYCLE.
12. Again, make sure the following procedures do not run:
 - RMM housekeeping activity at the DR site
 - Short on scratch RMM procedures at the DR site

Tasks after the DR test

After the test is finished, we have a set of tapes in the TS7700 Virtualization Engine that belong to test activities. We need to decide what to do with these tapes. As the test ends, the RMM database and VOLCAT will be destaged (as is all the data used in the test), but in the MI database, the tapes remain defined: One will be in master status and the others in scratch status.

Our guidance differs if the tapes are not needed anymore, or if the tapes will be used for future test activities.

In the first case (tapes are not needed anymore), perform the following steps:

1. Stop the RMM address space and subsystem and, using ISMF 2.3 (at the DR site), return to scratch all private cartridges.
2. After all of the cartridges are in the scratch status, again, using ISMF 2.3 (at the DR site), eject all the cartridges. Keep in mind that MI only accepts 1000 eject commands at one time. If you have to eject a high amount of cartridges, this process will be time consuming.

In the second case (tapes will be used in the future), run only step 1. Cartridges remain in the scratch status and are ready for use.

Note: Though cartridges in MI remain ready to use, you must ensure that the next time you create the test environment that these cartridges are defined to RMM and VOLCAT, otherwise you cannot use them.

10.7.4 Scenario 4 for a three-cluster grid

In this scenario, we have a three-cluster grid. In general, two of the cluster should be on a production site, and they can have high availability locally. From the DR point of view, the differences are not major compared to the two Grid procedures we described earlier.

We assume the following information:

- ▶ That Grid links are not broken.
- ▶ That the production site will be running everyday jobs as usual.
- ▶ That the DR site must not affect the production site at all.
- ▶ That the DR site is ready to start if a real disaster happens.

Figure 10-26 on page 774 helps us understand the environment and the main tasks to do in this DR situation.

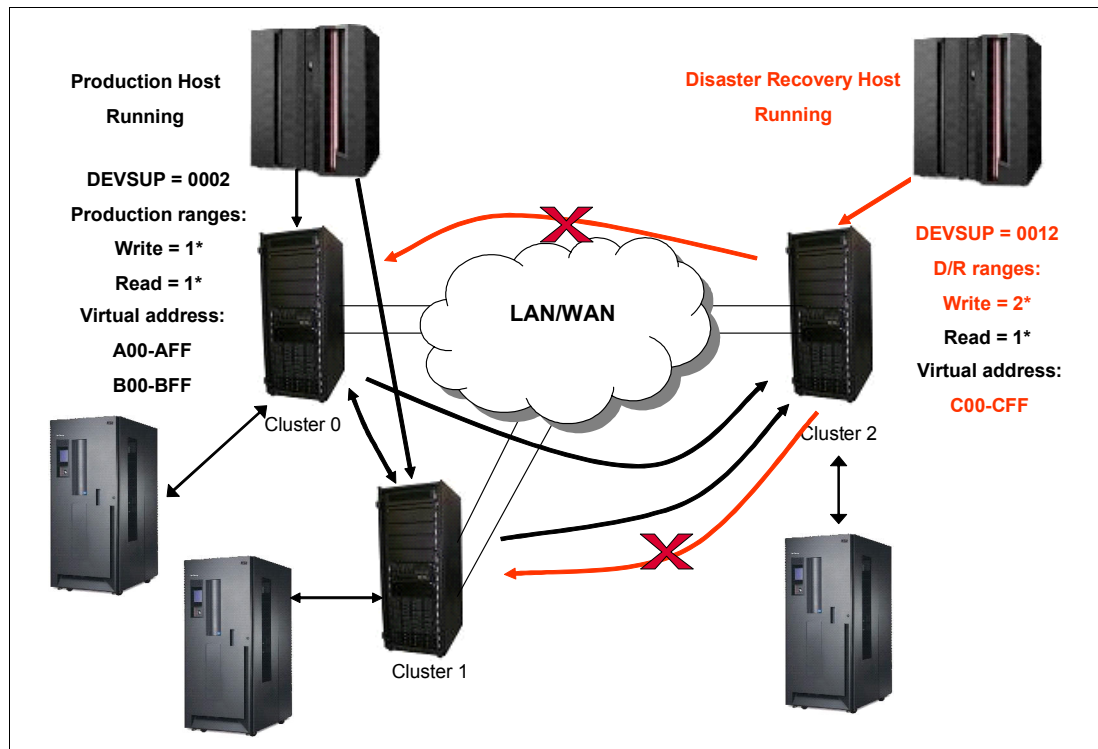


Figure 10-26 Disaster recovery environment: three clusters and links not broken

Note the following information about the figure:

- ▶ The production site can write and read its usual cartridges (in this case, 1*).
- ▶ The production site can write in any address in Cluster 0 or 1.
- ▶ The DR site can read production cartridges (1*) but cannot write on this range. We have to create a new range for this purpose (2*). This new range must not be accessible by the production site.
- ▶ We must ensure that no production tapes can be modified in any way by DR site systems.
- ▶ We must ensure that the production site does not rewrite tapes that will be needed during the DR test.
- ▶ We must not waste resources copying cartridges from the DR site to the production site.

Issues

Be aware of the following issues:

- ▶ You should not run the HSKP process at the production site, or you can run it without the EXPROC parameter in RMM. In other TMSs, we need to stop the return-to-scratch process, if possible. If not, we must stop the whole daily process. To avoid problems with scratch shortage, you can add more logical volumes.
- ▶ If you run HSKP with the EXPROC (or daily process in other TMSs) parameter in the production site, you cannot expire volumes that are needed in the DR test. If you fail to do so, and the TS7700 Virtualization Engine sees that volumes as scratch and with fast ready category on, the TS7700 Virtualization Engine presents the volume as a scratch volume, and you lose the data on the cartridge.
- ▶ Again, be sure that the HSKP or short-on-scratch procedures are deactivated at the DR site.

Tasks before the DR test

Before you do a DR test on the TS7700 Virtualization Engine Grid, we must prepare the environment and do some tasks that will allow us to run the test without complications or affecting our production site.

Perform the following steps:

1. Plan and decide upon the scratch categories needed at the DR site (1*).
2. Plan and decide upon the VOLSER ranges that will be used to write at the DR site (2*).
3. Modify the production site PARMLIB RMM member EDGRMMxx:
 - a. Include REJECT ANYUSE (2*) to prevent the production site from using or accepting the insertion of 2* cartridges.
 - b. In your Tape Management System, disable the CBRUXENT exit before inserting cartridges in the DR site.
4. Plan and decide upon the virtual address used at the DR site (C00-CFF).
5. Insert additional scratch virtual volumes at the production site to ensure that during the DR test that production cartridges can return to scratch but are not rewritten.
6. Plan and define a new Management Class with copy policies on NR in the MI at the DR site, for example, NOCOPY.
7. Remove the Fast-Ready attribute for the production scratch category at the DR site TS7700 Virtualization Engine. Do this for the duration of the DR test.

Tasks during the DR test

After starting the DR system, but before DR itself can start, you must change several things to be ready to use tapes from the DR site. Usually, the DR system is started using a “clone” image of the production system, so we need to alter some values and definitions to customize the DR site.

Perform the following steps:

1. Modify DEVSUPxx in SYS1.PARMLIB at the DR site and define the scratch category for DR.
2. Use the DEVSER QLIB,CATS command at the DR site to change dynamically scratch categories. See “DEVSERV QUERY LIBRARY command” on page 282 for more information.
3. Modify the test PARMLIB RMM member EDGRMMxx at the DR site:
 - a. Include REJECT OUTPUT (1*) to allow only read activity against production cartridges.
 - b. If you have another TMS product, ask your software provider for a similar function. We believe that there are not similar functions in other TMSs.
4. Modify test PARMLIB RMM member EDGRMMxx at the DR site and delete REJECT ANYUSE(2*) to allow write and insertion activity of 2* cartridges.
5. Define a new SMS MC (NOCOPY) in SMS CDS at the DR site.
6. Modify the MC ACS routine at the DR site. All the writes must be directed to MC NOCOPY.
7. Restart the SMS configuration at the DR site.

8. Insert a new range (2*) of cartridges from the MI at the DR site. Ensure that all the cartridges are inserted in the DR TS7700 Virtualization Engine so the ownership of these cartridges belong to the TS7700 Virtualization Engine at the DR site.
 - If you have RMM, your cartridges will be defined automatically to TCDB and RMM.
 - If you have another TMS, check with the OEM software provider. In general, to add cartridges to other TMSs, you need to stop them.
9. Modify the DFSMSHsm at the DR site:
 - a. Mark all HSM ML2 cartridges as full by using the DELVOL MARKFULL HSM command.
 - b. Run HOLD HSM RECYCLE.
10. Again, be sure the following procedures are not running:
 - a. RMM housekeeping activity at the DR site
 - b. Short on scratch RMM procedures at the DR site

Tasks after the DR test

After the test is finished, we have a set of tapes in the TS7700 Virtualization Engine that belong to test activities. We need to decide what to do with these tapes. As the test ends, the RMM database and VOLCAT will be destaged (and all the data used in the test), but the tapes remain defined in MI database: One will be in the master status, and the others in scratch status.

Our guidance differs if the tapes will not be needed anymore, or if the tapes will be used in the future for future test activities.

In the first case (tapes are not needed anymore), perform the following steps:

1. Stop the RMM address space and subsystem and, using ISMF 2.3 (at the DR site), return to scratch all private cartridges.
2. After all of the cartridges are in scratch status, use ISMF 2.3 (at the DR site) to eject all the cartridges. Keep in mind that MI can only accept 1000 eject commands at one time. If you have to eject a high amount of cartridges, this process will be time consuming

In the second case (tapes will be used in the future), run only step 1. Cartridges remain in scratch status and ready for use.

Note: Though cartridges in MI remain ready to use, you must ensure that the next time you create the test environment that these cartridges are defined to RMM and VOLCAT. Otherwise, you cannot use them.

10.8 A real disaster

We must now clarify what a real disaster means. For example, if we have a hardware issue that, for example, stops the TS7700 Virtualization Engine for 12 hours, is this a real disaster? It depends.

We are sure that for a bank, during the batch window, and without any other alternatives to bypass a 12 hour TS7700 Virtualization Engine outage, this can be similar to a real disaster.

However, if this bank has a three-cluster grid (two local and one remote), the same situation is less dire, because the batch window can continue accessing the second local TS7700 Virtualization Engine.

Because no set of fixed answers exists for all situations, you must carefully and clearly define which situations can be considered real disasters, and which actions to perform for all possible situations.

As we explain in 10.7, “Disaster recovery testing detailed procedures” on page 763, several differences exist between a DR test situation and a real disaster situation. In a real disaster situation, we do not have to do anything to be able to use the DR TS7700 Virtualization Engine, which makes our task easier.

However, this “easy-to-use” capability does not mean that we have all the cartridge data copied to the DR TS7700 Virtualization Engine. If our copy mode is RUN, we only need to consider “in-flight” tapes that are being created when the disaster happens. We must rerun all these jobs to recreate tapes for the DR site.

But, if our copy mode is DEFERRED, we must realize that we have tapes that are not copied yet. To know what tapes are not copied, we can go to MI in the DR TS7700 Virtualization Engine and find cartridges that are already in the copy queue. After we have this information we can, using our Tape Management System, discover what data sets are missing, and we can re-run the jobs to recreate these data sets at the DR site.

Figure 10-27 indicates a real disaster situation.

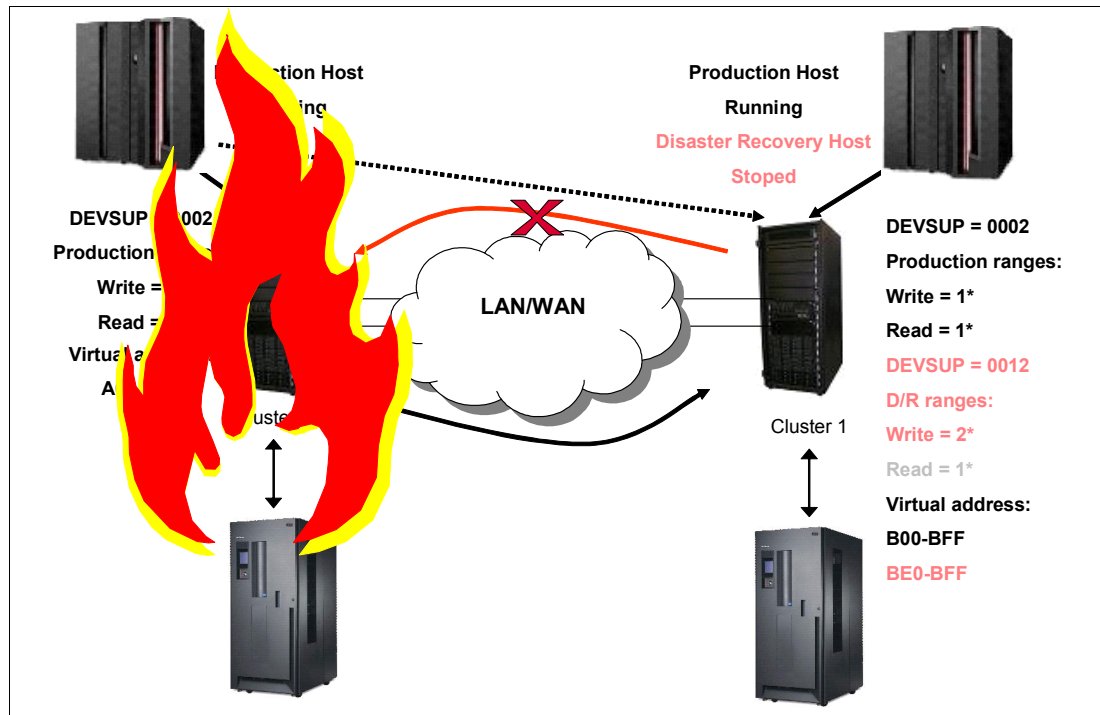


Figure 10-27 Real disaster situation

In a real disaster situation, we lose the whole primary site. This means that we need to start our production systems at the disaster recovery site. To do this, we need to have a copy of all your information not only on tape, but all DASD data copied to the DR site.

After you are able to start z/OS partitions, from the TS7700 Virtualization Engine perspective, you must be sure that your HCD configuration “sees” the DR TS7700 Virtualization Engine. Otherwise, you will not be able to put the TS7700 Virtualization Engine online.

We must change ownership takeover also. To perform that task, go to the MI interface and allow ownership takeover for read and write.

All the other changes you did in your DR test are not needed now. Production tapes ranges, scratch categories, SMS definitions, RMM inventory, and so on, are in a real configuration that are in DASD copied from the primary site. Now we use the real configuration.

We suggest some changes, because of the special situation that a disaster merits:

- Change your Management Class to obtain a dual copy of each tape created after the disaster.
- Depending on situation, you can consider using the Copy Export capability to move one of the copies outside the DR site.

After you are in a stable situation at the DR site, you need to start the tasks required to recover your primary site or create a new one. The old DR site is now the production site, so you need to create a new DR site, which is beyond the scope of this book.

Appendixes

In this part, we offer management and operation information for your TS7700.

This part contains the following appendixes:

- ▶ Appendix A, “Feature codes” on page 781
- ▶ Appendix B, “IBM Virtualization Engine TS7700 implementation steps” on page 797
- ▶ Appendix C, “TS3500 and TS7700 checklists” on page 803
- ▶ Appendix D, “JES3 examples and information” on page 821
- ▶ Appendix E, “IBM Virtualization Engine TS7700 statistical record format” on page 839
- ▶ Appendix F, “Sample JCLs” on page 889
- ▶ Appendix H, “DEVSERV QLIB command” on page 923
- ▶ Appendix G, “Library Manager volume categories” on page 913



A

Feature codes

In this appendix, we list all feature codes that are related to the installation of IBM Virtualization Engine TS7700. We provide several quick-reference lists, and then a more detailed explanation.

Note: This appendix provides a general description of feature codes and to where they apply. When planning for an upgrade, see the customer information center:

<http://www.ibm.com/support/docview.wss?rs=1166&uid=ssg1S7001552>

Feature code (FC) lists

This section lists feature code description for the TS7700 Virtualization Engine according to component, machine type and model.

3952 F05 features

The features are as follows:

- ▶ FC1903, Dual AC power
- ▶ FC1904, Redundant AC power
- ▶ FC2730, TS3000 System Console
- ▶ FC5626, Plant install 3957-VEA
- ▶ FC5628, Plant install 3957-V06
- ▶ FC5635, Plant install 3956-CS8
- ▶ FC5638, Plant install 3956-CC6
- ▶ FC5640, Plant install 3956-CC8
- ▶ FC5641, Plant install 3956-CX7
- ▶ FC5642, Field install 3956-CX7
- ▶ FC5646, Plant install 3956-XS7
- ▶ FC5647, Field install 3956-XS7
- ▶ FC5648, Plant install 3956-CX6
- ▶ FC5649, Field install 3956-CX6
- ▶ FC5759, Integrated control path
- ▶ FC7312, TS7740 Base frame
- ▶ FC7322, TS7720 Base frame
- ▶ FC7323, TS7720 Storage expansion frame
- ▶ FC9323 Expansion frame attachment
- ▶ FC9954, NEMA L6-30 power cord
- ▶ FC9955, RS 9750 DP power cord
- ▶ FC9956, IEC 309 power cord
- ▶ FC9957, PDL 4.3 power cord
- ▶ FC9958, Korean 4.3-m power cord
- ▶ FC9959, Unterminated power cord
- ▶ FC9966, Unterminated power cord (China)

3494 B10 and B20 VTS, 3957-V06, and 3957-VEA Server features

The 3494 B10 and B20 VTS uses FC3488, 4 Gb Fibre Channel Switch. The 3957-V06 uses the following features:

- ▶ FC0201, 9-micron LC/LC 31-meter
- ▶ FC0202, 9-micron LC/SC 31-meter
- ▶ FC0203, 50-micron LC/LC 31-meter
- ▶ FC0204, 50-micron LC/SC 31-meter
- ▶ FC0205, 62.5-micron LC/LC 31-meter
- ▶ FC0206, 62.5-micron LC/SC 31-meter
- ▶ FC0521, Functional enhancement field
- ▶ FC1030, 1-Gb grid copper connection
- ▶ FC1031, 1-Gb optical SW connection
- ▶ FC1032, 1-Gb grid dual port copper connection
- ▶ FC1033, 1-Gb grid dual port optical SW connection
- ▶ FC2714, Console expansion
- ▶ FC2715, Console attachment

- ▶ FC3441, FICON short-wavelength attachment
- ▶ FC3442, FICON long-wavelength attachment
- ▶ FC3443, FICON 10-km long-wavelength attachment
- ▶ FC3461, 8 GB Memory upgrade, field
- ▶ FC4015, Grid enablement
- ▶ FC4016, Remove Cluster from Grid
- ▶ FC4017, Cluster cleanup
- ▶ FC5240, Dual port Fibre Channel host bus adapter
- ▶ FC5267, 1 TB cache enablement
- ▶ FC5268, 100 MBps increment
- ▶ FC9000, Mainframe attachment
- ▶ FC9218, Attach to 3494 LM
- ▶ FC9219, Attach to TS3500
- ▶ FC9350, Plant install TS7700 Server in 3952 F05
- ▶ FC9351, Field merge 3957-V06 in 3952 F05
- ▶ FC9461, 8 GB Memory upgrade, plant
- ▶ FC9700, No factory cables
- ▶ FC9900, Encryption configuration

The 3957-VEA Server features are as follows:

- ▶ FC0201, 9-micron LC/LC 31-meter
- ▶ FC0202, 9-micron LC/SC 31-meter
- ▶ FC0203, 50-micron LC/LC 31-meter
- ▶ FC0204, 50-micron LC/SC 31-meter
- ▶ FC0205, 62.5-micron LC/LC 31-meter
- ▶ FC0206, 62.5-micron LC/SC 31-meter
- ▶ FC0521, Functional enhancement field
- ▶ FC1032, 1-Gb grid dual port copper connection
- ▶ FC1033, 1-Gb grid dual port optical SW connection
- ▶ FC2714, Console expansion
- ▶ FC2715, Console attachment
- ▶ FC3441, FICON short-wavelength attachment
- ▶ FC3442, FICON long-wavelength attachment
- ▶ FC3443, FICON 10-km long-wavelength attachment
- ▶ FC3461, 8 GB Memory upgrade, field
- ▶ FC4015, Grid enablement
- ▶ FC4016, Remove Cluster from Grid
- ▶ FC4017, Cluster cleanup
- ▶ FC5240, Dual port Fibre Channel host bus adapter
- ▶ FC5268, 100 MBps increment
- ▶ FC9000, Mainframe attachment
- ▶ FC9350, Plant install TS7700 Server in 3952 F05
- ▶ FC9268, Plant install 100 MBps throughput
- ▶ FC9461, 8 GB Memory upgrade, plant
- ▶ FC9700, No factory cables

3956-CC6, 3956-CC7, 3956-CC8, and 3956-CS8 Cache Controller features

Feature codes that are available for the TS7700 3956 cache controller are listed by model.

3956-CC6 Cache Controller features

The features are as follows

- ▶ FC7120, 1.7 TB fibre storage
- ▶ FC9230, Attach to 3957-V06
- ▶ FC9352, Plant install a TS7700 Cache Controller in a 3952 F05

3956-CC7 Cache Controller features

The features are as follows

- ▶ FC7403, Enable first expansion drawer
- ▶ FC7404, Enable fourth expansion drawer

3956-CC8 Cache Controller features

The features are as follows

- ▶ FC7123, 9.6 TB Fibre storage
- ▶ FC7403, Enable first expansion drawer
- ▶ FC9352, Plant install a TS7700 Cache Controller in a 3952 F05

3956-CS8 Cache Controller features

The features are as follows

- ▶ FC7114, 32TB SATA storage
- ▶ FC9352, Plant install a TS7700 Cache Controller in a 3952 F05

TS7700 Cache Drawer features 3956-CX6, 3956-CX7, and 3956-XS7

Feature codes available for the Cache Drawer are listed by model in this section.

3956-CX6 Cache Drawer features

The features are as follows

- ▶ FC6000, Intraframe fibre cable to 3956-CC6
- ▶ FC7120, 1.7 TB fibre storage
- ▶ FC9354, Plant install a TS7700 Cache Drawer in a 3952 F05
- ▶ FC9355, Field merge a TS7700 Cache Drawer in a 3952 F05

3956-CX7 Cache Drawer features

The features are as follows

- ▶ FC7123, 9.6 TB Fibre storage
- ▶ FC9354, Plant install a TS7700 Cache Drawer in a 3952 F05
- ▶ FC9355, Field merge a TS7700 Cache Drawer in a 3952 F05

3956-XS7 Cache Drawer features

The features are as follows

- ▶ FC7114, 32TB SATA storage
- ▶ FC9354, Plant install a TS7700 Cache Drawer in a 3952 F05
- ▶ FC9355, Field merge a TS7700 Cache Drawer in a 3952 F05

TS7700 Virtualization Engine feature code descriptions

In this section, we describe feature codes for the TS7700 Virtualization Engine. For more details, see the information center:

<http://www.ibm.com/support/docview.wss?rs=1166&uid=ssg1S7001552>

Feature codes are as follows:

► FC0201, 9-micron LC/LC 31-meter

FC0201, 9-micron LC/LC 31-meter, provides a 31-meter (101 ft.), 9-micron, single mode FICON cable with LC Duplex connectors on both ends for connection between the long wavelength FICON channel feature and a host system or a fibre component long wavelength channel with a 9-micron fibre cable LC Duplex adapter.

► FC0202, 9-micron LC/SC 31-meter

FC0202, 9-micron LC/SC 31-meter, provides a 31-meter, (101 ft.), 9-micron, single mode FICON cable with and LC Duplex connector on one end for connection to the long wavelength FICON channel feature, and an SC Duplex connector on the other end for connection to a host system or fibre component long wavelength channel with a 9-micron fibre cable SC Duplex adapter.

► FC0203, 50-micron LC/LC 31-meter

FC0203, 50-micron LC/LC 31-meter, provides a 31-meter (101 ft.), 50-micron, single mode FICON cable with LC Duplex connectors on both ends for connection between the long wavelength FICON channel feature and a host system or a fibre component long wavelength channel with a 50-micron fibre cable LC Duplex adapter.

► FC0204, 50-micron LC/SC 31-meter

FC0204, 50-micron LC/SC 31-meter, provides a 31-meter, (101 ft.), 50-micron, single mode FICON cable with and LC Duplex connector on one end for connection to the long wavelength FICON channel feature, and an SC Duplex connector on the other end for connection to a host system or fibre component long wavelength channel with a 50-micron fibre cable SC Duplex adapter.

► FC0205, 62.5-micron LC/LC 31-meter

FC0205, 62.5-micron LC/LC 31-meter, provides a 31-meter (101 ft.), 62.5-micron, single mode FICON cable with LC Duplex connectors on both ends for connection between the long wavelength FICON channel feature and a host system or a fibre component long wavelength channel with a 62.5-micron fibre cable LC Duplex adapter.

► FC0206, 62.5-micron LC/SC 31-meter

FC0206, 62.5-micron LC/SC 31-meter, provides a 31-meter, (101 ft.), 62.5-micron, single mode FICON cable with and LC Duplex connector on one end for connection to the long wavelength FICON channel feature, and an SC Duplex connector on the other end for connection to a host system or fibre component long wavelength channel with a 62.5-micron fibre cable SC Duplex adapter.

► FC0521, Functional enhancement field

FC0521, Functional enhancement field, provides a one-time update to the microcode of an installed TS7700 Virtualization Engine to provide enhanced functions contained in the latest level of functional microcode firmware support. Newer microcode levels may be required when adding new functions.

Note: When this feature is installed on a TS7740 Virtualization Engine, it includes the minimum level of library manager microcode required to support the current TS7700 Virtualization Engine microcode level.

► FC1030, 1-Gb grid copper connection ¹

FC1030, 1-Gb grid copper connection, provides an Ethernet 1000BaseT adapter for grid communication between TS7740 Virtualization Engines. You must supply your own Ethernet cables when FC4015, Grid enablement, is installed. Cat-5 Ethernet cables are not allowed. This adapter has an RJ-45 connector and conforms to the IEEE 802.3ab 1000Base-T standard. The Ethernet cables used with the adapter must be Cat-5e or Cat-6, though Cat-6 is preferred. It supports distances of up to 100 meters using four pairs of Cat-5e or Cat-6 balanced copper cabling.

► FC1031, 1-Gb optical SW connection¹

FC1031, 1-Gb optical SW connection, provides a 1-Gb shortwave adapter for grid communication between TS7740 Virtualization Engines. It replaces RPQ #Q8B3409. This adapter has an LC Duplex connector for attaching a 50.0- or 62.5-micron, multi-mode fibre cable. This is a standard shortwave (850 nm) adapter that conforms to the IEEE 802.3z standards. It supports distances of 2 to 260 meters for 62.5-micron cable and 2 to 550 meters for 50.0-micron cable. Multi-mode fibre cables must be used with the adapter when FC4015, Grid enablement, is installed.

► FC1032, 1-Gb grid dual port copper connection

FC1032, 1-Gb grid dual port copper connection, provides a dual port, 1-Gbps 10/100/1000Base-TX adapter for grid communication between TS7700 Virtualization Engines with a single port enabled. This adapter has an RJ-45 connector for attaching Cat-5e or Cat-6 cables. This adapter conforms to the IEEE 802.3ab 1000Base-T standard. It supports distances up to 100 meters using four pairs of Cat-5e or Cat-6 balanced copper cabling. You must supply Cat-5e or Cat-6 cables for use with the adapter when FC4015, Grid enablement is installed. This feature is only supported by TS7700 Virtualization Engine microcode level 8.5.0.xx or later.

► FC1033, 1-Gb grid dual port optical SW connection

FC1033, 1-Gb grid dual port optical SW connection, provides a dual port, 1-Gbps Ethernet shortwave adapter for grid communication between TS7700 Virtualization Engines with a single port enabled. This adapter has an LC Duplex connector for attaching 50-micron or 62.5-micron, multimode fibre cable. This is a standard shortwave (850 nm) adapter that conforms to the IEEE 802.3z standards. It supports distances of 260 meters for 62.5-micron and 550 meters for 50.0-micron cable. Multi-mode fibre cables must be used with the adapter when FC4015, Grid enablement is installed. This feature is only supported by TS7700 Virtualization Engine microcode level 8.5.0.xx or later.

► FC1903, Dual AC power

FC1903, Dual AC power, provides a second power distribution unit (PDU) to allow connection to independent power sources. This feature supplies two power cords when ordered.

¹ This feature has been withdrawn.

- ▶ FC1904, Redundant AC power
FC1904, Redundant AC power provides power switching to allow a TS7720 Virtualization Engine or a TS7740 Virtualization Engine redundant attachment to the service network.
- ▶ FC2714, Console expansion²
FC2714, Console expansion, provides an Ethernet hub and product attachment cable for expanding the number of components that can attach to the TSSC that is provided by FC2713, System console for service
- ▶ FC2715, Console attachment²
FC2715, Console attachment, provides a cable for attaching a TS7740 Server to the Ethernet hub provided by FC2714, Console expansion, FC2718, FC2720, TS3000 System Console, or by FC2721, Rack mountable TS3000 System Console. A single TSSC facility can support a maximum of forty instances of this feature.
- ▶ FC2719, Console upgrade
FC2719, Console upgrade, provides a second Ethernet network interface card to allow redundant attachment to the service network. This feature provides a memory upgrade to 2 GB total RAM and a second Ethernet card to allow redundant connections into the service network. This feature only applies to consoles shipped with FC2718, FC2720, TS3000 System Console, or by FC2721, Rack mountable TS3000 System Console. It is required for use with any TS7700 Virtualization Engine using features FC2714, Console expansion, FC2715, Console attachment, or FC2720, TS3000 System Console. Support for this feature on the 3957-V06 was withdrawn on December 5, 2008. To use this feature, order it against the 3953 F05 Tape Frame.
- ▶ FC2720, TS3000 System Console
FC2720, TS3000 System Console, provides the IBM TotalStorage Master Console, an Ethernet hub, and a cable and connectors to enable remote enhanced service connection of a TS7740 Virtualization Engine to an IBM-supplied modem. The Ethernet hub provides 14 additional connections for cables supplied with FC2714, Console expansion, or FC2715, Console attachment. You should specify FC2720, TS3000 System Console on the first unit in an installation connected to a master console facility. Support for this feature on the 3957-V06 was withdrawn on December 5, 2008. To use this feature, order it against the 3953 F05 Tape Frame.
- ▶ FC2721, Rack mountable TS3000 System Console
FC2721, Rack mountable TS3000 System Console, provides the rack mountable TS3000 System Console, an Ethernet switch for the Master Console and a cable and connector for connection of one subsystem.
- ▶ FC2730, TS3000 System Console
FC2730, TS3000 System Console, provides a rack mount version of the TS3000 System Console for installation in a TS7720 Virtualization Engine or TS7740 Virtualization Engine 3952 F05 frame or a customer provided rack. FC2730, TS3000 System Console provides a 1U server, keyboard, display, mouse, and Ethernet switch.
- ▶ FC3441, FICON short-wavelength attachment
FC3441, FICON short-wavelength attachment, provides one short-wavelength 4-Gbps FICON adapter with an LC duplex connector for attachment to a FICON host system short-wavelength channel using a 50- or 62.5-micron multi-mode fibre cable. At 4 Gbps, the maximum fibre cable length allowed by 50-micron cable is 150 m (492 ft.); by 62.5-micron cable is 55 m (180 ft.). Each 4-Gbps FICON attachment can support up to 256 logical channels.

² If you have FC2713, System console for service, you must order FC2720, TS3000 System Console with the TS7740 Server (3957-V06) or FC2721 with 3953 F05 Tape Frame.

- ▶ **FC3442, FICON long-wavelength attachment**
FC3442, FICON long-wavelength attachment, provides one long-wavelength 4-Gbps FICON adapter with an LC duplex connector for attachment of a TS7700 Virtualization Engine to a FICON host system long-wavelength channel using a 9-micron single-mode fibre cable. The maximum fibre cable length is 4 km (2.48 mi.). Each 4-Gbps FICON attachment can support up to 256 logical channels.
- ▶ **FC3443, FICON 10-km long-wavelength attachment**
FC3443, FICON 10-km long-wavelength attachment, provides one long-wavelength 4-Gbps FICON adapter with an LC duplex connector for attachment to a FICON host system long-wavelength channel using a 9-micron single-mode fibre cable. The maximum fibre cable length is 10 km (6.21 mi.). Each 4-Gbps FICON attachment can support up to 256 logical channels.
- ▶ **FC3461, 8 GB Memory upgrade, field**
FC3461, 8 GB Memory upgrade, field, delivers a field-installable TS7700 Server memory upgrade to 16 GB RAM.
- ▶ **FC3488, 4 Gb Fibre Channel Switch**
FC3488, 4 Gb Fibre Channel Switch, provides a 4 Gb Fibre Channel switch to attach to 3592 Tape Drives in the 3584 Tape Library. The 4 Gb Fibre Channel switch has dual power connection for optional attachment to separate power supplies. The Fibre Channel switch types must be the same (2 Gb or 4 Gb) and cannot be intermixed within FC4888, Switch mount kit.
- ▶ **FC4015, Grid enablement**
FC4015, Grid enablement, provides a key to enable the communication function that allows a TS7700 Virtualization Engine to communicate with other TS7700 Virtualization Engines in a grid. Each TS7700 Virtualization Engine must have this feature to be able to participate in a grid configuration.³
- ▶ **FC4016, Remove Cluster from Grid**
FC4016, Remove Cluster from Grid, delivers instructions for a one-time process to remove a TS7700 Cluster from a TS7700 Grid. You must order this feature before you can perform a cluster cleanup (FC4017, Cluster cleanup) on any TS7700 Cluster configured to participate in a TS7700 Grid. If a TS7700 Cluster is removed from a TS7700 Grid, cleaned up using FC4017, Cluster cleanup, and then joined to a new TS7700 Grid, another instance of FC4016, Remove Cluster from Grid, is required to remove the cluster from the new Grid.
- ▶ **FC4017, Cluster cleanup**
FC4017, Cluster cleanup, facilitates a one-time cluster cleanup to clean the database, delete logical volumes from the Tape Volume Cache, and remove configuration data for host connections from a TS7700 Cluster. If the cluster is a member of a TS7700 Grid, the cluster must first be removed from the Grid using FC4016, Remove Cluster from Grid. If another cleanup is required on this TS7700 Cluster, another instance of FC4017, Cluster cleanup, is required.

³ If all clusters in a grid configuration already have FC4015, Grid enablement installed, contact your IBM System Service Representative to properly set up, connect, and configure the grid environment.

► FC4888, Switch mount kit

FC4888, Switch mount kit, provides the mounting hardware for up to two Fibre Channel switches on the 3953 F05 Tape Frame that contains a 3592 J70 Controller or is attached to a VTS. It includes the required mounting hardware, bifurcated power cables, and instructions for installing up to two Fibre Channel switches in the 3953 F05 Tape Frame. The Fibre Channel switch types must be the same (2-Gb or 4-Gb) and cannot be intermixed within FC4888, Switch mount kit.

► FC4897, Reinstall 4Gb FC Switch

FC4897, Reinstall 4Gb FC Switch, provides a customer-owned 4-Gb Fibre Channel switch to attach to 3592 Tape Drives in the 3584 Tape Library. The 4-Gb Fibre Channel switch has dual power connection for optional attachment to separate power supplies. The Fibre Channel switch types must be the same (2-Gb or 4-Gb) and cannot be intermixed within FC4888, Switch mount kit.

► FC5240, Dual port Fibre Channel host bus adapter

FC5240, Dual port Fibre Channel host bus adapter, installs two Fibre Channel interface cards in the TS7700 Server (3957-V06 or 3957-VEA) and provides two 31 Meter, 50 μ Fibre Channel cables to connect the TS7700 Server to the Fibre Channel switch.

When ordered for the TS7720 Server (3957-VEA), this feature connects the disk arrays in the 3952 Storage Expansion Frame with the 3957-VEA.

For situations where customer-supplied Fibre Channel cables connect the TS7700 Virtualization Engine to the Fibre Channel switches, and a 4 Gbps rate is expected, the following maximum length restrictions apply:

- 50 micron, 500 MHz multimode Fibre Channel cables cannot exceed 150 meters.
- 62.5 micron, 200 MHz Fibre Channel cables cannot exceed 70 meters.

► FC5267, 1 TB cache enablement

FC5267, 1 TB cache enablement, delivers a key to enable a 1 TB increment of disk cache to store logical volumes. Enabling a cache increment does not guarantee that amount of additional cache capacity. The amount of additional cache capacity provided is limited by the capacity of the underlying physical cache installed.⁴

► FC5268, 100 MBps increment

FC5268, 100 MBps increment, delivers a key to enable an additional 100-MBps increment of potential Peak Data Throughput. Enabling a data throughput increment does not guarantee that the overall TS7700 Virtualization Engine performs at that data throughput level. However, installation of the maximum allowed feature codes results in unrestricted peak data throughput capabilities.

► FC5626, Plant install 3957-VEA

FC5626, Plant install 3957-VEA, allows the factory installation of a TS7720 Server (3957-VEA) into a new 3952 Tape Frame. This feature occurs on the 3952 Tape Frame order.

You must also order FC9350, Plant install TS7700 Server in 3952 F05 against the 3957-VEA when ordering this feature.

► FC5628, Plant install 3957-V06

FC5628, Plant install 3957-V06, allows plant installation of a TS7740 3957-V06 into a new 3952 Tape Frame. This feature occurs on the 3952 Tape Frame order.

You must also order FC9350, Plant install TS7700 Server in 3952 F05, against the TS7740 3957-V06 when ordering this feature.

⁴ This feature is available only with the TS7740 Virtualization Engine.

► FC5635, Plant install 3956-CS8

FC5635, Plant install 3956-CS8, allows the factory installation of a TS7720 Cache Controller (3956-CS8) into a new 3952 Tape Frame. This feature occurs on the 3952 Tape Frame order.

You must also order FC9352, Plant install a TS7700 Cache Controller in a 3952 F05 against the 3956-CS8 when ordering this feature.

► FC5638, Plant install 3956-CC6

FC5638, Plant install 3956-CC6, allows factory installation of a TS7740 Cache Controller (3956-CC6) into a new 3952 Tape Frame. This feature occurs on the 3952 Tape Frame order.

You must also order FC9352, Plant install a TS7700 Cache Controller in a 3952 F05, against the 3956-CC6 when ordering this feature.

► FC5640, Plant install 3956-CC8

FC5640, Plant install 3956-CC8, allows the factory installation of a TS7740 Cache Controller (3956-CC8) into a new 3952 Tape Frame. This feature occurs on the 3952 Tape Frame order.

You must also order FC9352, Plant install a TS7700 Cache Controller in a 3952 F05 against the 3956-CC8 when ordering this feature.

► FC5641, Plant install 3956-CX7

FC5641, Plant install 3956-CX7, allows the factory installation of a TS7740 Cache Drawer (3956-CX7) into a new 3952 Tape Frame. This feature occurs on the 3952 Tape Frame order.

You must also order FC9354, Plant install a TS7700 Cache Drawer in a 3952 F05 against the 3956-CX7 when ordering this feature.

► FC5642, Field install 3956-CX7

FC5642, Field install 3956-CX7, allows the field installation of a TS7740 Cache Drawer (3956-CX7) into an existing customer 3952 Tape Frame. This feature occurs on the 3952 Tape Frame order.

You must also order FC7312, TS7740 Base frame against the 3952 Tape Frame and FC9355, Field merge a TS7700 Cache Drawer in a 3952 F05 against the 3956-CX7 when ordering this feature.

► FC5646, Plant install 3956-XS7

FC5646, Plant install 3956-XS7, allows the factory installation of a TS7720 Cache Drawer (3956-XS7) into a new 3952 Tape Frame. This feature occurs on the 3952 Tape Frame order⁵.

► FC5647, Field install 3956-XS7

FC5647, Field install 3956-XS7, allows the field installation of a TS7720 Cache Drawer (3956-XS7) into an existing 3952 Tape Frame. This feature occurs on the 3952 Tape Frame order⁵.

⁵ You must also order FC7322, TS7720 Base frame, or FC7323, TS7720 Storage expansion frame with the 3952 Tape Frame and FC9355, Field merge a TS7700 Cache drawer in a 3952 F05 with the 3956-XS7 when ordering this feature

► FC5648, Plant install 3956-CX6⁶

FC5648, Plant install 3956-CX6, allows plant installation of a TS7740 Cache Drawer (3956-CX6) into a new 3952 Tape Frame. This feature occurs on the 3952 Tape Frame order.

You must also order FC9354, Plant install a TS7700 Cache Drawer in a 3952 F05, against the 3956-CX6 when ordering this feature.

► FC5649, Field install 3956-CX6

FC5649, Field install 3956-CX6, allows the field installation of a TS7740 Cache Drawer (3956-CX6) into an existing customer 3952 Tape Frame. This feature occurs on the 3952 Tape Frame order.

You must also order FC9355, Field merge a TS7700 Cache Drawer in a 3952 F05 against the 3956-CX6 when ordering this feature.

► FC5759, Integrated control path

FC5759, Integrated control path, provides the router cables necessary to create the control path between the TS7740 Virtualization Engine and the Library Manager.

► FC6000, Intraframe fibre cable to 3956-CC6

FC6000, Intraframe fibre cable to 3956-CC6, provides one 1m fibre cable and is also used to connect TS7740 Cache Drawers.

► FC6003, Intraframe fibre cable to 3957-V06

FC6003, Intraframe fibre cable to 3957-V06, provides one 2m fibre cable.

► FC7114, 32TB SATA storage

FC7114, 32TB SATA storage, installs a complete set of 16 2000GB, 7200rpm, SATA disk drive modules in the TS7720 Cache Controller (3956-CS8), providing 32TB of unformatted disk storage capacity. Usable capacity is approximately 22TB when used with RAID 6 formatting.

► FC7120, 1.7 TB fibre storage

FC7120, 1.7 TB fibre storage, installs a complete set of 16 Fibre Channel-capable disk drive modules in the TS7740 Cache Controller (3956-CC6) or, TS7740 Cache Drawer (3956-CX6) providing 1.7 TB of usable storage capacity.

► FC7123, 9.6 TB Fibre storage

FC7123, 9.6 TB Fibre storage, installs a complete set of 16 600-GB, 15 thousand RPM, Fibre Channel-capable disk drive modules in the TS7740 Cache Controller (3956-CC8), providing 9TB of unformatted storage capacity. Usable capacity is 6.86 TB when used with RAID 5 formatting.

► FC7312, TS7740 Base frame

FC7312, TS7740 Base frame, identifies this 3952 Tape Frame as the base unit for the TS7740 Virtualization Engine.

► FC7322, TS7720 Base frame

FC7322, TS7720 Base frame identifies a 3952 Tape Frame as the base unit for a TS7720 Virtualization Engine.

► FC7323, TS7720 Storage expansion frame

FC7323, TS7720 Storage expansion frame, identifies a 3952 Tape Frame as the disk storage expansion unit for a TS7720 Virtualization Engine.

⁶ This feature has been withdrawn.

- ▶ FC7403, Enable first expansion drawer

FC7403, Enable first expansion drawer, is required on a TS7740 Cache Controller (3956-CC7) to allow attachment of the first TS7740 Cache Drawer (3956-CX7) in its storage string. If no cache drawers are attached to the cache controller, this feature is not required.

- ▶ FC7404, Enable fourth expansion drawer

FC7404, Enable fourth expansion drawer, is required on a TS7740 Cache Controller (3956-CC7) to allow attachment of the fourth TS7740 Cache Drawer (3956-CX7) in its storage string. If fewer than three cache drawers are attached to the cache controller, this feature is not required.

- ▶ FC9000, Mainframe attachment

FC9000, Mainframe attachment, is a specify feature that indicates attachment to one of the following host platforms:

- z/OS 1.4 and later
- z/VM, 5.4.0 and later
- z/VSE, 4.1 and later
- zTPF, V1.1 and later
- TPF, 4.1 and later

- ▶ FC9013, Attach to TS7700⁷

FC9013, Attach to TS7700, is a specify feature that indicates attachment to a TS7700 Virtualization Engine. This feature cannot be used with TS7740 Virtualization Engines operating microcode levels of 8.5.0.x or later.

- ▶ FC9217, Attach to 3953 LM⁷

FC9217, Attach to 3953 LM, is a specify feature that indicates that the TS7740 Server (3957-V06) attaches to a 3953 L05 Library Manager. This feature cannot be used with TS7740 Virtualization Engines operating microcode code levels of 8.5.0.x or later.

- ▶ FC9218, Attach to 3494 LM

FC9218, Attach to 3494 LM, is a specify feature that indicates that the TS7740 Server (3957-V06) attaches to a Library Manager in a 3494 Tape Library.

- ▶ FC9219, Attach to TS3500

FC9219, Attach to TS3500 enables the TS7740 Server (3957-V06) to attach directly to the TS3500 Tape Library, without the use of the 3953 L05 Library Manager.

The 4-Gb Fibre Channel features are required in the 3584 Tape Library frames containing the back-end drives.

- ▶ FC9230, Attach to 3957-V06

FC9230, Attach to 3957-V06, is a specify feature that indicates a TS7740 Cache Controller (3956-CC6) is attached to a TS7740 Server (3957-V06). This feature occurs on the 3956-CC6 order.

You must also order FC9352, Plant install a TS7700 Cache Controller in a 3952 F05, against the IBM Virtualization Engine TS7700 and FC5638, Plant install 3956-CC6, against the 3952 Tape Frame when ordering this feature.

⁷ This feature was withdrawn on December 5, 2008.

- ▶ FC9268, Plant install 100 MBps throughput

FC9268, Plant install 100 MBps throughput, delivers a key to enable the first 100-MBps increment of potential peak data throughput on the TS7720 Server (3957-VEA). Enabling a data throughput increment does not guarantee that the overall TS7720 Virtualization Engine performs at that data throughput level. However, installation of the maximum allowed number of features results in unrestricted peak data throughput capabilities.

This feature is restricted to new factory installation and represents the first 100-MBps increment.

- ▶ FC9323 Expansion frame attachment

FC9323 Expansion frame attachment indicates that a 3952 Tape Frame used with a TS7720 Virtualization Engine is a 3952 Base Frame and is attached to a 3952 Storage Expansion Frame.

- ▶ FC9350, Plant install TS7700 Server in 3952 F05

FC9350, Plant install TS7700 Server in 3952 F05, allows the plant installation of a TS7700 Server (3957-V06 or 3957-VEA) into a new 3952 Tape Frame. This feature occurs on the TS7700 Server order. You must also order FC5628, Plant install 3957-V06, against the 3952 Tape Frame when ordering this feature for the 3957-V06.

- ▶ FC9351, Field merge 3957-V06 in 3952 F05

FC9351, Field merge 3957-V06 in 3952 F05, allows the field merge of a TS7740 Server (3957-V06) into an existing installed 3952 Tape Frame. This feature occurs on the 3957-V06 order. You must also order FC5729, Field install 3957-V06, against the 3952 Tape Frame. when ordering this feature.

- ▶ FC9352, Plant install a TS7700 Cache Controller in a 3952 F05

FC9352, Plant install a TS7700 Cache Controller in a 3952 F05, allows the factory installation of a TS7700 Cache Controller (3956-CC8, or 3956-CS8) into a new 3952 Tape Frame.

- If installing a 3956-CC8, you must also order FC5640, Plant install 3956-CC8 with the 3952 Tape Frame.
- If installing a 3956-CS8, you must also order FC5635, Plant install 3956-CS8 with the 3952 Tape Frame.

- ▶ FC9353, Field merge in 3952 F05

FC9353, Field merge in 3952 F05, allows the field merge of a TS7720 Cache Controller into a customer installed 3952 Tape Frame.

This feature must be installed on the TS7720 Cache Controller and FC5637 must be ordered against the 3952 Tape Frame.

- ▶ FC9354, Plant install a TS7700 Cache Drawer in a 3952 F05

FC9354, Plant install a TS7700 Cache Drawer in a 3952 F05 allows the factory installation of a TS7700 Cache Drawer (3956-CX6, 3956-CX7, or 3956-XS7) into a new 3952 Tape Frame.

- If installing a 3956-CX6, you must also order FC5648, Plant install 3956-CX6 with the 3952 Tape Frame.
- If installing a 3956-CX7, you must also order FC5641, Plant install 3956-CX7 with the 3952 Tape Frame.
- If installing a 3956-XS7, you must also order FC5646, Plant install 3956-XS7 with the 3952 Tape Frame.

- ▶ FC9355, Field merge a TS7700 Cache Drawer in a 3952 F05

FC9355, Field merge a TS7700 Cache Drawer in a 3952 F05, is a specify code that allows the field merge of a TS7700 Cache Drawer (3956-CX6, 3956-CX7, or 3956-XS7) 3952 Tape Frame.

- If installing a 3956-CX6, you must also order FC5649, Field install 3956-CX6 with the 3952 Tape Frame.
- If installing a 3956-CX7, you must also order FC5642, Field install 3956-CX7 with the 3952 Tape Frame.
- If installing a 3956-XS7, you must also order FC5647, Field install 3956-XS7 with the 3952 Tape Frame.

- ▶ FC9461, 8 GB Memory upgrade, plant

FC9461, 8 GB Memory upgrade, plant, allows the factory installation of 16 GB RAM in the TS7700 Server.

- ▶ FC9700, No factory cables

FC9700, No factory cables, instructs the plant not to ship any FICON cables with the TS7700 Server (3957-VEA or 3957-V06).

- ▶ FC9900, Encryption configuration

Order FC9900, Encryption configuration when you use encryption with the TS7740 Virtualization Engine. This feature includes publication updates with information about how to enable and configure the library (virtual or real) to support encryption. This feature also provides an encryption key manager publication. Customer initiated procedures must be completed to enable and configure the TS7740 Virtualization Engine to support encryption with the TS1120 encryption capable tape drive. This feature occurs on the 3957-V06 order.

FC9900, Encryption configuration is only supported with encryption-capable TS1120 3592 E05 Tape Drives. FC9900, Encryption configuration, is not supported by 3592 J1A or 3590 Tape Drives.

- ▶ FC9954, NEMA L6-30 power cord

FC9954, NEMA L6-30 power cord, provides a National Electrical Manufacturers Association (NEMA) L6-30 non-watertight, 4.3-m (14-ft.) power cord, rated for 200 V ac to 208 V ac or 240 V ac and 24 A. This power cord is suggested for use in the USA, Canada, Latin America, and Japan. This feature occurs on the 3952 Tape Frame order.

- ▶ FC9955, RS 9750 DP power cord

FC9955, RS 9750 DP power cord, provides a Russellstoll 3750 DP, watertight, 4.3-m (14-ft.) power cord, rated for 200 V ac to 208 V ac or 240 V ac and 24 A. This power cord is suggested for use in the USA (highly recommended in Chicago, Illinois, to conform with local requirements), Canada, Latin America, and Japan. This feature occurs on the 3952 Tape Frame order.

- ▶ FC9956, IEC 309 power cord

FC9956, IEC 309 power cord, provides an International Electrotechnical Commission (IEC) 309 p+n+g, 32 A plug, 4.3-m (14-ft.) power cord, rated for 230 V ac and can safely carry 24 A in all countries. This power cord is suggested for use in Europe, the Middle East, and Africa. This feature occurs on the 3952 Tape Frame order.

► FC9957, PDL 4.3 power cord

FC9957, PDL 4.3 power cord, provides a PDL 4.3-m (14-ft.) power cord, rated for 230 V ac to 240 V ac and 24 A. This power cord is suggested for use in Australia and New Zealand. This feature occurs on the 3952 Tape Frame order.

► FC9958, Korean 4.3-m power cord

FC9958, Korean 4.3-m power cord, provides a NEMA L6-30 non-watertight, 4.3-m (14-ft.) power cord, rated for 200 V ac to 208 V ac or 240 V ac and 24 A, with a South Korean plug. This power cord is suggested for use in South Korea. This feature occurs on the 3952 Tape Frame order.

► FC9959, Unterminated power cord

FC9959, Unterminated power cord, provides an unterminated, non-watertight, 4.3-m (14-ft.) power cord, rated for 200 V ac to 208 V ac and 24 Amp, with IRAM and BSMI agency certifications. This is the recommended cord for Argentina and China. This feature occurs on the 3952 Tape Frame order.

► FC9966, Unterminated power cord (China)

FC9966, Unterminated power cord (China), provides an unterminated, non-watertight, 4.3-m (14-ft.) power cord, rated for 200 V ac to 208 V ac and 24 Amp, with CCC agency certifications. This is the recommended cord for China. This feature occurs on the 3952 Tape Frame order.



IBM Virtualization Engine TS7700 implementation steps

In this appendix, we document the typical implementation scenario for the TS7700 Virtualization Engine. Use these steps as guidance only; tailor the steps specifically for your site. Most of the steps are the same for IBM Virtualization Engine TS7740 and IBM Virtualization Engine TS7720 implementations, although steps related to physical tape do not apply to TS7720 Virtualization Engine configurations. The steps assume that the host system is DFSMS and covers a first-time TS7700 Virtualization Engine library installation.

For more details see the following DFSMS documentation:

- ▶ *z/OS DFSMS Object Access Method Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC35-0427
- ▶ *z/OS DFSMS Object Access Method Planning, Installation, and Storage Administration Guide for Object Support*, SC35-0426

These DFMS references are available at the following address:

<http://www.ibm.com/systems/z/os/zos/bkserv/r9pdf/#dfsms>

Now, we detail an example definition procedure for a first-time TS7700 Virtualization Engine tape library installation.

Perform the following steps:

1. Analyze the TS7700 Virtualization Engine targeted workloads to determine the appropriate TS7700 Virtualization Engine configuration for your environment. The BatchMagic tool can be helpful in making this determination. Have your IBM System Storage Specialist helping you through this step.

Important: This is the crucial step for configuring your TS7700 Virtualization Engine correctly and should be done in conjunction with your IBM System Storage Specialist.

2. Check the latest TS7700 Virtualization Engine and IBM System Storage TS3500 Tape Library Systems Assurance Product Review (SAPR) Guides. Again, your IBM System Storage Specialist can be instrumental in this task.
3. Order stacked volume media or labels.
4. Set up the HCD (see 5.2.1, “Defining devices through HCD for Cluster 0” on page 270 for more information).
5. Vary CHPIDs, paths, and TS7700 Virtualization Engine virtual drives online.
6. Review JES3 system considerations (see Appendix D, “JES3 examples and information” on page 821 for more information).
7. For DFSMSrmm:
 - Create VLPOOLS for logical volumes.
 - ADD logical volumes/location to DFRMM.
 - Check REJECT PARMs. With DFSMS V1R10 and later, the use of PARTITION and OPENRULE is recommended.
 - Refresh DFRMM.
8. For CA1:
 - If your planned logical volume range has not been defined to your Tape Management System (TMS), you must extend the TMC. Also, check the amount of free DSNBs available and add any, if required, that are concurrent with this extension.
 - Check that the SCRTCH PARM is set to YES in PPOPTION. When a tape becomes scratch or a DSNB record gets deleted, control is given to the security exit TMSUX2S. This is the scratch notification interface between CA1 and the 3953 Library Manager.
 - Depending on the parameters specified in PPOPTION, you might need to initialize the logical volumes after they have been inserted.

Note: Verify with the vendor of your tape management product to ensure that the level of software installed on your system will support the TS35000 and TS7700 Virtualization Engine.

9. Verify and, if necessary, modify the current SYS1.PARMLIB members:
 - IEFSSNxx
Add or update the OAM1 entry with the name of the initialization module (CBRINIT) executed at IPL.
 - SCHEDxx
The default program properties table, IEFSDPPT, included with z/OS has entries for all z/OS elements (including CBROAM). Therefore, the SCHEDxx entry no longer needs to be updated to include CBROAM.

- IGDSMSxx

If you want OAM to start automatically as part of SMS initialization, add the OAMPROC and OAMTASK optional parameters.

Note: If you use an independent storage vendor tape management system, it might require that the OAM address space be started after the tape management initialization. In this case, do not start OAM automatically. Check with the vendor of the tape management product.

- CONSOLxx

If you want to receive library messages at a specific console, update the CONSOLxx member referenced by IEASYSxx. This console name must also be defined to SMS and is set when you define the library through ISMF.

- DEVSUPxx

This member is used to hard partition your TS7700 Virtualization Engine by specifying category codes to provide a unique range of tapes for each system. If no category codes are specified, the defaults are assigned.

Important: Do not use the DEVSUPxx default categories. Avoiding the use of the defaults will make future partitioning of the library with other systems easier and more secure. See *Guide to Sharing and Partitioning IBM Tape Library Data*, SG24-4409 for additional information.

- COMMNDxx

If you want your library automatically brought online after each IPL, add the VARY SMS,LIBRARY command. Another area where you can bring your library automatically online is through ISMF when you define your library. Set the Initial Online Status to YES.

- GRSCNFxx

If the tape library will be shared among two or more systems in a DFSMS complex, a global resource serialization ring can be created to include all sharing systems. This step enables OAM to serialize the cartridge entry process.

- LOADxx (optional)

The default data set name for the TCDB is SYS1.VOLCAT.VGENERAL. If you want to use this name, no change is required to the LOADxx member. If you want to use a separate HLQ than the defined default, update columns 64 - 71 of the SYSCAT statement with the preferred HLQ.

- COFVLFxx

Add the volume catalogs to the IGGCAS class definitions where you have other ICF catalogs.

- ALLOCxx

Add tape automation policies.

- IECIOSxx

Set Missing Interrupt Handler values. This value should be 45 minutes for the TS7700 Virtualization Engine logical drives.

Note: For coding examples and more information regarding SYS1.PARMLIB updates, see *IBM TotalStorage 3494 Tape Library: A Practical Guide to Tape Drives and Tape Automation*, SG24-4632.

10. If your IBM System Service Representative (SSR) has installed your TS7700 Virtualization Engine, perform the following definitions at the management interface:

- Define stacked volume ranges (see 4.3.3, “Defining VOLSER ranges for physical (stacked) volumes” on page 203 for more information).
- Define Fast Ready categories (see 4.3.5, “Defining Fast Ready categories” on page 220 for more information). The Fast Ready category number you define is the same value as that defined for MEDIA1 or MEDIA2, as specified in the DEVSUPxx of your PARMLIB.
- Define the expired logical volume data policy (see 4.3.6, “Defining the logical volume expiration time” on page 222 for more information).
- Define TS7700 Virtualization Engine management policies. Set the following values:
 - Inhibit reclaim schedule
 - Reclaim threshold percentage
 - Free storage threshold

11. If stacked volume media is available, insert the physical TS7700 Virtualization Engine stacked volumes (see “Inserting physical volumes into the TS3500 Tape Library” on page 193 for more information).

12. Create TCDB.

When sizing the TCDB, calculate 275 bytes for each volume entry. For 10,000 logical volumes, 2.75 MB or about three cylinders of 3390 disk space is required. Allow sufficient room for growth.

Note: If you are sharing the tape library (non-hard partitioned), use IDCAMS to IMPORT CONNECT the VOLCAT to the other sharing systems.

13. Define the following security profiles:

- ISMF
- DFSMS constructs
- STGADMIN
- DFSMSrmm
- z/OS Operator Commands

Note: The following manuals are helpful in setting up your security environment:

- ▶ *z/OS DFSMSdfp Storage Administrator Reference*, SC35-0422
- ▶ *z/OS DFSMSrmm Implementation and Customization Guide*, SC26-7405
- ▶ *z/OS MVS Planning: Operation*, SC22-7601

14. Create the cataloged OAM procedure.

15. Start OAM. Upon startup, you might receive the following messages:

CBR1115I - No libraries defined to OAM.

CBR0094E - OAM has initialized without tape or object support.

These messages are informational and expected. They will be resolved when the library has been defined and the DFSMS SCDS activated.

Tip: Code RESTART=YES PARM into your OAM procedure. This PARM, upon an activation of the SCDS, automatically restarts OAM. This step is important when you have changed tape library related constructs. Without this PARM, you must restart OAM manually after an SCDS activation if you have changed tape library-related constructs. To restart, run the following command:

```
F OAM,RESTART
```

16. Define the following DFSMS constructs using ISMF:

- Tape library
- Data class
- Storage class
- Management Class
- Storage group

17. Create the following ACS routines for DFSMS-managed tape:

- Data class
- Storage class
- Management Class
- Storage group

18. Define the Tape Library using ISMF.

19. Translate the DFSMS ACS routines through ISMF.

20. Validate the DFSMS SCDS through ISMF.

21. Test the ACS routines through ISMF.

22. Activate the DFSMS SCDS.

Note: By previously varying the TS7700 Virtualization Engine drives online, upon activation of the SCDS, the library will come online automatically. You must have a TS7700 Virtualization Engine drive online before the library can come up. To bring a library online manually, issue the following command:

```
VARY SMS,LIB(library-name),ONLINE
```

After OAM is started, and the library is online, and the host has even a single path initialized for communication with the library, that host is capable of doing cartridge entry processing. The library drives do not have to currently be online for the path to be usable. If this library is being partitioned with other host systems, it is important that each system's tape management CBRUXENT first be coded and linked to prevent the host from entering volumes that do not belong to it.

23. Insert logical volumes (see "Inserting logical volumes" on page 240 for more information).

24. The TS7700 Virtualization Engine is now ready for use.

Tip: For help in navigating the ISMF screens, and detailed information pertaining to the definition of the DFSMS constructs, see the following publications:

- ▶ *IBM TotalStorage 3494 Tape Library: A Practical Guide to Tape Drives and Tape Automation*, SG24-4632
- ▶ *z/OS DFSMSdfp Storage Administrator Reference*, SC35-0422



TS3500 and TS7700 checklists

These checklists should be used when attaching a IBM Virtualization Engine TS7740 to a TS3500 Tape Library and TS7700 virtualization engine installation. For the most recent version, see the TS7700 Virtualization Engine TS3500 Tape Library Installation Worksheets that are part of the Solution Assurance documents for the TS7700 Virtualization Engine. You can obtain this document from the Customer Information Center (under Planning, Checklists) or from your IBM System Services Representative (SSR).

TS7700 Virtualization Engine installation worksheets

The following customer-supplied information (as applicable) will be required during the TS7700 Virtualization Engine installation. These worksheets must be filled out by the customer and an IBM Support team and given to the installation team.

Tip: Clusters that belong to the same grid will probably be installed at separate locations, attached to separate libraries (or even not using a Tape Library at all). Therefore, you might be looking to more than one physical Tape Library within the same worksheet. Be sure that the relationship between clusters and tape libraries is clearly defined in the worksheets.

These worksheets are based on worksheets in the TS7700 information center:

<http://publib.boulder.ibm.com/infocenter/ts7700/cust/index.jsp>

Additional Tape Control Units (3592-J70/C06) or Virtual Tape Subsystems (3494-B10/B20) might be attached to the same TS3500 Tape Library. The installation or configuration of these devices is not within the scope of this publication. In most cases, these devices must be installed using documentation provided with the devices.

The terms *Cluster 0*, *Cluster 1*, *Cluster 2*, and *Cluster 3* refer to the separate TS7700 Virtualization Engines participating in a grid (multi-cluster) configuration.

If you are installing a stand-alone (single-grid) TS7700 Virtualization Engine, only Cluster 0 information is needed. In a two-way grid, fill up for Clusters 0 and 1; for three-way grid, fill up for Clusters 0, 1, and 2.

Grid cluster descriptions

Table C-1 shows the grid cluster descriptions worksheet. Provide your SSR with a serial number and a description of which TS7700 Virtualization Engine is which cluster. The information in Table C-1 informs the SSR which TS7700 Virtualization Engine is designated as Cluster 0, 1, 2, or 3.

Complete the Description of location column by entering geographical information that is unique to each cluster. Enter as much information as possible; some suggestions for descriptions are the x, y coordinates within a data center, room number, floor number, building numbers, city, and so on.

Table C-1 Grid cluster descriptions worksheet

Cluster	Machine type-model	Serial number	Description of location
Cluster 0	3957-V06		
	3584-Lxx		
Cluster 1	3957-V06		
	3584-Lxx		
Cluster 2	3957-V06		
	3584-Lxx		
Cluster 3	3957-V06		
	3584-Lxx		

TS3500 Tape Library configuration information

Table C-2 on page 806 shows the TS3500 Tape Library configuration information worksheet. When filling out this table, make sure to correlate each Cluster to its Library (if applicable).

The following information is required:

- ▶ Network Ethernet configuration method
- ▶ Network Ethernet I/P address for TS3500 Tape Library Web Specialist interface
- ▶ TS3500 Ethernet host name
- ▶ Network Ethernet gateway
- ▶ Network Ethernet subnet mask
- ▶ Maximum number of physical cartridges
- ▶ Logical Library name for the Cluster's library
- ▶ Physical position and worldwide name (WWN) of each tape drive within TS3500 belonging to each Cluster.

Table C-2 TS3500 Tape Library configuration information worksheet

Field	Value	Notes
TS3500 Tape Library Ethernet Network Configuration Method (customer network)	Cluster 0: DHCP [] Fixed IP []	The network configuration method is specified by the customer's LAN administrator. It is either Fixed IP or Dynamic Host Configuration Protocol (DHCP). Fixed IP is recommended.
	Cluster 1: DHCP [] Fixed IP []	
	Cluster 2: DHCP [] Fixed IP []	
	Cluster 3: DHCP [] Fixed IP []	
TS3500 Tape Library Ethernet IP Address (customer network) Used for TS3500 Tape Library Web specialist access	Cluster 0: ____.____.____.____	If the Network Configuration Method is DHCP, then this field is not used. The Ethernet ports are 10/100 Mb only.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
	Cluster 3: ____.____.____.____	
TS3500 Tape Library Ethernet Hostname	Cluster 0:	The entry in this field is used to identify the machine in remote support logs. Even if host names are not typically used by the customer, this host name is still required. A typical host name is ATL1.
	Cluster 1:	
	Cluster 2:	
	Cluster 3:	
TS3500 Tape Library Ethernet Subnet Mask (customer network)	Cluster 0: ____.____.____.____	If the Network Configuration Method is DHCP, then this field is not used.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
	Cluster 3: ____.____.____.____	
TS3500 Tape Library Ethernet Gateway (customer network)	Cluster 0: ____.____.____.____	If the Network Configuration Method is DHCP, then this field is not used.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
	Cluster 3: ____.____.____.____	

TS3500 Tape Library Logical Library Name (for TS7740 Cluster attachment)	Cluster 0:	<p>Each IBM Virtualization Engine TS7740 must be connected to a single TS3500 Tape Library logical library. This logical library must have a name, which should have been assigned when the logical library was created. Record the logical library name (assign it if necessary). The logical library name will be needed when performing the following tasks:</p> <ul style="list-style-type: none"> ▶ Configuring the logical library. ▶ Obtaining the Starting Element Address for the logical library. ▶ Obtaining the physical position of tape drives within the logical library. ▶ Obtaining the WWNNs of those tape drives. ▶ Setting the Cartridge Assignment Policy. ▶ Configuring ALMS.
	Cluster 1:	
	Cluster 2:	
	Cluster 3:	
Max Cartridges	Cluster 0:	<p>This value defaults to the number of physical cartridge slots currently installed in the TS3500 Tape Library. You might want to set a different value in the following cases:</p> <ul style="list-style-type: none"> ▶ The customer wants to restrict the number of cartridges in the logical library to manage the cost of application licensing. ▶ The TS3500 Tape Library will be expanded at a later date and you want to avoid reconfiguring the host.
	Cluster 1:	
	Cluster 2:	
	Cluster 3:	

Important: FC0500 provides for library and tape drive microcode updates by an IBM System Service Representative) for cluster-attached TS3500 Tape Library. We strongly advise against customers performing library or tape drive microcode updates on their own in such an environment.

TS3500 Tape Library drive information worksheet

Table C-3 shows the TS3500 Tape Library drive information worksheet. Fill out the assigned tape drives for each cluster (you can get that information from 3584's Web Interface in two steps: under Drives - Drive Assignment screen, and World Wide Names).

Table C-3 TS3500 Tape Library drive information worksheet

Field	Value	Notes
<p>Tape Drive Physical Positions (Fn, Rnn) in the TS3500 Tape Library for the drives assigned to this TS7700 Virtualization Engine.</p> <p>Record the physical position and the last two digits of the worldwide node name (WWNN) for each tape drive. Circle the drives that will be control paths.</p> <p>Read the notes in the right column of this table for guidance in filling in the values.</p>	<p>Cluster 0: F=Frame, R=Row WWNN = worldwide node name.</p> <p>1. F____,R____ WWNN____ 2. F____,R____ WWNN____ 3. F____,R____ WWNN____ 4. F____,R____ WWNN____ 5. F____,R____ WWNN____ 6. F____,R____ WWNN____ 7. F____,R____ WWNN____ 8. F____,R____ WWNN____ 9. F____,R____ WWNN____ 10. F____,R____ WWNN____ 11. F____,R____ WWNN____ 12. F____,R____ WWNN____ 13. F____,R____ WWNN____ 14. F____,R____ WWNN____ 15. F____,R____ WWNN____ 16. F____,R____ WWNN____</p>	<p>List the drives for a single cluster in order, using the frame and row. The lowest numbered drive should be in the lowest numbered frame and row for its assign drive slots. This is not a requirement, but it might help avoid confusion when identifying drives during future troubleshooting.</p> <p>Distribute the tape drives across two TS3500 Tape Library frames to improve availability. Place eight drives (two control path drives) in one frame and either drives (two control path drives) in a second frame.</p> <p>To obtain the WWNN from the Operator Panel, select Menu → Vital Product Data → World Wide Node Names.</p>
	<p>Cluster 1: F=Frame, R=Row WWNN = worldwide node name.</p> <p>1. F____,R____ WWNN____ 2. F____,R____ WWNN____ 3. F____,R____ WWNN____ 4. F____,R____ WWNN____ 5. F____,R____ WWNN____ 6. F____,R____ WWNN____ 7. F____,R____ WWNN____ 8. F____,R____ WWNN____ 9. F____,R____ WWNN____ 10. F____,R____ WWNN____ 11. F____,R____ WWNN____ 12. F____,R____ WWNN____ 13. F____,R____ WWNN____ 14. F____,R____ WWNN____ 15. F____,R____ WWNN____ 16. F____,R____ WWNN____</p>	

	<p>Cluster 2:</p> <p>F=Frame, R=Row WWNN = worldwide node name.</p> <p>1. F____,R____WWNN____</p> <p>2. F____,R____WWNN____</p> <p>3. F____,R____WWNN____</p> <p>4. F____,R____WWNN____</p> <p>5. F____,R____WWNN____</p> <p>6. F____,R____WWNN____</p> <p>7. F____,R____WWNN____</p> <p>8. F____,R____WWNN____</p> <p>9. F____,R____WWNN____</p> <p>10. F____,R____WWNN____</p> <p>11. F____,R____WWNN____</p> <p>12. F____,R____WWNN____</p> <p>13. F____,R____WWNN____</p> <p>14. F____,R____WWNN____</p> <p>15. F____,R____WWNN____</p> <p>16. F____,R____WWNN____</p>	
	<p>Cluster 3:</p> <p>F=Frame, R=Row WWNN = worldwide node name.</p> <p>1. F____,R____WWNN____</p> <p>2. F____,R____WWNN____</p> <p>3. F____,R____WWNN____</p> <p>4. F____,R____WWNN____</p> <p>5. F____,R____WWNN____</p> <p>6. F____,R____WWNN____</p> <p>7. F____,R____WWNN____</p> <p>8. F____,R____WWNN____</p> <p>9. F____,R____WWNN____</p> <p>10. F____,R____WWNN____</p> <p>11. F____,R____WWNN____</p> <p>12. F____,R____WWNN____</p> <p>13. F____,R____WWNN____</p> <p>14. F____,R____WWNN____</p> <p>15. F____,R____WWNN____</p> <p>16. F____,R____WWNN____</p>	

Media volume serial range

Complete Table C-4 with the requested information by the installation team. There might be one or many media volume serial ranges, so complete as many rows as apply to your system.

Fill out the Cluster numbers, the associated volume serial ranges, along with the media type for each range. Also fill out the Distributed Library Sequence Numbers that are associated to each cluster in the proper field, and the planned Home Pool for each range (if you want, a pool separation; if no requirement exists, the default Scratch Pool is 00). You can mark if Encryption will be adopted for an specific pool or not (not valid for pool 00).

Table C-4 Media volume serial range

Cluster (0-3)	From	To	Media Type	Distributed Library Sequence Number	Home Pool	Encrypted Pool (Y/N)

The following notes apply Table C-4 on page 810:

- ▶ From, To: This range contains the bar code label VOLSERS of all the cartridges assigned to a single TS7740 Virtualization Engine. As an example, if cartridges assigned to a TS7740 Virtualization Engine have barcode labels in the range from A00000JA - A00500JA, then record the following information:
 - From: A00000
 - To: A00500
 - Media Type: The type is indicated by the last two characters of the 8-character bar code label VOLSER on the cartridges. As an example, if the cartridges are labeled 123456JA, then the media type is JA. The rules are as follows:
 - JA and JJ tape cartridges are supported and can be mixed in a TS7700 Virtualization Engine.
 - JB tape cartridges are also supported (and can be mixed with JA and JJ tape cartridges) if all of the tape drives associated with the TS7700 Virtualization Engine are 3952 E05/E06/EU6 Tape Drives and no E05 Tape Drives are in J1A emulation mode.
 - No other tape types are currently supported for use with the TS7740 Virtualization Engine.
 - If at least one tape cartridge that is associated with a TS7700 Virtualization Engine has been written by a 3592 E05/EU5 Tape Drives that is *not* in 3592 J1A emulation mode, then the TS7700 Virtualization Engine will no longer support any 3592 J1A Tape Drives or any 3592 E05/EU5 Tape Drive that is in J1A emulation mode. After the 3592 “E05 Native” mode has been used, you may not return to J1A mode. The reason is because the 3592 J1A Tape Drives cannot read or write a tape cartridge written in E05 mode and the TS7700 Virtualization Engine does not currently support mixed 3592 J1A and E05 Tape Drives (unless all E05 Tape Drives are in J1A emulation mode. To use E05 Tape Drives in native mode, the drives must be set to E05 native mode.
 - The capacity of a JJ tape cartridge is 60 GB if written by a J1A drive (or an E05 drive working in J1A emulation mode) or 100 GB if written by an E05 drive in native mode. The capacity of a JA tape cartridge is 300 GB if written by a J1A drive (or an E05 drive working in J1A emulation mode) or 500 GB if written by an E05 drive in native mode. The capacity of a JB cartridge is 700 GB.
 - TS7700 Virtualization Engine Feature Code 9900 (Encryption) requires that all tape drives are in E05 Native mode. The Encryption feature is *not* compatible with J1A emulation mode.

Using these rules, determine whether you want the tape drives that are attached to the TS7700 Virtualization Engine to be in J1A emulation mode or in E05 Native mode. Record this in the Tape Drive Format entry in Table C-6 on page 815.

- ▶ (Distributed) Library Sequence Number: The Distributed Library Sequence Number is assigned by the customer’s administrator. This five-character name is used as an identifier for a specific cluster and the associated Library a Grid configuration. This identifier is specified in the TS7700 Virtualization Engine configuration. *It is required even if the TS7700 Virtualization Engine is not in a Grid configuration.*

Note: Each TS7700 Virtualization Engine Cluster must have a single, unique value for the Distributed Library Sequence Number. For the TS7700 Virtualization Engine, a typical value is the last 5 digits of the 3952-F05 frame serial number.

- Home Pool (also called Scratch Pool): You might have assigned a Home Pool value. If one has not been set, the default value is 00.

Note: A *pool* is a group of physical tape cartridges. A *scratch pool* is a group of cartridges that are considered to be scratch, meaning they are ready for use by any write job.

- Encrypted pool: Encryption can be controlled by pools. If you want, mark the pool enabled for encryption.

Note: Certain requirements must be met before enabling Encryption. See the customer information center and Chapter 4, “Hardware implementation” on page 175 and Chapter 5, “Software implementation” on page 2635 for more details about encryption.

TS7700 Virtualization Engine configuration information

Table C-5 shows the TS7700 Virtualization Engine configuration information worksheet. Fill in the required information in the proper fields:

- Library Sequence Numbers (composite and distributed)
- Disk cache type
- IP addresses (virtual, primary, alternate and NTP)
- Subnet mask
- Gateway
- Tape drive format
- Network speed settings

Table C-5 TS7700 Virtualization Engine configuration information worksheet

Field	Value	Notes
Composite Library Sequence Number		This five-character name must be the same on all clusters (peers) within the same Grid. This identifier is specified in the TS7700 Virtualization Engine configuration. It is required even if the machine is not in a Grid configuration. The Composite Library Sequence Number must differ from the Distributed Library Sequence number specified.
Distributed Library Sequence Number	Cluster 0: _____	The Distributed Library Sequence Number is assigned by the customer's administrator. This is a five-hex character name that is used as an identifier for a specific cluster. Each TS7740 Virtualization Engine must have a single, unique value for the Distributed Library Sequence Number. For the TS7740 Virtualization Engine, a typical value is the last 5 digits of the 3952-F05 frame serial number. The Distributed Library Sequence Number must differ from the Composite Library Sequence number specified.
	Cluster 1: _____	
	Cluster 2: _____	
	Cluster 3: _____	

Disk Cache Type	Cluster 0: _____	<p>The TS7720 Cache uses a 3956-CS7/XS7 cache, which is a DS4000cache type.</p> <p>The TS7740 Cache uses a 3956-CC6/CX6 or 3956-CC7/CX7 cache, which is a DS6000™ or a DS4000® cache type.</p>
	Cluster 1: _____	
	Cluster 2: _____	
	Cluster 3: _____	
Customer IP 1 (Virtual)	Cluster 0: ____.____.____.____	<p>Used for TS7700 Virtualization Engine web management interface.</p> <p>This number is a virtual IP that is not associated with a physical cable. It communicates through the Primary IP, and automatically fails over to the Alternate IP when required.</p>
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
	Cluster 3: ____.____.____.____	
Customer IP 2 (Primary)	Cluster 0: ____.____.____.____	<p>This is the IP address used to connect to the TS7700 Virtualization Engine through the internal primary network. This IP address must not be used by the customer unless the Virtual IP is inaccessible.</p> <p>The Ethernet ports are 10/100 Mb only.</p>
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
	Cluster 3: ____.____.____.____	
Customer IP 3 (Alternate)	Cluster 0: ____.____.____.____	<p>This IP address is used to connect to the TS7700 Virtualization Engine through the internal alternate network. This IP address must not be used by the customer unless the Virtual IP is inaccessible.</p> <p>The Ethernet ports are 10/100 Mb only.</p>
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
	Cluster 3: ____.____.____.____	
Customer Gateway	Cluster 0: ____.____.____.____	<p>This is used with the virtual, primary, and alternate customer IP addresses.</p>
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
	Cluster 3: ____.____.____.____	

Customer Subnet Mask	Cluster 0: ____.____.____.____	This is used with the virtual, primary, and alternate customer IP addresses.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
	Cluster 3: ____.____.____.____	
NTP server IP address (if used) Using the NTP server ensures that all components have consistent time settings.	____.____.____.____	The TCP/IP address you obtain from the customer is either the NTP server at the customer site (if the customer maintains one locally), or an Internet server. Use of an Internet server assumes that the customer allows access to the Internet on the NTP services port (TCP/IP port 123).
Tape Drive Format	Cluster 0: J1A Emulation Mode <input type="checkbox"/> E05 Native Mode <input type="checkbox"/>	To determine whether to use J1A emulation mode or E05 Native mode, refer to the rules listed under Media Type.
	Cluster 1: J1A Emulation Mode <input type="checkbox"/> E05 Native Mode <input type="checkbox"/>	
	Cluster 2: J1A Emulation Mode <input type="checkbox"/> E05 Native Mode <input type="checkbox"/>	
	Cluster 3: J1A Emulation Mode <input type="checkbox"/> E05 Native Mode <input type="checkbox"/>	

TSSC remote support information worksheet

Table C-6 shows the TSSC remote support information worksheet. Fill out the following required information, such as the following information:

- ▶ Dial-out prefixes
- ▶ Call Home modem number
- ▶ Call Home Ethernet IP, subnet and gateway
- ▶ Customer name
- ▶ Customer contact information

Table C-6 TSSC remote support information

Field	Value	Notes
Dial Out Prefix (if required)	Cluster 0:	If you need to dial a prefix to get an outside line, specify it here. As an example, if you need to dial 9 to get an outside line, record a 9 in this field.
	Cluster 1:	
	Cluster 2:	
	Cluster 3:	
Call-in Modem Number	Cluster 0:	This telephone number is assigned to the analog phone line where the Call Home modem will be connected.
	Cluster 1:	
	Cluster 2:	
	Cluster 3:	
Customer Company Name	Cluster 0:	This value is sent out with the Call Home records.
	Cluster 1:	
	Cluster 2:	
	Cluster 3:	
Customer Contact Phone Number	Cluster 0:	This is the phone number that IBM contacts when Call Home indicates a failure.
	Cluster 1:	
	Cluster 2:	
	Cluster 3:	
Off Shift Customer Contact Phone Number	Cluster 0:	This is the alternate phone number that IBM contacts when Call Home indicates a failure.
	Cluster 1:	
	Cluster 2:	
	Cluster 3:	

<p>IP Address (customer network) Used for optional Ethernet Call Home</p> <p>Call Home through an Ethernet connection greatly improves the speed of the data transfer, making the data available to support personnel more quickly. At this time, the Call Home through an Ethernet connection is only used for Call Home (outbound). Enable Call Home to improve data collection when a problem occurs. Call Home through an Ethernet connection provides better throughput and lower cost than Call Home through a modem connection.</p>	Cluster 0:	<p>This is an optional (but good practice) IP address on the customer network that will allow the TSSC to Call Home using a high speed connection through the customer network to the Internet. Ports 443 (HTTPS) and 80 (HTTP) on this IP address must be free for outbound traffic through the customer network to the Internet.</p> <p>The modem is still needed for Call In (remote support)</p>
	Cluster 1:	
	Cluster 2:	
	Cluster 3:	
<p>Subnet Mask (for optional Ethernet Call Home)</p>	Cluster 0:	<p>This is the subnet mask associated with the IP Address (for optional Ethernet Call Home).</p>
	Cluster 1:	
	Cluster 2:	
	Cluster 3:	
<p>Gateway (for optional Ethernet Call Home)</p>	Cluster 0:	<p>This is the gateway associated with the IP Address (for optional Ethernet Call Home).</p>
	Cluster 1:	
	Cluster 2:	
	Cluster 3:	

Grid local address

The following notes apply to Table C-7 on page 817:

- ▶ If the TS7700 Virtualization Engine you are installing is a stand-alone machine (not part of a Grid configuration), leave the table blank.
- ▶ The Grid interfaces are the 1 Gb Internet connections between clusters, allowing them to
- ▶ The primary and alternate Grid interfaces must be on separate subnets. As an example, 10.10.1.n for the primary interface and 10.11.1.n for the alternate interface. Notice that the second set of octets must be different. If the Grid interfaces are direct connected (without using Ethernet switches), then using separate subnets is required.

The Grid interfaces require connections using the following TCP/IP ports:

- ▶ 7 (Ping)
- ▶ 9 (Discard Service - for bandwidth measuring tools)
- ▶ 123 (Network Time Protocol)
- ▶ 350 (Distributed Library to Distributed Library File Transfer)
- ▶ 1415 (WebSphere message queues Grid to Grid)
- ▶ 1416 (WebSphere message queue HDM to HDM)

The following TCP/IP ports are also useful in service scenarios if allowed:

- ▶ 20 (FTP)
- ▶ 21 (FTP)
- ▶ 23 (Telnet)

Table C-7 Grid local address

Field	Value	Notes
Primary Grid Interface IP address	Cluster 0: ____.____.____.____	The Primary Grid Interface is the 1 Gb Ethernet Adapter located in slot C4 of the 3957-V06. The Primary Grid Interface at each cluster will connect to the Primary Grid Interface at each of the other clusters in the same Grid.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
	Cluster 3: ____.____.____.____	
Primary Grid Interface Network Mask	Cluster 0: ____.____.____.____	The network masks accompany an IP address and the two values work together. Applying the subnet mask to an IP address splits the address into two parts, an extended network address and a host address
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
	Cluster 3: ____.____.____.____	
Primary Grid Interface Gateway	Cluster 0: ____.____.____.____	If a gateway is not used, leave this field blank. If using crossover cables, you must not specify a gateway.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
	Cluster 3: ____.____.____.____	
Alternate Grid Interface IP address	Cluster 0: ____.____.____.____	The Alternate Grid Interface is the 1 Gb Ethernet Adapter located in slot C5 of the 3957-V06. The Alternate Grid Interface at each cluster will connect to the Alternate Grid Interface at each of the other clusters in the same Grid.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
	Cluster 3: ____.____.____.____	

Alternate Grid Interface Network Mask	Cluster 0: ____.____.____.____	Specify the subnet mask to use with the alternate Grid IP addresses.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
	Cluster 3: ____.____.____.____	
Alternate Grid Interface Gateway	Cluster 0: ____.____.____.____	If a gateway is not used, leave this field blank. If using crossover cables, you must not specify a gateway.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
	Cluster 3: ____.____.____.____	

TSSC Grid configuration information

The following notes apply to Table C-8 on page 819:

- ▶ If the TS7700 Virtualization Engine you are installing is a stand-alone machine (not part of a Grid configuration), leave the table blank.
- ▶ If you will be using the Grid in a Cascade Deferred style, you do not need the Autonomic Ownership Takeover Manager (AOTM) and you can leave the table blank.

Note: Cascade Deferred means one or more clusters do not have host FICON connections enabled in normal operation. There is no need to use AOTM on a cluster that does not have host FICON connections enabled in normal operation.

- ▶ See “The Autonomic Ownership Takeover Manager (AOTM)” of the *Virtualization Engine TS7700 Installation Roadmap for use with IBM Systems Storage TS3500, IBM 3584 Tape Library*, PN 23R7608 for more information about AOTM before you continue. Do *not* attempt to configure AOTM, but use the information to make an informed decision about whether to use AOTM.
- ▶ If you do not want to use the AOTM, leave the table blank.
- ▶ The TSSC Grid interface is only used for the AOTM.
- ▶ Each cluster can be configured to use AOTM to provide ownership takeover for one cluster.

The AOTM requires the following TCP/IP ports:

- ▶ 7 (Ping)
- ▶ 80 (HTTP)

Table C-8 TSSC Grid configuration information

Field	Value	Notes
TSSC Grid Interface IP Address	Cluster 0: ____.____.____.____	The TSSC Grid Interface is used to allow the TSSC at one cluster to communicate with the TSSC at one other cluster. This is required if the AOTM will be used.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
	Cluster 3: ____.____.____.____	
TSSC Grid Interface Subnet Mask	Cluster 0: ____.____.____.____	Specify the subnet mask to use with the grid interface IP address.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
	Cluster 3: ____.____.____.____	
TSSC Grid Interface Gateway	Cluster 0: ____.____.____.____	Specify the gateway IP address to use with the grid interface IP addresses.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
	Cluster 3: ____.____.____.____	
Takeover Mode	Cluster 0: ROT <input type="checkbox"/> WOT <input type="checkbox"/>	The Takeover Mode must be either Read Only Takeover (ROT) or Write Only Takeover (WOT).
	Cluster 1: ROT <input type="checkbox"/> WOT <input type="checkbox"/>	
	Cluster 2: ROT <input type="checkbox"/> WOT <input type="checkbox"/>	
	Cluster 3: ROT <input type="checkbox"/> WOT <input type="checkbox"/>	
Grace Period (Minutes)	Cluster 0:	When a cluster detects that another cluster in the same Grid has failed, it will wait for the number of minutes specified as the Grace Period before it attempts to take over the volumes of the failed cluster. A typical Grace Period is 25 minutes.
	Cluster 1:	
	Cluster 2:	
	Cluster 3:	

Retry Period (Minutes)	Cluster 0:	The retry period is the number of minutes between attempts to take over ownership of the volumes associated with a failed cluster. A typical Retry Period is five minutes. In most cases, this will be the same for both clusters.
	Cluster 1:	
	Cluster 2:	
	Cluster 3:	



D

JES3 examples and information

This appendix presents configuration examples and other JES3 considerations. It describes the necessary parameters and the considerations if more than one device type is installed in a library in a JES3 environment.

We provide two examples:

- ▶ Two libraries with an intermix of native drives of 3592-J1A and 3592-E05 installed
- ▶ Three libraries with an intermix of 3592-J1A, 3592-E05, 3592-E06/EU6, single cluster IBM Virtualization Engine TS7700 Grid, and multi-cluster TS7700 Virtualization Engine Grid installed

These examples provide all of the necessary information to install any possible configuration in an IBM System Storage TS3500 Tape Library. For more basic information about the products in these scenarios, see the following publications:

- ▶ *z/OS JES3 Initialization and Tuning Guide*, SA22-7549
- ▶ *z/OS JES3 Initialization and Tuning Reference*, SA22-7550

JES3 support for system-managed tape

This section describes JES3 TS3500 Tape Library support with DFSMS. The primary purpose of this support is to maintain JES3 resource allocation and share tape allocations. For detailed information, see *z/OS JES3 Initialization and Tuning Reference*, SA22-7550.

DFSMS has support that provides JES3 allocation with the appropriate information to select a TS3500 Tape Library device by referencing device strings with a common name among systems within a JES3 complex.

To set up a TS3500 Tape Library in a JES3 environment, perform the following steps:

1. Define library device groups (LDGs). Prepare the naming conventions in advance. Clarify all the names for the library device groups that you need.
2. Include the esoteric names from step 1 in the hardware configuration definition (HCD) and activate the new EDT.
3. Update the JES3 INISH deck:
 - a. Define all devices in the TS3500 Tape Library through DEVICE statements.
 - b. Set JES3 device names through the SETNAME statement.
 - c. Define which device names are subsets of other device names through the high water mark setup name (HWSNAME) statement.

All TS3500 Tape Library units can be shared between processors in a JES3 complex. They must also be shared among systems within the same SMSplex.

Note: Tape drives in the TS3500 Tape Library cannot be used by JES3 dynamic support programs (DSPs).

Define all devices in the libraries through DEVICE statements. All TS3500 Tape Library tape drives within a complex must be either JES3-managed or non-JES3-managed. Do not mix managed and non-managed devices. Mixing might prevent non-managed devices from use for new data set allocations and reduce device eligibility for existing data sets. Allocation failures or delays in the job setup will result.

Neither JES3 or DFSMS verifies that a complete and accurate set of initialization statements is defined to the system. Incomplete or inaccurate TS3500 Tape Library definitions can result in jobs failing to be allocated.

Library device groups

Library device groups (LDGs) isolate the TS3500 Tape Library drives from other tape drives in the complex. They allow JES3 main device scheduler (MDS) allocation to select an appropriate set of library-resident tape drives.

The DFSMS JES3 support requires LDGs to be defined to JES3 for SETNAME groups and HWSNAME names in the JES3 initialization statements. During converter/interpreter (C/I) processing for a job, the LDG names are passed to JES3 by DFSMS for use by MDS in selecting library tape drives for the job. Unlike a JES2 environment, a JES3 operating environment requires the specification of esoteric unit names for the devices within a library. These unit names are used in the required JES3 initialization statements.

Important: Even if the LDG definitions are defined as esoterics in the HCD, they are not used in the JCL. There is no need for any UNIT Parameter in JES3 JCL for libraries. The allocation goes through the automatic class selection (ACS) routines. Coding a UNIT Parameter might cause problems.

The only need for coding the LDG definition in HCD as an esoteric name is the HWSNAME definitions in the JES3 INISH deck.

Each device within a library must have exactly four special esoteric names associated with it. These names are:

- ▶ The *complex-wide name* is always LDGW3495. It allows you to address every device and device type in every library.
- ▶ The *library-specific name* is an eight character string composed of LDG prefixing the five digit library identification number. It allows you to address every device and device type in that specific library.
- ▶ The *complex-wide device type*, shown in Table D-1, defines the various device types that are used. It contains a prefix of LDG and a device type identifier. It allows you to address a specific device type in every tape library.

Table D-1 Library device groups: Complex-wide device type specifications

Device type	Complex-wide device type definition
3490E	LDG3490E
3592-J1A	LDG359J
3592-E05	LDG359K
3592-E05 encryption-enabled	LDG359L
3592-E06 encryption-enabled	LDG359M

- ▶ A *library-specific device type name*, which is an eight-character string, starts with a different prefix for each device type followed by the five digit library device number, as shown in Table D-2.

Table D-2 Library device groups: Library-specific device types

Device type	Library-specific device type	Content
3490E	LDE + library number	All 3490E in lib xx
3592-J1A	LDJ + library number	All 3592 Model J1A in lib xx
3592-E05	LDK + library number	All 3592 Model E05 in lib xx
3592-E05	LDL + library number	3592 Model E05 encryption-enabled in lib xx
3592-E06	LDM + library number	3592 Model E06 encryption-enabled in lib xx

It also allows you to address a specific device type in a specific tape library. In a single-cluster grid and a Multi Cluster TS7700 Virtualization Engine Grid installed in two physical libraries, there is still only one library-specific device name. The LIBRARY-ID of the Composite Library is used.

Updating the JES3 INISH deck

To allow JES3 to allocate the appropriate device, you must code several definitions:

- ▶ **DEVICE** statements
- ▶ **SETNAME** statements
- ▶ **HWSNAME** (high water mark setup) statements

These statements are described in detail in the following sections.

DEVICE statement: Defining I/O devices for TS3500 Tape Libraries

Use the **DEVICE** format to define a device so that JES3 can use it. A device statement (Figure D-1) must be defined for each string of TS3500 Tape Library drives in the complex. **XTYPE** specifies a one to eight character name, which is given by the user. There is no default or specific naming convention for this statement. This name is used in other JES3 init statements to group the devices together for certain JES3 processes (for example, allocation). Therefore, it is necessary that all the devices with the same **XTYPE** belong to:

- ▶ The same library
- ▶ The same device type

The letters **CA** in the **XTYPE** definition tell us that this is a **CARTRIDGE** device.

```
*/  Devices 3592-J1A, 3592-E05, and 3592-E06 in Library 1 ...../*  
  
DEVICE,XTYPE=(LB13592J,CA),XUNIT=(1100,*ALL,,OFF),numdev=4  
DEVICE,XTYPE=(LB13592K,CA),XUNIT=(1104,*ALL,,OFF),numdev=4  
DEVICE,XTYPE=(LB13592M,CA),XUNIT=(0200,*ALL,,OFF),numdev=4
```

Figure D-1 *DEVICE statement sample*

Note: TS3500 Tape Library tape drives cannot be used as support units by JES3 dynamic support programs (DSPs). Therefore, do not specify **DTYPE**, **JUNIT**, and **JNAME** parameters on the **DEVICE** statements. No check is made during initialization to prevent TS3500 Tape Library drives from definition as support units, and no check is made to prevent the drives from allocation to a DSP if they are defined. Any attempt to call a tape DSP by requesting a TS3500 Tape Library fails, because the DSP is unable to allocate a TS3500 Tape Library drive.

SETNAME statement

The **SETNAME** statement is used for proper allocation in a JES3 environment. For tape devices, it tells JES3 which tape device belongs to which library. The **SETNAME** statement specifies the relationships between the **XTYPE** values (coded in the **DEVICE** Statement) and the **LDG** names (Figure D-2). A **SETNAME** statement must be defined for each unique **XTYPE** in the device statements.

```
SETNAME,XTYPE=LB1359K,  
      NAMES=(LDGW3495,LDGF4001,LDG359K,LDKF4001)  
      Complex Library Complex Library  
      Wide Specific Wide Specific  
      Library Library Device Device  
      Name Name Name Name
```

Figure D-2 *SETNAME rules*

The rules for the SETNAME statement are:

- ▶ Each SETNAME statement has one entry from each LDG category.
- ▶ The complex-wide library name must be included in all statements.
- ▶ A library-specific name must be included for XTYPEs within the referenced library.
- ▶ The complex-wide device type name must be included for all XTYPEs of the corresponding device type in the complex.
- ▶ A library-specific device type name must be included for the XTYPE associated with the devices within the library.

Note: Do not specify esoteric and generic unit names, such as 3492, SYS3480R, and SYS348XR. Also, never use esoteric names, such as TAPE and CART.

High watermark setup names (HWSNAME)

Use the HWSNAME statement to define which device names are subsets of other device names. You must specify all applicable varieties. Therefore, HWSNAME coding is not easily understood. To clarify the subject, we introduce only the rules and coding and discuss the details in the configuration examples. The statement is:

HWSNAME, TYPE=(*groupname*, {*altname*})

Note the following explanations:

- ▶ *groupname*: Specifies a device type valid for a high watermark setup.
- ▶ *altname*: Specifies a list of valid user-supplied or IBM-supplied device names. These are alternate units to be used in device selection.

Consider the following example:

```
HWSNAME, TYPE=(LDGW3495,LDGF4001,LDGF4006,LDG359J,LDG359K,LDG359M,  
LDJF4001,LDKF4001,LDKF4006)
```

The rules for LDG HWSNAME statements are:

- ▶ The complex-wide library name, LDGW3495, must include all other LDG names as alternates.
- ▶ The library-specific name must include all LDG names for the corresponding library as alternates. When all tape devices of a type within the complex are within a single 3494 Tape Library, the complex-device type name must also be included as an alternate name.
- ▶ The complex-wide device type name must include all library-specific device type names. When all devices of one type in the complex are within a single TS3500 Tape Library, the complex-wide device type name is equivalent to that library name. In this case, you need to also specify the library name as an alternate.
- ▶ The library-specific device type name must be included. Alternate names can be specified as follows:
 - When all drives within the TS3500 Tape Library have the same device type, the library-specific device type name is equivalent to the library name. In this case, you need to specify the library-specific name as an alternate.
 - When these drives are the only drives of this type in the complex, the complex-wide device type name is equivalent to the library-specific device type name.

Make sure that all valid alternate names are specified.

First configuration example

The first example includes different native tape drives in two separate Tape Libraries. Figure D-3 shows a JES3 complex with two TS3500 Tape Libraries attached to it. Library 1 has a LIBRARY-ID of F4001 and a mix of 3592-J1A and 3592-E05 drives installed. Library 2 has a LIBRARY-ID of F4006 and only 3592-E05 models installed. The 3592-E05 drives are not encryption-enabled.

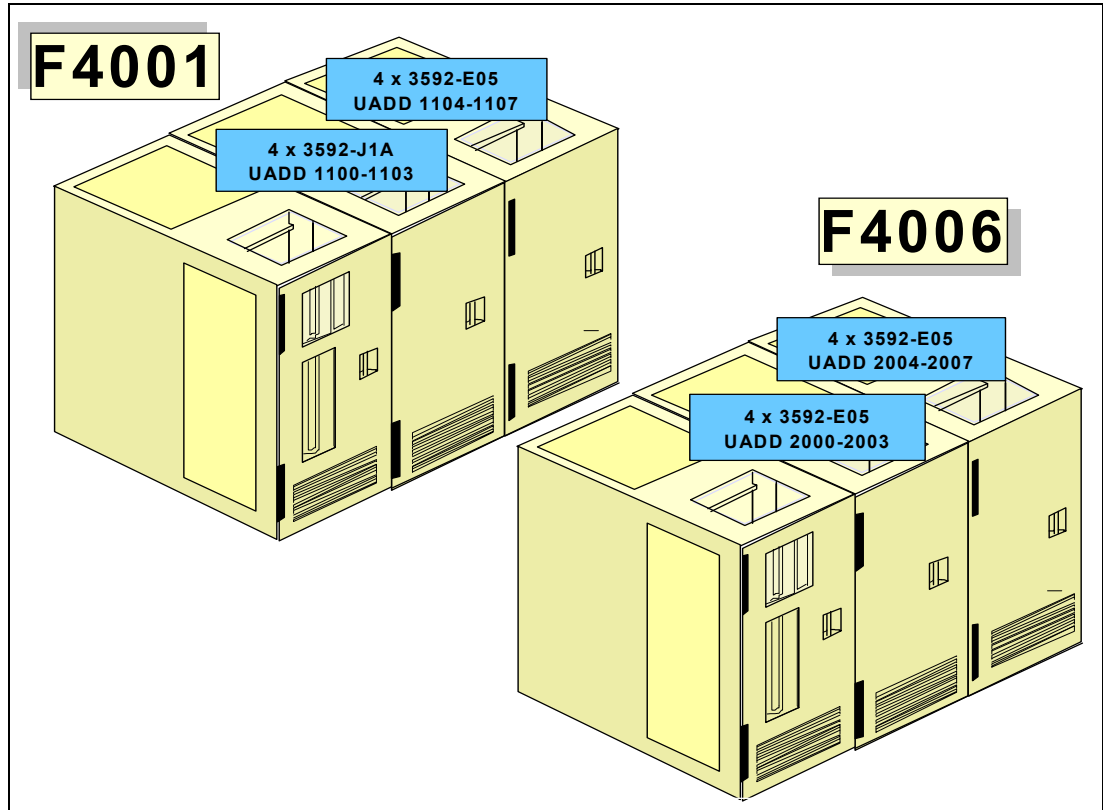


Figure D-3 First JES3 configuration example

LDG definitions necessary for the first example

Table D-3 shows all the LDG definitions needed in HCD. There are a total of eight esoterics to define.

Table D-3 LDG definitions for the first configuration example

LDG definition	Value of LDG	Explanation
Complex-wide name	LDGW3495	Standard name, appears once
Library-specific name	LDGF4001 LDGF4006	One definition for each library
Complex-wide device type	LDG359J LDG359K	One definition for each installed device type: Represents the 3592-J1A devices Represents the 3592-E05 devices
Library-specific device type	LDJF4001 LDKF4001 LDKF4006	One definition for each device type in each library: Represents the 3592-J1A in library F4001 Represents the 3592-E05 in library F4001 Represents the 3592-E05 in library F4006

Device statements needed for this configuration

In our examples, we use a naming convention for XTYPE that contains the library (LB1, LB2) in the first three digits and then the device type (see Figure D-4). A naming convention for XTYPE is not mandatory, but it makes it easier to use the JES3 INISH deck.

```
*/ Devices 3592-J1A and 3592-E05 in Library 1 ...../*  
  
DEVICE,XTYPE=(LB13592J,CA),XUNIT=(1000,*ALL,,OFF),numdev=4  
DEVICE,XTYPE=(LB13592K,CA),XUNIT=(1104,*ALL,,OFF),numdev=4  
  
*/ Devices 3592-E05 Encryption-Enabled in Library 2 ...../*  
  
DEVICE,XTYPE=(LB23592L,CA),XUNIT=(2000,*ALL,,OFF),numdev=8
```

Figure D-4 First configuration example: Device type definition sample

SETNAME statements needed for this configuration

Figure D-5 includes all the SETNAME statements for the first configuration example.

```
SETNAME,XTYPE=(LB13592J,CA),NAMES=(LDGW3495,LDGF4001,LDG359J,LDJF4001)  
  
SETNAME,XTYPE=(LB13592K,CA),NAMES=(LDGW3495,LDGF4001,LDG359K,LDKF4001)  
  
SETNAME,XTYPE=(LB23592L,CA),NAMES=(LDGW3495,LDGF4006,LDG359K,LDKF4006)
```

Figure D-5 First configuration example: SETNAME definition sample

You need three SETNAME statements, because you have:

- ▶ One library with two different device types = Two SETNAME statements
- ▶ One library with one device type = One SETNAME statement

Note: For definition purposes, encryption-enabled and non-encryption-enabled drives are considered two different device types. In the first example, all 3592 Tape Drives are not encryption-enabled.

HWSNAME statement needed for this configuration

The HWSNAME definition is tricky and needs your attention, which is why we explain every statement shown in Figure D-6. If you are not experienced in JES3, read carefully through the explanation.

```
HWSNAME,TYPE=(LDGW3495,LDGF4001,LDGF4006,LDG359J,LDG359K,LDJF4001,LDKF4001,LDKF4006)1
HWSNAME,TYPE=(LDGF4001,LDJF4001,LDKF4001,LDG359J)2
HWSNAME,TYPE=(LDGF4006,LDKF4006)3
HWSNAME,TYPE=(LDJF4001,LDG359J)4
HWSNAME,TYPE=(LDG359J,LDJF4001)5
HWSNAME,TYPE=(LDG359K,LDKF4001,LDGF4006,LDKF4006)6
```

Figure D-6 HWSNAME definition sample

The following numbers, which correspond to the numbers in Figure D-6, explain the statements in the sample:

1. All LDG definitions are a subset of the complex-wide name.
2. LDG359J is a subset of library F4001 (LDGF4001), because the other library only has 3592-E05 installed.
3. All 3592-E05s in library F4006 (LDKF4006) are a subset of library F4006. LDG359K will not be specified, because there are also 3592-E05s installed in the other library.
4. All 3592-J1As (LDG359J) are a subset of the 3592-J1A in library F4001, because no other 3592-J1As are installed.
5. All 3592-J1As in library F4001 (LDJF4001) are a subset of 3592-J1A, because no other 3592-J1As are installed.
6. All 3592-E05s in library F4001 (LDKF4001) are a subset of 3592-E05. LDGF4006 (the entire library with the ID F4006) is a subset of 3592-E05, because only 3592-E05s are installed in this library.

Second configuration example

Figure D-7 on page 829 shows a JES3 configuration with three TS3500 Tape Libraries attached to it. Library 1 has a LIBRARY-ID of F4001, a mix of 3592-J1A and 3592-E05 drives that are not encryption-enabled, and one TS7700 Virtualization Engine of a Multi Cluster TS7700 Virtualization Engine Grid (Distributed Library) installed. The Multi Cluster TS7700 Virtualization Engine Grid has a Composite Library LIBRARY-ID of 47110.

Library 2 has a LIBRARY-ID of F4006 and a mix of encryption-enabled and non-encryption-enabled 3592-E05 drives installed, which is also the reason why you might need to split a string of 3592-E05 drives. Library 2 is also the second Distributed Library for the multi-cluster grid with Composite Library LIBRARY-ID 47110.

Library 3 has a LIBRARY-ID of 22051 and only a TS7700 Virtualization Engine installed with a Composite Library LIBRARY-ID of 13001.

Figure D-7 does not show the actual needed configuration for the TS7700 Virtualization Engine configurations regarding the numbers of frames, controllers, and the back-end drives. We have limited the drives and frames displayed to those drives and frames actually needed for the host definitions.

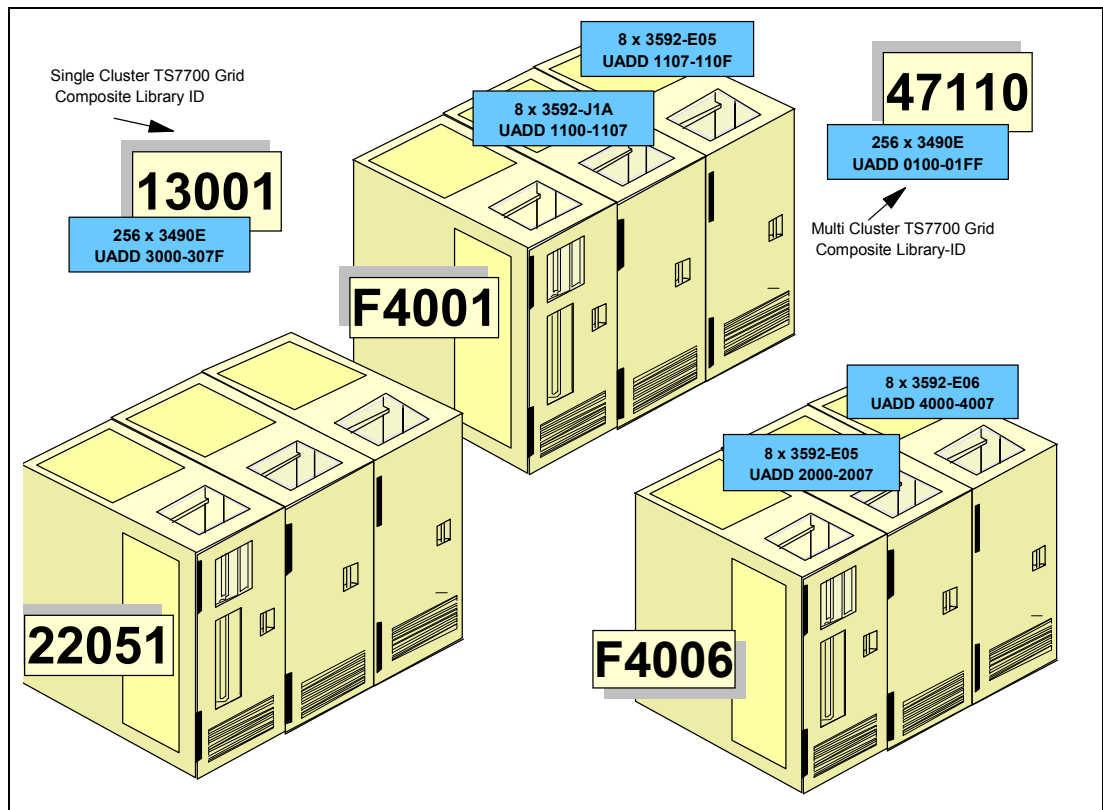


Figure D-7 Second JES3 configuration example

LDG definitions needed for the second configuration example

Table D-4 shows all the LDG definitions needed in the hardware configuration definition of the second configuration example.

Table D-4 LDG definitions for the second configuration example

LDG definition	Value for LDG	Explanations
Complex-wide name	LDGW3495	Standard name, which appears once.
Library-specific name	LDGF4001 LDGF4006 LDG13001 LDG47110	One definition for each library and for each Single Cluster TS7700 Virtualization Engine Grid. For a Single or Multi Cluster TS7700 Virtualization Engine Grid, only the Composite Library LIBRARY-ID is specified.
Complex-wide device type	LDG3490E LDG359J LDG359K LDG359L LDG359M	One definition for each installed device type: <ul style="list-style-type: none"> ▶ Represents the devices in Multi Cluster TS7700 Virtualization Engine Grid ▶ Represents the 3592-J1A ▶ Represents the 3592-E05 ▶ Represents the 3592-E05 with Encryption ▶ Represents the 3592-E06
Library-specific device type	LDE13001 LDE47110 LDJF4001 LDKF4001 LDLF4006 LDMF4006	One definition for each device type in each library, except for the Multi Cluster TS7700 Virtualization Engine Grid: <ul style="list-style-type: none"> ▶ Represents the virtual drives in the Single Cluster TS7700 Virtualization Engine Grid in library 22051 ▶ Represents the virtual drives in the Multi Cluster TS7700 Virtualization Engine Grid in libraries F4001 and F4006 ▶ Represents the 3592-J1A in library F4001 ▶ Represents the 3592-E05 in library F4001 ▶ Represents the encryption-enabled 3592-E05 in library F4006 ▶ Represents the 3592-E06 in library F4006

Device statement needed for the second configuration example

Figure D-8 shows all of the device statements for the second configuration example. The comment statements describe to which library the devices belong.

```
*/ Devices 3592-J1A and 3592-E05 in Library F4001 ...../*  
  
DEVICE,XTYPE=(LB13592J,CA),XUNIT=(1100,*ALL,,OFF),numdev=8  
DEVICE,XTYPE=(LB13592K,CA),XUNIT=(1107,*ALL,,OFF),numdev=8,  
  
*/ Devices 3592-E06 and 3592-E05 in Library F4006...../*  
  
DEVICE,XTYPE=(LB2359M,CA),XUNIT=(4000,*ALL,,OFF),numdev=8  
DEVICE,XTYPE=(LB2359L,CA),XUNIT=(2000,*ALL,,OFF),numdev=8  
  
*/ Devices Single Cluster TS7700 Grid in library 22051 ...../*  
  
DEVICE,XTYPE=(LB3GRD1,CA),XUNIT=(3000,*ALL,,OFF),numdev=256  
  
*/ Devices Multi Cluster TS7700 grid in libraries F4001 and F4006...../*  
ADDRSORT=NO  
  
DEVICE,XTYPE=(LB12GRD,CA),XUNIT=(0110,*ALL,S3,OFF)  
DEVICE,XTYPE=(LB12GRD,CA),XUNIT=(0120,*ALL,S3,OFF)  
DEVICE,XTYPE=(LB12GRD,CA),XUNIT=(0130,*ALL,S3,OFF)  
DEVICE,XTYPE=(LB12GRD,CA),XUNIT=(0140,*ALL,S3,OFF)  
DEVICE,XTYPE=(LB12GRD,CA),XUNIT=(0111,*ALL,S3,OFF)  
DEVICE,XTYPE=(LB12GRD,CA),XUNIT=(0121,*ALL,S3,OFF)  
DEVICE,XTYPE=(LB12GRD,CA),XUNIT=(0131,*ALL,S3,OFF)  
DEVICE,XTYPE=(LB12GRD,CA),XUNIT=(0141,*ALL,S3,OFF)  
;;;;;;  
DEVICE,XTYPE=(LB12GRD,CA),XUNIT=(01FF,*ALL,S3,OFF)
```

Figure D-8 DEVICE statements for the second configuration example

Consideration: If you code NUMDEV in a PTP VTS environment, the workload balancing from the CX1 controllers does not work. Therefore, you must specify each device as a single statement and specify ADDRSORT=NO to prevent JES3 from sorting them.

The same restriction applies to the virtual devices of the clusters of a multi-cluster grid configuration; if you want to balance the workload across the virtual devices of all clusters, do not code the NUMDEV parameter.

SETNAME statements needed for the second configuration example

Figure D-9 includes all the necessary SETNAME statements for the second configuration example.

```
SETNAME,XTYPE=LB1359J,NAMES=(LDGW3495,LDGF4001,LDG359J,LDJF4001)  
SETNAME,XTYPE=LB1359K,NAMES=(LDGW3495,LDGF4001,LDG359K,LDKF4001)  
SETNAME,XTYPE=LB2359L,NAMES=(LDGW3495,LDGF4006,LDG359L,LDKF4006)  
SETNAME,XTYPE=LB2359M,NAMES=(LDGW3495,LDGF4006,LDG359M,LDMF4006)  
SETNAME,XTYPE=LB3GRD1,NAMES=(LDGW3495,LDG13001,LDG22051,LDG3490E,LDE22051,LDE13001)  
SETNAME,XTYPE=LB12GRD,NAMES=(LDGW3495,LDG47110,LDG3490E,LDE47110)
```

Figure D-9 SETNAME statement values for the second example

High watermark setup name statements for the second example

Figure D-10 shows the high watermark setup name statements for the second configuration example.

```
HWSNAME,TYPE=(LDGW3495,LDGF4001,LDGF4006,LDG13001,LDG47110,LDG3490E,  
              LDG359J,LDG359K,LDG359L,LDG359M,LDE13001,LDE47110,LDJF4001,  
              LDKF4001,LDLF4006,LDMF4006)  
HWSNAME,TYPE=(LDGF4001,LDJF4001,LDKF4001)  
HWSNAME,TYPE=(LDGF4006,LDLF4006,LDMF4006)  
HWSNAME,TYPE=(LDG47110,LDE47110)  
HWSNAME,TYPE=(LDG13001,LDE13001)  
HWSNAME,TYPE=(LDG3490E,LDE47110,LDE13001)  
HWSNAME,TYPE=(LDG359J,LDJF4001)  
HWSNAME,TYPE=(LDG359K,LDKF4001)  
HWSNAME,TYPE=(LDG359L,LDLF4006)  
HWSNAME,TYPE=(LDG359M,LDMF4006)
```

Figure D-10 High watermark setup statements for the second example

Processing changes

Although no JCL changes are required, a few processing restrictions and limitations are associated with using the TS3500 Tape Library in a JES3 environment:

- ▶ JES3 spool access facility (SPAF) calls are not used.
- ▶ Two calls, one from the prescan phase and the other call from the locate processing phase, are made to the new DFSMS/MVS support module, as shown in Figure D-11 on page 833.
- ▶ The main device scheduler (MDS) processing phases, system select and system verify, are not made for tape data sets.
- ▶ The MDS verify phase is bypassed for TS3500 Tape Library mounts, and mount processing is deferred until job execution.

Figure D-11 on page 833 shows the JES3 processing phases for C/I and MDS. The processing phases shown include the support for system-managed direct access storage device (DASD) data sets.

The major differences between TS3500 Tape Library deferred mounting and tape mounts for non-library drives are:

- ▶ Mounts for non-library drives by JES3 are only for the first use of a drive. Mounts for the same unit are issued by z/OS for the job. All mounts for TS3500 Tape Library drives are issued by z/OS.
- ▶ If all mounts within a job are deferred, because there are no non-library tape mounts, that job is not included in the setup depth parameter (SDEPTH).
- ▶ MDS mount messages are suppressed for the TS3500 Tape Library.




Figure D-11 JES3 C/I and MDS processing phases

JES3/DFSMS processing

DFSMS is called by the z/OS interpreter to:

- Update the scheduler work area (SWA) for DFSMS tape requests
- Call automatic class selection (ACS) exits for construct defaults

DFSMS/MVS system-managed tape devices are not selected by using the UNIT parameter in the JCL. For each DD request requiring a TS3500 Tape Library unit, a list of device pool names is passed, and from that list, an LDG name is assigned to the DD request, which results in an LDG name passed to JES3 MDS for that request. Device pool names are never known externally.

Selecting UNITNAMEs

For a DD request, the LDG selection is based on the following conditions:

- ▶ When all devices in the complex are eligible to satisfy the request, the complex-wide LDGW3495 name is used.
- ▶ When the list of names contains names of all devices of one device type in the complex, the corresponding complex-device type name (for example, LDG3490E) must be used.
- ▶ When the list of names contains all subsystems in one TS3500 Tape Library, the library-specific LDG name (in our examples, LDGF4001, LDGF4006, and so on) is used.
- ▶ When the list contains only subsystems for a specific device type, within one TS3500 Tape Library, the LDG device type library name (in our example, LDKF4001, and so on) is used.

New or modified data sets

For new data sets, ACS directs the allocation by providing Storage Group, Storage Class, and Data Class. When the Storage Group specified by ACS is defined in the active DFSMS configuration as a tape Storage Group, the request is allocated to a TS3500 Tape Library tape drive.

DFSMS-managed DISP=MOD data sets are assumed to be new update locate processing. If a catalog locate determines that the data set is *old* by the VOLSER specified, a new LDG name is determined based on the rules for old data sets.

Old data sets

Old data set allocations are directed to a specific TS3500 Tape Library when the volumes containing the data set are located within that TS3500 Tape Library. For old data sets, the list is restricted to the TS3500 Tape Library that contains the volumes.

DFSMS catalog processing

JES3 catalog processing determines all of the catalogs required by a job and divides them into two categories:

- ▶ DFSMS-managed user catalogs
- ▶ JES3-managed user catalogs

DFSMS catalog services, a subsystem interface call to catalog locate processing, is used for normal locate requests. DFSMS catalog services is invoked during locate processing. It invokes SVC 26 for all existing data sets when DFSMS is active. Locates are required for all existing data sets to determine whether they are DFSMS-managed, even if VOL=SER= is present in the DD statement. If the request is for an old data set, catalog services determine whether it is for a library volume. For multivolume requests that are system-managed, a check is made to determine whether all volumes are in the same library.

DFSMS VOLREF processing

DFSMS VOLREF services are invoked during locate processing if VOL=REF= is present in a DD statement for each data set that contains a volume reference to a cataloged data set. DFSMS VOLREF services determine whether the data set referenced by a VOL=REF= parameter is DFSMS-managed. Note that VOL=REF= now maps to the same Storage Group for a DFSMS-managed data set, but not necessarily to the same volume. DFSMS VOLREF services also collect information about the job's resource requirements.

The TS3500 Tape Library supports the following features:

- ▶ Identifies the DDs that are TS3500 Tape Library-managed mountable entries
- ▶ Obtains the associated device pool names list
- ▶ Selects the LDG that best matches the names list
- ▶ Provides the LDG name to JES3 for setup
- ▶ Indicates to JES3 that the mount is deferred until execution

Fetch messages

While TS3500 Tape Library cartridges are mounted and demounted by the library, fetch messages to an operator are unnecessary and can be confusing. With this support, all fetch messages (IAT5110) for TS3500 Tape Library requests are changed to be the non-action informational USES form of the message. These messages are routed to the same console destination as other USES fetch messages. The routing of the message is based on the UNITNAME.

JES3 allocation and mounting

JES3 MDS controls the fetching, allocation, and mounting of the tape volumes requested in the JCL for each job to be executed on a processor. The scope of MDS tape device support is complex-wide, unlike z/OS job resource allocation, whose scope is limited to one processor. Another difference between JES3 MDS allocation and z/OS allocation is that MDS considers the resource requirements for all the steps in a job for all processors in a loosely coupled complex. z/OS allocation considers job resource requirements one step at a time in the executing processor.

MDS processing also determines which processors are eligible to execute a job based on resource availability and connectivity in the complex.

z/OS allocation interfaces with JES3 MDS during step allocation and dynamic allocation to get the JES3 device allocation information and to inform MDS of resource deallocations. z/OS allocation is enhanced by reducing the allocation path for mountable volumes. JES3 supplies the device address for the TS3500 Tape Library allocation request through an SSI request to JES3 during step initiation when the job is executing under the initiator. This support is not changed from previous releases.

DFSMS/MVS and z/OS provide all of the TS3500 Tape Library support except for the interfaces to JES3 for MDS allocation and processor selection.

JES3 MDS continues to select tape units for the TS3500 Tape Library. MDS no longer uses the UNIT parameter for allocation of tape requests for TS3500 Tape Library requests. DFSMS/MVS determines the appropriate LDG name for the JES3 setup from the Storage Group and Data Class assigned to the data set, and replaces the UNITNAME from the JCL with that LDG name. Because this action is done after the ACS routine, the JCL-specified UNITNAME is available to the ACS routine. This capability is used to disallow JCL-specified LDG names. If LDG names are permitted in the JCL, the associated data sets must be in a

DFSMS tape environment. Otherwise, the allocation fails, because an LDG name restricts allocation to TS3500 Tape Library drives that can be used only for system-managed volumes.

Note: An LDG name specified as a UNITNAME in JCL can be used only to filter requests within the ACS routine. Because DFSMS/MVS replaces the externally specified UNITNAME, it cannot be used to direct allocation to a specific library or library device type.

All components within z/OS and DFSMS/MVS request tape mounting and demounting inside a TS3500 Tape Library. They call a DFP service, library automation communication services (LACS), instead of issuing a write to operator (WTO), which is done by z/OS allocation, so all mounts are deferred until job execution. The LACS support is called at that time.

MDS allocates an available drive from the available unit addresses for LDGW3495. It passes that device address to z/OS allocation through the JES3 allocation SSI. At data set OPEN time, LACS are used to mount and verify a scratch tape. When the job finishes with the tape, either CLOSE or deallocation issues a demount request through LACS, which removes the tape from the drive. MDS does normal breakdown processing and does not need to communicate with the TS7700.

Multi-cluster grid considerations

With JES3, the Device Allocation Assist (DAA) function is not applicable, because JES3 handles tape device allocation itself and does not use standard MVS allocation algorithms. In a multi-cluster grid configuration, careful planning of the Copy Consistency Points and the Copy Override Settings can help avoid unnecessary copies in the Grid and unnecessary traffic on the Grid links.

Consider the following aspects, especially if you are using a multi-cluster grid with more than two clusters and if not all clusters contain copies of all logical volumes:

- Retain Copy Mode setting

If you do not copy logical volumes to all of the clusters in the grid, JES3 may, for a specific mount, select a drive that does not have a copy of the logical volume. If Retain Copy Mode is not enabled on the mounting cluster, an unnecessary copy might be forced according to the Copy Consistency Points that are defined for this cluster in the Management Class.

- Copy Consistency Point (CCP)

Copy Consistency Point has one of the largest influences on which cluster's cache is used for a mount. The CCP of Rewind/Unload (R) takes precedence over a CCP of Deferred (D). For example, assuming each cluster has a consistent copy of the data, if a virtual device on Cluster 0 is selected for a mount and the CCP is RD then the CL0 cache will be selected for the mount. However, if the CCP is DR, then CL1's cache will be selected.

For workload balancing, consider specifying "DD" rather than "RD". This will more evenly distribute the workload to both clusters in the grid.

If the CCPs for a cluster are DD, then other factors are used to determine which cluster's cache to use. The "Prefer local cache for fast ready mount requests" and "Prefer local cache for non-fast ready mounts" overrides will cause the cluster that received the mount request to be the cache used to access the volume.

- ▶ Cluster Families

If there are more than two clusters in the grid, consider defining Cluster Families. Especially in multi-site configurations with larger distances between the sites, defining one cluster family per site can help reduce the grid link traffic between both sides.

- ▶ Copy Override settings

Remember that all Copy Override settings such as “Prefer local cache for fast ready mount requests” and “Prefer local cache for non-fast ready mounts” always apply to the entire cluster, where the CCPs defined in the Management Class can be different and tailored according to workload requirements.

You can find detailed information about these settings and other workload considerations in Chapter 4, “Hardware implementation” on page 175 and Chapter 9, “Performance and monitoring” on page 589



IBM Virtualization Engine TS7700 statistical record format

This appendix is based on the *IBM Virtualization Engine TS7700 Series Statistical Data Format Version 1.6 White Paper* and provides the definition of the TS7700 Virtualization Engine statistics records. Version 1.7 has no differences from 1.6 document. The original white paper can be found at the IBM Techdocs website at the following address:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100829>

Introduction

The IBM TotalStorage 3494 Virtual Tape Server (VTS) revolutionized the way IBM eServer™ zSeries® customers utilized their tape resources. To help the customer monitor the performance of the VTS, various statistics were developed. These statistics appeared in two forms

- ▶ Hourly records sent to the hosts, known as SMF94 records
- ▶ Periodic real-time statistics available through the Tape Library's Web specialist or in an XML file for use by the Total Productivity Center (TPC)

For the next generation of VTS, the TS7700 Virtualization Engine, the statistics design has been revisited. The useful statistics from the VTS have been retained and new statistics relevant for the TS7700 Virtualization Engine have been added. Also, both point-in-time (PIT) and historical statistics are recorded. The point-in-time records present data from the most recent interval and provide speedometer like statistics. The historical statistics provide statistics where historical trends can be observed. These statistics are available to a host through the Bulk Volume Information Retrieval (BVIR) facility. See the *IBM Virtualization Engine TS7700 Series Bulk Volume Information Retrieval Function User's Guide* for more information concerning the BVIR facility as it relates to the TS7700 Virtualization Engine statistics. You can find this guide on the web at the following address:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/Techdocs>

Note that the statistical records described within this white paper may contain records, fields, or field expansions that are not supported in the current release of the product. In general, those fields have been identified in the field descriptions, but you should also validate the availability of the function with your IBM representative.

Changes introduced to Bulk Volume Information Retrieval (BVIR) by different Licensed Internal Code levels up to R1.7 are listed below:

- ▶ V 1.0
 - Initial release.
- ▶ V 1.1
 - Updated text to note that JB media is now supported.
- ▶ V1.5
 - Updated with new statistics delivered in Release 1.5
 - Various clarifications
- ▶ V1.6
 - Removed unpopulated fields and changed some fields to reserved:
 - hNode Grid Historical Record
 - hNode Grid Point-In-Time (PIT) Record
 - Various clarifications added, including:
 - “Awaiting Replication to Available Clusters” field in hNode HSM Historical Record
 - The maximum number of clusters in a Grid is now four.
- ▶ V1.7
 - No updates

Overview

This appendix outlines the information that the TS7700 Virtualization Engine returns for the request of statistical information. The information is divided into point-in-time (PIT) and historical (HIS) data:

- ▶ The PIT information is intended to supply information about what the system is doing the instant the request is made to the system. This information is not persistent on the system, and is updated on a 15-second interval. This information focuses on the individual components of the system and their current activity.
- ▶ The HIS information is intended to help with capacity planning, and tracking system use over an extended period of time. The information focuses more on the system as a whole, and the movement of data through the system. This information is kept on the system for 90 days, and is collected on a 15-minute interval basis.

With the support for Cluster-Paired configured systems, some of the HSM records account for the fact that the two clusters are sharing the back-end tape drive, library, and media resources. For the PIT HSM records, an indicator specifies which cluster is in control of the physical drive. While both clusters have access to each drive, and therefore create records for all of them, only one cluster actively controls that drive at a time. Historical records for clusters in a pair configuration will report activity on a per cluster basis (that is, individually).

Record types and sizes

Table E-1 describes the record types, the length of each record, and the number of records produced.

Table E-1 Record types and sizes

Data type (hexadecimal)	Description	Record length (in bytes)	Number of records
x01	vNode Virtual Device Point-in-Time Record	96 + (number of virtual devices in this cluster x 32)	1 per vNode
x02	vNode Adapter Point-in-Time Record	384	1 per vNode
x20	vNode Virtual Device Historical Record	192	1 per vNode
x21	vNode Adapter Historical Record	384	1 per vNode
x10	hNode HSM Point-in-Time Record	96 + (number of physical libraries attached to this cluster x 1568)	1 per hNode
x11	hNode Grid Point-in-Time Record	96 + (number of clusters in the Grid x 129)	1 per hNode
x30	hNode HSM Historical Record	1152	1 per hNode
x31	Reserved		

Data type (hexadecimal)	Description	Record length (in bytes)	Number of records
x32	hNode Library Historical Record	8256	1 per physical library attached to a Cluster
x33	hNode Grid Historical Record	96 + (number of clusters in the Grid x 256)	1 per hNode

vNode Virtual Device Point-in-Time (PIT) record

The vNode Point-In-Time record has the following nested structure:

- ▶ Header
- ▶ General Information Container
- ▶ Virtual Device Container
 - Virtual Device 0 information
 - Virtual Device 1 information
 - Virtual Device xxx information

Table E-2 shows the layout of vNode Virtual Device Point-in-Time record.

Table E-2 vNode Virtual Device Point-in-Time record

Bytes	Name	Description	When data is sampled or updated
0-1	Length	This 2 byte hexadecimal field contains the length of this record. The length includes these 2 bytes.	
2	Version	This 1 byte hexadecimal field contains the version of the data presented in this record. Initially the version is set to x02.	
3	Data Type	This 1 byte hexadecimal field indicates the type of data contained in this record. For this record, the value is set to x01, indicating this is a vNode Virtual Device Point-In-Time record.	
4	Node ID	This 1 byte hexadecimal field indicates the vNode ID that this interval's data represents. Valid values are x00 – x0F.	
5	Cluster ID	This 1 byte hexadecimal field indicates the Cluster ID of which this vNode is a part. Valid values are x00 – x07.	
6-7	Interval Duration	This 2 byte hexadecimal field indicates the interval in seconds that this interval's data was taken over.	
8-11	Time Stamp	This 4 byte hexadecimal field indicates the end time of the interval this data was taken over. This value is the time in seconds since the beginning of the Epoch (00:00:00 UTC, January 1, 1970).	
12-15	Machine Type	This 4 byte EBCDIC field contains this node's machine type. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "3957".	
16-18	Machine Model	This 3 byte EBCDIC field contains this node's machine model. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "V06".	
19-26	Machine Serial Number	This 8 character EBCDIC field contains the serial number of this node. This field is left justified and padded with EBCDIC blanks. The format is XX-YYYYY, where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed.	

Bytes	Name	Description	When data is sampled or updated
27-34	VE Code Level	This 8 byte hexadecimal field contains the code level of the TS7700 Virtualization Engine. The 8 bytes are actually four 2 byte fields. Each 2 byte field represents a portion of the code level. The code level is expressed as Version. Release. Modification.Fix in a decimal form. For example, the code level of 8.0.0.104 would be represented in the 8 bytes as x0008000000000068.	
35-39	Grid Library Sequence Number	This 5 character EBCDIC field contains the Library Sequence Number of the Grid (Composite) library.	
40-44	Distributed Library Sequence Number	This 5 character EBCDIC field contains the Distributed Library Sequence Number for this Distributed Library ID.	
45-63	Reserved	All bytes are set to x00.	
vNode PIT General Information Container Bytes 64-95 These fields provide information concerning the configuration of the vNode whose data is being reported.			
64	Node State	This 1 byte hexadecimal field indicates the state of the vNode at the end of this interval. Possible values are: x00: Offline x01: Online x02: Going Offline x03: Going Online x04-xFE: Reserved xFF: Node not operational	
65-66	Configured Maximum Throughput	This 2 byte hexadecimal field contains the maximum throughput for this vNode. The value is expressed in MBps (1 MB = 1024x bytes). This field is set to x0000 if there is no restriction for the maximum throughput. This is the value at the end of the interval.	Set to x0000 for first release.
67-68	Installed Virtual Devices	This 2 byte hexadecimal field indicates the number of installed virtual devices in this vNode. This field can be used to determine how many Virtual Device containers will be attached to this record. This is the value at the end of the interval.	
69-95	Reserved	All bytes set to x00.	

Bytes	Name	Description	When data is sampled or updated
Virtual Device Container Bytes 96 and up (Number of Installed Virtual Devices x 32 bytes/device). For example, if there are 128 installed virtual devices, there will be 128 sets of data with 32 bytes, each of which totals 4096 bytes. There is a maximum of 256 virtual devices per vNode. This next segment of the record contains one set of data for each virtual device installed in the vNode as defined in bytes 67-68 above. Each set of data contains 32 bytes. The following fields define the 32 bytes of data and are numbered starting with byte 0. The first virtual device's data can be found in bytes 96-127, the second device's data can be found in bytes 128-159, and so on.			
0-9	Mounted Volume	<p>This 10 byte EBCDIC field contains the VOLSER of the logical volume, if any, that is currently mounted in the virtual device or was most recently mounted in the device.</p> <p>For device mount states that indicate a mount in progress or is mounted (x01, x02, x03, x06; see byte 11), this field will contain the VOLSER of the volume that is in the process of being mounted.</p> <p>For device mount states that indicate a device is not mounted or in the process of being mounted (x00, x04, x05; see byte 11), this field will contain the VOLSER of the last successfully mounted volume, if any, or will be filled with EBCDIC blanks.</p> <p>This field is left-justified and padded with EBCDIC blanks.</p> <p>This is the value at the end of the interval.</p>	Updated when the logical volume changes in the virtual device.
10	Cluster Access Point	<p>This 1 byte hexadecimal field indicates the Cluster ID that is sourcing or has most recently sourced the logical volume for a mount. In the case of the most recently sourced cluster, the current device mount state will indicate "Device unloaded, failed or cancelled" (see byte 11). Valid values for this field are x00 – x07.</p> <p>The subsystem has the ability to access a logical volume in any cluster from any cluster, where a logical volume is sourced from is based on a set of criteria including volume consistency, access policies, and so forth.</p> <p>This is the value at the end of the interval.</p>	Updated for each virtual mount.

Bytes	Name	Description	When data is sampled or updated
11	Device Mount State	<p>This 1 byte hexadecimal field indicates the mount state of the virtual device.</p> <p>Valid values are as follows:</p> <ul style="list-style-type: none"> ▶ x00: Device unloaded. This indicates the previous mount completed successfully and the virtual device is now unloaded. This is also reported when the virtual device has not mounted a volume and is not in the process of performing a mount. ▶ x01: Mount request being processed (DE owed, PLF received). ▶ x02: Mount accepted and DE given (Initial status received from Library Manager). ▶ x03: Mount in progress (Device mount command received from Library Manager). ▶ x04: Mount failed. This indicates the previous mount failed (error reported by device) and the virtual device is unloaded. ▶ x05: Mount cancelled. This indicates the previous mount request was cancelled before the mount occurred (Demount received before Unload). The virtual device is now unloaded. ▶ x06: Mounted <p>All other values are reserved.</p> <p>For a non-configured or non-installed device, this field will indicate Device unloaded.</p> <p>This is the value at the end of the interval.</p>	Updated when the virtual device mount state changes.

Bytes	Name	Description	When data is sampled or updated
12-13	Device Flags	<ul style="list-style-type: none"> ▶ x8000: The device is ready. It provides information as to the mount state of the virtual drive. If the bit is set, then a virtual volume is loaded into the drive and is ready for host I/O activity. ▶ x4000: The device is write-protected. This flag indicates the logical write protect state of the virtual drive/virtual volume. This prevents any modifications to a virtual volume loaded into this drive. ▶ x2000: Write data is in the buffer. This flag indicates that the control unit has write data for this virtual volume held in the controller's memory that is not committed to the virtual device. This is normal for streaming write data to the device. ▶ x1000: Write mode. The last IO to the device was a write operation. Indicates that the virtual device is in write mode. The device enters write mode upon the first write I/O operation, and stays in this mode until a non-write command (position change, read, and so on) is encountered. ▶ x0800: Volume is in the Logical End of Partition (LEOP). Indicates that the current block position of the virtual tape is in the LEOP region of the virtual volume, and the host should be performing end of volume processing on the virtual volume. ▶ x0400: Volume is at Beginning of Tape (BOT). Indicates the virtual volume loaded is currently at the beginning of the virtual volume. This would be the tape position after a mount request, or if the host repositioned the volume back to block 0. ▶ x0200: The device is fenced. The control unit presents the unit check status with the associated sense data, indicating ERA 47, Volume Fenced, if a condition has occurred that has resulted in the loss of volume integrity because of lost positioning or assignment. The control unit prevents further access to the tape volume by generating deferred unit checks with associated sense data, indicating ERA 47, Volume Fenced, for all eligible commands until the condition is reset or until the cartridge is unloaded. The original condition that subsequently caused ERA 47, Volume Fenced, to be presented has already been indicated by a previous unit check and associated sense data. ▶ x0100: The device is in stand-alone mount mode. This indicates that the virtual volume in this virtual drive was loaded as part of an operator request through the management interface, and not through a connected host. Stand-alone mounts are required to be able to IPL the operating system image from a virtual tape device, without the operating system having to have tape mount capabilities. <p>All other bits are reserved. For a non-configured or non-installed device, this field will be set to x0000. This is the value at the end of the interval.</p>	Updated when the virtual device mount state changes.

Bytes	Name	Description	When data is sampled or updated
14	Buffer CCR Conditions	<p>This 1 byte hexadecimal field contains the number of times the virtual device had to use Channel Command Retry (CCR) on the channel because of a buffer condition during this interval. If more than 255 buffer CCRs occur during the interval, this field will indicate 255 buffers CCRs (xFF).</p> <p>This field can be used in conjunction with the Device Flags field, Write mode bit, to determine if the device was in write mode (buffer full condition) or read mode (buffer empty condition) during the interval.</p> <p>This value is reset to 0 at the beginning of the interval.</p> <p>This field is not accurate at this time.</p>	The count is incremented when a buffer CCR condition occurs.
15-16	Channel Bytes Read	<p>This 2 byte hexadecimal field contains the number of bytes transferred from the virtual device to the channel for a read from this device. The value is reported in increments of 100 KB (100 x 1024). Any residual data will cause the value to be rounded up to the next higher value.</p> <p>This value is reset to 0 at the beginning of the interval</p>	The count is incremented for each read of a logical volume's data.
17-18	Channel Bytes Written	<p>This 2 byte hexadecimal field contains the number of bytes transferred from the channel to the virtual device for a write to this device. The value is reported in increments of 100 KB (100 x 1024). Any residual data will cause the value to be rounded up to the next higher value. This value is reset to 0 at the beginning of the interval.</p>	The count is incremented for each write of data to a logical volume.
19-31	Reserved	All bytes set to x00.	

vNode Adapter Point-In-Time (PIT) record

The vNode Point-In-Time record has the following nested structure:

- ▶ Header
- ▶ Adapter Container
 - Adapter 0 general information
 - Port 0 information
 - Port 1 information
 - Adapter 1 general information
 - Port 0 information
 - Port 1 information
 - Adapter 2 general information
 - Port 0 information
 - Port 1 information
 - Adapter 3 general information
 - Port 0 information
 - Port 1 information

Table E-3 on page 848 shows the record layout and field descriptions. Figure E-1 on page 850 shows sampling points.

Table E-3 vNode Adapter Point-In-Time record

Bytes	Name	Description	When data is sampled or updated
0-1	Length	This 2 byte hexadecimal field contains the length of this record. The length includes these 2 bytes.	
2	Version	This 1 byte hexadecimal field contains the version of the data presented in this record. Initially the version is set to x02.	
3	Data Type	This 1 byte hexadecimal field indicates the type of data contained in this record. For this record, the value is set to x02, indicating that this is a vNode Adapter Point-In-Time record.	
4	Node ID	This 1 byte hexadecimal field indicates the vNode ID that this interval's data represents. Valid values are x00 – x0F.	
5	Cluster ID	This 1 byte hexadecimal field indicates the Cluster ID of which this vNode is a part. Valid values are x00 – x07.	
6-7	Interval Duration	This 2 byte hexadecimal field indicates the interval in seconds that this interval's data was taken over.	
8-11	Time Stamp	This 4 byte hexadecimal field indicates the end time of the interval this data was taken over. This value is the time in seconds since the beginning of the Epoch (00:00:00 UTC, January 1, 1970).	
12-15	Machine Type	This 4 byte EBCDIC field contains this node's machine type. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "3957".	
16-18	Machine Model	This 3 byte EBCDIC field contains this node's machine model. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "V06".	
19-26	Machine Serial Number	This 8 character EBCDIC field contains the serial number of this node. This field is left justified and padded with EBCDIC blanks. The format is XX-YYYYY, where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed.	
27-34	VE Code Level	This 8 byte hexadecimal field contains the code level of the TS7700 Virtualization Engine. The 8 bytes are actually four 2 byte fields. Each 2 byte field represents a portion of the code level. The code level is expressed as Version.Release.Modification.Fix in a decimal form. For example, the code level of 8.0.0.104 would be represented in the 8 bytes as x0008000000000068.	
35-39	Grid Library Sequence Number	This 5 character EBCDIC field contains the Library Sequence Number of the Grid (Composite) library.	
40-44	Distributed Library Sequence Number	This 5 character EBCDIC field contains the Distributed Library Sequence Number for this Distributed Library ID.	
45-63	Reserved	All bytes set to x00.	

Bytes	Name	Description	When data is sampled or updated
Adapter Container Bytes 64-383 (4 sets of data x 80 bytes/set = 320 bytes) This next set of bytes contains information for up to 4 host bus adapters (HBA). For each adapter, there is data for up to 2 ports on the adapter. Each set of data contains 80 bytes. The following fields define the 80 bytes of data and are numbered starting with byte 0. The first set of data is for adapter 0, the second for adapter 1, and so on. The first HBA's data can be found in bytes 64-143, the second adapter's data can be found in bytes 144-223, and so on.			
0	Adapter Type	This 1 byte hexadecimal field identifies the type of Host Bus Adapter (HBA). The possible values are: <ul style="list-style-type: none"> ▶ x00: No adapter installed. ▶ x01-x08: Reserved. ▶ x09: FICON – 1 Port (Arctic Circle). ▶ x0A: FICON – 2 Port (Yukon). All other values are reserved. This is the value at the end of the interval.	
1	Adapter State	This 1 byte hexadecimal field identifies the current state of the adapter. The possible values are: <ul style="list-style-type: none"> ▶ x00: No adapter is installed. ▶ x01: The adapter is online. ▶ x02: The adapter is offline. ▶ x03: The adapter is not working. ▶ x04: The adapter is reloading itself. ▶ x05: The adapter is in a Check1 condition. All other values are reserved. This is the value at the end of the interval.	
2	HBA Drawer	This 1 byte hexadecimal field indicates in which drawer the HBA is located. The possible values are: <ul style="list-style-type: none"> ▶ x00: The HBA is in the left drawer when looking at the drawers from the back, which is the side the cables plug into. ▶ x01: The HBA is in the right drawer when looking at the drawers from the back, which is the side the cables plug into. All other values are reserved. This is the value at the end of the interval.	
3	HBA Slot Number	This 1 byte hexadecimal field indicates the physical slot number of the HBA within its drawer. This is the value at the end of the interval.	
4-15	Reserved	All bytes set to x00.	
Adapter-Port Container Relative bytes 16 – 79 (2 sets of data x 32 bytes/set = 64 bytes) This next set of bytes contains information for up to 2 ports on the HBA. Each set of data contains 32 bytes. The following fields define the 32 bytes and are numbered starting with byte 0. The first port's data can be found in relative bytes 16-47 and the second port's data can be found in relative bytes 48-79.			
0-1	RCD Interface ID	This 2 byte hexadecimal field contains the internal ID of the HBA port that is reported in the Read Configuration Data (RCD), General Node Element Qualifier (NEQ) record. This is the value at the end of the interval.	
2-4	Reserved	All bytes are set to x00.	

Bytes	Name	Description	When data is sampled or updated
5-8	Bytes Read by the Channel	This 4 byte hexadecimal field contains the number of bytes transferred to the channel from this HBA port as part of a read operation. This is the value after the data has been decompressed by the HBA. The value is reported in increments of 4 KB (4 x 1024). Any residual data will cause the value to be rounded up to the next higher value. In Figure E-1 on page 850, Bytes Read by the Channel is indicated by the <i>A</i> label. This value is reset to 0 at the beginning of the interval.	The count is incremented for each block of data read.
9-12	Bytes Written by the Channel	This 4 byte hexadecimal field contains the number of bytes transferred from the channel to this HBA port as part of a write operation. This is the value before the effect of the HBA compression. The value is reported in increments of 4 KB (4 x 1024). Any residual data will cause the value to be rounded up to the next higher value. In Figure E-1 on page 850, Bytes Written by the Channel is indicated by the <i>B</i> label. This value is reset to 0 at the beginning of the interval.	The count is incremented for each block of data written.
13-16	Bytes Read from the Virtual Devices	This 4 byte hexadecimal field contains the number of bytes transferred from virtual devices to this HBA port as part of a read operation. The value is for data previously compressed by the HBA. The value is reported in increments of 4 KB (4 x 1024). Any residual data will cause the value to be rounded up to the next higher value. In Figure E-1 on page 850, Bytes Read from Virtual Devices is indicated by the <i>C</i> label. This value is reset to 0 at the beginning of the interval.	The count is incremented for each block of data read.
17-20	Bytes Written to Virtual Devices	This 4 byte hexadecimal field contains the number of bytes transferred to virtual devices from this HBA port as part of a write operation. The value is for data compressed by the HBA. The value is reported in increments of 4 KB (4 x 1024). Any residual data will cause the value to be rounded up to the next higher value. In Figure E-1 on page 850, Bytes Written to Virtual Devices is indicated by the <i>D</i> label. This value is reset to 0 at the beginning of the interval.	The count is incremented for each block of data written.
21-31	Reserved	All bytes set to x00.	

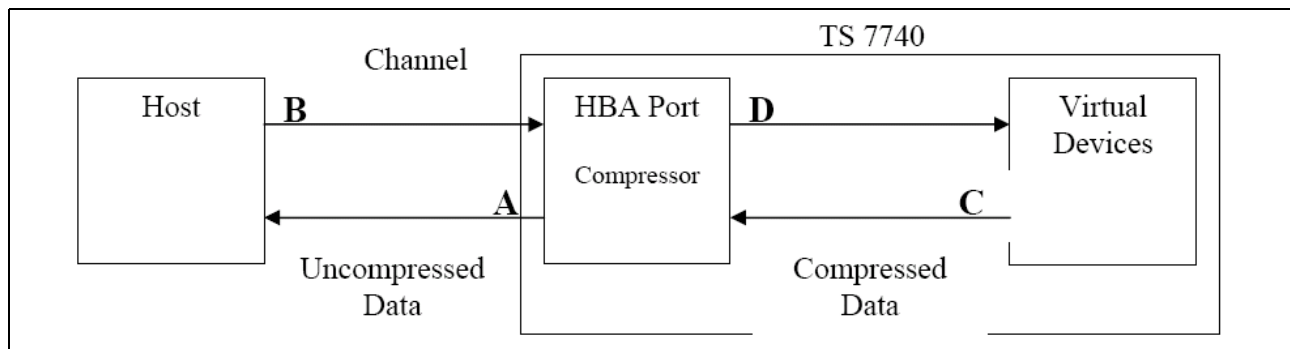


Figure E-1 Sampling points

vNode Virtual Device Historical record

The vNode Virtual Device Historical record has the following nested structure:

- Header
- Virtual Device Container

Table E-4 shows the layout of the vNode Virtual Device Historical record.

Table E-4 vNode Virtual Device Historical record

Bytes	Name	Description	When data is sampled or updated
0-1	Length	This 2 byte hexadecimal field contains the length of this record. The length includes these 2 bytes.	
2	Version	This 1 byte hexadecimal field contains the version of the data presented in this record. Initially the version is set to x01.	
3	Data Type	This 1 byte hexadecimal field indicates the type of data contained in this record. For this record, the value is set to x20, indicating that this is a vNode Virtual Device Historical record.	
4	Node ID	This 1 byte hexadecimal field indicates the vNode ID that this interval's data represents. Valid values are x00 – x0F.	
5	Cluster ID	This 1 byte hexadecimal field indicates the Cluster ID of which this vNode is a part. Valid values are x00 – x07.	
6-7	Interval Duration	This 2 byte hexadecimal field indicates the interval in seconds since this interval's data was taken over.	
8-11	Time Stamp	This 4 byte hexadecimal field indicates the end time of the interval this data was taken over. This value is the time in seconds since the beginning of the Epoch (00:00:00 UTC, January 1, 1970)	
12-15	Machine Type	This 4 byte EBCDIC field contains this node's machine type. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "3957".	
16-18	Machine Model	This 3 byte EBCDIC field contains this node's machine model. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "V06".	
19-26	Machine Serial Number	This 8 character EBCDIC field contains the serial number of this node. This field is left justified and padded with EBCDIC blanks. The format is XX-YYYYY, where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed.	
27-34	VE Code Level	This 8 byte hexadecimal field contains the code level of the TS7700 Virtualization Engine. The 8 bytes are actually four 2 byte fields. Each 2 byte field represents a portion of the code level. The code level is expressed as Version.Release.Modification.Fix in a decimal form. For example, the code level of 8.0.0.104 would be represented in the 8 bytes as x0008000000000068.	

Bytes	Name	Description	When data is sampled or updated
35-39	Grid Library Sequence Number	This 5 character EBCDIC field contains the Library Sequence Number of the Grid (Composite) library.	
40-44	Distributed Library Sequence Number	This 5 character EBCDIC field contains the Distributed Library Sequence Number for this Distributed Library ID.	
45-63	Reserved	All bytes set to x00.	
vNode Virtual Device Container Bytes 64-191			
64-65	Installed Virtual Devices	This 2 byte hexadecimal field indicates the number of installed virtual devices in this vNode. This is the value at the end of the interval.	
66-69	Virtual Device Type	This 4 byte EBCDIC field indicates the device type emulated by the virtual devices. Initially this is set to "3490". This field is left justified and padded with blanks.	
70-72	Virtual Device Model	This 3 byte EBCDIC field indicates the device model emulated by the virtual devices. Initially this is set to "C2A". This field is left justified and padded with blanks. This is the value at the end of the interval.	
73-80	Channel Blocks Written 1-2048 byte range	This 8 byte hexadecimal field indicates the number of channel blocks written to all the virtual devices in this vNode that had a size of between 1 and 2048 bytes inclusive for this interval. This value is reset to 0 at the beginning of the interval.	The count is incremented for each block of data written that fits the size.
81-88	Channel Blocks Written 2049-4096 byte range	This 8 byte hexadecimal field indicates the number of channel blocks written to all the virtual devices in this vNode that had a size of between 2049 and 4096 bytes inclusive for this interval. This value is reset to 0 at the beginning of the interval.	The count is incremented for each block of data written that fits the size.
89-96	Channel Blocks Written 4097-8192 byte range	This 8 byte hexadecimal field indicates the number of channel blocks written to all the virtual devices in this vNode that had a size of between 4097 and 8192 bytes inclusive for this interval. This value is reset to 0 at the beginning of the interval.	The count is incremented for each block of data written that fits the size.
97-104	Channel Blocks Written 8193-16384 byte range	This 8 byte hexadecimal field indicates the number of channel blocks written to all the virtual devices in this vNode that had a size of between 8193 and 16384 bytes inclusive for this interval. This value is reset to 0 at the beginning of the interval.	The count is incremented for each block of data written that fits the size.
105-112	Channel Blocks Written 16385-32768 byte range	This 8 byte hexadecimal field indicates the number of channel blocks written to all the virtual devices in this vNode that had a size of between 16385 and 32768 bytes inclusive for this interval. This value is reset to 0 at the beginning of the interval.	The count is incremented for each block of data written that fits the size.

Bytes	Name	Description	When data is sampled or updated
113-120	Channel Blocks Written 32769-65536 byte range	This 8 byte hexadecimal field indicates the number of channel blocks written to all the virtual devices in this vNode that had a size of between 32769 and 65536 bytes inclusive for this interval. This value is reset to 0 at the beginning of the interval.	The count is incremented for each block of data written that fits the size.
121-128	Channel Blocks Written above 65536 byte range	This 8 byte hexadecimal field indicates the number of channel blocks written to all the virtual devices in this vNode that had a size of 65537 or higher for this interval. This value is reset to 0 at the beginning of the interval.	The count is incremented for each block of data written that fits the size.
129-130	Configured Maximum Throughput	This 2 byte hexadecimal field contains the current maximum throughput for this vNode. The value is expressed in MBps (1 MB = 1024x1024 bytes). This field is set to x0000 if there is no restriction for the maximum throughput. This is the value at the end of the interval.	This will be set to x0000 for the first release.
131-132	Minimum Virtual Devices Mounted	This 2 byte hexadecimal field indicates the minimum number of virtual devices that were mounted at the same time over the interval.	The count of mounted virtual devices is sampled every 15 seconds. The min/max/avg is updated over the interval.
133-134	Maximum Virtual Devices Mounted	This 2 byte hexadecimal field indicates the maximum number of virtual devices that were mounted at the same time over the interval.	The count of mounted virtual devices is sampled every 15 seconds. The min/max/avg is updated over the interval.
135-136	Average Virtual Devices Mounted	This 2 byte hexadecimal field indicates the average number of virtual devices that were mounted at the same time over the interval. The average is calculated by recording the number of mounted devices on a periodic basis and then averaging it over the interval.	The count of mounted virtual devices is sampled every 15 seconds. The min/max/avg is updated over the interval.
137-191	Reserved	All bytes are set to x00.	

vNode Adapter Historical record

The vNode Adapter Historical record has the following nested structure:

- ▶ Header
- ▶ Adapter 0 Container
 - Adapter-Port 0 Container
 - Adapter-Port 1 Container
- ▶ Adapter 1 Container
 - Adapter-Port 0 Container
 - Adapter-Port 1 Container
- ▶ Adapter 2 Container
 - Adapter-Port 0 Container
 - Adapter-Port 1 Container
- ▶ Adapter 3 Container
 - Adapter-Port 0 Container
 - Adapter-Port 1 Container

Table E-5 shows the vNode Adapter Historical record layout. Figure E-2 on page 857 shows sampling points.

Table E-5 vNode Adapter Historical record

Bytes	Name	Description	When data is sampled or updated
0-1	Length	This 2 byte hexadecimal field contains the length of this record. The length includes these 2 bytes.	
2	Version	This 1 byte hexadecimal field contains the version of the data presented in this record. Initially the version is set to x01.	
3	Data Type	This 1 byte hexadecimal field indicates the type of data contained in this record. For this record, the value is set to x21, indicating that this is a vNode Adapter Historical record.	
4	Node ID	This 1 byte hexadecimal field indicates the vNode ID that this interval's data represents. Valid values are x00 – x0F.	
5	Cluster ID	This 1 byte hexadecimal field indicates the Cluster ID of which this vNode is a part. Valid values are x00 – x07.	
6-7	Interval Duration	This 2 byte hexadecimal field indicates the interval in seconds that this interval's data was taken over.	
8-11	Time Stamp	This 4 byte hexadecimal field indicates the end time of the interval this data was taken over. This value is the time in seconds since the beginning of the Epoch (00:00:00 UTC, January 1, 1970)	
12-15	Machine Type	This 4 byte EBCDIC field contains this node's machine type. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "3957".	
16-18	Machine Model	This 3 byte EBCDIC field contains this node's machine model. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "V06".	

Bytes	Name	Description	When data is sampled or updated
19-26	Machine Serial Number	This 8 character EBCDIC field contains the serial number of this node. This field is left justified and padded with EBCDIC blanks. The format is XX-YYYYY, where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed.	
27-34	VE Code Level	This 8 byte hexadecimal field contains the code level of the TS7700 Virtualization Engine. The 8 bytes are actually four 2 byte fields. Each 2 byte field represents a portion of the code level. The code level is expressed as Version.Release.Modification.Fix in a decimal form. For example, the code level of 8.0.0.104 would be represented in the 8 bytes as x0008000000000068.	
35-39	Grid Library Sequence Number	This 5 character EBCDIC field contains the Library Sequence Number of the Grid (Composite) library.	
40-44	Distributed Library Sequence Number	This 5 character EBCDIC field contains the Distributed Library Sequence Number for this Distributed Library ID.	
45-63	Reserved	All bytes are set to x00.	
vNode Adapter Container Bytes 64-383 (4 sets x 80 bytes/set = 320 bytes) This next set of bytes contains information for up to 4 host bus adapters (HBA). For each adapter, there is data for up to 2 ports. Each set of data contains 80 bytes. The following fields define the bytes of data and are numbered starting with byte 0. The first set of data is for the adapter 0, the second for adapter 1, and so on. The first adapter's data can be found in bytes 64-143, the second adapter's data can be found in bytes 144-223, and so on.			
0	Adapter Type	This 1 byte hexadecimal field identifies the type of Host Bus Adapter (HBA) this data uses. The possible values are: <ul style="list-style-type: none"> ▶ x00: No adapter installed. ▶ x02-x08: Reserved. ▶ x09: FICON – 1 Port (Arctic Circle). ▶ x0A: FICON – 2 Port (Yukon). All other values are reserved. This is the value at the end of the interval.	
1	Adapter State	This 1 byte hexadecimal field identifies the current state of the adapter. The possible values are: <ul style="list-style-type: none"> ▶ x00: No adapter is installed. ▶ x01: The adapter is online. ▶ x02: The adapter is offline. ▶ x03: The adapter is not working. ▶ x04: The adapter is reloading itself. ▶ x05: The adapter is in a Check1 condition. All other values are reserved. This is the value at the end of the interval.	

Bytes	Name	Description	When data is sampled or updated
2	HBA Drawer	This 1 byte hexadecimal field indicates in which drawer the HBA is located. The possible values are: <ul style="list-style-type: none"> ► x00: The HBA is in the left drawer when looking at the drawers from the back which is the side the cables plug into. ► x01: The HBA is in the right drawer when looking at the drawers from the back which is the side the cables plug into. All other values are reserved. This is the value at the end of the interval.	
3	HBA Slot Number	This 1 byte hexadecimal field indicates the physical slot number of the HBA within its drawer. This is the value at the end of the interval.	
4-15	Reserved	All bytes set to x00.	
vNode Adapter-Port Container Relative bytes 16-79 (2 sets x 32 bytes/set = 64 bytes) This next set of bytes contains information for up to 2 ports on the HBA. Each set of data contains 32 bytes. The following fields define the 32 bytes and are numbered starting with byte 0. The first port's data can be found in relative bytes 16-47 and the second port's data can be found in relative bytes 48-79.			
0-1	RCD Interface ID	This 2 byte hexadecimal field contains the internal ID of the HBA port that is reported in the Read Configuration Data (RCD), General Node Element Qualifier (NEQ) record. This is the value at the end of the interval.	
2	Maximum Data Rate	This 1 byte hexadecimal field indicates the maximum data rate the FICON port is capable of at the end of the interval. The value is reported in Gbps.	
3	Actual Data Rate	This 1 byte hexadecimal field indicates the actual data rate of the FICON port at the end of the interval. The value is reported in Gbps. A value of x00 in this field indicates that auto-negotiate is enabled.	
4-7	Bytes Read By the Channel	This 4 byte hexadecimal field contains the number of bytes transferred to the channel from this HBA port as part of a read operation. This is the value after the data has been decompressed by the HBA. The value is reported in increments of 4 KB (4 x 1024). Any residual data will cause the value to be rounded up to the next higher value. In Figure E-2 on page 857, Bytes Read by the Channel is indicated by the <i>A</i> label. This value is reset to 0 at the beginning of the interval.	The count is incremented for each block of data read.
8-11	Bytes Written by the Channel;	This 4 byte hexadecimal field contains the number of bytes transferred from the channel to this HBA port as part of a write operation. This is the value before the effect of the HBA compression. The value is reported in increments of 4 KB (4 x 1024). Any residual data will cause the value to be rounded up to the next higher value. In Figure E-2 on page 857, Bytes Written by the Channel is indicated by the <i>B</i> label. This value is reset to 0 at the beginning of the interval.	The count is incremented for each block of data written.

Bytes	Name	Description	When data is sampled or updated
12-15	Bytes Read from Virtual Devices	This 4 byte hexadecimal field contains the number of bytes transferred from virtual devices to this HBA port as part of a read operation. The value is for data previously compressed by the HBA. The value is reported in increments of 4 KB (4 x 1024). Any residual data will cause the value to be rounded up to the next higher value. In Figure E-2 on page 857, Bytes Read from Virtual Devices is indicated by the <i>C</i> label. This value is reset to 0 at the beginning of the interval.	The count is incremented for each block of data read.
16-19	Bytes Written to Virtual Devices	This 4 byte hexadecimal field contains the number of bytes transferred to virtual devices from this HBA port as part of a write operation. The value is for data compressed by the HBA. The value is reported in increments of 4 KB (4 x 1024). Any residual data will cause the value to be rounded up to the next higher value. In Figure E-2 on page 857, Bytes Written to Virtual Devices is indicated by the <i>D</i> label. This value is reset to 0 at the beginning of the interval.	The count is incremented for each block of data written.
20	Selective Resets	This 1 byte hexadecimal field indicates the number of selective resets this port received during the interval. This field is set to xFF when the number of selective resets is greater than 255 for the interval. This value is reset to 0 at the beginning of the interval.	The count is incremented for each selective reset.
21	System Resets	This 1 byte hexadecimal field indicates the number of system resets this port received during the interval. This field is set to xFF when the number of system resets is greater than 255 for the interval. This value is reset to 0 at the beginning of the interval.	The count is incremented for each system reset.
22-31	Reserved	All bytes are set to x00.	

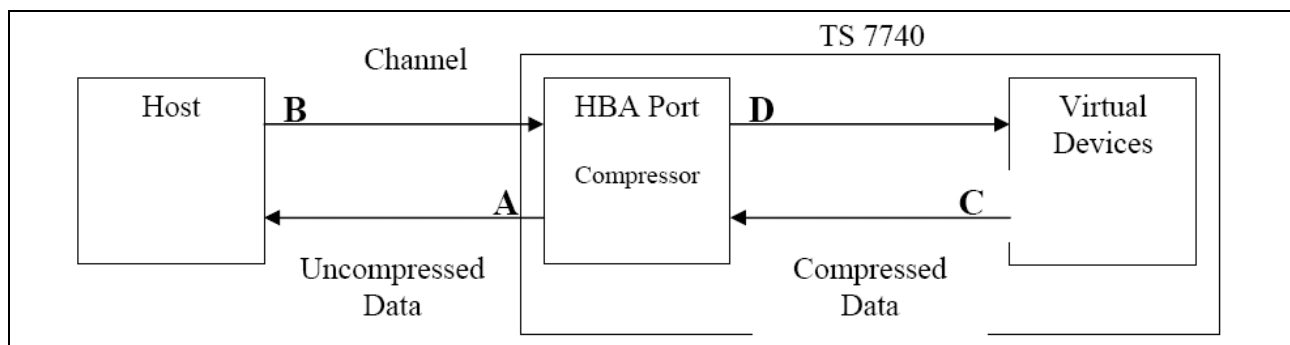


Figure E-2 Sampling points

hNode HSM Point-In-Time (PIT) record

The hNode HSM Point-In-Time record has the following nested structure:

- ▶ Header
- ▶ HSM Container
 - HSM-Library 0 Container
 - HSM-Library-Physical Device 0 Container
 - .
 - .
 - .
 - HSM-Library-Physical Device 31Container
 - HSM-Library 1 Container (If installed)
 - Same sub-containers as HSM-Library 0 Container
 - HSM-Library 2 Container (If installed)
 - Same sub-containers as HSM-Library 0 Container
 - HSM-Library 3 Container (If installed)
 - Same sub-containers as HSM-Library 0 Container

See Table E-6 for the hNode HSM Point-In-Time record format and values.

Table E-6 hNode HSM Point-in-Time record

Bytes	Name	Description	When data is sampled or updated
0-1	Length	This 2 byte hexadecimal field contains the length of this record. The length includes these 2 bytes.	
2	Version	This 1 byte hexadecimal field contains the version of the data presented in this record. Initially the version is set to x02.	
3	Data Type	This 1 byte hexadecimal field indicates the type of data contained in this record. For this record, the value is set to x10, indicating that this is an hNode HSM Point-in-Time record.	
4	Node ID	This 1 byte hexadecimal field indicates the hNode ID that this interval's data represents. Valid values are x00 – x01.	
5	Cluster ID	This 1 byte hexadecimal field indicates the Cluster ID of which this hNode is a part. Valid values are x00 – x07.	
6-7	Interval Duration	This 2 byte hexadecimal field indicates the interval in seconds that this interval's data was taken over.	
8-11	Time Stamp	This 4 byte hexadecimal field indicates the end time of the interval this data was taken over. This value is the time in seconds since the beginning of the Epoch (00:00:00 UTC, January 1, 1970).	
12-15	Machine Type	This 4 byte EBCDIC field contains this node's machine type. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "3957".	
16-18	Machine Model	This 3 byte EBCDIC field contains this node's machine model. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "V06".	

Bytes	Name	Description	When data is sampled or updated
19-26	Machine Serial Number	This 8 character EBCDIC field contains the serial number of this node. This field is left justified and padded with EBCDIC blanks. The format is XX-YYYYY, where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed.	
27-34	VE Code Level	This 8 byte hexadecimal field contains the code level of the TS7700 Virtualization Engine. The 8 bytes are actually four 2 byte fields. Each 2 byte field represents a portion of the code level. The code level is expressed as Version.Release.Modification.Fix in a decimal form. For example, the code level of 8.0.0.104 would be represented in the 8 bytes as x0008000000000068.	
35-39	Grid Library Sequence Number	This 5 character EBCDIC field contains the Library Sequence Number of the Grid (Composite) library.	
40-44	Distributed Library Sequence Number	This 5 character EBCDIC field contains the Distributed Library Sequence Number for this Distributed Library ID	
45-63	Reserved	All bytes are set to x00.	
HSM Container Bytes 64 and up. This container provides information concerning HSM related items. The total length depends upon the number of physical libraries attached to this hNode.			
64-65	Recalls in Queue	This 2 byte hexadecimal field contains the current number of queued recall operations at the end of the interval.	The count is updated as recalls added or removed from the queue.
66-67	Pre-migrates in Queue	This 2 byte hexadecimal field contains the current number of queued premigrate operations at the end of the interval.	The count is updated with the current value every 30 seconds.
68-71	Host Write Throttle	This 4 byte hexadecimal field contains the copy throttling value over the interval. The value is reported in thousandths of a second. This is the value at the end of the interval.	The count is updated with the current value every 30 seconds.
76	Number of Physical Libraries	This 1 byte hexadecimal field indicates the number of physical libraries to which this hNode is attached. This field can be used to determine how many HSMLibrary containers will be attached to this record. This is the value at the end of the interval. Note: Only one physical library is supported at this time.	
77-80	Deferred Copy Throttle	This 4 byte hexadecimal field contains the current copy throttling value over the interval. This is for throttling where the cluster was prioritizing Host and Immediate Copies over sourcing Deferred Copies because of a constrained resource. The value is reported in thousandths of a second. This is the value at the end of the interval.	
81-95	Reserved	All bytes are set to x00.	

Bytes	Name	Description	When data is sampled or updated
HSM-Library Container Bytes 96 and up (Number of Physical Libraries x 1568 bytes/set). For example, if there is just one physical library attached to this node, there will be 1 set of data with 1568 bytes. There is a maximum of 4 physical libraries attached to a single Cluster. This next segment of the record contains one set of data for each physical library attached to this hNode as defined in byte 76. Each set of data contains 1568 bytes. The data for the first library is found in bytes 96-1663; the second library's data (if the library exists) is found in bytes 1664-3231, and so on. These fields contain information concerning the underlying automation.			
0-15	Library Sequence Number	This 16 byte EBCDIC field indicates the Library Sequence Number of the underlying automation. This field is left justified and padded with EBCDIC blanks. This is the value at the end of the interval.	
16-31	Reserved	All bytes are set to x00.	
HSM-Library-Physical Device Container Relative bytes 32-415 (32 sets of device data x 48 bytes/set = 1536 bytes). For a system with just 1 class of physical device (homogeneous), the data for the devices is found in the first 16 sets of data. For a system with 2 classes of physical devices (heterogeneous), the data for one class of drives is found in the first 16 sets of data, and the data for the second class of devices is found in the second 16 sets of data. There are 32 sets of data per possible library, one for each of 32 possible physical devices per physical library. The data for the first device of the first device class is in relative bytes 0-47; the second device's data of the first device class is in bytes 48-95, and so on. For the second device class, the first device's data is found in relative bytes 768-815, the second device's data is in bytes 816-863, and so forth.			
0	Device Class ID	This one byte hexadecimal field contains this device's device class identifier. ▶ x00: No device installed. ▶ x11: 3590 Model B1A. ▶ x13: 3590 Model E1A. ▶ x14: 3590 Model H1A. ▶ x20: 3592 Model J1A. This also includes a 3592-E05 that is emulating a 3592-J1A device. ▶ x22: 3592 Model E05. This is for a 3592-E05 that is behaving as a 3590-E05. ▶ x23: 3592 Model E05 (encryption configured). ▶ x24: 3592 Model E06. All other values are reserved. This is the value at the end of the interval.	
1-10	Physical Volume	This 10 byte EBCDIC field contains the VOLSER of the volume that is loaded in this device at the end of the interval, if any. This field is left justified and padded with EBCDIC blanks. This field is filled with EBCDIC blanks when there is no volume in the device.	The VOLSER is updated when there is a change.
11	Volume Pool	This 1 byte hexadecimal field indicates the pool associated with the volume, if any, loaded in the device at the end of the interval. This field is set to x00 when there is not a volume loaded in the device. Values 1-32 are also valid in this field.	The pool is updated when there is a change.

Bytes	Name	Description	When data is sampled or updated
12	Device State	This 1 byte hexadecimal field indicates the state of the device at the end of the interval. Possible values are: <ul style="list-style-type: none"> ▶ x00: Device is online. ▶ x01: Device is offline. All other values are reserved.	The state is updated when there is a change.
13	Device Role	This 1 byte hexadecimal field contains the role of the device at the end of the interval. Possible values are: <ul style="list-style-type: none"> ▶ x00: Idle ▶ x01: Recalling ▶ x02: Migrating ▶ x03: Reclaiming ▶ x04: Importing ▶ x05: Exporting ▶ x06: Data Security Erase All other values are reserved.	The role is updated when there is a change.
14-23	Logical Volume	This 10 byte EBCDIC field contains the VOLSER of the logical volume, if any, that is being processed by the physical device at the end of the interval. This field is left justified and padded with EBCDIC blanks. This field contains all EBCDIC blanks when there is not a logical volume being processed at the end of the interval.	The VOLSER is updated when there is a change.
24-27	Data read from Device	This 4 byte hexadecimal field indicates the number of bytes transferred from the physical device during this interval. The value is reported in increments of 100 KB (100 x 1024). Any residual data will cause the value to be rounded up to the next higher value. This value is reset to 0 at the beginning of the interval.	The count is incremented after every read.
28-31	Data Written to Device	This 4 byte hexadecimal field indicates the number of bytes transferred to the physical device during this interval. The value is reported in increments of 100 KB (100 x 1024). Any residual data will cause the value to be rounded up to the next higher value. This value is reset to 0 at the beginning of the interval.	The count is incremented after each write.
32-47	Reserved	All bytes are set to x00.	

hNode Grid Point-In-Time (PIT) record

The hNode Grid Point-In-Time (PIT) record has the following nested structure:

- ▶ Header
- ▶ Grid Container
 - Grid-Cluster 0 Container
 - Grid-Cluster 1 Container (If installed)
 - Grid-Cluster 2 Container (If installed)
 - Grid-Cluster 3 Container (If installed)
 - Grid-Cluster 4 Container (If installed)
 - Grid-Cluster 5 Container (If installed)
 - Grid-Cluster 6 Container (If installed)
 - Grid-Cluster 7 Container (If installed)

Table E-7 shows the record data format and values.

Table E-7 hNode Grid PIT Record record

Bytes	Name	Description	When data is sampled or updated
0-1	Length	This 2 byte hexadecimal field contains the length of this record. The length includes these 2 bytes.	
2	Version	This 1 byte hexadecimal field contains the version of the data presented in this record. Initially the version is set to x02.	
3	Data Type	This 1 byte hexadecimal field indicates the type of data contained in this record. For this record, the value is set to x10, indicating that this is an hNode HSM Point- In-Time record.	
4	Node ID	This 1 byte hexadecimal field indicates the hNode ID that this interval's data represents. Valid values are x00 – x01.	
5	Cluster ID	This 1 byte hexadecimal field indicates the Cluster ID of which this hNode is a part. Valid values are x00 – x07.	
6-7	Interval Duration	This 2 byte hexadecimal field indicates the interval in seconds that this interval's data was taken over.	
8-11	Time Stamp	This 4 byte hexadecimal field indicates the end time of the interval this data was taken over. This value is the time in seconds since the beginning of the Epoch (00:00:00 UTC, January 1, 1970).	
12-15	Machine Type	This 4 byte EBCDIC field contains this node's machine type. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "3957".	
16-18	Machine Model	This 3 byte EBCDIC field contains this node's machine model. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "V06".	
19-26	Machine Serial Number	This 8 character EBCDIC field contains the serial number of this node. This field is left justified and padded with EBCDIC blanks. The format is XX-YYYYY where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed.	

Bytes	Name	Description	When data is sampled or updated
27-34	VE Code Level	This 8 byte hexadecimal field contains the code level of the TS7700 Virtualization Engine. The 8 bytes are actually four 2 byte fields. Each 2 byte field represents a portion of the code level. The VE code level is expressed as Version.Release.Modification.Fix in a decimal form. For example, the code level of 8.0.0.104 would be represented in the 8 bytes as x0008000000000068.	
35-39	Grid Library Sequence Number	This 5 character EBCDIC field contains the Library Sequence Number of the Grid (Composite) library.	
40-44	Distributed Library Sequence Number	This 5 character EBCDIC field contains the Distributed Library Sequence Number for this Distributed Library ID.	
45-63	Reserved	All bytes are set to x00.	
Grid Container Bytes 64 and up. This container provides information concerning the Grid aspects of this cluster. The length of the message depends on the number of clusters in the Grid. There are 32 fixed bytes plus 32 bytes for each cluster in the Grid.			
64-67	Immediate Copy Queue	This 4 byte hexadecimal field indicates the number of logical volumes in the immediate copy queue targeted for this cluster at the end of the interval.	The count is updated every 5 minutes.
68-71	Deferred Copy Queue	This 4 byte hexadecimal field indicates the number of logical volumes in the deferred copy queue targeted for this cluster at the end of the interval.	The count is updated every 5 minutes.
72-75	Active Copies	This 4 byte hexadecimal field indicates the number of active copies that are targeting this cluster at the end of the interval.	The count is updated every 5 minutes.
76-78	Reserved	Reserved for future use.	
79	Number of Clusters	This 1 byte hexadecimal field indicates the number of clusters in the Grid. This field can be used to determine how many Grid-cluster containers will be attached to this record. There is a maximum of 8 clusters. This is the value at the end of the interval.	
80-95	Reserved	All bytes are set to x00.	

Bytes	Name	Description	When data is sampled or updated
Grid-Cluster Container Bytes 96 and up (Number of Clusters x 129 bytes/set). For example, if there are 3 clusters in the Grid there will be 3 sets of data. There is a maximum of 8 clusters in a Grid. This next segment of the record contains one set of data for each cluster in the Grid as defined in byte 79. Each set of data contains 129 bytes. The data for the first cluster (Cluster 0) can be found in bytes 96-223; the second cluster's (Cluster 1) data can be found in bytes 224-351, and so on.			
0	Cluster Link State	This 1 byte hexadecimal field indicates the state of the link between this cluster and the others in the Grid at the end of the interval. Possible values are: <ul style="list-style-type: none"> ► x00: All links are fully operational. There are no detected error conditions between the clusters in the Grid. For a Grid with only a single cluster, this is the value always reported in this field. ► x01: Degraded. One or more of the logical communication paths between the clusters is not operational. ► x02: Failed. All of the logical communication paths between this Cluster and the other Clusters in the Grid are not working. All other values are reserved. This field represents the overall link state as determined by the cluster. The actual connection may consists of several physical connections and underlying infrastructure, which is transparent to the system.	This is updated when the link state changes.
1-32	Data transferred into a Cluster's Cache from other Clusters	This field contains eight 4 byte fields indicating the number of bytes transferred to this cluster's cache from other clusters as part of a remote file access during this interval. The first 4 byte field indicates the number of bytes transferred into this cluster's cache from Cluster 0. The second 4 byte field indicates the number of bytes transferred from Cluster 1, and so on. A value of x00000000 is indicated for a cluster's own value and for non-existent clusters. The value is reported in increments of 100 KB (100 x 1024). Any residual data will cause the value to be rounded up to the next higher value. This value is reset to 0 at the beginning of the interval.	The count is incremented every time a block of data is written to this cluster.
33-64	Data transferred from a Cluster's Cache to Other Clusters	This field contains eight 4 byte fields indicating the number of bytes transferred from this cluster's cache from other clusters as part of a copy and remote file access during this interval. The first 4 byte field indicates the number of bytes transferred from this cluster's cache to Cluster 0. The second 4 byte field indicates the number of bytes transferred to Cluster 1, and so on. A value of x00000000 is indicated for a cluster's own value and for non-existent clusters. The value is reported in increments of 100 KB (100 x 1024). Any residual data will cause the value to be rounded up to the next higher value. This value is reset to 0 at the beginning of the interval.	The count is incremented every time a block of data is read from this cluster.
65-128	Reserved	All bytes are set to x00.	

hNode HSM Historical record

The hNode HSM Historical record has the following nested structure:

- ▶ HSM-Disaster Recovery Container
- ▶ HSM-Cache Container
 - HSM-Cache-Partition 0 Container
 - HSM-Cache-Partition-Preference Group 0 Container
 - HSM-Cache-Partition-Preference Group 1 Container
 - HSM-Cache-Partition 1 Container
 - HSM-Cache-Partition-Preference Group 0 Container
 - HSM-Cache-Partition-Preference Group 1 Container
 - HSM-Cache-Partition 2 Container
 - HSM-Cache-Partition-Preference Group 0 Container
 - HSM-Cache-Partition-Preference Group 1 Container
 - HSM-Cache-Partition 3 Container
 - HSM-Cache-Partition-Preference Group 0 Container
 - HSM-Cache-Partition-Preference Group 1 Container
 - HSM-Cache-Partition 4 Container
 - HSM-Cache-Partition-Preference Group 0 Container
 - HSM-Cache-Partition-Preference Group 1 Container
 - HSM-Cache-Partition 5 Container
 - HSM-Cache-Partition-Preference Group 0 Container
 - HSM-Cache-Partition-Preference Group 1 Container
 - HSM-Cache-Partition 6 Container
 - HSM-Cache-Partition-Preference Group 0 Container
 - HSM-Cache-Partition-Preference Group 1 Container
 - HSM-Cache-Partition 7 Container
 - HSM-Cache-Partition-Preference Group 0 Container
 - HSM-Cache-Partition-Preference Group 1 Container

Table E-8 shows the hNode HSM Historical record.

Table E-8 hNode HSM Historical record

Bytes	Name	Description	When data is sampled or updated
0-1	Length	This 2 byte hexadecimal field contains the length of this record. The length includes these 2 bytes.	
2	Version	This 1 byte hexadecimal field contains the version of the data presented in this record. Initially the version is set to x02.	
3	Data Type	This 1 byte hexadecimal field indicates the type of data contained in this record. For this record, the value is set to x30, indicating this is an hNode HSM Historical record.	
4	Node ID	This 1 byte hexadecimal field indicates the hNode ID that this interval's data represents. Valid values are x00 – x01.	
5	Cluster ID	This 1 byte hexadecimal field indicates the Cluster ID of which this hNode is a part. Valid values are x00 – x07.	
6-7	Interval Duration	This 2 byte hexadecimal field indicates the interval in seconds that this interval's data was taken over.	

Bytes	Name	Description	When data is sampled or updated
8-11	Time Stamp	This 4 byte hexadecimal field indicates the end time of the interval this data was taken over. This value is the time in seconds since the beginning of the Epoch (00:00:00 UTC, January 1, 1970).	
12-15	Machine Type	This 4 byte EBCDIC field contains this node's machine type. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "3957".	
16-18	Machine Model	This 3 byte EBCDIC field contains this node's machine model. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "V06".	
19-26	Machine Serial Number	This 8 character EBCDIC field contains the serial number of this node. This field is left justified and padded with EBCDIC blanks. The format is XX-YYYYY, where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed.	
27-34	VE Code Level	This 8 byte hexadecimal field contains the code level of the TS7700 Virtualization Engine. The 8 bytes are actually four 2 byte fields. Each 2 byte field represents a portion of the code level. The VE code level is expressed as Version.Release.Modification.Fix in a decimal form. For example, the code level of 8.0.0.104 would be represented in the 8 bytes as x0008000000000068.	
35-39	Grid Library Sequence Number	This 5 character EBCDIC field contains the Library Sequence Number of the Grid (Composite) library.	
40-44	Distributed Library Sequence Number	This 5 character EBCDIC field contains the Distributed Library Sequence Number for this Distributed Library ID.	
45-63	Reserved	All bytes are set to x00.	
HSM Container Bytes 64-1151. This record contains information pertinent to the TS7700 Storage Manager.			
HSM - Disaster Recovery Container Bytes 64-1151. This set of 32 bytes contains information concerning HSM Disaster Recovery (DR).			
64-73	Disaster Recovery Volser	This 10 byte EBCDIC field contains the VOLSER of the physical volume that contains the latest subsystem backup information. This field is set to all EBCDIC blanks if a disaster recovery volume does not exist. This field is left justified and padded with EBCDIC blanks. This is the value at the end of the interval. For the model VEA without a back-end tape, this field will be EBCDIC blanks.	The VOLSER is updated when it changes.
74-95	Reserved	All bytes are set to x00.	

Bytes	Name	Description	When data is sampled or updated
HSM – Cache Container Bytes 96 – 1151 This set of bytes contains information concerning the Tape Volume Cache (TVC) and this hNode.			
96	Active hNode	This 1 byte hexadecimal field indicates if this hNode was the active node in charge of Tape Volume Cache (TVC) management at the end of the interval being reported. Possible values are: <ul style="list-style-type: none"> ▶ x00: This is not the active hNode. ▶ x01: This is the active hNode. All other values are reserved.	This value is set to x00 for the first release.
97-100	TVC Size	This 4 byte hexadecimal field indicates the current size of the cluster's TVC in increments of 1 GB (1024 x 1024 x 1000). A TVC that is 1.7 TB in size will be reported as x000006A4 (1700 decimal). This field is filled in by the hNode that is currently in charge of TVC management. The hNode that is not in charge of the TVC management reports x00000000 in this field. Refer to the Active hNode field. This is the value at the end of the interval.	
101	Percent Host Write Throttle	This 1 byte hexadecimal field indicates the percentage of 30-second periods where there was at least one throttling value greater than zero. This field is completed by the hNode that is currently in charge of TVC management. The hNode that is not in charge of the TVC management reports x00 in this field. Refer to the Active hNode field. This value is recomputed at the end of the interval based on the data from the interval.	The current throttle value is sampled every 30 seconds. The average is computed over the interval.
102-105	Average Host Write Throttle	This 4 byte hexadecimal field indicates the average host write throttle value during the interval. The value presented is the average of the non-zero throttling values. This field is filled in by the hNode that is currently in charge of TVC management. The hNode that is not in charge of the TVC management reports x00000000 in this field. Refer to the Active hNode field. This value is recomputed at the end of the interval based on the data from the interval.	The current throttle value is sampled every 30 seconds. The average is computed over the interval.
106	Percent Write Throttle	This 1 byte hexadecimal field indicates the percentage of 30-second periods where there was at least one throttling value greater than zero and that copy was the predominant reason for throttling. This field is filled in by the hNode that is currently in charge of TVC management. The hNode that is not in charge of the TVC management reports x00 in this field. Refer to the Active hNode field. This value is recomputed at the end of the interval based on the data from the interval.	The current throttle value is sampled every 30 seconds. The average is computed over the interval.

Bytes	Name	Description	When data is sampled or updated
107-110	Average Copy Throttle	This 4 byte hexadecimal field indicates the average copy throttle value during the interval. The value presented is the average of the non-zero throttling values where copy was the predominant reason for throttling. The value is reported in thousandths of a second. This field is filled in by the hNode that is currently in charge of TVC management. The hNode that is not in charge of the TVC management reports x00000000 in this field. Refer to the Active hNode field. This value is recomputed at the end of the interval based on the data from the interval.	The current throttle value is sampled every 30 seconds. The average is computed over the interval.
111-114	Average Overall Throttle	This 4 byte hexadecimal field indicates the average of all throttling values during the interval. The calculation includes samples for periods where throttling was both zero and non-zero. The value is reported in thousandths of a second. This field is completed by the hNode that is currently in charge of TVC management. The hNode that is not in charge of the TVC management reports x00000000 in this field. Refer to the Active hNode field. This value is recomputed at the end of the interval based on the data from the interval.	The current throttle values are sampled every 30 seconds. The average is computed over the interval.
115	Percent Deferred Copy Throttle	This 1 byte hexadecimal field indicates the percentage of 30-second periods where deferred copy throttle from this hNode was active (non-zero). This field is completed by the hNode that is currently in charge of TVC management. The hNode that is not in charge of the TVC management reports x00 in this field. Refer to the Active hNode field. This value is recomputed at the end of the interval based on the data from the interval.	
116-119	Average Deferred Copy Throttle	This 4 byte hexadecimal field indicates the average deferred copy throttle value during the interval. The value presented is the average of 30-second intervals of the deferred copy throttle value over the historical record interval. The value is reported in thousandths of a second. This value is recomputed at the end of the interval based on the data from the interval.	
120-123	Base Deferred Copy Throttle	This 4 byte hexadecimal field indicates the base deferred copy throttle value that would be used if deferred copy throttle were to be applied. The value is reported in thousandths of a second.	
124-127	Pre-migration Throttle Threshold	This 4 byte hexadecimal field indicates the current threshold of the premigration throttle in increments of 1 GB (1024 x 1024 x 1000). A threshold that is 1.7 TB in size will be reported as x000006A4 (1700 decimal). This field represents amount of un-migrated data in cache, at which the system will begin throttling the host I/O to keep the value below this level. This field is completed by the hNode that is currently in charge of TVC management. The hNode that is not in charge of the TVC management reports x00000000 in this field. Refer to the Active hNode field. This is the value at the end of the interval.	

Bytes	Name	Description	When data is sampled or updated
HSM – Cache – Partition Container Bytes 128 – 1151. This next set of bytes contains information for up to 8 cache partitions for the cluster. Each set of data contains 128 bytes. The following fields define the 128 bytes and are numbered starting with byte 0. The first cache partition's data can be found in relative bytes 0-127 and the second cache partition's data can be found in relative bytes 128-255, and so forth. These fields are completed by the hNode that is currently in charge of TVC management. The hNode that is not in charge of the TVC management reports x00 in these fields. Note: At this time, the TS7700 Virtualization Engine only supports a single cache partition.			
0-3	Partition Size	This 4 byte hexadecimal field indicates the amount of cache assigned to this partition. The value is reported in increments of 1 GB (1024 x 1024 x 1000). This field is completed by the hNode that is currently in charge of TVC management. The hNode that is not in charge of the TVC management reports x00000000 in this field. Refer to the Active hNode field. This is the value at the end of the interval.	The size is updated when it changes.
4-5	Fast Ready Mounts	This 2 byte hexadecimal field indicates the number of mount requests completed using the Fast Ready method during this interval. A mount is accredited to the interval when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.	The count is incremented when the mount complete message is received.
6-9	Average Fast Ready Mount Time	This 4 byte hexadecimal field indicates the average time, in milliseconds, taken to complete Fast-Ready mounts during the interval. Mount time is accrued from the time the mount request is accepted by the system (PLF received, DE returned) until the x20 message is received from the Library Manager. The mount time is averaged into the interval's time when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.	The time is incremented for each mount and averaged at the end of the interval.
10-11	Cache Hit Mounts	This 2 byte hexadecimal field indicates the number of mount requests completed that the data was resident in the Tape Volume Cache (TVC) during this interval. A mount is accredited to the interval when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.	The count is incremented when the mount complete message is received.
12-15	Average Cache Hit Mount Time	This 4 byte hexadecimal field indicates the average time, in milliseconds, taken to complete Cache Hit mounts during the interval. Mount time is accrued from the time the mount request is accepted by the system (PLF received, DE returned) until the x20 message is received from the Library Manager. The mount time is averaged into the interval's time when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.	The time is incremented for each mount and averaged at the end of the interval.
16-17	Cache Miss Mounts	This 2 byte hexadecimal field indicates the number of mount requests completed that required recall from a stacked volume during this interval. A mount is accredited to the interval when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.	The count is incremented when the mount complete message is received.

Bytes	Name	Description	When data is sampled or updated
18-21	Average Cache Miss Mount Time	This 4 byte hexadecimal field indicates the average time, in milliseconds, taken to complete Cache Miss mounts during the interval. Mount time is accrued from the time the mount request is accepted by the system (PLF received, DE returned) until the x20 message is received from the Library Manager. The mount time is averaged into the interval's time when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.	The time is incremented for each mount and averaged at the end of the interval.
22-31	Reserved	All bytes are set to x00.	
HSM – Cache – Partition – Preference Group Container Relative bytes 32 – 127 (2 x 48 bytes = 96 bytes). These bytes are relative to the HSM – Cache - Partition Container. This next set of bytes contains information for 2 preference groups for the cache partition. Each set of data contains 48 bytes. The following fields define the 48 bytes and are numbered starting with byte 0. The first preference group's (PG0) data can be found in bytes 0-47 and the second preference group's (PG1) data can be found in relative bytes 48-95. Because these fields are part of the HSM-Cache-Partition container, these fields are completed by the hNode that is currently in charge of TVC management. The hNode that is not in charge of the TVC management reports x00 in these fields.			
0-3	Virtual Volumes in Cache	This 4 byte hexadecimal field contains the number of virtual volumes in the Tape Volume Cache (TVC) partition that are assigned to the preference group that this data uses. This is the value at the end of the interval.	The count is updated every 10 minutes.
4-7	Data Resident in Cache	This 4 byte hexadecimal field contains the amount of data in the TVC partition whose volumes are assigned to the preference that this data uses. The value is reported in increments of 1 MB (1024 x 1024). Any residual data will cause the value to be rounded up to the next higher value. This is the value at the end of the interval.	The count is updated every 10 minutes.
8-11	4 Hour Average Cache Age	This 4 byte hexadecimal field contains the average age, in minutes, of the oldest logical volume in cache, excluding outliers, from the previous 4 hourly samples. Each hourly sample discards "outliers" that are small numbers of logical volumes that are not representative of the cache as a whole. This value is for volumes that were assigned to the preference group that this data uses. Each sample is rounded up to the nearest minute. This value is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.	The age is updated once an hour, on the hour.
12-15	Volumes Migrated Last 4 Hours	This 4 byte hexadecimal field contains the number of virtual volumes migrated from the cache partition over the past 4 hours that are assigned to the preference group that this data uses. This data is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.	The count is updated once an hour, on the hour.

Bytes	Name	Description	When data is sampled or updated
16-19	48 Hour Average Cache Age	This 4 byte hexadecimal field contains the average age, in minutes, of the oldest logical volume in cache, excluding outliers, from the previous 48 hourly samples. Each hourly sample discards "outliers" that are small numbers of logical volumes that are not representative of the cache as a whole. This value is for volumes that were assigned to the preference group that this data uses. Each sample is rounded up to the nearest minute. This value is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.	The age is updated once an hour, on the hour.
20-23	Volumes Migrated Last 48 Hours	This 4 byte hexadecimal field contains the number of virtual volumes migrated from the cache partition over the past 48 hours that are assigned to the preference group that this data uses. This data is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.	The count is updated once an hour, on the hour.
24-27	35 Day Average Cache Age	This 4 byte hexadecimal field contains the average age, in minutes, of the oldest logical volume in cache, excluding outliers, from the previous 35 days worth of hourly samples. Each hourly sample discards "outliers" that are small numbers of logical volumes that are not representative of the cache as a whole. This value is for volumes that were assigned to the preference group that this data uses. Each sample is rounded up to the nearest minute. This value is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.	The age is updated once an hour, on the hour.
28-31	Volumes Migrated Last 35 Days	This 4 byte hexadecimal field contains the number of virtual volumes migrated from the cache partition over the past 35 days that are assigned to the preference group that this data uses. This data is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.	The count is updated once an hour, on the hour.
32-35	Unmigrated Data	This 4 byte hexadecimal field contains the amount of data in the TVC partition whose volumes are assigned to this preference group, and are not yet migrated to physical tape (cache only). The value is reported in increments of 1 MB (1024 x 1024). Any residual data will cause the value to be rounded up to the next higher value. This is the value at the end of the interval. For the model VEA, or for a partition that is defined to not migrate out to tape, all active data should be in the unmigrated state.	
36-39	Awaiting Replication to available Clusters	This 4 byte hexadecimal field contains the amount of data in the TVC partition whose volumes are assigned to this preference group, and are awaiting replication to other available clusters. Data to be replicated to clusters that are either not available (service or offline) or are blocked from receiving copies (Host Console Request) are not counted. This field depicts data that resides in cache. Data to be replicated that exists on tape only is not included. The value is reported in increments of 1 MB (1024 x 1024). Any residual data will cause the value to be rounded up to the next higher value. This is the value at the end of the interval.	
40-47	Reserved	All bytes are set to x00.	

hNode Library Historical record

This hNode Library Historical record has the following nested structure:

- ▶ Library Device Type 0 Usage Container
- ▶ Library Device Type 1 Usage Container
- ▶ Library Device Type 2 Usage Container
- ▶ Library Device Type 3 Usage Container
- ▶ Library-Pooling Container
 - Library-Pooling-CSP Media Type 0 Container
 - Library-Pooling-CSP Media Type 1 Container
 - Library-Pooling-CSP Media Type 2 Container
 - Library-Pooling-CSP Media Type 3 Container
 - Library-Pooling-CSP Media Type 4 Container
 - Library-Pooling-CSP Media Type 5 Container
 - Library-Pooling-CSP Media Type 6 Container
 - Library-Pooling-CSP Media Type 7 Container
 - Library-Pooling-GUP 1 Container
 - Library-Pooling-GUP-Media Type 0 Container
 - Library-Pooling-GUP-Media Type 1 Container
 - Library-Pooling-GUP-Media Type 2 Container
 - Library-Pooling-GUP-Media Type 3 Container
 - Library-Pooling-GUP-Media Type 4 Container
 - Library-Pooling-GUP-Media Type 5 Container
 - Library-Pooling-GUP-Media Type 6 Container
 - Library-Pooling-GUP-Media Type 7 Container
 - Library-Pooling-GUP-Reclaim Container
 - Library-Pooling-GUP-Properties Container
- ▶ Library-Pooling-GUP 2 Container
 - Same sub-containers as GUP 1
- ▶ .
- ▶ .
- ▶ .
- ▶ Library-Pooling-GUP 32 Container
 - Same sub-containers as GUP 1

Table E-9 shows the hNode Library Historical record.

Table E-9 hNode Library Historical record

Bytes	Name	Description	When data is sampled or updated
0-1	Length	This 2 byte hexadecimal field contains the length of this record. The length includes these 2 bytes.	
2	Version	This 1 byte hexadecimal field contains the version of the data presented in this record. Initially the version is set to x02.	
3	Data Type	This 1 byte hexadecimal field indicates the type of data contained in this record. For this record, the value is set to x32, indicating that this is an hNode Library Historical record.	
4	Node ID	This 1 byte hexadecimal field indicates the hNode ID that this interval's data represents. Valid values are x00 – x01.	
5	Cluster ID	This 1 byte hexadecimal field indicates the Cluster ID of which this hNode is a part. Valid values are x00 – x07.	

Bytes	Name	Description	When data is sampled or updated
6-7	Interval Duration	This 2 byte hexadecimal field indicates the interval in seconds that this interval's data was taken over.	
8-11	Time Stamp	This 4 byte hexadecimal field indicates the end time of the interval this data was taken over. This value is the time in seconds since the beginning of the Epoch (00:00:00 UTC, January 1, 1970).	
12-15	Machine Type	This 4 byte EBCDIC field contains this node's machine type. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "3957".	
16-18	Machine Model	This 3 byte EBCDIC field contains this node's machine model. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "V06".	
19-26	Machine Serial Number	This 8 character EBCDIC field contains the serial number of this node. This field is left justified and padded with EBCDIC blanks. The format is XX-YYYYY where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed.	
27-34	VE Code Level	This 8 byte hexadecimal field contains the code level of the TS7700 Virtualization Engine. The 8 bytes are actually four 2 byte fields. Each 2 byte field represents a portion of the code level. The VE code level is expressed as Version.Release.Modification.Fix in a decimal form. For example, the code level of 8.0.0.104 would be represented in the 8 bytes as x0008000000000068.	
35-39	Grid Library Sequence Number	This 5 character EBCDIC field contains the Library Sequence Number of the Grid (Composite) library.	
40-44	Distributed Library Sequence Number	This 5 character EBCDIC field contains the Distributed Library Sequence Number for this Distributed Library ID.	
45-63	Reserved	All bytes are set to x00.	
Library Container Bytes 64-255 (192 bytes). This record contains information pertinent to the Library operations associated with this hNode.			
64-69	Library Machine Type	This 6 character EBCDIC field contains the machine type of the underlying automation. The field is left justified and padded with EBCDIC blanks. Initially this field will be set to "3494" or "3584". This is the value at the end of the interval.	
70-72	Library Model Number	This 3 character EBCDIC field contains the model number of the underlying automation. The field is left justified and padded with EBCDIC blanks. Initially this field will be set to "L10" when attached to a 3494 or "L22" when attached to a 3584. This is the value at the end of the interval.	
73-75	Library Manufacturer	This 3 byte EBCDIC field contains an abbreviation of the manufacturer of the underlying automation. The field is left justified and padded with EBCDIC blanks. Initially this will be set to "IBM". This is the value at the end of the interval.	

Bytes	Name	Description	When data is sampled or updated
76-77	Library Plant of Manufacture	This 2 byte EBCDIC field contains the library plant of manufacture of the underlying automation. The field is left justified and padded with EBCDIC blanks. This is the value at the end of the interval.	
78-93	Library Sequence Number	This 16 byte EBCDIC field contains the sequence number of the underlying automation. The field is left justified and padded with EBCDIC blanks. This is the value at the end of the interval.	
94-127	Reserved	All bytes are set to x00.	
Library – Tape Device Usage (TDU) Container Bytes 128-255 (4 sets of data x 32 bytes/set = 128 bytes). This container contains 4 sets of data. The 4 sets of data allow up to 4 device types/models to be attached to the hNode. Each set of data contains 32 bytes. The following fields define these bytes and are numbered starting with 0. Data for the first device type is found in bytes 128-159, the data for the second device type is found in bytes 160-191, and so on. Each hNode reports this data from its perspective.			
0	Device Class ID	This one byte hexadecimal field contains this device's device class identifier. Possible values are: <ul style="list-style-type: none"> ▶ x00: No device is installed. ▶ x11: 3590 Model B1A. ▶ x13: 3590 Model E1A. ▶ x14: 3590 Model H1A. ▶ x20: 3592 Model J1A. This also includes a 3592-E05 that is emulating a 3592-J1A device. ▶ x22: 3592 Model E05. This is for a 3592-E05 that is behaving as a 3590-E05. ▶ x23: 3592 Model E05 (encryption configured). ▶ x24: 3592 Model E06. All other values are reserved. This is the value at the end of the interval.	
1	Installed Physical Devices	This 1 byte hexadecimal field contains the number of physical devices, of the device class indicated, that are available for use at the end of the interval.	
2	Available Physical Devices	This 1 byte hexadecimal field contains the number of physical devices, of the device class indicated, that are available for use at the end of the interval.	The count is updated when the number of available physical devices changes.
3	Maximum Physical Devices Mounted	This 1 byte field contains the maximum number of physical devices, of the device class indicated, that were concurrently mounted during the interval. This value is reset to 0 at the beginning of the interval.	The count is updated when the number of mounted physical devices changes.
4	Minimum Physical Devices Mounted	This 1 byte field contains the minimum number of physical devices, of the device class indicated, that were concurrently mounted during the interval. This value is reset to 0 at the beginning of the interval.	The count is updated when the number of mounted physical devices changes.

Bytes	Name	Description	When data is sampled or updated
5	Average Physical Devices Mounted	<p>This 1 byte field contains the average number of physical devices, of the device class indicated, that were concurrently mounted during the interval. The average is calculated by recording the number of mounted devices on a periodic basis then averaging it over the interval.</p> <p>This value is reset to 0 at the beginning of the interval.</p>	The count of mounted physical devices is sampled every 15 seconds. The average is computed over the interval.
6-7	Maximum Physical Mount Time	<p>This 2 byte field contains the maximum time, in seconds, that it took to complete the execution of a mount request for a physical device, of the device class indicated, over the interval. Mount time is accrued from the time the mount request is sent to the Library Manager (LM) until the mount complete is received from the LM. The mount time is accredited to the interval it was completed.</p> <p>This value is reset to 0 at the beginning of the interval.</p>	The time for each mount is examined when the mount completes.
8-9	Minimum Physical Mount Time	<p>This 2 byte field contains the minimum time, in seconds, that it took to complete the execution of a mount request for a physical device, of the device class indicated, over the interval. Mount time is accrued from the time the mount request is sent to the Library Manager (LM) until the mount complete is received from the LM. The mount time is accredited to the interval it was completed.</p> <p>This value is reset to 0 at the beginning of the interval.</p>	The time for each mount is examined when the mount completes.
10-11	Average Physical Mount Time	<p>This 2 byte field contains the average time, in seconds, that it took to complete the execution of a mount request for a physical device, of the device class indicated, over the interval. Mount time is accrued from the time the mount request is sent to the Library Manager (LM) until the mount complete is received from the LM. The mount time is accredited to the interval it was completed.</p> <p>This value is reset to 0 at the beginning of the interval.</p>	The time for each mount is examined when the mount completes.
12-13	Physical recall Mounts	<p>This 2 byte hexadecimal field contains the number of physical mount requests completed by the library during the interval to satisfy recall mounts for the device class indicated. A mount is accredited to the interval the x20 message is received from the Library Manager.</p> <p>This value is reset to 0 at the beginning of the interval.</p>	The count is incremented each time a recall mount is completed.
14-15	Physical Pre-Migrate Mounts	<p>This 2 byte hexadecimal field contains the number of physical mount requests completed by the library during the interval to satisfy premigrate mounts for the device class indicated. A mount is accredited to the interval the x20 message is received from the Library Manager.</p> <p>This value is reset to 0 at the beginning of the interval.</p>	The count is incremented each time a premigrate mount is completed.

Bytes	Name	Description	When data is sampled or updated
16-17	Physical Reclaim Mounts	This 2 byte hexadecimal field contains the number of physical mount requests completed by the library during the interval to satisfy reclaim mounts for the device class indicated. A mount is accredited to the interval the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.	The count is incremented each time a reclaim mount is completed.
18-19	Physical Security Data Erase Mounts	This 2 byte hexadecimal field contains the number of physical mount requests completed by the library during the interval to satisfy Security Data Erase mounts for the device class indicated. A mount is accredited to the interval the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.	The count is incremented each time a Security Data Erase mount is completed.
20-31	Reserved	All bytes are set to x00.	
Library - Pooling Container Bytes 256-8255 The data from each hNode will be the same within a cluster.			
Library - Pooling – Common Scratch Pool (CSP) Media Container Bytes 256-319 (8 sets of data x 8 bytes/set = 64 bytes). These fields contain 8 sets of data describing the physical types in the CSP. The first media type is referred to as Media Type 0; the second is referred to as Media Type 1 and so on. Each set of data contains 8 bytes. The following fields define these bytes and are numbered starting with 0. The data for media type 0 can be found in bytes 256-263; the data for media type 1 can be found in bytes 264-271, and so forth.			
0	Physical Media Type	This 1 byte hexadecimal field contains the identifier for the media type associated with following common scratch pool volume counts. The value is recorded at the end of the interval. Here are values for this field: <ul style="list-style-type: none"> ▶ x00: No media type for this set of data. ▶ x10: 3590 J media. ▶ x11: 3590 K media. ▶ x20: 3592 JA media. ▶ x21: 3592 JW media (reserved for future use). ▶ x22: 3592 JJ media. ▶ x23: 3592 JR media (reserved for future use). ▶ x24: 3592 JB media. ▶ x25: 3592 JX media (reserved for future use). All other values are reserved. This is the value at the end of the interval.	The type is updated every 2 minutes.
1-2	Physical Media Count	This two byte hexadecimal field contains the number of scratch stacked volumes, of the type identified, assigned to the common scratch pool. This is the value at the end of the interval.	The count is updated every 2 minutes.
3-7	Reserved	All bytes are set to x00.	

Bytes	Name	Description	When data is sampled or updated
Library - Pooling – General Use Pool (GUP) Container Bytes 320-8255 (32 sets of data x 248 bytes/set = 7936 bytes). These fields contain 32 sets of data describing each of the 32 General Use Pools. Each set of data contains 248 bytes. The following fields define these bytes and are numbered starting with 0. The data for GUP 1 can be found in bytes 320-567; the data for GUP 2 can be found in bytes 568- 815, and so forth.			
0-3	Active Logical Volume	This 4 byte hexadecimal field contains the number of logical volume images resident in the volume pool at the end of the interval. The number is updated and the reported value is the snapshot of that count when statistics are calculated at the end of the interval. To be considered resident in a pool, the logical volume must be on one of the physical volumes assigned to the pool. Cache resident only volumes, although assigned to the pool, are not included.	The count is updated every 60 minutes.
4-7	Active Data	This 4 byte hexadecimal field contains the number of megabytes of logical volume image data managed in the volume pool. The number is updated dynamically and the reported value is the snapshot of that count when statistics are calculated at the end of the interval. To be considered resident in a pool, the logical volume must be on one of the physical volumes assigned to the pool. Cache resident only volumes, although assigned to the pool, are not included. The value is reported in increments of 1 MB (1024x1024). Any volume with a count of less than 1 MB is rounded up to 1 MB.	The count is updated every 60 minutes.
8-11	Data Written to Pool	This 4 byte hexadecimal field represents the number bytes written to the media associated with this pool during the last interval. This is data premigrated from the Tape Volume Cache (TVC), and does not include data moved as part of reclamation. The value is reported in increments of 1 MB (1024x1024). This value is reset to 0 at the beginning of the interval.	The count is incremented when a logical volume's data is written to a pool.
12-15	Data Read from Pool	This 4 byte hexadecimal field represents the number bytes read from the media associated with this pool during the last interval. This is recall data written to the Tape Volume Cache (TVC), and does not include data moved as part of reclamation. The value is reported in increments of 1 MB (1024x1024). This value is reset to 0 at the beginning of the interval.	The count is incremented when a logical volume's data is read from a pool.

Bytes	Name	Description	When data is sampled or updated
16	Device Class	<p>This 1 byte hexadecimal field indicates the device class identifier for the pool. Possible values are:</p> <ul style="list-style-type: none"> ▶ x00: No device. ▶ x11: 3590 Model B1A. ▶ x13: 3590 Model E1A. ▶ x14: 3590 Model H1A. ▶ x20: 3592 Model J1A. This also includes a 3592-E05 that is emulating a 3592-J1A device. ▶ x22: 3592 Model E05. This is for a 3592-E05 that is behaving as a 3590-E05. ▶ x23: 3592 Model E05 (encryption configured). ▶ x24: 3592 Model E06. <p>All other values are reserved. This is the value at the end of the interval.</p>	
17-18	Average time since data began expiring on volumes.	<p>This 2 byte hexadecimal field contains the average age, in days, of the residual data that resides on the stacked volumes assigned to the pool. Statically assigned and borrowed volumes are included in this calculation. A physical volume has residual data on it if it is not full. This value is calculated based on the date a volume transitions to not full and the current date.</p> <p>This data does not include stacked volumes that are in Read-Only- Recovery or are Unavailable. The data for this field is reported at the end of the interval.</p>	The average age is updated every 360 minutes.
19-20	Maximum time since data began expiring on volumes.	<p>This 2 byte hexadecimal field contains the maximum age, in days, of the residual data that resides on the stacked volumes assigned to the pool. Statically assigned and borrowed volumes are included in this calculation. A physical volume has residual data on it if it is not full. This value is calculated based on the date a volume transitions to not full and the current date. This data does not include stacked volumes that are in Read-Only- Recovery or are Unavailable. The data for this field is reported at the end of the interval.</p>	The maximum age is updated every 360 minutes.
21-22	Average Age of Full Private Volumes	<p>This 2 byte hexadecimal field contains the average age, in days, of private stacked volumes in the pool. Statically assigned and borrowed volumes are included in this calculation. This value is calculated based when the volume is marked as full until it is reclaimed.</p> <p>This data does not include stacked volumes that are in Read-Only- Recovery or are Unavailable. The data for this field is reported at the end of the interval.</p>	The average age is updated every 360 minutes.
23-24	Maximum Age of Full Private Volumes	<p>This 2 byte hexadecimal field contains the maximum age, in days, of private stacked volumes in the pool. Statically assigned and borrowed volumes are included in this calculation. This value is calculated based when the volume is marked as full until it is reclaimed.</p> <p>This data does not include stacked volumes that are in Read-Only- Recovery or are Unavailable.</p>	The maximum age is updated every 360 minutes.
23-51	Reserved	All bytes are set to x00.	

Bytes	Name	Description	When data is sampled or updated
Library - Pooling – GUP - Media Container Relative bytes 32-223 (8 sets of data x 24 bytes/set = 192 bytes). These bytes are relative to the Library - Pooling – GUP Container. These fields contain 8 sets of data describing up to 8 physical media types in the pool. The first media type is referred to as Media Type 0; the second is referred to as Media Type 1 and so on. Each set of data contains 24 bytes. The following fields define these bytes and are numbered starting with 0. The data for media type 0 can be found in relative bytes 32-55; the data for media type 1 can be found in relative bytes 56-79, and so forth.			
0	Physical Media Identifiers	This 1 byte hexadecimal field contains the identifier for the media type associated with following general use pool (GUP) volume counts. Here are the values for this field: <ul style="list-style-type: none"> ▶ x00: No media type for this set of data. ▶ x10: 3590 J media. ▶ x11: 3590 K media. ▶ x20: 3592 JA media. ▶ x21: 3592 JW media (reserved for future use). ▶ x22: 3592 JJ media. ▶ x23: 3592 JR media (reserved for future use). ▶ x24: 3592 JB media. ▶ x25: 3592 JX media (reserved for future use). All other values are reserved. This is the value at the end of the interval.	
1-2	Scratch Volume Count	This 2 byte hexadecimal field indicates the number of scratch stacked volumes associated with this pool of the media type indicated. The count includes volumes statically assigned to the pool at the end of the interval. It does not include any volumes that are borrowed from the CSP or are waiting to be erased because of a Security Data Erase. This data does not include stacked volumes that are in Read-Only- Recovery or are Unavailable. This is the value at the end of the interval.	The count is updated every 60 minutes.
3-4	Private Volume Count	This 2 byte hexadecimal field indicates the number of private stacked volumes associated with this pool of the media type indicated. The count includes volumes statically assigned to the pool at the end of the interval. It does not include any volumes that are borrowed from the CSP or are waiting to be erased because of a Security Data Erase. This data does not include stacked volumes that are in Read-Only- Recovery or are Unavailable. This is the value at the end of the interval.	The count is updated every 60 minutes.

Bytes	Name	Description	When data is sampled or updated
5-6	Waiting for Security Data Erase	<p>This 2 byte hexadecimal field indicates the number of stacked volumes associated with this pool of the media type indicated that are waiting for Security Data Erase. The count includes volumes statically assigned to the pool and any volumes borrowed from the CSP.</p> <p>This data does not include stacked volumes that are in Read-Only- Recovery or are Unavailable.</p> <p>This is the value at the end of the interval.</p>	The count is updated every 60 minutes.
7-8	Borrowed Scratch Volume Count	<p>This 2 byte hexadecimal field indicates the number of stacked volumes associated with this pool of the media type indicated that are in scratch status that have been borrowed from the CSP.</p> <p>This data does not include stacked volumes that are in Read-Only- Recovery or are Unavailable.</p> <p>This is the value at the end of the interval.</p>	The count is updated every 60 minutes.
9-10	Borrowed Private Volume Count	<p>This 2 byte hexadecimal field indicates the number of stacked volumes associated with this pool of the media type indicated that are in private status that have been borrowed from the CSP.</p> <p>This data does not include stacked volumes that are in Read-Only- Recovery or are Unavailable.</p> <p>This is the value at the end of the interval.</p>	The count is updated every 60 minutes.
11-12	Read Only Recovery Volume Count	<p>This 2 byte hexadecimal field indicates the number of stacked volumes associated with this pool of the media type indicated that are in read only recovery status. The count includes stacked volumes that are both statically assigned to the pool and are borrowed stacked volumes.</p> <p>This is the value at the end of the interval.</p>	The count is updated every 60 minutes.
13-14	Unavailable Volume Count	<p>This 2 byte hexadecimal field indicates the number of stacked volumes associated with this pool of the media type indicated that are in unavailable status. The count includes stacked volumes that are both statically assigned to the pool and are borrowed stacked volumes. This is the value at the end of the interval.</p>	The count is updated every 60 minutes.
15-23	Reserved	All bytes are set to x00.	
Pooling – GUP - Reclaim Container Relative bytes 224-239 (1 set of data x 16 bytes/set = 16 bytes). These bytes are relative to the Library - Pooling – GUP Container. These fields contain Reclaim information for this GUP.			
224	Reclaim Threshold	This 1 byte hexadecimal field contains the reclaim threshold percentage for the pool as defined at the end of the interval.	The threshold is updated when its value changes.

Bytes	Name	Description	When data is sampled or updated
225	Reclaim Pool	This 1 byte hexadecimal field contains the reclaim pool for the pool as defined at the end of the interval.	The pool is updated when its value changes.
226-227	Last Access Policy	This 2 byte hexadecimal field indicates, in days, when a physical volume is eligible for reclamation based on last access. A volume is eligible for reclamation when the number of days specified has elapsed since any data on the volume has been accessed because of a recall. If this field contains a value of 0, it is not used as criteria for reclaim. This is the value at the end of the interval.	The policy is updated when its value changes.
228-229	Last Written Policy	This 2 byte hexadecimal field indicates, in days, when a physical volume is eligible for reclamation based on when it was last written to. A volume is eligible for reclamation when the number of days specified has elapsed since any data has been written to the volume. If this field contains a value of 0, it is not used as criteria for reclaim. This is the value at the end of the interval.	The policy is updated when its value changes.
230-231	Last Data Invalidation Policy	This 2 byte hexadecimal field indicates, in days, when a physical volume is eligible for reclamation based on when data was last invalidated on it. A volume is eligible for reclamation when the number of days specified has elapsed since any data has invalidated on the volume. If this field contains a value of 0, it is not used as criteria for reclaim. This is the value at the end of the interval.	The policy is updated when its value changes.
232	Minimum Active Data Percentage Policy	This 1 byte hexadecimal field indicates the minimum active data percentage a physical volume's active data must fall below before it can be reclaimed using the days since last data invalidation reclamation policy. This is the value at the end of the interval.	The policy is updated when its value changes.
233-234	Force Erasure Policy	This 2 byte hexadecimal field indicates the number of days before a physical volume must complete the erase process. The time starts when the first data is invalidated on the volume. Supported values are 0 to 365. If this field contains a value of 0, it is not used as criteria for reclaim. When this field contains a non-zero value, the pool is operating in the secure data erasure mode. Any volume that is reclaimed in the pool is physically erased before being returned to scratch. This is the value at the end of the interval.	The policy is updated when its value changes.
235-239	Reserved	All bytes are set to x00.	

Bytes	Name	Description	When data is sampled or updated
Pooling – GUP - Properties Container Relative bytes 240-247 (1 set of data x 8 bytes/set = 8 bytes). These bytes are relative to the Library - Pooling – GUP Container. These fields contain pool property information for this GUP.			
240	Pool Type	This 1 byte hexadecimal field indicates the type of the pool. The valid values are: <ul style="list-style-type: none"> ► x80: Reserved. ► x40: Reserved. ► 20: Export Pool All other values are reserved. This is the value at the end of the interval.	The pool type is updated when its value changes.
241	Return Borrowed Volumes	This 1 byte hexadecimal field indicates whether volumes borrowed from the CSP should be returned to the CSP after they are no longer needed by the pool. The valid values are: <ul style="list-style-type: none"> ► x00: Do not return borrowed volumes. ► x01: Return borrowed volumes. All other values are reserved. This is the value at the end of the interval.	The policy is updated when its value changes.

Bytes	Name	Description	When data is sampled or updated
242	First Media Types to Borrow	<p>This 1 byte bit-mapped field indicates the first media types to be borrowed from the CSP if additional scratch physical volumes are needed by the pool. When a bit is set to one (1), then the media type defined by the corresponding Library-Pooling-GUP-Media – Physical Media Identifier set can be used as one of the first media types to borrow from the CSP. When a bit is set to zero (0), the corresponding physical media type, if any, is not one of the first media types to borrow. For example, if the first Physical Media Identifier (first set of 24 bytes) of the Library-Pooling-GUP-Media container indicates J media, the second (second set of 24 bytes) indicates K media, the third (third set of 24 bytes) indicates JA media, and the fourth (fourth set of 24 bytes) indicates JJ media, a value of x20 in this field would indicate the media type identified in the third set of 24 bytes should be borrowed first. In this example, JA media will be borrowed first. A value of x00 indicates borrowing is turned off for this pool. Valid values are:</p> <ul style="list-style-type: none"> ▶ x80: The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the first set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a first choice. ▶ x40: The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the second set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a first choice. ▶ x20: The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the third set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a first choice. 	The media type is updated when its value changes.
		<ul style="list-style-type: none"> ▶ x10: The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the fourth set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a first choice. ▶ x08: The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the fifth set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a first choice. ▶ x04: The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the sixth set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a first choice. ▶ x02: The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the seventh set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a first choice. ▶ x01: The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the eighth set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a first choice. <p>This is the value at the end of the interval.</p>	

Bytes	Name	Description	When data is sampled or updated
243	Second Media Types to Borrow	<p>This 1 byte bit-mapped field indicates the second media types to be borrowed from the CSP if additional scratch physical volumes are needed by the pool. When a bit is set to one (1), then the media type defined by the corresponding Library-Pooling-GUP-Media – Physical Media Identifier set can be used as one of the second media types to borrow from the CSP. When a bit is set to zero (0), the corresponding physical media type, if any, is not one of the second media types to borrow. For example, if the first Physical Media Identifier (first set of 24 bytes) of the Library-Pooling-GUP-Media container indicates J media, the second (second set of 24 bytes) indicates K media, the third (third set of 24 bytes) indicates JA media, and the fourth (fourth set of 24 bytes) indicates JJ media, a value of x10 in this field would indicate the media type identified in the fourth set of 24 bytes should be borrowed second. In this example, JJ media will be borrowed second. A value of x00 indicates a second media type to borrow is not specified. Valid values are:</p> <ul style="list-style-type: none"> ▶ x80: The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the first set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a second choice. ▶ x40: The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the second set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a second choice. ▶ x20: The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the third set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a second choice. ▶ x10: The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the fourth set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a second choice. ▶ x08: The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the fifth set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a second choice. 	The media type is updated when its value changes.
		<ul style="list-style-type: none"> ▶ x04: The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the sixth set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a second choice. ▶ x02: The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the seventh set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a second choice. ▶ x01: The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the eighth set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a second choice. <p>This is the value at the end of the interval.</p>	
244-247	Reserved	All bytes are set to x00.	

hNode Grid Historical record

The hNode Grid Historical record has the following nested structure:

- ▶ Header
- ▶ Grid Container
 - Grid-Cluster 0 Container
 - Grid-Cluster 1 Container (If installed)
 - Grid-Cluster 2 Container (If installed)
 - Grid-Cluster 3 Container (If installed)
 - Grid-Cluster 4 Container (If installed)
 - Grid-Cluster 5 Container (If installed)
 - Grid-Cluster 6 Container (If installed)
 - Grid-Cluster 7 Container (If installed)

Table E-10 shows the hNode Grid Historical record.

Table E-10 hNode Grid Historical record

Bytes	Name	Description	When data is sampled or updated
0-1	Length	This 2 byte hexadecimal field contains the length of this record. The length includes these 2 bytes.	
2	Version	This 1 byte hexadecimal field contains the version of the data presented in this record. Initially the version is set to x02.	
3	Data Type	This 1 byte hexadecimal field indicates the type of data contained in this record. For this record, the value is set to x10, indicating that this is an hNode HSM Point-in-Time record.	
4	Node ID	This 1 byte hexadecimal field indicates the hNode ID that this interval's data represents. Valid values are x00 – x01.	
5	Cluster ID	This 1 byte hexadecimal field indicates the Cluster ID of which this hNode is a part. Valid values are x00 – x07.	
6-7	Interval Duration	This 2 byte hexadecimal field indicates the interval in seconds that this interval's data was taken over.	
8-11	Time Stamp	This 4 byte hexadecimal field indicates the end time of the interval this data was taken over. This value is the time in seconds since the beginning of the Epoch (00:00:00 UTC, January 1, 1970).	
12-15	Machine Type	This 4 byte EBCDIC field contains this node's machine type. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "3957".	
16-18	Machine Model	This 3 byte EBCDIC field contains this node's machine model. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "V06".	
19-26	Machine Serial Number	This 8 character EBCDIC field contains the serial number of this node. This field is left justified and padded with EBCDIC blanks. The format is XX-YYYYY, where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed.	

Bytes	Name	Description	When data is sampled or updated
27-34	VE Code Level	This 8 byte hexadecimal field contains the code level of the TS7700 Virtualization Engine. The 8 bytes are actually four 2 byte fields. Each 2 byte field represents a portion of the code level. The VE code level is expressed as Version.Release.Modification.Fix in a decimal form. For example, the code level of 8.0.0.104 would be represented in the 8 bytes as x0008000000000068.	
35-39	Grid Library Sequence Number	This 5 character EBCDIC field contains the Library Sequence Number of the Grid (Composite) library.	
40-44	Distributed Library Sequence Number	This 5 character EBCDIC field contains the Distributed Library Sequence Number for this Distributed Library ID.	
45-63	Reserved	All bytes are set to x00.	
Grid Container Bytes 64-95. The data from each hNode will be the same within a cluster.			
64-67	Logical Volumes for Copy	This 4 byte hexadecimal field indicates the number of logical volumes that are scheduled to be copied to this cluster. This is the value at the end of the interval.	The count is updated every 5 minutes.
68-71	Data to Copy	This 4 byte hexadecimal field indicates the amount of data that is scheduled to be copied to this cluster. This represents the amount of data contained in the logical volumes that are scheduled to be copied. The value is reported in increments of 1 MB (1024x1024). Any residual data is rounded up to 1 MB. This is the value at the end of the interval.	The count is updated every 5 minutes.
72-75	Average Deferred Queue Age	This 4 byte hexadecimal field indicates the average age, in seconds, of the logical volumes in the deferred copy queue destined to be copied to this cluster. This is the value at the end of the interval.	The age is updated every 5 minutes.
76-79	Average Immediate Queue Age	This 4 byte hexadecimal field indicates the average age, in seconds, of the logical volumes in the immediate copy queue destined to be copied to this cluster. This is the value at the end of the interval.	The age is updated every 5 minutes.
80	Number of Clusters	This 1 byte hexadecimal field indicates the number of clusters in the Grid. This field can be used to determine how many Grid-cluster containers will be attached to this record. There is a maximum of 8 clusters. This is the value at the end of the interval. Note: A maximum of four clusters is supported at this time.	
81-95	Reserved	All bytes are set to x00.	

Bytes	Name	Description	When data is sampled or updated
Grid-Cluster Container Bytes 96 and up (Number of clusters x 256 bytes/set). For example, if there are 3 clusters in the Grid, there will be 3 sets of data. There is a maximum of 8 clusters. This next segment of the record contains one set of data for each cluster in the Grid as defined in byte 80. Each set of data contains 256 bytes. The data for the first cluster (Cluster 0) can be found in bytes 96-351; the second cluster's (Cluster 1) data can be found in bytes 352-447, and so forth.			
0-3	Data Transferred into a Cluster's Cache from other Clusters as part of a Copy Operation	A 4 byte field indicating the number of bytes transferred to this cluster's cache from another cluster as part of a copy operation during this interval. The cluster transferred from is determined by the relative position of the container segment within the record. A value of x00000000 is indicated for a cluster's own value. The value is reported in increments of 1 MB (1024x1024). Any residual data is rounded up to 1 MB. This value is reset to 0 at the beginning of the interval.	The count is incremented every time a block of data is transferred to this cluster.
4-31	Reserved	All bytes are set to x00.	
32-35	Data Transferred From a Cluster's Cache To Other Clusters as part of a Copy Operation	A 4 byte field indicating the number of bytes transferred from this cluster's cache to another cluster as part of a copy operation during this interval. The cluster transferred to is determined by the relative position of the container segment within the record. A value of x00000000 is indicated for a cluster's own value. The value is reported in increments of 1 MB (1024x1024). Any residual data is rounded up to 1 MB. This value is reset to 0 at the beginning of the interval.	The count is incremented every time a block of data is transferred.
36-63	Reserved	All bytes are set to x00.	
64-67	Logical Mounts Directed to other Clusters	A 4 byte field indicating the number of logical mounts from all vNodes in this cluster that were satisfied by accessing another cluster. The cluster receiving the mounts is determined by the relative position of the container segment within the record. A value of x00000000 is indicated for a cluster's own value. A logical mount is counted when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.	The count is incremented when a logical mount is completed.
68-95	Reserved	All bytes are set to x00.	
96-99	Data Transferred into a Cluster's Cache from other Clusters as part of a Remote Write Operation	A 4 byte field indicating the number of bytes transferred to this cluster's cache from another cluster as part of a remote write operation during this interval. The source cluster is determined by the relative position of the container segment within the record. A value of x00000000 is indicated for a cluster's own value. The value is reported in increments of 1 MB (1024x1024). Any residual data is rounded up to 1 MB. This value is reset to 0 at the beginning of the interval.	The count is incremented every time a block of data is transferred.
100-127	Reserved	All bytes are set to x00.	

Bytes	Name	Description	When data is sampled or updated
128-131	Data Transferred from a Cluster's Cache To Other Clusters as part of a Remote Read operation	A 4 byte field indicating the number of bytes transferred from of this cluster's cache to another cluster as part of a remote read operation during this interval. The Cluster doing the remote read is determined by the relative position of the container segment within the record. A value of x00000000 is indicated for a cluster's own value. The value is reported in increments of 1 MB (1024x1024). Any residual data is rounded up to 1 MB. This value is reset to 0 at the beginning of the interval.	The count is incremented every time a block of data is transferred.
132-255	Reserved	All bytes are set to x00.	



Sample JCLs

In this appendix, we provide the JCL of sample jobs to collect statistical data and perform volume reporting. You may use the sample jobs to perform the following tasks:

- ▶ Obtain statistical data from the IBM Virtualization Engine TS7700 using Bulk Volume Information Retrieval (BVIR) jobs
- ▶ Analyze Point-in-Time and Historical statistics records obtained through BVIR with VEHSTATS
- ▶ Create a Copy Export Volume
- ▶ Support data migration into the TS7700

Notes: You can find tailored JCL to run BVIR jobs and to analyze the data using VEHSTATS in the IBMTTOOLS libraries. To access the IBM Tape Tools, go to the following address:

<ftp://ftp.software.ibm.com/storage/tapetool/>

For the most current information about BVIR, see the *IBM Virtualization Engine TS7700 Series Bulk Volume Information Retrieval Function User's Guide*, located at the following address:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101094>

BVIR jobs to obtain historical data

The following JCL can be used to request BVIR data:

BVIRHSTS	Requests the BVIR Historical data for one or more days and writes them to the SMF log file. These SMF Records can then be used as input to VEHSTATS. See Example F-1.
BVIRHSTU	Requests the BVIR Historical data for one or more days and writes them to a disk data set with RECFM=U. The output of this job can then be used as input to VEHSTATS. See Example F-2 on page 892.
BVIRHSTV	Requests the BVIR Historical data for one or more days and writes them to a disk data set with RECFM=VB. The output of this job can then be used as input to VEHSTATS. See Example F-3 on page 893.

These jobs are also available as members in *userid*.IBMT00LS.JCL after you have installed the IBMT00LS.exe on your host. See 9.11, “IBM Tape Tools” on page 706.

After you have run one of these jobs, you can create a variety of reports using VEHSTATS. See “VEHSTATS Reports” on page 904.

BVIRHSTS

Example F-1 lists the JCL in *userid*.IBMT00LS.JCL member BVIRHSTS.

Example F-1 BVIRHSTS JCL to obtain BVIR Historical data

```
//ITS01  JOB  CONSOLE,
//          MSGCLASS=H,MSGLEVEL=(1,1),CLASS=B,
//          TIME=1440,REGION=2M
//*
/*JOBPARM SYSAFF=*
//*
/* THIS IS A TWO JOB MEMBER TO ACCOMMODATE EITHER JES2 OR JES3.
/* BVIR DATA WILL ONLY BE WRITTEN TO A TAPE AFTER THE INITIAL
/* DISMOUNT AND RE-MOUNT FOR READ.
//*
/* IF YOU HAVE MULTIPLE TS7740 GRIDS, YOU MUST HAVE A SEPARATE
/* STORAGE GROUP DEFINED FOR EACH GRID IN ORDER TO REQUEST STATISTICS
/* FROM EACH ONE.  USE AN ACS ROUTINE TO SELECT THE TARGET GRID.
//*
/* THIS JOB ISSUES THE BULK VOLUME INFORMATION (BVIR) REQUEST FOR
/* HISTORICAL STATISTICS FROM THE GRID ASSOCIATED WITH THE VIRTUAL
/* DRIVE ADDRESS USED.  THE BVIR FEATURE MUST BE ACTIVATED ON THE
/* VTS RECEIVING THE REQUEST.  THE FINAL OUTPUT IS DIRECTED TO SMF.
/* HISTORICAL STATISTICS FOR ALL CLUSTERS IN A GRID ARE RETURNED
/* FROM A SINGLE BVIR REQUEST TO ANY OF THE CLUSTERS.
/* NEXT, RUN VEHSTATS TO GET REPORTS.
//*
//PUTBVIR PROC USERHLQ=USERID,          HI-LEVEL FOR USER DATA FILES
//          TOOLHLQ=TOOLID,              HLQ FOR LOAD AND CNTL
//          SITE=SITENAME,              2ND LEVEL QUALIFIER
//          GRIDID=GRID#,              GRID SERIAL NUMBER TO BE PART OF DSN
//          UNIT=VTAPE                  UNITNAME ON THIS VTS
//*
//STEP1   EXEC PGM=IEFBR14
//DEL1    DD  UNIT=(&UNIT,,DEFER),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..BVIRTAPE
```

```

//*
//STEP2 EXEC PGM=GETHIST          ISSUE HISTORICAL STATS REQUEST
//STEPLIB DD DISP=SHR,DSN=&TOOLHLQ..IBMTTOOLS.LOAD
//SYSLIST DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//BVIRREQ DD DSN=&USERHLQ..&SITE..#&GRIDID..BVIRTAPE,
//          UNIT=&UNIT,LABEL=(,SL),DISP=(NEW,CATLG),
//          DCB=(RECFM=F,LRECL=80,BLKSIZE=80,TRTCH=NOCOMP)
// PEND
//*
//RUNPROC EXEC PUTBVIR
//STEP2.SYSCNTL DD DISP=SHR,DSN=&TOOLHLQ..IBMTTOOLS.JCL(EXPIRE)
// DD *
*
* SEE MEMBER, VEHDATES, FOR MORE DETAIL ON DATES
*
SDATE= 01SEP2010;          USE HERE AS DDMONYEAR
EDATE= 30SEP2010;          USE HERE AS DDMONYEAR
*SDATE= TODAY- 1;          THIS FORMAT PULLS STATS FROM PREVIOUS DAY
*EDATE= TODAY;
*SDATE= LASTWEEK;          OR LASTMONTH WITH + OR - OPTIONS ALSO
*
* SEE MEMBER, VEHDATES, FOR MORE DETAIL ON DATES
*
//*
//ITS01 JOB CONSOLE,
//          MSGCLASS=H,MSGLEVEL=(1,1),CLASS=B,
//          TIME=1440,REGION=2M
//*
/*JOBPARM SYSAFF=*
//*
//COPYBVIR PROC USERHLQ=USERID,          HI-LEVEL FOR USER DATA FILES
//          TOOLHLQ=TOOLID,          HLQ FOR LOAD AND CNTL
//          SITE=SITENAME,          2ND LEVEL QUALIFIER
//          GRIDID=GRID#          GRID SERIAL NUMBER TO BE PART OF DSN
//*
//STEP3 EXEC PGM=CPYHIST APF LIBRARY NEEDED IF WRITING TO SMF
//STEPLIB DD DISP=SHR,DSN=USER.AUTH.LIB
//SYSLIST DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//RECLIST DD SYSOUT=*
//SYSUT1 DD DSN=&USERHLQ..&SITE..#&GRIDID..BVIRTAPE,
//          DCB=(RECFM=U,BLKSIZE=22000),
//          DISP=(OLD,DELETE)
//SYSUT2 DD DSN=NULLFILE,DCB=(RECFM=U,BLKSIZE=22000),
//          DISP=(NEW,CATLG),SPACE=(CYL,(40,25),RLSE),UNIT=SYSDA
// PEND
//*
//RUNPROC EXEC COPYBVIR
//STEP3.SYSCNTL DD DISP=SHR,DSN=&TOOLHLQ..IBMTTOOLS.JCL(EXPIRE)
// DD *
SMFNUM = 194;  USER SELECTABLE SMF # FOR QUARTER HOUR STATISTICS
* THE SMF TIME STAMP WILL BE THE QUARTER HOUR DATE/TIME ORIGINALLY
* WRITTEN EVEN IF PULLED SEVERAL DAYS LATER.
//

```

BVIRHSTU

Example F-2 lists the JCL in *userid.IBMTTOOLS.JCL* member BVIRHSTU.

Example F-2 BVIRHSTU JCL to obtain BVIR Historical data

```
//ITS01  JOB  CONSOLE,
//          MSGCLASS=H,MSGLEVEL=(1,1),CLASS=B,
//          TIME=1440,REGION=2M
//*
/*JOBPARM SYSAFF=*
//*
/** THIS IS A TWO JOB MEMBER TO ACCOMODATE EITHER JES2 OR JES3.
/** BVIR DATA WILL ONLY BE WRITTEN TO A TAPE AFTER THE INITIAL
/** DISMOUNT AND RE-MOUNT FOR READ.
/**
/** IF YOU HAVE MULTIPLE TS7740 GRIDS, YOU MUST HAVE A SEPARATE
/** STORAGE GROUP DEFINED FOR EACH GRID IN ORDER TO REQUEST STATISTICS
/** FROM EACH ONE.  USE AN ACS ROUTINE TO SELECT THE TARGET GRID.
/**
/** THIS JOB ISSUES THE BULK VOLUME INFORMATION (BVIR) REQUEST FOR
/** HISTORICAL STATISTICS FROM THE VTS ASSOCIATED WITH THE VIRTUAL
/** DRIVE ADDRESS USED.  THE BVIR FEATURE MUST BE ACTIVATED ON THE
/** VTS RECEIVING THE REQUEST.  THE FINAL OUTPUT IS WRITTEN TO A
/** DISK DATA SET AS RECFM=U.
/** HISTORICAL STATISTICS FOR ALL CLUSTERS IN A GRID ARE RETURNED
/** FROM A SINGLE BVIR REQUEST TO ANY OF THE CLUSTERS.
/** NEXT, RUN VEHSTATS TO GET REPORTS.
/**
//PUTBVIR  PROC USERHLQ=USERID,          HI-LEVEL FOR USER DATA FILES
//          TOOLHLQ=TOOLID,              HLQ FOR LOAD AND CNTL
//          SITE=SITENAME,              2ND LEVEL QUALIFIER
//          GRIDID=GRID#,              GRID SERIAL NUMBER TO BE PART OF DSN
//          UNIT=VTAPE                  UNITNAME ON THIS VTS
//*
//STEP1    EXEC PGM=IEFBRI4
//DEL1     DD  UNIT=(&UNIT,,DEFER),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..&GRIDID..BVIRTAPE
//*
//STEP2    EXEC PGM=GETHIST              ISSUE HISTORICAL STATS REQUEST
//STEPLIB  DD  DISP=SHR,DSN=&TOOLHLQ..IBMTTOOLS.LOAD
//SYSLIST  DD  SYSOUT=*
//SYSUDUMP DD  SYSOUT=*
//BVIRREQ  DD  DSN=&USERHLQ..&SITE..&GRIDID..BVIRTAPE,
//          UNIT=&UNIT,LABEL=(,SL),DISP=(NEW,CATLG),
//          DCB=(RECFM=F,LRECL=80,BLKSIZE=80,TRTCH=NOCOMP)
//  PEND
//*
//RUNPROC  EXEC PUTBVIR
//STEP2.SYSCNTL DD DISP=SHR,DSN=&TOOLHLQ..IBMTTOOLS.JCL(EXPIRE)
//          DD  *
*
*  SEE MEMBER, VEHDATES, FOR MORE DETAIL ON DATES
*
  SDATE= 01SEP2010;          USE HERE AS DDMONYEAR
  EDATE= 30SEP2010;          USE HERE AS DDMONYEAR
*SDATE= TODAY- 1;          THIS FORMAT PULLS STATS FROM PREVIOUS DAY
*EDATE= TODAY;
*SDATE= LASTWEEK;          OR LASTMONTH WITH + OR - OPTIONS ALSO
*
```

```

* SEE MEMBER, VEHDATES, FOR MORE DETAIL ON DATES
/**
//ITS01  JOB  CONSOLE,
//      MSGCLASS=H,MSGLEVEL=(1,1),CLASS=B,
//      TIME=1440,REGION=2M
/**
/*JOBPARM SYSAFF=*
/**
//COPYBVIR PROC USERHLQ=USERID,      HI-LEVEL FOR USER DATA FILES
//      TOOLHLQ=TOOLID,      HLQ FOR LOAD AND CNTL
//      SITE=SITENAME,      2ND LEVEL QUALIFIER
//      GRIDID=GRID#,      GRID SERIAL NUMBER TO BE PART OF DSN
//      SDATE=YYMMDD,      YYMMDD BEGINNING DATE
//      EDATE=YYMMDD      YYMMDD ENDING DATE
/**
//STEP1  EXEC  PGM=IEFBR14
//DEL2   DD  UNIT=SYSDA,DISP=(MOD,DELETE),SPACE=(TRK,1),
//        DSN=&USERHLQ..&SITE..#&GRIDID..HSTU.D&SDATE..D&EDATE
/**
//STEP3  EXEC  PGM=CPYHIST
//STEPLIB DD  DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.LOAD
//SYSLIST DD  SYSOUT=*
//SYSPRINT DD  SYSOUT=*
//RECLIST DD  SYSOUT=*
//SYSUT1  DD  DSN=&USERHLQ..&SITE..#&GRIDID..BVIRTAPE,
//            DCB=(RECFM=U,BLKSIZE=22000),DISP=(OLD,DELETE)
//SYSUT2 DD  DSN=&USERHLQ..&SITE..#&GRIDID..HSTU.D&SDATE..D&EDATE.,
//            DCB=(RECFM=U,BLKSIZE=22000),UNIT=SYSDA,
//            DISP=(NEW,CATLG),SPACE=(CYL,(40,25),RLSE)
//      PEND
/**
//RUNPROC EXEC COPYBVIR,SDATE=YYMMDD,EDATE=YYMMDD  HERE AS YYMMDD
//STEP3.SYSCNTL DD DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.JCL(EXPIRE)
/**

```

BVIRHSTV

Example F-3 lists the JCL in *userid*.IBMTOOLS.JCL member BVIRHSTV.

Example F-3 BVIRHSTV JCL to obtain BVIR Historical data

```

//ITS01  JOB  CONSOLE,
//      MSGCLASS=H,MSGLEVEL=(1,1),CLASS=B,
//      TIME=1440,REGION=2M
/**
/*JOBPARM SYSAFF=*
/**
/** THIS IS A TWO JOB MEMBER TO ACCOMODATE EITHER JES2 OR JES3.
/** BVIR DATA WILL ONLY BE WRITTEN TO A TAPE AFTER THE INITIAL
/** DISMOUNT AND RE-MOUNT FOR READ.
/**
/** IF YOU HAVE MULTIPLE TS7740 GRIDS, YOU MUST HAVE A SEPARATE
/** STORAGE GROUP DEFINED FOR EACH GRID IN ORDER TO REQUEST STATISTICS
/** FROM EACH ONE.  USE AN ACS ROUTINE TO SELECT THE TARGET GRID.
/**
/** THIS JOB ISSUES THE BULK VOLUME INFORMATION (BVIR) REQUEST FOR
/** HISTORICAL STATISTICS FROM THE VTS ASSOCIATED WITH THE VIRTUAL
/** DRIVE ADDRESS USED.  THE BVIR FEATURE MUST BE ACTIVATED ON THE
/** VTS RECEIVING THE REQUEST.  THE FINAL OUTPUT IS WRITTEN TO A

```

```

/** DISK DATA SET AS RECFM=VB.
/** HISTORICAL STATISTICS FOR ALL CLUSTERS IN A GRID ARE RETURNED
/** FROM A SINGLE BVIR REQUEST TO ANY OF THE CLUSTERS.
/** NEXT, RUN VEHSTATS TO GET REPORTS.
/**
/**PUTBVIR  PROC USERHLQ=USERID,      HI-LEVEL FOR USER DATA FILES
/**      TOOLHLQ=TOOLID,      HLQ FOR LOAD AND CNTL
/**      SITE=SITENAME,      2ND LEVEL QUALIFIER
/**      GRIDID=GRID#,      GRID SERIAL NUMBER TO BE PART OF DSN
/**      UNIT=VTAPE      UNITNAME ON THIS VTS
/**
/**STEP1    EXEC PGM=IEFBRI4
/**DEL1     DD  UNIT=(UNIT,,DEFER),DISP=(MOD,DELETE),
/**          DSN=&USERHLQ..&SITE..#&GRIDID..BVIRTAPE
/**
/**STEP2    EXEC PGM=GETHIST      ISSUE HISTORICAL STATS REQUEST
/**STEPLIB  DD  DISP=SHR,DSN=&TOOLHLQ..IBMTTOOLS.LOAD
/**SYSLIST  DD  SYSOUT=*
/**SYSUDUMP DD  SYSOUT=*
/**BVIRREQ  DD  DSN=&USERHLQ..&SITE..#&GRIDID..BVIRTAPE,
/**          UNIT=&UNIT,LABEL=(,SL),DISP=(NEW,CATLG),
/**          DCB=(RECFM=F,LRECL=80,BLKSIZE=80,TRTCH=NOCOMP)
/**      PEND
/**
/**RUNPROC  EXEC PUTBVIR
/**STEP2.SYSCNTL DD DISP=SHR,DSN=&TOOLHLQ..IBMTTOOLS.JCL(EXPIRE)
/**      DD *
*
* SEE MEMBER, VEHDATES, FOR MORE DETAIL ON DATES
*
SDATE= 01SEP2010;      USE HERE AS DDMONYEAR
EDATE= 30FEB2010;      USE HERE AS DDMONYEAR
*SDATE= TODAY- 1;      THIS FORMAT PULLS STATS FROM PREVIOUS DAY
*EDATE= TODAY;
*SDATE= LASTWEEK;      OR LASTMONTH WITH + OR - OPTIONS ALSO
*
* SEE MEMBER, VEHDATES, FOR MORE DETAIL ON DATES
/**
/**ITS01    JOB  CONSOLE,
/**          MSGCLASS=H,MSGLEVEL=(1,1),CLASS=B,
/**          TIME=1440,REGION=2M
/**
/***JOBPARM SYSAFF=*
/**
/**COPYBVIR PROC USERHLQ=USERID,      HI-LEVEL FOR USER DATA FILES
/**      TOOLHLQ=TOOLID,      HLQ FOR LOAD AND CNTL
/**      SITE=SITENAME,      2ND LEVEL QUALIFIER
/**      GRIDID=GRID#,      GRID SERIAL NUMBER TO BE PART OF DSN
/**      SDATE=YYMMDD,      YYMMDD BEGINNING DATE
/**      EDATE=YYMMDD      YYMMDD ENDING DATE
/**
/**STEP1    EXEC PGM=IEFBRI4
/**DEL2     DD  UNIT=SYSDA,DISP=(MOD,DELETE),SPACE=(TRK,1),
*
/**          DSN=&USERHLQ..&SITE..#&GRIDID..HSTV.D&SDATE..D&EDATE
/**
/**STEP3    EXEC PGM=CPYHIST
/**STEPLIB  DD  DISP=SHR,DSN=&TOOLHLQ..IBMTTOOLS.LOAD
/**SYSLIST  DD  SYSOUT=*
/**SYSPRINT DD  SYSOUT=*

```



```

//RECLIST DD SYSOUT=*
//SYSUT1 DD DSN=&USERHLQ..&SITE..#&GRIDID..BVIRTAPE,
//          DCB=(RECFM=U,BLKSIZE=22000),DISP=(OLD,DELETE)
//SYSUT2 DD DSN=&USERHLQ..&SITE..#&GRIDID..HSTV.D&SDATE..D&EDATE.,
//          DCB=(RECFM=VB,BLKSIZE=22000,LRECL=21996),UNIT=SYSDA,
//          DISP=(NEW,CATLG),SPACE=(CYL,(40,25),RLSE)
// PEND
//*
//RUNPROC EXEC COPYBVIR,SDATE=YMMDD,EDATE=YMMDD HERE AS YMMDD
//STEP3.SYSCNTL DD DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.JCL(EXPIRE)
//*
//

```

Additional BVIR reporting

The IBM Tape Tools libraries also provide JCL for additional reports that can be requested directly from the TS7700. You can also find the JCL in the IBM Tape Tools libraries.

Volume Map report

The Volume Map report shows the relationship between physical and logical volumes. In a TS7700 Virtualization Engine Grid configuration, separate requests must be issued to each cluster to obtain the map of all physical volumes.

Note: If this is an IBM Virtualization Engine TS7720 cluster, the following record is returned:

'NOT SUPPORTED IN A DISK-ONLY TS7700 VIRTUALIZATION ENGINE'

Example F-4 shows the JCL to obtain the Volume Map report, which is also contained in *userid*.IBMTOOLS.JCL member BVIRVTS.

Example F-4 JCL to obtain the Volume Map report

```

//ITS01 JOB CONSOLE,
//      MSGCLASS=H,MSGLEVEL=(1,1),CLASS=B,
//      TIME=1440,REGION=2M
/*JOBPARM SYSAFF=*
/*
/* TO INSURE THAT BVIR REQUESTS ARE SATISFIED FROM THE PROPER
/* CLUSTER, YOU SHOULD HAVE MANAGEMENT CLASSES FOR RNN(CLO),NRN(CL1)
/* AND NNR(CL2) SO THAT ONLY THE TARGET CLUSTER WILL HAVE A COPY OF
/* THE BVIR VOLUME. YOU CAN'T CONTROL WHERE THE INITIAL SCRATCH
/* MOUNT IS SATISFIED, BUT YOU CAN CONTROL WHICH TVC THE VOLUME WILL
/* BE IN WHEN THE SUBSEQUENT SPECIFIC MOUNT IS DONE. THE SPECIFIC
/* MOUNT COLLECTS THE BVIR INFORMATION, NOT THE INITIAL SCRATCH MOUNT.
/*
/* IF YOU HAVE MULTIPLE TS7740 GRIDS, YOU MUST HAVE A SEPARATE
/* STORAGE GROUP DEFINED FOR EACH GRID IN ORDER TO ALLOCATE ON THE
/* DESIRED GRID. USE AN ACS ROUTINE TO SELECT THE TARGET GRID.
/*
/* IF YOU ARE RUNNING JES3, YOU MUST RUN STEP3 AS A SEPARATE JOB
/* IN ORDER TO FORCE THE DISMOUNT OF THE TAPE IN STEP2. BVIR DATA
/* WILL ONLY BE WRITTEN TO A TAPE AFTER THE INITIAL DISMOUNT AND
/* RE-MOUNT FOR READ.

```

```

/**
/** THIS JOB ISSUES THE BULK VOLUME INFORMATION REQUEST (BVIR) FOR
/** ALL VIRTUAL VOLUMES BELONGING TO THE VTS ASSOCIATED WITH THE
/** VIRTUAL DRIVE ADDRESS USED. THE BVIR FEATURE MUST BE ACTIVATED
/** ON THE VTS RECEIVING THE REQUEST. VTS CODE 7.4 OR GREATER NEEDED
/** IN ORDER TO PROVIDE THE VIRTUAL VOLUME SIZE.
/** IF YOU ARE RUNNING AGAINST A PTP AND GETTING DATA FOR THE BVIRRPT
/** JOB, YOU NEED TO RUN THIS JOB TWICE, ONCE FOR EACH VTS.
/**
/** NEXT, RUN BVIRRPT TO GET REPORTS OR BVIRMCB TO SEE HOW MANY
/** PHYSICALS WERE USED TO CONTAIN A LIST OF LOGICAL VOLSERS.
/** OR, RUN THE PRESTAGE JOB TO CAUSE A LIST OF VOLSERS TO BE STAGED.
/**
/**BVIRVTS PROC USERHLQ=USERID,          HI-LEVEL FOR USER DATA FILES
/**      TOOLHLQ=TOOLID,          HLQ FOR LOAD AND CNTL
/**      SITE=SITENAME,          2ND LEVEL QUALIFIER
/**      TYPE=,                  JCL WILL REQUEST TYPE LATER
/**      MC=,                    DIRECT TO SPECIFIC VTS IN PTP
/**      VTSID=CLO,              USE CLO, CL1, CL2, ETC TO BE PART OF DSN
/**      UNIT=VTAPE              UNITNAME ON THIS VTS
/**
/**STEP1  EXEC PGM=IEFBRI4
/**DEL1   DD  UNIT=(&UNIT,,DEFER),DISP=(MOD,DELETE),
/**          DSN=&USERHLQ..&SITE..&VTSID..BVIR&TYPE
/**DEL2   DD  UNIT=SYSDA,DISP=(MOD,DELETE),SPACE=(TRK,1),
/**          DSN=&USERHLQ..&SITE..&VTSID..&TYPE.FILE
/**
/**STEP2  EXEC PGM=IEBGENER
/**SYSPRINT DD  SYSOUT=*
/**SYSUT1  DD  DISP=SHR,DSN=&TOOLHLQ..IBMTTOOLS.CNTL(BVIR&TYPE)
/**SYSUT2  DD  DSN=&USERHLQ..&SITE..&VTSID..BVIR&TYPE,MGMTCLAS=&MC,
/**          UNIT=&UNIT,LABEL=(,SL),DISP=(NEW,CATLG),
/**          DCB=(RECFM=FB,BLKSIZE=80,LRECL=80,TRTCH=NOCOMP)
/**SYSIN   DD  DUMMY
/**
/**STEP3  EXEC PGM=IEBGENER
/**SYSPRINT DD  SYSOUT=*
/**SYSUT1  DD  DSN=&USERHLQ..&SITE..&VTSID..BVIR&TYPE,
/**          DISP=OLD
/**SYSUT2  DD  DSN=&USERHLQ..&SITE..&VTSID..&TYPE.FILE,
/**          DISP=(NEW,CATLG),SPACE=(CYL,(1,3)),UNIT=SYSDA,
/**          DCB=(DSORG=PS,RECFM=FB,LRECL=80,BLKSIZE=0)
/**SYSIN   DD  DUMMY
/**      PEND
/**
/**GETVOLS EXEC BVIRVTS,TYPE=VOL,VTSID=CLO,MC=MCVTSO
/**
/**

```

Cache Contents report

The Cache Contents report shows what cartridges reside in cache of one specific cluster.

You can use the same JCL as shown in Example F-4 on page 895 for the cache report by replacing the last statement written in bold with the statement listed in Example F-5, which creates a report for Cluster 0.

Example F-5 JCL to obtain Cache Contents report

```
//GETCACHE EXEC BVIRVTS,TYPE=CACH,VTSID=CLO,MC=MCVTSO
```

Change the following parameters to obtain this report from each of the clusters in the Grid:

- ▶ VTSID=
- ▶ MC=

Note: Cache contents report refers to the specific cluster to which the request volume was written. In a TS7700 Virtualization Engine Grid configuration, separate requests must be issued to each cluster to obtain the cache contents of all of the clusters.

Copy Audit report

Prior to a Cluster removal from a Grid, the Copy Audit request can be used to ensure that all logical volumes are copied in the remaining Grid members, or in a condition where one of the Grid members is no longer available (a site disaster or a test procedure) and where it is necessary to determine which volumes (if any) on the remaining TS7700 Virtualization Engines do not have a valid copy.

To obtain the Copy Audit report, you can use the same job as shown in Example F-4 on page 895, simply replace the last statement written in bold by the statement shown in Example F-6 and update the following parameters:

- ▶ VTSID=
- ▶ MC=

Example F-6 JCL to obtain Copy Audit report

```
//GETAUD EXEC BVIRVTS,TYPE=AUD,VTSID=C00,MC=
```

Volume Status report

The Volume Status report shows the logical volumes status in the cluster and within the Grid. Example F-7 shows the JCL that is used to obtain the Volume Status report. The JCL is also available in member BVIRMES in *userid*.IBMT00LS.JCL after you have installed the IBM tape tools.

Example F-7 JCL to obtain Volume Status report

```
//ITS01      JOB  CONSOLE,
//          MSGCLASS=H,MSGLEVEL=(1,1),CLASS=B,
//          TIME=1440,REGION=2M
/*JOBPARM  SYSAFF=*
/*
/*
/* TO INSURE THAT BVIR REQUESTS ARE SATISFIED FROM THE PROPER
```

```

/** CLUSTER, YOU SHOULD HAVE MANAGEMENT CLASSES FOR RNN(CLO), NRN(CL1)
/** AND NNR(CL2) SO THAT ONLY THE TARGET CLUSTER WILL HAVE A COPY OF
/** THE BVIR VOLUME. YOU CAN'T CONTROL WHERE THE INITIAL SCRATCH
/** MOUNT IS SATISFIED, BUT YOU CAN CONTROL WHICH TVC THE VOLUME WILL
/** BE IN WHEN THE SUBSEQUENT SPECIFIC MOUNT IS DONE. THE SPECIFIC
/** MOUNT COLLECTS THE BVIR INFORMATION, NOT THE INITIAL SCRATCH MOUNT.
/**
/** IF YOU HAVE MULTIPLE TS7740 GRIDS, YOU MUST HAVE A SEPARATE
/** STORAGE GROUP DEFINED FOR EACH GRID IN ORDER TO ALLOCATE ON THE
/** DESIRED GRID. USE AN ACS ROUTINE TO SELECT THE TARGET GRID.
/**
/** IF YOU ARE RUNNING JES3, YOU MUST RUN STEP3 AS A SEPARATE JOB
/** IN ORDER TO FORCE THE DISMOUNT OF THE TAPE IN STEP2. BVIR DATA
/** WILL ONLY BE WRITTEN TO A TAPE AFTER THE INITIAL DISMOUNT AND
/** RE-MOUNT FOR READ.
/**
/** THIS JOB ISSUES THE BULK VOLUME INFORMATION REQUEST (BVIR) FOR
/** ALL VIRTUAL VOLUMES BELONGING TO THE VTS ASSOCIATED WITH THE
/** VIRTUAL DRIVE ADDRESS USED. THE BVIR FEATURE MUST BE ACTIVATED
/** ON THE VTS RECEIVING THE REQUEST. THIS IS FOR TS7740 ONLY.
/** IF YOU ARE RUNNING AGAINST A PTP AND GETTING DATA FOR THE PTPSYNC
/** JOB, YOU NEED TO RUN THIS JOB TWICE, ONCE FOR EACH VTS.
/**
/**BVIRMES PROC USERHLQ=USERID,      HI-LEVEL FOR USER DATA FILES
/**      TOOLHLQ=TOOLID,      HLQ FOR LOAD AND CNTL
/**      SITE=SITENAME,      2ND LEVEL QUALIFIER
/**      MC=,      DIRECT TO SPECIFIC VTS OR CLUSTER
/**      VTSID=,      USE CLO, CL1, CL2, ETC TO BE PART OF DSN
/**      UNIT=VTAPE      UNITNAME ON THIS VTS
/**
/**STEP1  EXEC PGM=IEFBRI4
/**DEL1   DD  UNIT=(&UNIT,,DEFER),DISP=(MOD,DELETE),
/**          DSN=&USERHLQ..&SITE..&VTSID..BVIRMES
/**DEL2   DD  UNIT=SYSDA,DISP=(MOD,DELETE),SPACE=(TRK,1),
/**          DSN=&USERHLQ..&SITE..&VTSID..MESFILE
/**
/**STEP2  EXEC PGM=IEBGENER
/**SYSPRINT DD  SYSOUT=*
/**SYSUT1  DD  DISP=SHR,DSN=&TOOLHLQ..IBMTTOOLS.CNTL(BVIRMES)
/**SYSUT2  DD  DSN=&USERHLQ..&SITE..&VTSID..BVIRMES,MGMTCLAS=&MC,
/**          UNIT=&UNIT,LABEL=(,SL),DISP=(NEW,CATLG),
/**          DCB=(RECFM=FB,BLKSIZE=80,LRECL=80,TRTCH=NOCOMP)
/**SYSIN   DD  DUMMY
/**
/**STEP3  EXEC PGM=IEBGENER
/**SYSPRINT DD  SYSOUT=*
/**SYSUT1  DD  DSN=&USERHLQ..&SITE..&VTSID..BVIRMES,DISP=(OLD,DELETE),
/**          DCB=(DSORG=PS,RECFM=U,BLKSIZE=640)
/**SYSUT2  DD  DSN=&USERHLQ..&SITE..&VTSID..MESFILE,
/**          UNIT=SYSDA,SPACE=(640,(500000,200000),RLSE),
/**          DISP=(,CATLG),DCB=(DSORG=PS,RECFM=U,BLKSIZE=640)
/**SYSIN   DD  DUMMY
/**      PEND
/**
/**GETVOLS EXEC BVIRMES,VTSID=CLO,MC=
/**
/**

```

Physical volume status

You may use BVIRPOOL to create an unformatted snapshot of the status of physical volumes. The sample JCL is listed in Example F-8.

Note: If this is a TS7720 Virtualization Engine cluster, the following record is returned:

'NOT SUPPORTED IN A DISK-ONLY TS7700 VIRTUALIZATION ENGINE'

Example F-8 BVIRPOOL Sample JCL

```
//ITS01      JOB  CONSOLE,
//           MSGCLASS=H,MSGLEVEL=(1,1),CLASS=B,
//           TIME=1440,REGION=2M
//*
//* THIS JOB ISSUES THE BULK VOLUME INFORMATION (BVIR) REQUEST FOR
//* MEDIA POOLING STATISTICS FROM THE VE ASSOCIATED WITH THE VIRTUAL
//* DRIVE ADDRESS USED.  THE JOBS DEFAULTS TO DOING A WTO WITH THE
//* REPORT OUTPUT IN ADDITION TO THE POOLRPT.
//* TO INSURE THAT BVIR REQUESTS ARE SATISFIED FROM THE PROPER
//* CLUSTER, YOU SHOULD HAVE MANAGEMENT CLASSES FOR RNN(CLO), NRN(CL1)
//* AND NNR(CL2) SO THAT ONLY THE TARGET CLUSTER WILL HAVE A COPY OF
//* THE BVIR VOLUME.  YOU CAN'T CONTROL WHERE THE INITIAL SCRATCH
//* MOUNT IS SATISFIED, BUT YOU CAN CONTROL WHICH TVC THE VOLUME WILL
//* BE IN WHEN THE SUBSEQUENT SPECIFIC MOUNT IS DONE. THE SPECIFIC
//* MOUNT COLLECTS THE BVIR INFORMATION, NOT THE INITIAL SCRATCH MOUNT.
//*
//* IF YOU HAVE MULTIPLE TS7740 GRIDS, YOU MUST HAVE A SEPARATE
//* STORAGE GROUP DEFINED FOR EACH GRID IN ORDER TO ALLOCATE ON THE
//* DESIRED GRID.  USE AN ACS ROUTINE TO SELECT THE TARGET GRID.
//*
//BVIRPOOL PROC USERHLQ=USERID,      HI-LEVEL FOR USER DATA FILES
//      TOOLHLQ=TOOLID,      HLQ FOR LOAD AND CNTL
//      MC=,      DIRECT TO SPECIFIC CLUSTER IN GRID
//      UNIT=B29M2C36      UNITNAME ON THIS VE
//*
//STEP1      EXEC PGM=IEFBRI4
//DEL1       DD  UNIT=(&UNIT,,DEFER),DISP=(MOD,DELETE),SPACE=(TRK,1),
//              DSN=&USERHLQ..BVIRPOOL.REQUEST
//*
//STEP2      EXEC PGM=IEBGENER
//SYSPRINT   DD  SYSOUT=*
//SYSUT1     DD  DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.CNTL(BVIRPOOL)
//SYSUT2     DD  DSN=&USERHLQ..BVIRPOOL.REQUEST,MGMTCLAS=&MC,
//              UNIT=&UNIT,LABEL=(,SL),DISP=(NEW,CATLG),
//              DCB=(RECFM=F,LRECL=80,BLKSIZE=80,TRTCH=NOCOMP)
//SYSIN      DD  DUMMY
//*
//STEP3      EXEC PGM=BVIRPOOL
//STEPLIB    DD  DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.LOAD
//SYSLIST     DD  SYSOUT=*
//BVIRIN     DD  DSN=&USERHLQ..BVIRPOOL.REQUEST,
//              DCB=(RECFM=U,BLKSIZE=24000),
//              DISP=(OLD,DELETE)
//POOLRPT    DD  SYSOUT=*
//          PEND
//*
//GETPOOL    EXEC BVIRPOOL
//STEP3.SYSCNTL DD  DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.JCL(EXPIRE)
//          DD  DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.JCL(DOWTO)
```

```
//          DD  *
*DLSER= FRSER  TOSER;  CHANGE FROM ONE VALUE TO ANOTHER FOR REPORTS
//*
//
```

Physical Volume Status report

For a formatted report of the physical volume status, you can use BVIRPHY which is listed in Example F-9. The output of BVIRPHY is then processed by BVIRPRPT (see “Physical Volume and Pool Status Report Writer” on page 902) to produce a fully formatted report.

Example F-9 Sample JCL for BVIRPHY

```
//ITS01      JOB  CONSOLE,
//          MSGCLASS=H,MSGLEVEL=(1,1),CLASS=B,
//          TIME=1440,REGION=2M
//*
//*
//*  USED TO GET PHYSICAL BACK-END VOLUME STATUS.  ALL OR INDIVIDUAL.
//*  JOB BVIRPRPT IS USED TO INTERPRET THE RECORD.
//*
//*  TO INSURE THAT BVIR REQUESTS ARE SATISFIED FROM THE PROPER
//*  CLUSTER, YOU SHOULD HAVE MANAGEMENT CLASSES FOR RNN(CLO), NRN(CL1)
//*  AND NNR(CL2) SO THAT ONLY THE TARGET CLUSTER WILL HAVE A COPY OF
//*  THE BVIR VOLUME.  YOU CAN'T CONTROL WHERE THE INITIAL SCRATCH
//*  MOUNT IS SATISFIED, BUT YOU CAN CONTROL WHICH TVC THE VOLUME WILL
//*  BE IN WHEN THE SUBSEQUENT SPECIFIC MOUNT IS DONE. THE SPECIFIC
//*  MOUNT COLLECTS THE BVIR INFORMATION, NOT THE INTIIAL SCRATCH MOUNT.
//*
//*  IF YOU HAVE MULTIPLE TS7740 GRIDS, YOU MUST HAVE A SEPARATE
//*  STORAGE GROUP DEFINED FOR EACH GRID IN ORDER TO ALLOCATE ON THE
//*  DESIRED GRID.  USE AN ACS ROUTINE TO SELECT THE TARGET GRID.
//*
//*  IF YOU ARE RUNNING JES3, YOU MUST RUN STEP3 AS A SEPARATE JOB
//*  IN ORDER TO FORCE THE DISMOUNT OF THE TAPE IN STEP2.  BVIR DATA
//*  WILL ONLY BE WRITTEN TO A TAPE AFTER THE INITIAL DISMOUNT AND
//*  RE-MOUNT FOR READ.
//*
//*  THIS JOB ISSUES THE BULK VOLUME INFORMATION REQUEST (BVIR) FOR
//*  ALL PHYSICAL VOLUMES BELONGING TO THE VTS ASSOCIATED WITH THE
//*  VIRTUAL DRIVE ADDRESS USED.  THE BVIR FEATURE MUST BE ACTIVATED
//*  ON THE VTS RECEIVING THE REQUEST.  THIS IS FOR TS7740 ONLY.
//*
//BVIRPHY PROC USERHLQ=USERID,          HI-LEVEL FOR USER DATA FILES
//          TOOLHLQ=TOOLID,              HLQ FOR LOAD AND CNTL
//          SITE=SITENAME,              2ND LEVEL QUALIFIER
//          MC=,                          DIRECT TO SPECIFIC VTS OR CLUSTER
//          VTSID=,                       USE CLO, CL1, CL2, ETC TO BE PART OF DSN
//          UNIT=VTAPE                    UNITNAME ON THIS VTS
//*
//STEP1      EXEC PGM=IEFBR14
//DEL1       DD  UNIT=(&UNIT,,DEFER),DISP=(MOD,DELETE),
//              DSN=&USERHLQ..&SITE..&VTSID..BVIRPHY
//DEL2       DD  UNIT=SYSDA,DISP=(MOD,DELETE),SPACE=(TRK,1),
//              DSN=&USERHLQ..&SITE..&VTSID..PHYFILE
//*
//STEP2      EXEC PGM=IEBGENER
//SYSPRINT   DD  SYSOUT=*
//SYSUT1     DD  DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.CNTL(BVIRPHY)
```

```
//SYSUT2 DD DSN=&USERHLQ..&SITE..&VTSID..BVIRPHY,MGMTCLAS=&MC,
//          UNIT=&UNIT,LABEL=(,SL),DISP=(NEW,CATLG),
//          DCB=(RECFM=FB,BLKSIZE=80,LRECL=80,TRTCH=NOCOMP)
//SYSIN DD DUMMY
//*
//STEP3 EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSUT1 DD DSN=&USERHLQ..&SITE..&VTSID..BVIRPHY,DISP=(OLD,DELETE),
//          DCB=(DSORG=PS,RECFM=U,BLKSIZE=320)
//SYSUT2 DD DSN=&USERHLQ..&SITE..&VTSID..PHYFILE,
//          UNIT=SYSDA,SPACE=(320,(500000,200000),RLSE),
//          DISP=(,CATLG),DCB=(DSORG=PS,RECFM=U,BLKSIZE=320)
//SYSIN DD DUMMY
// PEND
//*
//GETVOLS EXEC BVIRPHY,VTSID=CLO,MC=
//*
```

Physical volume pool status report

For a formatted report of the physical volume pool status, you can use BVIRPLNN which is listed in Example F-10. The output of BVIRPLNN is then processed by BVIRPRPT (see “Physical Volume and Pool Status Report Writer” on page 902) to produce a fully formatted report.

Example F-10 Sample JCL for BVIRPLNN

```
//ITS01 JOB CONSOLE,
//          MSGCLASS=H,MSGLEVEL=(1,1),CLASS=B,
//          TIME=1440,REGION=2M
//*
//*
//* THIS JOB ISSUES THE BULK VOLUME INFORMATION (BVIR) REQUEST FOR
//* MEDIA POOLING STATISTICS FROM THE VE ASSOCIATED WITH THE VIRTUAL
//* DRIVE ADDRESS USED.
//*
//* TO INSURE THAT BVIR REQUESTS ARE SATISFIED FROM THE PROPER
//* CLUSTER, YOU SHOULD HAVE MANAGEMENT CLASSES FOR RNN(CLO), NRN(CL1)
//* AND NNR(CL2) SO THAT ONLY THE TARGET CLUSTER WILL HAVE A COPY OF
//* THE BVIR VOLUME. YOU CAN'T CONTROL WHERE THE INITIAL SCRATCH
//* MOUNT IS SATISFIED, BUT YOU CAN CONTROL WHICH TVC THE VOLUME WILL
//* BE IN WHEN THE SUBSEQUENT SPECIFIC MOUNT IS DONE. THE SPECIFIC
//* MOUNT COLLECTS THE BVIR INFORMATION, NOT THE INTIIAL SCRATCH MOUNT.
//*
//* IF YOU HAVE MULTIPLE TS7740 GRIDS, YOU MUST HAVE A SEPARATE
//* STORAGE GROUP DEFINED FOR EACH GRID IN ORDER TO ALLOCATE ON THE
//* DESIRED GRID. USE AN ACS ROUTINE TO SELECT THE TARGET GRID.
//*
//* IF YOU ARE RUNNING JES3, YOU MUST RUN STEP3 AS A SEPARATE JOB
//* IN ORDER TO FORCE THE DISMOUNT OF THE TAPE IN STEP2. BVIR DATA
//* WILL ONLY BE WRITTEN TO A TAPE AFTER THE INITIAL DISMOUNT AND
//* RE-MOUNT FOR READ.
//*
//BVIRPOOL PROC USERHLQ=USERID, HI-LEVEL FOR USER DATA FILES
//          TOOLHLQ=TOOLID, HLQ FOR LOAD AND CNTL
//          SITE=SITENAME, 2ND LEVEL QUALIFIER
//          VTSID=CLO, USE CLO, CL1, CL2, ETC TO BE PART OF DSN
//          POOL=, TWO DIGIT POOL NUMBER REQUESTED
//          MC=, DIRECT TO SPECIFIC CLUSTER IN GRID
```

```

//          UNIT=VTAPE                UNITNAME ON THIS VE
//*
//STEP1    EXEC PGM=IEFBRI4
//DEL1     DD  UNIT=(&UNIT,,DEFER),DISP=(MOD,DELETE),SPACE=(TRK,1)
//          DSN=&USERHLQ..&SITE..&VTSID..BVIRPOOL.NUM&POOL
//DEL2     DD  UNIT=(&UNIT,,DEFER),DISP=(MOD,DELETE),SPACE=(TRK,1)
//          DSN=&USERHLQ..&SITE..&VTSID..POOL&POOL
//*
//STEP2    EXEC PGM=IEBGENER
//SYSPRINT DD  SYSOUT=*
//SYSUT1   DD  DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.CNTL(BVIRPL&POOL)
//SYSUT2   DD  DSN=&USERHLQ..&SITE..&VTSID..BVIRPOOL.NUM&POOL,
//          MGMTCLAS=&MC,
//          UNIT=&UNIT,LABEL=(,SL),DISP=(NEW,CATLG),
//          DCB=(RECFM=F,LRECL=80,BLKSIZE=80,TRTCH=NOCOMP)
//SYSIN    DD  DUMMY
//*
//STEP3    EXEC PGM=IEBGENER
//SYSPRINT DD  SYSOUT=*
//SYSUT1   DD  DSN=&USERHLQ..&SITE..&VTSID..BVIRPOOL.NUM&POOL,
//          DCB=(RECFM=U,BLKSIZE=320),
//          DISP=(OLD,DELETE)
//SYSUT2   DD  DSN=&USERHLQ..&SITE..&VTSID..POOL&POOL,
//          DCB=(RECFM=U,BLKSIZE=320),DISP=(NEW,CATLG),
//          SPACE=(TRK,(10,10),RLSE)
//SYSIN    DD  DUMMY
//  PEND
//*  REQUEST AS MANY AS YOU CURRENTLY USE
//GETPOL00 EXEC BVIRPOOL,POOL=00
//GETPOL01 EXEC BVIRPOOL,POOL=01
//GETPOL02 EXEC BVIRPOOL,POOL=02
//GETPOL03 EXEC BVIRPOOL,POOL=03
.
. Same for Pools 4 - 29
.
/*GETPOL30 EXEC BVIRPOOL,POOL=30
/*GETPOL31 EXEC BVIRPOOL,POOL=31
/*GETPOL32 EXEC BVIRPOOL,POOL=32
/*
//

```

Physical Volume and Pool Status Report Writer

BVIRPRPT uses the output of BVIRPLNN or BVIRPHY to produce formatted reports of physical volume and pool status. Example F-11 shows the sample JCL.

Example F-11 Sample JCL for BVIRPRPT

```

//ITS01     JOB  CONSOLE,
//          MSGCLASS=H,MSGLEVEL=(1,1),CLASS=B,
//          TIME=1440,REGION=2M
//*
/*  PROCESSES THE OUTPUT FILES FROM BVIRPHY AND BVIRPLNN JOBS.
/*
/*  IF USING RECLAIMGB FOR COPY-EXPORT VOLUMES, MODULE BVIRPRPT MUST
/*  BE IN AN APF LIBRARY.
/*
//BVIRPRPT  PROC TOOLHLQ=TOOLID,  HLQ FOR LOAD AND CNTL
//          USERHLQ=USERID,      HLQ FOR USER DATA

```



```

//      SITE=SITENAME,          2ND LEVEL QUALIFIER
//      VTSID=CLO      USE CLO, CL1, CL2, ETC TO BE PART OF DSN
// *
//TOOLSTEP EXEC PGM=BVIRPRPT
//STEPLIB DD DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.LOAD
//SYSUDUMP DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//SYSLIST DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//SORTWK01 DD UNIT=SYSDA,SPACE=(CYL,(25,25))
//SORTWK02 DD UNIT=SYSDA,SPACE=(CYL,(25,25))
//SORTWK03 DD UNIT=SYSDA,SPACE=(CYL,(25,25))
//SORTWK04 DD UNIT=SYSDA,SPACE=(CYL,(25,25))
//SORTWK05 DD UNIT=SYSDA,SPACE=(CYL,(25,25))
//SORTIN DD UNIT=SYSDA,SPACE=(CYL,(20,50),RLSE),
//      DCB=(RECFM=FB,BLKSIZE=0)
//SORTOUT DD UNIT=SYSDA,SPACE=(CYL,(20,50),RLSE),
//      DCB=(RECFM=FB,BLKSIZE=0)
//POOLRPT DD SYSOUT=*          LRECL=170
// PEND
//RUNJOB EXEC BVIRPRPT
// *
//TOOLSTEP.SYSCNTL DD DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.JCL(EXPIRE)
//      DD *
CUSTOMER= FOR TITLE LINE BVIRPRPT; <= 1-50 CHAR
*DETAIL= N;      N MEANS DO NOT SHOW SECOND DETAIL LINE WITH TIMES
*LINES= 65535;
*
*   IF RECLAIMGB IS USED FOR COPY-EXPORT VOLUMES,
*   MODULE BVIRPRPT MUST BE IN APF LIBRARY
*   LIBNAME IS THE DISTRIBUTED LIBRARY NAME THAT THE CUSTOMER ASSIGNED
*   THROUGH THE DFSMS LIBRARY DEFINE PANEL.
*LIBNAME= ABCDE;
*LRDELAY= 7;      SECONDS DELAY BETWEEN LIBRARY REQUEST COMMANDS
*RECLAIMGB= 140;  RECLAIM IF LESS ACTIVE GIGABYTES THAN NNN VALUE
* ABOVE COMMENTED MEANS 0 GB SO NOTHING GETS RECLAIMED, JUST REPORTED
* IF USED A LIBRARY REQUEST,LIBNM,COPYEXP,RECLAIM  COMMAND IS ISSUED
MAXGB= 200; IF RECLAIMING, LIMIT THE AMOUNT BROUGHT BACK TO CACHE
DAYSAGO = 3; PVOL MUST HAVE BEEN WRITTEN >N DAYS AGO FOR RECLAIM
*CONSOLENAME= XXXXXXXX; OBTAINED FROM D C OPERATOR COMMAND
*      USE 00000000; WHEN NO CONSOLE DEFINED FOR LPAR
*INCVOL= LO*;      INCLUDE ONLY THESE VOLERS
*INCVOL= 010000 010999; INCLUDE THIS RANGE
*EXCVOL= 475000 475999; EXCLUDE THIS RANGE
// *
// * PICK YOUR BVIRIN FILE
// *BVIRIN DD DISP=SHR,DSN=&USERHLQ..&SITE..&VTSID..PHYFILE
// * YOU CAN CONCATINATE MULTIPLE POOL FILES FROM BVIRPLNN JOB
//BVIRIN DD DISP=SHR,DSN=&USERHLQ..&SITE..&VTSID..POOL00
//      DD DISP=SHR,DSN=&USERHLQ..&SITE..&VTSID..POOL01
//      DD DISP=SHR,DSN=&USERHLQ..&SITE..&VTSID..POOL02
// *      DD DISP=SHR,DSN=&USERHLQ..&SITE..&VTSID..POOL03
.
. Same for Pools 4 - 29
.
// *      DD DISP=SHR,DSN=&USERHLQ..&SITE..&VTSID..POOL30
// *      DD DISP=SHR,DSN=&USERHLQ..&SITE..&VTSID..POOL31
// *      DD DISP=SHR,DSN=&USERHLQ..&SITE..&VTSID..POOL32

```

VEHSTATS Reports

VEHSTATS jobs process the history files that are created by BVIRHSTS, BVIRHSTU, BVIRHSTV. The JCL, which had been provided in member VEHSTATS, has been split into three jobs, depending on how you want to view or save your reports:

VEHSTSO	Writes reports directly to SYSOUT (this is the old VEHSTATS).
VEHSTPS	Writes final reports to a single physical sequential file where the reports are written with DISP=MOD.
VEHSTPO	Writes final reports to a PDSE where each report is a separate member.

These three VEHSTATS jobs can also be found in *userid*.IBMT00LS.JCL. Example F-12 lists the sample JCL for VEHSTPO.

Example F-12 VEHSTPO Sample JCL

```
//ITS01      JOB  CONSOLE,
//           MSGCLASS=H,MSGLEVEL=(1,1),CLASS=B,
//           TIME=1440,REGION=2M
//*
//*  VEHSTATS WILL ONLY RUN IN Z/OS.
//*
//*  VIRTUALIZATION ENGINE HISTORICAL STATISTICS REPORTING
//*  MODIFIED VERSION OF VEHSTATS TO WRITE REPORTS OF VARYING LENGTHS
//*  TO A SINGLE OUTPUT REPORT PDSE TO ACCOMMODATE THE WIDEST REPORT.
//*  RUN THE BVIRHST(U/V/S) JOB FIRST TO GET THE STATISTICS FILE(S).
//*  FOR LONGER DESCRIPTION OF FIELD NAMES SEE IBMT00LS.JCL(ORDERV12)
//*
//VEHSTATS PROC TOOLHLQ=TOOLID,      HLQ FOR LIBRARIES
//  USERHLQ=USERID,      FOR THE INPUT BVIR FILE
//  SITE=SITENAME,      2ND LEVEL QUALIFIER
//  ORDER=ORDERV12,      DEFAULT ORDER STATEMENTS FOR GRAPHING PACKAGE
//*  ORDER=ORDERALL,      ALL AVAILABLE ORDER STATEMENTS
//  RECL=260,      260 IS WIDE ENOUGH FOR ALL REPORTS AND 22 CLUSTERS
//*  ON COMPARE REPORT.  ADD 11 FOR EACH CLUSTER ABOVE 22
//  BLK=5200,      EVEN MULTIPLE OF RECL
//  ID=RUN1,      LAST NODE FOR REPORT FILE
//  GRIDID=12345      ID FOR REPORTING SYSTEM
//*
//DELETE EXEC PGM=IEFBR14
//HOURFLAT DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..HOURFLAT.TXT,
//          DCB=(RECFM=FB,BLKSIZE=0)
//DAYHSMRY DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..DAYHSMRY.TXT,
//          DCB=(RECFM=FB,BLKSIZE=0)
//WEKHSMRY DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..&SITE..#&GRIDID..WEKHSMRY.TXT,
//          DCB=(RECFM=FB,BLKSIZE=0)
//OUTRPTS DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DCB=(RECFM=FBA,LRECL=&RECL.,BLKSIZE=&BLK.),
//          DSN=&USERHLQ..&SITE..#&GRIDID..RPTPDS.&ID
//*
//ALLOC EXEC PGM=IEFBR14
//OUTRPTS DD UNIT=SYSDA,DISP=(,CATLG),SPACE=(CYL,(5,5,10)),
//          DCB=(RECFM=FBA,LRECL=&RECL.,BLKSIZE=&BLK.),DSNTYPE=LIBRARY,
//          DSN=&USERHLQ..&SITE..#&GRIDID..RPTPDS.&ID
//*
```

```

//RPTSTEP EXEC PGM=VEHSTATS,REGION=OM,PARM='FILEOUT'
//STEPLIB DD DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.LOAD
//SYSLIST DD SYSOUT=* CONTROL PARAMETERS USED
//RECLIST DD DUMMY,SYSOUT=* DETAIL LIST OF BVIR RECORD TIME STAMPS
//H20VIRT DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H20VIRT
//H21ADP00 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H21ADP00
//H21ADP01 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H21ADP01
//H21ADP02 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H21ADP02
//H21ADP03 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H21ADP03
//H21ADPXX DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H21ADPXX
//H21ADPSU DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H21ADPSU
//H30TVC1 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H30TVC1
//H31IMEX DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H31IMEX
//H32TDU12 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H32TDU12
//H32TDU34 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H32TDU34
//H32CSP DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H32CSP
//H32GUP01 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H32GUP01
//H32GUP03 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H32GUP03
//H32GUP05 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H32GUP05
//H32GUP07 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H32GUP07
//H32GUP09 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H32GUP09
//H32GUP11 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H32GUP11
//H32GUP13 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H32GUP13
//H32GUP15 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H32GUP15
//H32GUP17 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H32GUP17
//H32GUP19 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H32GUP19
//H32GUP21 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H32GUP21
//H32GUP23 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H32GUP23
//H32GUP25 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H32GUP25
//H32GUP27 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H32GUP27
//H32GUP29 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H32GUP29
//H32GUP31 DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H32GUP31

```

```

//H33GRID DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&H33GRID
//HOURXFER DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&HOURXFER
//DAYXFER DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&DAYXFER
//AVGRDST DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&AVGRDST
//DAYSMRY DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&DAYSMRY
//MONSMRY DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&MONSMRY
//COMPARE DD UNIT=SYSDA,SPACE=(CYL,(5,5)),DISP=(,PASS),
// DSN=&&COMPARE
//WEKHSRMRY DD UNIT=SYSDA,SPACE=(TRK,(3,5),RLSE),DISP=(,CATLG),
// DSN=&USERHLQ..&SITE..#&GRIDID..WEKHSRMRY.TXT,
// DCB=(RECFM=FB,BLKSIZE=0)
//DAYHSRMRY DD UNIT=SYSDA,SPACE=(TRK,(3,5),RLSE),DISP=(,CATLG),
// DSN=&USERHLQ..&SITE..#&GRIDID..DAYHSRMRY.TXT,
// DCB=(RECFM=FB,BLKSIZE=0)
//HOURFLAT DD UNIT=SYSDA,SPACE=(CYL,(90,9),RLSE),DISP=(,CATLG),
// DSN=&USERHLQ..&SITE..#&GRIDID..HOURFLAT.TXT,
// DCB=(RECFM=FB,BLKSIZE=0)
//SORTIN DD UNIT=(SYSDA,1),SPACE=(CYL,(300,100)),
// DSN=&&SORTIN,DCB=(RECFM=VB,LRECL=12000,BLKSIZE=0)
//SORTOUT DD UNIT=(SYSDA,1),SPACE=(CYL,(300,100)),
// DSN=&&SORTED,DCB=(RECFM=VB,LRECL=12000,BLKSIZE=0)
//SORTWK01 DD UNIT=SYSDA,SPACE=(CYL,(200,100))
//SORTWK02 DD UNIT=SYSDA,SPACE=(CYL,(200,100))
//SORTWK03 DD UNIT=SYSDA,SPACE=(CYL,(200,100))
//SORTWK04 DD UNIT=SYSDA,SPACE=(CYL,(200,100))
//SORTWK05 DD UNIT=SYSDA,SPACE=(CYL,(200,100))
//SORTWK06 DD UNIT=SYSDA,SPACE=(CYL,(200,100))
//SORTWK07 DD UNIT=SYSDA,SPACE=(CYL,(200,100))
//SYSOUT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
// PEND
//RUNRPTS EXEC VEHSTATS
//*****
//*-DATES ARE ENTERED 'DDMMYYYY', E.G., 15JUN2008 *
//* ALTERNATIVELY, 'TODAY', 'TODAY- NN', AND 'TODAY+ NN' MAY BE GIVEN*
//* E.G. SDATE= TODAY- 7; NN CAN'T EXCEED 360. *
//RPTSTEP.SYSCNTL DD DISP=SHR,DSN=&TOOLHLQ..IBMTTOOLS.JCL(EXPIRE)
// DD DISP=SHR,DSN=&USERHLQ..&SITE..IBMTTOOLS.JCL(&ORDER)
// DD *
*****
* FILL IN THE FOLLOWING RECORDS AS APPROPRIATE: *
*****
CUSTOMER= TITLENAME VEHSTATS; 1-50 CHAR
* IF THE CLUSTERS IN YOUR GRID ARE NOT SEQUENTIAL, USE THE DEFDL
* PARAMETER TO DEFINE WHICH ONES ARE ACTUALLY PRESENT.
*DEFDL= H2909 0; CLUSTER SERIAL AND CLUSTER NUMBER
*DEFDL= H2918 2; CLUSTER SERIAL AND CLUSTER NUMBER
*DEFDL= H2906 3; CLUSTER SERIAL AND CLUSTER NUMBER
*EUROFORMAT; USE COMMA INSTEAD OF PERIOD FOR FRACTIONAL NUMBERS
*GERMANMONTH; USE GERMAN MONTHS, MAI, OKT, DEZ FOR GERMAN EXCEL 2003
*SINGLESPACE; USE SINGLE SPACE BETWEEN FIELDS IN FLAT FILES
*ONEHEADING; ONLY ONE HEADING ON FLAT FILES, NOT BETWEEN CLUSTERS
*NOFILLER; DO NOT WRITE FILLR LINES TO DAYHSRMRY
*IGNOREHEADER; DO NO WRITE ID HEADER TO HOURFLAT FILE

```

```

QUEAGEMINUTES; REPORT DEF & RUN QUEUE AGE AS MINUTES, NOT SECONDS
REPORT= HRS HDSUM COM; HOURLY ROLL-UP, COMPARE, AND FLAT FILE SUMMARY
*   = QTR      REQUEST 15 MINUTE REPORTING AS GENERATED BY TS7740
*   = HRS      REQUEST HOURLY ROLL-UP REPORTING
*   = GRID     SUMMARIZES ALL CLUSTERS WITHIN GRID
*   = COMPARE  REQUEST SIDE BY SIDE CLUSTER COMPARISON
*   = HDSUM    DAILY SUMMARY FLAT FILE - HORIZONTAL 1 DAY/LINE
*   = DXFR     FOR DAILY ON DEMAND TRANSFER REPORTING
*UTCMINUS= 07;  ADJUST UTC TO LOCAL TIME WEST OF GREENWICH
*UTCPLUS= 02;   ADJUST UTC TO LOCAL TIME EAST OF GREENWICH
*
* SEE MEMBER, VEHDATES, FOR MORE DETAIL ON DATES
*
*SDATE= 14JAN2009;  START DATE FOR OUTPUT REPORTING
*SDATE= TODAY- 1;   REPORT JUST YESTERDAY'S DATA
*SDATE= LASTWEEK;   REPORT JUST LAST WEEK'S ACTIVITY
*STIME= 00:05;      START TIME FOR OUTPUT REPORTING
*EDATE= 15JAN2009;  END DATE FOR OUTPUT REPORTING
*EDATE= TODAY- 1;   REPORT JUST YESTERDAY'S DATA
*EDATE= LASTWEEK;   REPORT JUST LAST WEEK'S ACTIVITY
*ETIME= 00:05;      END TIME FOR OUTPUT REPORTING
*SELECTDOW= FRI;    LIMITS HOURFLAT TO JUST THIS DOW
*
* SEE MEMBER, VEHDATES, FOR MORE DETAIL ON DATES
*
LINES= 58;  LINES= 999 TO PUT DAYSMRY & MONSMRY ON SINGLE PAGE BREAK
*
*   A MICRO CODE UPGRADE CHANGED THE SERIAL NUMBER BEING REPORTED.
*   YOU CAN EITHER CHANGE THE OLD TO MATCH THE NEW OR THE NEW TO
*   MATCH THE OLD VALUE.
*DLSER= FRSER TOSER;  CHANGE FROM ONE VALUE TO ANOTHER FOR REPORTS
*
*   THE INITIAL GRID SERIAL WAS BINARY 0, BUT APPEARED ON THE
*   REPORTS AS A VALUE OF ?????. YOU CAN CHANGE THE ????? TO THE
*   NEW VALUE SO OLD AND NEW DATA WILL APPEAR AS THE SAME GRID.
*GRIDSER= ????? TOSER;  CHANGE BINARY 0 TO NEW GRID SERIAL NUMBER
SMFNUM = 194;  USER SELECTABLE SMF # FOR STATSMF DATA
*VTSNUM = SERNO;  SELECT JUST THIS CLUSTER TO MAKE IT EASIER TO WORK
*                  WITH FLAT FILES AND GRAPHING PACKAGE
*
* THE ORDER STATEMENTS DETERMINE WHICH FIELDS WILL BE REPORTED IN THE
* DAYSMRY, MONSMRY, HOURFLAT, DAYHSMRY, AND WEKHSRMRY REPORTS AND WHAT
* ORDER THEY WILL APPEAR IN.
* PICK AND CHOOSE FROM THIS LIST AND RE-ARRANGE TO FIT YOUR NEEDS.
*
* IBMTOOLS.JCL(ORDERV12) IS THE DEFAULT MEMBER OR YOU CAN CREATE YOUR
* OWN MEMBER WITH YOUR FIELDS AND SEQUENCE.
*
/**
/** ACTIVATE ONE OR MORE OF THE FOLLOWING DD STATEMENTS FOR YOUR
/** DATA DEPENDING ON WHICH BVIRHST(U/V/S) JOB WAS USED TO COLLECT
/** THE STATISTICS
/** ACTIVATE THE FOLLOWING YOU USED BVIRHSTU
/**STATSU DD DISP=SHR,
/**      DSN=&USERHLQ..&SITE..&GRIDID..HSTU.D090205.D090205
/** ACTIVATE THE FOLLOWING YOU USED BVIRHSTV
/**STATSVB DD DISP=SHR,
/**      DSN=&USERHLQ..&SITE..&GRIDID..HSTV.D090728.D090730
/** ACTIVATE THE FOLLOWING YOU USED BVIRHSTS
/**STATSMF DD DISP=SHR, RECORDS WILL BE SELECTED BASED ON SMFNUM

```

```

/*      DSN=&USERHLQ..&SITE..#&GRIDID..SMF194
/*
//COPYRPTS  PROC RPT=          WHICH REPORT TO COPY
//COPYRPT  EXEC  PGM=COPY2PDS,PARM='&RPT.'
//STEPLIB  DD   DISP=SHR,DSN=*.RUNRPTS.RPTSTEP.STEPLIB
//SYSUDUMP  DD   SYSOUT=*
//INRECS   DD   DISP=(OLD,DELETE),DSN=&&&RPT.
//OUTRECS  DD   DISP=SHR,DSN=*.RUNRPTS.ALLOC.OUTRPTS
//  PEND
/*
/*      COMMENT LINES BELOW IF YOU DON'T WANT THOSE REPORTS KEPT
//H2OVIRT  EXEC  COPYRPTS,RPT=H2OVIRT
//H21ADP00 EXEC  COPYRPTS,RPT=H21ADP00
//H21ADP01 EXEC  COPYRPTS,RPT=H21ADP01
//H21ADP02 EXEC  COPYRPTS,RPT=H21ADP02
//H21ADP03 EXEC  COPYRPTS,RPT=H21ADP03
//H21ADPXX EXEC  COPYRPTS,RPT=H21ADPXX
//H21ADPSU EXEC  COPYRPTS,RPT=H21ADPSU
//H30TVC1  EXEC  COPYRPTS,RPT=H30TVC1
//H31IMEX  EXEC  COPYRPTS,RPT=H31IMEX
//H32TDU12 EXEC  COPYRPTS,RPT=H32TDU12
//H32TDU34 EXEC  COPYRPTS,RPT=H32TDU34
//H32CSP   EXEC  COPYRPTS,RPT=H32CSP
//H32GUP01 EXEC  COPYRPTS,RPT=H32GUP01
//H32GUP03 EXEC  COPYRPTS,RPT=H32GUP03
//H32GUP05 EXEC  COPYRPTS,RPT=H32GUP05
//H32GUP07 EXEC  COPYRPTS,RPT=H32GUP07
//H32GUP09 EXEC  COPYRPTS,RPT=H32GUP09
//H32GUP11 EXEC  COPYRPTS,RPT=H32GUP11
//H32GUP13 EXEC  COPYRPTS,RPT=H32GUP13
//H32GUP15 EXEC  COPYRPTS,RPT=H32GUP15
//H32GUP17 EXEC  COPYRPTS,RPT=H32GUP17
//H32GUP19 EXEC  COPYRPTS,RPT=H32GUP19
//H32GUP21 EXEC  COPYRPTS,RPT=H32GUP21
//H32GUP23 EXEC  COPYRPTS,RPT=H32GUP23
//H32GUP25 EXEC  COPYRPTS,RPT=H32GUP25
//H32GUP27 EXEC  COPYRPTS,RPT=H32GUP27
//H32GUP29 EXEC  COPYRPTS,RPT=H32GUP29
//H32GUP31 EXEC  COPYRPTS,RPT=H32GUP31
//H33GRID  EXEC  COPYRPTS,RPT=H33GRID
//HOURXFER EXEC  COPYRPTS,RPT=HOURXFER
//DAYXFER  EXEC  COPYRPTS,RPT=DAYXFER
//AVGRDST  EXEC  COPYRPTS,RPT=AVGRDST
//DAYSMRY  EXEC  COPYRPTS,RPT=DAYSMRY
//MONSMRY  EXEC  COPYRPTS,RPT=MONSMRY
//COMPARE  EXEC  COPYRPTS,RPT=COMPARE
//

```

Export List Volume sample JCL

In this section, we provide sample JCLs to create an Export List Volume. Example F-13 shows how to create the three necessary files. An example of how to use an Export List Volume as part of the Export Copy function is in 5.3.2, “Implementing Copy Export” on page 288.

Example F-13 serves as input for an Copy Export function, which enables you to do off-site vaulting of physical volumes from a TS7700 Virtualization Engine.

Example F-13 Creation of Export List Volume

```
//*****  
  
/* FILE 1: EXPORT LIST  
//*****  
//STEP1 EXEC PGM=IEBGENER  
//SYSPRINT DD SYSOUT=*  
//SYSIN DD DUMMY  
//SYSUT2 DD DSN=HILEVELQ.EXPLIST,MGMTCLAS=MCNOCOPY,  
// UNIT=VTS1,DISP=(NEW,KEEP),LABEL=(1,SL),  
// VOL=(,RETAIN),  
// DCB=(RECFM=FB,BLKSIZE=80,LRECL=80,TRTCH=NOCOMP)  
//SYSUT1 DD *  
EXPORT LIST 03  
EXPORT PARAMETERS PHYSICAL POOL TO EXPORT:09  
OPTIONS1,COPY,EJECT  
/*  
//*****  
/* FILE 2: RESERVED FILE  
//*****  
//STEP2 EXEC PGM=IEBGENER,COND=(4,LT)  
//SYSPRINT DD SYSOUT=*  
//SYSIN DD DUMMY  
//SYSUT2 DD DSN=HILEVELQ.RESERVED,MGMTCLAS=MCNOCOPY,  
// UNIT=VTS1,DISP=(NEW,KEEP),LABEL=(2,SL),  
// VOL=(,RETAIN,REF=*.STEP1.SYSUT2),  
// DCB=*.STEP1.SYSUT2  
//SYSUT1 DD *  
RESERVED FILE  
/*  
//*****  
/* FILE 3: EXPORT STATUS FILE  
//*****  
//STEP3 EXEC PGM=IEBGENER,COND=(4,LT)  
//SYSPRINT DD SYSOUT=*  
//SYSIN DD DUMMY  
//SYSUT2 DD DSN=HILEVELQ.EXPSTATS,MGMTCLAS=MCNOCOPY,  
// UNIT=VTS1,DISP=(NEW,CATLG),LABEL=(3,SL),  
// VOL=(,,REF=*.STEP1.SYSUT2),  
// DCB=*.STEP1.SYSUT2  
//SYSUT1 DD *  
EXPORT STATUS 01  
/*
```

JCL for TS7700 Virtualization Engine migration scenarios

In this section, we include several JCL and REXX examples that can help you with the TS7700 Virtualization Engine migration scenarios, described in the previous sections.

Using EDGUTIL to validate TCDB inconsistencies

The JCL in Example F-14 can help you identify inconsistencies in the TCDB.

Example F-14 Verify information in RMM CDS, Library Manager database, and TCDB

```
//EDGUTIL EXEC PGM=EDGUTIL,PARM='VERIFY(ALL,VOLCAT)'  
//SYSPRINT DD SYSOUT=*  
//MASTER DD DSN=your.rmm.database.name,DISP=SHR  
//VCINOUT DD UNIT=3390,SPACE=(CYL,(900,500))
```

After running EDGUTIL, you receive information about all volumes with conflicting information. Resolve discrepancies before the migration. For more information about this utility, see *z/OS V1 DFSMSrmm Implementation and Customization Guide*, SC26-7405. The job must be run before the migration starts.

IDCAMS example to delete a library definition in TCDB

The JCL in Example F-15 can help you delete library information in the TCDB.

Example F-15 Delete a library in the TCDB

```
//STEP1 EXEC PGM=IDCAMS  
//SYSPRINT DD SYSOUT=*  
//SYSIN DD *  
DELETE (vtsname) -  
LIBRARYENTRY
```

IDCAMS example to list all entries in TCDB

The JCL in Example F-16 can help you list all entries in TCDB.

Example F-16 JCL to list all entries in the TCDB

```
//STEP1 EXEC PGM=IDCAMS  
//SYSPRINT DD SYSOUT=*  
//SYSIN DD *  
LISTC VOLUMEENTRIES(V*) LIBRARY(vtsname)
```

IDCAMS example for change the TCDB

The JCL in Example F-17 can help you change the TCDB for scratch or private volumes in the TS7700 Virtualization Engine.

Example F-17 JCL for changing the TCDB to a new TS7700 Virtualization Engine

```
//*****  
//**** Change TCDB for a scratch volume to a new TS7700 ****  
//*****  
//TCDBSCR EXEC PGM=IDCAMS  
//SYSPRINT DD SYSOUT=*  
//SYSIN DD *  
  ALTER Vvolser VOLENTY LIBRARYNAME(TSname) USEATTRIBUTE(SCRATCH)  
//*****  
//**** Change TCDB entry for a private volume to a new TS7700 ****  
//**** Also change sname to the one used (same as on the VTS)****  
//*****  
//TCDBPRIV EXEC PGM=IDCAMS  
//SYSPRINT DD SYSOUT=*  
//SYSIN DD *  
  ALTER Vvolser VOLUMEENTRY LIBRARYNAME(TSname) -  
    USEATTRIBUTE(PRIVATE) STORAGEGROUP(sname)
```

JCL to change volumes in RMM

The JCL in Example F-18 can help you change volumes in RMM for the TS7700 Virtualization Engine.

Example F-18 JCL for changing volumes in DFSMS/RMM to a new TS7700 Virtualization Engine

```
//PROCESS EXEC PGM=IKJEFT01,DYNAMNBR=25,  
//          TIME=100  
//ISPLOG DD DUMMY  
//SYSPRINT DD SYSOUT=*  
//SYSTSPRT DD SYSOUT=*  
//SYSTSIN DD *  
  RMM CV volser LOCATION(TSname) CMOVE FORCE
```

Even if you specify the FORCE parameter, it takes effect only when necessary. This parameter requires you to be authorized to use a specific RACF Facility class named STGADMIN.EDG.FORCE. Verify that you have the required authorization.

REXX EXEC to update the library name

Example F-19 provides a sample REXX EXEC program for updating the library name for volumes in the TCDB.

Example F-19 REXX EXEC for updating the library name in the TCDB

```
/* REXX */  
/*****/  
/* ALTERVOL */  
/* */  
/* Usage: ALTERVOL DSN(volserlist) LIB(libname) */
```

```

/*                                                    */
/*      Before this EXEC is run, you must create the  */
/*      input data set "volserlist". The LISTCAT command */
/*      can be used to generate the list of volumes   */
/*      to be altered to an output data set.         */
/*                                                    */
/*      LISTCAT VOLUMEENTRIES(V*)                    */
/*              LIBRARY(sourcelib)                   */
/*              OUTFILE(ddname)                      */
/*                                                    */
/*      The list generated has the following format:  */
/*      VOLUME-ENTRY----Vvolser                      */
/*      For command specifics, refer to "Access Method */
/*      Services for the Integrated Catalog Facility". */
/*                                                    */
/*      For each volume in the "volserlist" specified, */
/*      the library name in the volume record is updated */
/*      to the library name specified on the invocation. */
/*                                                    */
/*      ALTER Vvolser VOLUMEENTRY LIBRARYNAME(libname) */
/*      *****/
Arg parms
Dsn=''; Lib=''
If pos('DSN(',parms)>0 then do
  parse var parms front 'DSN(' dsn ')' back
  parms = front || back
end
If pos('LIB(',parms)>0 then do
  parse var parms front 'LIB(' lib ')' back
  parms = front || back
end
If dsn='' | lib='' then do
  'Usage:  ALTERVOL DSN(volserlist) LIB(libname) '
  exit 4
end
/* *****/
/* Get volume serials from source input dsn          */
/* *****/
Address TSO "FREE FI(INDD)"
Address TSO "ALLOCATE FI(INDD) DA("dsn") SHR"
Address TSO "EXECIO * DISKR INDD (STEM X."
Alter1 = "ALTER '"
Alter2 = "' VOLUMEENTRY LIBRARYNAME("lib")"
Volumes = 0
Do N=1 to X.0
  If Pos("VOLUME-ENTRY----",x.n)>0 then do
    Volumes = Volumes + 1
    Parse var x.n "VOLUME-ENTRY----" volser .
    Address TSO Alter1||volser||Alter2
  end
end
End
Say "Lines Read:      " format(x.0,9)
Say "Volumes Altered: " format(Volumes,9)
Address TSO "EXECIO * DISKR INDD (FINIS"
Address TSO "FREE FI(INDD)"
Exit 0

```



Library Manager volume categories

Table G-1 lists all the default Library Manager volume categories, the platforms on which they are used, and their definitions. As defaults, these categories can be modified to address the operational needs of the system environment.

Note: z/OS users may define any category from 0x0001 to 0xFEFF (0x0000 and 0xFFxx cannot be used) with the DEVSUPxx member SYS1.PARMLIB. The appropriate member must be pointed to by IEASYSxx. For z/OS environments, use categories 0x1000 to 0xF000 to avoid potential conflicts with requirements of other operating systems.

Table G-1 Library Manager volume categories

Category (in hex)	Used by	Definition
0000	Null Category	This pseudo-category is used in certain library commands to specify that the category that is already associated with the volume is to be used by default, or that no category is specified. Use of the null category does not affect the volume's order within the category to which it is assigned. No volumes are associated with this category.
0001	DFSMS/MVS	Indicates scratch MEDIA1. MEDIA1 is a standard-capacity cartridge system tape.
0002	DFSMS/MVS	Indicates scratch MEDIA2. MEDIA2 is an enhanced-capacity cartridge system tape.
0003	DFSMS/MVS	Indicates scratch MEDIA3. MEDIA3 is the IBM TotalStorage Enterprise 3590 High Performance Tape Cartridge.
0004	DFSMS/MVS	Indicates scratch MEDIA4. MEDIA4 is the IBM TotalStorage Enterprise 3590 Extended High Performance Tape Cartridge.
0005	DFSMS/MVS	Indicates scratch MEDIA5. MEDIA5 is the IBM TotalStorage Enterprise Tape Cartridge 3592 DATA.

Category (in hex)	Used by	Definition
0006	DFSMS/MVS	Indicates scratch MEDIA6. MEDIA6 is the IBM TotalStorage Enterprise Tape Cartridge 3592 WORM.
0007	DFSMS/MVS	Indicates scratch MEDIA7. MEDIA7 is the IBM TotalStorage Enterprise Tape Cartridge 3592 ECONOMY.
0008	DFSMS/MVS	Indicates scratch MEDIA8. MEDIA8 is the IBM TotalStorage Enterprise Tape Cartridge 3592 ECONOMY WORM.
0009	DFSMS/MVS	Indicates scratch MEDIA9. MEDIA8 is the IBM System Storage Extended Enterprise Tape Cartridge 3592.
000A	DFSMS/MVS	Indicates scratch MEDIA10. MEDIA 10 is the IBM System Storage Extended Enterprise Tape Cartridge 3592 Economy.
000B to 000D	DFSMS/MVS	Reserved.
000E	DFSMS/MVS	Indicates an error volume. Volumes in this category are scratch volumes for which the software detected an error during processing.
000F	DFSMS/MVS	Indicates a private volume. Volumes in this category contain user data or are assigned to a user.
0010 to 007F	DFSMS/MVS	Reserved. These volume categories can be used for library partitioning.
0080	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCH0.
0081	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCH1.
0082	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCH2.
0083	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCH3.
0084	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCH4.
0085	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCH5.
0086	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCH6.
0087	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCH7.

Category (in hex)	Used by	Definition
0088	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCH8.
0089	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCH9.
008A	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCHA.
008B	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCHB.
008C	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCHC.
008D	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCHD.
008E	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCHE.
008F	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCHF.
0090 to 009F	-	Currently not assigned.
00A0	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH00.
00A1	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH01.
00A2	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH02.
00A3	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH03.
00A4	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH04.
00A5	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH05.
00A6	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH06.
00A7	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH07.
00A8	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH08.

Category (in hex)	Used by	Definition
00A9	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH09.
00AA	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH10.
00AB	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH11.
00AC	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH12.
00AD	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH13.
00AE	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH14.
00AF	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH15.
00B0	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH16.
00B1	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH17.
00B2	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH18.
00B3	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH19.
00B4	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH20.
00B5	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH21.
00B6	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH22.
00B7	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH23.
00B8	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH24.
00B9	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH25.
00BA	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH26.
00BB	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH27.
00BC	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH28.
00BD	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH29.

Category (in hex)	Used by	Definition
00BE	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH30.
00BF	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH31.
00C0 to 00FF	-	Currently not used.
0100	OS/400® (MLDD)	Indicates that the volume has been assigned to category *SHARE400. Volumes in this category can be shared with all attached IBM System i and AS/400® systems.
0101	OS/400 (MLDD)	Indicates that the volume has been assigned to category *NOSHARE. Volumes in this category can be accessed only by the OS/400 system that assigned it to the category.
0102 to 012B	-	No assignment to a specific host system. These categories can be dynamically assigned by the Library Manager at the request of a host.
012C	TSM for AIX	Indicates a private volume. Volumes in this category are managed by Tivoli Storage Manager (TSM).
012D	TSM for AIX	Indicates an IBM 3490 scratch volume. Volumes in this category are managed by TSM.
012E	TSM for AIX	Indicates an IBM 3590 scratch volume. Volumes in this category are managed by TSM.
012F to 0FF1	-	No assignment to a specific host system. These categories can be dynamically assigned by the Library Manager at the request of a host.
0FF2	BTLS	Indicates a scratch volume. Volumes in this category belong to the optional scratch pool SCRTCH2.
0FF3	BTLS	Indicates a scratch volume. Volumes in this category belong to the optional scratch pool SCRTCH3.
0FF4	BTLS	Indicates a scratch volume. Volumes in this category belong to the optional scratch pool SCRTCH4.
0FF5	BTLS	Indicates a scratch volume. Volumes in this category belong to the optional scratch pool SCRTCH5.
0FF6	BTLS	Indicates a scratch volume. Volumes in this category belong to the optional scratch pool SCRTCH6.
0FF7	BTLS	Indicates a scratch volume. Volumes in this category belong to the optional scratch pool SCRTCH7.
0FF8	BTLS	Indicates a scratch volume. Volumes in this category belong to the optional scratch pool SCRTCH8.
0FF9 to 0FFE	-	No assignment to a specific host system. These categories can be dynamically assigned by the Library Manager at the request of a host.
0FFF	BTLS	Indicates a scratch volume. Volumes in this category belong to the default scratch pool used by BTLS. Note: If you are planning to migrate to DFSMS/MVS, you should use this default scratch category only.

Category (in hex)	Used by	Definition
1000 to F00D	-	No assignment to a specific host system. These categories can be dynamically assigned by the Library Manager at the request of a host.
F00E	BTLS	Indicates a volume in error. Volumes are assigned to the error category during demount if the volume serial specified for demount does not match the external label of the volume being demounted.
F00F to FEFF	-	No assignment to a specific host system. These categories can be dynamically assigned by the Library Manager at the request of a host.
FF00	All	Insert category. When a tape volume is added to an automated tape library, the library reads the external label on the volume, creates an inventory entry for the volume, and assigns the volume to the insert category. This category may be updated by operator interaction through Librarian Workstation Support.
FF01	Virtual Tape Server and IBM Virtualization Engine TS7700	Stacked Volume Insert category for a Virtual Tape Server & TS7700 Virtualization Engine. A volume is set to this category when its volume serial number is in the range specified for stacked volumes for any VTS library partition.
FF02	Virtual Tape Server	Stacked Volume Scratch category 0 for a Virtual Tape Server and TS7700 Virtualization Engine. This category is reserved for future use for scratch stacked volumes.
FF03	Virtual Tape Server	Stacked Volume Scratch category 1 for a Virtual Tape Server and TS7700 Virtualization Engine. This category is used by the VTS for its scratch stacked volumes. This category is not used if LIC is 527 or later.
FF04	Virtual Tape Server and TS7700 Virtualization Engine	Stacked Volume Private category for a Virtual Tape Server and TS7700 Virtualization Engine. This category includes both scratch and private volumes (since VTS LIC level 527).
FF05	Virtual Tape Server and TS7700 Virtualization Engine	Stacked Volume Disaster Recovery category for a Virtual Tape Server and TS7700 Virtualization Engine. A volume is set to this category when its volume serial number is in the range specified for stacked volumes for any VTS library partition <i>and</i> the Library Manager is in disaster recovery mode.
FF06	Virtual Tape Server and TS7700 Virtualization Engine	This category is used by the VTS as a temporary category for disaster recovery. After a stacked volume in category FF05 is processed, it is put into this category. This category is also used by the PFE tool called "movedata" as a temporary category.
FF07	Virtual Tape Server and TS7700 Virtualization Engine	This category is reserved for future hardware functions.

Category (in hex)	Used by	Definition
FF08	Virtual Tape Server	This category is used by the VTS to indicate a Read-Only-Recovery Stacked Volume with active data cannot be recovered.
FF09 to FF0F	-	Reserved for future hardware functions.
FF10	Library Manager	Convenience-Eject category. When a tape volume is assigned to the convenience-eject category, it becomes eject pending and the Library Manager queues the tape volume to be moved to a convenience output station. When the volume is delivered to an output station, it is deleted from the Library Manager's inventory. Note: Logical volumes cannot be ejected from the library. They can be deleted or exported.
FF11	Library Manager	Bulk-Eject category. Set when the Library Manager accepts an eject request. The volume becomes eject pending and is queued to be moved to the high capacity output station. When the cartridge accessor delivers the volume to the output rack, it is deleted from the Library Manager's inventory. Note: Logical volumes cannot be ejected from the library. They can be deleted or exported.
FF12	Virtual Tape Server	Export-Pending category. A logical volume to be exported is assigned to this category at the beginning of a Virtual Tape Server export operation. Logical volumes in this category are considered in use. Any attempt by a host to mount, audit, or change the category of a volume fails.
FF13	Virtual Tape Server	Exported category. Set when the Virtual Tape Server has exported the logical volume. The attached hosts are notified when volumes are assigned to this category. Any attempt by a host to mount, audit, or change the category of a volume fails, except a Library Set Volume Category order assigning the volume to the purge-volume category.
FF14	Virtual Tape Server	Import category. Stacked volumes that contain logical volumes to import into the Virtual Tape Server are assigned to this category by an operator at the Library Manager, after they are entered into the library through the convenience I/O station and placed in the Unassigned category.
FF15	Virtual Tape Server	Import-Pending category. Logical volumes to be imported from a stacked volume are added to the Library Manager inventory and assigned to this category when the Virtual Tape Server starts importing them. At completion, successfully imported volumes are assigned to the insert category (FF00). The attached hosts are then notified of volumes assigned to the insert category. Any host attempt to use a volume assigned to this category will be failed.

Category (in hex)	Used by	Definition
FF16	Virtual Tape Server and TS7700 Virtualization Engine	Unassigned Category. Volumes are assigned to this category by the Library Manager whenever volumes are added to the library through the convenience I/O station and the library contains one or more VTS subsystems that have the Import/Export functions installed and enabled. Manual intervention is required to assign the cartridges to the proper category. For exported stacked volumes, this is the import category (FF14).
FF17	Virtual Tape Server	Export-Hold category. Physical volumes are assigned to this category on completion of processing for an export stacked volume.
FF18 & FF19	-	Reserved for library. These categories are reserved for future hardware functions.
FF20	PTP Virtual Tape Server and TS7700 Virtualization Engine	Corrupted-Token Volume Category. In a Peer to Peer VTS, volumes are assigned to this category when it has been determined that the tokens associated with the volume have been corrupted. This is to prevent the volume from being selected by a category mount request.
FF21 to FFF5	-	Reserved for library. These categories are reserved for future hardware functions.
FFF4	Library Manager	3592 Cleaner Volume. Cleaner volumes for 3592 type devices in the library are assigned to this category automatically.
FFF5	Library Manager	3592 Service Volume. Volumes are assigned to this category by the Library Manager when it detects that a volume has a unique service cartridge VOLSER and a media type compatible with a 3592 device.
FFF6	Library Manager	3590-Service-Volume Category. Volumes are assigned to this category by the Library Manager when it detects that a volume has a unique service cartridge VOLSER and a media type compatible with a 3590 device.
FFF7 & FFF8	-	Reserved for library. These categories are reserved for internal library functions.
FFF9	Library Manager	3490-Service-Volume Category. Volumes are assigned to this category by the Library Manager when it detects that a volume has a unique service cartridge VOLSER and a media type compatible with a 3490 device.
FFFA	Library Manager	Manually-Ejected Category. Volumes are assigned to this category when they are removed from the library under the control of an operator, not the control program. Volumes in this category are no longer available for any other operations except purge-volume category assignment.

Category (in hex)	Used by	Definition
FFFB	Library Manager	<p>Purge-Volume Category.</p> <p>When this category is specified in a Perform Library Function command with the Library Set Volume Category order and the volume is either in the misplaced state, is assigned to the exported category, or is assigned to the manually-ejected category, the specified VOLSER's record is deleted from the inventory.</p> <p>No volumes are associated with this category.</p>
FFFC	Library Manager	<p>Unexpected-Volume Category.</p> <p>This category is reserved for future use.</p>
FFFD	Library Manager	<p>3590-Cleaner-Volume Category.</p> <p>Cleaner volumes for 3590 type devices in the library are assigned to this category automatically.</p>
FFFE	Library Manager	<p>3490-Cleaner-Volume Category.</p> <p>Cleaner volumes for 3490 type devices in the library are assigned to this category automatically.</p>
FFFF	Library Manager	<p>Volser-Specific Category.</p> <p>This category is for general use by programming except that any Library Mount request to this category must be for a specific VOLSER and not based on the category only.</p>



DEVSERV QLIB command

This appendix lists the syntax and parameter explanations copied from the cover letter of the PTFs for APAR OA24965 (see Example H-1 on page 924) with DEVSERV QLIB command enhancements added in IBM Virtualization Engine TS7700 R1.5.

Additional support has been provided in APAR OA07505 (see Example H-2 on page 930).

The DEVSERV QLIB command can be used to:

- ▶ Request a list of Tape Library subsystems that are defined to the host. Libraries are listed by serial number (Library ID).
- ▶ Request a list of devices within a library. Devices are listed by device number and the library port for each device is displayed.
- ▶ Validate the connection status of devices in a library, for example, devices that are connected to the host.
- ▶ Delete an improperly defined library control block in preparation for an IODF activation.
- ▶ Issue a diagnostic state save order to a library when requested by the IBM Service Center. For the state save, you also need to apply the PTF for APAR OA09599.

Important: Do not use this **state save** command for just testing purposes. It will impact the performance of your VTS/ATL, because it consumes time to take the dump in the hardware.

- ▶ When using the DEVSERV QLIB command to display the subsystems (port-IDs) and drives associated with the specified Library ID, if the Library ID specified is for a composite library, the command will now display the distributed library IDs associated with the composite library.

Tip: You can use DEVSERV QLIB,? to get the complete syntax of the command:

IEE459I 13.16.57 DEVSERV QLIB 040

DEVSERV 'QLIB' command syntax:

```
DS QL,libid(,filter)
DS QL,LIST(,filter)
DS QL,LISTALL(,filter)
DS QL,libid,QUEUE
DS QL,LIST,QUEUE
DS QL,dddd,SS
DS QL,DDR
DS QL,IEA438I
DS QL,CATS|CATS(XXX*)
```

Example: H-1 APAR OA24965 text

APAR OA24965 text

DOCUMENTATION:

This new function APAR adds support to the DEVSERV command for a new Query Library option.

Use the Query Library option of the DEVSERV command to:

- * Request a list of tape library subsystems that are defined to the host. Libraries are listed by serial number (Library ID).
- * Request a list of devices within a library. Devices are listed by device number and the library port for each device is displayed.
- * Validate the connection status of devices in a library. For example, devices that are connected to the host.
- * Delete an improperly defined library control block in preparation for an IODF activate.
- * Issue a diagnostic state save order to a library when requested by the IBM Service Center.

Query Library can be abbreviated QLIB or QL and supports the following parameters:

```
DS QL,LIST(,filter)
DS QL,LISTALL(,filter)
DS QL,libid(,filter)
DS QL,dddd,SS
```

Parameters:

LIST- Indicates that QLIB should display a list of the ACTIVE Library IDs (ACTIVE is the default). You can optionally generate a list of INACTIVE Library IDs or QUEUE'd library orders. LIST uses the sub-parameters ACTIVE, INACTIVE, and QUEUE.

- LISTALL- Produces a detailed list of all libraries, including the devices and port-ids within each library. LISTALL uses the sub-parameters ACTIVE and INACTIVE (ACTIVE is the default).
- libid- List information for the library with serial number 'libid'. The 'libid' parameter uses sub-parameters ACTIVE, INACTIVE, VALIDATE, QUEUE and DELETE. ACTIVE is the default.
- dddd- Indicates that the request is either for the library that contains device dddd, or is for the device dddd itself. A sub-parameter is required when dddd is specified. dddd uses the sub-parameter SS.
- ?- Causes QLIB to display the command syntax.

Sub-Parameters:

- ACTIVE- Displays information about the library configuration that is currently in use by the system.
- INACTIVE- Displays information about the library configuration that will become active following the next IODF activate. The INACTIVE configuration is similar to ACTIVE, but may contain additional devices or libraries.
- VALIDATE- Displays the INACTIVE configuration. However, before the configuration is displayed, I/O is issued to each device in the configuration to validate the devices connectivity to the host.
- DELETE- Indicates that QLIB should delete the INACTIVE control blocks for library LIBID, but not affect the existing ACTIVE library definition. The DELETE command is used to remove incorrectly defined library control blocks so they can be rebuilt. DEVSERV DELETE provides an alternative to the method described in information APAR II09065, which requires two IODF activates.

The DEVSERV QLIB method is as follows:

1. Use QLIB DELETE to delete all of the devices from the incorrect control blocks.
2. Use QLIB LIST to display that the INACTIVE control blocks have been deleted.
3. Use ACTIVATE IODF to redefine the devices.
4. Use QLIB LIST to display that the ACTIVE control blocks are properly defined.

Note: the steps above assume that library devices are HCD defined with LIBID and PORTID. Using LIBID and PORTID enables the activate in step 3 (above) to build library control blocks. If

LIBID and PORTID are not defined, then the following alternate method must be used:

1. Use QLIB DELETE to delete all of the devices from the incorrect control blocks.
2. Attempt to vary ONLINE each device in the library. Each VARY should fail with message:

IEA437I TAPE LIBRARY DEVICE(dddd), ACTIVATE IODF
IS REQUIRED

3. Each VARY attempt in the previous step should add a device to the INACTIVE configuration. Use QLIB LIST to list the INACTIVE configuration and verify that devices are configured correctly. If there are configuration errors, correct them and begin at step 1.
4. Use ACTIVATE IODF to rebuild the ACTIVE configuration. This step replaces the currently ACTIVE configuration with the INACTIVE configuration. This step also rebuilds the allocation EDT's.
5. Use QLIB LIST to display that the ACTIVE control blocks are properly defined.

QUEUE- Lists the library orders that are waiting to be completed. Such orders include:

MOUNT, DEMOUNT, EJECT and AUDIT

When an order completes, the library notifies the host and the order is removed from the queue. The QL display can list orders for all libraries, or can be limited to a single library.

SS- Indicates that QLIB should issue a diagnostic state save to the library containing device dddd. This command is intended to be used at the request of IBM Support Center. For example, SS can be used to diagnose a hardware error that results in a mount failure message. Automated Operator code can extract the failing device number from the failure message, then insert the device in a QLIB SS command.

Examples-

DS QL,LIST

IEE459I 13.59.01 DEVSERV QLIB 478

The following are defined in the ACTIVE configuration:
10382 15393

DS QL,10382

IEE459I 13.59.09 DEVSERV QLIB 481

The following are defined in the ACTIVE configuration:

LIBID PORTID DEVICES

10382	04	0940	0941	0942	0943	0944	0945	0946	0947
		0948	0949	094A	094B	094C	094D	094E	094F

03	09A0	09A1	09A2	09A3	09A4	09A5	09A6	09A7
	09A8	09A9	09AA	09AB	09AC	09AD	09AE	09AF
02	09D0	09D1	09D2	09D3	09D4	09D5	09D6	09D7
	09D8	09D9	09DA	09DB	09DC	09DD	09DE	09DF
01	F990	F991	F992	F993	F994	F995	F996	F997
	F998	F999	F99A	F99B	F99C	F99D	F99E	F99F

DISTRIBUTED LIBID(S)

AAAAA BBBB

DS QL,10382,DELETE

*04 REPLY 'YES' TO DELETE THE INACTIVE CONFIGURATION FOR
LIBRARY 10382, ANY OTHER REPLY TO QUIT.

IEF196I Reply 'YES' to delete the INACTIVE configuration for
library 10382, any other reply to quit.

R 4,YES

IEE459I 14.01.19 DEVSERV QLIB 490

Inactive configuration for library 10382 successfully deleted

COMMENTS:

CROSS REFERENCE-MODULE/MACRO NAMES TO APARS

IGUDSL01 OA07505

CROSS REFERENCE-APARS TO MODULE/MACRO NAMES

OA07505 IGUDSL01

THE FOLLOWING MODULES AND/OR MACROS ARE AFFECTED BY THIS PTF:

MODULES

IGUDSL01

LISTEND

*/.

++ HOLD(UA17546) SYS FMID(HDZ11G0) REASON(DOC) DATE(05098)

COMMENT

(This new function APAR adds support to the DEVSERV command for
a new Query Library option.

Use the Query Library option of the DEVSERV command to:

- * Request a list of tape library subsystems that are defined to the host. Libraries are listed by serial number (Library ID).
- * Request a list of devices within a library. Devices are listed by device number and the library port for each device is displayed.
- * Validate the connection status of devices in a library. For example, devices that are connected to the host.
- * Delete an improperly defined library control block in preparation for an IODF activate.
- * Issue a diagnostic state save order to a library when requested by the IBM Service Center.

Query Library can be abbreviated QLIB or QL and supports the following parameters:

DS QL,LIST(,filter)
DS QL,LISTALL(,filter)
DS QL,libid(,filter)

DS QL,dddd,SS

Parameters:

- LIST- Indicates that QLIB should display a list of the ACTIVE Library IDs (ACTIVE is the default). You can optionally generate a list of INACTIVE Library IDs or QUEUE'd library orders. LIST uses the sub-parameters ACTIVE, INACTIVE, and QUEUE.
- LISTALL- Produces a detailed list of all libraries, including the devices and port-ids within each library. LISTALL uses the sub-parameters ACTIVE and INACTIVE (ACTIVE is the default).
- libid- List information for the library with serial number 'libid'. The 'libid' parameter uses sub-parameters ACTIVE, INACTIVE, VALIDATE, QUEUE and DELETE. ACTIVE is the default.
- dddd- Indicates that the request is either for the library that contains device dddd, or is for the device dddd itself. A sub-parameter is required when dddd is specified. dddd uses the sub-parameter SS.
- ?- Causes QLIB to display the command syntax.

Sub-Parameters:

- ACTIVE- Displays information about the library configuration that is currently in use by the system.
- INACTIVE- Displays information about the library configuration that will become active following the next IODF activate. The INACTIVE configuration is similar to ACTIVE, but may contain additional devices or libraries.
- VALIDATE- Displays the INACTIVE configuration. However, before the configuration is displayed, I/O is issued to each device in the configuration to validate the devices connectivity to the host.
- DELETE- Indicates that QLIB should delete the INACTIVE control blocks for library LIBID, but not affect the existing ACTIVE library definition. The DELETE command is used to remove incorrectly defined library control blocks so they can be rebuilt. DEVSERV DELETE provides an alternative to the method described in information APAR II09065, which requires two IODF activates.

The DEVSERV QLIB method is as follows:

1. Use QLIB DELETE to delete all of the devices from the incorrect control blocks.
2. Use QLIB LIST to display that the INACTIVE control blocks have been deleted.
3. Use ACTIVATE IODF to redefine the devices.
4. Use QLIB LIST to display that the ACTIVE control blocks are properly defined.

Note: the steps above assume that library devices are HCD defined with LIBID and PORTID. Using LIBID and PORTID enables the activate in step 3 (above) to build library control blocks. If LIBID and PORTID are not defined, then the following alternate method must be used:

1. Use QLIB DELETE to delete all of the devices from the incorrect control blocks.
2. Attempt to vary ONLINE each device in the library. Each VARY should fail with message:

```
IEA437I TAPE LIBRARY DEVICE(dddd), ACTIVATE IODF
      IS REQUIRED
```

3. Each VARY attempt in the previous step should add a device to the INACTIVE configuration. Use QLIB LIST to list the INACTIVE configuration and verify that devices are configured correctly. If there are configuration errors, correct them and begin at step 1.
4. Use ACTIVATE IODF to rebuild the ACTIVE configuration. This step replaces the currently ACTIVE configuration with the INACTIVE configuration. This step also rebuilds the allocation EDT's.
5. Use QLIB LIST to display that the ACTIVE control blocks are properly defined.

QUEUE- Lists the library orders that are waiting to be completed. Such orders include:

MOUNT, DEMOUNT, EJECT and AUDIT

When an order completes, the library notifies the host and the order is removed from the queue. The QL display can list orders for all libraries, or can be limited to a single library.

SS- Indicates that QLIB should issue a diagnostic state save to the library containing device dddd. This command is intended to be used at the request of IBM Support Center. For example, SS can be used to diagnose a hardware error that results in a mount failure message. Automated Operator code can extract the failing device number from the failure message, then insert the device in a QLIB SS command.

Examples-

DS QL,LIST

IEE459I 13.59.01 DEVSERV QLIB 478

The following are defined in the ACTIVE configuration:
10382 15393

DS QL,10382

IEE459I 13.59.09 DEVSERV QLIB 481

The following are defined in the ACTIVE configuration:

LIBID PORTID DEVICES

10382	04	0940	0941	0942	0943	0944	0945	0946	0947
		0948	0949	094A	094B	094C	094D	094E	094F
	03	09A0	09A1	09A2	09A3	09A4	09A5	09A6	09A7
		09A8	09A9	09AA	09AB	09AC	09AD	09AE	09AF
	02	09D0	09D1	09D2	09D3	09D4	09D5	09D6	09D7
		09D8	09D9	09DA	09DB	09DC	09DD	09DE	09DF
	01	F990	F991	F992	F993	F994	F995	F996	F997
		F998	F999	F99A	F99B	F99C	F99D	F99E	F99F

DISTRIBUTED LIBID(S)

AAAAA BBBB

DS QL,10382,DELETE

*04 REPLY 'YES' TO DELETE THE INACTIVE CONFIGURATION FOR
LIBRARY 10382, ANY OTHER REPLY TO QUIT.

IEF196I Reply 'YES' to delete the INACTIVE configuration for
library 10382, any other reply to quit.

R 4,YES

IEE459I 14.01.19 DEVSERV QLIB 490

Inactive configuration for library 10382 successfully deleted).

APAR OA07505 (Example H-2 on page 930) has additional support.

Example: H-2 APAR OA07505 text

This new function apar adds support to the DEVSERV command for
a new Query Library option.

.

Use the Query Library option of the DEVSERV command to:

- * Request a list of tape library subsystems that are defined to
the host. Libraries are listed by serial number (library-id).
- * Request a list of devices within a library. Devices are listed
by device number and the library port for each device is
displayed.
- * Validate the connection status of devices in a library. For
example, devices that are connected to the host.
- * Delete an improperly defined library control block in
preparation for an IODF activate.
- * Issue a diagnostic state save order to a library when
requested by the IBM Service Center.

.

Query Library can be abbreviated QLIB or QL and supports the
following parameters:

DS QL,LIST(,filter)

DS QL,LISTALL(,filter)

DS QL,libid,(filter)

DS QL,libid,QUEUE

DS QL,LIST,QUEUE

DS QL,dddd,SS

Parameters:

LIST: Indicates that QLIB should display a list of the ACTIVE library-ids (the default).

You can optionally generate a list of INACTIVE INACTIVE library-ids or QUEUE'd library orders. LIST uses the sub-parameters ACTIVE, INACTIVE, and QUEUE.

LISTALL: Produces a detailed list of all libraries, including the devices and port-ids within each library. LISTALL uses the sub-parameters ACTIVE and INACTIVE.

libid: List detailed configuration information for the library with serial number 'libid'.

libid: Indicates that the request is for a specific library. LIBID uses the sub-parameters ACTIVE, INACTIVE, VALIDATE, QUEUE, and DELETE.

dddd: A library device
Indicates that the request is either for the library that contains device dddd, or is for the device dddd itself. A sub-parameter is required when DDDD is specified. DDDD uses the sub-parameter SS.

?: Causes QLIB to display the command syntax.

Sub-Parameters:

ACTIVE: Displays information about the library configuration that is currently in use by the system.

INACTIVE: Displays information about the library configuration that will become active following the next IODF activate. The INACTIVE configuration is similar to ACTIVE, but may contain additional devices or libraries.

VALIDATE: Displays the same information as the INACTIVE configuration. However, before the configuration is displayed, I/O is issued to each device in the configuration to validate connectivity to the host.

DELETE: Indicates that QLIB should delete the INACTIVE control blocks for library LIBID and not affect the existing ACTIVE library definition. The DELETE command is used to remove incorrectly defined library control blocks so they can be rebuilt. DEVSERV DELETE provides an alternative to the method described in information APAR II09065, which requires two IODF activates. The DEVSERV QLIB method is as follows:

1. Use QLIB DELETE to delete all of the devices from the incorrect control blocks.
2. Use QLIB LIST to display that the INACTIVE control blocks have been deleted.
3. Use ACTIVATE IODF to redefine the devices.
4. Use QLIB LIST to display that the ACTIVE control blocks are properly defined.

QUEUE: Lists the library orders that are waiting to be completed. Such orders include:
MOUNT,DEMOUNT,EJECT and AUDIT
When an order completes, the library notifies the host and the order is removed from the queue. The QLIB display can list orders for all libraries, or can be limited for a single library.

SS: Indicates that QLIB should issue a diagnostic state save to the library containing device DDDD. This command is intended to be used at the request of IBM Support Center. For example, SS can be used to diagnose a hardware error that results in a mount failure message. Automated Operator code can extract the failing device number from the failure message, then insert the device in a QLIB SS command.

Examples:

DS QL,LIST

IEE459I 13.59.01 DEVSERV QLIB 478

The following are defined in the ACTIVE configuration:

10382 15393

.

DS QL,10382

IEE459I 13.59.09 DEVSERV QLIB 481

The following are defined in the ACTIVE configuration:

LIBID	PORTID	DEVICES
10382	04	0940 0941 0942 0943 0944 0945 0946 0947 0948 0949 094A 094B 094C 094D 094E 094F
	03	09A0 09A1 09A2 09A3 09A4 09A5 09A6 09A7 09A8 09A9 09AA 09AB 09AC 09AD 09AE 09AF
02		09D0 09D1 09D2 09D3 09D4 09D5 09D6 09D7 09D8 09D9 09DA 09DB 09DC 09DD 09DE 09DF
01		F990 F991 F992 F993 F994 F995 F996 F997 F998 F999 F99A F99B F99C F99D F99E F99F

.

DS QL,10382,DELETE

*04 Reply 'YES' to delete the INACTIVE configuration for library 10382, any other reply to quit.

IEF196I Reply 'YES' to delete the INACTIVE configuration for library 10382, any other reply to quit.

R 4,YES

IEE459I 14.01.19 DEVSERV QLIB 490

Inactive configuration for library 10382 successfully deleted

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 publications

For information about ordering these publications, see “How to get Redbooks” on page 935. Note that some of the documents referenced here might be available in softcopy only.

- ▶ *Continuous Availability - Systems Design Guide*, SG24-2085
- ▶ *Continuous Availability S/390 Technology Guide*, SG24-2086
- ▶ *DFSMSHsm Primer*, SG24-5272
- ▶ *Fiber Saver (2029) Implementation Guide*, SG24-5608
- ▶ *FICON Native Implementation and Reference Guide*, SG24-6266
- ▶ *Guide to Sharing and Partitioning IBM Tape Library Data*, SG24-4409
- ▶ *IBM System Storage Tape Library Guide for Open Systems*, SG24-5946
- ▶ *IBM System z Connectivity Handbook*, SG24-5444
- ▶ *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789
- ▶ *IBM TotalStorage 3494 Tape Library: A Practical Guide to Tape Drives and Tape Automation*, SG24-4632
- ▶ *Introduction to IBM S/390 FICON*, SG24-5176
- ▶ *Introduction to SAN Distance Solutions*, SG24-6408

Other publications

These publications are also relevant as further information sources:

- ▶ *DFSMSdfp Utilities*, SC26-7414
- ▶ *DFSMSHsm Storage Administration Guide*, SC35-0421
- ▶ *DFSMS/VM Function Level 221 Removable Media Services User's Guide and Reference*, SC35-0141
- ▶ *IBM Encryption Key Manager component for the Java platform Introduction, Planning, and User's Guide*, GA76-0418
- ▶ *IBM System Storage 3953 Library Manager Model L05 Operator Guide*, GA32-0448
- ▶ *IBM System Storage 3953 Tape System Introduction and Planning Guide*, GA32-0557
- ▶ *IBM System Storage Tape System 3592 Introduction and Planning Guide*, GA32-0464
- ▶ *IBM System Storage TS1120 Tape Drive and Controller Introduction and Planning Guide*, GA32-0555
- ▶ *IBM System Storage TS1120 Tape Drive and Controller Operator Guide*, GA32-0556

- ▶ *IBM System Storage TS3500 Tape Library Introduction and Planning Guide*, GA32-0559
- ▶ *IBM System Storage TS3500 Tape Library Operator Guide*, GA32-0560
- ▶ *IBM TotalStorage 3953 Tape Frame Model F05 and Library Manager Model L05 Operator Guide*, GA32-0473
- ▶ *IBM TotalStorage Enterprise Tape System 3592 Operators Guide*, GA32-0465
- ▶ *IBM TotalStorage TS3500 Operator Guide*, GA32-0468
- ▶ *IBM Virtualization Engine TS7700 Series Introduction and Planning Guide*, GA32-0567
- ▶ *Implementing System Managed Storage*, SC26-3123
- ▶ *Lights Out! Advanced Tape Automation Using VM/ESA*, GG24-4347
- ▶ *MVS Initialization and Tuning Reference*, SA22-7592-08
- ▶ *VM/ESA DFSMS/VM Removable Media Services User's Guide and Reference*
- ▶ *z/OS DFSMSdfp Storage Administrator Reference*, SC35-0422
- ▶ *z/OS DFSMSdss Storage Administration Reference*, SC35-0424
- ▶ *z/OS DFSMS OAM Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC35-0427
- ▶ *z/OS DFSMSrmm Guide and Reference*, SC26-7404
- ▶ *z/OS DFSMSrmm Implementation and Customization Guide*, SC26-7405
- ▶ *z/OS JES3 Initialization and Tuning Reference*, SA22-7550
- ▶ *z/OS MVS Planning: Operation*, SC22-7601
- ▶ *z/OS V1R3.0-V1R4.0 DFSMS Using Magnetic Tapes*, SC26-7412
- ▶ *z/OS V1R7.0 MVS System Commands*, SA22-7627
- ▶ *z/VM V5R1.0 DFSMS/VM Planning Guide*, SC24-6089
- ▶ *z/VM V5R2.0 DFSMS/VM Removable Media Services*, SC24-6090
- ▶ *z/VM V5R2.0 Running Guest Operating Systems*, SC24-6115
- ▶ *z/VM V5R2.0 DFSMS/VM Storage Administration*, SC24-6091
- ▶ *z/VSE System Administration Guide*, SC33-8224
- ▶ *z/VSE System Macros Reference*, SC33-8230

Online resources

These websites are also relevant as further information sources:

- ▶ Common Information Model (CIM)
<http://www.dmtf.org/standards/cim/>
- ▶ FICON Director support matrix
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/FQ116133>
- ▶ IBM Business Continuity and Recovery Services
<http://www.ibm.com/services/continuity>
- ▶ IBM Global Services home page
<http://www.ibm.com/services>
- ▶ Technote Performance Considerations for a Cascaded FICON Director Environment
<http://www-1.ibm.com/servers/eserver/zseries/library/techpapers/pdf/gm130237.pdf>
- ▶ TS7700 Information Center
<http://publib.boulder.ibm.com/infocenter/ts7700ic/v1r0/index.jsp>
- ▶ Web-Based Enterprise Management (WBEM)
<http://www.dmtf.org/standards/wbem/>

How to get Redbooks

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, and order hardcopy Redbooks, at this website:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

Symbols

\$\$INDEX file 161

Numerics

1 Gb Grid Dual Port Copper Connection (FC1032) 110, 112
1 TB Cache Enablement (FC5267) 109, 112
1.7 TB Fibre Storage (FC7120) 110
100 MBps Increment (FC5268) 109
16-device group 139
256-bit Advanced Encryption Standard (AES-256) 153
3480 163
3490E 29, 75, 139
3490E format 57
3490E Integrated Cartridge 28
3490E tape control units 23, 28
3490E tape drives 10
3490E tape technology 10
3490E volumes 57
3494 Library Manager 35, 176, 198, 321, 352–353, 360
3494 Tape Library 7–9, 35, 108, 175, 321, 327, 353, 357, 364, 373–374, 383, 399, 792, 825
3584 111, 191
 L22 51
 Library Manager 179
 logical library 180–181
 Operator Panel 572
 physical volumes 204
 zSeries host logical libraries 204
3584 Model L23 111
3584 Specialist 179, 629
3584/3953 305
3590 164
3590 Tape Drives 54, 350, 353, 377, 794
3592 48, 190, 192, 306, 412
3592 (Model S24) 116
3592 cartridges 51, 164
3592 drive family 118
3592 drives 118
3592 E05 119
3592 family 53
3592 JA cartridges 222
3592 Media Types JJ 48
3592 Model E05 110–111
3592 Model E05 formats 53
3592 Model E05 native mode 112
3592 Model E06 110
3592 Model J1A 53, 110–111, 124, 156
3592 Model J1A Tape Drive 53
3592 Model J1A Tape Drives 48
3592 Model J70 120
3592 Tape Drive 48
3592 Tape Drives 7–8, 48–49, 110
3592 Tape Drives (FC5240) 109

3592 WORM 75
3592-E05 49
3592-E05 Tape Drives 53
3592-E06 data format 49
3592-formatted data 350
3592-J1A 28, 48–49, 52–54, 108, 118–119, 188, 352, 821, 823–824, 826–828, 830, 860, 874, 878
3592-J1A drives 49, 119
3592-J1A emulation mode 119
3592-J70 52
3592-J70 Controller 190
3593 Library 195
3952 cache expansion frame 17
3952 F05 (FC9350) 110, 112
3952 F05 (FC9352) 113
3952 Model F05 Storage Expansion Frame 38
3952 Model F05 Tape Base Frame 38
3952 Storage Expansion Frame 43
3952 Tape Base Frame 36, 43–44
3952 Tape Frame 36, 46, 125–126
3952 Tape Frame Model F05 109, 111–113
3952-F05 108–109
3952-F05 frame 37
3953 35, 55, 178, 187
 F05 934
 L05 934
3953 frame 110
3953 Library Manager 195, 323, 353, 360, 363, 365–366, 369, 373
 IBM 3584 Tape Library 203
 VOLSER ranges 203, 205
3953 Model F05 110–111
3953 Model L05 111
3953 Tape System 176
3953-F05 54–55, 120, 320–321, 324, 360
3953-F05 frame 110
3956 CC8 5
3956 CS7 (FC5636) 112
3956 Model CC8 44
3956 Model CS7 109
3956 Model CS8 38
3956 Model CS8 Cache Controller RAID 6 38
3956 Model CS8s 43
3956 Model CX7 44, 112
3956 Model CX7 Cache Drawers 112
3956 model CX7s 44
3956 Model XS7 38, 43
3956-CC6/CX6 cache 122
3956-CC7/CX7 122
3956-CX6 123
3957 Model V06 44, 110
3957 Model VEA 38, 113
3957 Model VEA features FC2714 110, 113
3957 VEA 122
3957 VEA (FC5626) 112

- 50 Micron LC/LC 31 meter fibre cable 117
- 62.5 Micron LC/LC 31 meter fibre cable 117
- 9 micron 117
- 9 micron diameter single-mode optical fiber 133
- 9 Micron LC/LC 117
- 9 Micron LC/SC 31 meter 117

A

- ABACKUP 385
- ABARS 384–386
- AC power 126
- accessor failure 583
- ACS 96, 140, 168
 - constructs 362, 365, 368, 371, 374
 - routines 59, 140, 264, 286, 288, 292, 296, 352, 358–359, 363, 377, 379–380, 383, 390, 394, 801
- active control data set (ACDS) 140
- active data 70, 144, 216, 919
 - minimum amount 147
 - threshold 215
 - threshold percentage 70
- Active Data Distribution 695
- active log
 - DASD data set 403
 - data 403
- active volume list 70
- Adjacent Frame Support for Back-End Drives 111
- Advanced Library Management System
 - See ALMS
- Advanced Policy Management 68
- affinity 75, 78
- allocation processing 85
- ALMS 51, 179–181, 188, 197
 - window 180
- alternate source category 66–67
- analyze workload 161
- AOTM 78–80, 138
- Application-Managed encryption 160
- architectural capability 82
- archive log 402
- ASMI 131
- Asynchronous Operation Manager (AOM) 140
- ATL 168, 567, 918
 - environment 168
- attach to 3957 V06 (FC9230) 110
- AUTOBACKUP 387
- AUTODUMP 392
- Automated Library 168
- automated tape library
 - See ATL
- automatic utilization 166
- Autonomic Ownership Takeover Manager
 - See AOTM

B

- B20 VTS model 350
- back end 10, 284
- back-end Fibre Channel switch 111
- backup copy database 583

- backup pool 146
- Backup/Restore 258
- BADBLKSZ 161
- bandwidth 118
- bar code label 197
- batch window efficiencies 166
- battery backup units 47
- beginning of tape (BOT) 74
- blank character 685
- BMCNTL.BIN 161
- BMPACKT 162
- borrowing capability 216
- borrow-return policy 214
- BOT 74–75
- bottlenecks 163
- Brocade 135–136
- buffer credit 133
- Bulk Volume Information Retrieval
 - See BVIR
- BVIR 31, 35, 71, 103, 148, 215, 220, 574, 710, 889
 - binary format 711
 - request 654
 - response 628
 - data 656, 665
 - upper area 696

C

- cabling infrastructure 133
- Cache Controller attachment 42
- Cache Controller Model CC7 112
- cache enablement features 123
- cache increments 123
- cache management 62
- Cache Management Policies 27, 339
- cache removal 259
- Cache Removal Policies 3, 343
- cache residency 15, 89, 92
- cache subsystem 16
- Call Home function 35
- CAP 178, 191–196, 203–204, 343, 376, 471, 572–573
- CARS (Compliance, Availability, Retention, and Security) 13
- Cartridge Assignment Policy
 - See CAP
- cartridge capacity 166
 - utilization 164
- cartridge slots 183
- cartridge system tape
 - See CST
- cascaded switches 132
- Cat 5e cables 132
- Cat 6 cables 118, 129, 132
- categories assignment 141
- category FF04 65
- CC8 7, 18
- CCP 83, 85
 - policies 84
- CDS 299, 377, 396, 576, 910
- CE/DE 136
- change management policies 103

- Channel Control Word (CCW) 284
- channel extender interfaces 136
- cleaner cartridge 198
- cleaning cartridge 195, 198
- Cluster 0 91, 247, 668
 - copy consistency points 94
 - Node 0 32
 - virtual devices 91
 - virtual drive 92
- Cluster 1 92, 669
 - deferred setting 92
 - virtual devices 91
 - virtual drive 93
- cluster availability 87
- cluster CCP 85
- cluster cleanup 100
 - feature 100–101
 - process 13, 103
- cluster copy consistency point 87
- cluster families 5, 127, 250
- Cluster Grid 108, 130
- Cluster outage 332
- cluster removal 100
- CNT/Inrange 136
- CNTLUNIT definition 136
- code level 353, 653, 655
- collocation 397
- commands
 - LIBRARY 280
 - LIBRARY EJECT 540
- common Grid 77
- Common Scratch Pool (CSP) 64–65, 216, 573, 656, 713
- Communication Grid 114
- complex-wide device type 823
- complex-wide name 823
- compliance 13
- Composite Library 18, 23, 25–26, 31, 35, 57, 68, 72, 75, 78, 139, 202, 264, 266–269, 276–277, 280–281, 283, 286, 288, 297, 314, 329, 331, 334–335, 337, 343, 353–354, 358, 360, 362–363, 365–366, 368–372, 375, 410, 563–564, 580, 722, 812, 823, 828, 830
 - LIBRARY command 280
- Composite Library ID 23, 200, 202, 567
- compressibility 143
- compression 148
- compression ratio 29, 88, 145, 161, 388, 637
- concurrent system upgrades 122
- Consistency Configuration 91, 93
- consistency point 86–87, 91, 103, 137
 - of RUN 96
 - settings 93
- Consistency Point policy 26
- Console Attachment (FC2715) 110, 113
- Console Expansion (FC2714) 110, 113
- Console Launcher 119
- Console Request function 148
- construct parameters 141
- control path 190–191
- control path drive 190
- control path drive definition 191

- control unit (CU) 139, 265, 270
- cooperative replication 624
- Coordinated Universal Time (UTC) 32
- copper connections 124
- copper networks 118
- Copy Consistency point 314
- copy consistency point 26, 82–83, 85, 87–90, 92–94, 97, 99, 138, 152, 249, 332–333, 336–339, 574, 665, 726, 740
 - Management Class 754
 - of Deferred 83
 - of RUN 86, 96–97
 - specified management class 152, 249
- copy consistency point settings 91
- Copy Consistency Policy 60
- Copy Count Override 90
- copy count target 85
- Copy Count value 90
- Copy Export 4, 55, 66, 70–72, 87, 145–146, 166, 208, 286, 288–294, 296–302, 384–386, 395–396, 462, 520, 534, 536, 546, 562, 603, 606–607, 609, 651, 681, 704, 725, 741–751, 778, 889, 909
 - function 70–71, 87, 146, 166
 - job 72
 - operation 70–72, 146, 289–291, 294–295, 297, 299–302
 - pool 66, 72
 - sets 72
 - volumes 71
- Copy Export operation 384, 651, 704, 744, 748
- Copy Export state 146
- Copy Exported physical volumes 166
- Copy Exported state 71
- Copy Failure Reporting 15
- Copy Management 58
- Copy Operation 67, 381, 421, 638
- copy policy management requirements 82
- Copy Time Management 58
- count column 691
- CPU 33
- CPU cycles 166
- Cr value 147
- Create Logical Library 182
- critical data 146
- cross-cluster mount 85
- cross-site Grid 132
- cryptography function 73, 153
- CSP 65
- CST 29, 143, 390, 392–393, 456, 560–561, 671
 - emulated cartridges 142
 - emulation 143
- customer service level 137
- customer-configured policy 80
- CX7 Cache Drawer 47
- CX7s 44

D

- D23 frame 111
- D32 51
- DAA 21, 84, 94–95
- DAA function 94–95

- daily production workload 144
- DASD 42, 58, 142, 166, 377, 383–384, 387–388, 393, 399, 401–402, 542, 832
- data cartridge 197
- Data Cartridges window 197
- Data Class 29, 177, 287, 377, 801
 - ACS routine 378
 - construct 30, 75
 - definition 148
 - Name 29, 543
 - parameter Compaction 148
 - policy name 543
 - storage construct 142
- Data Class ACS
 - routines 378
- data encryption solutions 153
- data erase function 69, 149–150
- Data Facility Storage Management Subsystem
 - See DFSMS
- Data Key 153
- data loss 753
- data management 6, 258
- data protection 153
- data record 660, 685
- data replication functionality. 12
- data set 140, 165, 380–381, 658, 661, 799
 - existing policies 761
 - minimum size 391
- data sets, old 834
- data type 378, 688–689
 - response data 688
- database backup 398
- DB2 catalog 380
- DB2 database 30–31
 - full image copy 403
- dddd,SS (DS) 926, 930
- DDMs 43–44, 47
- DDR2 memory 46
- dedicated pool 150
- Default Encryption Keys 73–74
- Default key 74
- default Storage Class 60
- deferred (DD) 138
- Deferred copy consistency point 87, 96, 333
- deferred copy mode 26, 58, 336, 339, 593, 717
- deferred copy throttling value 76
- define
 - fast ready 220
 - space reclamation 212
- defining a logical library 179, 188, 190–191, 203, 236, 245
- Delete Expired Volume Data 146, 222
- Delete Expired Volume Data setting 146, 222
- Dense Wave Division Multiplexers (DWDMs) 134
- Device Allocation Assist 15
- Device Allocation Assist (DAA) 94
- device class 285
 - definition 380, 398
 - IBM-supplied default 285
 - MIH default 285
- device dddd 925–926
- device drivers 153
- device service
 - eject requests 576
- device service (DS) 282–283, 576
- Device State Change attentions 31
- DEVICE statement 824, 827
- device type 273, 566, 654–655
 - 348X 567
 - command displays information 566
- DEVSERV command 924
 - Query Library option 924, 927
 - to find LIBPORT-ID 268
- DEVSUPnn 222
- DEVSUPxx 148, 220, 286, 913
- DEVSUPxx member 220
- DFSMS 57, 59, 96, 137, 140–142, 154, 159, 168, 170, 177, 200–201, 222, 232, 260, 263, 266, 286, 288, 292, 304–305, 351, 353, 385–386, 390, 399, 540, 576, 584, 665, 720–721, 740, 797–801, 822, 832–835, 903, 911, 913–915, 917
- DFSMS Automatic Class Selection routines 97
- DFSMS catalog processing 834
- DFSMS catalog services 834
- DFSMS construct 73
- DFSMS constructs 153
- DFSMS library definitions 353
- DFSMS OAM 67
- DFSMS SCDS 800–801
- DFSMS tape solution 168
- DFSMS VOLREF processing 835
- DFSMS/MVS 177, 286
 - define composite libraries 287
 - define distributed libraries 287
- DFSMSdss 378, 398–399, 401
- DFSMSShsm 377–378, 380, 382, 384, 386–394, 396–398, 402, 405, 654, 714
- DFSMSShsm tape data
 - volume count 388
- DFSMSrmm 144, 170, 286, 299, 379–382, 395–396, 405, 575–576, 798, 800, 910
 - synchronize with TCDB 576
- DFSORT 379
- direct-attached drives 156
- Director Address 136–137
 - hex 137
 - ranges 136
- Director GUI 136
- Director GUI setup 136
- Director ID 136
- Director Port 137
- Director Switch ID 136
- Director Switch ID ranges 136
- disaster 166
- disaster recovery (DR) 12, 97, 166–167, 277, 386, 401, 404, 583, 725, 740, 752, 755, 763, 918
 - planning 739
 - solution 98
 - temporary category 918
- disaster recovery cluster 99

- discontinuance 103
- disk cache 108
- Disk Cache controller 25
- disk-based cache 10
- dismount operations 176
- display library information 541
- DISPLAY SMS 279–280, 542, 580, 719
 - sample output 280
- DISPLAY SMS,OAM 579
- disruptive 343
- distributed libraries 18
- Distributed Library 25–26, 31, 200, 202–203, 266–267, 269, 280–281, 286, 288, 329, 334–335, 337, 353–354, 358, 362, 365, 368, 563, 580, 664, 689, 698, 703, 810–812, 816, 828, 843, 848, 852, 855, 859, 863, 866, 873, 886
- Distributed Library ID 202
- distributed library ID 139
- distributed library names 203
- Drive Assignment 187
- Drive Assignment window 191
- drive cleaning 197
- drive concurrency 164
- Drive level Default key definitions 74
- drive swap 582
- DS QL 282–283, 924, 927
- DSP 822, 824, 827
- Dual AC Power (FC1903) 109, 112
- Dual AC Power supply feature 36
- Dual Cluster Grid 19, 23, 128, 130, 264, 638, 640
 - tape drives 264
- Dual Control Path Facility 353
- Dual Copy 65–66, 145–146, 402
- Dual Optical SW Connection (FC1033) 110
- dual-ported 1 Gbps Ethernet adapters 37
- dual-ported copper 124
- DUMP 377, 392, 399, 405
- DWDM 97, 99–100, 133–134
- Dynamic Grid Load Balancing 26–27
- Dynamic Grid Network Load Balancing enhancements 129

E

E05

- format 48
- formatted tapes 49, 118
- models 53
- native mode 119

E06

- data format 49
- format 118

EDGINERS program 67

education 167

effective Grid cache size 95

Eight-Character-VOLSER support 178

EJECT logical volume 575

eject stacked volume function 69

EKM 3, 153, 156, 160

- addresses 157, 236
- configuration file 156

- must be installed first 156

ELC 21, 48, 54–56

ELC Library Manager 360

electronic call home 119

Eligible Device Table (EDT) 275

empty scratch volume 143

empty stacked volume 148

empty storage cells 193

Emulated 3490E volumes 57

emulated tape volumes 57

emulated volumes 11

emulation mode 48, 119, 135–136

encrypted data keys 149

encrypted physical volume 149

encryption 13, 132, 153

encryption certificates 74

encryption characteristics 73

Encryption Configuration (FC9900) 112

encryption infrastructure 74

Encryption Key Manager

See EKM

encryption keys method 69

encryption method 160

encryption mode 160

encryption parameters 73

encryption settings 157

encryption solution 153

encryption support 73, 153

encryption-capable E05 drives 156

encryption-capable tape drives 156

encryption-capable. 160

encryption-enabled 160

end of file (EOF) 662

Enhanced Cache Removal Policies 3

enhanced capacity cartridge system tape (ECCST) 29, 390, 671

Enhanced Library Controller attachment 39

Enhanced Removal Policies 104

enhanced statistical reporting 4

Enterprise Automated Tape Library (ETL) 410

Enterprise Library Controller 108

Enterprise Tape

Cartridge 64–65, 913–914

Library Specialist 628

EOV 380, 389, 393, 396

erase, functionality 149

erasure reclamation policy 150

EREPMDR 161

ESCON 32, 50, 133, 284, 354, 359, 362, 365, 368, 372, 375

director 141

Ethernet 629

Ethernet connection 77, 93, 730

Ethernet IP address 197

Ethernet router 36

ETL Specialist 60, 411, 629

logical volume 262

status 424, 429

TS7700 Virtualization Engine Management pages 631

- expire hold 68
- Expire Time 222–224
- expired data 149
- expired logical volume 68
- expired volume attribute 68
- expired volume management 68
- expired volumes 68
- Export List File Volume 289, 291, 298, 748
- Exported physical volumes 71
- exported physical volumes 71
- export-hold category 72
- External Director connections 137
- external ELC 55

F

- failover scenario 730
- false-end-of-batch 643
- family affinity 87
- Fast Ready 67, 223
 - explanation 220
- Fast Ready attribute 28, 68, 220
- Fast Ready Categories 138, 142, 221–222
- Fast Ready Categories Table 221
- Fast Ready category 61, 68
- Fast Ready Category Configuration 56
- Fast Ready definition 220
- Fast Ready Mount 92
- Fast Ready Mounts 91
- fast-ready attribute 28, 310, 391, 578
- fault-tolerant disk subsystem. 57
- FC1950 Power Distribution Units 111
- FC9900 License Key 160
- FC9954 109
- feature codes 125
- fetch messages 835
- fiber optic cable 133
- Fibre Channel 133
 - architecture 133, 266
 - node 133
 - port 133
 - processor 44
 - processor cards 43, 47
- Fibre Channel Protocol 33
- Fibre Channel Switches 110
- FICON 8, 17, 28–29, 32, 50, 88, 91, 95, 110–111, 113–114, 116–117, 121, 132–134, 137, 141–142, 164, 177, 244, 264–266, 269, 277, 317, 331, 335, 338, 342, 345–347, 351, 354, 359, 362, 365, 368, 372, 375, 591, 593, 595, 597–600, 604, 714, 730–731, 783, 785, 787–788, 794, 818, 849, 855–856
 - adapters 42, 46, 117, 122
 - attachments 114
 - channel 97, 109, 137, 142, 264, 740
 - channel attachments 23
 - channel extenders 135
 - channels 116
 - extended distances 134
- FICON 10 km Longwave Attachment (FC3443) 110
- FICON Director 134, 272, 730
 - port 354

- FICON Directors 134–135, 137
- FICON Longwave Attachment (FC3442) 111, 114
- FICON Long-Wavelength adapter 117
- FICON Long-Wavelength attachment 122
- FICON Long-Wavelength channels 117
- FICON Shortwave Attachment (FC3441) 110
- FICON Short-Wavelength attachment 117, 122
- FICON Short-Wavelength channels 117
- FICON Switches 134
- FICON-attached 141
- Field Delimiter 664
- first choice 65, 381
- first-in first-out (FIFO) 63
- flexible replication policies 166
- floating deferred site 90
- flow control 133
- Force Local Copy setting 62
- four drawer cache 124
- four-cluster grid 13, 19, 21, 31, 34, 90, 125, 130–131, 139, 199–200, 248, 264, 267, 276–277, 320, 327, 340–341, 343–344, 347, 599, 621–622, 625–626, 635–637, 639–640, 740, 764
 - configuration 100
 - configurations 15, 127
- four-cluster hybrid grid 21, 342, 345, 600, 738
- four-cluster TS7700 Virtualization Engine configuration 84
- fragment 166
- frame shuttle 135
- Free-space Management 58
- full cache copy throttling 76

G

- Gbit Ethernet link 130–131
- GDPS 90–91, 152, 250, 393, 725
- General Node 16
- Geographically Dispersed Parallel Sysplex
 - See GDPS
- giant magnetoresistance (GMR) 53
- Global Default key definitions 74
- global marketplace 13
- Global Technology Services (GTS) 382
- gNode 16, 34
- graphical user interface (GUI) 119
- Greenwich Mean Time (GMT) 32
- Grid 12, 60–63, 77, 79, 82–85, 88, 90, 95, 100, 103, 118, 122, 137–139, 142, 165
 - adapters 117
 - clusters 57, 103
 - configuration 3, 12–13, 18, 23, 33, 58, 63, 70, 82, 88, 100, 103–104, 124, 137–139, 635, 637, 725–726, 739
 - Cluster 1 93
 - consistency point settings 93
 - other clusters 12, 18
 - TS7700 Virtualization Engine
 - Cluster 726–727
 - connection 124
 - definition 326
 - enablement 26, 111
 - environment 129

- links 137
- Grid capacity 103
- Grid component 30
- Grid configuration
 - TS7740 Virtualization Engine Cluster 666
- Grid Connection adapters 111, 114
- Grid Enablement (FC4015) 111, 114
- Grid environment 26, 65
- Grid implementation 104
- Grid interconnection 79
- Grid interconnections 95
- Grid join capability 104
- grid layer 33
- Grid links 27
- Grid mount points 82
- Grid Network 118, 326
 - connectivity 118
- Grid network 4, 18, 57, 79, 88
- Grid Network Load Balancing 15
- Grid subsystem 32
- Grid total copy count 85
- Grid WAN connection 21
- Grid. 60

H

- HA1 Frame 583
- hardware configuration definition 28, 139, 175, 263, 269
- Hardware I/O configuration definition 177
- Hash 74
- HBA 133, 637
- HCD 28, 139, 170, 175, 177, 202, 263–264, 266–274, 276, 305, 329, 334–335, 337, 342–343, 345–347, 350–351, 354, 359, 362, 365–366, 368, 372, 375, 778, 798, 822–823, 827, 925, 929
 - dialog 266
- HCD/IOCP 139, 177
- HD frame 115
- HD frame model S24 116
- HD slots 116
- HDR1 record 67
- Head Overcoat technology 53
- Health & Monitoring 203
- hierarchical storage management (HSM) 41, 45
- Hierarchical Storage Manager (HSM) 381, 386, 654
- high availability 12, 51, 137
- high availability solutions 97
- high capacity media 11
- High Density Capacity on Demand (HD CoD) feature 116
- High Voltage Differential (HVD) 50
- high watermark setup names 825
- high-capacity tape cartridges 57
- Historical Statistics 35
- HLQs 163
- hNode 16, 34, 81, 591, 654–655
- Hold function 224
- Home Pool 205
- homogeneous multi cluster 38
- homogenous TS7720 Virtualization Engine configurations 60
- horizontally scaling 34

- host attachment cables 117
- host bus adapter
 - See HBA
- host channel capabilities 122
- Host Command Line Query Capabilities 61
- host compression definition 148
- host connectivity 96
- Host Console Request function 58, 77, 148, 215
- Host Copy Control 77
- host definitions 177
- host I/O operation 731
- host system 263, 285, 395, 400, 917
 - 3490E tape drives 285
- Host System Attachment 117
- hot spares 113
- Hot Tokens 31
- HSM 41, 45, 65, 386, 389–390, 394, 653–655, 713, 716, 840–841, 858–860, 862, 865–867, 869–870, 885
- HWSNAME 825
- hybrid configurations 60, 124, 339, 742
- hybrid four-cluster grid 343
- Hybrid Grid 11, 15, 19, 21, 248, 281, 339–342, 345, 445, 522, 596–598, 600, 623, 625, 627, 675, 738, 741
 - definition of 521
 - implemented by R1.6 11
- Hydra Storage Manager
 - See HSM

I

- I/O Expansion drawers 108–109
- I/O operation 21, 83, 152, 250, 258, 277
- I/O station 51, 191, 193–194, 197–198, 568, 573
 - tape library 51
- I/O Station window 197
- I/O TVC 21, 63, 67, 82–83, 85, 87–90, 93, 754
 - PG0 assignments 63
 - selection 87, 89
 - selection criteria 62
- IART 60
- IART value 60
- IBM 3490 control unit 28
- IBM 3490E
 - device 284
 - I/O 28
- IBM 3490E Tape Drives 57
- IBM 3494 models Lxx 120
- IBM 3494 Tape Library 9, 175, 177
- IBM 3592 51
- IBM 3592 media 192
- IBM 3592 Model EU6 Tape Drive 118
- IBM 3592-J1A 28
- IBM 3592-J1A Tape Drive 52
- IBM 3952 Storage Expansion Frame 36
- IBM 3952 Tape Base Frame 37
- IBM 3952 Tape Base Frame Model F05 36
- IBM 3952 Tape Frame 36, 109, 125
- IBM 3953 52, 108, 352
- IBM 3953 Library Manager 267, 582
- IBM 3953 Tape System 48, 108, 176
- IBM 3956 Controller 131

- IBM Business Continuity 169
- IBM ETL Specialist 205
- IBM Information Infrastructure 13
 - solutions 13
- IBM Integrated Technology Services 168
- IBM Multi-Path Architecture 50, 52
- IBM POWER 5++ System 45
- IBM service intervention 224
- IBM SSR 55, 57, 81
- IBM SSR settings 82
- IBM standard label 28
- IBM Support Center 35
- IBM System p 108–109
- IBM System Service Representative (SSR) 31, 127, 178, 187, 259, 412, 567, 754, 800
- IBM System Storage TS1120 Tape Drive 52
- IBM System Storage TS1120 Tape Drive (3592-E05) 53
- IBM System Storage TS1130 Tape Drive 7
- IBM System Storage TS3500 7
- IBM System Storage TS3500 Tape Library 8
- IBM System x 50
- IBM System z 50, 176
 - environment 177–178
 - host 49, 132, 139, 188, 264
 - host support 188
 - logical library 194
 - operator command 284
 - production work 187
 - software 263
 - tape drive attachment 188
- IBM System z environment 187–188
- IBM System z host 136, 179
- IBM System z software 8
- IBM tape encryption solutions 153
- IBM Tape Library 141, 274
- IBM tape management system 311
- IBM tape product 168
- IBM Tivoli Storage Manager 379–380, 396–398
- IBM TotalStorage 3592-J1A 54
- IBM Virtualization Engine 3957 131
- IBM Virtualization Engine TS7700 3
- IBM VTS 350
- IBM z/OS 37
- IBM z/TPF 37
- IBMJCL.BIN JCL 161
- IBMLoad.BIN 161
- IBMLZ1 algorithm 29
- IBMPAT.BIN 161
- ICL 28
- IDRC 393, 403, 567
- IECIOSxx 284–285
- image copy 378, 403
- immediate mode copy throttling 76
- IMS 285, 404
- INACTIVE configuration 926–927
- inboard director (ID) 136, 139
- inboard Director Switch ID 136
- incremental data throughput 122
- independent logical subnets 130
- in-flight frames 133

- information availability requirements 13
- information compliance 13
- information retention periods 13
- information security threats 13
- Information Services 13
- infrastructure initiatives 4
- infrastructure to provide (I/P) 166
- Inhibit Reclaim Schedule 150, 215, 217–219
- inhibited reclamation 150
- Initial Access Response Time (IART) 60
- initial program load (IPL) 399
- input/output definition file (IODF) 267, 274, 354
- Insert category 66, 918–919
- Insert Notification 194, 196, 573
- installation
 - post-installation tasks 178
 - SYS1.PARMLIB 286
- integrated cartridge loader (ICL) 28, 540
- Integrated Control Path (FC5759) 109, 112
- Integrated Enterprise Library Controller 108–109
- Integrated Library Manager 26, 35, 56
- Integrated Library Manager support 54
- Interchange 166
- interface data (ID) 653, 655
- Internal Director connections 137
- internal Gbit Ethernet Link 130
- internal volume ID 146
- intervention required
 - display on LM 582
- Intraframe Fibre Channel cable (FC6003) 110
- intranet 128
- intranet traffic 132
- invalid logical volume data 149
- invalidated logical volumes 70
- Inventory Upload 204
- IODF 137, 264, 267, 274–275, 282, 330, 334, 338, 345, 360, 363, 366, 369, 372, 375, 923–931
- IOSTATS 709–710
- IP address 74, 128–129, 197, 538
- IP addresses per cluster 129
- IPL 264–265, 267, 274, 282, 285, 399–401, 798–799, 846
- IPv6 117
- ISMF 26, 143, 202, 268, 330, 343, 359, 365, 368, 383, 385, 395, 410, 540, 542, 770, 773, 776, 799–801
 - ISMF EJECT line operator command 383
 - Library Define window 143
 - screen 286, 383, 575
- IT infrastructure 168
- IT life cycle 168
- ITSM 396

J

- J1A
 - drives 156–157
 - formatted tape 49
- J1A Emulation mode 49, 52, 124, 146
 - TS7700 Virtualization Engine run 52
- J70 52–53, 120, 366, 373, 412, 789, 804
- JA cartridge (MEDIA5) 119

- JA cartridges 146
- JA, JB (JJ) 64
- JA-ETC 205
- JB 64
- JB cartridge (MEDIA9) 119
- JB tape cartridges 49
- JB-ETCL 205
- JCL 11, 148, 161, 164, 166, 295–296, 305–306, 330, 377–379, 383, 399, 405, 658, 706, 709, 712, 823, 832–833, 835, 889–893, 895, 897, 899, 901–904, 906–907, 909–911
- JES 2 84
- JES3 85
 - allocation and mounting 835
 - device pool 822
 - DSP 822
 - initialization deck 822
- JES3 INISH deck 823–824
- JES3/DFSMS processing 833
- JJ cartridge (MEDIA7) 119
- JJ cartridges 65
- JJ-EETC 205
- job completion 247
 - valid copy 247
- job stream 161
- job throttling 88
- joining cluster 103

K

- key labels 74
- key management request 74
- key manager 74, 156
- key manager services 74

L

- L22 51
- L23 110
- LAN 19
- LAN PCI Library Manager 353
- LAN ports 37
- LAN/WAN 118, 132
 - IP 118
- larger core diameter 133
- laser 133
- law of diminishing returns 33
- Lc 147
- LC Duplex
 - connections 117
 - connector 117
- LC Duplex connector 37
- Lc value 147
- LCU 266
 - virtual devices 276
- LCU definitions 359, 362, 366, 368, 372, 375
- LDG 822, 825
 - LDD 825
 - LDE 825
- least recently used
 - See LRU

- LIBPORT-ID 139
- library control system (LCS) 140, 541
- library device group (LDG) 822
- LIBRARY EJECT command 540
- Library Emulation 52, 54
- Library Emulation Capable 63
- Library Manager 21, 64, 224, 285, 411, 798
 - category 203
 - console 205
 - Fast Ready categories 222
 - functions 176
 - multiple logical scratch pools 260
 - VOLSER ranges 178
- Library Manager database 350, 358, 360, 362, 910
- Library Name 82, 286, 308, 353, 542
- Library operator window 179
- library outages 4
- Library partition 17
- Library Sequence Number 203
- library slot capacity 148
- Library Specialist 179
- Library-ID 25, 139, 267, 275, 287, 329, 334, 337, 343, 354, 359, 362, 365, 368, 371–372, 375, 565, 660, 665
- Library-Managed 160
- library-specific device type name 823
- library-specific name 823
- LIBSERV JCL 305–306
- licensed functional characteristics 104
- Licensed Internal Code 121
- Linear Tape Open (LTO) 49, 192
- link extenders 133
- Local Cache 89, 151, 248, 754
- local cluster 86, 122, 130, 165, 754
- Local Cluster for Fast Ready Mounts override 138
- local copy 151, 248–249, 402, 404, 754
- local site
 - Cluster 0 92
 - only connection 95
- local TS7700 Virtualization Engine 98
- local TS7700 Virtualization Engine Cluster 96
- local TVC 89–93, 754
 - valid copy 94
- locality factor 92
- locally attached TS7700 Virtualization Engine cluster 94
- logical block ID. 29
- logical control unit
 - See LCU
- logical library 178, 183–184, 186, 191–192, 197, 203–204, 264, 267, 541
 - Cartridge Assignment Policy 191
 - control path drive 191
 - drive assignment window 186
 - existing TS3500 Tape Library concept 52
 - first four drives 191
 - in TS3500 Tape Library 196
 - library ID 264
 - library sequence number 264
 - name 182, 186
 - partition 193, 197, 572
 - separate applications 52

- sharing between hosts 264
- starting SCSI element address 183
- tape drives 52, 187
- tape units 541
- view 411
- logical mount request 165
- logical volume 11, 21, 118, 131, 141–143, 145–147, 149–150, 153, 156, 164, 205, 216, 223, 305, 352, 437, 726, 730, 798, 800
 - current copy 67
 - data 148
 - data movement 88
 - dual copies 66
 - EJECT 575
 - following number 143
 - host request 78
 - insert time 260
 - large numbers 575–576
 - management 54
 - maximum number 242, 447
 - orphan 585
 - ownership takeover 80
 - physical location 64
 - previous use 69
 - primary pool 66
 - recall time 401
 - recent instance 70
 - relative order 686
 - scratch volume recovery 575
 - secondary pool 66
 - size 143
 - sizes 142–143
 - small number 683
 - status 424, 429
 - total number 142, 242, 447
 - unused capacity 397
 - valid copy 247
 - valid version 247
 - written data 391
- logical volume copy count 84
- Logical Volume Removal Policy 61
- logical volumes 63, 74
- logical WORM 127, 142, 233
 - support 15
- Longwave Attachment (FC3443) 111, 114
- Long-Wavelength Attachment 117
- Long-Wavelength Attachment Feature (FC3442 and FC3443) 117
- low level format 103
- low scratch state 216
- Low Voltage Differential (LVD) 50
- low-latency directors 137
- LPAR 64
- LRU 11–12, 27, 67
 - algorithm 59, 95
 - order 61
- LTO 192
- LTO Ultrium tape cartridges 50–51
- Lv value 147
- LWORM 75–76

- compliant 75
- volumes 75–76

M

- machine type 655
- mainframe 4
- Mainframe Attach (FC9000) 109
- Manage Logical Libraries window 182
- Management Class 66, 82–83, 85, 87–92, 385, 573, 658
 - ACS routines 377
 - actions 72
 - consistency point settings 85
 - construct 72
 - copy consistency point definitions 87
 - copy consistency points 754
 - copy mode 760
 - copy requirements 82
 - definition 87–88, 754
 - definition window 82
 - explicit direction 82
 - locality 87
 - MCPOOLDR 85
 - MCPROD01 87
 - MCPROD02 88
 - Name 82
 - policy name 543
 - possible Cluster Copy Data Consistency Points 247
 - storage construct 82
 - window 245
- Management Classes 56, 82, 87, 227–230, 245–246, 288, 336, 338, 345–346, 450, 481–483, 528, 530, 760
- Management Console 128
- management interface xv, 3–4, 21–22, 29–30, 33, 35, 37, 52, 54–56, 59, 61, 64–67, 72–75, 77–79, 82, 85–87, 89, 118, 128–131, 138, 140–141, 146, 151–152, 157, 159, 167, 176, 178, 188–189, 192, 194–195, 198, 200, 203, 205–206, 212, 214–215, 218, 220, 222, 225–226, 228, 245, 249–250, 255–256, 258, 260–261, 267, 288, 297–298, 304, 311, 319–320, 334, 336, 338, 345–346, 353, 361, 363, 367, 370, 373–374, 376, 400, 409–410, 415, 417–420, 424, 429, 431–432, 436, 438, 440, 460, 467, 474, 480, 482, 496, 498, 500–502, 505–506, 508, 511–512, 514, 519–521, 532–534, 574, 579, 582–583, 585, 606, 628–629, 631–632, 635–640, 645, 649, 652, 727–728, 735, 743, 745–747, 749–751, 754, 756–757, 763, 765, 800, 813, 846
 - access 130
 - IP address 129, 131
 - login page 415
 - ownership takeover 735
 - Storage Group constructs 64
- management policies 27, 103, 165, 176, 336, 339, 347, 352, 716, 740, 800
- manual mode 542, 583
- Manual Operations menu 198
- manual takeover mode 79
- mask 130
- Master Console 412
- Master Console for Service (FC2713) 110, 113
- MAXCAPacity 398

- Maximum Cartridges menu 183
- maximum number 144, 391, 730
- maximum number of cartridge slots 183
- MBps 222
- McDATA 136
- MDR (Miscellaneous Data Records) 161
- media 142
- Media Type 192, 287, 692
 - default volume size 390
 - duplicate VOLSER ranges 192
 - empty physical volume counts 148, 574
 - filling state 692
 - physical media 696
 - scratch threshold 143
 - volume ranges 394
- MEDIA1 29
- MEDIA2 29, 143
- memory expansion 3
- memory upgrade 318
- merge process 350
- migrated volumes 11
- migration 11
- MIH 58, 177, 264, 284–285, 334, 338, 343, 359, 362, 366, 369, 372, 375, 646, 799
 - settings 177
 - time 284
 - value 284–285, 343
- minimum retention time 61
- minute interval
 - statistics reports operations 652
- Miscellaneous Data Records (MDR) 161
- Missing Interrupt Handler
 - See MIH
- ML2 386, 390–391, 394
- Model CC8 Cache Controllers 18
- Model CX7 Cache Drawers 18
- Model L23 52
- Model S24 116
- Model S24 frame 51
- Model S54 frame 51
- Model TS1120 53
- Model XS7 Cache Drawers 18
- Modify Pool Properties 214
- monolithic system 34
- mount command 86
- MOUNT FROM CATEGORY 220
- mount process 92
- mount vNode 88, 90, 93
- mounted tape volume I/O operations 89
- mounting cluster 66
- MOUNTRetention 398
- Move function 69
- move process 350
- moving data 376
- Multi Cluster TS7720 56
- multi-cluster configuration 18, 31, 57, 88, 103–104, 109
- multi-cluster grid 13, 19, 23, 25–26, 29, 34, 62, 77, 91–93, 108, 129–130, 135, 139, 151, 166, 243, 246, 248, 258–259, 267, 284, 286, 295, 304, 314, 320, 326, 340, 357, 393–394, 396, 403, 410, 436, 563, 567–568, 596,

- 636, 638–640, 652–655, 658, 710, 717, 739, 764, 828, 831, 836
 - configuration 12–13, 77, 81–82, 100, 134, 410, 684
 - copy control 288
 - environment 656, 710
 - installation 199
 - operation 276
 - TS7700 Virtualization Engine 286
 - installation 286
- multifile volumes 166
- multi-mode 133
- Multi-Path Architecture 52
- multi-platform environment 187
- multiple Grid configuration 138
- multiple logical volumes 164
- multi-volume set 143, 389
- MVS 540–541, 565, 719
 - operator commands 540

N

- NAT 128
- native (E05) mode 49, 52, 156
- native drive 164, 352
 - IDRC compression 393
 - recommendations 398
- native longwave Fibre Channel transmitters 133
- native mode 157
- Native z/VSE 305, 915
- network cables 129
- network management team 131
- Network Time Protocol
 - See NTP
- No Copy 86, 90
- No copy consistency point 333
- No Copy option 82
- No Factory Cables (FC9700) 110
- non-concurrent upgrades 123
- non-encryption-capable 118
- non-fast ready 87, 90
 - mount 91, 94
 - mount request 94
- non-zero
 - delete expired data parameter value 149
 - Expire Time 223–224
 - setting 150
- NOSTACK 392
- NTP 32, 567
 - server 32

O

- OAM 60, 67, 140, 170, 177, 241–242, 260, 280–281, 286, 288, 330, 334–335, 338, 345, 360, 363, 366, 369, 372, 375, 384, 390, 401–402, 410, 541–544, 576–577, 579–581, 719, 722, 729, 799–801
 - defining new tape library 177
 - documentation resources 286
 - eject logical volume 576
 - object 401
 - uses LCS 140

Object Access Method

See OAM

old data sets 834

On Demand business 168

ONLINE cluster 31

ONLINE peer 31

open systems 192

attachment 52

environment 187

hosts 52

hosts attachment 52

partition 187

operator console 410

Operator Interventions 204, 225

OPM 60, 149, 260, 389, 543

setup 2 800

OS/390 operator commands

LIBRARY EJECT 540

Out of Cache Resources

condition 61

state 61–62

out of scratch situation 148

Outboard Policy Management

See OPM

override settings 34, 89, 92–93

ownership takeover 78, 96, 523, 672, 740

further information 760

mode 55, 138

P

packet loss 138

Panic Reclamation mode 216

Panic scratch state 216

Parity Disks 113

PARMLIB 148, 162, 220, 282, 284, 286, 334, 338, 343, 359, 362, 366, 369, 372, 375, 389–390, 395, 798, 800, 913

PARMLIB definition 148

pause mode 568

Pc value 147

PCI adapter slots 42, 46

PDS 387

peer cluster 66, 77

peer clusters 66

peer TS7700 Virtualization Engine 62

peer TS7740 Virtualization Engine clusters 66

Peer-to-Peer 18, 31–33, 35, 139–140, 277, 286, 304, 357, 370–371, 373, 576, 579, 713

Peer-to-Peer Copy Modes 92–93

Peer-to-Peer VTS 26, 32–33, 81–82, 87–89, 92–93, 352, 370–372, 376, 576, 831

Deferred Copy Mode 26

Deferred Copy/Preferred Mode 93

Immediate Copy/Balanced Mode 92

performance

components 604

Performance Scaling 119

performance throughput capacity 122

PG0

See Preference Group 0

PG1

See Preference Group 1

physical cartridge 29, 191, 193, 216, 739

insertion 193

physical device 52

physical drive 10, 25, 28, 217, 392

internal allocation 602

TS7700 Virtualization Engine 394

physical library 25, 541, 656

operational error 541

physical media

counts 148

pool 215, 657

physical stacked volumes 74, 141

physical tape 11, 51, 152, 249, 754

attachment 11

cartridge 67

client data 397

device 656

drive 10–11, 57, 716

library 583

only resident 152, 249

separation 397

several logical volume 11

physical volume 11, 87, 142, 146, 392, 419, 683, 686, 920

active data 146

being reclaimed 149

content management 176

following number 148

logical volume 164

pools 68, 73, 141

ranges 204

Physical Volume Pooling 64

Physical Volume Pools selection 73

Physical Volume Ranges window 196

piggy-back recall 386

Plant Install 3956 CC6 109

Plant Install 3957 V06 109

Plant Install in 3952 F05 (FC9352) 110

Point-In-Time (PIT) 652

Point-in-Time Statistics 35, 653

binary reports 710

policies 140

policy-based cache management 59

pool

assignments 63

definition of 812

scratch pool definition 812

statistics 150

pool 00 64, 574, 690

pooling by expiration 150

Port Address 137

port type 137

post-installation tasks 178

predefined VOLSER range 205

Prefer Local Cache for non-Fast Ready Mounts 89

Prefer Local Cluster for Fast Ready Mount 138

Prefer Local Cluster for Non-Fast Ready Mounts 138

preference group 57, 656

- Preference Group 0 61–63, 67, 95
 - volumes 59
- Preference Group 1 63
 - private volume 61
- preference level 0 58–60
- preference level 0 volume 58–59
- preference level 1 volume 58–59
- pre-installation planning 177
- premigrated 66, 72, 74
- premigration 11, 29, 66, 150
- Premigration Management 58
- premigration priority threshold 76
- premigration throttling threshold 76
- pre-owned volumes 78
- PRESTAGE 336, 339, 666, 896
- prevent reclamation 215
- prior generation 81, 87, 163, 584
 - B10 685
 - PTP VTS 81
 - VTS 81
- private stacked volumes 65
- production systems 137
- PtP VTS
 - See Peer-to-Peer
- Pv value 147

Q

- QLIB List 275, 929
- Query Library (QL) 282–283, 927
- query tape (QT) 268

R

- R1.4 166
- R1.5 108, 110, 121
 - code 124
- R1.6 12–13, 111, 115, 175
- R1.7 104
 - cluster 3
- rack-optimized server 120
- RAID 33
 - controller 57
 - group 57
 - implementation 57
- RAID 5 10, 44
- RAID 6 43, 113
- RAM 33
- random pattern 149
- Read Configuration Data (RCD) 284
- read ownership takeover (ROT) 78, 666, 672, 728, 735
 - mode 138
- read/write ownership takeover mode 138
- read-only processing 582
- recall takeaway 386
- recalled volume 63
- receiving port 133
- reclaim
 - inhibit schedule 218
 - policies 216
 - scheduled time 148

- tasks 216
- threshold 216
- threshold setting 147
- Reclaim Threshold Percentage 147, 220
- reclaimed physical volumes 68
- reclamation 69, 71, 147–148, 150, 216, 218, 220
 - Inhibit Reclaim Schedule 215
 - interruption 218
 - level 214
 - methods 216
 - operation 74
 - policies 65
 - prevent 215
 - process 70, 148, 150, 220
 - setting 147
 - tasks 146, 213
 - threshold 65, 68
- reclamation process 69
- reconciliation process 70
- record block sizes 122
- record types 654, 691
 - 14, 15, 21, and 30 161
- recovery
 - objectives 62
 - scenarios 582
 - time 57, 259, 739
- Recovery Point Objective (RPO) 72
- RECYCLE 378, 380, 394
- Redbooks website 935
 - Contact us xviii
- redundant Fibre Channel switches 108
- redundant network routers 108–109
- reinstall 103
- reliability 4
- remote
 - cluster 58, 130, 584
 - virtual devices 91
 - support capability 119
 - tape device 136
 - tape volume cache accesses 138
 - TS7700 Virtualization Engine 92
 - TS7700 Virtualization Engine Cluster 98
 - TS7700 Virtualization Engine Clusters 67
 - TVC 89, 93, 574, 754
- Remote Copy Management Facility (RCMF) 752
- removal policy 60
- removal threshold 61
- replication 132
- request volume 661–662
- RESPONSE Data 665
- response record 660, 663
 - human readable header 687
 - total length 689
- Retain Copy Mode setting 83–84, 339
- retention 13
- return-to-scratch processing 68, 142–143
- Rewind/Unload 58, 86, 247
 - clusters 90
 - command 83, 87, 89–90, 333
 - command completion 83

- command processing 86
- consistency point 90
- copy consistency point 83
- copy consistency point requirement 87
- operation 26
- time 66, 83
- RJ45 77
- RMM 170, 299, 328–329, 350–351, 358, 360, 365, 368–369, 371–372, 374, 376, 379–380, 395–396, 575–576, 910–911
- RMM activity 396
- RMM CDS 396, 576
 - correct size 396
- RMS function 304
- router configuration 131
- Rp value 147
- RUN 26, 58, 62, 85–86, 90–92, 94, 96, 137, 151–152, 229, 247–249, 284, 295, 302, 314, 332–333, 336, 344, 394, 437, 445, 482, 494, 530, 560, 563, 574, 593, 605–606, 615, 618–619, 669–670, 699–700, 710, 715, 717–718, 726, 732, 735–738, 740, 754, 777, 890, 892, 894–896, 898, 900–901, 904, 907
- RUN consistency points 90
- RUN copy consistency point 333
- RUN,RUN (RR) 91

S

- S/390 116
- S/390 servers 117
- S1120 Tape Drives 156
- S54 frame 51
- TS7700 Virtualization Engine 58
- TS7720 Virtualization Engine Cache Controller 43
- TS7740 Virtualization Engine Server 46
- sample JCL 909
- SATA Cache Controller 113–114, 123
- SC Duplex 117
- SCDS 287–288, 292, 330, 334, 338, 359–360, 363, 366, 368, 372, 375, 800–801
- SCDS Name 287
- Schedules table 218
- scratch cartridge 149
- scratch categories 69, 103, 141, 149, 215
- scratch logical volume 143
- scratch mount 69
- scratch mount category 222
- scratch mount times 166
- scratch physical volumes 147
- scratch pool 148, 260, 391, 398, 573, 696
 - return volumes 689
- scratch state 217
- scratch volume 28, 142, 146, 220, 280, 404, 540, 666, 761, 914
 - candidate list 66
 - nominal number 143
- SCSI controllers 42, 46
- SCSI element address 197
- secondary copy pool 289–291
- secondary key manager 74
- secondary pool 66, 70, 87

- Secure Data Erase 68, 149–150
- Secure Data Erase function 68
- Selective Dual Copy 66
- Selective Dual Copy function 65–66
- Selective Write Protect 3, 5
- Selective write protect for DR testing 5, 104
- SEQUENCE number 264, 664
- sequential addresses 131
- SERIAL number 274, 566, 655, 925
- Service (QoS) 137
- Service mode 31, 78, 521, 728
- Service Prep mode 31, 78, 319, 334, 337–338, 568
- SETNAME statement 824
- SETSYS 388–391, 393
- shortwave optical adapters 124
- shortwave optical fiber 132
- shortwave/longwave 122
- Short-Wavelength adapter 117
- shredding encryption keys 69, 149
- shredding user data 101
- Single Cluster 23, 25, 139, 199, 227, 264, 274, 286, 328, 343, 345, 357, 361–364, 366–367, 369–370, 376, 410, 535, 567, 594–595, 616–617, 808, 821, 830
 - configuration 70, 129
 - configuration requirements 108–109
 - hardware components 410
 - TS7700 Virtualization Engine 118
- Single Cluster Grid 19, 23, 25, 32, 34, 65, 166, 199, 248, 277, 284, 286, 300, 330, 342, 353, 357–359, 361–368, 370, 372, 375, 387, 742, 823
 - configuration 23
 - stand-alone TS7700 Virtualization Engine 202
- single cluster TS7720 Virtualization Engine 38
- single library 139
- single point of failure 31, 96, 132
- single use feature 101
- single-port adapters 77
- SMC commands 55
- SMF 162
- SMF records 160–162, 378
- SMIT 80, 258–259, 334, 338, 345–346, 361, 363, 367, 370, 373, 376
- SMS 26, 29, 140, 170, 232, 259, 264–267, 276, 279–281, 303, 329–330, 334–335, 337–338, 342–343, 345–347, 350–351, 358–360, 362–363, 365, 368, 371–372, 375, 378–380, 384, 389–390, 394, 399, 540–543, 566, 579–581, 719–720, 722–723, 740, 751, 769, 772, 775, 778, 799, 801
 - classes 141
 - construct 64, 264, 389
- SMS Tape Design 168
- SMS-managed 140
- SMS-managed z/OS environment 192
- SNMP 239
- software requirements 139
- solutions
 - Conversion (SC) 143
- source cluster 137
- source TS7740 71
- space reclamation 212

- spanning 143
- Specialist web interface 196
- SSIC 135
- SSR 35, 55, 57, 80–82, 89, 101, 119, 127, 169, 179, 187, 198, 200, 259, 266–268, 319–320, 325, 329, 334, 337–338, 341, 343, 345–347, 351, 355–356, 358, 362, 365, 368, 371, 374, 412, 414, 433, 520, 582–585, 644, 648–649, 665, 728, 730–737, 739, 745, 803, 805 SSR) 31
- stack multiple files 166
- stacked cartridge 90
- stacked cartridges 70
- stacked media 141
- stacked physical volume 165
- stacked volume 11, 27, 63–65, 67–69, 73, 166
 - add 583
 - function 150
- Stacked Volume Pool properties 65
- stacked volume pools 30
- stacked volumes 30, 64–65, 216, 298, 395, 919
- stacking 11
- stand-alone mode 400
 - source device 401
- Stand-Alone Services 399, 401
 - installation procedure 399
- stand-alone TS7700 23
- stand-alone TS7700 Virtualization Engine 23
- status of read-only 150
- storage administrators 167
- Storage Class (SC) 59–61, 63, 140, 177, 230, 581, 801
 - ACS routine 140, 377–378
 - attribute 60
 - cache management effect 259
 - construct 62–63, 259, 344, 556–557, 647, 676, 762
 - construct names 63
 - definition 60
 - name 59–60, 543
 - policy name 543
 - preference 76
- Storage Expansion Frame Cache Controller 41
- Storage Group (SG) 63, 65, 140–141, 389, 542, 801
 - construct name 74
 - constructs 63
 - definitions 359, 365, 368, 371, 375
 - distributed libraries 288
 - library names 542
 - separate pools 72
 - state 138
- Storage Management Subsystem 543
- storage pool 64, 153, 397
- subnet mask 128, 130–131
- subsystem performance 153
- support for a four-cluster hybrid grid configuration 5
- Support Line 168
- Supported Workflow Management Enhancements 76
- SW fiber connections 124
- switch port 133
- SYNCDEV 393
- SYS1.PARMLIB 162, 220
 - definitions 286

- sysplex 137, 141
- System p server 33
- System Storage Interoperation Center (SSIC) 135
- System Storage TS1130 Tape Drive 53
- System z host 91–92
- System-Managed encryption 160
- system-managed storage
 - See SMS
- system-managed tape 140, 382, 575
 - environment 29, 540
- Systems Assurance Product Review (SAPR) 798

T

- takeover mode 67
- takeover scenario, TSSCs (TEST) 739, 741
- tape addresses 4
- tape cartridge 387, 401, 570
 - slow operation 403
- tape controller 187
- Tape Copy Service 168
- tape data 28, 146, 387, 658
 - host interaction 28
- tape drive failure 582
- tape drives 14, 23, 384, 582
- tape frame 36
 - input power requirements 126
- Tape Library 127, 163, 179, 205, 285, 409, 797
 - automated system management 263
 - other cartridges 30
 - Storage Administration Guide 286, 409, 540, 720–721, 934
- Tape Library Specialist web interface 181, 197
- Tape Management Catalog 163
- Tape Management System (TMS) 308, 379, 576, 754, 798
- tape media technology 11
- Tape Mount Management (TMM) 402
- tape processing jobs 164
- tape subsystems 163
- Tape Table-of-Contents (TTOC) 389
- tape virtualization 11, 16
- Tape Virtualization for mainframe 3
- tape volume 11, 27, 57, 140, 308, 381, 576, 590, 602, 606, 918
 - ISMF EJECT line operator 576
 - storage conventions 657
- Tape Volume Cache
 - See TVC
- Tape Volume Cache Management 57
- Tape Write Immediate 28
- TAPECOMP 161
- TAPECOPY 387, 392–393
- TAPESPANSIZE 388, 391
- TAPEVOL 401
- target TS7740 Virtualization Engine 73
- TCDB 265, 286, 311, 328–329, 350–351, 353, 358, 360, 362–363, 365–368, 371–372, 374–375, 540, 542–543, 576, 755, 799–800, 910–911
- TCP/IP 16, 33, 118
 - address 129–130

- connections 134
- protocol 129
- third IP 128–129
- three-cluster configuration of TS7720 Virtualization Engines 39
- three-cluster Grid 13, 19–20, 25, 39, 80–81, 90, 98–99, 125, 130–131, 248, 264, 276, 326–327, 330–340, 346, 597, 619–621, 698, 727, 737, 740, 764, 773, 777
 - configuration 25, 276, 331–332, 619
- three-cluster Grid configuration 104
- three-cluster hybrid grid 339, 347
 - configuration 339
- threshold value 147
- throttling controls 76
- throttling mechanism 76
- time period 222
- time statistic 652, 657
- timeout protection 313
- Tivoli Lifecycle Key Manager 73, 153
 - server 73
- Tivoli Storage Manager 379–380, 396–398, 917
- TMS 161–162
- TMS catalog information 162
- tokens 30–31
- total cost of ownership (TCO) 4, 6, 14
- tower server (IBM System x 206m) 120
- Transaction Processing Facility (TPF) 140, 177, 310
- TS1120 48, 50, 53, 73–74
- TS1120 drives 52
- TS1120 tape controller 55
- TS1120 Tape Drive 48, 52, 118–119, 157, 160, 353, 410, 537
 - encryption capable tape drives 112
 - Model E05 156
 - native mode 119
- TS1120 Tape Drive Model E05 53
- TS1120 Tape Drives 48–49, 73
- TS1130 48, 52
- TS1130 Tape Drive 48–49, 53, 118, 153
 - 3592-E06 118
 - E06 119
 - Model E06 156
 - native mode 119
- TS1130 Tape Drives 48–49
- TS3000 Service Console 36
- TS3000 Service Console. 36
- TS3000 System Console 37, 40, 79
- TS3000 System Console (TSSC) 37, 79–81, 108, 111, 113, 117, 119–120, 131, 167, 728
 - external 109
 - FC2721 110, 113
 - FC2730 and FC2719 110, 113
 - network 131
 - with Keyboard/Display 109
- TS3000 System Consoles 79
- TS3500 17, 28, 36
- TS3500 Tape Library 15, 25, 48–52, 64, 107, 111, 115, 124, 140, 175–180, 191–193, 195–197, 204–205, 367, 401, 410–411, 583, 628–631, 730, 798
 - 3584 Tape Library 110

- 3953 52, 176
- attachment 52
- Cleaning Cartridge Usage 539
- definitions 178
- FC9219 110
- frames 190
- High Density frame 115
- IBM System z attachment 52
- level 193
- logical library partition 193, 573
- Models D23 51
- needs cleaning 197
- operator panel 179
- operators 160
- physical cartridges 205
- TS7700 Virtualization Engine owned physical volumes 178
- unassigned volumes 195
- TS3500 Tape Library and its four to 16 IBM 3592 Tape Drives and two FC switches 36
- TS3500 Tape Library frame 50
- TS3500 Tape Library frames 50
- TS3500 Tape Library High Density frame 51
- TS3500 Tape Library Model D23 50
- TS3500 Tape Library Model D53 51
- TS3500 Tape Library Model L23 50
- TS3500 Tape Library Model L53 50
- TS3500 Tape Library Model S24 51
- TS3500 Tape Library Model S54 51
- TS3500 Tape Library Specialist 30, 167, 176, 178–179, 195, 197, 539, 629
 - Welcome page 411
 - Welcome window 181
 - windows 187
- TS3500 Tape Library Specialist web interface 196
- TS3500 Tape Library Web Specialist 188
- TS3500 Tape Library. 20
- TS7700 3, 5, 15, 22, 34, 54, 62, 77, 80, 83, 86, 96–97
- TS7700 cluster 61
- TS7700 Clusters 9
- TS7700 Grid 94
- TS7700 Virtualization 60
- TS7700 Virtualization Engine 3–4, 6–7, 9–16, 21–22, 27–30, 32–38, 52, 54–60, 62–63, 66–67, 76, 79, 86, 88–90, 102, 107–109, 112, 116, 118–119, 121–123, 125–126, 128–129, 131–134, 136, 139, 141–146, 160–161, 163–168, 175–177, 191, 204–205, 215–216, 220, 222, 224, 267, 287, 304, 309, 377, 384, 389, 391, 397, 401, 405, 410, 415, 574, 590, 654, 660, 725–726, 729–730, 734, 753, 798, 918
 - activity 215, 711
 - availability characteristics 126
 - back-end cartridges 193
 - Base Frame 109
 - Cache 164
 - hit 165
 - Cluster 86, 124–125, 128, 130–131, 138, 165, 436, 632, 638, 640, 726
 - communication paths 728

- Gb Ethernet link 131
 - in a multi-cluster grid 166
 - tape volume cache 726
 - virtual device addresses 726
- Cluster 0 735
- Cluster 1 733
- component upgrades 120
- configuration 120
- control path drives 191
- default behavior 753
- enhancements 5
- environment 119
- environmental conditions 126
- FC4871 111
- Gbit Ethernet links 131
- Grid 130, 137, 165
 - configuration 122, 138, 165–166, 640, 726, 755
 - failover test scenario 729
 - subsystem 579
- Hold function 280
- host definition 266
- implementation 168, 176
- installation 169
- internal management functions 216
- key benefits 164
- LAN/WAN requirements 128
- local 166
- logical volume 142
- main management concepts 56
- management interface 128–129, 195
 - definitions 193
- modular design 33
- multi-cluster grid 130
- multi-cluster grid environment 129–130
- multiple ranges 144
- NAT routers 129
- node (3957 Model V06 or VEA) 117
- nomenclature 6
- online drive 801
- overall performance 713
- overall throughput 165
- performance 132
- physical tape drives 28
- physical volumes 148
- pre-installation 117, 169
- primary attributes 166
- R1.3 149
- reconfiguration 187
- remote 166
- same volume 665
- Series 729
- Server 8, 110–111, 114, 122, 124
- significant distance 665
- single or multi-cluster 9
- software definition 177
- software failure 585
- Stand-Alone 124
- subsystem 28, 108, 146, 578, 604
 - Shared Resources 604
- system requirements 125

- Tape Volume Cache 57
- throughput capability 266
- upgrade options 121
- virtual 3490E tape drives 264
- virtual drives 274
- Virtualization Interface management interface 178
- TS7700 Virtualization Engine architecture 16, 19, 35
- TS7700 Virtualization Engine Cache Controller 17
- TS7700 Virtualization Engine Cache Modules 36
- TS7700 Virtualization Engine Cluster 17, 78, 86, 97, 103
- TS7700 Virtualization Engine cluster 55
- TS7700 Virtualization Engine Clusters 19, 58, 63, 67, 78–79, 95–97, 326, 330, 332–333, 340, 416, 483, 489, 521, 666, 726–728, 732–733, 736, 739
- TS7700 Virtualization Engine configuration 36
- TS7700 Virtualization Engine database 35
- TS7700 Virtualization Engine family 6
- TS7700 Virtualization Engine Grid 101
- TS7700 Virtualization Engine Grid configuration 96
- TS7700 Virtualization Engine Grid configurations 19
- TS7700 Virtualization Engine Integrated Library Manager 56
- TS7700 Virtualization Engine logical volumes 55
- TS7700 Virtualization Engine Management 35
- TS7700 Virtualization Engine management interface 35, 54, 72, 78, 85
- TS7700 Virtualization Engine management interface (MI) 55
- TS7700 Virtualization Engine node 36, 78
- TS7700 Virtualization Engine Policy Construct definitions 55
- TS7700 Virtualization Engine R1.5 35, 54, 67
- TS7700 Virtualization Engine R1.6 39
- TS7700 Virtualization Engine R1.6 Grid configurations 61
- TS7700 Virtualization Engine R1.7 5, 15, 19, 36, 44, 56, 61, 73, 103
- TS7700 Virtualization Engine R1.7 clusters 103
- TS7700 Virtualization Engine R1.7 Grid 18
- TS7700 Virtualization Engine replication priority 63
- TS7700 Virtualization Engine resource 76
- TS7700 Virtualization Engine Server 37, 41
- TS7700 Virtualization Engine Service 55
- TS7700 Virtualization Engine settings 92–93
- TS7700 Virtualization Engine statistics 35
- TS7700 Virtualization Engine subsystem 35
- TS7700 Virtualization Engine Tape Volume Cache 57
- TS7700 Virtualization Engine Two-Cluster Grid configuration 97
- TS7700 Virtualization Engine, Release 1.6 75
- TS7700 Virtualization Engines 37, 91, 100
- TS7700 Virtualization Engines at R1.7 104
- TS7720 38, 61, 84
- TS7720 cache 61
- TS7720 cluster 61
- TS7720 Logical Volume Removal Policy 60
- TS7720 Storage Expansion Frame 18
- TS7720 Virtualization Engine 6, 8–9, 11–13, 19–20, 27, 36–39, 41, 44, 54, 56, 61–62, 109, 111–115, 122–123, 127, 176

- 3957-VEA 122
- configuration 10, 113
- Disk Only Server 109
- Disk Only solution 167
- Frame (FC7322) 112
- SATA Cache Controller 113
- SATA Cache Controller Model CS7 8
- SATA Cache Drawers 114
- Server 113–114
 - 3957 Model VEA 112
 - Model VEA 8
- TS7720 Virtualization Engine Base Frame 18, 40
- TS7720 Virtualization Engine Cache 43
- TS7720 Virtualization Engine cache configurations 60
- TS7720 Virtualization Engine Cache Controller 43
- TS7720 Virtualization Engine Cache Controllers 43
- TS7720 Virtualization Engine Cache Drawer 44
- TS7720 Virtualization Engine Cache Model CS8 Controller 36
- TS7720 Virtualization Engine Cache Model XS7 Drawers 36
- TS7720 Virtualization Engine clusters 61
- TS7720 Virtualization Engine configuration 27
- TS7720 Virtualization Engine configurations 57
- TS7720 Virtualization Engine disk-only solution 36
- TS7720 Virtualization Engine Grids 39
- TS7720 Virtualization Engine Model CS8 Cache Controllers 36
- TS7720 Virtualization Engine Model XS7 Cache Drawers 36
- TS7720 Virtualization Engine R1.6 PGA1 104
- TS7720 Virtualization Engine R1.7 18, 38
- TS7720 Virtualization Engine SATA Cache Controller 38
- TS7720 Virtualization Engine SATA Cache Drawers 38
- TS7720 Virtualization Engine Server 41–42
- TS7720 Virtualization Engine server 36
- TS7720's deep cache 95
- TS7720 Virtualization Engine 57
- TS7740 Virtualization Engine 6–8, 11–12, 21, 25, 28, 30, 36–37, 47–48, 52–53, 55–59, 63–65, 67–74, 77, 84, 89, 107–109, 111–112, 118–119, 123–125, 127–128, 147–150, 153, 156–157, 160, 164–167, 176–179, 187–188, 191, 193, 196–197, 204, 259, 269, 350, 589, 667, 687
 - attached drives 160
 - back-end 4 Gb switches 111
 - back-end drive 119
 - back-end drive encryption 153
 - backstore 153
 - cache 165–166
 - cache capacity 11
 - cache subsystems 18
 - Cluster 95, 97
 - FICON channels 97
 - cluster 18
 - components 45
 - configuration 48, 57, 111
 - encryption implementation plan 160
 - four drawer 124
 - frame 177

- Grid configuration 125
- hardware 122
- historical statistics 150
- Integrated Library Manager 194
- internal processes 212
- internal storage management software 30
- logical library 192, 194
 - partitions 195
- management interface 167
- Model 3957-V06 122
- node 10–11, 30
- node, storage management software 30
- partition 176
- peers 129
- physical volume cartridge 115
- processing 148
- R1.6 127, 193
- R1.7 configuration 44
- ranges 193
- scratch cartridge 148
- subsystem 115, 192
- tracks 150
- versus TS7720 Virtualization Engine 741
- virtual volume 27
- VTs 187
- TS7740 Virtualization Engine at R1.5 55
- TS7740 Virtualization Engine Cache Controller 7, 46–47, 110
- TS7740 Virtualization Engine Cache Drawer 7, 47, 112, 123
 - Model CX7 112
- TS7740 Virtualization Engine Cache subsystem 124
- TS7740 Virtualization Engine Cluster 65, 195, 666, 740
- TS7740 Virtualization Engine data 66
- TS7740 Virtualization Engine Grid
 - configuration 13
- TS7740 Virtualization Engine partition 52
- TS7740 Virtualization Engine R1.7 48, 52, 104
- TS7740 Virtualization Engine Release 1.6 48
- TS7740 Virtualization Engine Server 45, 108, 110–111, 116
 - 3957 Model V06 109–111
 - Model V06 7
- TS7740 Virtualization Engine tracks 69
- TS7740 Virtualization Engine TVC data 58
- TS7740 Virtualization Engine's database 73
- TS7740 Virtualization Engines 52, 54
- TVC 10–12, 21, 27–30, 57–59, 61–63, 67–70, 74, 82–83, 85, 87–88, 90–91, 93, 101, 148, 165
 - candidate 61–62
 - local 89
 - management 27
 - read hits 166
 - selection 12, 61–62
 - selection processing 67
- two production clusters 137
- two-cluster grid 19–20, 24, 26–27, 31, 62, 90, 95, 97–98, 104, 125, 129–131, 243–244, 247, 264, 295, 326–338, 340, 343, 345, 357, 370–371, 373, 375–376, 387, 617–618, 730, 739, 741, 763, 767, 770

- configuration 20, 24, 95–96, 137, 327, 370, 741
- in multi-cluster grid 13
- two-cluster TS7700 Virtualization Engine Grid 125

U

- Ultra2/Wide High 50
- Ultrium (LTO) tape cartridges 51
- unassigned category 196
- unassigned drives 186
- unassigned physical volumes 204
- unencrypted 149
- unencrypted physical volume 149
- ungrid 100
- Unit Control Block (UCB) 275
- UNITNAME 822, 827, 834
- Universal Time (UTC) 218
- upgrade Grid capacity 103
- upgrade window 124
- URL line 197, 538
 - library URL 197
- User roles 22
- User Security Enhancements 127

V

- valid copy 92, 152, 249, 668, 729, 754
- VEHSTATS 656
- Version 3.1.2 140
- virtual device 10, 18, 78, 139, 266, 270, 400, 668, 729–730
 - first set 270
 - IPL volume 401
 - large number 10
 - second set 271
 - TS7700 Virtualization Engine
 - Cluster 727
 - TS7740 Virtualization Engine
 - Cluster 78
 - virtual volume 12
- virtual devices 139
- virtual drive 28, 96, 141, 288, 398, 713
 - maximum number 730
 - preferred path 28
- virtual drives 10
- virtual IP address 128–129
- virtual machine (VM) 307–308
 - z/VM guest 307
- virtual mount 122
- Virtual Tape Controller
 - See VTC
- virtual tape devices 137
- virtual tape drives 10
- Virtual Tape Server
 - See VTS
- virtual tape subsystem 10
- virtual volume 11–12, 60, 141, 311, 391, 604
 - full capacity 398
 - initial creation 247
 - insert processing 303
 - remaining space 391

- virtual volume data 27
- virtualization 11
- Virtualization Engine 4
- Virtualization Engine FICON cables 113
 - FC0201 through FC0206 110
- Virtualization Engine Hybrid Grid 15
- Virtualization Engine R1.6 members 103
- Virtualization Engine, 11
- virtualization node
 - See vNode
- vital information 13
- VM guest (VGS) 308, 310
- VM support 140
- VM system 308
- VM tape library 140
- VM/ESA
 - DGTVCNTL DATA 305
 - LIBRCMS 308
 - LIBSERV-VGS 308
 - RMCONFIG DATA 305
 - specific mount 305
 - VGS 308
 - VSE/ESA guest 308
- VMA 162
- vNode 16, 21, 34, 57, 81–83, 87–90, 93, 636
 - cluster 90
- VOLSER 28–29, 66, 68, 142, 144, 178, 191–192, 194, 203–205, 242, 260, 262, 282, 297, 311, 343, 379, 381–383, 393, 395–396, 428, 436–437, 439, 442, 444, 447–449, 451–452, 455–456, 464, 466, 468, 470–472, 474–475, 478–479, 526, 535, 561, 572–573, 575, 584, 630, 655, 667–668, 677, 681, 684, 686, 698–699, 811, 834, 844, 860, 866, 920–921
 - entry fields 205
 - range 65, 178, 204–205, 758, 761
 - ranges 205
- VOLSER Ranges Table 203–205
- volume
 - category, defining 913
 - copy policy management 82
 - copying 137
 - counts by media type 148
 - map 658
 - ownership 77
 - ownership takeover 79
 - premigration requests 67
 - processing times 166
- Volume Copy Retention Group 344
- Volume Copy Retention Time 344
- Volume Pool 216
- Volume Removal 254
- Volume Removal Policy 12
- volume Rewind/Unload time 83
- Volume Status
 - Information 665
 - request 665
- volumes tagged for erasure 150
- VSAM 405
- VSE Guest 915
 - Server 308

Server Considerations 308

VSE/ESA

guest 308

RMSMASTR 308

tape management 310

VTC 32, 35, 81, 89

VTs xvi, 3, 18, 26–27, 31–35, 53, 60, 67, 81–82, 87–88, 92–93, 139–140, 143, 187–188, 205–206, 277, 293, 304, 325, 349–355, 357–362, 364, 366, 368–373, 376, 399, 456, 568, 573, 576, 579, 582, 584, 590–591, 593, 652, 659, 661, 664, 666, 678, 683, 696, 710, 719, 782, 789, 831, 840, 890, 892–893, 896, 898, 900, 911, 918–920, 923

architecture 32

cache 355

capabilities 33

database 353, 355, 360–361, 363

partition 195, 351

systems 163

VTs B10 357, 371

VTs Model B10 350

VTs Tape Volume Cache 360, 363

W

WAN 118, 132

connection 132

paths 118

ports 37

topology 132

web browser 179

web interface 197

web-based management interface 77

Welcome Page 197

wireless protocols 168

work item 186

Workflow Management controls 38

Workflow Management Enhancements 76

workflow management functions 62

World Wide Identifier 75

WORM 75, 119

WORM media 75

WORM tape 75

wrapped data keys 153

write appends 83

Write Once Read Many

See WORM

Write Ownership Takeover (WOT) 78, 728

Write-Mount Count 76

Write-Protect Mode 492

Z

z/OS 140–141, 177, 263, 537, 565

environment 67

facilities 77

host 94, 203

Host Console Request function 63

host software 67

workload 60

z/OS DFSMS 222, 540, 720–721

z/OS V1R11 95

z/VM 140, 177

z/VM V5.1.0 140

z/VSE 140, 177

z/VSE System Macros (Z/VM) 140, 264, 304

zSeries

host 132

host FICON channel extenders 132

logical library 573



IBM Virtualization Engine TS7700 with R1.7

(1.5" spine)
1.5" <-> 1.998"
789 <-> 1051 pages



IBM Virtualization Engine TS7700 with R1.7



Redbooks®

New cache options for capacity and performance

Improved disaster recover testing capabilities

Enhanced hybrid multi-cluster grid management

The IBM Virtualization Engine TS7700 provides tape virtualization for the IBM System z environment. The IBM Virtualization Engine TS7700 offers improved performance and capacity to help lower the total cost of ownership for tape processing. It introduces a modular, scalable, and high-performing architecture for mainframe tape virtualization. The TS7700 Virtualization Engine is available as the IBM Virtualization Engine TS7740 R1.7, which brings significant enhancements over its previous release.

The TS7700 Virtualization Engine integrates the advanced performance, capacity, and data integrity design of the 3592 Tape Drives, the IBM industry-leading tape technology, with high-performance disks and an advanced IBM System p server to form a storage hierarchy, managed by robust storage management firmware with extensive self-management capability. It includes functions, such as advanced policy management to control physical volume pooling, cache management, dual copy, dual copy across a Grid network, and copy mode control. The TS7700 Virtualization Engine offers enhanced statistical reporting and a single, standards-based management interface for TS7700 Virtualization Engine management.

This IBM Redbooks publication covers the TS7740 Virtualization Engine Release 1.7 and the IBM Virtualization Engine TS7720, includes information about supported tape drives and libraries, and discusses tasks involved with planning, implementing, and operating the TS7700 Virtualization Engine, and also migration considerations.

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

SG24-7712-02

ISBN 0738435260